Breaking the Twinkle Authenticated Encryption Scheme and Analyzing Its Underlying Permutation

Group work at Asian Symmetric Key (ASK) workshop 2024

Debasmita Chakraborty	TU Graz (Austria)
Hosein Hadipour	RUB (Germany)
Anup Kumar Kundu	ISI (India)
Mostafizar Rahman	Kyoto University (Japan)
Prathamesh Ram	IIT Bhilai (India)
Yu Sasaki	NTT (Japan) and NIST Associate (US)
Dilip Sau	IIT Kharagpur (India)
Aman Sinha	NTU (Singapore)

August 13, 2025, SAC2025@Toronto

Summary

- This talk presents cryptanalysis for Twinkle-family, a low-latency AE or MAC proposed at IACR CiC 2024.

- Mode-analysis:
  - broke claimed security when confidentiality is higher than integrity
  - nonce-respect, recovering c-bit authentication key, $c \in \{512,1024\}$, with $O(2^t)$ queries, where $t \in \{64,128\}$ is a tag size.

- Primitive-analysis:
  - analyzed an internal permutation (1280-bit state, SPN structure)
  - developed automated tool for several approaches
  - improved the attacks by designers by using differential-linear distinguisher

Motivation of Twinkle

NTT

- The goal is to provide memory protection, by designing domain-specific AE and MACs tailored for system-level security.

- **Twinkle-AE**: Cache-line encryption (target of this paper)
  - The cache line (plaintext) size, c, is either 512 bits or 1024 bits.
  - The tag size, t, is either 64 or 128 bits.
- **Twinkle-PA**: Pointer authentication
  - Input size is 128 bits, a 64-bit pointer address and a 64-bit context.
  - The tag size is at most 128 bits, and can be truncated: $1 \leq t \leq 128$.

- How to optimize designs for low latency in those use cases?

Design Approach of Twinkle-AE

- use of large keys
- single call of a large PRF
- PRF takes nonce, so every invocation derives a new random string.
  - Use a part of PRF output for encryption (one-time pad)
  - Use a part of PRF output for authentication: (Wegman-Carter MAC)

parallel

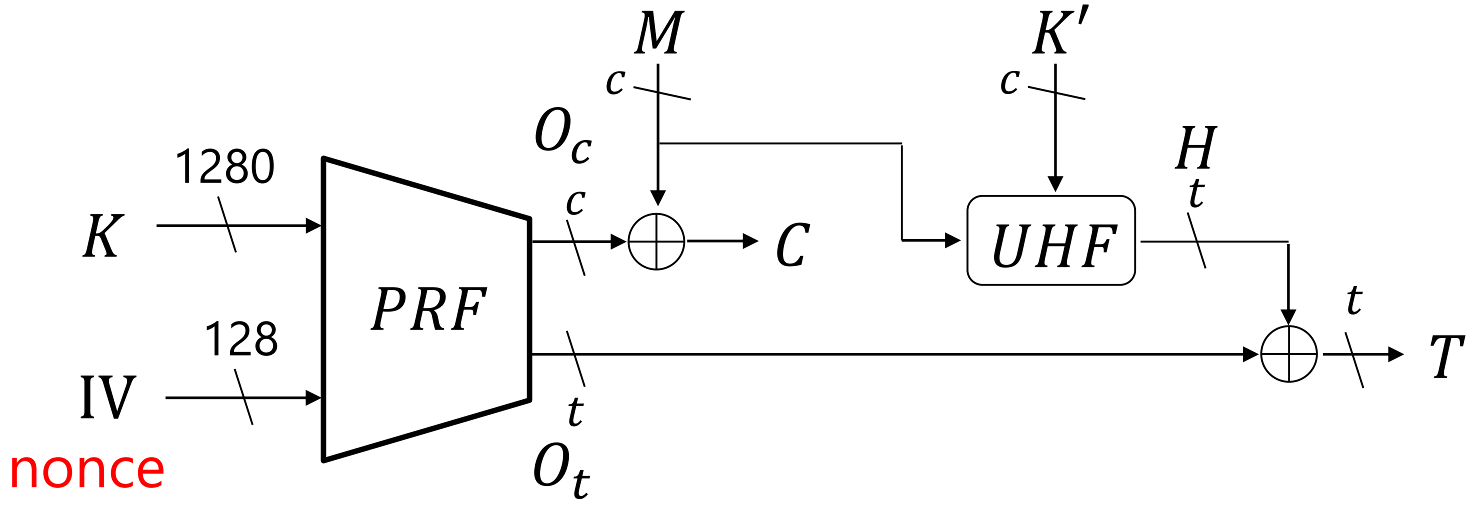

$$UHF: \sum_{i=0}^{c/t-1} M_i \otimes K_i'$$

parallel

Design of Twinkle-PRF

- Even-Mansour construction thanks to a large key.
  - $\mathcal{F}_I$ is almost 10 copies of IV: fast and parallel.
  - $\mathcal{F}_o$ is almost truncation: fast.
  - R is a parameter depending on the confidentiality level.

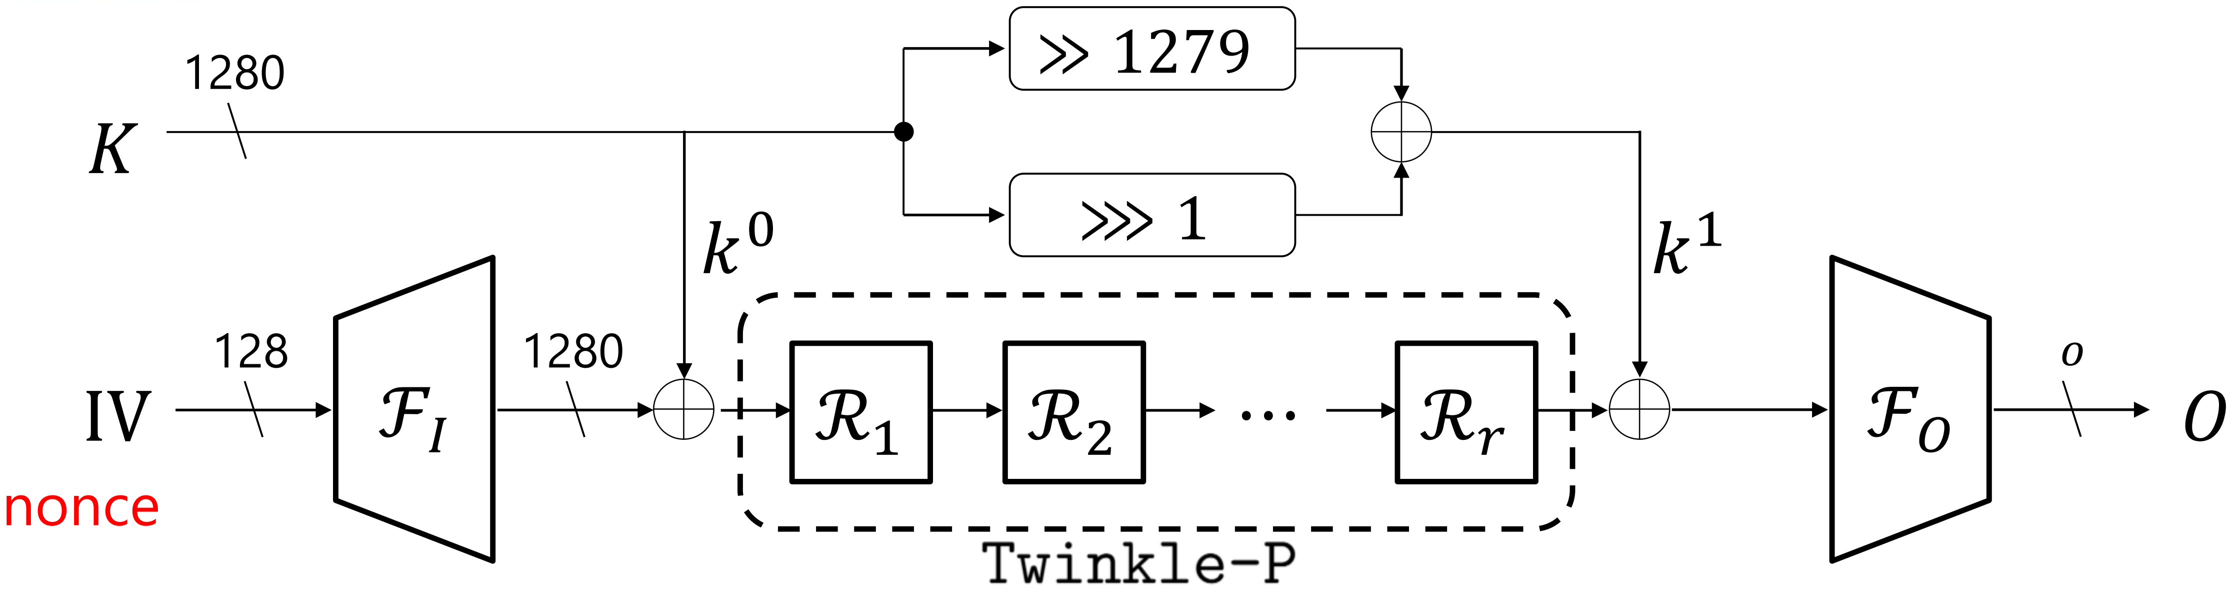

Claimed Security of Twinkle-AE

NTT

Cache line size

Integrity = tag size

Versions	Confidentiality	Integrity (t)
Twinkle-AE-512a	128	64
Twinkle-AE-512b	128	128
Twinkle-AE-512c	256	128
Twinkle-AE-1024a	128	64
Twinkle-AE-1024b	128	128
Twinkle-AE-1024c	256	128

Table 3: Twinkle-AE Versions and Security in Bits

higher confidentiality

- 4 schemes claim higher confidentiality than integrity.
- Confidentiality / Integrity is not defined in the original paper.

Generic Attacks on Twinkle-AE: Authentication Key Recovery

High-level Ideas

NTT

- Twinkle-AE is secure as long as the assumption is held.

  In every invocation, PRF output is random.

- However, in some parameters,
  - claimed confidentiality level is higher than that of integrity:
    allowing something more than exhaustive guesses on the tag
  - claimed confidentiality level is bigger than the nonce size:
    nonce-repeat is inevitable

- For such parameters, authentication key is recovered with $O(2^t)$ queries, and then universal forgery is possible.

Nonce-Respect Authentication Key-Recovery

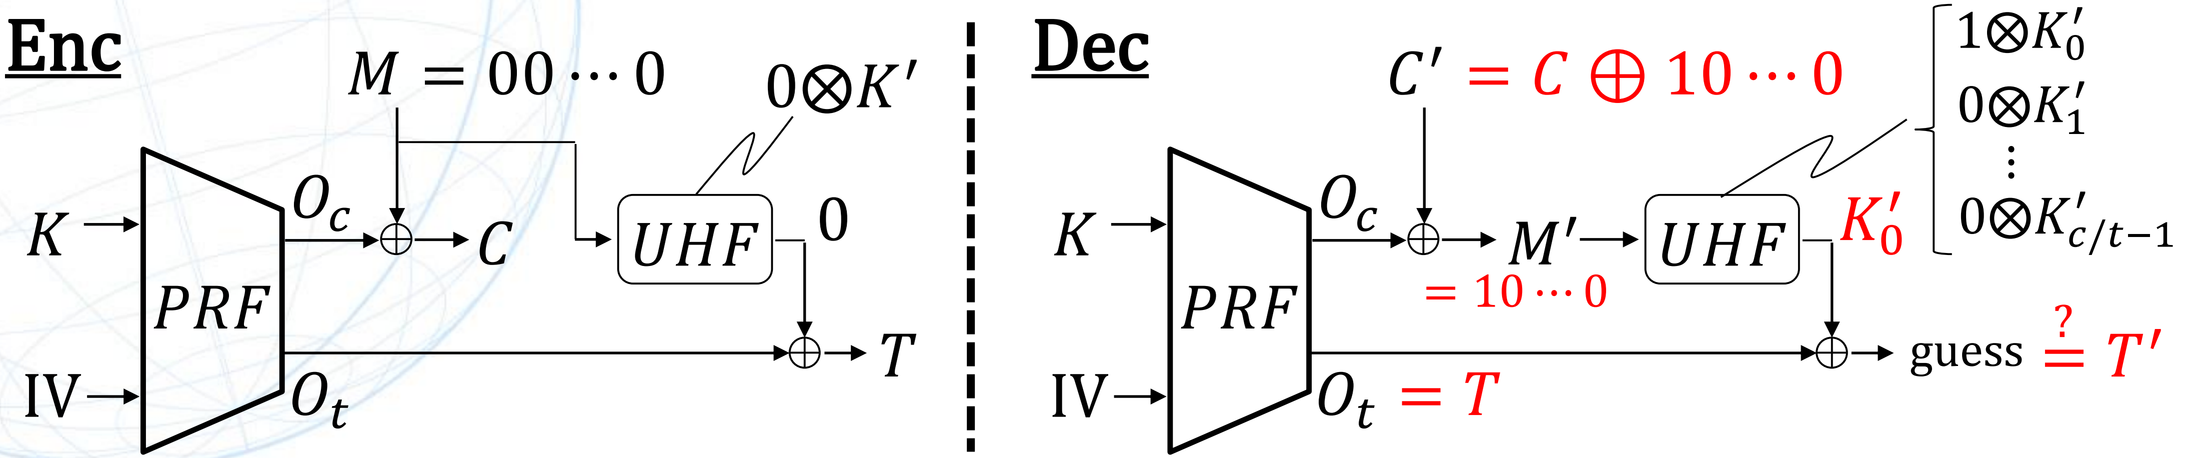

1. Set $M \leftarrow 0$. For some N, make Enc query of (N, M) to obtain (C, T).
   Since $0\otimes K' = 0$ for any K', (C, T) reveals PRF's output: $O_c = C, O_t = T$.
2. Set $C' \leftarrow C \oplus 10\cdots0$ to ensure Dec results for $N, C' = 10\cdots0$.
   Since $1\otimes K'_0 = K'_0$ for any K'_0, it ensures $UHF(M'\otimes K') = K'_0$.
3. Guess T' for all $O(2^t)$ choices, and make Dec query of (N, C', T').
   If verification succeeds, $K'_0 = T' \oplus O_t$.

More Impact

Universal forgery after recovering K'

- For any M, with $O(1)$ cost, the adversary can find a (N, C, T) such that the decryption result is M.
- breaks confidentiality w.r.t. IND-CCA2

Nonce-misuse K' recovery with $O(1)$ cost

- Query $(C, T) \leftarrow Enc(N, M = 00000000)$.
- Query $(C', T') \leftarrow Enc(N, M' = 10000000)$.

Then, $K'_0 = T' \oplus O_t$

Attacking key-commitment security with $O(1)$ cost

- Straightforward. Note that the designers didn't claim this security.

Cryptanalysis on the Underlying Permutation

We analyze the security of Twinkle-P as a standalone permutation, regardless of how it is used in the mode of operation.

Description of Twinkle-P

SPN-Style Round Function (total: 18.5, 9.5, or 5 rounds)

- SubBytes
- LaneRotation0
- MixSlice
- LaneRotation1
- AddConstant

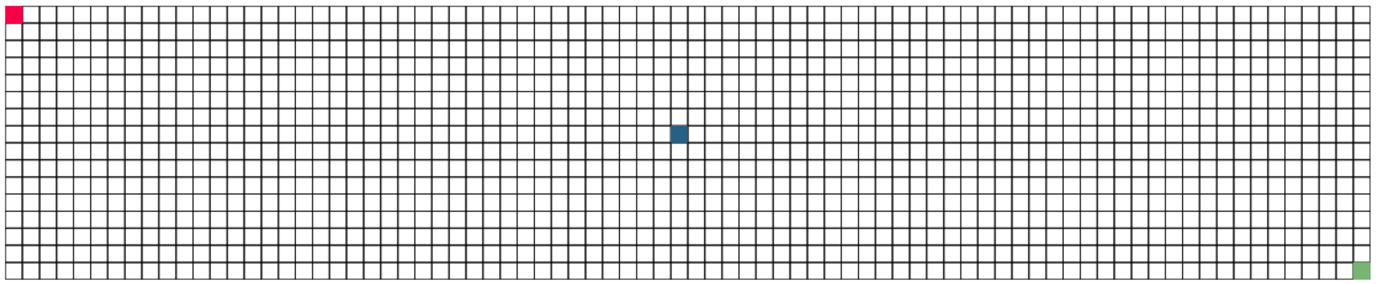

(b) 2D representation of the Twinkle internal state

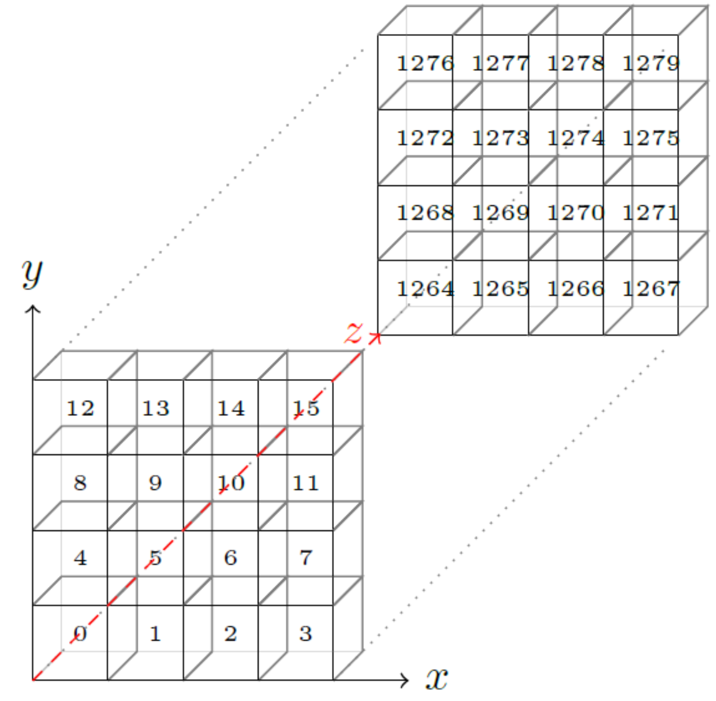

(a) 3D representation of a the Twinkle internal state

Summary of Attack Results

- Main challenge: the large size (1280-bit state)
- investigated attacks by developing automated tools

Distinguisher	#Rounds	#Distinguishers	Attack complexity	Ref.
Differential	4	–	$> 2^{58}$	[32]
Linear	4	1	2^{60}	[32]
Truncated Differential	3.5	1	$2^{7.4}$	[32]
Differential-Linear	**4**	80	2	subsection 5.6
	5	80	$2 \cdot \mathbf{2^{5.70}}$	subsection 5.6
	6	80	$2 \cdot \mathbf{2^{73.32}}$	subsection 5.6
Impossible Differential	4	$80 \cdot 2^{1820}$	–	subsection 5.3
	5	$80 \cdot 2^{1148}$	–	subsection 5.3
	6	$80 \cdot 2^{356}$	–	subsection 5.3
Zero-Correlation Linear	4	$80 \cdot 2^{1278}$	–	subsection 5.4
	5	$80 \cdot 2^{1140}$	–	subsection 5.4
	6	$80 \cdot 2^{16}$	–	subsection 5.4
Integral	3	80	2	subsection 5.4
	4	80	2^4	subsection 5.4
	5	80	2^{12}	subsection 5.5

Table 1: Summary of distinguishers for Twinkle-P

practical, implemented

best attack

Tools Used to Evaluate Each Attack

NTT

- Modeling S-box: **S-box Analyzer** from [ToSC22,ToSC24]
- Imp Diff / Zero-correlation: We implemented two approaches.
  - **negative CP model** [Eurocrypt17]
  - **positive CP model** [Eurocrypt23,ToSC24]
- Integral (division property)
  - **MILP-based model** [FSE16,Asiacrypt16]
- Differential Linear
  - **Technique by Hadipour et al.** [CRYPTO24]
  - 5-round attack with a complexity of $2 \cdot 2^{5.70}$ was experimentally verified.
  - 6-round attack with a complexity of $2 \cdot 2^{73.32}$ is the current best attack.

Conclusion

- This talk: Cryptanalysis for Twinkle from mode and primitive.

  Mode

  - recover c-bit auth key with $O(2^t)$ queries, where $c \in \{512,1024\}, t \in \{64,128\}$
  - $O(1)$ in nonce misuse, inevitable when confidentiality is larger than IV size.

  Primitive

  - analyzed the internal permutation by developing automated tools
  - 5-round practical attack and 6-round theoretical attack by differential-linear

- Our attacks do not work for 2 schemes with balanced conf-int.
- The attacked 4 schemes can also be secure if the claimed confidentiality is compromised to be equal to the *integrity*.

Thank you for your attention!!