A note on "Tweakable HCTR: A BBB Secure Tweakable Enciphering Scheme"

Mustafa Khairallah

Department of Electrical and Information Technology, Lund University, Lund, Sweden
khairallah@ieee.org

Abstract. Tweakable HCTR is an tweakable enciphering proposed by Dutta and Nandi in Indocrypt 2018. It provides beyond birthday bound security when each tweak value is not used too frequently. More importantly for this note, its security bound degrades linearly with the maximum input length. We show in this note that this is not true by showing a single query distinguisher with advantage $O(l^2/2^n)$ where l is the length of that query. The distinguisher does not break the beyond-birthday-bound claim but gives higher advantage than the claimed bound. After disclosing this flaw publicly, the authors of [ABPV21] have pointed out that they also discovered this flaw in their paper earlier, yet their attack seems to have been unnoticed by the designers. Thus, this note should serve now as a confirmation of their analysis rather than a new observation.

Keywords: Tweakable HCTR · BBB · Birthday Bound · Enciphering

1 Introduction

Tweakable HCTR [DN18, DN19] is an enciphering scheme that can be used as a variable-length Tweakable Wide Block Cipher (TWBC). It targets being Beyond Birthday Bound (BBB) secure when each tweak value is not used too frequently. The scheme is depicted in Figure 1.

The scheme has a security bound as follows (approximating the hash advantages and considering a random TBC):

$$2(\mu-1)\frac{q+\sigma}{2^n}+\frac{2q\sigma+q^2}{2^{2n}}+2\max\{\frac{(\mu-1)ql}{2^n}+\frac{q\sigma}{2^{2n}},\frac{\sigma}{2^n}\},$$

where μ is the maximum tweak multiplicity, l is the maximum input length in blocks, q is the number of queries, σ is the total number of blocks queried and n is the block size of the underlying TBC. The bound above is assuming $|H_2| \geq n$.

2 Bound for a single query

Consider an adversary that performs a single query of length l. In this case, $q = \mu = 1$ and $\sigma = l$. Then, the bound becomes

$$\frac{2l}{2^{2n}}+\frac{1}{2^{2n}}+2\max\{\frac{l}{2^{2n}},\frac{l}{2^n}\}\leq\frac{3l}{2^n}.$$

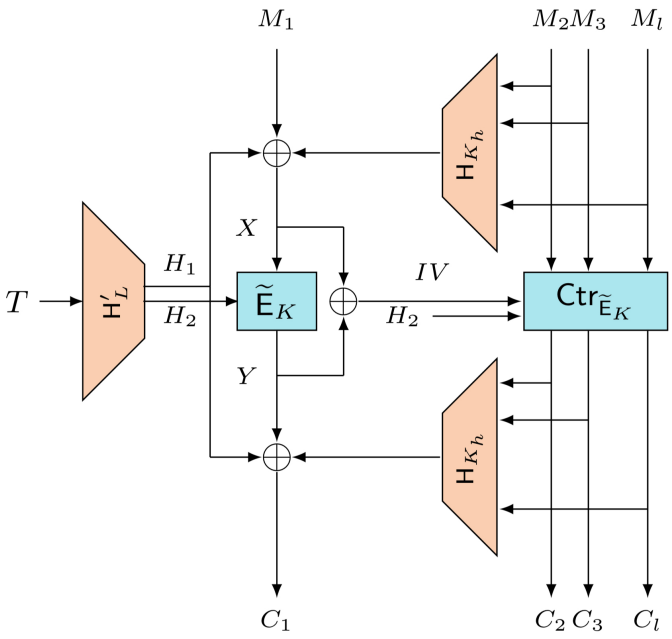

Figure 1: The Tweakable HCTR Scheme [DN19].

3 Contradicting Distinguisher

Consider an adversary that chooses a tweak T and the plaintext 0^{ln} for some $l > 24$. The adversary makes only one encryption query and gets the output C such that

$$C = C_1 \| C_2 \| \cdots C_l.$$

The adversary outputs 0 if the elements of the set $\{C_2, C_3, \cdots C_l\}$ are distinct.

In the real world, we can see how these elements are selected in line 6 of the LHS algorithm of [DN19, Figure 3.2]:

$$C_j = 0^n \oplus \tilde{E}_K(H_2, IV \oplus \bar{j}),$$

where $\bar{j}$ is a binary representation of the integer j. Thus, $IV \oplus \bar{j} \neq IV \oplus \bar{i}$ for all $i \neq j$. Since all the TBC calls use the same key and tweak, they all use the same permutation and $\{C_2, C_3, \cdots C_l\}$ are always distinct.

In the ideal world, each n bit block of the ciphertext can be seen as an n-bit truncation of an ln-bit random permutation. Thus, with a small negligible error, the probability that the adversary outputs 0 is the probability that there is no collision between these random values. We know that

$$\Pr[\texttt{coll}] \geq \frac{(l-1)(l-2)}{2^{n+2}}.$$

which gives us an adversarial advantage that is

$$\geq \frac{l^2}{2^{n+3}}.$$

With $l > 24$, the bound contradicts the claimed security.

4 On the Proof

The proof of tweakable HCTR uses the H coefficient technique. The bad transcript analysis is given in [DN19, Section 4.1]. The authors state:

> The underlying principle for identifying the bad events is that
> *if hash of two tweak value happens to collide in two different invocations of the cipher, then the block cipher input and output must not collide.*

Indeed, we see that the authors give five bad events all of which depend on events that require at least two queries to occur. It seems that events that require one query have been missed. The designers explained in direct communications that by extending the bad event $\mathsf{B3}$ to cover block cipher calls in the same query, it will cover this distinguisher, albeit with updated probability calculation.

5 Implication

The distinguisher does not nullify the BBB claim when the l is reasonably bounded, albeit with a worse bound. It drops to birthday bound when l is unbounded, a property that is shared with several BBB-secure enciphering schemes. However, the proof would require reworking to include the missing bad events.

A solution to remove this quadratic drop is to include the counter in the tweak, instead of xoring it to the IV/plaintext of the TBC, but that would increase the tweak size of the TBC and incur a significant overhead if the TBC is constructed from block ciphers. In the original construction, the tweak is updated once per query, while in this case, it needs to be updated for every call.

Acknowledgement

I would like to thank Avijit Dutta and Mridul Nandi for confirming the observation, the missing term in the bound and the missing bad event, and Bart Mennink for the discussions. This was discovered while on a research visit to Radboud University. The author is supported by the Wallenberg-NTU Presidential Postdoctoral Fellowship. I would like to also thank Amit Singh Bhati for pointing to [ABPV21].

References

[ABPV21] Elena Andreeva, Amit Singh Bhati, Bart Preneel, and Damian Vizár. 1, 2, 3, fork: Counter mode variants based on a generalized forkcipher. *IACR Transactions on Symmetric Cryptology*, 2021(3):1–35, Sep. 2021.

[DN18] Avijit Dutta and Mridul Nandi. Tweakable hctr: A bbb secure tweakable enciphering scheme. In *International Conference on Cryptology in India*, pages 47–69. Springer, 2018.

[DN19] Avijit Dutta and Mridul Nandi. Tweakable hctr: A bbb secure tweakable enciphering scheme. Cryptology ePrint Archive, Paper 2019/1324, 2019. https://eprint.iacr.org/2019/1324.