Extended Analysis of Key Committing Security of HCTR2

Donghoon Chang[1,2] Yukihito Hiraga[3] Kazuhiko Minematsu[4,5] Nicky Mouha[6]
Yusuke Naito[7] Yu Sasaki[1,8] **Rentaro Shiba**[7,9] Takeshi Sugawara[3]

[1] Associate of NIST
[2] FWI, Fairfax, USA
[3] The University of Electro-Communications, Tokyo, Japan
[4] NEC Corporation, Kawasaki, Japan
[5] The University of Osaka, Suita, Japan
[6] KeyCryptic, Bethesda, USA
[7] Mitsubishi Electric Corporation, Kanagawa, Japan
[8] NTT Social Informatics Laboratories, Tokyo, Japan
[9] Nagoya University, Nagoya, Japan

Selected Areas in Cryptography (SAC 2026)
University of Ottawa, 26 August 2026

Cryptographic Accordions

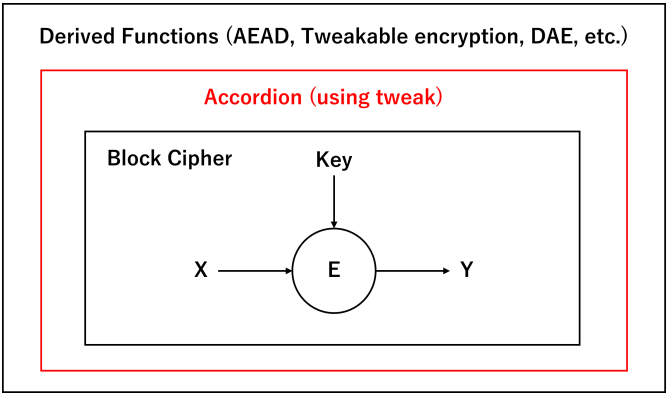

What the middle layer adds

An accordion is a variable-input-length, length-preserving tweakable permutation (VIL-SPRP) constructed from a block cipher with a security proof.

Robust AE by Accordion + Encode-then-Encipher

Encode-then-Encipher principle

message encoding + length-preserving enciphering = robust authenticated encryption

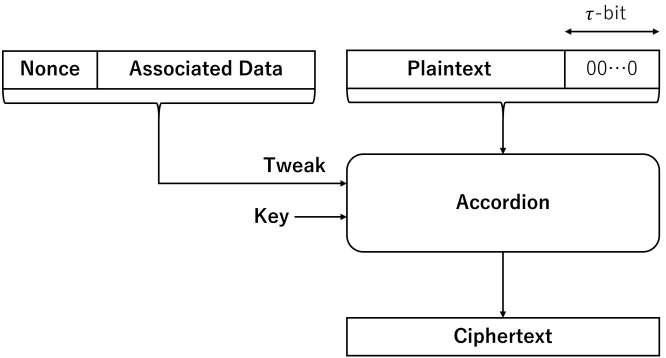

P. Rogaway and T. Shrimpton, EUROCRYPT 2006; Y. L. Chen et al., NIST IR 8552, 2025.

HCTR2: Secure, Efficient, and Widely Deployed

Deployed

Developed by Google; integrated into the Linux kernel and Android file-based encryption.

Efficient

Hash–encrypt–hash using AES, POLYVAL, and XCTR, with hardware acceleration.

NIST relevance

NIST proposes HCTR2-based accordions and asks for additional properties such as key commitment.

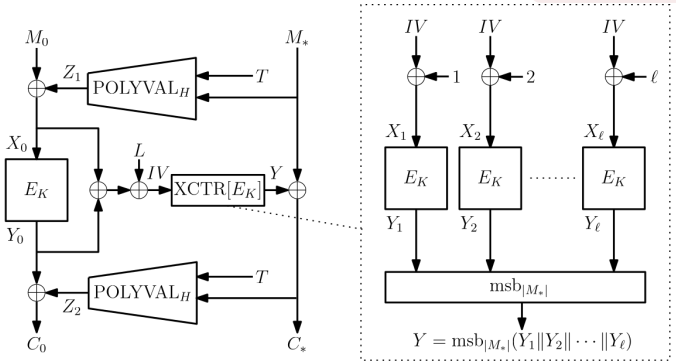

EtE-HCTR2: three plaintext formats

Encode-then-Encipher adds redundancy to the plaintext, then feeds the **entire formatted string** to HCTR2.

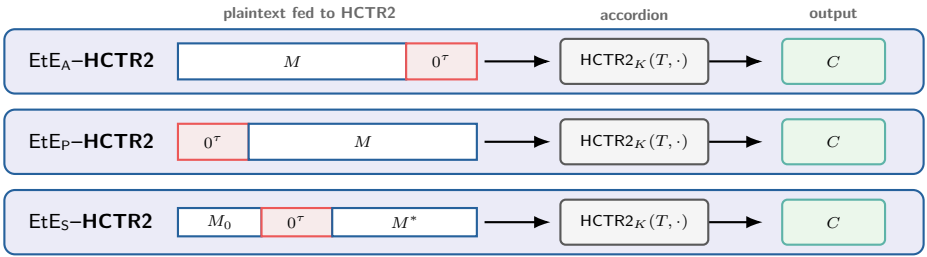

Authentication mechanism

On decryption, the encoding is accepted only if the recovered plaintext contains the required 0^τ redundancy.

Key-committing security of EtE-HCTR2

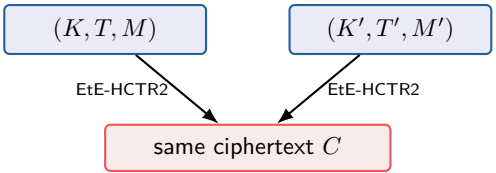

$$F := \text{EtE-HCTR2},$$
$$F_K(T, M) = F_{K'}(T', M') = C, \qquad K \neq K'.$$
$$(T, T' : \text{ tweaks})$$

- CMT_k: the tweaks may differ.
- CMT_k^*: additionally require $T = T'$.

Why quantum attacks fit the Q1 setting

The adversary chooses K and K', and all attack processes can be performed offline. Thus, no quantum access to a keyed construction is required.

What we analyze

Attack upper bounds, the role of the IV mask L, and quantum speedups with and without qRAM.

P. Farshim et al., ToSC 2017(1); S. Menda et al., EUROCRYPT 2023; NIST IR 8552.

Previous work: Chang et al. (2026)

At CRYPTO 2026, Chang et al. established the main classical baseline for the three EtE-HCTR2 formats.

Best classical attacks before this work

target	time
EtE_A	$\mathcal{O}\left(2^{\max\{\tau/3,\, \tau-n\}}\right)$
EtE_P	$\mathcal{O}\left(2^{\tau/2}\right)$
EtE_S	$\mathcal{O}\left(2^{\tau/2}\right)$

The $2^{\tau/3}$ attack on EtE_A uses a three-phase divide-and-conquer strategy.

What remained unresolved

1. Does the IV mask L matter for key commitment?
2. Can the information-theoretic attacks be tightened toward the known bounds?
3. How much quantum speedup is possible, with or without qRAM?

This paper extends that analysis in all three directions

- $\mathcal{O}(1)$ attack without L
- tighter information-theoretic attacks
- quantum attacks obtained by classical attacks of Chang et al.

Contributions

Focus 1: IV mask L

$$L = E_k(\mathsf{bin}_n(1))$$
$$IV = X_0 \oplus Y_0 \oplus L$$

Removing L exposes a key-controlled correlation between the left block-cipher call and XCTR.

without L $\implies$ $\mathcal{O}(1)$ attack.

Paper Result: Information-theoretic attacks

One XCTR key-stream block can be reused across many message lengths, creating more collision sources per BC call.

$2^{\tau/2+2} \longrightarrow 2^{(\tau-\log_2 n)/2+5/2}$.

Focus 2: Quantum attacks

Quantumly accelerate Chang et al.'s divide-and-conquer attack and re-optimize the split across adjacent XCTR blocks.

with qRAM : $2^{\tau/4}$,

without qRAM : $2^{2\tau/7}$.

Outline

1. **The Role of the IV Mask L**
2. **Quantum Attacks**

The mask L is redundant for AE security—but not for commitment

Conventional AE analysis

The key is secret. The adversary cannot freely control the polynomial-hash output. With suitable padding, prior work proves the same AE security even without L.

Key commitment

The adversary chooses the keys and therefore knows:

- $H = E_K(0^n)$ (Hash key)
- $L = E_K(1^n)$ (XCTR input mask)

It can invert POLYVAL and force correlations between the leftmost block-cipher call and an XCTR call.

Without L, the first key-stream block is a two-call function

Let $c := \mathrm{bin}_n(1)$, the counter value used in the first XCTR block.

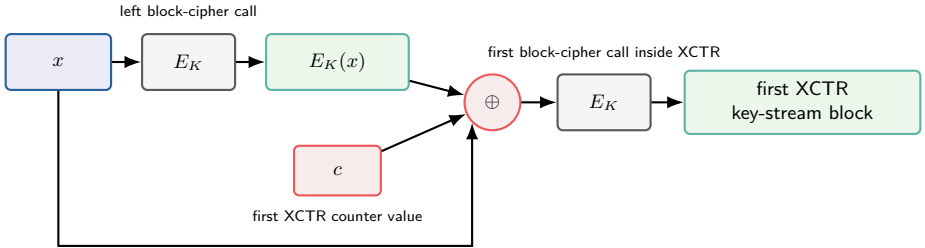

The only function needed for the attack

$$g_K(x) := E_K(x \oplus E_K(x) \oplus c)\,.$$

This is the first XCTR key-stream block when the IV mask L is removed.

What the mask changes

With L, the corresponding input is

$$x \oplus E_K(x) \oplus L \oplus c.$$

The extra L prevents the cancellation used next.

For any chosen key, one decryption forces the output c

Choose the preimage of the counter value

For an attacker-chosen key K, set

$$x_K := E_K^{-1}(c).$$

Then

$$\begin{aligned} g_K(x_K) &= E_K(x_K \oplus E_K(x_K) \oplus c) \\ &= E_K(x_K \oplus c \oplus c) \\ &= E_K(x_K) = c. \end{aligned}$$

No collision search

Every chosen key has an efficiently computable input whose first XCTR key-stream block is the same constant c.

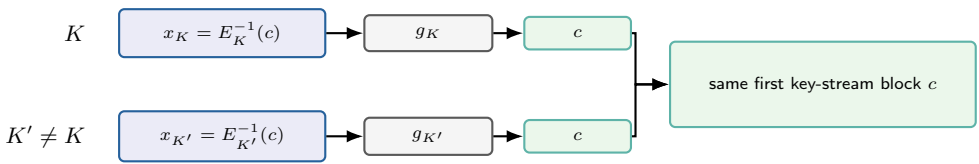

Complete the constant-time committing collision

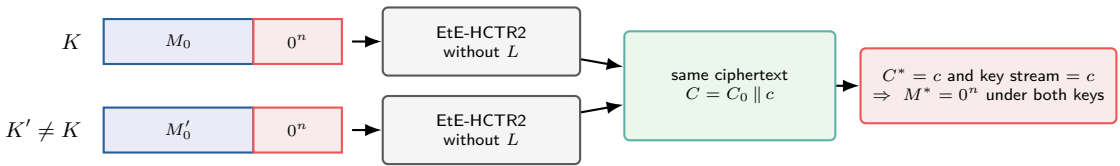

1. Choose two keys	2. Satisfy the EtE format	3. Align the left block
Choose any $K \neq K'$ and compute the two preimages x_K and $x_{K'}$ from the previous slide.	Set $C^* = c$. Both right plaintext blocks decrypt to 0^n, which contains the required τ zero bits for $\tau \leq n$.	Fix C_0, invert the second POLYVAL layer to choose the tweak(s), and recover M_0, M_0' from the first hash layer.

Result: $\mathcal{O}(1)$ BC calls for EtE_A-HCTR2 and EtE_S-HCTR2 without L ($\tau \leq n$).

Outline

1. The Role of the IV Mask L
2. **Quantum Attacks**

Quantum attacks: roadmap and memory assumptions

Generic Grover search is not enough

Searching for a second key whose decryption has valid τ-bit redundancy costs

$$\mathcal{O}(2^{\tau/2}).$$

This is

- slower than the EtE_A classical attack $\mathcal{O}(2^{\tau/3})$,
- only equal to the EtE_P classical attack $\mathcal{O}(2^{\tau/2})$.

We therefore accelerate the *existing structured attacks*, rather than the trivial key search.

Two settings used in the main talk

With qRAM: a precomputed classical list can be read coherently in superposition.

Without qRAM: the list is stored in ordinary classical RAM; the quantum computation uses only polynomially many working qubits.

A stronger writable-quantum-memory model appears only in the backup and is not needed for the main explanation.

What happens for the two formats?

EtE_P: directly replace its collision search by a quantum collision algorithm.

EtE_A: re-optimize the split across two adjacent XCTR blocks.

G. Brassard et al., LATIN 1998; A. Chailloux et al., ASIACRYPT 2017.

Classical starting point: the padding spans two XCTR blocks

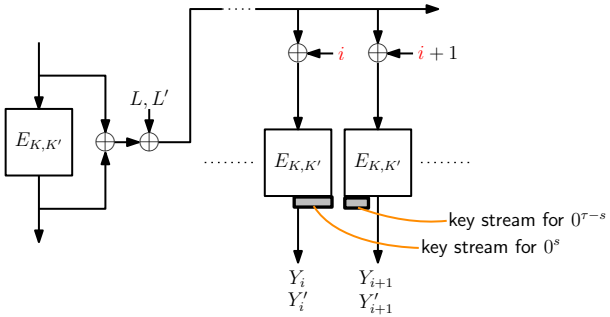

Required collision

The final 0^τ is placed across adjacent key-stream blocks:

$$\mathsf{lsb}_s(Y_i) = \mathsf{lsb}_s(Y'_i)$$

and

$$\mathsf{msb}_{\tau-s}(Y_{i+1}) = \mathsf{msb}_{\tau-s}(Y'_{i+1}).$$

Chang et al.'s classical split

They choose

$$s = \frac{2\tau}{3}, \qquad \tau - s = \frac{\tau}{3}.$$

Phase 1 handles block i; Phase 2 searches the counter i so that the adjacent block also collides.

Phase 1: collide s bits before choosing the counter

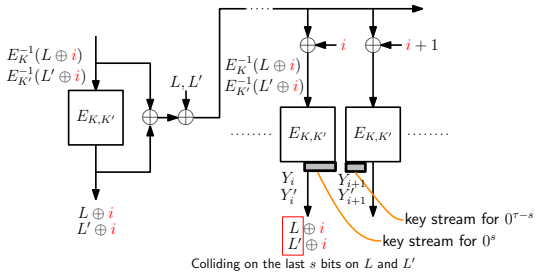

Key-pair condition

Find distinct K, K' such that

$$\mathsf{lsb}_s(L) = \mathsf{lsb}_s(L'), \qquad L = E_K(\mathsf{bin}_n(1)).$$

The core observation gives

$$\mathsf{lsb}_s(Y_i) = \mathsf{lsb}_s(Y'_i) \quad \text{for every counter } i.$$

Cost

This is an s-bit collision problem:

$$T_1 = \begin{cases} 2^{s/2} & \text{classical,} \\ 2^{s/3} & \text{BHT with qRAM.} \end{cases}$$

Phase 2 searches the adjacent block; Phase 3 reaches it

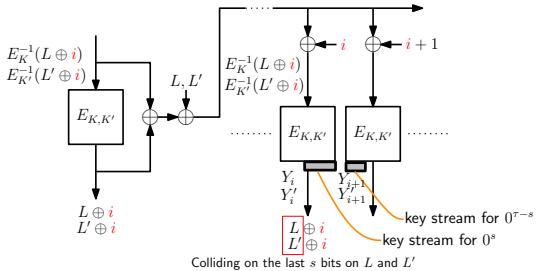

Phase 2: search the counter i

Search for

$$\mathsf{msb}_{\tau-s}(Y_{i+1}) = \mathsf{msb}_{\tau-s}(Y'_{i+1}).$$

Thus

$T_2 = 2^{\tau-s}$ classically, $T_2 = 2^{(\tau-s)/2}$ with Grover.

Phase 3: explicit output

A successful counter is typically as large as

$$i = \mathcal{O}(2^{\tau-s}).$$

The attack must materialize the plaintext through block $i+1$, so

$$T_3 = \mathcal{O}(2^{\tau-s}),$$

with no generic quantum speedup.

A direct quantum speedup is not optimal

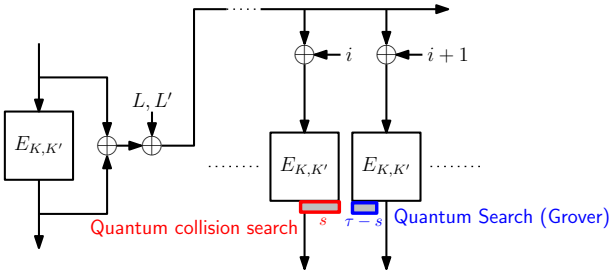

Quantum costs for a general split

$T_1 = 2^{s/3}$ (BHT),

$T_2 = 2^{(\tau-s)/2}$ (Grover),

$T_3 = 2^{\tau-s}$ (explicit output).

Keep the classical split $s = 2\tau/3$?

$$T_1 = 2^{2\tau/9}, \qquad T_2 = 2^{\tau/6}, \qquad T_3 = 2^{\tau/3}.$$

Phase 3 still dominates, so merely replacing the two searches by quantum algorithms is not optimal.

Our idea: re-optimize the split s

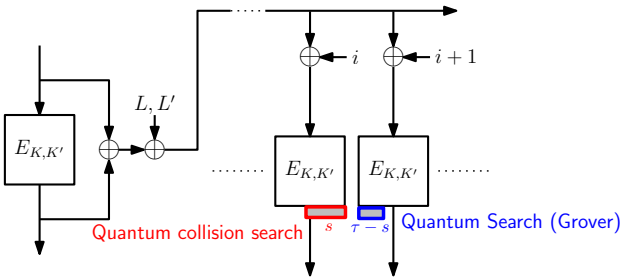

Move work from the adjacent block into Phase 1

Increasing s

- makes quantum collision search more expensive,
- leaves fewer bits for Grover search,
- reduces the counter range and explicit output.

Optimization target

$$T = T_{\text{coll}}(s) + 2^{(\tau-s)/2} + 2^{\tau-s}.$$

The optimal s depends on whether Phase 1 uses qRAM-assisted or qRAM-free collision search.

With qRAM: $T = M = \mathcal{O}(2^{\tau/4})$

Use the Brassard–Høyer–Tapp (BHT) collision algorithm; its precomputed classical list is readable in superposition.

For a BHT precomputed list of size 2^t,

$$M = 2^t,$$

$$T = 2^t + 2^{(s-t)/2} + 2^{(\tau-s)/2} + 2^{\tau-s}.$$

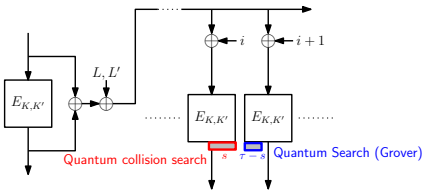

Joint optimum

$$t = \frac{\tau}{4}, \qquad s = \frac{3\tau}{4},$$

$$T = M = \mathcal{O}(2^{\tau/4}).$$

Interpretation

Phase 1 now covers $3\tau/4$ bits; only $\tau/4$ bits remain in the adjacent block and in the explicit-output exponent.

Range

$$s \le n \implies \tau \le \frac{4n}{3}.$$

G. Brassard, P. Høyer, and A. Tapp, LATIN 1998.

Without qRAM: $T = \mathcal{O}(2^{2\tau/7})$

Use the Chailloux–Naya-Plasencia–Schrottenloher (CNS) collision algorithm; the list is stored in ordinary classical RAM.

Use CNS collision search in Phase 1:

$$T_{\text{coll}} = 2^{t-r/2} + 2^{(s-t-1)/2}\left(2^{r/2} + 2^{t-r}\right),$$

$$M_{\text{class}} = 2^{t-r}.$$

The joint optimum is

$$s = \frac{5\tau}{7}, \qquad t = \frac{3\tau}{7}, \qquad r = \frac{2\tau}{7}.$$

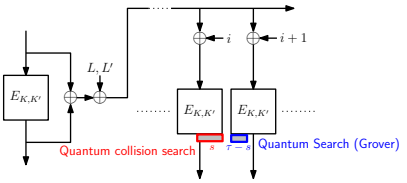

Same attack location, different Phase-1 tool

The blue s-bit collision is found without exponential qRAM; the red adjacent-block condition is still handled by Grover search.

Result

$$T = \mathcal{O}(2^{2\tau/7}), \qquad M_{\text{class}} = \mathcal{O}(2^{\tau/7}), \qquad \tau \leq \frac{7n}{5}.$$

Large τ: use many colliding key pairs

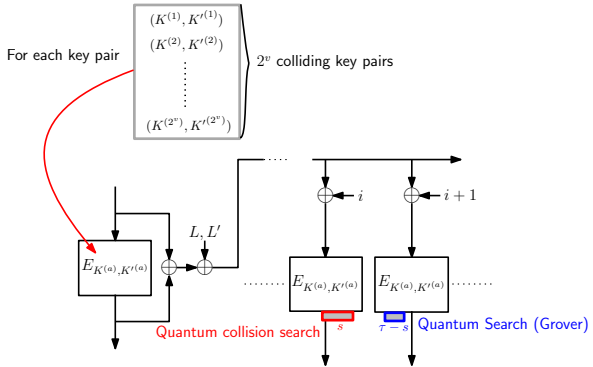

Phase 1 output

Collect 2^v colliding key pairs on the full n-bit value L and store them.

Phase 2 improvement

For each counter i, test all pairs. The hit probability becomes

$$2^{v-(\tau-n)},$$

so the counter range and Phase-3 cost shrink to

$$2^{\tau-n-v}.$$

Setting	Time	Memory
with qRAM	$2^{(2\tau-n)/5}$	$2^{(2\tau-n)/5}$
without qRAM	$2^{2(2\tau-n)/9}$	$2^{(2\tau-n)/9}$

Quantum Attacks Summary

Target	Model	Time	Memory	Range
EtE_A-HCTR2	classical	$2^{\max\{\tau/3,\tau-n\}}$	–	$\tau < 2n$
	with qRAM	$2^{\tau/4}$	$2^{\tau/4}$	$\tau \leq 4n/3$
	without qRAM	$2^{2\tau/7}$	$2^{\tau/7}$	$\tau \leq 7n/5$
	with qRAM	$2^{(2\tau-n)/5}$	$2^{(2\tau-n)/5}$	$4n/3 \leq \tau < 2n$
	without qRAM	$2^{2(2\tau-n)/9}$	$2^{(2\tau-n)/9}$	$7n/5 \leq \tau < 2n$
EtE_P-HCTR2	classical	$2^{\tau/2}$	–	$\tau \leq n$
	with qRAM	$2^{\tau/3}$	$2^{\tau/3}$	$\tau \leq n$
	without qRAM	$2^{2\tau/5}$	$2^{\tau/5}$	$\tau \leq n$

- EtE_P-HCTR2 follows by direct quantum collision search.
- "Without qRAM" means ordinary classical RAM plus polynomially many working qubits.

Takeaways

1. **The IV mask L.** It may be unnecessary for conventional AE security, yet it prevents a constant-time key-commitment attack by breaking a key-controlled block-cipher correlation.
2. **Information-theoretic attacks.** Exploiting every usable message length yields many collision sources per BC call and narrows the gap to the known bounds.
3. **Quantum attacks.** The best attack is not obtained by replacing classical searches one by one; the phase split and memory parameters must be re-optimized, and useful speedups remain without qRAM.

Scope and Limitations

These results are attack upper bounds. This work does not establish matching quantum lower bounds or a full quantum security characterization.

Questions?

References I

[1] Y. L. Chen, M. Davidson, M. Dworkin, J. Kelsey, Y. Sasaki, M. S. Turan, D. Chang, N. Mouha, and A. Thompson.
Requirements for Cryptographic Accordions.
NIST IR 8552, 2025.

[2] P. Rogaway and T. Shrimpton.
A Provable-Security Treatment of the Key-Wrap Problem.
EUROCRYPT 2006.

[3] P. Crowley, N. Huckleberry, and E. Biggers.
Length-Preserving Encryption with HCTR2.
IACR ePrint 2021/1441.

[4] P. Farshim, C. Orlandi, and R. Roșie.
Security of Symmetric Primitives under Incorrect Usage of Keys.
ToSC 2017(1).

[5] S. Menda, J. Len, P. Grubbs, and T. Ristenpart.
Context Discovery and Commitment Attacks.
EUROCRYPT 2023.

[6] Y. L. Chen, Y. Hiraga, N. Mouha, Y. Naito, Y. Sasaki, and T. Sugawara.
Beyond-Birthday-Bound Security with HCTR2.
ASIACRYPT 2025.

[7] G. Brassard, P. Høyer, and A. Tapp.
Quantum Cryptanalysis of Hash and Claw-Free Functions.
LATIN 1998.

References II

[8] A. Chailloux, M. Naya-Plasencia, and A. Schrottenloher.
An Efficient Quantum Collision Search Algorithm and Implications on Symmetric Cryptography.
ASIACRYPT 2017.

[9] X. Bonnetain, J. Loyer, A. Schrottenloher, and Y. Shen.
A Tight Quantum Algorithm for Multiple Collision Search.
IACR ePrint 2025/1690.

Appendix: memory-model terminology and large-τ conditions

Plain-language setting	Paper	Complexity	Key-length condition
with qRAM	QACM	$T = M = 2^{(2\tau-n)/5}$	$k \geq (\tau + 2n)/5$
without qRAM	CACM	$T = 2^{2(2\tau-n)/9}$, $M = 2^{(2\tau-n)/9}$	$k \geq (\tau + n)/3$
stronger writable qRAM	QAQM	$T = M = 2^{(2\tau-n)/5}$	$(3\tau + n)/10 \leq k \leq n$

- **QACM:** a classically prepared table supports coherent read access.
- **CACM:** an ordinary classical table; checking it inside the quantum search is more expensive.
- **QAQM:** a quantum-maintained data structure supporting coherent updates; this is a strictly stronger assumption and is omitted from the main talk.

G. Brassard et al., LATIN 1998; A. Chailloux et al., ASIACRYPT 2017; X. Bonnetain et al., ePrint 2025/1690.