FHE.org 2026 ▪ TAIPEI ▪ MARCH 8, 2026

# ITERATED HASH FUNCTIONS OVER ENCRYPTED DATA

## NEWS FROM THE FRONT

OLIVIER BERNARD    MARC JOYE

# Motivation

**Goal: Efficient Hash Functions Evaluable under FHE**

- FHE cost dominated by:
  - algebraic depth
  - bootstrapping (PBS) count
  - circuit regularity
- **Lightweight block ciphers** attractive for FHE
- **PRINCEv2: 64-bit block, 128-bit key, low algebraic complexity**

**Challenge**

**64-bit block size** insufficient for **128-bit collision resistance** using standard DBL constructions

# Iterated Hashing

## Merkle–Damgård Construction

1. Compression function $F \colon \{0,1\}^L \times \{0,1\}^\ell \to \{0,1\}^L$
2. Hash of a padded message $M = m_1 \| \ldots \| m_t$ is computed as

$$\begin{cases} h_0 = \mathrm{IV} \\ h_i = F(h_{i-1}, m_i) \quad \text{for } i = 1, \ldots, t \end{cases}$$

3. Then $H(M) = h_t$

- Hash output length $L \rightsquigarrow$ collision resistance $\approx 2^{L/2}$
- 128-bit security $\rightsquigarrow L \geq 256$

# Collision Resistance Requirement

**Problem Statement**

Given:

$$E_K : \{0,1\}^n \to \{0,1\}^n \quad \text{with } n = 64,\ |K| = 2n$$

Need:

$$\text{Collision resistance} \approx 2^{128}$$

**Observation:**

- Double-block-length (DBL) → output size $L = 2n$ → birthday bound $2^{64}$
- Require **quadruple-block-length (QBL) compression** (output size $L = 4n$)

**Issue:** Counter-4DM requires $4n$-bit key → incompatible with **(**$n, 2n$**)** ciphers

# Multi-Block-Length Constructions

**Counter-*b*DM [AFL+14]**

$$K_i = h_{i-1}^{(2)} \| h_{i-1}^{(3)} \| \cdots \| h_{i-1}^{(b)} \| m_i$$

For $j = 1, \ldots, b$:

$$h_i^{(j)} = E_{K_i}\left(h_{i-1}^{(1)} \oplus (j-1)\right) \oplus h_{i-1}^{(1)}$$

Remarks:

- Works generically for $b \geq 2$
- Ensures distinct inputs via counter constants $(j - 1)$

**Problem:**

$$|K_i| = 4n$$

(For $b = 4$, requires 256-bit key when $n = 64$)

# Multi-Block-Length Constructions

**Counter-$b$DM [AFL+14]**

$$K_i = h_{i-1}^{(2)} \| h_{i-1}^{(3)} \| \cdots \| h_{i-1}^{(b)} \| m_i$$

For $j = 1, \ldots, b$:

$$h_i^{(j)} = E_{K_i}(h_{i-1}^{(1)} \oplus (j-1)) \oplus h_{i-1}^{(1)}$$

Remarks:

- Works generically for $b \geq 2$
- Ensures distinct inputs via counter constants $(j-1)$

**Problem:**

$$|K_i| = 4n$$

(For $b = 4$, requires 256-bit key when $n = 64$)

**Core Challenge**

We only have:

$$E_K : (n, 2n)$$

Need behavior of:

$$(n, 4n)$$

**Strategy:**

- Extend effective key space
- Preserve FHE efficiency
- Avoid structural collision collapse

# Tweaking PRINCEv2

**Extend effective key size without increasing master key**

**Emulate $(n, 4n)$ cipher from $(n, 2n)$**

- Use tweak inputs (TWEAKEY framework)
- Inject chaining words + message into tweak/key
- Add 4 extra rounds for security margin

**Result:**

$$\widehat{E}_K(T_0, T_1, X)$$

⇝ Behaves like effective $4n$-bit key

# Tweaking PRINCEv2

**Extend effective key size without increasing master key**

**Emulate $(n, 4n)$ cipher from $(n, 2n)$**

- Use tweak inputs (TWEAKEY framework)
- Inject chaining words + message into tweak/key
- Add 4 extra rounds for security margin

**Result:**

$$\widehat{E}_K(T_0, T_1, X)$$

⇝ Behaves like effective $4n$-bit key

**Proposition (Orthomorphisms)**

Define:

$$\Theta_k(x) = (x \ggg k) \oplus (x \lll (n-k))$$

If $2 \leq 2k < n$, then:

$$\Theta_k \text{ and } \Theta_k \oplus \mathrm{id}$$

are permutations

Crucial for:

- Round tweak permutation
- Nibble-aligned implementation

# Modified QBL Compression (1/2)

## Main Construction

Define constants:

$$C_0 = 0\,, \quad C_1, C_2, C_3 \neq 0$$

with

$$C_1 \oplus C_2 \oplus C_3 \neq 0$$

### Compression

1. Let

$$K_i = h_{i-1}^{(2)} \| h_{i-1}^{(3)} \| h_{i-1}^{(4)} \| m_i$$

2. For $j = 1, \ldots, 4$:

$$h_i^{(j)} = E_{K_i}\left(h_{i-1}^{(1)} \oplus C_{j-1}\right) \oplus h_{i-1}^{(1)}$$

**Crucial change:** Avoid structural cancellation caused by $C_1 \oplus C_2 \oplus C_3 = 0$

# Modified QBL Compression (2/2)

## Main Construction

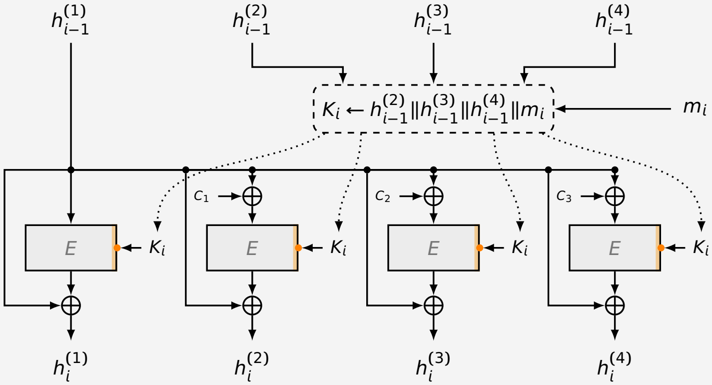


$(C_1 \oplus C_2 \oplus C_3 \neq 0)$

# Security Model

## Ideal-cipher model

- Independent random permutation per key
- Forward + inverse queries

## Super-query technique

- Handles regimes near $2^n$
- $2^{n-1}$ query entries for a fixed key $K$ imply all remaining queries for that key for free (**super query**)
- Query $E_K(X)$ gives $E_K(X \oplus C_j)$ for free

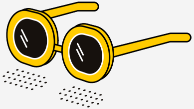

## Collision events

Three events:

1. NormalQueryWin
2. SuperQueryWin
3. **SameQueryWin** (internal collisions)

Dominant term in practice:

Internal collisions

# Collision Bound

**Final bound**

Let $N = 2^n$

$$\mathbf{Adv}_H^{\mathrm{coll}}(q) \leq 3 \cdot 2^{10} \frac{q^2}{N^4} + 7 \cdot 2^6 \frac{q}{N^3}$$

## Why Counter-4DM Is Weaker

- Counter-4DM uses:

$$C_j = j - 1$$

- Hence $C_1 \oplus C_2 \oplus C_3 = 0$
    - Internal equations collapse from 3 independent relations to 2
    - $\Pr[\mathrm{SameQueryWin}] \leq 6 \cdot 2^4 \frac{q}{N^2}$
- Improvement over Counter-4DM: $\frac{q}{N^3}$ vs $\frac{q}{N^2}$
- Linear term dominates for $q \leq 2^n$

# QBL–MDC Compression

**Native QBL Construction #1**

- Two Hirose-like cores
- Explicit MDC-2-like cross-branch diffusion

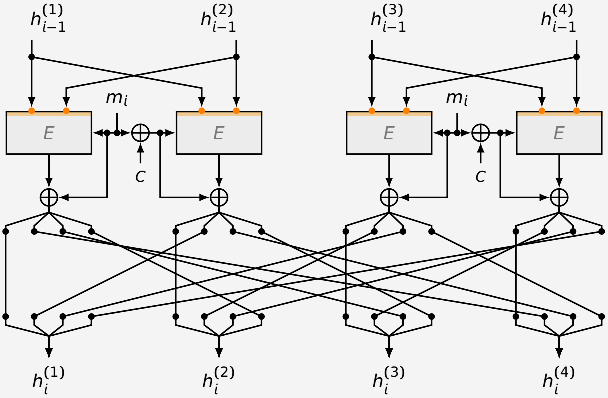

# QBL–TDM Compression

## Native QBL Construction #2

- Inspired by Tandem-DM
- Two nonlinear subkey derivations

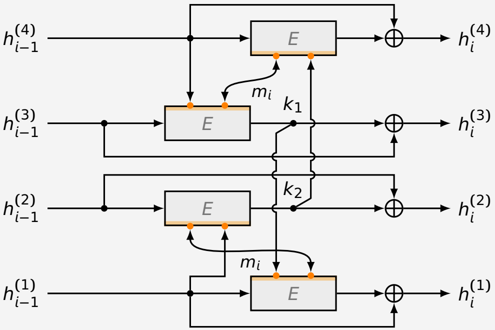

# FHE Implementation Strategy

**Hybrid LUT approach:** Flexible output formats

- Instead of fixed 4 → 4 LUTs, allow more flexible outputs
  - Extract shifted (pairs of) bits directly
  - Recombine via inexpensive ciphertext additions
- Effect
  - Reduced PBS count & depth
  - Better plumbing across layers

**Layer engineering:** Optimized circuit structure and data flow

**Key schedule in FHE:** (iterated hashing setting ≠ transciphering setting)

- Incorporate round constants into specialized S-box LUTs
- Effect: Key schedule is virtually free

# Validation & Results

## PRINCEv2 Block Cipher

- **Experimental setup**
  - Amazon AWS `hpc7a.96xlarge` with 64 cores
  - `PARAM_MESSAGE_2_CARRY_2_KS_PBS_GAUSSIAN_2M128`
  - Failure probability $2^{-128}$

### Results

| Number of cores | 1 | 2 | 4 | 8 | 16 | 32 | 64 |
|---|---|---|---|---|---|---|---|
| Timing (s) | 33.123 | 19.527 | 10.236 | 5.132 | 2.701 | 1.492 | 0.776 |

- Excellent scalability with respect to the number of cores
- Low latency: 776 ms per call (64 cores)
- Competitive single-core performance

# Conclusion

## Summary of Key Contributions

- **Quadruple-block-length compression** tailored to a lightweight $(n, 2n)$ cipher
  - careful constant selection & orthomorphisms for FHE
  - improved internal-collision bound
- **Native QBL candidates** for $(n, 2n)$ constructions (open for analysis)
- **FHE-aware compression design**
  - optimized homomorphic implementation of PRINCEv2
  - sub-second TFHE evaluation at $2^{-128}$ failure probability

🕶 Read the full paper at ePrint 2026/309

# Contact and Links

marc@zama.ai

zama.org

github.com/zama-ai

zama.org/community