Refining Probabilistic-Linearization TIDA for 5-Round SHA3-384 Collision Attacks

Luhan Yan[1] Zhenzhen Bao[1,2,3(✉)] Huina Li[1(✉)]

[1]Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing, China

[2]Zhongguancun Laboratory, Beijing, China

[3]State Key Laboratory of Cryptography and Digital Economy Security, Tsinghua University, Beijing, China

SAC 2026 — Ottawa, Canada
August 27, 2026

Outline

1. Background & Motivation
2. Contribution 1: SA-PIDS
3. Contribution 2: Revised Complexity Analysis
4. Updated Attack on 5-Round SHA3-384
5. Conclusion

Outline

1. Background & Motivation
2. Contribution 1: SA-PIDS
3. Contribution 2: Revised Complexity Analysis
4. Updated Attack on 5-Round SHA3-384
5. Conclusion

SHA-3 and the KECCAK permutation

SHA-3 family (FIPS 202, 2015 [5])

- Built from the KECCAK permutation [1] KECCAK-$f[1600]$, 24 rounds.
- *Sponge* of width $b = r + c = 1600$.
- SHA3-d: $c = 2d$, suffix 01.

Round function $R = \iota \circ \chi \circ \pi \circ \rho \circ \theta$

- θ, ρ, π, ι: *linear*; only ι breaks z-translation invariance.
- χ: the only nonlinear layer, 5-bit S-box $a_i \leftarrow a_i \oplus (a_{i+1} \oplus 1)\, a_{i+2}$.

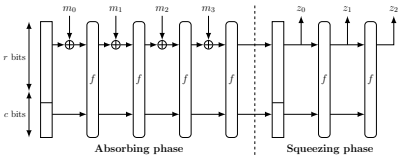

Instances	r	c	d
SHA3-224	1152	448	224
SHA3-256	1088	512	256
SHA3-384	**832**	**768**	**384**
SHA3-512	576	1024	512
SHAKE128	1344	256	*var.*
SHAKE256	1088	512	*var.*

Collision attacks on SHA-3: where we stand

Function	Rounds	Complexity	Method	Reference
SHA3-224 / -256	5	practical	Differential	Song *et al.* [7], Guo *et al.* [3]
SHA3-224 / -256	5	$2^{96.67}$	Internal diff.	Zhang *et al.* [8]
SHAKE128	5	practical	Differential	Qiao *et al.* [6]
SHA3-384	4	$2^{59.64}$	Differential	Huang *et al.* [4]
	5	$2^{170.73}$	Internal diff. + PL-TIDA	Zhang *et al.* [8]
	5	$\mathbf{2^{164.11}}$	**Internal diff. + PL-TIDA**	**This work**
SHA3-512	4	$2^{225.29}$	Internal diff.	Zhang *et al.* [8]
SHAKE256	6	$2^{232.29}$	Internal diff.	Zhang *et al.* [8]

Internal differentials and the squeeze attack

Symmetric state with period i [2]: $\mathbf{A}[x][y][z] = \mathbf{A}[x][y][(z+i) \bmod w]$ ($i \mid w$; here i=32, so $\mathbf{A} = (\hat{\mathbf{A}}, \hat{\hat{\mathbf{A}}})$).
Internal difference:

$$\Delta(\mathbf{v}) = \hat{\mathbf{A}} \oplus \hat{\hat{\mathbf{A}}}.$$

Tracks structure *inside* one state (not between two). Works on KECCAK: θ, ρ, π, χ preserve z-periodicity; only the sparse ι breaks it.

Squeeze attack (Dinur *et al.* [2]):

- Structured input subset $S' \to$ small output subset D' with prob. p.
- If $p^2 > q = |D'|/|D|$, collision costs $\sqrt{q|D|}/p$.

$\Rightarrow$ **Internal-diff. char. + subset birthday search** — combined into a 3-stage framework (next slide).

The 3-stage attack framework (Zhang *et al.* [8])

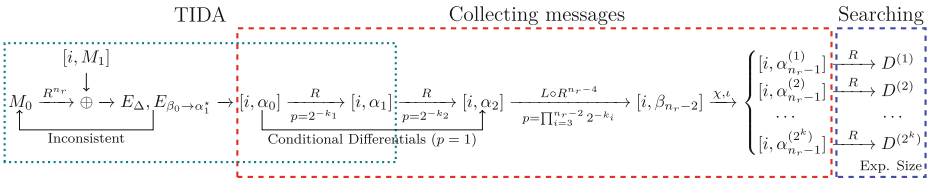

Prerequisites. An internal differential characteristic starting from the target difference α_1^* (searched *offline*); **two-block** messages $M_0 \| M_1$.

1. **TIDA**: build connector linear equation system E_Δ; sample M_0, find a valid β_0, solve for an M_1-subspace.
2. **Collect**: solve the transition systems, propagate the M_1-subspaces through the remaining rounds, store outputs in collision subsets.
3. **Search**: birthday search inside the collision subsets until a collision is found.

Both our contributions concern Stage 1 (TIDA) only.

The baseline: probabilistic-linearization TIDA

Connector construction = linear algebra around one χ layer.

- Non-active S-box ($\delta_{\text{out}} = 0$): force $\delta_{\text{in}} = 0$ via 5 linear equations.
- Active S-box ($\delta_{\text{out}} \neq 0$): instead of fixing δ_{in} to a 2-D affine subspace (deterministic, many eqs.), **relax** to a higher-dimensional subspace $U \subseteq \mathbb{F}_2^5$ from MDST[1] (3-D or 4-D).

Trade-off:

- Fewer linear equations $\Rightarrow$ smaller capacity subsystem $E_{\mathcal{C}} \Rightarrow$ cheaper M_0 sampling.
- But not every $\delta_{\text{in}} \in U$ is valid for $\delta_{\text{out}} \Rightarrow$ a probability factor.

PIDS produces a probabilistic input-difference system E_Δ with

$$p_1^* = \prod_{a \in \mathcal{A}} \underbrace{\frac{\#\{\delta \in U_a \colon \delta \to \delta_{\text{out}}^{(a)}\}}{|U_a|}}_{\text{local diff. density}}.$$

MDST = maximum-difference-density subspace table: for each δ_{out}, the affine subspaces $U \subseteq \mathbb{F}_2^5$ with the highest fraction of input differences that can produce δ_{out}.

Two gaps in the current PL-TIDA

Gap 1 — PIDS construction is a **global** problem, but solved *locally*

Greedy-PIDS picks each S-box's subspace to not raise $\mathrm{rank}(E_{\mathcal{C}})$ *immediately*. But constraints interact only after the linear layer + Gaussian elimination: a locally good choice can still inflate $E_{\mathcal{C}}$, shrink the remaining space, or hurt p_1^*.

Gap 2 — evaluating p_1: **imprecise** proxy, **infeasible** estimation

$$\underbrace{p_1^* = \frac{|\mathcal{V}|}{|\mathcal{S}_\Delta|}}_{\text{global density over } \textit{all} \text{ solutions}} \quad \neq \quad \underbrace{p_1 = \Pr\left[\, \mathcal{S}_{\mathcal{R},\mathcal{C}^*} \text{ contains a valid } \beta_0 \;\middle|\; \Delta_{\mathcal{C}}^* \in \mathcal{S}_{\mathcal{C}} \,\right].}_{\text{what the attack actually samples}}$$

- *The proxy is imprecise*: projection *folds* several valid assignments onto the same capacity value (at most $|\mathcal{S}_{\mathcal{R}}|$ each) $\Rightarrow p_1^* \leq p_1$, i.e. it **overestimates** the complexity.
- *Direct estimation*: sampling valid input differences needs $\sim 1/p_1$ trials $\Rightarrow$ **inapplicable at low density**.

Our contributions

1. **SA-PIDS** — a *simulated-annealing* search for affine-subspace assignments that
   - treats PIDS as a **global multi-objective optimization**: $\mathrm{rank}(E_{\mathcal{C}})$, p_1^*, and the remaining solution-space dimension, jointly;
   - escapes the local optima that trap Greedy-PIDS.
2. **A structural evaluation of the true connector probability p_1.**
   - Formal relation $p_1 = p_1^* \cdot \dfrac{|\mathcal{S}_{\mathcal{R}}|}{\bar{\kappa}}$, with $\bar{\kappa}$ the avg. projection multiplicity.
   - **Harmonic-mean estimator** of $\bar{\kappa}$ that corrects sampling bias.

Symbiosis — the two contributions **unlock** each other

SA-PIDS opens a new (low-density) region $\longleftrightarrow$ **structural evaluation of p_1** makes that region *quantifiable*

neither alone moves the attack; together: $\mathbf{-6.6}$ bits.

5-round SHA3-384 (same characteristic)	Zhang *et al.* [8]	**This work**
Total complexity	$2^{170.73}$	$\mathbf{2^{164.11}}$

Outline

1. Background & Motivation
2. Contribution 1: SA-PIDS
3. Contribution 2: Revised Complexity Analysis
4. Updated Attack on 5-Round SHA3-384
5. Conclusion

PIDS as a global combinatorial optimization

Search space. For each active S-box $a \in \mathcal{A}$, pick one subspace $U_a \in \mathcal{C}_a = \mathrm{MDST}[\delta_{\mathrm{out}}^{(a)}]$. A *state* is an assignment

$$\mathcal{S} : \ \mathcal{A} \to \bigcup_a \mathcal{C}_a.$$

Why global? The score of an assignment $\mathcal{S}$ is only visible *after*:

1. building E_Δ from all chosen subspaces,
2. pulling constraints back through the KECCAK linear layer,
3. Gaussian elimination $\Rightarrow E_{\mathcal{C}},\ E_{\mathcal{R}}$.

Three coupled objectives ((s_1, s_2, p_1^*) all log-scale):

- $s_1 = \mathrm{rank}(E_{\mathcal{C}})$: drives the cost of sampling a good M_0 ($\propto 2^{s_1}$).
- p_1^*: local density of valid β_0.
- s_2: $\log_2$ of the remaining solution space after fixing $\Delta_{\mathcal{C}}$ (i.e., $\dim$ of the solution space of $E_{\mathcal{R}}$).

SA-PIDS: simulated annealing over subspace assignments

Simulated annealing: key idea

Take a random walk on $\mathcal{S}$: each step applies a random change to the current $\mathcal{S}$; **accept worse moves with probability** $\exp(-\Delta/T)$ (Δ = score change), so the walk can *escape local optima*; slowly cool T from *hot* (explore broadly) to *cold* (settle greedily).

Four ingredients.

- **Warm start** from Greedy-PIDS (a feasible assignment).
- **Neighborhood**: replace the subspace of *one* active S-box, $\mathcal{S}'(a) \in \mathcal{C}_a \setminus \{\mathcal{S}(a)\}$.
- **Metropolis rule**: if $\Delta = \text{Score}(\mathcal{S}') - \text{Score}(\mathcal{S}) < 0$ accept; else accept with prob. $\exp(-\Delta/T)$.
- **Cooling**: $T \leftarrow \lambda T$ after I_0 moves; keep the best state. (T_0=500, I_0=50 moves/level, I_1=500 levels, λ=0.9.)

Evaluation function.

$$\text{Score}(\mathcal{S}) \;=\; s_1 \;-\; \log_2 p_1^* \;-\; s_2$$

Outline

1. Background & Motivation
2. Contribution 1: SA-PIDS
3. Contribution 2: Revised Complexity Analysis
4. Updated Attack on 5-Round SHA3-384
5. Conclusion

$p_1^* \neq p_1$: the projection gap

Setup. E_Δ splits as $(\Delta_\mathcal{R}, \Delta_\mathcal{C})$-variables into

$$E_\mathcal{C}(\Delta_\mathcal{C}) \quad \text{and} \quad E_\mathcal{R}(\Delta_\mathcal{R}, \Delta_\mathcal{C}).$$

For each M_0 producing $\Delta_\mathcal{C}^* \in \mathcal{S}_\mathcal{C}$, the remaining system $E_{\mathcal{R},\mathcal{C}^*}$ has solution space $\mathcal{S}_{\mathcal{R},\mathcal{C}^*}$ of *constant* size $|\mathcal{S}_\mathcal{R}|$.

Define (with the natural projection $\mathrm{proj}(\Delta_\mathcal{R}, \Delta_\mathcal{C}) = \Delta_\mathcal{C}$):

- $\mathcal{V} = \mathcal{S}_\Delta \cap \mathcal{S}_N =$ valid solutions.
- $\mathcal{V}_\mathcal{C} = \mathrm{proj}(\mathcal{V}) \subseteq \mathcal{S}_\mathcal{C} =$ *images* of valid solutions under the projection.
- $\bar{\kappa} = |\mathcal{V}|/|\mathcal{V}_\mathcal{C}| =$ average *projection multiplicity*.

Two relevant probabilities:

$$p_1^* = \frac{|\mathcal{V}|}{|\mathcal{S}_\Delta|} \text{ — global density;} \qquad p_1 = \frac{|\mathcal{V}_\mathcal{C}|}{|\mathcal{S}_\mathcal{C}|} \text{ — what the attack actually samples.}$$

$\mathcal{S}_\Delta$ = affine solution space of E_Δ; $\mathcal{S}_N$ = solution set of the exact nonlinear system; $\mathcal{S}_\mathcal{C}$ = solution set of the capacity subsystem $E_\mathcal{C}$.

A structural formula for p_1

Using $|\mathcal{V}_\mathcal{C}| = |\mathcal{V}|/\bar{\kappa}$ (definition of $\bar{\kappa}$) and $|\mathcal{S}_\Delta| = |\mathcal{S}_\mathcal{R}| \cdot |\mathcal{S}_\mathcal{C}|$ (affine-space decomposition),

$$p_1 = \frac{|\mathcal{V}_\mathcal{C}|}{|\mathcal{S}_\mathcal{C}|} = \underbrace{\frac{|\mathcal{V}|}{|\mathcal{S}_\Delta|}}_{p_1^*} \cdot \underbrace{\frac{|\mathcal{S}_\mathcal{R}|}{\bar{\kappa}}}_{\text{effective remaining freedom}}$$

Effective remaining degrees of freedom

$d_{\text{eff}} = \log_2 |\mathcal{S}_\mathcal{R}| - \log_2 \bar{\kappa}$ — the part of $|\mathcal{S}_\mathcal{R}|$ that *genuinely* contributes to p_1.

Estimating $\bar{\kappa}$ via the harmonic mean

Two-stage experiment.

- *Outer*: draw a valid $\beta_0 \Rightarrow \Delta_{\mathcal{C}}^* \in \mathcal{V}_{\mathcal{C}}$ (prob. $\propto N_j$).
- *Inner*: count $N^{(\Delta_{\mathcal{C}}^*)} = \#$ valid β_0 in $\mathcal{S}_{\mathcal{R},\mathcal{C}^*}$.

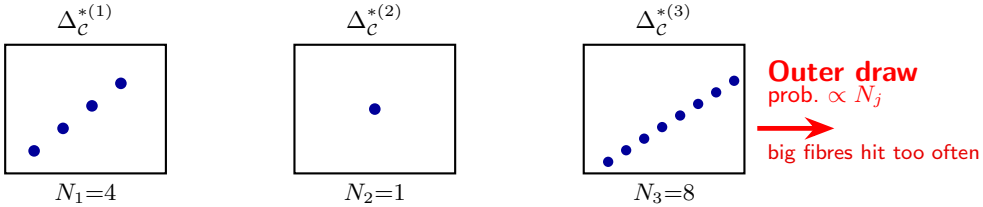

Arithmetic mean: samples weighted by $N_j \Rightarrow$ **over-estimates** $\bar{\kappa}$.

Harmonic mean: $1/N_j$ weights correct the size bias $\Rightarrow$ **consistent** estimator of $\bar{\kappa}$.

By the law of large numbers, $\widehat{\kappa}_{\mathrm{harm}} \xrightarrow{m\to\infty} \bar{\kappa}$. *(Full derivation: backup slide.)*

Outline

1. Background & Motivation
2. Contribution 1: SA-PIDS
3. Contribution 2: Revised Complexity Analysis
4. Updated Attack on 5-Round SHA3-384
5. Conclusion

The updated attack on 5-round SHA3-384

Same characteristic as Zhang et al. [8] (period $i = 32$; $(k_2, k_3, k_4) = (25, 18, 16)$; 77 active S-boxes in α_1^*).

Connector from SA-PIDS (with $\dim(E_{\mathcal{R}}) \leq 28$ threshold so $\bar{\kappa}$ is estimable)[2]:

$$p_1^* = 2^{-91.19}, \quad |E_{\mathcal{C}}| = 85 \ (p_0 = 2^{-85}), \quad \dim(E_{\mathcal{R}}) = 28.$$

Estimate $\bar{\kappa}$ (2^{14} outer samples, exhaustive inner):

$$\widehat{\kappa}_{\text{harm}} = 2^{5.34} \ \Rightarrow \ p_1 = 2^{-91.19} \cdot 2^{28-5.34} = \mathbf{2^{-68.53}}.$$

Remaining factors (as in Zhang *et al.* [8]): $p_2 = 2^{-10.58}$; message collection $\approx 2^{151.38}$.

$$\text{TIDA cost} \ = \ 2^{85+68.53+10.58} \ \approx \ \mathbf{2^{164.11}}.$$

[2] $p_0 = 2^{-|E_{\mathcal{C}}|}$: per-try success probability when sampling an M_0 that satisfies the capacity subsystem $E_{\mathcal{C}}$.

Stage-by-stage comparison

Metric	Zhang *et al.* [8]	**This work**	Metric	Zhang *et al.* [8]	**This work**
$\lvert E_{\mathcal{C}}\rvert$	90	**85**	$\dim$ sol. of $E_{\mathcal{R}}$	10	**28**
$-\log_2 p_1^*$	76.97	91.19	Verifying β_0 ($\log_2$)	92.19	107.11
$-\log_2 p_1$	70.15	**68.53**	Message collection	151.38	151.38
$-\log_2 p_2$	10.58	10.58	**Final complexity**	$2^{170.73}$	$\mathbf{2^{164.11}}$

- Two small gains compound: -5 bits in $E_{\mathcal{C}}$ *and* ~ 1.6 bits in p_1.
- Both improvements target the **connector phase**; useful when (and only when) it dominates — exactly the SHA3-384 case.
- Source: `github.com/KeccakInternalDiff/keccak-internal-differential`

Outline

1. Background & Motivation
2. Contribution 1: SA-PIDS
3. Contribution 2: Revised Complexity Analysis
4. Updated Attack on 5-Round SHA3-384
5. Conclusion

Conclusion

Two refinements of probabilistic-linearization TIDA.

1. **SA-PIDS**: global simulated-annealing search over affine-subspace assignments, with a score jointly trading off $E_{\mathcal{C}}$, p_1^* and the remaining freedom.
2. **Structural p_1 evaluation**: $p_1 = p_1^* \cdot |\mathcal{S}_{\mathcal{R}}|/\bar{\kappa}$, with a consistent harmonic-mean estimator of $\bar{\kappa}$.

Result

On 5-round SHA3-384, using the same characteristic in [8]:

$$2^{170.73} \longrightarrow \mathbf{2^{164.11}} \qquad (\approx -6.6 \text{ bits})$$

Limitations and open problem

Limitations.

- Only helps when the *connector* phase is the bottleneck — for SHA3-224/-256 and SHAKE128, the bottleneck of existing attacks lies in other stages.
- SHA3-512: smaller rate $\Rightarrow$ tighter degree-of-freedom budget.
- SHAKE256 (6-round): message collection and search dominate.

Open problem

Design a subspace selection that **directly** optimises $d_{\text{eff}} = \log_2 |\mathcal{S}_{\mathcal{R}}| - \log_2 \bar{\kappa}$, instead of the proxy score used by SA-PIDS?

Backup — derivation of the harmonic-mean identity

Recall the size-biased outer draw: $\Pr[\text{draw } \Delta_{\mathcal{C}}^{*(j)}] = N_j/|\mathcal{V}|$. Then

$$\mathbb{E}\left[\frac{1}{N_j}\right] \;=\; \sum_{j=1}^{|\mathcal{V}_{\mathcal{C}}|} \Pr[\text{draw } \Delta_{\mathcal{C}}^{*(j)}] \cdot \frac{1}{N_j} \;=\; \sum_{j=1}^{|\mathcal{V}_{\mathcal{C}}|} \frac{N_j}{|\mathcal{V}|} \cdot \frac{1}{N_j} \;=\; \sum_{j=1}^{|\mathcal{V}_{\mathcal{C}}|} \frac{1}{|\mathcal{V}|} \;=\; \frac{|\mathcal{V}_{\mathcal{C}}|}{|\mathcal{V}|}.$$

Taking reciprocals,

$$\left(\mathbb{E}\left[\frac{1}{N_j}\right]\right)^{-1} \;=\; \frac{|\mathcal{V}|}{|\mathcal{V}_{\mathcal{C}}|} \;=\; \bar{\kappa}.$$

By the law of large numbers, the empirical harmonic mean $\widehat{\kappa}_{\mathrm{harm}} = \left(\frac{1}{m}\sum_j 1/N_j\right)^{-1}$ converges to $\bar{\kappa}$ as $m \to \infty$.

References I


Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche.
The KECCAK reference, 2011.


Itai Dinur, Orr Dunkelman, and Adi Shamir.
Collision attacks on up to 5 rounds of SHA-3 using generalized internal differentials.
In *FSE 2013*, volume 8424 of *LNCS*, 2014.


Jian Guo, Guohong Liao, Guozhen Liu, Meicheng Liu, Kexin Qiao, and Ling Song.
Practical collision attacks against round-reduced SHA-3.
Journal of Cryptology, 33, 2020.


Senyang Huang, Orna Agmon Ben-Yehuda, Orr Dunkelman, and Alexander Maximov.
Finding collisions against 4-round SHA-3-384 in practical time.
IACR Trans. Symm. Cryptol., 2022.


National Institute of Standards and Technology.
SHA-3 standard: Permutation-based hash and extendable-output functions, 2015.


Kexin Qiao, Ling Song, Meicheng Liu, and Jian Guo.
New collision attacks on round-reduced Keccak.
In *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, 2017.

References II

Ling Song, Guohong Liao, and Jian Guo.
Non-full sbox linearization: Applications to collision attacks on round-reduced Keccak.
In *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, 2017.

Zhongyi Zhang, Chengan Hou, and Meicheng Liu.
Probabilistic linearization: Internal differential collisions in up to 6 rounds of SHA-3.
In *CRYPTO 2024, Part IV*, volume 14923 of *LNCS*, 2024.

Thank you!

Questions & discussion