Tight Security of TNT and Beyond

Attacks, Proofs and Possibilities for the Cascaded LRW Paradigm

Ashwin Jha[1], Mustafa Khairallah[2,3], Mridul Nandi[4], and Abishanka Saha[4]

[1]CISPA Helmholtz Center for Information Security, Saarbrücken, Germany
letterstoashwin@gmail.com
[2]Seagate Research Group, Singapore, Singapore
[3]Lund University, Lund, Sweden
khairallah@ieee.org
[4]Indian Statistical Institute, Kolkata, India
mridul.nandi@gmail.com,sahaa.1993@gmail.com

Abstract. Liskov, Rivest and Wagner laid the theoretical foundations for tweakable block ciphers (TBC). In a seminal paper, they proposed two (up to) birthday-bound secure design strategies — LRW1 and LRW2 — to convert any block cipher into a TBC. Several of the follow-up works consider cascading of LRW-type TBCs to construct beyond-the-birthday bound (BBB) secure TBCs. Landecker et al. demonstrated that just two-round cascading of LRW2 can already give a BBB security. Bao et al. undertook a similar exercise in context of LRW1 with TNT — a three-round cascading of LRW1 — that has been shown to achieve BBB security as well. In this paper, we present a CCA distinguisher on TNT that achieves a non-negligible advantage with $O(2^{n/2})$ queries, directly contradicting the security claims made by the designers. We provide a rigorous and complete advantage calculation coupled with experimental verification that further support our claim. Next, we provide new and simple proofs of birthday-bound CCA security for both TNT and its single-key variant, which confirm the tightness of our attack. Furthering on to a more positive note, we show that adding just one more block cipher call, referred as 4-LRW1, does not just re-establish the BBB security, but also amplifies it up to $2^{3n/4}$ queries. As a side-effect of this endeavour, we propose a new abstraction of the cascaded LRW-design philosophy, referred to as the LRW+ paradigm, comprising two block cipher calls sandwiched between a pair of tweakable universal hashes. This helps us to provide a modular proof covering all cascaded LRW constructions with at least 2 rounds, including 4-LRW1, and its more established relative, the well-known CLRW2, or more aptly, 2-LRW2.

Keywords: TNT, LRW1, 4-LRW1, CLRW2, birthday-bound attack

Version History: An abridged version of this paper appears in IACR-EUROCRYPT 2024. This is the full version.

This article is an amalgamation and extension of prior work of the same authors. Concretely, it combines and significantly extends the contents of IACR ePrint articles 2023/1212 (by Khairallah), and 2023/1233 (by Jha, Nandi, and Saha) that appeared in August 2023 on closely related topics into a single edited document. This article should be seen as a successor of both these IACR ePrint articles.

1 Introduction

Tweakable Block Cipher or TBC is a highly versatile symmetric-key primitive that has found applications in almost all verticals of modern information security, including encryption schemes [8], message authentication codes [21], authenticated encryption [28,39], and even leakage resillience [43]. The popularity of TBCs is largely credited to the simplicity of TBC-based constructions, and more importantly, comparatively simpler proofs of beyond-the-birthday bound (BBB) security.

In a seminal paper [31] at CRYPTO 2002, Liskov, Rivest, and Wagner (LRW) formalized the notion of tweakable block ciphers (TBCs), although the high-level idea already appeared in some AES candidates such as Hasty Pudding [42] and Misty [13]. Over the years, the design landscape of TBCs has changed progressively. The design of a TBC mainly falls into one of the two categories: ad hoc designs based on well-established primitive design paradigms, or provably secure designs based on block ciphers or cryptographic permutations. In recent years, the popularity of ad-hoc designs has gained momentum with the advent of the TWEAKEY framework [22], its chief example being Deoxys-TBC [23], Skinny [6] and Qarma [2]. These designs are built from scratch, and their security mainly depends on cryptanalysis. On the other hand, the security of provably secure designs is directly linked to the security of the underlying primitives, such as a block cipher, a permutation, or a pseudorandom function. Some prominent examples include LRW's original constructions [31] LRW1 and LRW2, XEX [41] by Rogaway, and its extensions by Chakraborty and Sarkar [9], Minematsu [35], and Granger et al. [16]. Note that all these schemes are inherently birthday bound secure due to detectable internal collisions.

Cascading LRW2: Landecker et al. were the first to notice [30] that a cascading of two independent instances of LRW2 results in a BBB secure TBC construction. They proved that 2-round cascaded LRW2 is secure up to approx. $2^{2n/3}$ CCA queries, where n denotes the block size in bits. The initial proof was flawed [40], and superseded by a corrected proof by both Landecker et al. and Procter [40]. The construction was later found [33,25] to be tightly secure up to $2^{3n/4}$ CCA queries. For any arbitrary $r \geq 2$-round independent cascading of LRW2, denoted r-LRW2, Lampe and Seurin proved [29] CCA security up to approx. $2^{\frac{rn}{r+2}}$ queries.

Cascading LRW1: The idea to cascade LRW1 came quite later in [4], where Bao et al. showed that 3-round cascading of LRW1, referred as TNT, is CCA secure up to $2^{2n/3}$ queries. The design is highly appreciated in the community for its simple design and high provable security guarantee. In fact, the CPA security was later improved to $2^{3n/4}$ queries, essentially matching the bound for 2-round LRW2. Since this later result, it is widely believed that the CPA improvement carries over to the CCA setting, as well. For the more general case of arbitrary $r \geq 3$, denoted r-LRW1, Zhang et al. proved [44] CCA security up to approx. $2^{\frac{r-1}{r+1}n}$ queries.

We remark that the aforementioned LRW-based constructions are all studied under the standard assumption on the pseudorandomness of the underlying block ciphers. However, several good constructions are also based on cryptographic permutations [11,12] and even rekeying[1] of block ciphers [35,32,24]. We skip a detailed discussion on these ideal model constructions since the focus here is specific to the LRW design paradigm. We encourage the readers to see [33,25] for a more inclusive discussion on ideal model constructions.

1.1 Motivation

The primary motivation behind this work is a peculiar non-random behavior exhibited by TNT in the CCA setting.

Suppose $\boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \boldsymbol{\pi}_3$ are three independent random permutations of $\{0,1\}^n$. The TNT construction (see Fig. 1.1) based on $\boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \boldsymbol{\pi}_3$ is a TBC with n-bit tweak and n-bit block input, defined by the mapping

$$(t, m) \xmapsto{\mathsf{TNT}} \boldsymbol{\pi}_3(t \oplus \boldsymbol{\pi}_2(t \oplus \boldsymbol{\pi}_1(m))).$$

As can be noticed by the definition of TNT, it has a peculiar property, that

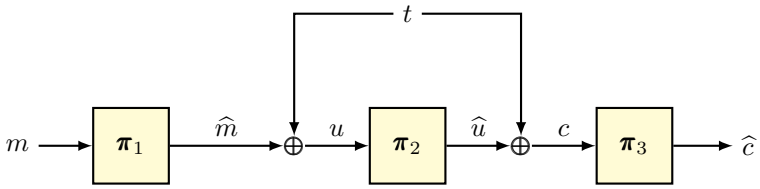

Fig. 1.1: The TNT construction [4].

we refer as the *final-block cancellation* property. Specifically, suppose we have a triple $(t, m, \widehat{c})$ such that $\mathsf{TNT}(t, m) = \widehat{c}$. Then, it is easy to see that any inverse query of the form $(t', \widehat{c})$ would result in a cancellation of the call to $\boldsymbol{\pi}_3$, and this is independent of the tweak values t and $t' = t \oplus \delta$. Essentially, the construction boils down to the one in Fig. 1.2. Let's call it $\mathsf{TNT}_{\delta,m}$ for some fixed $\delta \neq 0^n$ and $m \in \{0,1\}^n$. For a fixed m, we have $u_1 \oplus u_2 = t_1 \oplus t_2$. Now, suppose the adversary can find a pair of tweaks (t_1, t_2) such that there is a collision at the output, i.e.,

$$(m'_1 = m'_2) \iff (\widehat{m}'_1 = \widehat{m}'_2) \iff (u'_1 \oplus u'_2 = t_1 \oplus t_2)$$

So, an output collision happens if and only if $u'_1 \oplus u'_2 = u_1 \oplus u_2$. Interestingly, for $\mathsf{TNT}_{\delta,m}$, we have the following property:

$$(\widehat{u}_1 \oplus \widehat{u}_2 = \delta) \implies (u'_1 \oplus u'_2 = u_1 \oplus u_2),$$

[1] These constructions are generally analyzed in the ideal cipher model.

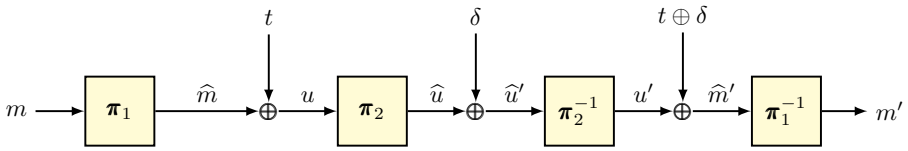

Fig. 1.2: TNT with final-block cancellation.

which implies that there are two sources of collisions in $\mathsf{TNT}_{\delta,m}$. A collision happens whenever $\widehat{u}_1 \oplus \widehat{u}_2 = \delta$, or $\widehat{u}_1 \oplus \widehat{u}_2 \neq \delta$ and $u'_1 \oplus u'_2 = u_1 \oplus u_2$. This indicates that one can expect more number of collisions (roughly double) in $\mathsf{TNT}_{\delta,m}$ as compared to a random function.

1.2 Contributions

Our contributions are threefold:

1. Birthday-bound CCA Attack on TNT: In section 3, we start by formalizing the aforementioned non-random behavior of TNT. We show (see section 3.1) that the expected number of output collisions for $\mathsf{TNT}_{\delta,m}$ is approximately twice the expected number for $\widetilde{\boldsymbol{\pi}}_{\delta,m}$, where $\widetilde{\boldsymbol{\pi}}$ is an n-bit uniform random permutation with n-bit tweaks. Our analysis strongly indicates a global non-random phenomenon that can be detected in roughly $O(2^{n/2})$ CCA queries. We establish this assertion by giving a fully scalable CCA distinguisher. We provide a rigorous analysis for the query complexity and advantage of our distinguisher, which shows that the distinguisher has an advantage expression of $1 - O(2^n/q^2)$, where q denotes the number of CCA queries. We provide details (see section 3.3) for efficient implementation and verification of our attacks, including results for an attack on TNT-GIFT-64, the TNT instantiation using GIFT-64 block cipher.
   Since the attack clearly contradicts the security claims of the designers of TNT, we study their security proof in section 4 and identify a bug, where a random variable is erroneously assumed to have a uniform distribution, leading to an overestimation of the security.
   See [27] and [26] for two alternative analyses of the attack. The former employs random permutations statistics to estimate the number of collisions and the latter directly bounds the probability of collisions in the two worlds. The analysis in this paper is more comprehensive and leads to a scalable advantage, but all three analyses come to the same conclusion: TNT can be broken in birthday bound queries!
2. Birthday-bound CCA Security of TNT: In section 5, we provide a simple proof of birthday-bound CCA security for TNT. Note that the CCA security bound also follows from the results in [44]. Nevertheless, given the flaws in TNT's original analysis, we believe that multiple security proofs using different techniques will lead to greater confidence in the revised security claim. In addition to the original TNT, we also analyze the single-keyed

variant of TNT, and show that it retains the same level of CCA security as well.

3. A Generalization of Cascaded LRW Paradigm: In a more abstract direction, in section 6, we present a generalized view of the cascaded LRW design strategy for any arbitrary number of rounds $r \geq 2$, called the LRW+ construction. It consists of two block cipher calls sandwiched between a pair of tweakable universal hashes. We show that as long as the tweakable hashes are sufficiently[2] universal, the LRW+ construction is CCA secure up to $2^{3n/4}$ queries. Note that LRW+ encompasses both 2-LRW2 and 4-LRW1. Thus, as a direct side-effect of our analysis, in section 6.2, we show that 2-LRW2 and 4-LRW1 are CCA secure up to $2^{3n/4}$ queries. In case of 2-LRW2, our bound matches the tight analysis in [25], and in case of 4-LRW1, our bound matches a concurrent result [15] by Datta et al.

Table 1.1: Summary of security bounds for LRW-based construction. We have assumed all hash functions to be 2^{-n}-(XOR) universal. The bottom four rows present our results. LRW+ generalizes both 2-LRW2 and 4-LRW1. So the bound on LRW+ implies similar bounds for 2-LRW2 and 4-LRW1.

Construction	BC calls	Hash calls	Security bound	Tightness
LRW1 [31]	1	0	$2^{n/2}$ (CPA) [31]	✓
LRW2 [31]	1	1	$2^{n/2}$ [31]	✓
3-LRW1 (TNT [18])	3	0	$2^{2n/3}$ [18]	(flawed)
4-LRW1	4	0	$2^{3n/4}$ [15]	–
2-LRW2 (CLRW2 [30])	2	2	$2^{3n/4}$ [25]	✓ [33]
r-LRW1 [44]	r odd	0	$2^{\frac{r-1}{r+1}n}$ [44]	–
	r even		$2^{\frac{r-2}{r}n}$	–
r-LRW2 [29]	r odd	r	$2^{\frac{r-1}{r+1}n}$ [29]	–
	r even		$2^{\frac{r}{r+2}n}$	–
3-LRW1 (TNT)	3	0	$2^{n/2}$	✓
1k-TNT	3	0	$2^{n/2}$	✓
LRW+	2	2	$2^{3n/4}$	–
4-LRW1	4	0	$2^{3n/4}$	–

Note that the result on LRW+ directly shows that r-LRW1 is at least $3n/4$-bit secure for any $r \geq 4$, improving on the results for $r \leq 8$. Similarly, for r-LRW2 it shows at least $3n/4$-bit security for any $r \geq 2$, improving on the

[2] Having approx. 2^{-n}-AU bound.

results for $r \leq 6$. See Table 1.1 for a summary of the state-of-the-art on the security of cascaded LRW constructions.

Comparison with [15]: Concurrently, Datta et al. also proposed [15] an improved bound for 4-LRW1 that matches our $2^{3n/4}$ bound. Both the proofs follow the proof strategy [25] used for 2-LRW2 by Jha and Nandi, although ours is in a more general form (analyzing LRW+) that applies to all the cascaded LRW constructions with two or more block cipher calls.

1.3 Impact of Our Birthday-bound Attack

As mentioned before, the authors of [4] claimed the CCA security of TNT to be $2n/3$ bits. In Asiacrypt 2020, the authors of [18] conjectured that the CCA security of TNT is probably $3n/4$ bits. In [44], the authors have stated:

> A natural open problem is the exact security of r-LRW1. Unlike r-LRW2, the exact security of r-LRW1 for $r = 3$ already appears challenging, and might require new proof approaches.

We believe this work answers a critical research question of both practical and theoretical implications. On one hand, it studies the exact security of an efficient construction that has several practical applications. On the other hand, it offers another cautionary tale on how to use statistical proof techniques such as the χ^2 method.[3]

Additionally, the attack applies to practical instances of TNT: TNT-AES in [4] and TNT-SM4-128 in [19]. The authors of [19] also introduced TNT-SM4-32, where the tweak size is limited to 32 bits. Our distinguisher requires $O(2^{n/2})$ tweaks, where $n = 128$ in case of TNT-SM4. Hence, the distinguisher directly applies to TNT-SM4-128, which has a tweak size of 128 bits. It does not directly apply to TNT-SM4-32, since the tweak space is too small. However, since our distinguisher breaks the BBB security proof in [4], the exact security of TNT-SM4-32 and whether it has BBB security is an open question.

We note that in Eurocrypt 2023, a full-round distinguisher on TNT-AES using truncated boomerang attacks was presented in [5]. However, the attack is particular to TNT-AES and requires almost 2^n queries. Our attack applied to any 128-bit instantiation of TNT, including TNT-AES, requires $\leq 2^{69}$ queries to have an almost 100% success rate, making it the best-known distinguisher for any 128-bit TNT variant, without relying on the properties of the underlying block cipher. We sum up all known distinguishers on TNT-AES in Table 1.2, which indicates that our distinguisher is not only theoretical but outperforms all cryptanalytic efforts on TNT, so far.

2 Preliminaries

Notational Setup: For $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, 2, \ldots, n\}$, $\{0,1\}^n$ denotes the set of bit strings of length n, and $\mathsf{Perm}(n)$ denotes the set of all per-

[3] Refer to [7] for another example of erroneously estimated distributions.

Table 1.2: Known distinguishers against TNT-AES. CCA stands for adaptive Chosen Ciphertext Adversary. NCPA stands for Non-adaptive Chosen Plaintext Adversary. **Rounds** is the number of AES rounds in $\boldsymbol{\pi}_1$, $\boldsymbol{\pi}_2$ and $\boldsymbol{\pi}_3$, respectively. $\star$ means any number of rounds. Generic attacks do not rely on any AES properties and apply to TNT instantiated with any 128-bit block cipher. 2^{69} is the complexity for which our attack is expected to have almost 100% success rate, while 2^{68} is expected to have 99% success rate.

Ref.	Type	Data	Time	Adversary	Rounds
[4]	Boomerang	2^{126}	2^{126}	CCA	$\star - 5 - \star$
[18]	Impossible Differential	$2^{113.6}$	$2^{113.6}$	NCPA	$5 - \star - \star$
[18]	Generic	$2^{99.5}$	$2^{99.5}$	NCPA	$\star - \star - \star$
[5]	Truncated Boomerang	2^{76}	2^{76}	CCA	$\star - 5 - \star$
[5]	Truncated Boomerang	2^{87}	2^{87}	CCA	$5 - 5 - \star$
[5]	Truncated Boomerang	$2^{127.8}$	$2^{127.8}$	CCA	$\star - 6 - \star$
This paper	Generic	$\leq 2^{69}$	$\leq 2^{69}$	CCA	$\star - \star - \star$

mutations over $\{0,1\}^n$. For $\tau, n \in \mathbb{N}$, $\widetilde{\mathsf{Perm}}(\tau, n)$ denotes the set of all families of permutations $\pi_t := \pi(t, \cdot) \in \mathsf{Perm}(n)$, indexed by $t \in \{0,1\}^\tau$. Any $\widetilde{\pi} \in \widetilde{\mathsf{Perm}}(\tau, n)$ is referred as a (τ, n)-tweakable permutation.

For $n, r \in \mathbb{N}$, such that $n \geq r$, we define the falling factorial $(n)_r := n!/(n-r)! = n(n-1)\cdots(n-r+1)$, and define $(n)_0 := 1$.

For $q \in \mathbb{N}$, x^q denotes the q-tuple $(x_1, x_2, \ldots, x_q)$, and in this context, $\mathtt{M}(x^q)$ and $\mathtt{S}(x^q)$ respectively denote the multiset and set corresponding to $\{x_i : i \in [q]\}$. For a set $\mathcal{I} \subseteq [q]$ and a q-tuple x^q, $x^{\mathcal{I}}$ denotes the tuple $(x_i)_{i \in \mathcal{I}}$. For a pair of tuples x^q and y^q, (x^q, y^q) denotes the 2-ary q-tuple $((x_1, y_1), \ldots, (x_q, y_q))$. An n-ary q-tuple is defined analogously. For $q \in \mathbb{N}$, for any set $\mathcal{X}$, $(\mathcal{X})_q$ denotes the set of all q-tuples with distinct elements from $\mathcal{X}$. For $q \in \mathbb{N}$, a 2-ary tuple (x^q, y^q) is called permutation compatible, denoted $x^q \nleftrightarrow y^q$, if $x_i = x_j \iff y_i = y_j$. Extending notations, a 3-ary tuple (t^q, x^q, y^q) is called tweakable permutation compatible, denoted by $(t^q, x^q) \nleftrightarrow (t^q, y^q)$, if $(t_i, x_i) = (t_j, x_j) \iff (t_i, y_i) = (t_j, y_j)$. For any tuple $x^q \in \mathcal{X}^q$, and for any function $f : \mathcal{X} \to \mathcal{Y}$, $f(x^q)$ denotes the tuple $(f(x_1), \ldots, f(x_q))$. We use shorthand notation $\exists^*$ to represent the phrase "there exists distinct".

Unless stated otherwise, upper and lower case letters denote variables and values, respectively, and Serif font letters are used to denote random variables. For a finite set $\mathcal{X}$, $\mathsf{X} \leftarrow\!\$\, \mathcal{X}$ denotes the uniform and random sampling of X from $\mathcal{X}$. We write $\mathsf{X}^q \xleftarrow{\text{wor}} \mathcal{X}$ to denote WOR (without replacement sampling) of a q-tuple X^q from the set $\mathcal{X}$, where $|\mathcal{X}| \geq q$ is obvious. More precisely, $\mathsf{X}^q \leftarrow\!\$\, (\mathcal{X})_q$.

We will use the following proposition, which is a slight variation of [17, Lemma 6].

Proposition 2.1. *Let* R_0 *and* R_1 *be two random variables with variances* σ_0^2 *and* σ_1^2*, respectively, and suppose their expectations follow the relation* $\mathsf{Ex}(\mathsf{R}_0) \geq$

$\mu_0 \geq \mu_1 \geq \mathsf{Ex}(\mathsf{R}_1)$, *for some* $\mu_0 \geq \mu_1 \geq 0$. *Then, for* $\mu = (\mu_0 + \mu_1)/2$, *we have*

$$|\mathsf{Pr}(\mathsf{R}_0 > \mu) - \mathsf{Pr}(\mathsf{R}_1 > \mu)| \geq 1 - \frac{4(\sigma_0^2 + \sigma_1^2)}{(\mu_0 - \mu_1)^2}.$$

When $\mathsf{Ex}(\mathsf{R}_0) = \mu_0$ and $\mathsf{Ex}(\mathsf{R}_1) = \mu_1$, we get back [17, Lemma 6]. A proof of this proposition can be derived using a similar approach as used in the proof of [17, Lemma 6]. We provide a short alternate proof in Supplementary Material A by using the Bienaymé-Chebyshev inequality.

2.1 (Tweakable) Block Ciphers and Random Permutations

A (κ, n)-block cipher with key size κ and block size n is a family of permutations $E \in \widetilde{\mathsf{Perm}}(\kappa, n)$. For $k \in \{0,1\}^\kappa$, we denote $E_k(\cdot) := E(k, \cdot)$, and $E_k^{-1}(\cdot) := E^{-1}(k, \cdot)$. A (κ, τ, n)tweakable block cipher with key size κ, tweak size τ and block size n is a family of permutations $\widetilde{E} \in \widetilde{\mathsf{Perm}}((\kappa, \tau), n)$. For $k \in \{0,1\}^\kappa$ and $t \in \{0,1\}^\tau$, we denote $\widetilde{E}_k(t, \cdot) := \widetilde{E}(k, t, \cdot)$, and $\widetilde{E}_k^{-1}(t, \cdot) := \widetilde{E}^{-1}(k, t, \cdot)$. Throughout this paper, we fix $\kappa, \tau, n \in \mathbb{N}$ as the key size, tweak size, and block size, respectively, of the given (tweakable) block cipher.

We say that $\boldsymbol{\pi}$ is an (ideal) random permutation on block space $\{0,1\}^n$ to indicate that $\boldsymbol{\pi} \leftarrow\!\!\$\ \mathsf{Perm}(n)$. Similarly, we say that $\widetilde{\boldsymbol{\pi}}$ is an (ideal) tweakable random permutation on tweak space $\{0,1\}^\tau$ and block space $\{0,1\}^n$ to indicate that $\widetilde{\boldsymbol{\pi}} \leftarrow\!\!\$\ \widetilde{\mathsf{Perm}}(\tau, n)$.

2.2 Security Definition

In this paper, we assume that the distinguisher is non-trivial, i.e. it never makes a duplicate query, and it never makes a query for which the response is already known due to some previous query. Let $\mathbb{A}(q, t)$ be the class of all non-trivial distinguishers limited to q oracle queries, and t computations.

In our analyses, especially security proofs, it will be convenient to work in the information-theoretic setting. Accordingly, we always skip the boilerplate hybrid steps and often assume that the adversary is computationally unbounded, i.e., $t = \infty$, and deterministic. A computational equivalent of all our security proofs can be easily obtained by a simple hybrid argument.

IND-CCA Security: The IND-CCA advantage of distinguisher $\mathbf{A}$ against $\widetilde{E}$ instantiated with a key $\mathsf{K} \leftarrow\!\!\$\ \{0,1\}^\kappa$ is defined as

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\widetilde{E}}(\mathbf{A}) = \mathbf{Adv}_{\widetilde{E}^{\pm}; \widetilde{\boldsymbol{\pi}}^{\pm}}(\mathbf{A}) := \left|\mathsf{Pr}\left(\mathbf{A}(\widetilde{E}_{\mathsf{K}}^{\pm}) = 1\right) - \mathsf{Pr}\left(\mathbf{A}(\widetilde{\boldsymbol{\pi}}^{\pm}) = 1\right)\right|. \quad (1)$$

The IND-CCA security of $\widetilde{E}$ is defined as

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\widetilde{E}}(q, t) := \max_{\mathbf{A} \in \mathbb{A}(q,t)} \mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\widetilde{E}}(\mathbf{A}).$$

2.3 The Expectation Method

Let $\mathbf{A}$ be a computationally unbounded and deterministic distinguisher that tries to distinguish between two oracles $\mathcal{O}_0$ and $\mathcal{O}_1$ via black box interaction with one of them. We denote the query-response tuple of $\mathbf{A}$'s interaction with its oracle by a transcript ω. This may also include any additional information the oracle chooses to reveal to the distinguisher at the end of the query-response phase of the game. We denote by Θ_1 (res. Θ_0) the random transcript variable when $\mathbf{A}$ interacts with $\mathcal{O}_1$ (res. $\mathcal{O}_0$). The probability of realizing a given transcript ω in the security game with an oracle $\mathcal{O}$ is known as the *interpolation probability* of ω with respect to $\mathcal{O}$. Since $\mathbf{A}$ is deterministic, this probability depends only on the oracle $\mathcal{O}$ and the transcript ω. A transcript ω is said to be *attainable* if $\mathsf{Pr}\left(\Theta_0 = \omega\right) > 0$. The expectation method [20] (stated below) is a generalization of Patarin's H-coefficients technique [37], which is quite useful in obtaining improved bounds in many cases [20,25].

Lemma 2.1 (Expectation Method [20]). *Let Ω be the set of all transcripts. For some $\epsilon_{\mathsf{bad}} \geq 0$ and a non-negative function $\epsilon_{\mathsf{ratio}} : \Omega \to [0, \infty)$, suppose there is a set $\Omega_{\mathsf{bad}} \subseteq \Omega$ satisfying the following:*

- $\mathsf{Pr}\left(\Theta_0 \in \Omega_{\mathsf{bad}}\right) \leq \epsilon_{\mathsf{bad}}$;
- *For any $\omega \notin \Omega_{\mathsf{bad}}$, ω is attainable and* $\dfrac{\mathsf{Pr}\left(\Theta_1 = \omega\right)}{\mathsf{Pr}\left(\Theta_0 = \omega\right)} \geq 1 - \epsilon_{\mathsf{ratio}}(\omega)$.

Then for any distinguisher $\mathbf{A}$ trying to distinguish between $\mathcal{O}_1$ and $\mathcal{O}_0$, we have the following bound on its distinguishing advantage:

$$\mathbf{Adv}_{\mathcal{O}_1;\mathcal{O}_0}(\mathbf{A}) \leq \epsilon_{\mathsf{bad}} + \mathsf{Ex}\left(\epsilon_{\mathsf{ratio}}(\Theta_0)\right).$$

When $\epsilon_{\mathsf{ratio}}$ is a constant function, we get the H-coefficients technique.

3 Birthday-bound Attack on TNT

We consider the TNT construction in an information-theoretic setting. Accordingly, we instantiate TNT based on three independent uniform random permutations $\boldsymbol{\pi}_1$, $\boldsymbol{\pi}_2$, and $\boldsymbol{\pi}_3$ of $\{0,1\}^n$. Recall that, the TNT construction is defined by the mapping

$$(t, m) \xmapsto{\mathsf{TNT}} \boldsymbol{\pi}_3(t \oplus \boldsymbol{\pi}_2(t \oplus \boldsymbol{\pi}_1(m))), \tag{2}$$

For some non-zero $\delta \in \{0,1\}^n$ and $m \in \{0,1\}^n$, consider the function $\mathcal{O}_{\delta,m} : \{0,1\}^n \to \{0,1\}^n$, associated to each n-bit tweakable permutation $\mathcal{O}$ with n-bit tweak, defined by the mapping

$$t \xmapsto{\mathcal{O}_{\delta,m}} \mathcal{O}^{-1}(t \oplus \delta,\ \mathcal{O}(t,\ m)). \tag{3}$$

We are only interested in $\widetilde{\boldsymbol{\pi}}_{\delta,m}$ and $\mathsf{TNT}_{\delta,m}$ where $\widetilde{\boldsymbol{\pi}}$ is a tweakable uniform random permutation of $\{0,1\}^n$ with n-bit tweaks.

Suppose $\widetilde{\boldsymbol{\pi}}_{\delta,m}$ is executed over q distinct inputs $(t_1, \ldots, t_q)$. Observe that, for any valid choice of $(t_1, \ldots, t_q)$, $\widetilde{\boldsymbol{\pi}}$ is executed at most twice for any tweak t_i.

Thus, one can expect $\widetilde{\boldsymbol{\pi}}_{\delta,m}(\cdot)$ to be almost uniform and independent, and thus, indistinguishable from a uniform random function $\boldsymbol{\rho} : \{0,1\}^n \to \{0,1\}^n$ for a large range of q. In fact, as long as

$$\widetilde{\boldsymbol{\pi}}(t_i, m) \neq \widetilde{\boldsymbol{\pi}}(t_j, m) \text{ for all } i \neq j \text{ such that } t_j = t_i \oplus \delta,$$

$\widetilde{\boldsymbol{\pi}}_{\delta,m}$ can be shown to be indistinguishable from $\boldsymbol{\rho}$ up to $O(2^n)$ queries. More importantly, as we show in the following discussion, one can easily show that the $\widetilde{\boldsymbol{\pi}}_{\delta,m}$ is almost identical to $\boldsymbol{\rho}$ in terms of the number of output collisions.

$\mathsf{TNT}_{\delta,m}$, on the other hand, exhibits a rather peculiar and interesting property. Apparently, $\mathsf{TNT}_{\delta,m}$ is more prone to collisions as compared to $\widetilde{\boldsymbol{\pi}}_{\delta,m}$, which results in a direct IND-CCA distinguisher for TNT. A formal distinguisher with complete advantage calculation appears later in section 3.2. We first demonstrate the biased behavior by comparing the number of output collisions for $\mathsf{TNT}_{\delta,m}$ and $\widetilde{\boldsymbol{\pi}}_{\delta,m}$.

3.1 Comparing the Number of Collision Pairs in $\widetilde{\boldsymbol{\pi}}_{\delta,m}$ and $\mathsf{TNT}_{\delta,m}$

Fix some non-negative integer $q \leq 2^n$. Fix a set $\mathcal{T} = \{t_1, \ldots, t_q\} \subseteq \{0,1\}^n$ of size q, an $m \in \{0,1\}^n$, and a non-zero $\delta \in \{0,1\}^n$. Let $\mathcal{O}$ be a tweakable permutation (which is either $\widetilde{\boldsymbol{\pi}}$ in the ideal world or TNT in the real world). We compute $\mathsf{M}'_i = \mathcal{O}_{\delta,m}(t_i)$ by making a forward query $\mathcal{O}(t_i, m) := \widehat{\mathsf{C}}_i$, followed by a backward query $\mathsf{M}'_i = \mathcal{O}^{-1}(t_i \oplus \delta, \widehat{\mathsf{C}}_i)$. We write $\mathtt{COLL}(\mathcal{O}_{\delta,m})$ to denote the number of pairs (i,j), $i < j$ such that $\mathsf{M}'_i = \mathsf{M}'_j$.

Analyzing $\mathsf{coll}_{\mathrm{id}} := \mathtt{COLL}(\widetilde{\boldsymbol{\pi}}_{\delta,m})$: For any $i \neq j \in [q]$, let $\mathbb{1}_{i,j}$ denote the indicator random variable corresponding to the event: $\mathsf{M}'_j = \mathsf{M}'_i$. Then, using linearity of expectation, we have

$$\mathsf{Ex}\left(\mathsf{coll}_{\mathrm{id}}\right) = \sum_{i<j\in[q]} \mathsf{Ex}\left(\mathbb{1}_{i,j}\right) = \sum_{i<j\in[q]} \mathsf{Pr}\left(\mathbb{1}_{i,j}\right), \tag{4}$$

where we abused the notation slightly to use $\mathbb{1}_{i,j}$ to denote the event $\mathbb{1}_{i,j} = 1$. Let $\sim$ be a relation on $[q]$, such that for all $i \neq j \in [q]$, $i \sim j$ if and only if $t_i = t_j \oplus \delta$. Note that $\sim$ is symmetric. Suppose there are ν pairs (t_i, t_j), $i < j$ such that $t_i \sim t_j$. Clearly, $\nu \leq q/2$. Now, we can split the right-hand side of (4) as follows:

$$\sum_{i<j\in[q]} \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) = \sum_{\substack{i<j\in[q] \\ i\sim j}} \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) + \sum_{\substack{i<j\in[q] \\ i\not\sim j}} \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) \tag{5}$$

Case $i \not\sim j$: We must have $\{t_i, t_j\} \cap \{t_i \oplus \delta, t_j \oplus \delta\} = \emptyset$. Thus, the two calls to $\widetilde{\boldsymbol{\pi}}_{\delta,m}$ corresponding to the i-th and j-th queries result in exactly 2 calls to $\widetilde{\boldsymbol{\pi}}$ and 2 calls $\widetilde{\boldsymbol{\pi}}^{-1}$, each with a distinct tweak than others. Hence, the outputs of $\widetilde{\boldsymbol{\pi}}_{\delta,m}$ on inputs t_i and t_j are mutually independent and uniformly distributed in $\{0,1\}^n$. Thus, for any $i \not\sim j$, we have

$$\mathsf{Pr}\left(\mathbb{1}_{i,j}\right) = \frac{1}{2^n}, \tag{6}$$

which results in

$$\sum_{\substack{i<j\in[q]\\ i\not\sim j}} \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) = \left(\binom{q}{2} - \nu\right)\frac{1}{2^n}, \tag{7}$$

Case $i \sim j$: In this case we have $t_i = t_j \oplus \delta$. Let $\mathtt{F}_{i,j}$ be the event that $\widetilde{\boldsymbol{\pi}}(t_i, m) = \widetilde{\boldsymbol{\pi}}(t_j, m)$. Then, we have $\mathsf{M}'_i = \mathsf{M}'_j = m$. Since, $t_i \neq t_j$, $\mathsf{Pr}\left(\mathtt{F}_{i,j}\right) = 2^{-n}$. So, for any $i \sim j$, we have

$$\begin{aligned}\mathsf{Pr}\left(\mathbb{1}_{i,j}\right) &= \mathsf{Pr}\left(\mathbb{1}_{i,j} \wedge \mathtt{F}_{i,j}\right) + \mathsf{Pr}\left(\mathbb{1}_{i,j} \wedge \neg\mathtt{F}_{i,j}\right)\\ &= \mathsf{Pr}\left(\mathtt{F}_{i,j}\right) + \mathsf{Pr}\left(\mathbb{1}_{i,j} \wedge \neg\mathtt{F}_{i,j}\right)\\ &= \frac{1}{2^n} + \mathsf{Pr}\left(\mathbb{1}_{i,j} \wedge \neg\mathtt{F}_{i,j}\right),\end{aligned}$$

which immediately gives

$$\frac{1}{2^n} \leq \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) \leq \frac{1}{2^n} + \mathsf{Pr}\left(\mathbb{1}_{i,j} \mid \neg\mathtt{F}_{i,j}\right) \leq \frac{1}{2^n} + \frac{1}{2^n - 1}. \tag{8}$$

Note that the last inequality follows from the observation that given $\neg\mathtt{F}_{i,j}$, outputs of $\widetilde{\boldsymbol{\pi}}^{-1}(t_i \oplus \delta)$ and $\widetilde{\boldsymbol{\pi}}^{-1}(t_j \oplus \delta)$ are sampled independently from a set of size exactly $2^n - 1$. This further results in

$$\frac{\nu}{2^n} \leq \sum_{\substack{i<j\in[q]\\ i\sim j}} \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) \leq \nu\left(\frac{1}{2^n} + \frac{1}{2^n - 1}\right). \tag{9}$$

Using (4), (5), (7), (9), and $\nu \leq q/2$ we have

$$\binom{q}{2}\frac{1}{2^n} \leq \mathsf{Ex}\left(\mathsf{coll}_{\mathrm{id}}\right) \leq \binom{q}{2}\frac{1}{2^n} + \frac{q}{2^n}. \tag{10}$$

Analyzing $\mathsf{coll}_{\mathrm{re}} := \mathtt{COLL}(\mathsf{TNT}_{\delta,m})$: The analysis of $\mathtt{COLL}(\mathsf{TNT}_{\delta,m})$ is a bit more subtle and interesting. Fig. 3.1 gives a pictorial view of the i-th execution of $\mathsf{TNT}_{\delta,m}$. Clearly, the respective calls to $\boldsymbol{\pi}_3$ and its inverse cancel out each

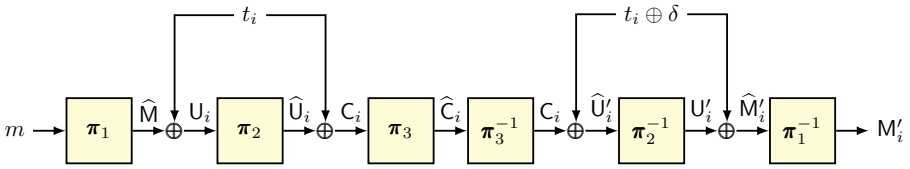

Fig. 3.1: The execution trace for $\mathsf{TNT}_{\delta,m}$ on input t_i.

other, resulting in the compressed view illustrated in Fig. 3.2.

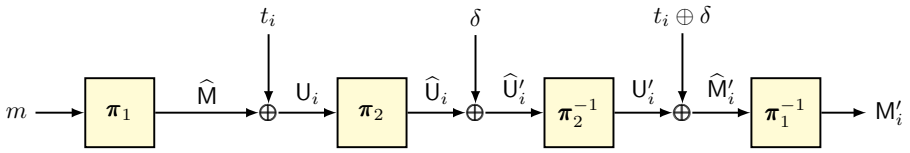

Fig. 3.2: The effective execution trace for $\mathsf{TNT}_{\delta,m}$ on input t_i.

Note that for any $i, j \in [q]$, $\mathsf{U}_i \oplus \mathsf{U}_j = t_i \oplus t_j$. Now, fix a pair of inputs (t_i, t_j) such that there is a collision at the output, i.e.,

$$(\mathsf{M}'_i = \mathsf{M}'_j) \iff (\widehat{\mathsf{M}}'_i = \widehat{\mathsf{M}}'_j) \iff (\mathsf{U}'_i \oplus \mathsf{U}'_j = t_i \oplus t_j) \iff (\mathsf{U}'_i \oplus \mathsf{U}'_j = \mathsf{U}_i \oplus \mathsf{U}_j),$$

and let $\mathbb{1}_{i,j}$ denote the corresponding indicator random variable. Observe that $\mathsf{TNT}_{\delta,m}$, has the following interesting property:

$$(\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j = \delta) \implies (\mathsf{U}'_i \oplus \mathsf{U}'_j = \mathsf{U}_i \oplus \mathsf{U}_j = t_i \oplus t_j),$$

which implies that there are two sources of collisions in $\mathsf{TNT}_{\delta,m}$. A collision happens whenever

1. $\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j = \delta$, or
2. $\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j \neq \delta$ and $\mathsf{U}'_i \oplus \mathsf{U}'_j = t_i \oplus t_j$.

From this one can easily get a good upper and lower bound on the expected number of collisions in the real world. Using linearity of expectation, we have

$$\mathsf{Ex}\left(\mathsf{coll}_{\mathrm{re}}\right) = \sum_{i<j\in[q]} \mathsf{Ex}\left(\mathbb{1}_{i,j}\right) = \sum_{i<j\in[q]} \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) \tag{11}$$

Further, from the above discussion, we have

$$\begin{aligned}
\mathsf{Pr}\left(\mathbb{1}_{i,j}\right) &= \mathsf{Pr}\left(\mathbb{1}_{i,j} \wedge \widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j = \delta\right) + \mathsf{Pr}\left(\mathbb{1}_{i,j} \wedge \widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j \neq \delta\right) \\
&= \mathsf{Pr}\left(\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j = \delta\right) + \mathsf{Pr}\left(\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j \neq \delta\right) \\
&\qquad \times \mathsf{Pr}\left(\mathsf{U}'_i \oplus \mathsf{U}'_j = t_i \oplus t_j \mid \widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j \neq \delta\right) \\
&= \frac{1}{2^n - 1} + \left(1 - \frac{1}{2^n - 1}\right) \\
&\qquad \times \mathsf{Pr}\left(\mathsf{U}'_i \oplus \mathsf{U}'_j = t_i \oplus t_j \mid \widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j \neq \delta\right),
\end{aligned} \tag{12}$$

Note that $\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j \neq \delta$ implies that $\mathsf{U}'_i, \mathsf{U}'_j \notin \{\mathsf{U}_i, \mathsf{U}_j\}$. Now, fix a valid choice for $(\mathsf{U}_i, \mathsf{U}_j, \widehat{\mathsf{U}}_i, \widehat{\mathsf{U}}_j)$, say $(u_i, u_j, \widehat{u}_i, \widehat{u}_j)$. Then, the number of valid choices for $(\mathsf{U}'_i, \mathsf{U}'_j)$ that satisfy the equation $\mathsf{U}'_i \oplus \mathsf{U}'_j = t_i \oplus t_j$, are all $(x, x \oplus t_i \oplus t_j)$ pairs such that

$$x \in \{0,1\}^n \setminus (\{u_i, u_j\} \cup \{u_i \oplus t_i \oplus t_j, u_j \oplus t_i \oplus t_j\})$$

But, observe that $\{u_i, u_j\} = \{u_i \oplus t_i \oplus t_j, u_j \oplus t_i \oplus t_j\}$ by definition, for any valid choice of (u_i, u_j). Therefore, the number of valid $(x, x \oplus t_i \oplus t_j)$ is exactly $2^n - 2$.

Furthermore, this counting is independent of the choice of $(u_i, u_j, \widehat{u}_i, \widehat{u}_j)$, whence it holds unconditionally. Now, each such choice for $(\mathsf{U}'_i, \mathsf{U}'_j)$ occurs with at most $1/(2^n-2)(2^n-3)$ probability, as they are sampled from $\{0,1\}^n \setminus \{\mathsf{U}_i, \mathsf{U}_j\}$ in a WOR (without replacement) manner. Then, using (12), we have

$$\begin{aligned}\mathsf{Pr}\left(\mathbb{1}_{i,j}\right) &= \frac{1}{2^n-1} + \left(1 - \frac{1}{2^n-1}\right) \times \frac{1}{2^n-3} \\ &= \frac{1}{2^n-1} + \frac{1}{2^n-3} - \frac{1}{(2^n-1)(2^n-3)} \\ &= \frac{2}{2^n} + \frac{1}{2^n(2^n-1)} + \frac{3}{2^n(2^n-3)} - \frac{1}{(2^n-1)(2^n-3)}\end{aligned}$$

Using (11), we immediately have

$$\mathsf{Ex}\left(\mathsf{coll}_{\mathrm{re}}\right) = \binom{q}{2}\left(\frac{1}{2^n-1} + \frac{1}{2^n-3} - \frac{1}{(2^n-1)(2^n-3)}\right) \geq \binom{q}{2}\frac{2}{2^n}, \tag{13}$$

and on comparing this with (10), we can conclude that

$$\mathsf{Ex}\left(\mathsf{coll}_{\mathrm{re}}\right) \approx 2\mathsf{Ex}\left(\mathsf{coll}_{\mathrm{id}}\right).$$

This clearly indicates that the occurrence of collisions in $\mathsf{TNT}_{\delta,m}$ is approximately twice that of $\widetilde{\boldsymbol{\pi}}_{\delta,m}$.

3.2 The Collision Counting Distinguisher

Based on the observations from the preceding section, we now present a formal distinguisher, called $\mathbf{A}^*$, in Algorithm 1.

Fix a message $m \in \{0,1\}^n$, a set $\mathcal{T} = \{t_1, \ldots, t_q\} \subseteq \{0,1\}^n$ of size q, and a $\delta \neq 0^n$. Let $\theta(q,n)$ be some non-negative function of q and n, which will be defined later in the course of analysis.

Let $\mathcal{O}^{\pm}$ be the oracle $\mathbf{A}^*$ is interacting with. Then, $\mathbf{A}^*$ works by collecting $\mathsf{M}'_i = \mathcal{O}_{\delta,m}(t_i)$ for all $t_i \in \mathcal{T}$ in a multiset $\mathcal{M}$. As shown in the preceding section and Algorithm 1, this can be easily done by a pair of encryption-decryption queries for each $i \in [q]$. After this, $\mathbf{A}^*$ counts the number of collisions in $\mathcal{M}$ using the function `collCount`. If the number of collisions is greater than $\theta(q,n)$, the distinguisher returns 1, otherwise, it returns 0.

Note that the exact implementation of `collCount` is not relevant for the forthcoming advantage calculation. So, we postpone a discussion on its implementation and resulting time and space complexity analysis to section 3.3, where we also provide experimental verification for $\mathbf{A}^*$.

However, it is amply evident that the space complexity of the attack is $O(q)$, i.e., dominated by the query complexity. Further, looking ahead momentarily, one can implement `collCount` in such a way that it runs in time $O(q \log_2 q)$. Other than this, $\mathbf{A}^*$ only makes $2q$ calls to $\mathcal{O}$, thus the overall time complexity is also in $O(q \log_2 q)$.

Algorithm 1 Algorithmic description of $\mathbf{A}^*(\mathcal{O}^{\pm})$. Note that `collCount` is an abstract function that counts the number of collisions in a multiset.

$m \leftarrow 0^n$ ▷ m can be initialized to any constant
$\delta \leftarrow 1^n$ ▷ δ can be initialized to any non-zero constant
$\mathcal{T} \leftarrow \{t_1, \ldots, t_q\}$ ▷ a set of q fixed but distinct tweaks
$\mathcal{M} \leftarrow \emptyset$ ▷ an empty multiset
for $i = 1 \ldots q$ **do**
    $\widehat{\mathsf{C}}_i \leftarrow \mathcal{O}(t_i, m)$
    $\mathsf{M}'_i \leftarrow \mathcal{O}^{-1}(t_i \oplus \delta, \widehat{\mathsf{C}}_i)$ ▷ ln. 6 and 7 together give $\mathcal{O}_{\delta,m}(t_i)$
    $\mathcal{M} \leftarrow \mathcal{M} \cup \{\mathsf{M}'_i\}$
$\mathtt{COLL}(\mathcal{O}_{\delta,m}) \leftarrow \mathtt{collCount}(\mathcal{M})$
if $\mathtt{COLL}(\mathcal{O}_{\delta,m}) > \theta(q,n)$ **then**
    return 1
else
    return 0

Define

$$\mu_{\mathrm{re}} := \binom{q}{2}\frac{2}{2^n} \qquad \mu_{\mathrm{id}} := \binom{q}{2}\frac{1}{2^n} + \frac{q}{2^n}.$$

Then, from (10) and (13), we have that $\mathsf{Ex}\left(\mathtt{COLL}(\mathsf{TNT}_{\delta,m})\right) \geq \mu_{\mathrm{re}} \geq \mu_{\mathrm{id}} \geq \mathsf{Ex}\left(\mathtt{COLL}(\widetilde{\boldsymbol{\pi}}_{\delta,m})\right)$, whenever $q \geq 3$.

Theorem 3.1. *For $n \geq 4$, $10 \leq q \leq 2^n$, and $\theta(q,n) = (\mu_{\mathrm{re}} + \mu_{\mathrm{id}})/2$, we have*

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\mathit{TNT}}(\mathbf{A}^*) \geq 1 - 371\frac{2^n}{q^2}.$$

Specifically, for $q \geq 28 \times 2^{\frac{n}{2}}$, $\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\mathit{TNT}}(\mathbf{A}^*) \geq 0.5$.

Proof. Recall that $\mathsf{coll}_{\mathrm{id}} = \mathtt{COLL}(\widetilde{\boldsymbol{\pi}}_{\delta,m})$ and $\mathsf{coll}_{\mathrm{re}} = \mathtt{COLL}(\widetilde{\boldsymbol{\pi}}_{\delta,m})$. Let $\sigma_s^2 := \mathsf{Var}\left(\mathsf{coll}_s\right)$, for all $s \in \{\mathrm{id}, \mathrm{re}\}$. In addition, whenever necessary, we also reuse the notations and definitions from the expectation calculation given in section 3.1. Now, we have

$$\begin{aligned}
\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\mathsf{TNT}}(\mathbf{A}^*) &= \left|\mathsf{Pr}\left(\mathbf{A}^*(\mathsf{TNT}_{\delta,m}) = 1\right) - \mathsf{Pr}\left(\mathbf{A}^*(\widetilde{\boldsymbol{\pi}}_{\delta,m}) = 1\right)\right| \\
&= \left|\mathsf{Pr}\left(\mathsf{coll}_{\mathrm{re}} > \theta(q,n)\right) - \mathsf{Pr}\left(\mathsf{coll}_{\mathrm{id}} > \theta(q,n)\right)\right| \\
&\geq 1 - \frac{4(\sigma^2_{\mathrm{re}} + \sigma^2_{\mathrm{id}})}{(\mu_{\mathrm{re}} - \mu_{\mathrm{id}})^2}. \qquad (14)
\end{aligned}$$

where the last inequality follows from Proposition 2.1. We make the following claim on σ^2_{re} and σ^2_{id}.

Claim 3.1. For $n \geq 4$, $10 \leq q \leq 2^n$, we have

$$\sigma^2_{\mathrm{id}} \leq \frac{4q^2}{2^n} \qquad \sigma^2_{\mathrm{re}} \leq \frac{11q^2}{2^n}$$

A proof of this claim is available in Supplementary Material C. Next, from (10) and (13), we have

$$(\mu_{\text{re}} - \mu_{\text{id}})^2 \geq \left(\binom{q}{2} \frac{2}{2^n} - \binom{q}{2} \frac{1}{2^n} - \frac{q}{2^n} \right)^2$$
$$\geq \binom{q}{2}^2 \frac{1}{2^{2n}} \left(1 - \frac{1}{q}\right)^2 \geq 0.162 \frac{q^4}{2^{2n}} \tag{15}$$

where the last inequality follows from $q \geq 10$. The result then follows from (14), Claim 3.1, and (15). □

Remark 3.1. Note that the constant in Theorem 3.1 is a bit loose for the sake of simplicity. It is likely that this constant can be improved by a more tighter estimation or a more sophisticated concentration inequality. Indeed, in the next section, we show that in practical applications the advantage might already be close to 0.8 when the number of queries is close to $4 \times 2^{\frac{n}{2}}$.

With that being said, it's important to highlight that our attack demonstrates full scalability. In other words, as the value of q approaches 2^n, the advantage becomes close to 1.

3.3 Experimental Verification

We have implemented the collision counting Algorithm 1 for different values for n. We have implemented two variants of the `collCount` function of the algorithm, which include various optimizations to make the attack practical. The first variant is an adversary without space complexity and with time complexity $O(q)$, and is given in Algorithm 2. The second is for a space-optimized adversary, with space complexity $O(q)$ and time complexity $O(q \log_2(q))$, described in Algorithm 3. For the underlying random permutations, we used generated using Python NumPy's `shuffle` and `argsort` functions, to generate and invert a permutation, respectively. We generated permutations of sizes 16, 20, 24, 28 and 32 bits and performed the distinguishing attack on each generated permutation. Results where taken over an average of $1,000 \sim 10,000$ random generations (each consisting of 3 independent permutations). In the ideal world, random values are sampled, since the tweaks are never repeated and lazy sampling can be used. Table 3.1 includes the average number of collisions for $n = 16$ and $n = 20$. The distinguisher reaches 16 expected collisions in the real world $4\times$ faster than the distinguisher in [18] for $n = 16$ and $16\times$ faster for $n = 20$.

Algorithm 1 is expected to have twice as many collisions in the real world as in the ideal world. $\theta(q, n)$ is set to:

$$\theta(q, n) = 2^{2d-1} + 2^{2d-2}$$

when $q = 2^{n/2+d}$, which is roughly 1.5 times the expected number of collisions in the ideal case.

Algorithm 2 An implementation of $\texttt{collCount}(\mathcal{M})$ from Algorithm 1 with no memory limitations. Here, $\mathcal{M}$ (the multiset of outputs) is assumed to be an array of size q.

```
1: for x ∈ {0,1}^n do
2:     L[x] ← 0
3: coll ← 0
4: for i ∈ {1,...,q−1} do
5:     x ← M[i]
6:     coll ← coll + L[x]
7:     L[x] ← L[x] + 1
```

Algorithm 3 An implementation of $\texttt{collCount}(\mathcal{M})$ from Algorithm 1 with memory limited to $O(q)$. Here, $\mathcal{M}$ (the multiset of outputs) is assumed to be an array of size q.

```
1: M ← sort(M)
2: rep ← 1
3: coll ← 0
4: x ← M[1]
5: for i ∈ {2,...,q} do
6:     if x = M[i] then
7:         coll ← coll + rep
8:         rep ← rep + 1
9:     else
10:        rep ← 1
11:        x ← M[i]
```

Table 3.1: Average number of collisions using random permutations.

n	16					
$\log_2(q)$	6	7	8	9	10	11
real	0.06	0.27	0.96	3.72	15.62	63.59
ideal	0.023	0.12	0.48	1.98	7.91	31.17
n	20					
$\log_2(q)$	8	9	10	11	12	13
real	0.073	0.203	1.02	4.01	15.69	63.63
ideal	0.023	0.11	0.47	1.94	7.92	32.57

We also calculated the success rate, which is the number of successful distinguishing attempts over the total number of attempts, for different values of q and $\theta(q,n)$. This is equivalent to the advantage in Theorem 3.1. Table 3.2 shows the success rate for the different parameters. The distinguisher reaches $\geq 85\%$ with $q = 2^{n/2+2}$ and 99% success rate with $q = 2^{n/2+3}$. The attack complexities are $2^{n/2+3}$ and $2^{n/2+4}$, respectively, since each iteration includes two queries to the construction. For large n, the factors 2^3 and 2^4 are small. With complexity $2^{n/2+5}$, we get a success rate of almost 100%, and an attack that breaks the security claim for In practice, $n \geq 64$. The complexity of the distinguisher is compared to known TNT distinguishers with $n = 128$ in Table 1.2.

Note that our experimental estimations closely match the advantage curve obtained through theoretical analysis, up to a change in constant. In fact, we get

a more optimistic constant in experimental results. In particular, we estimate that the advantage is around

$$1 - 2\frac{2^n}{q^2},$$

but the discrepancy is expected since the theoretical advantage is more conservative and bound to be a bit loose for the sake of simplicity.

We also calculated the success rate, which is the number of successful distinguishing attempts over the total number of attempts, for different values of q and $\theta(q, n)$. Table 3.2 shows the success rate for the different parameters. The distinguisher reaches $\geq 85\%$ with $q = 2^{n/2+2}$ and 99% success rate with $q = 2^{n/2+3}$. The attack complexities are $2^{n/2+3}$ and $2^{n/2+4}$, respectively, since each iteration includes two queries to the construction. For large n, the factors 2^3 and 2^4 are small. With complexity $2^{n/2+5}$, we get a success rate of almost 100%, and an attack that breaks the security claim for In practice, $n \geq 64$. The complexity of the distinguisher is compared to known TNT distinguishers with $n = 128$ in Table 1.2.

Table 3.2: The success rate achieved for different values of n and q.

n	q	$\theta(q,n)$	**Success Rate**	q	$\theta(q,n)$	**Success Rate**
16	10	12	87.2%	11	48	99%
20	12	12	86.6%	13	48	99%
24	14	12	90%	15	48	99%
28	16	12	85%	17	48	99%
32	18	12	87.5%	19	48	99%

On the Time-Memory Trade-off. Algorithm 2 runs in time $O(q)$, with space complexity $O(2^n)$. This is sufficient and provides optimal time complexity for information-theoretic (unbounded) adversaries. On the other hand, Algorithm 3 is more geared towards linear space complexity. Its time complexity is dominated by the `sort` function, which can be executed with time complexity $O(q \log_2(q))$ using merge-sort. The space complexity is dominated by the size of the list L which is $O(q)$.

In practice, while we assume that the cost of applying encryption is constant, executing q encryptions are decryptions is more costly than sort a list with q entries. However, the adversary will not actually execute the encryptions and decryptions themselves, but will request them from the challenger, and in that case, the time complexity of Algorithm 2 is indeed superior to that of Algorithm 3, since the former will be able to terminate shortly after all the queries are executed, while the later needs to execute the costly sorting operation. However, the exponential space complexity of Algorithm 2 makes it unsuitable for attacking practical instances of TNT. In Table 1.2, we provide the parameters for attacking TNT-AES using Algorithm 3, bounding both time and memory

Table 3.3: Results for an attack on TNT-GIFT-64.

n	64			
$\log_2(q)$	32	33	34	35
Average Number of Collisions	1	4	16	61
Time	3 hrs	3 hrs 40 mins	12 hrs 15 mins	20 hrs
CPU Time	5 hrs	10 hrs	28 hr 15 mins	72 hrs
Number of Cores	2	4	8	16
RAM	96 GB	192 GB	128 GB	192 GB
Disk Space	73 GB	146 GB	292 GB	583 GB

by $2^{n/2+5} = 2^{69}$. We ignore the $\log_2$ term in the time complexity since this is concerning the practical and not asymptotic performance, which is dominated by the encryptions and decryptions.

Attacking TNT-GIFT-64. We have implemented this variant to attack TNT instantiated with GIFT-64 [3]. We used the implementation of GIFT-64 described in [1] which can encrypt 2 blocks at the same time. We implemented the attack over 16 cores on an Intel Xeon E5-2630 CPU, each doing 2^{31} encryption calls and 2^{31} decryption calls (2^{30} calls $\times$ 2 blocks), generating 2^{35} blocks in total. This process took two hours (32 core-hours). In practice, the adversary is unlikely to be able to parallelize the queries, since that depends on the challenger.

Counting the collisions cannot be parallelized. It requires 40 minutes to count collisions in a set of 2^{32} blocks and 1 hour, 20 minutes in a set of 2^{33} blocks, generating on average 1 collision and 4 collisions, respectively. These results are reported in details in Table 3.3. We note that for $q \geq 2^{34}$, the attack uses less memory and significantly more time than the other cases. This is due to memory limitations, since the platform is limited by 256 GB, so the collision counting phase had to be optimized towards memory consumption, leading to a significant slow down. The time in Table 3.3 seems (at first glance) dominated by collision counting, which is contradictory to the statement we made earlier. However, it is to be noted that the collision counting part is serial in nature, while the TNT queries have been parallelized. For instance, performing the attack with 2^{35} complexity needs 68 core-hour, while counting needs 36 core-hour on a limited memory machine, but can be faster on a machine with more memory. In particular, we estimate that with memory of about 384 GB and 768 GB, we can run the attacks with 2^{34} and 2^{35} complexities in slightly more than 20 and 40 core-hours, respectively.

4 Spotting the Flaw in the BBB Security Proof of TNT

In [4], Bao et al. presented an IND-CCA security proof for TNT that contradicts our attack. This proof employs the χ^2 technique [14] — a relatively new proof technique — due to Dai et al.

In this section, we carefully revisit the security proof with the distinguisher $\mathcal{A}^*$, and identify an issue that involves a subtle, yet fundamental, case analysis. We temporarily switch to the notation of [4] to follow their proof approach. Namely, the random variable corresponding to plaintext is referred to as X. the random variable corresponding to ciphertext is referred to as Y and the random variable corresponding to the tweak is referred to as T. The rest of the random variables are related to the internal values of TNT and relate to the first three variables as: $\mathsf{S} = \boldsymbol{\pi}_1(\mathsf{M})$, $\mathsf{U} = \mathsf{T} \oplus \mathsf{S}$, $\mathsf{V} = \boldsymbol{\pi}_2(\mathsf{U})$, $\mathsf{W} = \mathsf{T} \oplus \mathsf{V}$ and $\mathsf{Y} = \boldsymbol{\pi}_3(\mathsf{W})$. For the l^{th} query to the construction, we define a set $\mathcal{Q}_l$ as the set of the first l queries $\{(\mathsf{T}_1, \mathsf{X}_1, \mathsf{Y}_1), \ldots, (\mathsf{T}_l, \mathsf{X}_l, \mathsf{Y}_l)\}$. We follow a slight abuse of notation utilized in [4]: we say $\mathsf{X} \in \mathcal{Q}_l$ to mean $\exists(\mathsf{T}, \mathsf{X}, \mathsf{Y}) \in \mathcal{Q}_l$, and similarly for Y. We define a random variable Inter as the vector of internal values in the first $l-1$ queries:

$$((\mathsf{S}_1, \ldots, \mathsf{S}_{l-1}), (\mathsf{U}_1, \ldots, \mathsf{U}_{l-1}), (\mathsf{V}_1, \ldots, \mathsf{V}_{l-1}), (\mathsf{W}_1, \ldots, \mathsf{W}_{l-1}))\,.$$

The main technique of the proof, from a high level point of view, works as follows:

- A deterministic distinguisher observes the first $l-1$ queries and selects whether the next query is a forward or inverse query as well as the tweak T_l and the plaintext X_l or ciphertext Y_l (M_l and C_l using our notations, respectively).
- Find the probability distribution of all the internal values of the construction given the first $l-1$ query. We call a set of possible vectors of internal values Inter.
- For each possible Inter, estimate the probability distribution of each possible response to query l.

The authors then analyze different possible cases and apply the χ^2 method on the resulting distribution.

In order to better understand the issue, we analyze our distinguisher in the flow of the security proof. Our distinguisher works as follows:

- If l is odd, it makes a forward query $(\mathsf{X}_0, \mathsf{T}_{l-2} + 1)$.
- If l is even, it makes a backward query $(\mathsf{Y}_{l-1}, \mathsf{T}_{l-1} \oplus \delta)$.

Let $(\mathsf{S}_o, \mathsf{U}_o, \mathsf{V}_o)$ are the output of $\boldsymbol{\pi}_1$, input of $\boldsymbol{\pi}_2$ and output of $\boldsymbol{\pi}_2$ in the last (odd) query $l-1$, and we estimate the probability, for a given $\mathsf{X}_i \in \mathcal{Q}_l$ where i is even, $\Pr[\mathsf{X}_l = \mathsf{X}_i]$.

Let $(\mathsf{S}_i, \mathsf{U}_i, \mathsf{V}_i)$ and $(\mathsf{S}_e, \mathsf{U}_e, \mathsf{V}_e)$ are the corresponding internal values of X_i and X_l, respectively. Then, we know that $\mathsf{V}_o \oplus \mathsf{V}_e = \delta$ and

$$\begin{aligned}\Pr[\mathsf{X}_l = \mathsf{X}_i] = \Pr[\mathsf{S}_e = \mathsf{S}_i] &= \Pr[\mathsf{U}_e \oplus \mathsf{T}_{l-1} \oplus \delta = \mathsf{U}_i \oplus \mathsf{T}_{i-1} \oplus \delta] \\ &= \Pr[\mathsf{U}_e \oplus \mathsf{U}_i = \mathsf{T}_{l-1} \oplus \mathsf{T}_{i-1}]\end{aligned}$$

Since X_0 is fixed for all odd queries, so is S_o. Thus, $\mathsf{U}_o \oplus \mathsf{T}_{l-1} = \mathsf{U}_{i-1} \oplus \mathsf{T}_{i-1}$. Therefore,

$$\Pr[\mathsf{U}_e \oplus \mathsf{U}_i = \mathsf{T}_{l-1} \oplus \mathsf{T}_{i-1}] = \Pr[\mathsf{U}_e \oplus \mathsf{U}_o = \mathsf{U}_i \oplus \mathsf{U}_{i-1}] \approx \frac{c}{2^n},$$

where c is a small positive integer constant. The security proof considers two possible cases such collisions may occur. The first is when U_e has appeared before in one of the previous queries, and the second is when it has never appeared before. They dubbed these two cases as class $\mathcal{A}$ and class $\mathcal{B}$ respectively. The collision can occur in either class $\mathcal{A}$ or class $\mathcal{B}$, which the proof bounds the probability of their probability for the l^{th} query by $4l/2^{2n}$ and $1/(2^n - l)$, respectively. Thus, our analysis deviates from the distribution assumed in [4]. In terms of the proof presented in [4], the event we are discussing belongs to case 5 (case 1 if we swap all the forward and backward queries). In this case, the authors claim [4, (9)]).

$$\Pr[\mathsf{X}_l = \mathsf{X}_i] \leq \frac{4l}{2^{2n}} + \frac{1}{2^n - l}$$

We argue that the distribution assumed for case 5/case 1 - class $\mathcal{A}$ erroneously underestimates the probability of certain bad events, and by changing the distribution to account for these bad events, the proof argumentation falls apart. Besides, it is not clear how to do so in the existing proof framework using the χ^2 method.

In particular, we look at the term $4l/2^{2n}$. The term stems from the following argument in [4]:

> It remains to bound $\Pr[\mathsf{Inter} \in \mathcal{A}|\mathcal{Q}_{l-1}]$. For this, note that once the values in Inter except for $(\mathsf{S}_l, \mathsf{W}_l)$ have been fixed, the number of choices for $(\mathsf{S}_l, \mathsf{W}_l)$ is at least $(2^n - \alpha(\mathcal{Q}_{l-1}))(2^n - \gamma(\mathcal{Q}_{l-1})) \geq 2^{2n}/4$, where $\alpha(\mathcal{Q}_{l-1}) \geq q \geq 2^n/2$ and $\gamma(\mathcal{Q}_{l-1}) \geq q \geq 2^n/2$ are the number of distinct values in $(\mathsf{S}_1, \ldots \mathsf{S}_{l-1})$ and $(\mathsf{W}_1, \ldots \mathsf{W}_{l-1})$. Out of these $\geq 2^{2n}/4$ choices, the number of choices that ensure the desired property $\mathsf{TNT}\,(\mathsf{T}_l, \mathsf{X}_l) = \mathsf{Y}_l$ is at most $l - 1$, which results from the following selection process: we first pick a pair of input-output $(\mathsf{U}_i, \mathsf{V}_i)$ with $i \leq l - 1$, and then set $\mathsf{S}_l = \mathsf{T}_l \oplus \mathsf{U}_i$ and $\mathsf{W}_l = \mathsf{T}_l \oplus \mathsf{V}_i$. Therefore, $\Pr[\mathsf{Inter} \in \mathcal{A}|\mathcal{Q}_{l-1}] \leq 4l/2^{2n}$, and thus the upper bound in this case is
>
> $$\frac{4l}{2^{2n}} + \frac{1}{2^n - l}.$$

Consider the first case of the collision in Figure 4.1. We note that if the triplet $(\alpha, \mathsf{S}_o, \mathsf{U}_o)$ is known, then the collision happens with probability 1, which puts it in class $\mathcal{A}$. Then, what remains is to calculate what is the probability that the adversary can force this collision, *i.e.,*

$$\Pr[\mathsf{Inter} \in \mathcal{A}|\mathcal{Q}_{l-1}] = \Pr[\mathsf{U}_e \oplus \mathsf{U}_o = \mathsf{T}_{l-1} \oplus \mathsf{T}_{i-1}|\mathcal{Q}_{l-1}],$$

where $\mathsf{T}_{l-1} = t_{l/2}$ and $\mathsf{T}_{i-1} = t_{i/2}$ are determined by the adversary during previous queries. This means than once U_o and all other values of U except U_e in Inter are fixed (both U_o and U_e belong to queries $i, j < l$), U_e has at most $2^n - \alpha(\mathcal{Q}_{l-1})$ choices where $\alpha(\mathcal{Q}_{l-1}) \leq q \leq 2^{n-1}$ is the number of distinct values

in $\{\mathsf{U}_1, \ldots \mathsf{U}_l\} \setminus \{\mathsf{U}_e\}$, and at most 1 of them ($\mathsf{U}_e = \mathsf{U}_o \oplus \alpha$) enforces the collision. In other words,

$$\Pr[\mathsf{Inter} \in \mathcal{A} | \mathcal{Q}_{l-1}] = \Pr[\mathsf{U}_e \oplus \mathsf{U}_o = \mathsf{T}_{l-1} \oplus \mathsf{T}_{i-1} | \mathcal{Q}_{l-1}] \geq \frac{1}{2^n - \alpha(\mathcal{Q}_{l-1})} \geq \frac{1}{2^n} \gg \frac{4l}{2^{2n}},$$

when $l \ll q$, contradicting [4, (9)]. This reflects in the final analysis of case 1 as follows: In [4, (11)], we take the maximum of two expressions. One is on the form $al/2^{2n}$ for a small constant a, and one is on the form $4l/2^{2n} + O(l/2^{2n})$. The term $4l/2^{2n}$ comes from [4, (9)]. However, if the term is on the form $O(1/2^n)$ instead of $4l/2^{2n}$, as suggested by our attack, then the maximum function in [4, (11)] would return $O(1/2^n) + O(l/2^{2n})$. Thus, the squared difference used in the χ^2 statistic becomes one the form

$$\left(\frac{1}{2^n - \alpha(\mathcal{Q}_{l-1})} - \frac{1}{2^n - \mu_l} + \frac{1}{2^n - l} \right)^2$$

or

$$\left(\frac{A2^n + B2^{2n}}{2^{3n}} \right)^2$$

for some constants A and B. Note that $\alpha(\mathcal{Q}_{l-1})$ is based on the probabilistic behaviour of the transcript, while μ_l is fully controlled by the adversary, and we cannot ensure that $\mu_l = \alpha(\mathcal{Q}_{l-1})$. Thus, the squared difference cannot be bounded tighter than $O(1/2^{2n})$. Besides,

$$\frac{1}{2^n - \alpha(\mathcal{Q}_{l-1})}$$

is a lower bound. The χ^2 statistic then becomes on the form

$$\sum \frac{O(1/2^{2n})}{O(1/2^n)} \approx \sum O(1/2^n) \approx 2^n \dot{O}(1/2^n) \approx O(1)$$

not leading to any meaningful security.

Note that the values of V_i and W_i for $i < l$ did not affect the behaviour of the collision or the probability that Inter is in class $\mathcal{A}$. It seems the ambiguity may stem from applying the χ^2 method to a primitive with two dependent functions ($\tilde{\pi}$ and its inverse). By cascading forward and backward queries, we managed to eliminate W_i for all $1 \leq i \leq q$ and the values of W_l do not matter for the attack. Similarly, by fixing the difference between V_o and V_e to a constant δ, we minimize the effect of their exact values on the attack.

5 Birthday-bound Security of TNT and Its Variant

In light of the above discussion, it is clear that the security of TNT is in limbo. One can rely on the IND-CCA bound by Zhang et al. to demonstrate the tightness of the proposed attacks. However, we observe that the generic bound in [44]

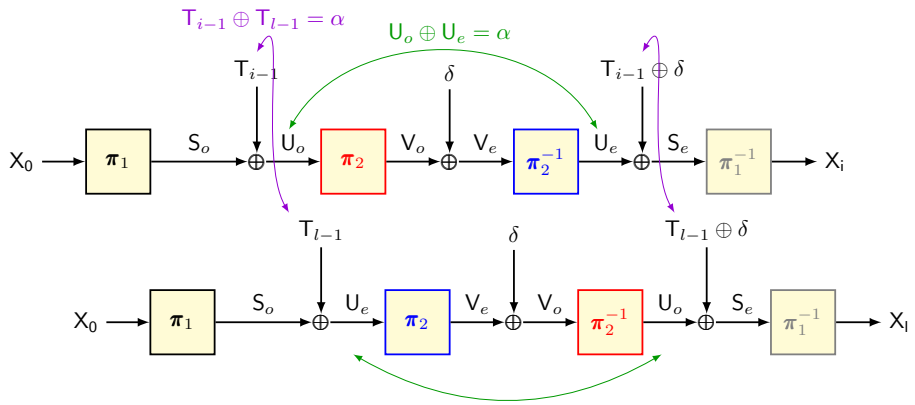

Fig. 4.1: A class $\mathcal{A}$ collision occurring in Algorithm 1. If the collision occurs, then permutation calls with the same color compute the same permutation point. Curved arrows represent difference relations (if the collision occurs with these internal values, each two connected nodes differ by α).

introduces some constant factors, and in general, an independent security proof, using a different proof technique, will instill greater confidence in the revised security claims of TNT.

Theorem 5.1. *Let $\boldsymbol{\pi}_1$, $\boldsymbol{\pi}_2$, and $\boldsymbol{\pi}_3$ be three independent random permutations of $\{0,1\}^n$. Then, for all $q \geq 1$, we have*

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\mathsf{TNT}}(q) \leq \frac{q^2}{2^n}.$$

In fact, we intend to prove a stronger version of Theorem 5.1, as stated in Theorem 5.2 below. Particularly, we show that even the single-keyed TNT, which we denote as 1k-TNT, is sufficient to achieve birthday-bound security. A proof of Theorem 5.1 is available in Supplementary Material F, for the sake of completeness.

Theorem 5.2. *Let $\boldsymbol{\pi}_1 = \boldsymbol{\pi}_2 = \boldsymbol{\pi}_3 = \boldsymbol{\pi}$, where $\boldsymbol{\pi}$ is a uniform random permutation of $\{0,1\}^n$. Then, for all $q \geq 1$, we have*

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\mathsf{1k\text{-}TNT}}(q) \leq \frac{8q^2}{2^n}.$$

Proof. The statement is vacuously true for $q \geq 2^{n/2}$. We will use the Expectation method (see Lemma 2.1) to prove the statement for $1 \leq q < 2^{n/2}$.

Let $\mathcal{O}_1$ and $\mathcal{O}_0$ be the oracles corresponding to 1k-TNT and a tweakable random permutation $\widetilde{\boldsymbol{\pi}}$, respectively. If $(\mathsf{T}_i, \mathsf{M}_i)$ is the encryption query with a

tweak T_i we write the response as $\widehat{\mathsf{C}}_i$. Similarly, if $(\mathsf{T}_i, \widehat{\mathsf{C}}_i)$ is the decryption query with a tweak T_i we write the response as M_i. After all queries have been made, the two oracles release some additional data to the adversary, who is obviously free to ignore this additional information, $\widehat{\mathsf{M}}^q$ and C^q.

In the real world, $\widehat{\mathsf{M}}^q$ and C^q correspond to the output of the first permutation and input of the third permutation, respectively, and thus they are well defined from the definition of 1k-TNT. The real-world transcript is thus defined as the tuple

$$\Theta_1 := (\mathsf{T}^q, \mathsf{M}^q, \widehat{\mathsf{C}}^q, \widehat{\mathsf{M}}^q, \mathsf{C}^q).$$

In the ideal system $\widetilde{\boldsymbol{\pi}}$, we sample $\widehat{\mathsf{M}}^q, \mathsf{C}^q$ as follows: For every $i \in [q]$,

1. $\widehat{\mathsf{M}}_i = \widehat{\mathsf{M}}_j$ whenever $\mathsf{M}_i = \mathsf{M}_j$ for $j < i$. Otherwise (for all $j < i$, $\mathsf{M}_j \neq \mathsf{M}_i$), we sample
$$\widehat{\mathsf{M}}_i \xleftarrow{\$} \{0,1\}^n \setminus \mathsf{S}(\widehat{\mathsf{M}}^{[i-1]}).$$
2. $\mathsf{C}_i = \mathsf{C}_j$ whenever $\widehat{\mathsf{C}}_j = \widehat{\mathsf{C}}_i$ for $j < i$. Otherwise (for all $j < i$, $\widehat{\mathsf{C}}_j \neq \widehat{\mathsf{C}}_i$), we sample
$$\mathsf{C}_i \xleftarrow{\$} \{0,1\}^n \setminus \mathsf{S}(\mathsf{C}^{[i-1]}).$$

The ideal world transcript is defined as

$$\Theta_0 := (\mathsf{T}^q, \mathsf{M}^q, \widehat{\mathsf{C}}^q, \widehat{\mathsf{M}}^q, \mathsf{C}^q).$$

Note that we use the same notation to denote the random variables in both worlds. However, their probability distributions will be unambiguously determined at the time of probability computations.

Bad Transcript and Its Analysis: Let $u^q := \widehat{m}^q \oplus t^q$, and $\widehat{u}^q := c^q \oplus t^q$. A transcript $(t^q, m^q, \widehat{c}^q, \widehat{m}^q, c^q)$ is called *bad* if and only if any of the following bad events occur:
bad_{1a}: $\exists i \neq j \in [q]$ such that $\widehat{m}_i = \widehat{c}_j$.
bad_{1b}: $\exists i \neq j \in [q]$ such that $c_i = m_j$.
bad_{2a}: $\exists i < j \in [q]$ such that $u_i = u_j$.
bad_{2b}: $\exists i < j \in [q]$ such that $\widehat{u}_i = \widehat{u}_j$.
bad_{3a}: $\exists i \neq j \in [q]$ such that $u_i = m_j$.
bad_{3b}: $\exists i \neq j \in [q]$ such that $\widehat{u}_i = \widehat{c}_j$.
bad_{4a}: $\exists i \neq j \in [q]$ such that $u_i = c_j$.
bad_{4b}: $\exists i \neq j \in [q]$ such that $\widehat{u}_i = \widehat{m}_j$.
Let Ω_{bad} denote the set of all bad transcripts. Then, using union bound, we have

$$\mathsf{Pr}\left(\Theta_0 \in \Omega_{\mathsf{bad}}\right) \leq \sum_{\substack{i \in [4] \\ \mathtt{s} \in \{a,b\}}} \mathsf{Pr}\left(\mathsf{bad}_{i\mathtt{s}}\right) \tag{16}$$

On the right-hand side, we bound the probability for $\mathsf{bad}_{i\mathtt{a}}$ for all $i \in [4]$. The $\mathtt{s} = b$ cases can be bounded analogously.

- $\mathsf{Pr}\left(\mathsf{bad}_{1a}\right) \leq \frac{q^2}{2^n}$. This follows from union bound: For a fixed choice of i and j, bad_{1a} happens with 2^{-n} probability, and there are q^2 such (i, j) pairs.

- $\mathsf{Pr}\left(\mathsf{bad}_{2a}\right) \leq \sum_{i<j} \mathsf{Pr}\left(\widehat{\mathsf{M}}_i + \mathsf{T}_i = \widehat{\mathsf{M}}_j + \mathsf{T}_j\right) \leq q^2/2^n$. This can be argued as follows: For fixed i and j, $\mathsf{Pr}\left(\widehat{\mathsf{M}}_i + \mathsf{T}_i = \widehat{\mathsf{M}}_j + \mathsf{T}_j\right) \leq 1/2^{n-1} \leq 2^{1-n}$, and there are $\binom{q}{2}$ such (i,j) pairs.
- $\mathsf{Pr}\left(\mathsf{bad}_{3a}\right) \leq \frac{q^2}{2^n}$. The argumentation is similar to the one for bad_{1a}.
- $\mathsf{Pr}\left(\mathsf{bad}_{4a}\right) \leq \frac{q^2}{2^n}$. This can be argued as follows: For any fixed i and j, bad_{4a} happens with at most 2^{-n} probability, and there are at most q^2 such (i,j) pairs.

Thus, on combining everything in (16), we have

$$\mathsf{Pr}\left(\Theta_0 \in \Omega_{\mathsf{bad}}\right) \leq \frac{8q^2}{2^n}.$$

Analysis of Good Transcripts: For a good transcript $\tau = (t^q, m^q, \widehat{c}^q, \widehat{m}^q, c^q)$, we know that $(m^q, \widehat{m}^q)$, $(c^q, \widehat{c}^q)$, and $(u^q, \widehat{u}^q)$ are permutation consistent non-overlapping input-output pairs and hence for the real world we have

$$\begin{aligned}\mathsf{Pr}\left(\Theta_1 = \omega\right) &= \mathsf{Pr}\left(\boldsymbol{\pi}(m^q) = \widehat{m}^q\right) \times \mathsf{Pr}\left(\boldsymbol{\pi}(u^q) = \widehat{u}^q\right) \times \mathsf{Pr}\left(\boldsymbol{\pi}(c^q) = \widehat{c}^q\right) \\ &= \frac{1}{(2^n)_{r+q+s}}\end{aligned}$$

where r and s denote the size of $\mathtt{S}(m^q)$ and $\mathtt{S}(\widehat{c}^q)$ respectively. In the ideal world, we have,

$$\mathsf{Pr}\left(\Theta_0 = \omega\right) = \mathsf{Pr}\left(\widetilde{\boldsymbol{\pi}}(t^q, m^q) = \widehat{c}^q\right) \times \frac{1}{(2^n)_r} \times \frac{1}{(2^n)_s} \leq \frac{1}{(2^n)_q} \times \frac{1}{(2^n)_r} \times \frac{1}{(2^n)_s},$$

where the final inequality follows from the fact that $\mathsf{Pr}\left(\widetilde{\boldsymbol{\pi}}(t^q, m^q) = \widehat{c}^q\right)$ maximizes when $t_i = t_j$ for all $1 \leq i < j \leq q$. Thus

$$\frac{\mathsf{Pr}\left(\Theta_1 = \omega\right)}{\mathsf{Pr}\left(\Theta_0 = \omega\right)} \geq \frac{(2^n)_q \times (2^n)_r \times (2^n)_s}{(2^n)_{q+r+s}} \geq 1$$

Now the result follows from the Expectation method by setting $\epsilon_{\mathsf{ratio}}$ to be a zero function.

6 The Generalized LRW Paradigm

In this section, we propose a generalized view of the cascaded LRW design that encompasses both cascaded LRW1 and cascaded LRW2 constructions. In addition, we identify some necessary properties to guarantee IND-CCA security up to $2^{3n/4}$ queries.

Almost XOR Universal Hash Function: A (τ, n)-hash function family $\mathcal{H}$, is a family of functions $\{h : \{0,1\}^\tau \to \{0,1\}^n\}$, keyed implicitly by the choice

of h. A (τ, n)-hash function family $\mathcal{H}$ is called an ϵ-almost XOR universal hash family (AXUHF) if for all $t \neq t' \in \{0,1\}^\tau$, and $\delta \in \{0,1\}^n$, we have

$$\mathsf{Pr}\left(\mathsf{H} \xleftarrow{\$} \mathcal{H} \ : \ \mathsf{H}(t) \oplus \mathsf{H}(t') = \delta\right) \leq \epsilon. \tag{17}$$

For the special case of $\delta = 0^n$, $\mathcal{H}$ is referred as an ϵ-AUHF.

THE LRW+ CONSTRUCTION: Let $\widetilde{\mathcal{H}}$ be a family of (τ, n)-tweakable permutations, and $\mathcal{H}$ be a (τ, n)-hash function family. Let $\widehat{\mathcal{H}} = (\widetilde{\mathcal{H}}^2 \times \mathcal{H})$, $(\widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H}) \leftarrow \mathsf{KG}\left(\widehat{\mathcal{H}}\right)$, and $(\boldsymbol{\pi}_1, \boldsymbol{\pi}_2) \xleftarrow{\$} \mathsf{Perm}(n)$, where $\mathsf{KG}\left(\widehat{\mathcal{H}}\right)$ is an efficient probabilistic algorithm that returns a random triple from $\widehat{\mathcal{H}}$.

The LRW+ construction is a (τ, n)- tweakable permutation family, defined by the following mapping (see Figure 6.1 for an illustration):

$$(t, m) \mapsto \widetilde{\mathbf{H}}_2^{-1}\left(t, \boldsymbol{\pi}_2\left(\mathbf{H}(t) \oplus \boldsymbol{\pi}_1\left(\widetilde{\mathbf{H}}_1\left(t, m\right)\right)\right)\right). \tag{18}$$

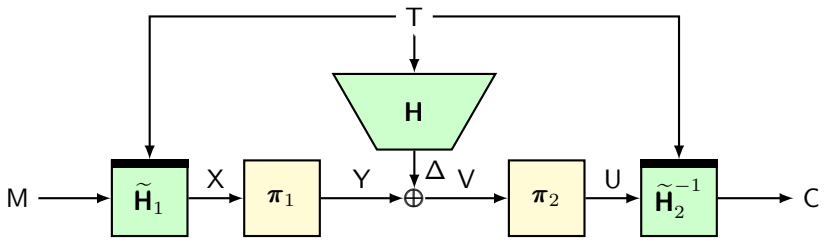

Fig. 6.1: The LRW+ construction.

6.1 Security of LRW+

We say that $\mathsf{KG}\left(\widehat{\mathcal{H}}\right)$ is a pairwise independent sampling mechanism or PISM, if $(\widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H}) \leftarrow \mathsf{KG}\left(\widehat{\mathcal{H}}\right)$ is a pairwise independent tuple.

We say that $\widetilde{\mathcal{H}}$ is an ϵ-almost universal tweakable permutation family (AUTPF) if and only if for all distinct $(t, m), (t', m') \in \{0,1\}^\tau \times \{0,1\}^n$,

$$\mathsf{Pr}\left(\widetilde{\mathbf{H}} \xleftarrow{\$} \widetilde{\mathcal{H}} \ : \ \widetilde{\mathbf{H}}(t, m) = \widetilde{\mathbf{H}}(t', m')\right) \leq \epsilon.$$

Theorem 6.1. *Let $\tau, n \in \mathbb{N}$, and $\epsilon_1, \epsilon_2 \in [0, 1]$. If $\widetilde{\mathcal{H}}$ and $\mathcal{H}$ are respectively ϵ_1-AUTPF and ϵ_2-AUHF, and $\mathsf{KG}\left(\widehat{\mathcal{H}}\right)$ is a PISM, then, for $q \leq 2^{n-2}$, we have*

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\mathsf{LRW+}}(q) \leq \epsilon(q, n),$$

where

$$\epsilon(q, n) = 2q^2\epsilon_1^{1.5} + \frac{4q^4\epsilon_1^2}{2^n} + \frac{32q^4\epsilon_1}{2^{2n}} + \frac{13q^4}{2^{3n}} + q^2\epsilon_1^2 + q^2\epsilon_1\epsilon_2 + \frac{2q^2}{2^{2n}}. \tag{19}$$

A proof of this theorem follows from a simple generalization of Jha and Nandi's (JN) proof [25] for 2-LRW2. In particular, the exact same strategy of using the expectation method with the JN adaptation of mirror theory [38,10] in the tweakable permutation settings works here as well. For the sake of completeness, we give the complete proof in Supplementary Material G.

Remark 6.1. The proof presented in [25] appears to have overlooked the analysis of a specific subset of transcripts, which in hindsight of our generalized analysis seems to be a minor issue. Indeed, our proof demonstrates that this omission does not impact the overall bound significantly, with any potential effects being limited to a small constant factor.

6.2 Instantiating LRW+

We show that any cascaded LRW construction with $r \geq 2$ rounds can be viewed as an instance of LRW+. Thus, they can be proven secure up to $2^{3n/4}$ queries provided the derived hash functions are close to 2^{-n}-universal. Note that it would be sufficient to define $\widetilde{\mathsf{H}}_1, \widetilde{\mathsf{H}}_2, \mathsf{H}, \boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$ for each construction. In the following discussion, let $\boldsymbol{\pi}'^r \xleftarrow{\$} \mathsf{Perm}(n)$ and $\mathsf{H}'^r \xleftarrow{\$} \mathcal{H}^r$, where $\mathcal{H}$ is an ϵ-AXUHF.

Cascaded LRW1. For $r \geq 2$, the r-LRW1$[\boldsymbol{\pi}^r]$ construction takes as input $(t, m) \in \{0,1\}^n \times \{0,1\}^n$ and returns $c \in \{0,1\}^n$, which is defined as follows: Let $y_0 = t \oplus m$ and for all $i \in [r]$:

$$x_i := t \oplus y_{i-1} \qquad \text{and} \qquad y_i := \boldsymbol{\pi}'_i(x_i),$$

and finally $c := y_r$. The inverse of r-LRW1 is analogously defined.

Cascaded LRW1 as an Instance of LRW+: For some $r \geq 2$, $r' := \lfloor r/2 \rfloor$, and any (t, m) such that r-LRW1$(t, m) = c$, define $\widetilde{\mathsf{H}}_1(t, m) := x_{r'}$, $\mathsf{H}(t) := t$, $\widetilde{\mathsf{H}}_2(t, c) := y_{r'+1}$, $\boldsymbol{\pi}_1 := \boldsymbol{\pi}'_{r'}$, and $\boldsymbol{\pi}_2 := \boldsymbol{\pi}'^{-1}_{r'+1}$.

Clearly, the LRW+ instance so defined is same as r-LRW1. We have the following corollary on the security of cascaded LRW1.

Corollary 6.1. *For $r \geq 4$, we have*

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{r\text{-}LRW1}(q) \leq \frac{2q^2}{(2^n-1)^{1.5n}} + \frac{49q^4}{(2^n-1)^3} + \frac{3q^2}{(2^n-1)^2}.$$

In particular, for $r = 4$, we have proved CCA security for 4-LRW1 up to $2^{3n/4}$ queries. A proof of this corollary is available in Supplementary Material H.1.

Cascaded LRW2. For $r \geq 1$, the r-LRW2$[\boldsymbol{\pi}^r, \mathsf{H}'^r]$ construction takes as input $(t, m) \in \{0,1\}^\tau \times \{0,1\}^n$ and returns $c \in \{0,1\}^n$, which is defined as follows: Let $y_0 = m$, H'_0 be a constant function that returns 0^n, and for all $i \in [r]$:

$$x_i := \mathsf{H}'_{i-1}(t) \oplus \mathsf{H}'_i(t) \oplus y_{i-1} \qquad \text{and} \qquad y_i := \boldsymbol{\pi}'_i(x_i),$$

and finally $c := \mathsf{H}'_r(t) \oplus y_r$. The inverse of r-LRW2 is analogously defined.

Cascaded LRW2 as an Instance of LRW+: For some $r \geq 2$, $r' = \lfloor r/2 \rfloor$, and any (t, k, m) such that r-LRW2$(t, m) = c$, define $\widetilde{\mathsf{H}}_1(t, m) := x_{r'}$, $\mathsf{H}(t) := \mathsf{H}'_{r'}(t) \oplus \mathsf{H}'_{r'+1}(t)$, $\widetilde{\mathsf{H}}_2(t, c) := y_{r'+1}$, $\boldsymbol{\pi}_1 := \boldsymbol{\pi}'_{r'}$, $\boldsymbol{\pi}_2 := \boldsymbol{\pi}'^{-1}_{r'+1}$.

Clearly, the LRW+ instance so defined is same as r-LRW2. We have the following corollary on the security of cascaded LRW2.

Corollary 6.2. *For $r \geq 2$, we have*

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{r\text{-}LRW2}(q) \leq 2q^2\epsilon^{1.5} + \frac{4q^4\epsilon^2}{2^n} + \frac{32q^4\epsilon}{2^{2n}} + \frac{13q^4}{2^{3n}} + 2q^2\epsilon^2 + \frac{2q^2}{2^{2n}}.$$

In particular, for $r = 2$, assuming $\epsilon = O\left(2^{-n}\right)$, we have reproved[4] the CCA security for 2-LRW2 up to $2^{3n/4}$ queries. A proof of this corollary is available in Supplementary Material H.2.

7 Conclusion and Future Directions

In this paper, we gave a birthday-bound CCA distinguisher on TNT, thereby completely invalidating its beyond-the-birthday bound security claims. Further, we showed that our attack is tight by reestablishing a birthday bound security for TNT and its single-keyed variant.

In addition, we showed that by adding just one more block cipher call, the security can be amplified to $3n/4$-bit even in the CCA setting. We note that our generalization of the cascaded LRW constructions could be of independent interest.

Open Problems: This work opens several new research avenues in (block cipher-based) TBC constructions. Some prominent problems that are worth exploring include:

1. Optimal LRW Construction for BBB Security: 4-LRW1 employs 4 calls of block cipher. Similarly, 2-LRW2 with block cipher based hash functions also requires 4 calls. This raises a natural question regarding their optimality. In other words, are 4 block cipher calls necessary for BBB security?
2. Reduced-key Version of 4-LRW1: 4-LRW1 needs 4 independent keys. Is it possible to reduce the number of keys from 4 to 3, or 2?
3. Exact Security of 4-LRW1: We do not have an attack against 4-LRW1. Neither Mennink's $O(\sqrt{n}2^{3n/4})$-distinguisher, nor any variant of our $O(2^{n/2})$-distinguisher seem to work. The additional permutation calls seem to help in avoiding these attack strategies. It would be interesting to see if there exists an attack that matches our bound, or if the construction is beyond $3n/4$-bit secure?

[4] The proof in [25] also has a minor issue, that leads to a slightly worse constant.

4. Security of short-tweak TNT: Our attack requires a tweak space of size roughly $2^{n/2}$. So it is natural to ask if TNT is still BBB secure when the tweak space size is much less than $2^{n/2}$?
5. Security of Longer Cascades of LRW: This is a long standing problem even for the more analyzed case of r-LRW2, for $r \geq 3$. The best bounds [11,44] that we have in this case are coupling-based. It is clear from our bound for LRW+ that these bounds are rather loose. It would be interesting to explore the possibility of better security bounds for the general case with a dedicated and more tighter analysis.

Acknowledgments: The authors would like to thank Chun Guo for his comments on the attacks presented on TNT. Ashwin Jha carried out this work under the framework of the French-German-Center for Cybersecurity, a collaboration of CISPA and LORIA. Mustafa Khairallah performed this work as part of the Seagate Research Group.

References

1. Adomnicai, A., Najm, Z., Peyrin, T.: Fixslicing: A new GIFT representation fast constant-time implementations of GIFT and GIFT-COFB on ARM cortex-m. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2020**(3), 402–427 (2020)
2. Avanzi, R.: The QARMA block cipher family. almost MDS matrices over rings with zero divisors, nearly symmetric even-mansour constructions with non-involutory central rounds, and search heuristics for low-latency s-boxes. IACR Trans. Symmetric Cryptol. **2017**(1), 4–44 (2017)
3. Banik, S., Pandey, S.K., Peyrin, T., Sasaki, Y., Sim, S.M., Todo, Y.: GIFT: A small present - towards reaching the limit of lightweight encryption. In: Cryptographic Hardware and Embedded Systems - CHES 2017, Proceedings. pp. 321–345 (2017)
4. Bao, Z., Guo, C., Guo, J., Song, L.: TNT: How to Tweak a Block Cipher. In: Advances in Cryptology - EUROCRYPT 2020, Proceedings, Part II. pp. 641–673 (2020)
5. Bariant, A., Leurent, G.: Truncated boomerang attacks and application to aes-based ciphers. In: Advances in Cryptology - EUROCRYPT 2023, Proceedings, Part IV. pp. 3–35 (2023)
6. 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: Advances in Cryptology - CRYPTO 2016, Proceedings, Part II. pp. 123–153 (2016)
7. Bhattacharya, S., Nandi, M.: A note on the chi-square method: A tool for proving cryptographic security. Cryptogr. Commun. **10**(5), 935–957 (2018)
8. Bhaumik, R., List, E., Nandi, M.: ZCZ - Achieving n-bit SPRP Security with a Minimal Number of Tweakable-Block-Cipher Calls. In: Advances in Cryptology - ASIACRYPT 2018, Proceedings, Part I. pp. 336–366 (2018)
9. Chakraborty, D., Sarkar, P.: A General Construction of Tweakable Block Ciphers and Different Modes of Operations. IEEE Trans. Information Theory **54**(5), 1991–2006 (2008)
10. Cogliati, B., Dutta, A., Nandi, M., Patarin, J., Saha, A.: Proof of Mirror Theory for a Wide Range of $\xi _{\max }$. In: Advances in Cryptology - EUROCRYPT 2023, Proceedings, Part IV. pp. 470–501 (2023)

11. Cogliati, B., Lampe, R., Seurin, Y.: Tweaking even-mansour ciphers. In: Advances in Cryptology - CRYPTO 2015, Proceedings, Part I. pp. 189–208 (2015). https://doi.org/10.1007/978-3-662-47989-6_9
12. Cogliati, B., Seurin, Y.: Beyond-birthday-bound security for tweakable even-mansour ciphers with linear tweak and key mixing. In: Advances in Cryptology - ASIACRYPT 2015, Proceedings, Part II. pp. 134–158 (2015). https://doi.org/10.1007/978-3-662-48800-3_6
13. Crowley, P.: Mercy: A Fast Large Block Cipher for Disk Sector Encryption. In: Fast Software Encryption - FSE 2000, Proceedings. pp. 49–63 (2000)
14. Dai, W., Hoang, V.T., Tessaro, S.: Information-theoretic indistinguishability via the chi-squared method. In: Advances in Cryptology - CRYPTO 2017, Proceedings, Part III. pp. 497–523 (2017)
15. Datta, N., Dey, S., Dutta, A., Mondal, S.: Cascading Four Round LRW1 is Beyond Birthday Bound Secure. IACR Cryptol. ePrint Arch. p. 1242 (2023), https://eprint.iacr.org/2023/1242
16. Granger, R., Jovanovic, P., Mennink, B., Neves, S.: Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption. In: Advances in Cryptology - EUROCRYPT 2016, Proceedings, Part I. pp. 263–293 (2016)
17. Gunsing, A., Bhaumik, R., Jha, A., Mennink, B., Shen, Y.: Revisiting the indifferentiability of the sum of permutations. IACR Cryptol. ePrint Arch. p. 840 (2023)
18. Guo, C., Guo, J., List, E., Song, L.: Towards Closing the Security Gap of Tweak-aNd-Tweak (TNT). In: Advances in Cryptology - ASIACRYPT 2020, Proceedings, Part I. pp. 567–597 (2020)
19. Guo, Z., Wang, G., Dunkelman, O., Pan, Y., Liu, S.: Tweakable SM4: how to tweak SM4 into tweakable block ciphers? J. Inf. Secur. Appl. **72**, 103406 (2023)
20. Hoang, V.T., Tessaro, S.: Key-Alternating Ciphers and Key-Length Extension: Exact Bounds and Multi-user Security. In: Advances in Cryptology - CRYPTO 2016, Proceedings, Part I. pp. 3–32 (2016)
21. Iwata, T., Minematsu, K., Peyrin, T., Seurin, Y.: ZMAC: A Fast Tweakable Block Cipher Mode for Highly Secure Message Authentication. In: Advances in Cryptology - CRYPTO 2017, Proceedings, Part III. pp. 34–65 (2017)
22. Jean, J., Nikolic, I., Peyrin, T.: Tweaks and Keys for Block Ciphers: The TWEAKEY Framework. In: Advances in Cryptology - ASIACRYPT 2014, Proceedings, Part II. pp. 274–288 (2014)
23. Jean, J., Nikolic, I., Peyrin, T., Seurin, Y.: The deoxys AEAD family. J. Cryptol. **34**(3), 31 (2021)
24. Jha, A., List, E., Minematsu, K., Mishra, S., Nandi, M.: XHX - A framework for optimally secure tweakable block ciphers from classical block ciphers and universal hashing. In: Progress in Cryptology - LATINCRYPT 2017, Revised Selected Papers. pp. 207–227 (2017). https://doi.org/10.1007/978-3-030-25283-0_12
25. Jha, A., Nandi, M.: Tight Security of Cascaded LRW2. J. Cryptol. **33**(3), 1272–1317 (2020)
26. Jha, A., Nandi, M., Saha, A.: Tight security of TNT: reinforcing khairallah's birthday-bound attack. IACR Cryptol. ePrint Arch. p. 1233 (2023)
27. Khairallah, M.: $\mathsf{CLRW1}^3$ is not Secure Beyond the Birthday Bound Breaking TNT with $O(2^{n/2})$ Queries. IACR Cryptol. ePrint Arch. p. 1212 (2023)
28. Krovetz, T., Rogaway, P.: The Software Performance of Authenticated-Encryption Modes. In: Fast Software Encryption - FSE 2011. Revised Selected Papers. pp. 306–327 (2011)

29. Lampe, R., Seurin, Y.: Tweakable Blockciphers with Asymptotically Optimal Security. In: Fast Software Encryption - FSE 2013, Revised Selected Papers. pp. 133–151 (2013)
30. Landecker, W., Shrimpton, T., Terashima, R.S.: Tweakable Blockciphers with Beyond Birthday-Bound Security. In: Advances in Cryptology - CRYPTO 2012, Proceedings. pp. 14–30 (2012)
31. Liskov, M.D., Rivest, R.L., Wagner, D.A.: Tweakable Block Ciphers. In: Advances in Cryptology - CRYPTO 2002, Proceedings. pp. 31–46 (2002)
32. Mennink, B.: Optimally secure tweakable blockciphers. In: Fast Software Encryption - FSE 2015, Revised Selected Papers. pp. 428–448 (2015)
33. Mennink, B.: Towards Tight Security of Cascaded LRW2. In: Theory of Cryptography - TCC 2018, Proceedings, Part II. pp. 192–222 (2018)
34. Mennink, B., Neves, S.: Encrypted Davies-Meyer and Its Dual: Towards Optimal Security Using Mirror Theory. In: Advances in Cryptology - CRYPTO 2017, Proceedings, Part III. pp. 556–583 (2017)
35. Minematsu, K.: Improved Security Analysis of XEX and LRW Modes. In: Selected Areas in Cryptography - SAC 2006, Revised Selected Papers. pp. 96–113 (2006)
36. Moch, A., List, E.: Parallelizable MACs Based on the Sum of PRPs with Security Beyond the Birthday Bound. In: Applied Cryptography and Network Security - ACNS 2019, Proceedings. pp. 131–151 (2019)
37. Patarin, J.: The "Coefficients H" Technique. In: Selected Areas in Cryptography - SAC 2008, Revised Selected Papers. pp. 328–345 (2008)
38. Patarin, J.: Introduction to Mirror Theory: Analysis of Systems of Linear Equalities and Linear Non Equalities for Cryptography. IACR Cryptol. ePrint Arch. p. 287 (2010)
39. Peyrin, T., Seurin, Y.: Counter-in-Tweak: Authenticated Encryption Modes for Tweakable Block Ciphers. In: Advances in Cryptology - CRYPTO 2016, Proceedings, Part I. pp. 33–63 (2016)
40. Procter, G.: A Note on the CLRW2 Tweakable Block Cipher Construction. IACR Cryptology ePrint Archive **2014**, 111 (2014)
41. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: Advances in Cryptology - ASIACRYPT 2004, Proceedings. pp. 16–31 (2004)
42. Schroeppel, R., Orman, H.: The Hasty Pudding Cipher. AES candidate submitted to NIST (1998), https://www.princeton.edu/~rblee/HPC/index.htm
43. Shen, Y., Peters, T., Standaert, F., Cassiers, G., Verhamme, C.: Triplex: an Efficient and One-Pass Leakage-Resistant Mode of Operation. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2022**(4), 135–162 (2022)
44. Zhang, Z., Qin, Z., Guo, C.: Just tweak! Asymptotically optimal security for the cascaded LRW1 tweakable blockcipher. Des. Codes Cryptogr. **91**(3), 1035–1052 (2023)

Supplementary Materials

A Proof of Proposition 2.1

Let $\overline{\mu} := (\mu_0 - \mu_1)/2$. Then, we have

$$\mu = \mu_0 - \overline{\mu} = \overline{\mu} + \mu_1.$$

Using Bienaymé-Chebyshev inequality, we have

$$\begin{aligned}\mathsf{Pr}\left(\mathsf{R}_0 > \mu\right) &= 1 - \mathsf{Pr}\left(\mathsf{R}_0 \leq \mu\right) \\ &\geq 1 - \mathsf{Pr}\left(\mathsf{R}_0 - \mu_0 \leq -\overline{\mu}\right) \\ &\geq 1 - \mathsf{Pr}\left(\mathsf{R}_0 - \mathsf{Ex}\left(\mathsf{R}_0\right) \leq -\overline{\mu}\right) \\ &\geq 1 - \mathsf{Pr}\left(\left|\mathsf{R}_0 - \mathsf{Ex}\left(\mathsf{R}_0\right)\right| \geq \overline{\mu}\right) \geq 1 - \frac{\sigma_0^2}{\overline{\mu}^2} \end{aligned} \tag{20}$$

and

$$\begin{aligned}\mathsf{Pr}\left(\mathsf{R}_1 > \mu\right) &\leq \mathsf{Pr}\left(\mathsf{R}_1 \geq \mu\right) \\ &\leq \mathsf{Pr}\left(\mathsf{R}_1 - \mu_1 \geq \overline{\mu}\right) \\ &\leq \mathsf{Pr}\left(\mathsf{R}_1 - \mathsf{Ex}\left(\mathsf{R}_1\right) \geq \overline{\mu}\right) \\ &\leq \mathsf{Pr}\left(\left|\mathsf{R}_1 - \mathsf{Ex}\left(\mathsf{R}_1\right)\right| \geq \overline{\mu}\right) \leq \frac{\sigma_1^2}{\overline{\mu}^2} \end{aligned} \tag{21}$$

The result then follows by subtracting (21) from (20). □

B Two Useful Inequalities From JN20

Definition B.1 ([25]). *For $r \geq s$, let $a = (a_i)_{i\in[r]}$ and $b = (b_j)_{j\in[s]}$ be two sequences over $\mathbb{N}$. We say that a compresses to b, if there exists a partition $\mathcal{P}$ of $[r]$ such that $\mathcal{P}$ contains exactly s cells, say $\mathcal{P}_1, \ldots, \mathcal{P}_s$, and $\forall i \in [s]$, $b_i = \sum_{j\in\mathcal{P}_i} a_j$.*

Proposition B.1 ([25]). *For $r \geq s$, let $a = (a_i)_{i\in[r]}$ and $b = (b_j)_{j\in[s]}$ be sequences over $\mathbb{N}$, such that a compresses to b. Then for any $n \in \mathbb{N}$, such that $2^n \geq \sum_{i=1}^{r} a_i$, we have $\prod_{i=1}^{r}(2^n)_{a_i} \geq \prod_{j=1}^{s}(2^n)_{b_j}$.*

Proposition B.2 ([25]). *For $r \geq 2$, let $c = (c_i)_{i\in[r]}$ and $d = (d_i)_{i\in[r]}$ be two sequences over $\mathbb{N}$. Let $a_1, a_2, b_1, b_2 \in \mathbb{N}$, such that $c_i \leq a_j$, $c_i + d_i \leq a_j + b_j$ for all $i \in [r]$ and $j \in [2]$, and $\sum_{i=1}^{r} d_i = b_1 + b_2$. Then, for any $n \in \mathbb{N}$, such that $a_j + b_j \leq 2^n$ for $j \in [2]$, we have $\prod_{i=1}^{r}(2^n - c_i)_{d_i} \geq (2^n - a_1)_{b_1}(2^n - a_2)_{b_2}$.*

C Proof of Claim 3.1

Preliminaries on Variance and Covariance: Recall that for any two indicator random variables $\mathbb{1}$ and $\mathbb{1}'$, the variance $\mathsf{Var}\left(\mathbb{1}\right)$ and covariance $\mathsf{Cov}\left(\mathbb{1}, \mathbb{1}'\right)$ are defined as:

$$\mathsf{Var}\left(\mathbb{1}\right) = \mathsf{Pr}\left(\mathbb{1}\right) - \mathsf{Pr}\left(\mathbb{1}\right)^2, \qquad \mathsf{Cov}\left(\mathbb{1}, \mathbb{1}'\right) = \mathsf{Pr}\left(\mathbb{1} \cdot \mathbb{1}'\right) - \mathsf{Pr}\left(\mathbb{1}\right) \cdot \mathsf{Pr}\left(\mathbb{1}'\right),$$

and for any random variable X that can be be written as a sum of indicator random variables $\sum_i \mathbb{1}_i$, we have

$$\mathsf{Var}\left(\mathsf{X}\right) = \sum_i \mathsf{Var}\left(\mathbb{1}_i\right) + \sum_{i\neq j} \mathsf{Cov}\left(\mathbb{1}_i, \mathbb{1}_j\right).$$

C.1 Upper Bounding $\mathsf{Var}\left(\mathsf{coll}_{\mathrm{id}}\right)$

Using the fact that $\mathsf{coll}_{\mathrm{id}} = \sum_{i<j} \mathbb{1}_{i,j}$, we have

$$\mathsf{Var}\left(\mathsf{coll}_{\mathrm{id}}\right) = \sum_{i<j} \mathsf{Var}\left(\mathbb{1}_{i,j}\right) + \sum_{\substack{i<j \\ k<\ell \\ \{i,j\}\neq\{k,\ell\}}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \tag{22}$$

Suppose there are ν pairs $i < j$ satisfying $i \sim j$, where we recall that $i \sim j$ if and only if $t_i = t_j \oplus \delta$.

Computing $\mathsf{Var}\left(\mathbb{1}_{i,j}\right)$. Recall that $\mathsf{Var}\left(\mathbb{1}_{i,j}\right) = \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) - \mathsf{Pr}\left(\mathbb{1}_{i,j}\right)^2$. We can have two cases, depending upon $i \sim j$, or not:
A. $i \not\sim j$: In this case, using (6), we have

$$\mathsf{Var}\left(\mathbb{1}_{i,j}\right) = \frac{1}{2^n} - \frac{1}{2^{2n}}.$$

B. $i \sim j$: In this case, using (8), we have

$$\mathsf{Var}\left(\mathbb{1}_{i,j}\right) \leq \frac{1}{2^n} + \frac{1}{2^n - 1} - \frac{1}{2^{2n}}.$$

By combining the two cases, we have

$$\sum_{i<j} \mathsf{Var}\left(\mathbb{1}_{i,j}\right) \leq \binom{q}{2}\frac{1}{2^n} + \frac{q}{2^n} - \binom{q}{2}\frac{1}{2^{2n}} \tag{23}$$

Computing $\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right)$. Recall that $\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) = \mathsf{Pr}\left(\mathbb{1}_{i,j} \cdot \mathbb{1}_{k,\ell}\right) - \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) \cdot \mathsf{Pr}\left(\mathbb{1}_{k,\ell}\right)$. We can have two cases, depending upon the size of $|\{i,j\} \cap \{k,\ell\}|$:
A. $|\{i,j\} \cap \{k,\ell\}| = 1$: Without loss of generality assume $j = k$, and consider the following subcases:

1. $\exists\, i'_1, i'_2 \in \{i,j,\ell\}$ such that $i'_1 \sim i'_2$: Note that there can be only one such (i'_1, i'_2) pair. We consider the case $i \sim \ell$, as the other two cases are relatively simpler (due to the independence of $\mathbb{1}_{i,j}$ and $\mathbb{1}_{j,\ell}$). Note that the event $\mathbb{1}_{i,j} \cdot \mathbb{1}_{j,\ell}$ is equivalent to $\mathbb{1}_{j,i} \cdot \mathbb{1}_{i,\ell}$, where of course we have abused the definition a bit as $j > i$. However, the meaning is still clear from the context. Now the events $\mathbb{1}_{j,i}$ and $\mathbb{1}_{i,\ell}$ are independent, since the j-th query uses distinct tweaks $(t_j, t_j + \delta)$. Thus, using (6) and (8), we have
$$\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{j,\ell}\right) \leq \frac{1}{2^n(2^n - 1)}.$$
2. $\forall i'_1, i'_2 \in \{i,j,k\}$, $i'_1 \not\sim i'_2$: The two events are independent and identically distributed, as all six tweaks are different. Thus, using (6), we have
$$\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{j,\ell}\right) \leq 0.$$

Now, there are at most $\nu(q-2) \le q^2/2$ triples (i,j,ℓ) that can satisfy case **A.1.**. Thus, we have

$$\sum_{\substack{i<j \\ k<\ell \\ |\{i,j\}\cap\{k,\ell\}|=1}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \le \frac{q^2}{2^{2n}} \tag{24}$$

B. $|\{i,j\} \cap \{k,\ell\}| = 0$: We handle this case depending upon the number of $(i'_1, i'_2) \in \{i,j,k,\ell\}$ such that $i'_1 \sim i'_2$. Let r be the number of such pairs. Note that r cannot be greater than 2. Thus, we have the following subcases:

1. $r = 0$: In this case, the two events are independent and identically distributed, as all eight tweaks are distinct. Thus, using (6), we have

$$\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \le 0.$$

2. $r = 1$: First, suppose $(i'_1, i'_2) \in \{(i,j),(k,\ell)\}$. Without loss of generality let $(i'_1, i'_2) = (i,j)$. Since $\{t_k, t_\ell, t_k \oplus \delta, t_\ell \oplus \delta\} \cap \{t_i, t_j\} = \emptyset$ and $k \not\sim \ell$, using (6) and (8), we get

$$\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \le \frac{1}{2^n(2^n-1)} \le \frac{2}{2^{2n}}.$$

Next, suppose $(i'_1, i'_2) \notin \{(i,j),(k,\ell)\}$. Without loss of generality, let $(i'_1, i'_2) = (i,k)$. Note that $\{t_j, t_\ell, t_j \oplus \delta, t_\ell \oplus \delta\} \cap \{t_i, t_k\} = \emptyset$. Then, by conditioning on the value of $(\mathsf{M}'_i, \mathsf{M}'_k)$, the event $\mathbb{1}_{i,j} \cdot \mathbb{1}_{k,\ell}$ holds with probability 2^{-2n}, whence using (6), we get

$$\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \le 0.$$

3. $r = 2$: Since there are at most $\nu^2 \le q^2/4$ choices for such quadruples, even a loose bound on the probability of $\mathsf{Pr}\left(\mathbb{1}_{i,j} \cdot \mathbb{1}_{k,\ell}\right)$ will suffice. In particular, we simply use $\mathsf{Pr}\left(\mathbb{1}_{i,j}\right)$. Using (8), we have

$$\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \le \frac{1}{2^n} + \frac{1}{2^n-1} - \frac{1}{2^{2n}}.$$

Finally, since there are $\nu q^2 \le q^3/2$ quadruples that satisfy **B.2.** and $\nu^2 \le q^2/4$ quadruples that satisfy **B.3.**, we get

$$\sum_{\substack{i<j \\ k<\ell \\ |\{i,j\}\cap\{k,\ell\}|=0}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \le \frac{3q^2}{2^{n+2}} + \frac{q^3}{2^{2n}} - \frac{q^2}{2^{2n+2}} \tag{25}$$

From (22)-(25) and $2 \le q \le 2^n$, we have

$$\mathsf{Var}\left(\mathsf{coll}_{\mathrm{id}}\right) \le \frac{4q^2}{2^n}. \tag{26}$$

C.2 Upper Bounding $\mathbf{Var}\,(\mathsf{coll}_{\mathrm{re}})$

The internal variables arising in the execution of $\mathsf{TNT}_{\delta,m}$ are represented by the notations from Fig. 3.2. In particular, we have $\widehat{\mathsf{M}} = \boldsymbol{\pi}_1(m)$, $\mathsf{U}_{i'} = \widehat{\mathsf{M}} \oplus t_{i'}$, $\widehat{\mathsf{U}}_{i'} = \boldsymbol{\pi}_2(\mathsf{U}_{i'})$, $\widehat{\mathsf{U}}'_{i'} = \widehat{\mathsf{U}}_{i'} \oplus \delta$, $\mathsf{U}'_{i'} = \boldsymbol{\pi}_2^{-1}(\widehat{\mathsf{U}}'_{i'})$, $\widehat{\mathsf{M}}'_{i'} = \mathsf{U}'_{i'} \oplus t_{i'}$, and $\mathsf{M}'_{i'} = \boldsymbol{\pi}_1^{-1}(\widehat{\mathsf{M}}'_{i'})$, for all $i' \in [q]$.

We have $\mathsf{coll}_{\mathrm{re}} = \sum_{i<j\in[q]} \mathbb{1}_{i,j}$, where $\mathbb{1}_{i,j}$ is the indicator random variable corresponding to the event $\mathsf{M}'_i = \mathsf{M}'_j$ in the real world. Recall that, for any $i \neq j \in [q]$, we have

$$\begin{aligned}\mathsf{Pr}\left(\mathbb{1}_{i,j}\right) &= \frac{1}{2^n-1} + \frac{1}{2^n-3} - \frac{1}{(2^n-1)(2^n-3)} \\ &= \frac{2}{2^n} - \frac{1}{2^n(2^n-1)} - \frac{3}{2^n(2^n-3)} - \frac{1}{(2^n-1)(2^n-3)}\end{aligned}$$

For simplicity we write $p := \mathsf{Pr}\left(\mathbb{1}_{i,j}\right)$. We will often employ the following inequalities

$$\frac{2}{2^n} \leq p \leq \frac{2}{2^n} + \frac{7}{2^{2n}}. \tag{27}$$

Now, we have

$$\mathsf{Var}\left(\mathsf{coll}_{\mathrm{re}}\right) = \sum_{i<j} \mathsf{Var}\left(\mathbb{1}_{i,j}\right) + \sum_{\substack{i<j \\ k<\ell \\ \{i,j\}\neq\{k,\ell\}}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \tag{28}$$

Computing $\mathbf{Var}\,(\mathbb{1}_{i,j})$. By definition, we have $\mathsf{Var}\left(\mathbb{1}_{i,j}\right) = p - p^2$, for any $i < j \in [q]$. Thus, using (27), we have

$$\sum_{i<j} \mathsf{Var}\left(\mathbb{1}_{i,j}\right) \leq \frac{q^2}{2^n} + \frac{2q^2}{2^{2n}} \tag{29}$$

Computing $\mathbf{Cov}\,(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell})$. We have

$$\begin{aligned}\mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) &= \mathsf{Pr}\left(\mathbb{1}_{i,j} \cdot \mathbb{1}_{k,\ell}\right) - \mathsf{Pr}\left(\mathbb{1}_{i,j}\right) \cdot \mathsf{Pr}\left(\mathbb{1}_{k,\ell}\right) \\ &= \mathsf{Pr}\left(\mathbb{1}_{i,j} \cdot \mathbb{1}_{k,\ell}\right) - p^2 \leq \mathsf{Pr}\left(\mathbb{1}_{i,j} \cdot \mathbb{1}_{k,\ell}\right) - \frac{4}{2^{2n}}\end{aligned} \tag{30}$$

where the last inequality follows from (27). So, from now on, we only have to handle the joint event $\mathbb{1}_{i,j,k,\ell} = \mathbb{1}_{i,j} \cdot \mathbb{1}_{k,\ell}$.

For the sake of simplicity, we perform the analysis, by conditioning on some arbitrary value of $\boldsymbol{\pi}_1(m)$, say $\widehat{m}$. Looking ahead, the final bounds will be independent of this choice, so the bounds hold unconditionally, and we take this fact for granted. Let $u_{i'} = \widehat{m} \oplus t_{i'}$, for all $i' \in [q]$. Then, $\mathsf{U}_{i'} = u_{i'}$.
As has been established before, the event $\mathbb{1}_{i,j}$ occurs, if and only if:

$\mathtt{E}_{i,j}$: $\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j = \delta$, or

$\mathtt{F}_{i,j}$: $\widehat{\mathsf{U}}_i \oplus \widehat{\mathsf{U}}_j \neq \delta$ and $\mathsf{U}'_i \oplus \mathsf{U}'_j = t_i \oplus t_j$.

Let $\mathtt{E}^2_{i,j,k,\ell}$, $\mathtt{EF}_{i,j,k,\ell}$, $\mathtt{FE}_{i,j,k,\ell}$, and $\mathtt{F}^2_{i,j,k,\ell}$ denote the joint events $(\mathtt{E}_{i,j} \cap \mathtt{E}_{k,\ell})$, $(\mathtt{E}_{i,j} \cap \mathtt{F}_{k,\ell})$, $(\mathtt{F}_{i,j} \cap \mathtt{E}_{k,\ell})$, and $(\mathtt{F}_{i,j} \cap \mathtt{F}_{k,\ell})$, respectively. Then, it is clear that $\mathbb{1}_{i,j,k,\ell}$ is a union of these four events.

The way we move forward is to count the number of all valid choices (or assignments), denoted by $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell, u'_i, u'_j, u'_k, u'_\ell)$ for $(\widehat{\mathsf{U}}_i, \widehat{\mathsf{U}}_j, \widehat{\mathsf{U}}_k, \widehat{\mathsf{U}}_\ell, \mathsf{U}'_i, \mathsf{U}'_j, \mathsf{U}'_k, \mathsf{U}'_\ell)$ that satisfy the event in focus. Then, the probability of the event is simply this count times $1/(2^n)_\alpha$, where α will denote a lower bound on the number of distinct elements in $\{u_i, u_j, u_k, u_\ell, u'_i, u'_j, u'_k, u'_\ell\}$ for the event in focus.

Now, we can have two cases depending upon $r := |\{i,j\} \cap \{k,\ell\}|$:

A. $r = 1$: Without loss of generality assume $j = k$. Then,

$$\begin{aligned}\mathsf{Pr}\left(\mathbb{1}_{i,j,j,\ell}\right) &= \mathsf{Pr}\left(\mathtt{E}^2_{i,j,j,\ell} \cup \mathtt{EF}_{i,j,j,\ell} \cup \mathtt{FE}_{i,j,j,\ell} \cup \mathtt{F}^2_{i,j,j,\ell}\right) \\ &\leq \mathsf{Pr}\left(\mathtt{E}^2_{i,j,j,\ell}\right) + \mathsf{Pr}\left(\mathtt{EF}_{i,j,j,\ell}\right) + \mathsf{Pr}\left(\mathtt{FE}_{i,j,j,\ell}\right) + \mathsf{Pr}\left(\mathtt{F}^2_{i,j,j,\ell}\right) \\ &= \mathsf{Pr}\left(\mathtt{EF}_{i,j,j,\ell}\right) + \mathsf{Pr}\left(\mathtt{FE}_{i,j,j,\ell}\right) + \mathsf{Pr}\left(\mathtt{F}^2_{i,j,j,\ell}\right) \qquad (31)\end{aligned}$$

where the last equality follows from the fact that $t_i \oplus \delta = t_j = t_\ell \oplus \delta$ if and only if $t_i = t_\ell$, which is impossible. We handle the three summands one by one:

1. Probability of $\mathtt{EF}_{i,j,j,\ell}$: Any valid choice $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_\ell, u'_i, u'_j, u'_j, u'_\ell)$ must satisfy

- $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_\ell)$ is pairwise distinct,
- $\widehat{u}_i \oplus \widehat{u}_j = \delta$ and $\widehat{u}_\ell \notin \{\widehat{u}_i, \widehat{u}_j\}$,
- $(u'_i, u'_j) = (u_j, u_i)$,
- $u'_\ell = u'_j \oplus t_j \oplus t_\ell \notin \{u'_i, u'_j, u_\ell\} = \{u_i, u_j, u_\ell\}$,
- (u'_i, u'_j, u'_ℓ) is pairwise distinct.

The first three conditions are obvious. In the fourth condition, $u'_\ell \neq u_\ell$ follows from $\delta \neq 0^n$. Now, $\widehat{u}_i$ has 2^n choices, $\widehat{u}_j = \widehat{u}_i \oplus \delta$, and $\widehat{u}_\ell \notin \{\widehat{u}_i, \widehat{u}_j\}$ has obviously $(2^n - 2)$ choices. Once we fix $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_\ell)$, u'_ℓ is fixed. Thus, there are at most $2^n(2^n - 2)$ choices. Further, from condition 3, we have $|\{u_i, u_j, u_\ell, u'_\ell\}| = 4$. Thus, each valid choice occurs with at most $1/(2^n)_4$ probability, as at least 4 variables are sampled in a WOR manner from $\{0,1\}^n$. Thus, we have

$$\begin{aligned}\mathsf{Pr}\left(\mathtt{EF}_{i,j,j,\ell}\right) &\leq \frac{1}{(2^n - 1)(2^n - 3)} \\ &\leq \frac{1}{2^{2n}}\left(1 + \frac{1}{2^n - 1}\right)\left(1 + \frac{3}{2^n - 3}\right) \\ &\leq \frac{1}{2^{2n}} + \frac{8}{2^{3n}} + \frac{12}{2^{4n}} \qquad (32)\end{aligned}$$

2. Probability of $\mathtt{FE}_{i,j,j,\ell}$: By symmetry, we have

$$\mathsf{Pr}\left(\mathtt{FE}_{i,j,j,\ell}\right) \leq \frac{1}{2^{2n}} + \frac{8}{2^{3n}} + \frac{12}{2^{4n}} \qquad (33)$$

3. Probability of $\mathtt{F}^2_{i,j,j,\ell}$: Any valid choice $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_j, \widehat{u}_\ell, u'_i, u'_j, u'_j, u'_\ell)$ must satisfy
- $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_\ell)$ is pairwise distinct,
- $\widehat{u}_i \oplus \widehat{u}_j \neq \delta$ and $\widehat{u}_j \oplus \widehat{u}_\ell \neq \delta$,
- $u'_i = u'_j \oplus t_i \oplus t_j \notin \{u_i, u_j\}$,
- $u'_\ell = u'_j \oplus t_j \oplus t_\ell \notin \{u_j, u_\ell\}$,
- (u'_i, u'_j, u'_ℓ) is pairwise distinct.

Now, $\widehat{u}_i \oplus \widehat{u}_\ell = \delta$ (which is possible) a valid assignment would have $(u'_i, u'_\ell) = (u_\ell, u_i)$. But, this implies that this assignment also satisfies $\mathtt{E}_{i,\ell}$. Accordingly, we refine the objective as

$$\mathsf{Pr}\left(\mathtt{F}^2_{i,j,j,\ell}\right) \leq \mathsf{Pr}\left(\mathtt{F}^2_{i,j,j,\ell} \cap \mathtt{E}_{i,\ell}\right) + \mathsf{Pr}\left(\mathtt{F}^2_{i,j,j,\ell} \mid \neg\mathtt{E}_{i,\ell}\right)$$

For the first summand we have at most $2^n(2^n-2)$ valid assignments, each occurring with at most $1/(2^n)_4$ probability, and for the second summand we have at most $2^n(2^n-1)(2^n-3)(2^n-4)$ valid assignments, each occurring with at most $1/(2^n)_6$ probability. Thus, we have

$$\mathsf{Pr}\left(\mathtt{F}^2_{i,j,j,\ell}\right) \leq \frac{2}{2^{2n}} + \frac{26}{2^{3n}} + \frac{92}{2^{4n}} \tag{34}$$

On combining (30)-(34), we get

$$\sum_{\substack{i<j \\ k<\ell \\ r=1}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \leq \frac{7q^3}{2^{3n}} + \frac{20q^3}{2^{4n}} \tag{35}$$

B. $r = 0$: In this case we have

$$\begin{aligned}\mathsf{Pr}\left(\mathbb{1}_{i,j,k,\ell}\right) &= \mathsf{Pr}\left(\mathtt{E}^2_{i,j,k,\ell} \cup \mathtt{EF}_{i,j,k,\ell} \cup \mathtt{FE}_{i,j,k,\ell} \cup \mathtt{F}^2_{i,j,k,\ell}\right) \\ &\leq \mathsf{Pr}\left(\mathtt{E}^2_{i,j,k,\ell}\right) + \mathsf{Pr}\left(\mathtt{EF}_{i,j,k,\ell}\right) + \mathsf{Pr}\left(\mathtt{FE}_{i,j,k,\ell}\right) + \mathsf{Pr}\left(\mathtt{F}^2_{i,j,k,\ell}\right) \end{aligned} \tag{36}$$

We handle the four summands one by one:

1. Probability of $\mathtt{E}^2_{i,j,k,\ell}$: Any valid choice $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell, u'_i, u'_j, u'_k, u'_\ell)$ must satisfy
- $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell)$ is pairwise distinct,
- $\widehat{u}_i \oplus \widehat{u}_j = \delta$ and $\widehat{u}_k \oplus \widehat{u}_\ell = \delta$,
- $(u'_i, u'_j, u'_k, u'_\ell) = (u_j, u_i, u_\ell, u_k)$,

Now, $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell)$ can be fixed in at most $2^n(2^n-2)$ ways, as fixing $\widehat{u}_i$ fixes $\widehat{u}_j$, and fixing $(\widehat{u}_i, \widehat{u}_j)$ leaves (2^n-2) choices for $\widehat{u}_k$ and this fixes $\widehat{u}_\ell$. With this the full assignment is fixed. Further, each such assignment holds with at most $1/(2^n)_4$ probability. Thus, we have

$$\mathsf{Pr}\left(\mathtt{E}^2_{i,j,k,\ell}\right) \leq \frac{1}{2^{2n}} + \frac{8}{2^{3n}} + \frac{12}{2^{4n}} \tag{37}$$

2. Probability of $\mathtt{EF}_{i,j,k,\ell}$: Any valid choice $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell, u'_i, u'_j, u'_k, u'_\ell)$ must satisfy

- $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell)$ is pairwise distinct,
- $\widehat{u}_i \oplus \widehat{u}_j = \delta$ and $\widehat{u}_k \oplus \widehat{u}_\ell \neq \delta$,
- $(u'_i, u'_j) = (u_j, u_i)$,
- $u'_\ell = u'_k \oplus t_k \oplus t_\ell \notin \{u_i, u_j, u_k, u_\ell\}$,
- $u'_k \notin \{u_i, u_j, u_k, u_\ell\}$,
- $(u'_i, u'_j, u'_k, u'_\ell)$ is pairwise distinct.

The fourth condition follows from $\delta \neq 0^n$, $\widehat{u}_\ell \neq \delta \oplus \widehat{u}_k$, and the fact that $u'_\ell = u_j$ (res. $u'_\ell = u_i$) would imply $\widehat{u}_\ell \oplus \delta = \widehat{u}_j = \widehat{u}_i \oplus \delta$ (res. $\widehat{u}_\ell \oplus \delta = \widehat{u}_i = \widehat{u}_j \oplus \delta$), which is impossible. Similar argument holds for condition 5. Thus, in this case, 6 distinct values are sampled in a WOR manner from $\{0,1\}^n$. There are at most $2^n(2^n-2)(2^n-3)(2^n-4)$ valid choices, each holding with at most $1/(2^n)_6$ probability. Thus, we have

$$\mathsf{Pr}\left(\mathtt{EF}_{i,j,k,\ell}\right) \leq \frac{1}{2^{2n}} + \frac{12}{2^{3n}} + \frac{20}{2^{4n}} \tag{38}$$

3. Probability of $\mathtt{FE}_{i,j,k,\ell}$: By symmetry, we have

$$\mathsf{Pr}\left(\mathtt{FE}_{i,j,k,\ell}\right) \leq \frac{1}{2^{2n}} + \frac{12}{2^{3n}} + \frac{20}{2^{4n}} \tag{39}$$

4. Probability of $\mathtt{F}^2_{i,j,k,\ell}$: Any valid choice $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_j, \widehat{u}_\ell, u'_i, u'_j, u'_j, u'_\ell)$ must satisfy

- $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell)$ is pairwise distinct,
- $\widehat{u}_i \oplus \widehat{u}_j \neq \delta$ and $\widehat{u}_k \oplus \widehat{u}_\ell \neq \delta$,
- $u'_i = u'_j \oplus t_i \oplus t_j \notin \{u_i, u_j\}$,
- $u'_j \notin \{u_i, u_j\}$,
- $u'_\ell = u'_j \oplus t_j \oplus t_\ell \notin \{u_j, u_\ell\}$,
- $u'_k \notin \{u_k, u_\ell\}$,
- $(u'_i, u'_j, u'_k, u'_\ell)$ is pairwise distinct.

Further, a valid choice also satisfies one of the following seven conditions:

i. $\widehat{u}_i \oplus \delta = \widehat{u}_k$, $\widehat{u}_j \oplus \delta = \widehat{u}_\ell$,
ii. $\widehat{u}_i \oplus \delta = \widehat{u}_\ell$, $\widehat{u}_j \oplus \delta = \widehat{u}_k$,
iii. $\widehat{u}_i \oplus \delta = \widehat{u}_k$, $\widehat{u}_j \oplus \delta \neq \widehat{u}_\ell$,
iv. $\widehat{u}_i \oplus \delta \neq \widehat{u}_k$, $\widehat{u}_j \oplus \delta = \widehat{u}_\ell$,
v. $\widehat{u}_i \oplus \delta = \widehat{u}_\ell$, $\widehat{u}_j \oplus \delta \neq \widehat{u}_k$,
vi. $\widehat{u}_i \oplus \delta \neq \widehat{u}_\ell$, $\widehat{u}_j \oplus \delta = \widehat{u}_k$,
vii. $\{\widehat{u}_i \oplus \delta, \widehat{u}_j \oplus \delta\} \cap \{\widehat{u}_k, \widehat{u}_\ell\}$.

Now, we can have one of the two subcases based on whether $\lambda := t_i \oplus t_j \oplus t_k \oplus t_\ell = 0^n$, or not:

a. $\lambda = 0^n$: Observe that, in this case, conditions iii-vi are not satisfiable. For instance, suppose $\widehat{u}_i \oplus \delta = \widehat{u}_k$. Then, $u'_j = u'_i \oplus t_i \oplus t_j = u_k \oplus t_k \oplus t_\ell = u_\ell$ which implies $\widehat{u}_j \oplus \delta = \widehat{u}_\ell$. Thus, only conditions i, ii, and vii are possible. Now, if condition i (res. condition ii) satisfies then we must have $u'_i = u_k$ (res. $u'_i = u_\ell$), $u'_j = u_\ell$ (res. $u'_j = u_k$). Thus, in both the cases fixing $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell)$ fixes the whole assignment. Further, $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell)$ can be fixed in at most

$2^n(2^n-2)$ ways, and each such assignment holds with at most $1/(2^n)_4$ probability. On the other hand, if condition vii satisfies then fixing $(\widehat{u}_i, \widehat{u}_j, \widehat{u}_k, \widehat{u}_\ell, u'_i, u'_k)$ fixes the full assignment. So, in this case we have at most $(2^n)_6$ choices, and each such choice holds with at most $1/(2^n)_8$ probability. Thus, when $t_i \oplus t_j = t_k \oplus t_\ell$, we have

$$\mathsf{Pr}\left(\mathtt{F}^2_{i,j,k,\ell}\right) \leq \frac{12}{2^{2n}} \tag{40}$$

b. $\lambda \neq 0^n$: Contrary to case **a.**, it can be easily verified that condition i and ii are not satisfiable in this case. Now, if condition iii-vi is satisfied, then there are at most $2^n(2^n-2)(2^n-3)$ valid choices, each holding with at most $1/(2^n)_6$ probability. On the other hand, if condition vii is satisfied, then there are at most $(2^n)_6$ valid choices and each choice occurs with $1/(2^n)_8$ probability. Thus, when $t_i \oplus t_j \neq t_k \oplus t_\ell$, we have

$$\mathsf{Pr}\left(\mathtt{F}^2_{i,j,k,\ell}\right) \leq \frac{1}{2^{2n}} + \frac{58}{2^{3n}} + \frac{168}{2^{4n}} \tag{41}$$

To summarize the above computations, we have

$$\sum_{\substack{i<j \\ k<\ell \\ r=0}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) = \sum_{\substack{i<j \\ k<\ell \\ r=0 \\ \lambda=0}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) + \sum_{\substack{i<j \\ k<\ell \\ r=0 \\ \lambda\neq 0}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \tag{42}$$

Using (30), (36)-(39), and (40), we have

$$\sum_{\substack{i<j \\ k<\ell \\ r=0 \\ \lambda=0}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \leq \frac{2q^3}{2^{2n}} + \frac{6q^3}{2^{3n}} + \frac{9q^3}{2^{4n}} \tag{43}$$

and using (30), (36)-(39) and (41), we have

$$\sum_{\substack{i<j \\ k<\ell \\ r=0 \\ \lambda\neq 0}} \mathsf{Cov}\left(\mathbb{1}_{i,j}, \mathbb{1}_{k,\ell}\right) \leq \frac{4q^4}{2^{3n}} + \frac{10q^4}{2^{4n}} \tag{44}$$

On combining (28), (29), (35), and (42)-(44), we have

$$\mathsf{Var}\left(\mathsf{coll}_{\mathrm{re}}\right) \leq \frac{q^2}{2^n} + \frac{2q^2}{2^{2n}} + \frac{13q^3}{2^{3n}} + \frac{29q^3}{2^{4n}} + \frac{2q^3}{2^{2n}} + \frac{4q^4}{2^{3n}} + \frac{10q^4}{2^{4n}} \tag{45}$$

The result follows from $q \leq 2^n$, and $n \geq 4$. □

D Some Results From JN20 on Hash Functions

Throughout this section, we fix $t^q = (t_1, \ldots, t_q) \in (\mathcal{T})_q$. Let $\mathcal{H}$ be a (τ, n)-hash function family with ϵ-AUHF property. A pair of distinct elements $t_i, t_j \in \mathtt{S}(t^q)$ is said to be *colliding* for a function $h \in \mathcal{H}$, if $h(t_i) = h(t_j)$. Then, for a randomly chosen hash function $\mathsf{H} \xleftarrow{\$} \mathcal{H}$, the probability of having at least one colliding pair in t^q is at most $\binom{q}{2} \cdot \epsilon$. This is straightforward from the union bound.

Lemma D.1 (Alternating Collisions Lemma [25]). *Suppose $\mathsf{H}_1, \mathsf{H}_2$ are two uniformly and independently drawn functions from an ϵ-AUHF $\mathcal{H}$ and $t^q \in (\mathcal{T})_q$. Then,*

$$\mathsf{Pr}\left(\exists^* i, j, k, l \in [q], \mathsf{H}_1(t_i) = \mathsf{H}_1(t_j) \wedge \mathsf{H}_1(t_k) = \mathsf{H}_1(t_l) \wedge \mathsf{H}_2(t_j) = \mathsf{H}_2(t_k)\right) \leq q^2 \epsilon^{1.5}.$$

Lemma D.2 (Alternating Events Lemma [25]). *Let $\mathsf{X}^q = (\mathsf{X}_1, \ldots, \mathsf{X}_q)$ be a q-tuple of random variables. Suppose for all $i < j \in [q]$, $\mathtt{E}_{i,j}$ are events associated with X_i and X_j, possibly dependent. Each event holds with probability at most ϵ. Moreover, for any distinct $i, j, k, l \in [q]$, $\mathtt{F}_{i,j,k,l}$ are events associated with X_i, X_j, X_k and X_l, which holds with probability at most ϵ'. Moreover, the collection of events $(\mathtt{F}_{i,j,k,l})_{i,j,k,l}$ is independent with the collection of event $(\mathtt{E}_{i,j})_{i,j}$. Then,*

$$\mathsf{Pr}\left(\exists^* i, j, k, l \in [q], \mathtt{E}_{i,j} \ \wedge \ \mathtt{E}_{k,l} \ \wedge \ \mathtt{F}_{i,j,k,l}\right) \leq q^2 \cdot \epsilon \cdot \sqrt{\epsilon'}$$

Let $\mathsf{X}^q = \mathsf{H}(t^q)$. We define an equivalence relation $\sim$ on $[q]$ as: $\alpha \sim \beta$ if and only if $\mathsf{X}_\alpha = \mathsf{X}_\beta$ (i.e. $\sim$ is simply the multicollision relation). Let $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_r$ denote those equivalence classes of $[q]$ corresponding to $\sim$, such that $\nu_i = |\mathcal{P}_i| \geq 2$ for all $i \in [r]$.

Lemma D.3 ([25]). *Let C denote the number of colliding pairs in X^q. Then, we have*

$$\mathsf{Ex}\left(\sum_{i=1}^{r} {\nu_i}^2\right) \leq 2q^2\epsilon.$$

Corollary D.1 ([36,25]). *Let $\nu_{\max} = \max\{\nu_i : i \in [r]\}$. Then, for some $a \geq 2$, we have*

$$\mathsf{Pr}\left(\nu_{\max} \geq a\right) \leq \frac{2q^2\epsilon}{a^2}.$$

E Extension of Patarin's Mirror Theory From JN20

We will use the Mennink and Neves interpretation [34] of mirror theory. Let $q \geq 1$ and let $\mathcal{L}$ be the system of linear equations

$$\{e_1 : Y_1 \oplus V_1 = \delta_1, \quad e_2 : Y_2 \oplus V_2 = \delta_2, \quad \ldots, \quad e_q : Y_q \oplus V_q = \delta_q\}$$

where Y^q and V^q are unknowns, and $\delta^q \in (\{0,1\}^n)^q$ are constants. In addition there are (in)equality restrictions on Y^q and V^q, which uniquely determine $\mathtt{S}Y^q$

and $\mathtt{S}V^q$. We assume that $\mathtt{S}(Y^q)$ and $\mathtt{S}(V^q)$, are indexed in an arbitrary order by the index sets $[q_Y]$ and $[q_V]$, where $q_Y = |\mathtt{S}(Y^q)|$ and $q_V = |\mathtt{S}(V^q)|$. This assumption is without any loss of generality as this does not affect the system $\mathcal{L}$. Given such an ordering, we can view $\mathtt{S}(Y^q)$ and $\mathtt{S}(V^q)$ as ordered sets $\{Y'_1, \ldots, Y'_{q_Y}\}$ and $\{V'_1, \ldots, V'_{q_V}\}$, respectively. We define two surjective index mappings:

$$\varphi_Y : \begin{cases} [q] \to [q_Y] \\ i \mapsto j \text{ if and only if } Y_i = Y'_j. \end{cases} \qquad \varphi_V : \begin{cases} [q] \to [q_V] \\ i \mapsto k \text{ if and only if } V_i = V'_k. \end{cases}$$

It is easy to verify that $\mathcal{L}$ is uniquely determined by $(\varphi_Y, \varphi_V, \delta^q)$, and vice-versa. Consider a labeled bipartite graph $\mathcal{G}(\mathcal{L}) = ([q_Y], [q_V], \mathcal{E})$ associated with $\mathcal{L}$, where $\mathcal{E} = \{(\varphi_Y(i), \varphi_V(i), \delta_i) : i \in [q]\}$, δ_i being the label of edge. Clearly, each equation in $\mathcal{L}$ corresponds to a unique labeled edge (assuming no duplicate equations). We give three definitions with respect to the system $\mathcal{L}$ using $\mathcal{G}(\mathcal{L})$.

Definition E.1 (cycle-freeness). *$\mathcal{L}$ is said to be cycle-free if and only if $\mathcal{G}(\mathcal{L})$ is acyclic.*

Definition E.2 ($\xi_{\max}$-component). *Two distinct equations (or unknowns) in $\mathcal{L}$ are said to be in the same component if and only if the corresponding edges (res. vertices) in $\mathcal{G}(\mathcal{L})$ are in the same component. The size of any component $\mathcal{C}$ in $\mathcal{L}$, denoted $\xi(\mathcal{C})$, is the number of vertices in the corresponding component of $\mathcal{G}(\mathcal{L})$, and the maximum component size is denoted by $\xi_{\max}(\mathcal{L})$ (or simply $\xi_{\max}$).*

Definition E.3 (non-degeneracy). *$\mathcal{L}$ is said to be non-degenerate if and only if there does not exist a path of even length at least 2 in $\mathcal{G}(\mathcal{L})$ such that the labels along the edges on this path sum up to zero.*

Isolated and Star Components: In an edge-labeled bipartite graph $\mathcal{G} = (\mathcal{Y}, \mathcal{V}, \mathcal{E})$, an edge (y, v, δ) is called *isolated* edge if both y and v have degree 1. A component $\mathcal{S}$ of $\mathcal{G}$ is called star, if $\xi(\mathcal{S}) \geq 3$ and there exists a unique vertex v in $\mathcal{S}$ with degree $\xi(\mathcal{S}) - 1$. We call v the center of $\mathcal{S}$. Further, we call $\mathcal{S}$ a $\mathcal{Y}$-$\star$ (res. $\mathcal{V}$-$\star$) component if its center lies in $\mathcal{Y}$ (res. $\mathcal{V}$).

Mirror Theory for Tweakable Permutation Setting. Consider a system of equation $\mathcal{L}$

$$\{e_1 : Y_1 \oplus V_1 = \delta_1, \quad e_2 : Y_2 \oplus V_2 = \delta_2, \quad \ldots, \quad e_q : Y_q \oplus V_q = \delta_q\},$$

such that each component in $\mathcal{G}(\mathcal{L})$ is either an isolated edge or a star. Let c_1, c_2, and c_3 denote the number of components of isolated, $\mathcal{Y}$-$\star$, and $\mathcal{V}$-$\star$ types, respectively. Let q_1, q_2, and q_3 denote the number of equations of isolated, $\mathcal{Y}$-$\star$, and $\mathcal{V}$-$\star$ types, respectively. Therefore, $c_1 = q_1$. Note that the equations in $\mathcal{L}$ can be arranged in any arbitrary order without affecting the number of solutions. For the sake of simplicity, we fix the ordering in such a way that all isolated edges occur first, followed by the star components. Let $(\delta'_1, \delta'_2, \cdots, \delta'_s)$ be an arbitrary

ordering of $\mathtt{S}(\delta^q)$, and for all $i \in [s]$, let ν_i denote the multiplicity of δ'_i in the multiset $\mathtt{M}(\delta^q)$, i.e., $s \leq q$ and $\sum_{i=1}^{s} \nu_i = q$.
In [25], Jha and Nandi proved the following result.

Theorem E.1 ([25]). *Let $\mathcal{L}$ be the system of linear equations as described above with $q < 2^{n-2}$ and $\xi_{\max} q \leq 2^{n-1}$. Then, the number of tuples $(y_1, \ldots, y_{q_Y}, v_1, \ldots, v_{q_V})$ that satisfy $\mathcal{L}$, denoted h_q, such that $y_i \neq y_j$ and $v_i \neq v_j$, for all $i \neq j$, satisfies:*

$$h_q \geq \left(1 - \frac{13q^4}{2^{3n}} - \frac{2q^2}{2^{2n}} - \left(\sum_{i=1}^{c_2+c_3} \eta_{c_1+i}^2\right) \frac{4q^2}{2^{2n}}\right) \times \frac{(2^n)_{q_1+c_2+q_3}(2^n)_{q_1+q_2+c_3}}{\prod_{i\in[s]}(2^n)_{\nu_i}},$$

where $\eta_j = \xi_j - 1$ and ξ_j denotes the size (number of vertices) of the j-th component, for all $j \in [c_1 + c_2 + c_3]$.

F Proof of Theorem 5.1

The statement is vacuously true for $q \geq 2^{n/2}$. We will use the Expectation method (see Lemma 2.1) to prove the statement for $1 \leq q < 2^{n/2}$.

Let $\mathcal{O}_0$ and $\mathcal{O}_1$ be the oracles corresponding to TNT and a tweakable random permutation $\widetilde{\boldsymbol{\pi}}$, respectively. If $(\mathsf{T}_i, \mathsf{M}_i)$ is the encryption query with a tweak T_i we write the response as $\widehat{\mathsf{C}}_i$. Similarly, if $(\mathsf{T}_i, \widehat{\mathsf{C}}_i)$ is the decryption query with a tweak T_i we write the response as M_i. After all queries have been made, the two oracles release some additional data to the adversary, who is obviously free to ignore this additional information, $\widehat{\mathsf{M}}^q$ and C^q.

In the real world, $\widehat{\mathsf{M}}^q$ and C^q correspond to the output of $\boldsymbol{\pi}_1$ and input of $\boldsymbol{\pi}_3$, respectively, and thus they are well defined from the definition of TNT. The real world transcript is thus defined as the tuple

$$\Theta_1 := (\mathsf{T}^q, \mathsf{M}^q, \widehat{\mathsf{C}}^q, \widehat{\mathsf{M}}^q, \mathsf{C}^q).$$

In the ideal system $\widetilde{\boldsymbol{\pi}}$, we sample $\widehat{\mathsf{M}}^q, \mathsf{C}^q$ as follows for all $i \in [q]$:

1. $\widehat{\mathsf{M}}_i = \widehat{\mathsf{M}}_j$ whenever $\mathsf{M}_i = \mathsf{M}_j$ for $j < i$. Otherwise (for all $j < i$, $\mathsf{M}_j \neq \mathsf{M}_i$), we sample
$$\widehat{\mathsf{M}}_i \xleftarrow{\$} \{0,1\}^n \setminus \mathsf{S}(\widehat{\mathsf{M}}^{[i-1]}).$$
2. $\mathsf{C}_i = \mathsf{C}_j$ whenever $\widehat{\mathsf{C}}_j = \widehat{\mathsf{C}}_i$ for $j < i$. Otherwise (for all $j < i$, $\widehat{\mathsf{C}}_j \neq \widehat{\mathsf{C}}_i$), we sample
$$\mathsf{C}_i \xleftarrow{\$} \{0,1\}^n \setminus \mathsf{S}(\mathsf{C}^{[i-1]}).$$

The ideal world transcript is defined as

$$\Theta_0 := (\mathsf{T}^q, \mathsf{M}^q, \widehat{\mathsf{C}}^q, \widehat{\mathsf{M}}^q, \mathsf{C}^q).$$

Note that we use the same notation to denote the random variables in both the worlds. However, their probability distributions will be unambiguously determined at the time of probability computations.

Bad Transcript and Its Analysis: Let $u^q := \widehat{m}^q \oplus t^q$, and $\widehat{u}^q := c^q \oplus t^q$. A transcript $(t^q, m^q, \widehat{c}^q, \widehat{m}^q, c^q)$ is called *bad* if and only if

- $\exists i < j \in [q]$ such that $u_i = u_j$; or
- $\exists i < j \in [q]$ such that $\widehat{u}_i = \widehat{u}_j$;

Let Ω_{bad} denote the set of all bad transcripts. Now, $\Theta_0 \in \Omega_{\mathsf{bad}}$ if either for some $i < j$, $\widehat{\mathsf{M}}_i + \mathsf{T}_i = \widehat{\mathsf{M}}_j + \mathsf{T}_j$ or $\mathsf{C}_i + \mathsf{T}_i = \mathsf{C}_j + \mathsf{T}_j$. It is easy to see that for any fixed $i < j$, $\mathsf{Pr}\left(\widehat{\mathsf{M}}_i + \mathsf{T}_i = \widehat{\mathsf{M}}_j + \mathsf{T}_j\right) \leq (2^n - 1)^{-1}$ and similarly for the other case. So, by using the union bound,

$$\mathsf{Pr}\left(\in \Omega_{\mathsf{bad}}\right) \leq \frac{q(q-1)}{2^n - 1} \leq \frac{q^2}{2^n}.$$

Analysis of Good Transcripts: For a good transcript $\tau = (t^q, m^q, \widehat{c}^q, \widehat{m}^q, c^q)$, we know that $(m^q, \widehat{m}^q)$, $(c^q, \widehat{c}^q)$, and $(u^q, \widehat{u}^q)$ are permutation consistent and hence for the real world we have

$$\begin{aligned}\mathsf{Pr}\left(\Theta_1 = \omega\right) &= \mathsf{Pr}\left(\boldsymbol{\pi}_1(m^q) = \widehat{m}^q\right) \times \mathsf{Pr}\left(\boldsymbol{\pi}_2(u^q) = \widehat{u}^q\right) \times \mathsf{Pr}\left(\boldsymbol{\pi}_3(c^q) = \widehat{c}^q\right) \\ &= \frac{1}{(2^n)_r} \times \frac{1}{(2^n)_q} \times \frac{1}{(2^n)_s}\end{aligned}$$

where r and s denote the size of $\mathsf{S}(m^q)$ and $\mathsf{S}(\widehat{c}^q)$ respectively. In the ideal world, we have,

$$\begin{aligned}\mathsf{Pr}\left(\Theta_0 = \omega\right) &= \mathsf{Pr}\left(\widetilde{\boldsymbol{\pi}}(t^q, m^q) = \widehat{c}^q\right) \times \frac{1}{(2^n)_r} \times \frac{1}{(2^n)_s} \\ &\leq \frac{1}{(2^n)_q} \times \frac{1}{(2^n)_r} \times \frac{1}{(2^n)_s},\end{aligned}$$

where the final inequality follows from the fact that $\mathsf{Pr}\left(\widetilde{\boldsymbol{\pi}}(t^q, m^q) = \widehat{c}^q\right)$ maximizes when $t_i = t_j$ for all $1 \leq i < j \leq q$. The result follows from the Expectation method by setting $\epsilon_{\mathsf{ratio}}$ to be a zero function.

G Proof of Theorem 6.1

Note that we are in the information-theoretic setting. In other words, we consider computationally unbounded distinguisher $\mathbf{A}$. Without loss of generality, we assume that $\mathbf{A}$ is deterministic and non-trivial.

G.1 Oracle Description

The two oracles of interest are: $\mathcal{O}_1$, the real oracle, that implements $\mathsf{LRW}+$; and, $\mathcal{O}_0$, the ideal oracle, that implements $\widetilde{\boldsymbol{\pi}} \xleftarrow{\$} \widetilde{\mathsf{Perm}}(\tau, n)$. We consider an extended version of these oracles, the one in which they release some additional information. We use notations analogously as given in Figure 6.1 to describe the transcript generated by $\mathbf{A}$'s interaction with its oracle.

Description of the real oracle, $\mathcal{O}_1$: The real oracle $\mathcal{O}_1$ faithfully runs $glrw$. We denote the transcript random variable generated by $\mathbf{A}$'s interaction with $\mathcal{O}_1$ by the usual notation Θ_1, which is an 11-ary q-tuple

$$(\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q, \mathsf{X}^q, \mathsf{Y}^q, \mathsf{V}^q, \mathsf{U}^q, \Delta^q, \widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H}),$$

defined as follows: The initial transcript consists of $(\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q)$, where for all $i \in [q]$:

T_i : i-th tweak value M_i : i-th plaintext value C_i : i-th ciphertext value,

where, $\mathsf{C}^q = \mathsf{LRW}+(\mathsf{T}^q, \mathsf{M}^q)$. At the end of the query-response phase $\mathcal{O}_1$ releases some additional information $(\mathsf{X}^q, \mathsf{Y}^q, \mathsf{V}^q, \mathsf{U}^q, \Delta^q, \widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H})$, such that for all $i \in [q]$:

- $(\mathsf{X}_i, \mathsf{Y}_i)$: i-th input-output pair for $\boldsymbol{\pi}_1$,
- $(\mathsf{V}_i, \mathsf{U}_i)$: i-th input-output pair for $\boldsymbol{\pi}_2$,
- Δ_i: i-th internal masking, $\widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H}$: are the hash keys.

Note that X^q, U^q, and Δ^q are completely determined by the hash keys $\widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H}$, and the initial transcript $(\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q)$. We include them anyhow for the sake of convenience.

Description of the ideal oracle, $\mathcal{O}_0$: The ideal oracle $\mathcal{O}_0$ has access to $\widetilde{\boldsymbol{\pi}}$. Since $\mathcal{O}_1$ releases some additional information, $\mathcal{O}_0$ must generate these values as well. The ideal transcript random variable Θ_0 is also an 11-ary q-tuple

$$(\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q, \mathsf{X}^q, \mathsf{Y}^q, \mathsf{V}^q, \mathsf{U}^q, \Delta^q, \widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H}),$$

defined below. The initial transcript consists of $(\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q)$, where for all $i \in [q]$:

T_i : i-th tweak value M_i : i-th plaintext value C_i : i-th ciphertext value,

where $\mathsf{C}^q = \widetilde{\boldsymbol{\pi}}(\mathsf{T}^q, \mathsf{M}^q)$. Once the query-response phase is over $\mathcal{O}_0$ first samples $(\widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H}) \leftarrow\!\$\ \mathsf{KG}\left(\widehat{\mathcal{H}}\right)$, and then computes $(\mathsf{X}^q, \mathsf{U}^q, \Delta^q)$, as follows:

$$\mathsf{X}^q := \widetilde{\mathbf{H}}_1(\mathsf{T}^q, \mathsf{M}^q) \qquad \mathsf{U}^q := \widetilde{\mathbf{H}}_2(\mathsf{T}^q, \mathsf{C}^q) \qquad \Delta^q := \mathbf{H}(\mathsf{T}^q).$$

Note that the conditional distributions of $(\mathsf{X}^q, \mathsf{U}^q, \Delta^q, \widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H})$, given $(\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q)$ is identical in both the worlds. This means that X^q, U^q, and Δ^q are defined honestly.

Given the partial transcript $\Theta_0' := (\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q, \mathsf{X}^q, \mathsf{U}^q, \Delta^q, \widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H})$ we wish to characterize the hash key $\widehat{\mathbf{H}} := (\widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H})$ as good or bad. We write $\widehat{\mathcal{H}}_{\mathsf{bad}}$ for the set of bad hash keys, and $\widehat{\mathcal{H}}_{\mathsf{good}} := \widehat{\mathcal{H}} \setminus \widehat{\mathcal{H}}_{\mathsf{bad}}$. We say that the hash key $\widehat{\mathbf{H}} \in \widehat{\mathcal{H}}_{\mathsf{bad}}$ (or $\widehat{\mathbf{H}}$ is bad) if and only if one of the following predicates is true:

1. $\mathtt{H}_1$: $\exists^* i, j \in [q]$ such that $\mathsf{X}_i = \mathsf{X}_j \wedge \mathsf{U}_i = \mathsf{U}_j$.
2. $\mathtt{H}_2$: $\exists^* i, j \in [q]$ such that $\mathsf{X}_i = \mathsf{X}_j \wedge \Delta_i = \Delta_j$.

3. $\mathtt{H}_3$: $\exists^* i, j \in [q]$ such that $\mathsf{U}_i = \mathsf{U}_j \wedge \Delta_i = \Delta_j$.
4. $\mathtt{H}_4$: $\exists^* i, j, k, \ell \in [q]$ such that $\mathsf{X}_i = \mathsf{X}_j \wedge \mathsf{U}_j = \mathsf{U}_k \wedge \mathsf{X}_k = \mathsf{X}_\ell$.
5. $\mathtt{H}_5$: $\exists^* i, j, k, \ell \in [q]$ such that $\mathsf{U}_i = \mathsf{U}_j \wedge \mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell$.
6. $\mathtt{H}_6$: $\exists k \geq 2^n/2q, \exists^* i_1, i_2, \ldots, i_k \in [q]$ such that $\mathsf{X}_{i_1} = \cdots = \mathsf{X}_{i_k}$.
7. $\mathtt{H}_7$: $\exists k \geq 2^n/2q, \exists^* i_1, i_2, \ldots, i_k \in [q]$ such that $\mathsf{U}_{i_1} = \cdots = \mathsf{U}_{i_k}$.

Case 1. $\widehat{\mathsf{H}}$ is bad: If the hash key $\widehat{\mathsf{H}}$ is bad, then Y^q and V^q values are sampled degenerately as $\mathsf{Y}_i = \mathsf{V}_i = 0$ for all $i \in [q]$. It means that we sample without maintaining any specific conditions, which will almost certainly lead to inconsistencies.

Case 2. $\widehat{\mathsf{H}}$ is good: To characterize the transcript corresponding to a good hash key, it will be useful to study a random bipartite edge-labeled graph associated with $(\mathsf{X}^q, \mathsf{U}^q, \Delta^q)$.

Definition G.1 (Transcript Graph). *A transcript graph $\mathcal{G} = (\mathcal{X}, \mathcal{U}, \mathcal{E})$ associated with $(\mathsf{X}^q, \mathsf{U}^q, \Delta^q)$, denoted $\mathcal{G}(\mathsf{X}^q, \mathsf{U}^q, \Delta^q)$, is an undirected bipartite graph, where $\mathcal{X} := \{(\mathsf{X}_i, 0) : i \in [q]\}$ and $\mathcal{U} := \{(\mathsf{U}_i, 1) : i \in [q]\}$ are the two partitions of the vertex-set, and $\mathcal{E} := \{((\mathsf{X}_i, 0), (\mathsf{U}_i, 1)) : i \in [q]\}$ denotes the edge-set. We also associate the label Δ_i with edge $((\mathsf{X}_i, 0), (\mathsf{U}_i, 1)) \in \mathcal{E}$.*

For all practical purposes we may drop the partition markers 0 and 1, for each vertex $(\mathsf{X}_i, 0) \in \mathcal{X}$ and $(\mathsf{U}_i, 1) \in \mathcal{U}$, as they can be easily distinguished from the context and notations. Note that the event $\mathsf{X}_i = \mathsf{X}_j$ and $\mathsf{U}_i = \mathsf{U}_j$, although extremely unlikely, will result in a parallel edge in $\mathcal{G}$. Finally, each edge $(\mathsf{X}_i, \mathsf{U}_i) \in \mathcal{E}$ corresponds to a query index $i \in [q]$, so we can equivalently view (and call) the edge $(\mathsf{X}_i, \mathsf{U}_i)$ as index (or query) i.

Consider the random transcript graph $\mathcal{G}(\mathsf{X}^q, \mathsf{U}^q)$ arising due to $\widehat{\mathsf{H}} \in \widehat{\mathcal{H}}_{\mathsf{good}}$. Lemma G.1 and Figure G.1 characterizes the different types of possible components in $\mathcal{G}(\mathsf{X}^q, \mathsf{U}^q)$.

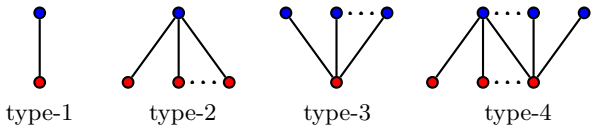

Fig. G.1: Enumerating all possible types of components of a transcript graph corresponding to a good hash key: type-1 is the only possible component of size = 1 edge; type-2 and type-3 are star components with center in $\mathcal{X}$ and $\mathcal{U}$, respectively; type-4 is the only possible component that is not isolated or star (can have degree 2 vertices in both $\mathcal{X}$ and $\mathcal{U}$). Note that the vertex-coloring is only for illustration purposes.

Lemma G.1. *The transcript graph $\mathcal{G}(\mathsf{X}^q, \mathsf{U}^q, \Delta^q)$ generated by a good hash key $\widehat{\mathsf{H}}$ has the following properties:*

1. *$\mathcal{G}$ is simple, acyclic and has no isolated vertices.*
2. *$\mathcal{G}$ has no two adjacent edges i and j such that $\Delta_i \oplus \Delta_j = 0$.*
3. *$\mathcal{G}$ has no component of size $\geq 2^n/2q$ edges.*
4. *$\mathcal{G}$ has no component such that it has 2 distinct degree 2 vertices in $\mathcal{X}$ or $\mathcal{U}$.*

In fact the all possible types of components in $\mathcal{G}$ are enumerated in Figure G.1.

The proof of Lemma G.1 is elementary and left as an exercise for the reader.

It should be noted that in [25], the authors do not explicitly address type-4 graphs. Instead, they focus on two specific subclasses, namely, type-4 and type-5 graphs [25]. Fortunately, this omission does not significantly impact either the security bound or the subsequent analysis.

In what follows, we describe the sampling of Y^q and V^q conditioned on the fact that $\widehat{\mathbf{H}} \in \widehat{\mathcal{H}}_{\mathsf{good}}$. We collect the indices $i \in [q]$ corresponding to the edges in all type-1, type-2, type-3, and type-4 components, in the index sets $\mathcal{I}_1$, $\mathcal{I}_2$, $\mathcal{I}_3$, and $\mathcal{I}_4$, respectively. Clearly, the five sets are disjoint, and $[q] = \mathcal{I}_1 \sqcup \mathcal{I}_2 \sqcup \mathcal{I}_3 \sqcup \mathcal{I}_4$. Let $\mathcal{I} = \mathcal{I}_1 \sqcup \mathcal{I}_2 \sqcup \mathcal{I}_3$. Consider a constrained system of equations

$$\mathcal{L} = \{Y_i \oplus V_i = \Delta_i \ : \ i \in \mathcal{I}\},$$

with the constraint

$$\phi : \ \mathsf{X}^q \leftrightsquigarrow Y^q \wedge \mathsf{U}^q \leftrightsquigarrow V^q.$$

The solution space for $\mathcal{L}$, satisfying the constraint ϕ, is precisely the set

$$\mathcal{S} = \{(y^{\mathcal{I}}, v^{\mathcal{I}}) : y^{\mathcal{I}} \leftrightsquigarrow \mathsf{X}^{\mathcal{I}} \wedge v^{\mathcal{I}} \leftrightsquigarrow \mathsf{U}^{\mathcal{I}} \wedge y^{\mathcal{I}} \oplus v^{\mathcal{I}} = \Delta^{\mathcal{I}}\}.$$

Given these definitions, the ideal oracle $\mathcal{O}_0$ samples $(\mathsf{Y}^q, \mathsf{V}^q)$ as follows:

- $(\mathsf{Y}^{\mathcal{I}}, \mathsf{V}^{\mathcal{I}}) \xleftarrow{\$} \mathcal{S}$, i.e., $\mathcal{O}_0$ uniformly samples one valid assignment from the set of all valid assignments for $\mathsf{Y}^{\mathcal{I}}$ and $\mathsf{V}^{\mathcal{I}}$.
- Let $\mathcal{G} \setminus \mathcal{C}_{\mathcal{I}}$ denote the subgraph of $\mathcal{G}$ after the removal of all type-1, type-2, and type-3 components. For each component $\mathcal{C}$ of $\mathcal{G} \setminus \mathcal{C}_{\mathcal{I}}$:
  - Suppose $(\mathsf{X}_i, \mathsf{U}_i) \in \mathcal{C}$ corresponds to an edge in $\mathcal{C}$, where both X_i and U_i have degree ≥ 2. Then, $\mathsf{Y}_i \xleftarrow{\$} \{0,1\}^n$ and $\mathsf{V}_i = \mathsf{Y}_i \oplus \Delta_i$.
  - For each edge $(\mathsf{X}_{i'}, \mathsf{U}_{i'}) \neq (\mathsf{X}_i, \mathsf{U}_i) \in \mathcal{C}$, either $\mathsf{X}_{i'} = \mathsf{X}_i$ or $\mathsf{U}_{i'} = \mathsf{U}_i$. Suppose, $\mathsf{X}_{i'} = \mathsf{X}_i$. Then, $\mathsf{Y}_{i'} = \mathsf{Y}_i$ and $\mathsf{V}_{i'} = \mathsf{Y}_{i'} \oplus \Delta_{i'}$. Now, suppose $\mathsf{U}_{i'} = \mathsf{U}_i$. Then, $\mathsf{V}_{i'} = \mathsf{V}_i$ and $\mathsf{Y}_{i'} = \mathsf{V}_{i'} \oplus \Delta_{i'}$.

At this point, $\Theta_0 = (\mathsf{T}^q, \mathsf{M}^q, \mathsf{C}^q, \mathsf{X}^q, \mathsf{Y}^q, \mathsf{V}^q, \mathsf{U}^q, \Delta^q, \widetilde{\mathbf{H}}_1, \widetilde{\mathbf{H}}_2, \mathbf{H})$ is completely defined. In this way we maintain both the consistency of equations of the form $\mathsf{Y}_i \oplus \mathsf{V}_i = \Delta_i$ (as in the case of real world), and the permutation consistency within each component, given that $\widehat{\mathbf{H}} \in \widehat{\mathcal{H}}_{\mathsf{good}}$. However, there might be collisions among Y or V values from different components.

G.2 Definition and Analysis of Bad Transcripts

Given the description of the transcript random variable corresponding to the ideal oracle we can define the set of transcripts Ω as the set of all tuples

$\omega = (t^q, m^q, c^q, x^q, y^q, v^q, u^q, \delta^q, \widetilde{h}_1, \widetilde{h}_2, h)$, where $t^q \in (\{0,1\}^\tau)^q$; $m^q, c^q, y^q, v^q \in (\{0,1\}^n)^q$; $\widehat{h} = (\widetilde{h}_1, \widetilde{h}_2, h) \in \widehat{\mathcal{H}}$; $x^q = \widetilde{h}_1(t^q, m^q)$; $u^q = \widetilde{h}_2(t^q, c^q)$; $\delta^q = h(t^q)$; and $(t^q, m^q) \rightsquigarrow (t^q, c^q)$.

Our bad transcript definition is inspired by two requirements:

1. Eliminate all x^q, u^q, and δ^q tuples such that both y^q and v^q are trivially restricted by way of linear dependence. For example, consider the condition $\mathtt{H}_2$. This leads to $y_i = y_j$, which would imply $v_i = y_i \oplus \delta_i = y_j \oplus \delta_j = v_j$. Assuming $i > j$, v_i is trivially restricted ($= v_j$) by way of linear dependence. This may lead to $u^q \not\rightsquigarrow v^q$ as u_i may not be equal to u_j.
2. Eliminate all x^q, u^q, y^q, v^q tuples such that $x^q \not\rightsquigarrow y^q$ or $u^q \not\rightsquigarrow v^q$.

Among the two, requirement 2 is trivial as $x^q \rightsquigarrow y^q$ and $u^q \rightsquigarrow v^q$ is always true for real world transcript. Requirement 1 is more of a technical one that helps in the ideal world sampling of y^q and v^q.

Bad Transcript Definition: Throughout the discussion, we consider the transcript

$$\omega = (t^q, m^q, c^q, x^q, y^q, v^q, u^q, \delta^q, \widehat{h})$$

to characterize the bad transcripts.

We first designate certain transcripts as bad depending upon the characterization of hash keys. Inspired by the ideal world description, we say that a hash key $\widehat{h} \in \widehat{\mathcal{H}}_{\mathsf{bad}}$ (or $\widehat{h}$ is bad) if and only if the following predicate is true:

$$\mathtt{H}_1 \vee \mathtt{H}_2 \vee \mathtt{H}_3 \vee \mathtt{H}_4 \vee \mathtt{H}_5 \vee \mathtt{H}_6 \vee \mathtt{H}_7.$$

We say that ω is *hash-induced* bad transcript, if $\widehat{h} \in \mathcal{H}_{\mathsf{bad}}$. We write this event as $\mathtt{BAD1}$, and by a slight abuse of notations,[5] we have

$$\mathtt{BAD1} = \bigcup_{i=1}^{7} \mathtt{H}_i. \tag{46}$$

This takes care of the first requirement. For the second one we have to enumerate all the conditions which might lead to $x^q \not\rightsquigarrow y^q$ or $u^q \not\rightsquigarrow v^q$. Since we sample degenerately when the hash key is bad, the transcript is *trivially inconsistent* in this case. For good hash keys, if $x_i = x_j$ (or $u_i = u_j$) then we always have $y_i = y_j$ (res. $v_i = v_j$); hence the inconsistency won't arise. So, given that the hash key is good, we say that ω is *sampling-induced* bad transcript, if one of the following conditions is true:
for some $\alpha \in [4]$ and $\beta \in \{\alpha, \ldots, 4\}$, we have

- $\mathtt{Ycoll}_{\alpha\beta}$: $\exists i \in \mathcal{I}_\alpha, j \in \mathcal{I}_\beta$, such that $x_i \neq x_j \wedge y_i = y_j$, and
- $\mathtt{Vcoll}_{\alpha\beta}$: $\exists i \in \mathcal{I}_\alpha, j \in \mathcal{I}_\beta$, such that $u_i \neq u_j \wedge v_i = v_j$,

where $\mathcal{I}_i$ is defined as before in section G.1. By varying α and β over all possible values, we get all 30 conditions which might lead to $x^q \not\rightsquigarrow y^q$ or $u^q \not\rightsquigarrow v^q$. Here we remark that some of these 30 conditions are never satisfied due to the sampling mechanism prescribed in section G.1. These are $\mathtt{Ycoll}_{11}$, $\mathtt{Ycoll}_{12}$, $\mathtt{Ycoll}_{13}$,

[5] We use the notation $\mathtt{H}_i$ to denote the event that the predicate $\mathtt{H}_i$ is true.

$\mathtt{Ycoll}_{22}$, $\mathtt{Ycoll}_{23}$, $\mathtt{Ycoll}_{33}$, $\mathtt{Vcoll}_{11}$, $\mathtt{Vcoll}_{12}$, $\mathtt{Vcoll}_{13}$, $\mathtt{Vcoll}_{22}$, $\mathtt{Vcoll}_{23}$, and $\mathtt{Vcoll}_{33}$. We listed them here only for the sake of completeness. We write the combined event that one of the 30 conditions hold as BAD2. Again by an abuse of notations, we have

$$\mathtt{BAD2} = \bigcup_{\alpha\in[4],\beta\in\{\alpha,\ldots,4\}} (\mathtt{Ycoll}_{\alpha\beta} \cup \mathtt{Vcoll}_{\alpha\beta}) \,. \tag{47}$$

Finally, a transcript ω is called bad, i.e. $\omega \in \Omega_{\mathsf{bad}}$, if it is either a hash-induced or a sampling-induced bad transcript. All other transcripts are called good. It is easy to see that all good transcripts are attainable (as required in the H-coefficient technique or the expectation method).

Bad Transcript Analysis: We analyze the probability of realizing a bad transcript in the ideal world. By definition, this is possible if and only if one of BAD1 or BAD2 occurs. So, we have

$$\begin{aligned}\epsilon_{\mathsf{bad}} = \mathsf{Pr}\left(\Theta_0 \in \Omega_{\mathsf{bad}}\right) &= \mathsf{Pr}_{\Theta_0}\left(\mathtt{BAD1} \cup \mathtt{BAD2}\right) \\ &\le \underbrace{\mathsf{Pr}_{\Theta_0}\left(\mathtt{BAD1}\right)}_{\epsilon_{\mathsf{h}}} + \underbrace{\mathsf{Pr}_{\Theta_0}\left(\mathtt{BAD2}\right)}_{\epsilon_{\mathsf{s}}} . \end{aligned} \tag{48}$$

Lemma G.2 upper bounds ϵ_{h} to $q^2\epsilon_1^2 + q^2\epsilon_1\epsilon_2 + 2q^2\epsilon_1^{1.5} + 16q^4\epsilon_1 2^{-2n}$ and Lemma G.3 upper bounds ϵ_{s} to $4q^4\epsilon_1^2 2^{-n}$. Substituting these values in (48), we get

$$\epsilon_{\mathsf{bad}} \le q^2\epsilon_1^2 + q^2\epsilon_1\epsilon_2 + 2q^2\epsilon_1^{1.5} + \frac{16q^4\epsilon_1}{2^{2n}} + \frac{4q^4\epsilon_1^2}{2^n}. \tag{49}$$

Lemma G.2. $\epsilon_{\mathsf{h}} \le q^2\epsilon_1^2 + q^2\epsilon_1\epsilon_2 + 2q^2\epsilon_1^{1.5} + \dfrac{16q^4\epsilon_1}{2^{2n}}$.

Proof. Using (46) and (48), we have

$$\epsilon_{\mathsf{h}} = \mathsf{Pr}\left(\mathtt{H}_1 \cup \mathtt{H}_2 \cup \mathtt{H}_3 \cup \mathtt{H}_4 \cup \mathtt{H}_5 \cup \mathtt{H}_6 \cup \mathtt{H}_7\right) \le \sum_{i=1}^{7} \mathsf{Pr}\left(\mathtt{H}_i\right).$$

$\mathtt{H}_1$ is true if for some distinct i, j both $\mathsf{X}_i = \mathsf{X}_j$, and $\mathsf{U}_i = \mathsf{U}_j$. Now $\mathsf{T}_i = \mathsf{T}_j \implies \mathsf{M}_i \ne \mathsf{M}_j$. Hence $\mathsf{X}_i \ne \mathsf{X}_j$ (since $\widetilde{\mathbf{H}}_1$ is a tweakable permutation) and $\mathtt{H}_1$ is not true. So suppose $\mathsf{T}_i \ne \mathsf{T}_j$. Then, using the fact that $\widetilde{\mathcal{H}}$ is an ϵ-AUHF and KG is a PISM, for a fixed i, j we get an upper bound of ϵ_1^2. Furthermore, we have at most $\binom{q}{2}$ pairs of (i, j). Thus, $\mathsf{Pr}\left(\mathtt{H}_1\right) \le \binom{q}{2}\epsilon_1^2$.

Following a similar line of argument one can bound $\mathsf{Pr}\left(\mathtt{H}_2\right) \le \binom{q}{2}\epsilon_1\epsilon_2$ and $\mathsf{Pr}\left(\mathtt{H}_3\right) \le \binom{q}{2}\epsilon_1\epsilon_2$.

In the remaining, we bound the probability of $\mathtt{H}_4$ and $\mathtt{H}_6$, while the probability of $\mathtt{H}_5$ and $\mathtt{H}_7$ can be bounded analogously. Now, $\mathtt{H}_4$ is true if for some pairwise distinct i, j, k, ℓ,

$$\widetilde{\mathbf{H}}_1(\mathsf{T}_i, \mathsf{M}_i) = \widetilde{\mathbf{H}}_1(\mathsf{T}_j, \mathsf{M}_j)\widetilde{\mathbf{H}}_2(\mathsf{T}_j, \mathsf{C}_j) \quad = \widetilde{\mathbf{H}}_2(\mathsf{T}_k, \mathsf{C}_k)\widetilde{\mathbf{H}}_1(\mathsf{T}_k, \mathsf{M}_k) = \widetilde{\mathbf{H}}_1(\mathsf{T}_\ell, \mathsf{M}_\ell).$$

Again, using the fact that KG is a PISM, we have that the second equation is independent of the other two equations. Using Lemma D.1, we have

$$\mathsf{Pr}\left(\mathtt{H}_4\right) \leq q^2\epsilon_1^{1.5}.$$

For $\mathtt{H}_6$, for some $i_1, \ldots, i_k$, we have

$$\mathsf{X}_{i_1} = \mathsf{X}_{i_2} = \cdots = \mathsf{X}_{i_k},$$

where $k \geq 2^n/2q$. Since, $(t_{i_j}, m_{i_j}) \neq (t_{i_l}, m_{i_l})$ for all $j \neq l$, we can apply Corollary D.1 to get

$$\mathsf{Pr}\left(\mathtt{H}_6\right) \leq \frac{8q^4\epsilon_1}{2^{2n}}.$$

Lemma G.3. $\epsilon_{\mathsf{s}} \leq \dfrac{4q^4\epsilon_1^2}{2^n}$.

Proof. Using (47) and (48), we have

$$\begin{aligned}\epsilon_{\mathsf{s}} &= \mathsf{Pr}\left(\bigcup_{\alpha\in[4],\beta\in\{\alpha,\ldots,4\}} \left(\mathtt{Ycoll}_{\alpha\beta} \cup \mathtt{Vcoll}_{\alpha\beta}\right)\right) \\ &\leq \sum_{\alpha\in[4]} \sum_{\beta\in\{\alpha,\ldots,4\}} \left(\mathsf{Pr}\left(\mathtt{Ycoll}_{\alpha\beta}\right) + \mathsf{Pr}\left(\mathtt{Vcoll}_{\alpha\beta}\right)\right).\end{aligned}$$

We bound the probabilities of the events on the right hand side in groups as given below:

1. Bounding $\sum_{\alpha\in[3],\beta\in\{\alpha,\ldots,3\}} \mathsf{Pr}\left(\mathtt{Ycoll}_{\alpha\beta}\right) + \mathsf{Pr}\left(\mathtt{Vcoll}_{\alpha\beta}\right)$: Recall that the sampling of Y and V values is always done consistently for indices belonging to $\mathcal{I} = \mathcal{I}_1 \sqcup \mathcal{I}_2 \sqcup \mathcal{I}_3$. Hence,

$$\sum_{\alpha\in[3],\beta\in\{\alpha,\ldots,3\}} \mathsf{Pr}\left(\mathtt{Ycoll}_{\alpha\beta}\right) + \mathsf{Pr}\left(\mathtt{Vcoll}_{\alpha\beta}\right) = 0, \tag{50}$$

2. Bounding $\sum_{\alpha\in[3]} \mathsf{Pr}\left(\mathtt{Ycoll}_{\alpha 4}\right) + \mathsf{Pr}\left(\mathtt{Vcoll}_{\alpha 4}\right)$: Let's consider the event $\mathtt{Ycoll}_{14}$, which translates to there exist indices $i \in \mathcal{I}_1$ and $j \in \mathcal{I}_4$ such that $\mathsf{X}_i \neq \mathsf{X}_j \wedge \mathsf{Y}_i = \mathsf{Y}_j$. Since $j \in \mathcal{I}_4$, there must exist $k, \ell \in \mathcal{I}_4 \setminus \{j\}$, such that one of the following happens

$$\begin{aligned}\mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell \\ \mathsf{U}_j = \mathsf{U}_k \wedge \mathsf{X}_k = \mathsf{X}_\ell \\ \mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_j = \mathsf{U}_\ell.\end{aligned}$$

We analyze the first case, while the other two cases can be similarly bounded. To bound the probability of $\mathtt{Ycoll}_{14}$, we can look at the joint event

$$\mathtt{E}: \ \exists i \in \mathcal{I}_1, \exists^* j, k, \ell \in \mathcal{I}_4, \text{ such that } \mathsf{Y}_i = \mathsf{Y}_j \wedge \mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell.$$

Note that the event $\mathsf{Y}_i = \mathsf{Y}_j$ occurs with exactly 2^{-n} probability conditioned on the event $\mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell$. Thus, we get

$$\begin{aligned}\mathsf{Pr}\left(\mathtt{E}\right) &= \mathsf{Pr}\left(\exists i \in \mathcal{I}_1, \exists^* j,k,\ell \in \mathcal{I}_4, \text{ such that } \mathsf{Y}_i = \mathsf{Y}_j \wedge \mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell\right) \\ &\leq \sum_{i \in \mathcal{I}_1} \sum_{j<k<\ell \in \mathcal{I}_4} \mathsf{Pr}\left(\mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell\right) \times \mathsf{Pr}\left(\mathsf{Y}_i = \mathsf{Y}_j \mid \mathsf{X}_j = \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell\right) \\ &\leq q\binom{q}{3}\frac{\epsilon_1^2}{2^n},\end{aligned}$$

where the last inequality follows from the AUHF property of $\widetilde{\mathcal{H}}$, the PISM property of KG, and the uniform randomness of Y_j. The probability of the other two cases are identically bounded, whence we get

$$\mathsf{Pr}\left(\mathtt{Ycoll}_{14}\right) \leq 3q\binom{q}{3}\frac{\epsilon_1^2}{2^n}.$$

We can bound the probabilities of $\mathtt{Ycoll}_{24}$, $\mathtt{Ycoll}_{34}$, and $\mathtt{Vcoll}_{\alpha 4}$ for all $\alpha \in [3]$ in a similar manner. So, we skip the argumentation for these cases, and summarize the probability for this group as

$$\sum_{\alpha \in [3]} \mathsf{Pr}\left(\mathtt{Ycoll}_{\alpha 4}\right) + \mathsf{Pr}\left(\mathtt{Vcoll}_{\alpha 4}\right) \leq \frac{3q^4\epsilon_1^2}{2^n}. \tag{51}$$

3. Bounding $\mathsf{Pr}\left(\mathtt{Ycoll}_{44}\right) + \mathsf{Pr}\left(\mathtt{Vcoll}_{44}\right)$: Consider the event $\mathtt{Ycoll}_{44}$, which translates to there exists distinct indices $i, j \in \mathcal{I}_4$ such that $\mathsf{X}_i \neq \mathsf{X}_j \wedge \mathsf{Y}_i = \mathsf{Y}_j$. Here as $i, j \in \mathcal{I}_4$, there must exist $k, \ell \in \mathcal{I}_4 \setminus \{j\}$ such that one of the following happens

$$\begin{aligned}\mathsf{X}_j &= \mathsf{X}_k \wedge \mathsf{U}_k = \mathsf{U}_\ell \\ \mathsf{U}_j &= \mathsf{U}_k \wedge \mathsf{X}_k = \mathsf{X}_\ell \\ \mathsf{X}_j &= \mathsf{X}_k \wedge \mathsf{U}_j = \mathsf{U}_\ell.\end{aligned}$$

The analysis of these cases is similar to 2 above. So, we skip it and provide the final bound

$$\mathsf{Pr}\left(\mathtt{Ycoll}_{44}\right) \leq 3q\binom{q}{3}\frac{\epsilon_1^2}{2^n}.$$

The probability of $\mathtt{Vcoll}_{44}$ can be bounded in a similar fashion.

$$\mathsf{Pr}\left(\mathtt{Ycoll}_{44}\right) + \mathsf{Pr}\left(\mathtt{Vcoll}_{44}\right) \leq \frac{q^4\epsilon_1^2}{2^n}. \tag{52}$$

The result follows by combining (50)-(52), followed by some simplifications. □

G.3 Good Transcript Analysis

From section G.1, we know the types of components present in the transcript graph corresponding to a good transcript ω are exactly as in Figure G.1. Let $\omega = (t^q, m^q, c^q, x^q, y^q, v^q, u^q, \delta^q, \widetilde{h}_1, \widetilde{h}_2, h)$ be the good transcript at hand. From the bad transcript description of section G.2, we know that for a good transcript $(t^q, m^q) \rightsquigarrow (t^q, c^q)$, $x^q \rightsquigarrow y^q$, $v^q \rightsquigarrow u^q$, and $y^q \oplus v^q = \delta^q$.

First, we add some new parameters with respect to ω to aid the remaining analysis.

For $i \in [4]$, let $c_i(\omega)$ and $q_i(\omega)$ respectively denote the number of components and number of indices (corresponding to the edges) of type-i in ω. Further, let $z_i^1(\omega)$ and $z_i^2(\omega)$ respectively denote the number of vertices from $\mathcal{X}$ and $\mathcal{U}$ in type-i components. Note that

- $z_1^1(\omega) = z_1^2(\omega) = q_1(\omega) = c_1(\omega)$;
- $z_2^1(\omega) = c_2(\omega)$, and $z_2^2(\omega) = q_2(\omega) \geq 2c_2(\omega)$;
- $z_3^2(\omega) = c_3(\omega)$, and $z_3^1(\omega) = q_3(\omega) \geq 2c_3(\omega)$; and
- $z_4^b(\omega) \geq 2c_4(\omega)$, for $b \in \{1, 2\}$, and $z_4^1 + z_4^2 = q_4 - c_4$.

In addition, for a good transcript $q = \sum_{i=1}^{5} q_i(\omega)$. For notational convenience, let $p_1 := z_1^1 + z_2^1 + z_3^1 = q_1 + c_2 + q_3$ and $p_2 := z_1^2 + z_2^2 + z_3^2 = q_1 + q_2 + c_3$.

Let $(t'_1, t'_2, \cdots, t'_r)$ be an arbitrary ordering of $\mathtt{S}(t^q)$, and for all $i \in [r]$, let μ_i denote the multiplicity of t'_i in the multiset $\mathtt{M}(t^q)$, i.e., $r \leq q$ and $\sum_{i=1}^{r} \mu_i = q$. In addition, let μ'_i denote the multiplicity of t'_i in the multiset $\mathtt{M}(t^{\mathcal{I}})$, i.e., $\sum_{i=1}^{r} \mu'_i = |\mathcal{I}|$.

Let $(\delta'_1, \delta'_2, \cdots, \delta'_s)$ be an arbitrary ordering of $\mathtt{S}(\delta^{\mathcal{I}})$, and for all $i \in [s]$, let ν_i denote the multiplicity of δ'_i in the multiset $\mathtt{M}(\delta^{\mathcal{I}})$, i.e., $s \leq |\mathcal{I}|$ and $\sum_{i=1}^{s} \nu_i = |\mathcal{I}|$.

For all these parameters, we will drop the ω qualification whenever it is understood from the context.

Interpolation probability for the real oracle: In the real oracle, $\widehat{\mathsf{H}} \leftarrow \mathsf{KG}\left(\widehat{\mathcal{H}}\right)$, $\boldsymbol{\pi}_1$ is called exactly $p_1 + z_4^1$ times and $\boldsymbol{\pi}_2$ is called exactly $p_2 + z_4^2$ times. Thus, we have

$$\mathsf{Pr}\left(\Theta_1 = \omega\right) = \mathsf{Pr}_{\mathsf{KG}}\left(\widehat{\mathsf{H}} = \widehat{h}\right) \times \frac{1}{(2^n)_{p_1 + z_4^1}} \times \frac{1}{(2^n)_{p_2 + z_4^2}}. \tag{53}$$

Interpolation probability for the ideal oracle: In the ideal oracle, the sampling is done in parts:

I. $\widetilde{\boldsymbol{\pi}}$ *sampling*: We have

$$\mathsf{Pr}\left(\widetilde{\boldsymbol{\pi}}(t^q, m^q) = c^q\right) \leq \frac{1}{\prod_{i=1}^{r} (2^n)_{\mu_i}}.$$

II. *Hash key sampling*: This is identical to the real world, and simply given by $\mathsf{Pr}_{\mathsf{KG}}\left(\widehat{\mathsf{H}} = \widehat{h}\right)$.

III. *Internal variables sampling*: The internal variables Y^q and V^q are sampled in two stages.

(A). *type*-1, *type*-2 *and type*-3 *sampling*: Recall the sets $\mathcal{I}_1$, $\mathcal{I}_2$, and $\mathcal{I}_3$, from section G.2. Consider the system of equation

$$\mathcal{L} = \{\mathsf{Y}_i \oplus \mathsf{V}_i = \delta_i \ : \ i \in \mathcal{I}\}.$$

From Figure G.1 we know that $\mathcal{L}$ is cycle-free and non-degenerate. Further, $\xi_{\max}(\mathcal{L}) \leq 2^n/2q$, since the transcript is good. So, we can apply Theorem E.1 to get a lower bound on the the number of valid solutions, $|\mathcal{S}(\mathcal{L})|$ for $\mathcal{L}$. Using the fact that $(\mathsf{Y}^{\mathcal{I}}, \mathsf{V}^{\mathcal{I}}) \leftarrow\$\ \mathcal{S}(\mathcal{L})$, and Theorem E.1, we have

$$\mathsf{Pr}\left((\mathsf{Y}^{\mathcal{I}}, \mathsf{V}^{\mathcal{I}}) = (y^{\mathcal{I}}, v^{\mathcal{I}})\right) \leq \frac{\prod_{i=1}^{s}(2^n)_{\nu_i}}{\zeta(\omega)(2^n)_{p_1}(2^n)_{p_2}},$$

where

$$\zeta(\omega) \geq \left(1 - \frac{13q^4}{2^{3n}} - \frac{2q^2}{2^{2n}} - \left(\sum_{i=1}^{c_2+c_3} \eta_{c_1+i}^2\right)\frac{4q^2}{2^{2n}}\right),$$

and η_i denotes the number of edges in the i-th component for all $i \in [c_1 + c_2 + c_3]$.

(B). *type*-4 *sampling*: For the remaining indices, one value is sampled uniformly for each of the components, i.e., we have

$$\mathsf{Pr}\left(\left(\mathsf{Y}^{[q]\setminus\mathcal{I}}, \mathsf{V}^{[q]\setminus\mathcal{I}}\right) = \left(y^{[q]\setminus\mathcal{I}}, v^{[q]\setminus\mathcal{I}}\right)\right) = \frac{1}{2^{nc_4}}.$$

By combining I, II, III, and rearranging the terms, we have

$$\mathsf{Pr}\left(\Theta_0 = \omega\right) \leq \mathsf{Pr}_{\mathsf{KG}}\left(\widehat{\mathbf{H}} = \widehat{h}\right) \times \frac{1}{\zeta(\omega)} \times \frac{\prod_{i=1}^{s}(2^n)_{\nu_i}}{\prod_{i=1}^{r}(2^n)_{\mu_i}(2^n)_{p_1}(2^n)_{p_2}2^{nc_4}}. \quad (54)$$

G.4 Ratio of Interpolation Probabilities

On dividing (53) by (54), and simplifying the expression, we get

$$\begin{aligned}
\frac{\mathsf{Pr}\left(\Theta_1 = \omega\right)}{\mathsf{Pr}\left(\Theta_0 = \omega\right)} &\geq \zeta(\omega) \cdot \frac{\prod_{i=1}^{r}(2^n)_{\mu_i}}{\prod_{i=1}^{s}(2^n)_{\nu_i}(2^n - p_1 - c_4)_{z_4^1 - c_4}(2^n - p_2)_{z_4^2}} \\
&\overset{1}{\geq} \zeta(\omega) \cdot \frac{\prod_{i=1}^{r}(2^n)_{\mu_i'}\prod_{i=1}^{r}(2^n - \mu_i')_{\mu_i - \mu_i'}}{\prod_{i=1}^{s}(2^n)_{\nu_i}(2^n - p_1 - c_4)_{z_4^1 - c_4}(2^n - p_2)_{z_4^2}} \\
&\overset{2}{\geq} \zeta(\omega) \cdot \frac{\prod_{i=1}^{r}(2^n - \mu_i')_{\mu_i - \mu_i'}}{(2^n - p_1 - c_4)_{z_4^1 - c_4}(2^n - p_2)_{z_4^2}} \Bigg\} A \\
&\overset{3}{\geq} \zeta(\omega). \qquad (55)
\end{aligned}$$

At inequality 1, we simply rewrite the numerator. Further, $r \geq s$, as number of distinct internal masking values is at most the number of distinct tweaks, and $\mathsf{S}(t^{\mathcal{I}})$ compresses to $\mathsf{S}(\delta^{\mathcal{I}})$. So, using Proposition B.1, we can justify inequality 2. At inequality 2, for $i \in \{2, 3, 4\}$, $c_i(\omega) > 0$ if and only if $r \geq 2$. Also, $\mu_i' \leq$

$c_1 + c_2 + c_3 \leq p_1 \leq p_1 + c_4$ and similarly $\mu'_i \leq p_2$ for all $i \in [r]$. Furthermore, $\mu_i \leq c_1 + c_2 + c_3 + 2c_4 \leq p_1 + z_4^1$, and similarly $\mu_i \leq p_2 + z_4^2$. Also, $\sum_{i=1}^{r} \mu_i - \mu'_i = q_4 = z_4^1 + z_4^2 - c_4$. Thus, A satisfies the conditions laid out in Proposition B.2, and hence $A \geq 1$. This justifies inequality 3.

We define $\epsilon_{\mathsf{ratio}} : \Omega \to [0, \infty)$ by the mapping

$$\epsilon_{\mathsf{ratio}}(\omega) = 1 - \zeta(\omega).$$

Clearly $\epsilon_{\mathsf{ratio}}$ is non-negative and the ratio of real to ideal interpolation probabilities is at least $1 - \epsilon_{\mathsf{ratio}}(\omega)$ (using (55)). Thus, we can use the expectation method to get

$$\mathbf{Adv}^{\mathsf{ind\text{-}cca}}_{\mathsf{LRW+}}(q) \leq \frac{2q^2}{2^{2n}} + \frac{13q^4}{2^{3n}} + \frac{4q^2}{2^{2n}} \mathsf{Ex}\left(\sum_{i=1}^{c_2+c_3} \eta^2_{c_1+i}\right) + \epsilon_{\mathsf{bad}}. \tag{56}$$

Let $\sim_1$ (res. $\sim_2$) be an equivalence relation over $[q]$, such that $\alpha \sim_1 \beta$ (res. $\alpha \sim_2 \beta$) if and only if $\mathsf{X}_\alpha = \mathsf{X}_\beta$ (res. $\mathsf{U}_\alpha = \mathsf{U}_\beta$). Now, each η_i random variable denotes the cardinality of some non-singleton equivalence class of $[q]$ with respect to either $\sim_1$ or $\sim_2$. Let $\mathcal{P}^1_1, \ldots, \mathcal{P}^1_k$ and $\mathcal{P}^2_1, \ldots, \mathcal{P}^2_{k'}$ denote the non-singleton equivalence classes of $[q]$ with respect to $\sim_1$ and $\sim_2$, respectively. Further, for $j \in [k]$ and $j' \in [k']$, let $\mathsf{n}_j = |\mathcal{P}^1_j|$ and $\mathsf{m}_{j'} = |\mathcal{P}^2_{j'}|$. Then, we have

$$\begin{aligned}\mathsf{Ex}\left(\sum_{i=1}^{c_2+c_3} \eta^2_{c_1+i}\right) &\leq \mathsf{Ex}\left(\sum_{j=1}^{k} \mathsf{n}_j^2\right) + \mathsf{Ex}\left(\sum_{j'=1}^{k'} \mathsf{m}_{k'}^2\right) \\ &\leq 4q^2\epsilon_1. \end{aligned} \tag{57}$$

where the first inequality follows from linearity, and the second inequality follows from Lemma D.3. Theorem 6.1 then follows from (49), (56), and (57). □

H Proofs Related to LRW+ Instances

H.1 Proof of Corollary 6.1

It is obvious that KG is a PISM and $\mathbf{H}$ is 0-AUHF. The result then follows from Theorem 6.1, once we establish that $\widetilde{\mathbf{H}}_1$ and $\widetilde{\mathbf{H}}_2$ are $(2^n - 1)^{-1}$-AUTPFs. We derive the AUTPF bound for $\widetilde{\mathbf{H}}_1$, and the corresponding bound for $\widetilde{\mathbf{H}}_2$ can be derived analogously.

Fix two arbitrary distinct pairs (t, m) and $(t', m') \in \{0,1\}^n \times \{0,1\}^n$, and let $\mathsf{X}_1 = m$, $\mathsf{X}'_1 = m'$, $\mathsf{Y}_{i-1} = \boldsymbol{\pi}'_{i-1}(\mathsf{X}_{i-1})$, $\mathsf{X}_i = t \oplus \mathsf{Y}_{i-1}$, $\mathsf{Y}'_{i-1} = \boldsymbol{\pi}'_{i-1}(\mathsf{X}'_{i-1})$, and $\mathsf{X}'_i = t' \oplus \mathsf{Y}'_{i-1}$, for $2 \leq i \leq r' - 1$.

Observe that, $(t, m) \neq (t', m')$ implies that $(\mathsf{X}_i, \mathsf{X}_{i+1}) \neq (\mathsf{X}'_i, \mathsf{X}'_{i+1})$, and $t = t'$ implies $\mathsf{X}_i \neq \mathsf{X}'_i$, as i-LRW1 is a tweakable permutation, for all $i \leq [r' - 1]$. Let $\delta = t \oplus t'$. We have

$$\mathsf{Pr}\left(\widetilde{\mathbf{H}}_1(t, m) = \widetilde{\mathbf{H}}_2(t', m')\right) = \mathsf{Pr}\left(\mathsf{X}_{r'} = \mathsf{X}'_{r'}\right)$$

$$\begin{aligned}
&\overset{1}{\leq} \mathsf{Pr}\left(\mathsf{X}_{r'} = \mathsf{X}'_{r'} \mid \mathsf{X}_{r'-1} \neq \mathsf{X}'_{r'-1}\right) \\
&= \mathsf{Pr}\left(\mathsf{Y}_{r'-1} \oplus \mathsf{Y}'_{r'-1} = \delta \mid \mathsf{X}_{r'-1} \neq \mathsf{X}'_{r'-1}\right) \\
&\overset{2}{\leq} \frac{1}{2^n - 1},
\end{aligned}$$

where inequality 1 follows from the aforementioned observation, and inequality 2 follows from the fact that $\mathsf{Y}_{r'-1}$ is chosen uniform at random from a set of size $(2^n - 1)$ conditioned on all other variables. □

H.2 Proof of Corollary 6.2

As in the case of r-LRW1, it is sufficient to show that $\widetilde{\mathbf{H}}_1$ is an ϵ-AUHF. The result then follows from Theorem 6.1.

Fix two arbitrary distinct pairs (t, m) and $(t', m') \in \{0,1\}^\tau \times \{0,1\}^n$, and let $\mathsf{X}_1 = m \oplus \mathbf{H}'_1(t)$, $\mathsf{X}'_1 = m' \oplus \mathbf{H}'_1(t')$, $\mathsf{X}_i = \mathbf{H}'_{i-1}(t) \oplus \mathbf{H}'_i(t) \oplus \boldsymbol{\pi}'_{i-1}(\mathsf{X}_{i-1})$, and $\mathsf{X}'_i = \mathbf{H}'_{i-1}(t') \oplus \mathbf{H}'_i(t') \oplus \boldsymbol{\pi}'_{i-1}(\mathsf{X}'_{i-1})$, for $2 \leq i \leq r' - 1$. Define

$$\mathsf{Z} := \mathbf{H}'_{r'-1}(t) \oplus \mathbf{H}'_{r'-1}(t') \oplus \boldsymbol{\pi}'_{r'-1}(\mathsf{X}_{r'-1}) \oplus \boldsymbol{\pi}'_{r'-1}(\mathsf{X}'_{r'-1})$$

Note that the two outputs collide only if $t \neq t'$. Then, we have

$$\begin{aligned}
\mathsf{Pr}\left(\widetilde{\mathbf{H}}_1(t, m) = \widetilde{\mathbf{H}}_2(t', m')\right) &= \mathsf{Pr}\left(\mathsf{X}_{r'} = \mathsf{X}'_{r'}\right) \\
&= \mathsf{Pr}\left(\mathbf{H}'_{r'}(t) \oplus \mathbf{H}'_{r'}(t') = \mathsf{Z}\right) \leq \epsilon,
\end{aligned}$$

where the inequality follows from the ϵ-AXUHF of $\mathbf{H}'_{r'}$ once we condition on Z. This completes the proof. □