Differential–Linear Distinguisher Search via Walsh Spectra over Low-Dimensional Mask Subspaces

Chenlu Zheng[1], Liu Zhang[2] and Zilong Wang[1]

[1] Xidian University, Xi'an, China
[2] Nanyang Technological University, Singapore, Singapore

Abstract

Searching for differential–linear (DL) distinguishers in ARX and AND–RX ciphers often requires empirical evaluation of multi-round middle correlations. Hourglass(-like) structures provide effective guidance by restricting middle differences and output masks to sparse forms. However, the corresponding low-weight output-mask position need not be the most favorable boundary for empirical middle-correlation evaluation. Exploring alternative boundaries may involve multi-bit output masks, whose numerous possible combinations make empirical evaluation expensive.

We propose a spectral DL search framework for efficiently evaluating multiple middle output masks. We use Hamming-weight-one correlations to select low-dimensional output-mask subspaces and apply the fast Walsh–Hadamard transform to jointly evaluate all masks in each selected subspace using shared samples. Experiments on 32-bit ciphers show that HW1-guided subspaces can recover DL candidates with correlation estimates comparable to those obtained through full-state spectral evaluation. We integrate this approach with round-split selection, automated outer-trail search, and independent fixed-key evaluation of the retained middle correlations. Experiments on SIMON, SIMECK, SPECK, and CHAM yield DL distinguishers with improved round coverage or estimated correlations. In particular, we obtain a 36-round SIMON128 distinguisher under single-trail outer-component evaluation and a 39-round result through refined outer-component evaluation. Our results further show that evaluating multi-bit masks at alternative middle-output boundaries can reveal more readily detectable correlations or reduce reliance on single-trail approximations.

Keywords: Differential–linear cryptanalysis, Walsh–Hadamard transform, Low-dimensional mask subspaces, Middle-correlation evaluation

1 Introduction

Differential–linear (DL) cryptanalysis, introduced by Langford and Hellman [LH94], combines differential propagation and linear approximations to construct cryptographic distinguishers. Bar-On et al. [BDKW19] introduced the differential–linear connectivity table (DLCT) at EUROCRYPT 2019 to characterize DL correlations associated with input differences and output linear masks. In the three-part framework, a cipher is decomposed as $E = E_2 \circ E_m \circ E_1$, where E_1 is the differential prefix, E_m the middle DL component,

Emails: clzheng@stu.xidian.edu.cn (Chenlu Zheng), liu.zhang@ntu.edu.sg (Liu Zhang), zlwang@xidian.edu.cn (Zilong Wang)

Date of this document: 2026-11-19.

and E_2 the linear suffix. The correlation of a complete DL distinguisher is typically approximated by combining the differential probability and linear correlation of the outer components with the DL correlation of the middle component.

Various automated approaches based on MILP, MIQCP, CP, and SAT have been developed for DL search [BGG+23, ZWH24, CXZQ24, HDE24, LZ26]. Recent advances have also provided new approaches to computing and approximating DL correlations. Hu, Niu, and Wang [HNW26] proposed a geometric approach to round-based DL correlation approximation, obtaining improved distinguishers for several ciphers, including PRESENT. The efficient realizations developed in their work primarily target S-box-based structures and certain large nonlinear functions. He et al. [HHN+26] developed a Walsh-transform-based formulation of DL analysis and derived correlation models for modular addition. Hadipour et al. [HMAV26] investigated high-performance Walsh–Hadamard transforms for exact DL correlation computation and key-dependence analysis. While these advances provide effective tools for particular cipher structures and settings, accurate and computationally tractable evaluation of multi-round middle DL correlations remains challenging for ARX and AND–RX ciphers, especially when a large number of difference–mask pairs must be considered.

To make middle-candidate exploration computationally tractable, several heuristic search strategies have been developed for ARX ciphers. Chen et al. [CBY23] proposed a semi-unconstrained search method at ASIACRYPT 2023, combining strong unbalanced bits (SUBs), iterative output-mask combinations, and meet-in-the-middle (MITM) techniques to explore multi-bit middle output masks. Their approach improved DL distinguishers for SPECK and LEA. Gong et al. [GWH+25] subsequently observed that several high-quality ARX DL distinguishers could be repartitioned into hourglass or hourglass-like structures. In a typical hourglass structure, both the middle input difference and output mask are single-bit, whereas an hourglass-like structure may have two adjacent active bits in the middle output mask. Exploiting these structural properties, Gong et al. developed search algorithms combining SUBs, bidirectional linear extension (BLE), and MITM techniques, obtaining improved DL distinguishers for several ARX ciphers. These studies demonstrate that structural guidance can substantially reduce the number of middle candidates requiring correlation evaluation.

Although hourglass(-like) structures provide effective guidance for DL search, the existence of such a representation does not necessarily imply that its low-weight output-mask position is the most favorable boundary for empirical middle-correlation evaluation. When the middle component ends at this position, restricting the output mask to a single bit can substantially reduce the candidate space. If it ends before this position, a shorter middle component with a multi-bit output mask may exhibit a stronger correlation that is easier to detect reliably with limited samples. Conversely, when the middle component extends beyond this position, additional linear propagation can be incorporated into direct empirical evaluation, reducing reliance on single-trail approximations when the sampling budget permits. For example, some multi-bit middle output masks identified by Chen et al. [CBY23] occur before the corresponding single-bit output-mask positions in the hourglass repartitions identified by Gong et al. [GWH+25].

Exploring alternative middle-output boundaries, however, often requires evaluating multi-bit output masks with many possible combinations of active positions. Even when restricted to relatively low Hamming weights, the number of candidate masks can be substantial. For ARX and AND–RX ciphers, multi-round middle correlations are often evaluated empirically using large numbers of samples. Evaluating candidate masks individually requires repeated correlation computations over the same output-difference samples, making the computational cost grow with both the number of masks and the sample size. Moreover, detecting weaker correlations requires larger sample sizes, further increasing the evaluation cost. Under a fixed computational budget, this creates a trade-off between

the range of output masks explored and the strength of correlations that can be reliably detected. This raises a central question: how can we efficiently evaluate the empirical DL correlations of many middle output masks for a fixed input difference, thereby supporting distinguisher search across different middle-output boundaries?

Our Contributions. To address this challenge, we develop a spectral approach to evaluating multiple middle output masks and integrate it into an automated three-part DL search framework. Our main contributions are summarized as follows. [3].

- **Efficient multi-mask correlation evaluation via low-dimensional Walsh spectra.** We propose an HW1-guided strategy for selecting low-dimensional output-mask subspaces and use the fast Walsh–Hadamard transform (FWHT) to jointly evaluate the empirical DL correlations of all masks in each selected subspace. By projecting shared output-difference samples onto selected coordinates, the method reduces the correlation-computation complexity from $O(N2^k)$ to $O(Nk + k2^k)$ for a subspace containing 2^k masks and N samples. Experiments on 32-bit ciphers show that low-dimensional subspaces can yield DL candidates with correlation estimates comparable to those obtained through full-state spectral evaluation.
- **An automated DL search framework with flexible middle-output boundaries.** We integrate low-dimensional spectral evaluation with round-split selection, middle-mask filtering, and automated outer-trail search to construct a three-part DL distinguisher search framework. Candidate discovery uses mixed-key sampling, followed by independent fixed-key evaluation of retained middle correlations. The framework exploits the sparse input differences suggested by hourglass(-like) structures without restricting the middle-output boundary to their low-weight mask positions. Our experiments show that alternative boundaries can yield more readily detectable middle correlations or reduce reliance on single-trail approximations.
- **Improved DL distinguishers and refined outer-component evaluation.** We apply the framework to several ARX and AND–RX ciphers, including SIMON, SIMECK, SPECK, and CHAM. Under single-trail outer-component evaluation, we obtain 22-, 29-, and 36-round DL distinguishers for `SIMON64`, `SIMON96`, and `SIMON128`, respectively, as well as 22- and 28-round results for `SIMECK48` and `SIMECK64`. Further refinement through differential-prefix aggregation and empirical linear-suffix evaluation yields 30- and 39-round distinguishers for `SIMON96` and `SIMON128`, respectively. Tables 1 and 2 compare our results with previous work. Additional experiments on PRESENT and KNOT demonstrate the applicability of the spectral evaluation method to other cipher structures.

The remainder of this paper is organized as follows. Section 2 introduces the necessary preliminaries for DL cryptanalysis and empirical correlation evaluation. Section 3 discusses hourglass(-like) structures, middle-output boundary selection, and the motivation for our approach. Section 4 presents the low-dimensional spectral evaluation method and the automated three-part DL search framework. Section 5 reports the experimental results and analysis. Finally, Section 6 concludes the paper.

2 Preliminaries

This section introduces the three-part differential–linear framework and the Walsh representation of DL correlations. Unless otherwise stated, differences and linear masks belong

[3] An anonymized repository for the code and experiments of this work is available at https://anonymous.4open.science/r/FSE_DL_FWHT-8F04/

Table 1: Comparison of differential–linear distinguishers under single-trail outer-component evaluation.

Cipher	R	(R_1, R_m, R_2)	Correlation	Reference
SIMON32	14	$(5, 6, 3)$	$2^{-17.49}$	[CXZQ24]
	14	$(3, 11, 0)$	$2^{-14.11}$	Section 5.4
SIMON48	18	$(7, 4, 7)$	$2^{-28.58}$	[CXZQ24]
	18	$(2, 12, 4)$	$2^{-22.69}$	Section 5.4
SIMON64	21	$(8, 5, 8)$	$2^{-36.91}$	[CXZQ24]
	22	$(5, 14, 3)$	$2^{-29.83}$	Section 5.4
SIMON96	26	$(9, 9, 8)$	$2^{-50.16}$	[CXZQ24]
	29	$(4, 17, 8)$	$2^{-46.16}$	Section 5.4
SIMON128	32	$(13, 10, 9)$	$2^{-65.61}$	[CXZQ24]
	36	$(8, 22, 6)$	$2^{-61.91}$	Section 5.4
SIMECK48	20	$(3, 14, 3)$	$2^{-21.89}$	[HDE24]
	22	$(3, 15, 4)$	$2^{-23.94}$	Section 5.4
SIMECK64	26	$(5, 16, 5)$	$2^{-30.35}$	[HDE24]
	28	$(4, 20, 4)$	$2^{-31.82}$	Section 5.4
SPECK48	12	$(3, 6, 3)$	$2^{-22.43}$	[HHN+26]
	12	$(2, 8, 2)$	$2^{-21.98}$	Section 5.4
SPECK96	16	$(4, 8, 4)$	$2^{-47.58}$	[GWH+25]
	16	$(3, 9, 4)$	$2^{-47.46}$	Section 5.4
SPECK128	18	$(5, 9, 4)$	$2^{-55.19}$	[GWH+25]
	18	$(5, 9, 4)$	$2^{-54.47}$	Section 5.4
CHAM64	46	$(6, 26, 14)$	$2^{-31.08}$	[Roh24]
	47	$(7, 33, 7)$	$2^{-29.91}$	Section 5.4
CHAM128	58	$(10, 32, 16)$	$2^{-58.86}$	[Roh24]
	60	$(10, 39, 11)$	$2^{-63.22}$	Section 5.4

to the corresponding binary vector spaces, and $\langle \cdot, \cdot \rangle$ denotes the standard inner product over $\mathbb{F}_2$.

2.1 Differential–Linear Distinguishers

We consider differential–linear (DL) distinguishers in the three-part framework of Bar-On et al. [BDKW19], as illustrated in Fig. 1. The target primitive is decomposed as

$$E = E_2 \circ E_m \circ E_1,$$

where E_1, E_m, and E_2 denote the differential prefix, the middle DL component, and the linear suffix, respectively. Let $\Delta_{\text{in}} \xrightarrow{E_1} \Delta_m$ have differential probability p, let $r(\Delta_m, \Gamma_m)$ denote the signed DL correlation of E_m, and let $\Gamma_m \xrightarrow{E_2} \Gamma_{\text{out}}$ have signed linear correlation q. We write r when Δ_m and Γ_m are clear from the context.

Under the conventional three-part approximation, the resulting DL correlation is estimated as

$$\text{Cor} \approx prq^2,$$

and its magnitude as

$$|\,\text{Cor}\,| \approx p|r|q^2.$$

Table 2: Comparison of differential–linear distinguishers with outer-path aggregation for the SIMON and SIMECK families.

Cipher	R	(R_1, R_m, R_2)	Correlation	Reference
SIMON32	14	$(7, 3, 4)$	$2^{-15.36}$	[CXZQ24]
	14	$(3, 11, 0)$	$2^{-14.11}$	Section 5.4
SIMON48	17	$(0, 12, 5)$	$2^{-20.34}$	[LZ26]
	18	$(2, 12, 4)$	$2^{-21.91}$	Section 5.5
SIMON64	21	$(0, 13, 8)$	$2^{-29.38}$	[LZ26]
	22	$(4, 13, 5)$	$2^{-28.27}$	Section 5.5
SIMON96	29	$(11, 14, 4)$	$2^{-46.67}$	[LZ26]
	30	$(9, 17, 4)$	$2^{-45.83}$	Section 5.5
SIMON128	37	$(13, 19, 5)$	$2^{-60.56}$	[LZ26]
	39	$(18, 21, 0)$	$2^{-61.90}$	Section 5.5
SIMECK48	21	$(0, 16, 5)$	$2^{-22.03}$	[LZ26]
	22	$(3, 15, 4)$	$2^{-23.94}$	Section 5.4
SIMECK64	26	$(0, 18, 8)$	$2^{-30.16}$	[LZ26]
	28	$(2, 20, 6)$	$2^{-29.93}$	Section 5.5

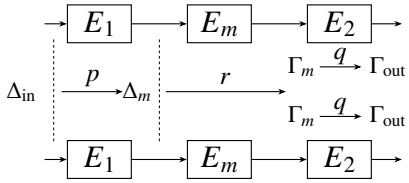

Figure 1: Three-part differential–linear distinguisher.

The data complexity of the corresponding statistical distinguisher scales as $|\operatorname{Cor}|^{-2} \approx p^{-2}|r|^{-2}q^{-4}$, up to a constant factor determined by the statistical test and the desired error probabilities.

In automated DL search, the differential prefix and linear suffix are commonly evaluated using their best single trails. Contributions from multiple trails sharing the same endpoints can also be incorporated to refine the correlation estimates.

2.2 Differential–Linear Correlation and Walsh Representation

Fix a middle input difference $\Delta_m \in \mathbb{F}_2^n$. For uniformly distributed $x \in \mathbb{F}_2^n$, let

$$\delta_x = E_m(x) \oplus E_m(x \oplus \Delta_m)$$

denote the corresponding output difference. For an output mask $\Gamma_m \in \mathbb{F}_2^n$, the middle DL correlation is defined as

$$r(\Delta_m, \Gamma_m) = \mathbb{E}_x \left[(-1)^{\langle \Gamma_m, \delta_x \rangle} \right]. \tag{1}$$

Let $P_{\Delta_m}(\delta) = \Pr_x[\delta_x = \delta]$ denote the output-difference distribution. Grouping Equation (1) by the output difference gives

$$r(\Delta_m, \Gamma_m) = \sum_{\delta \in \mathbb{F}_2^n} P_{\Delta_m}(\delta)(-1)^{\langle \Gamma_m, \delta \rangle}.$$

Define the Walsh–Hadamard matrix H_n by

$$H_n[\Gamma, \delta] = (-1)^{\langle \Gamma, \delta \rangle}, \quad \Gamma, \delta \in \mathbb{F}_2^n.$$

Writing

$$\mathbf{P}_{\Delta_m} = \big(P_{\Delta_m}(\delta)\big)_{\delta\in\mathbb{F}_2^n}, \quad \mathbf{r}_{\Delta_m} = \big(r(\Delta_m,\Gamma)\big)_{\Gamma\in\mathbb{F}_2^n},$$

we obtain

$$\mathbf{r}_{\Delta_m} = H_n\mathbf{P}_{\Delta_m}.$$

Thus, for a fixed input difference, the DL correlations of all output masks form the Walsh spectrum of a common output-difference distribution.

This representation is consistent with the classical relationship between the difference distribution table (DDT) and the differential–linear connectivity table (DLCT). Under the normalization of Bar-On et al. [BDKW19], twice a DLCT row equals the Walsh transform of the corresponding DDT row. We apply this relationship to multi-round middle components using empirically estimated output-difference distributions.

2.3 Empirical Differential–Linear Correlations

For a fixed Δ_m, let $x_1,\dots,x_N$ be independently and uniformly sampled inputs, and set $\delta_i = \delta_{x_i}$. The empirical DL correlation for an output mask Γ_m is

$$\widehat{r}(\Delta_m,\Gamma_m) = \frac{1}{N}\sum_{i=1}^{N}(-1)^{\langle\Gamma_m,\delta_i\rangle}. \tag{2}$$

Its variance is

$$\mathrm{Var}[\widehat{r}(\Delta_m,\Gamma_m)] = \frac{1-r(\Delta_m,\Gamma_m)^2}{N}.$$

Hence, when the true correlation is small, the sampling standard deviation is approximately $N^{-1/2}$.

For a fixed input difference, all output masks are evaluated using the same sampled output differences. Define the empirical histogram

$$C_{\Delta_m}(\delta) = |\{i : \delta_i = \delta\}|.$$

Then

$$\widehat{r}(\Delta_m,\Gamma_m) = \frac{1}{N}\sum_{\delta\in\mathbb{F}_2^n} C_{\Delta_m}(\delta)(-1)^{\langle\Gamma_m,\delta\rangle}.$$

Writing

$$\mathbf{C}_{\Delta_m} = \big(C_{\Delta_m}(\delta)\big)_{\delta\in\mathbb{F}_2^n}, \quad \widehat{\mathbf{r}}_{\Delta_m} = \big(\widehat{r}(\Delta_m,\Gamma)\big)_{\Gamma\in\mathbb{F}_2^n},$$

we obtain

$$\widehat{\mathbf{r}}_{\Delta_m} = \frac{1}{N}H_n\mathbf{C}_{\Delta_m}.$$

This provides the basis for the spectral evaluation of middle DL correlations in Section 4.

Fast Walsh–Hadamard Transform. The fast Walsh–Hadamard transform (FWHT) evaluates all 2^n empirical correlations in $O(n2^n)$ arithmetic operations and $O(2^n)$ memory, without explicitly constructing the Walsh–Hadamard matrix [FA76]. Each resulting coefficient equals the corresponding empirical correlation obtained by direct evaluation on the same samples.

Key Sampling. For keyed primitives, the correlations defined above are understood under a fixed master key unless otherwise specified. For a fixed master key K, the middle DL correlation is

$$r_K(\Delta_m, \Gamma_m) = \mathbb{E}_x \left[(-1)^{\langle \Gamma_m, E_{m,K}(x) \oplus E_{m,K}(x \oplus \Delta_m) \rangle} \right].$$

We distinguish the signed key-average correlation r_{mix} from the average fixed-key correlation magnitude r_{abs}:

$$r_{\text{mix}}(\Delta_m, \Gamma_m) = \mathbb{E}_K[r_K(\Delta_m, \Gamma_m)], \quad r_{\text{abs}}(\Delta_m, \Gamma_m) = \mathbb{E}_K[|r_K(\Delta_m, \Gamma_m)|].$$

By the triangle inequality,

$$|r_{\text{mix}}(\Delta_m, \Gamma_m)| \leq r_{\text{abs}}(\Delta_m, \Gamma_m).$$

Equality holds when the fixed-key correlations have the same sign over the key distribution. Otherwise, signed averaging may cause cross-key cancellation.

In our framework, mixed-key sampling is used for candidate discovery, while retained candidates are independently evaluated under fixed master keys, as detailed in Section 4.5.

3 Motivation and Approach Overview

3.1 Hourglass(-Like) Structures

Previous studies have observed that many high-quality DL distinguishers for ARX ciphers exhibit hourglass or hourglass-like structures. In an hourglass structure, a differential trail over E_1 connects the external input difference Δ_{in} to a single-bit middle input difference Δ_m, while a linear trail over E_2 connects a single-bit middle output mask Γ_m to the external output mask Γ_{out}. In the hourglass-like structure described by Gong et al., Δ_m remains single-bit, whereas Γ_m consists of two adjacent active bits [GWH+25]. Prior work on AND–RX ciphers such as `SIMON` and `SIMECK` has also used low-weight middle input differences and middle output masks to construct effective DL distinguishers [HDE24].

Restricting the middle input difference Δ_m and output mask Γ_m to these sparse forms substantially reduces the number of candidate difference–mask pairs requiring empirical correlation evaluation. This provides useful structural guidance for DL distinguisher search. Fig. 2 illustrates the hourglass structure.

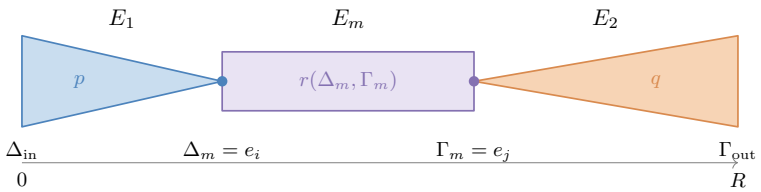

Figure 2: Illustration of the hourglass structure in a three-part DL distinguisher, where both Δ_m and Γ_m are single-bit.

However, the ability to repartition previously reported DL distinguishers into hourglass or hourglass-like structures does not imply that imposing these structural constraints is the most effective search strategy. Alternative round partitions may yield promising multi-bit middle output masks Γ_m, but exploring their possible combinations can substantially increase the computational cost of correlation evaluation.

3.2 Middle-Output Boundary Selection

We distinguish three cases according to the position of the evaluated output boundary of E_m relative to a round boundary at which the corresponding hourglass representation has a single-bit output mask. Figures 2, 3, and 4 illustrate these cases. The three partitions share the same differential prefix E_1, middle input difference $\Delta_m = e_i$, and total number of rounds. Only the allocation of rounds between E_m and E_2 differs. The reference boundary is not assumed to be known during the search.

Case 1: At a Single-Bit Output Mask. In this case, E_m ends at the single-bit output mask $\Gamma_m = e_j$, as shown in Fig. 2. Leurent restricted both Δ_m and Γ_m to single-bit values in his DL search for Chaskey [Leu16]. Gong et al. systematically repartitioned previously reported DL distinguishers, identified recurring hourglass and hourglass-like structures, and used these structures to guide their search [GWH+25].

Case 2: Before a Single-Bit Output Mask. In this case, E_m ends before the reference single-bit output-mask position. The evaluated middle output mask Γ_m has multiple active bits, while the reference boundary lies within the linear suffix E_2, as illustrated in Fig. 3.

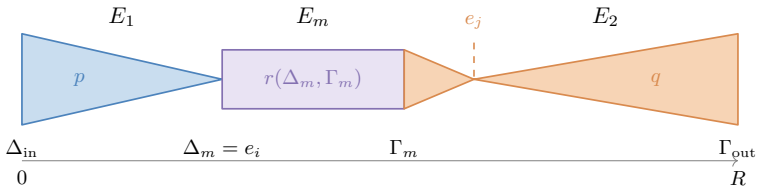

Figure 3: A middle component E_m ending before a reference single-bit output-mask position e_j. The evaluated output mask Γ_m has multiple active bits, while the reference boundary lies within E_2.

Chen et al. developed semi-unconstrained search to explore multi-bit middle output masks [CBY23], obtaining an 18-round DL distinguisher for `SPECK128`. Gong et al. subsequently showed that this distinguisher could be repartitioned into an hourglass structure [GWH+25]. Chen et al.'s $5 + 8 + 5$ partition uses $\Delta_m = [117]$ and $\Gamma_m = [5, 77]$, with an 8-round middle correlation of magnitude approximately $2^{-5.81}$. In the $5 + 9 + 4$ hourglass partition, the middle output mask becomes the single-bit mask $\Gamma_m = [5]$, while the 9-round middle correlation has magnitude approximately $2^{-11.19}$. Under the same statistical detection criterion, the ratio of the required sample sizes is approximately

$$\frac{N_9}{N_8} \approx \left(\frac{2^{-5.81}}{2^{-11.19}}\right)^2 = 2^{10.76}.$$

This comparison concerns the sampling cost of detecting the two middle correlations, rather than the data complexity of the complete distinguisher.

The shorter split yields a more readily detectable middle correlation in this example, at the potential cost of evaluating many multi-bit output masks during search. This trade-off motivates efficient multi-bit mask evaluation at alternative middle-output boundaries.

Case 3: After a Single-Bit Output Mask. In this case, E_m ends after the reference single-bit output-mask position, which lies within E_m. The evaluated middle output mask Γ_m may have multiple active bits, as illustrated in Fig. 4.

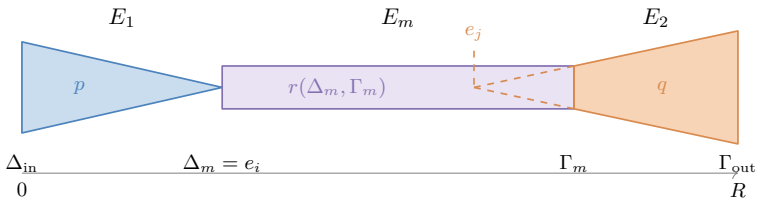

Figure 4: A middle component E_m extending beyond a reference single-bit output-mask position e_j. The additional rounds are included in the empirical evaluation of the middle correlation.

When the sampling budget permits, we can move rounds from E_2 into E_m and evaluate their contribution as part of the extended middle correlation, reducing reliance on a single linear trail. Extending E_m may weaken the middle correlation.

Motivation. The three cases illustrate that an hourglass or hourglass-like representation does not necessarily determine the most favorable boundary for empirical middle-correlation evaluation. A shorter E_m may yield a stronger correlation that is easier to detect, whereas a longer E_m can reduce reliance on single-trail approximations for the linear suffix when reliable empirical evaluation remains feasible. Both cases may involve multi-bit middle output masks Γ_m with many possible combinations of active positions.

Evaluating these masks individually can incur substantial computational cost, especially when weak correlations require large sample sizes. Under a fixed computational budget, the number of masks evaluated and the sample size used for each correlation estimate must therefore be balanced. This motivates efficient evaluation of empirical DL correlations for many output masks Γ_m under a fixed middle input difference Δ_m, while avoiding repeated correlation computations over shared samples.

3.3 Overview of the Spectral DL Search Framework

For the ARX and AND–RX ciphers considered here, we use single-bit middle input differences $\Delta_m = e_i$ and explore alternative middle-output boundaries. For a fixed E_m and Δ_m, we use Hamming-weight-one (HW1) correlations to select a low-dimensional set of output coordinates. We project the sampled output differences onto these coordinates and apply a fast Walsh–Hadamard transform (FWHT) to their histogram. A single transform yields the empirical DL correlations of all masks in the selected subspace, covering every combination of the selected coordinates. Each coefficient equals the estimate obtained by evaluating its mask directly on the same samples. When correlation computation dominates the cost, this joint evaluation can accommodate more masks or larger sample sizes within the same budget.

We integrate spectral evaluation into an automated three-part DL search. Preliminary searches identify promising round splits based on the detectability of middle correlations and the estimated quality of complete candidates. For each selected split, we combine the retained middle masks with differential-prefix and linear-suffix trails and rank the complete candidates by their three-part correlation estimates from mixed-key discovery. We then independently re-evaluate the middle correlations of the retained candidates under fixed master keys. Fig. 5 summarizes the procedure. Section 4 details its components and computational costs.

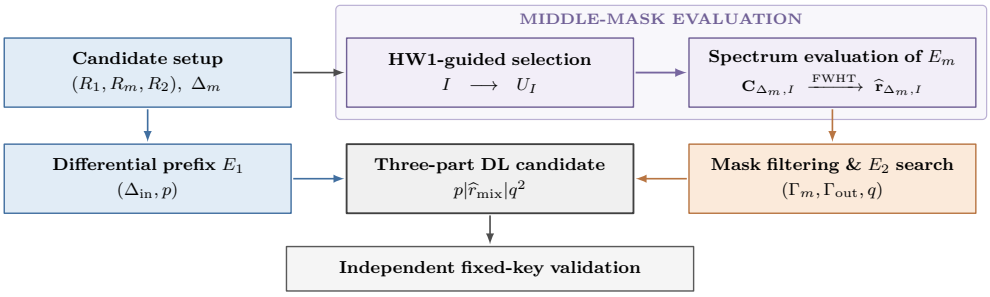

Figure 5: Overview of the proposed spectral DL search framework. HW1-guided subspace selection and FWHT-based middle-correlation evaluation are combined with outer-trail search, followed by independent fixed-key validation.

4 Spectral DL Search Framework

This section presents the proposed spectral DL search framework. We first describe the joint evaluation of middle correlations over low-dimensional output-mask subspaces and the HW1-guided strategy for selecting these subspaces. We then introduce middle-mask filtering and outer-trail completion, followed by round-split selection and the overall search procedure. Finally, we describe the sampling protocols for candidate discovery and independent fixed-key validation.

4.1 Low-Dimensional Spectral Evaluation

For a fixed input difference Δ_m, evaluating all 2^n output masks of E_m may be computationally infeasible for large state sizes. We therefore select k output coordinates, with $1 \le k \le n$, and use the Walsh spectrum of the projected output differences to compute jointly the empirical DL correlations of all masks supported on these coordinates.

Output-Mask Subspace and Output-Difference Projection. Let $I = \{i_0, \ldots, i_{k-1}\} \subseteq \{0, \ldots, n-1\}$ be the set of selected output coordinates, and let e_j denote the j-th standard basis vector. The corresponding mask subspace is

$$U_I = \operatorname{span}\{e_j : j \in I\} \subseteq \mathbb{F}_2^n, \qquad |U_I| = 2^k.$$

For $a \in \mathbb{F}_2^k$ and $\delta \in \mathbb{F}_2^n$, let $\Gamma_I(a) = \sum_{\ell=0}^{k-1} a_\ell e_{i_\ell}$ be the corresponding mask in U_I and $\pi_I(\delta) = (\delta_{i_0}, \ldots, \delta_{i_{k-1}})$ be the projected output difference. Then

$$\langle \Gamma_I(a), \delta \rangle = \langle a, \pi_I(\delta) \rangle.$$

Since this identity holds for each sample, projection preserves the empirical correlations of all masks in U_I.

Walsh Spectrum Evaluation. Let $\delta_1, \ldots, \delta_N$ be the output differences sampled from E_m under input difference Δ_m. In the discovery stage, we generate these samples using the mixed-key protocol in Section 4.5. For each sample, set $z_t = \pi_I(\delta_t)$ and let $C_{\Delta_m,I}(z) = |\{t : z_t = z\}|$ count the occurrences of each $z \in \mathbb{F}_2^k$. Grouping the samples by their projected values gives

$$\widehat{r}\big(\Delta_m, \Gamma_I(a)\big) = \frac{1}{N} \sum_{z \in \mathbb{F}_2^k} C_{\Delta_m,I}(z)(-1)^{\langle a,z \rangle}.$$

Collecting these correlations into a vector indexed by $a \in \mathbb{F}_2^k$ gives

$$\widehat{\mathbf{r}}_{\Delta_m,I} = \frac{1}{N} H_k \mathbf{C}_{\Delta_m,I}, \tag{3}$$

where H_k is the $2^k \times 2^k$ Walsh–Hadamard matrix and $\mathbf{C}_{\Delta_m,I}$ is the histogram vector. Thus, a single length-2^k FWHT computes the empirical correlations of all masks in U_I. When $k = n$, the same formula gives the full output-mask spectrum.

Computational Complexity. For a fixed Δ_m and coordinate set I, projecting N output differences and constructing the histogram require $O(Nk + 2^k)$ time, including histogram initialization. The subsequent FWHT requires $O(k2^k)$ time. Therefore, evaluating all 2^k masks in U_I has total complexity

$$T_{\text{spec}} = O(Nk + k2^k),$$

with $O(2^k)$ auxiliary memory when the histogram is constructed incrementally. When $k = n$, the projection can be omitted, giving a tighter time bound of $O(N + n2^n)$.

By comparison, evaluating the same 2^k masks individually over the same N samples requires $O(N2^k)$ correlation-computation operations. The FWHT therefore reduces the cost of evaluating multiple masks while yielding exactly the same empirical correlations for a given sample batch. These complexity bounds exclude sample generation, coordinate selection, and outer-trail search.

Runtime Comparison on `SPECK128/128`. We compare the proposed FWHT-based subspace evaluation with direct per-mask correlation evaluation for the nine-round middle component E_m of `SPECK128/128`, using $\Delta_m = [125]$. Both methods process $N = 2^{34}$ input pairs on four GPUs. Their measured runtimes include sample generation, evaluation of E_m, and the subsequent correlation computation. Samples are processed in batches, and neither method stores the complete set of output differences.

Excluding the separate coordinate-selection stage, the $k = 32$ FWHT procedure takes 168.66 seconds to project the output differences, construct the histogram, and evaluate all $2^{32} - 1$ nonzero masks in the selected subspace. Over three repetitions, direct streaming evaluation of randomly selected 128-bit masks takes 164.07 seconds for 11,264 masks and 174.37 seconds for 12,288 masks, giving an estimated crossover of approximately 11,720 masks. Thus, under the same runtime budget and sample-processing conditions, the FWHT evaluates over 2^{18} times as many masks.

4.2 HW1-Guided Mask Subspace Selection

For $k < n$, we follow the strong-unbalanced-bit strategy of Chen et al. [CBY23] and use the empirical correlations of Hamming-weight-one (HW1) masks to select output coordinates. For a given Δ_m, we take the indices of the k largest magnitudes:

$$I = \mathrm{TopK}_k \left(\left(|\widehat{r}(\Delta_m, e_i)| \right)_{i=0}^{n-1} \right), \tag{4}$$

where e_i is the single-bit mask at coordinate i and ties are resolved in a fixed coordinate order. The HW1 correlations determine only I. We then evaluate all 2^k masks in U_I using Equation (3).

For $k < n$, we use two independent batches of N input pairs each. The first batch is used to estimate all n HW1 correlations and select I, with $O(Nn)$ sample-level updates. In the second batch, we project each output difference onto I and accumulate the histogram for the FWHT. Both stages process output differences as they are generated, avoiding the need to store the full collection. The two stages use $2N$ input pairs in total, and each

spectral correlation is estimated from the second batch alone. When $k = n$, we evaluate the full spectrum directly from a single batch of N input pairs. Section 5.2 compares the selected subspaces under different dimension budgets.

4.3 Middle-Candidate Selection and Outer-Trail Completion

For a fixed three-part round split (R_1, R_m, R_2), we construct complete DL distinguisher candidates by combining empirical middle correlations with differential and linear trails over E_1 and E_2, respectively. The main search uses the best single trails for both outer components.

Input-Difference Selection and Prefix Completion. We consider all single-bit middle input differences $\Delta_m = e_i$, $0 \le i < n$. For each Δ_m, we use SAT-based search to find the best single differential trail over E_1 ending at Δ_m. This determines the external input difference Δ_{in} and the corresponding differential probability p.

Middle-Mask Filtering. For each Δ_m, we evaluate all masks in U_I with a single FWHT as described in Section 4.1. We then filter and rank the masks to limit the number of SAT-based linear-suffix searches.

1. **Correlation filtering.** We discard the zero mask and retain masks satisfying $|\widehat{r}(\Delta_m, \Gamma_m)| > \tau$, where τ is a threshold for excluding empirical correlations that are too weak relative to the sampling noise.
2. **Structure-dependent filtering.** The propagation properties of modular addition and bitwise AND motivate favoring low-weight masks for linear-suffix extension. For selected instances with $R_2 > 0$, we therefore impose a Hamming-weight bound $h_{\max}$ on Γ_m to reduce the subsequent linear-suffix search space.
3. **Ranking and candidate selection.** Among the remaining masks, we first retain up to B_{HW} masks with the largest $|\widehat{r}(\Delta_m, \Gamma_m)|$ at each permitted Hamming weight. We then select up to B_{cand} masks with the largest correlation magnitudes from the resulting pool.

The values of τ, $h_{\max}$, B_{HW}, and B_{cand} are specified in Section 5.1.

Linear-Suffix Search. For each retained middle output mask Γ_m, we use SAT-based search to find the best single linear trail over E_2 starting from Γ_m. Following the three-part single-trail approximation in Section 2.1, we estimate the correlation magnitude of a complete DL candidate as

$$|\widehat{\text{Cor}}| := p|\widehat{r}(\Delta_m, \Gamma_m)|q^2,$$

where q is the correlation of the selected linear trail. Here, $\widehat{r}(\Delta_m, \Gamma_m)$ is the mixed-key empirical correlation estimated during candidate discovery, as detailed in Section 4.5.

For each fixed round split, we retain up to B_{keep} complete candidates with the largest estimated correlation magnitudes. Once B_{keep} candidates have been retained, let Cor_{cut} denote the smallest estimated correlation magnitude among them. This cutoff is updated whenever a better candidate is found.

For a new candidate with known p and $|\widehat{r}(\Delta_m, \Gamma_m)|$, improving the retained set requires

$$p|\widehat{r}(\Delta_m, \Gamma_m)|q^2 > \text{Cor}_{\text{cut}}\,.$$

This gives a candidate-specific threshold

$$q_{\min} := \sqrt{\frac{\text{Cor}_{\text{cut}}}{p|\widehat{r}(\Delta_m, \Gamma_m)|}}, \qquad |q| > q_{\min}.$$

We use this threshold to restrict the SAT-based linear-suffix search. If no trail satisfies $|q| > q_{\min}$, the candidate cannot improve the current retained set and is discarded. Otherwise, the best qualifying trail provides the external output mask Γ_{out} and correlation q for updating the retained candidates.

Before B_{keep} complete candidates are available, no such threshold is imposed. When $R_2 = 0$, no linear-suffix search is needed; we set $\Gamma_{\text{out}} = \Gamma_m$ and $q = 1$. The value of B_{keep} is specified in Section 5.1.

The retained middle candidates can also be combined with aggregated contributions from the outer components without repeating the middle-spectrum evaluation. We investigate this extension in Section 5.5.

4.4 Round-Split Selection and Overall Search

Having described the search for a fixed three-part round split, we now present the preliminary procedure for selecting promising round splits and the subsequent search over all single-bit middle input differences.

Preliminary Round-Split Search. Rather than exploring all middle input differences for every possible round split, we first identify promising configurations (R_1, R_m, R_2) through a preliminary search. The procedure consists of three steps.

1. **Middle-length probing.** Starting from a prescribed middle length, we increase R_m one round at a time. At each length, we evaluate the spectra for $\mathcal{D}_{\text{probe}} = \{e_0, e_{n/4}, e_{n/2}, e_{3n/4}\}$ using Sections 4.1 and 4.2. We apply the same middle-mask filters and quotas as in the main search.

   For ciphers whose round function depends on the absolute round index, the complete distinguisher starts at round 0. In a split (R_1, R_m, R_2), E_m starts at R_1 and E_2 starts at $R_1 + R_m$. We therefore evaluate the middle spectra separately at each admissible start. For a round-position-independent middle model, we evaluate each (Δ_m, R_m) once at a standard start and reuse its spectrum across compatible splits.

   We stop at the first length with no retained candidate at any tested start. Let $R_m^{\max}$ be the largest tested length with a retained candidate. Since a shorter middle component can yield a stronger complete candidate, we take $\mathcal{M} = \{R_m^{\max}, R_m^{\max} - 1\}$.

2. **Outer-part extension.** For each $R_m \in \mathcal{M}$, we group middle candidates that can share an outer-trail search. Let $r_\star$ denote the largest empirical middle-correlation magnitude in the group, and let p and q denote the best single-trail factors for the prefix and suffix, respectively. We use $2^{-n/2}$ as the reference correlation magnitude, corresponding to a sample complexity on the scale of 2^n under Section 2.1.

   Since $p \le 1$ and $q^2 \le 1$, the group score is bounded above by $pr_\star$ and by $r_\star q^2$. In a nested extension with a fixed middle component, we stop extending the corresponding outer component when

   $$pr_\star < 2^{-n/2} \quad \text{for } E_1, \qquad r_\star q^2 < 2^{-n/2} \quad \text{for } E_2.$$

   For a suffix with a fixed start, removing the last round of a longer trail gives a valid shorter trail with at least the same correlation magnitude. For a round-position-independent prefix, removing the first round and renumbering gives the analogous result for the trail probability. For a round-position-dependent prefix, we solve each admissible R_1 independently at its actual round positions. We record timeouts and searches with unproven optimality as incomplete.

3. **Round-split construction.** For each retained middle pair (Δ_m, Γ_m), we combine compatible prefix and suffix lengths, matching the endpoints and round positions. We score each complete candidate individually by $p|\widehat{r}(\Delta_m, \Gamma_m)|q^2$ and include a split in $\mathcal{R}$ if at least one of its candidates scores above $2^{-n/2}$.

Overall Search Procedure. For each round split in $\mathcal{R}$, we apply the procedure in Section 4.3 to all n single-bit middle input differences $\Delta_m = e_i$, $0 \leq i < n$. We rank the resulting DL candidates first by total round count and then by estimated correlation magnitude. Appendix A provides pseudocode for the main search with $\mathcal{R}$ as input.

4.5 Discovery Sampling and Independent Validation

We use different sampling protocols for candidate discovery and independent validation. The former estimates signed mixed-key correlations for candidate selection, while the latter evaluates the correlation magnitudes of retained middle candidates under fixed master keys.

- **Discovery sampling.** We evaluate each input pair under an independently sampled master key and use the signed mixed-key estimate $\widehat{r}_{\mathrm{mix}}(\Delta_m, \Gamma_m)$ for candidate selection and ranking (see Section 2.3). For $k < n$, HW1-guided coordinate selection and spectral evaluation use independent batches under this protocol.

- **Independent fixed-key validation.** For each retained pair (Δ_m, Γ_m), we sample T independent master keys and evaluate N_{val} fresh input pairs per key, independently of the discovery samples. Let $\widehat{r}_{K_j}(\Delta_m, \Gamma_m)$ denote the signed empirical middle correlation under key K_j. We report the mean of the fixed-key correlation magnitudes

$$\widehat{r}_{\mathrm{abs}}(\Delta_m, \Gamma_m) = \frac{1}{T} \sum_{j=1}^{T} \left| \widehat{r}_{K_j}(\Delta_m, \Gamma_m) \right| .$$

  With single-trail outer factors, the three-part correlation magnitude is estimated as $p\,\widehat{r}_{\mathrm{abs}}(\Delta_m, \Gamma_m)q^2$.

Section 5.1 specifies T and N_{val}.

5 Experimental Evaluation

This section evaluates the proposed spectral DL search framework. We assess HW1-guided subspace selection against full-state spectra on 32-bit cipher instances and compare mixed-key discovery estimates with independent fixed-key measurements. We then report results for ARX and AND–RX ciphers, evaluate outer-path aggregation, and examine how the middle-output boundary relates to hourglass and hourglass-like structures. Supplementary experiments on SPN structures are provided in the appendix.

5.1 Experimental Setup

All experiments were conducted on a server with two Intel Xeon E5-2680 v4 processors, 503 GiB of memory, and seven NVIDIA GeForce RTX 3080 Ti GPUs with 12 GiB of memory each. Distributed spectrum evaluations at $k = 32$ typically used four GPUs, subject to availability. The software environment included Python 3.9.21, NumPy 2.0.2, PyTorch 2.7.1 (CUDA 11.8), Triton 3.3.1, and CryptoMiniSat 5.8.0. We primarily used the Open Cryptanalysis Platform (OCP) [Ope] with CryptoMiniSat for outer-trail search,

reusing existing models where available and implementing the remaining models through the OCP modeling interface.

Table 3 lists the common search and validation parameters. Each discovery batch contains N input pairs. We use two independent batches when $k < n$ and one batch when $k = n$, as described in Section 4.2.

Table 3: Common search and validation parameters.

Parameter	Setting
Maximum subspace dimension k	32
Discovery samples per batch N	2^{30} (32-bit), 2^{34} (otherwise)
Correlation threshold τ	$2^{-12.29}$ ($N = 2^{30}$), $2^{-14.29}$ ($N = 2^{34}$)
Per-HW mask budget B_{HW}	1000
Middle-mask budget B_{cand}	2000
HW bound $h_{\max}$	6 (selected instances with $R_2 > 0$)
Candidates per split B_{keep}	10
Validation keys T	50
Validation samples per key N_{val}	2^{30} (32-bit), 2^{34} (otherwise)

We choose the screening threshold τ such that $\tau\sqrt{N} \approx 6.55$. For a nonzero mask with zero true correlation, the normal approximation gives an exceedance probability of $2\Phi(-6.55) \approx 5.75 \times 10^{-11}$ for the magnitude of its empirical correlation, where Φ is the standard normal cumulative distribution function. If all nonzero masks have zero true correlation, the expected number of threshold exceedances in a full 2^{32}-point spectrum is approximately 0.25.

5.2 Evaluation of HW1-Guided Mask Subspace Selection

We evaluate the effectiveness of HW1-guided low-dimensional mask subspaces by comparing the quality of DL candidates obtained under different subspace dimensions. We focus on 32-bit cipher instances, for which full-state spectral evaluation over all 2^{32} output masks is computationally feasible. In addition to comparing the resulting three-part correlation estimates, we examine whether the selected subspaces contain the middle output masks found by full-state search.

We apply the search framework to `SPECK32`, `SIMON32`, and `SIMECK32`, with $k \in \{8, 16, 24, 32\}$. For $k < 32$, the subspace is constructed from the k output coordinates with the largest HW1 empirical correlation magnitudes. The configuration $k = 32$ evaluates the full output-mask space and serves as the reference. All configurations consider the 32 single-bit middle input differences, with $N = 2^{30}$, $\tau = 2^{-12.29}$, and the remaining search parameters held fixed. For $k < 32$, coordinate selection and spectral evaluation use independent batches of N input pairs, as described in Section 4.2. Table 4 reports representative discovery-stage three-part correlation estimates.

The results in Table 4 illustrate three representative outcomes.

- **Retention of the reference mask.** For `SPECK32` with the $2 + 5 + 4$ split, the coordinate of the reference mask [3] ranks third by HW1 correlation magnitude. All tested subspaces contain this mask and yield the same reported estimate of $2^{-15.83}$. For `SIMECK32` with the $3+11+0$ split, the support coordinates of the reference mask $[0, 1, 15, 16]$ have HW1 ranks $1, 4, 6, 11$. The smallest tested subspace containing the

Table 4: Discovery-stage three-part correlation estimates under different output-mask subspace dimensions.

Cipher	Split	$k = 8$	$k = 16$	$k = 24$	$k = 32$
SPECK32	$2 + 9 + 0$	–	–	–	$2^{-15.84}$
SPECK32	$2 + 5 + 4$	$2^{-15.83}$	$2^{-15.83}$	$2^{-15.83}$	$2^{-15.83}$
SIMON32	$3 + 11 + 0$	$2^{-17.04}$	$2^{-14.18}$	$2^{-14.03}$	$2^{-14.02}$
SIMECK32	$3 + 11 + 0$	$2^{-13.90}$	$2^{-13.87}$	$2^{-13.87}$	$2^{-13.87}$

Entries report $p|\widehat{r}_{\text{mix}}|q^2$; "–" indicates that no complete DL candidate was obtained for the corresponding configuration.

full support has $k = 16$, and the estimates for $k = 16, 24, 32$ all equal $2^{-13.87}$ at the reported precision.

- **Recovery of a comparable alternative.** For SIMON32 with the $3 + 11 + 0$ split, the support coordinates of the reference mask $[11, 12, 17, 29]$ have HW1 ranks $18, 20, 23, 31$, placing the mask outside the $k = 24$ subspace. At $k = 24$, the search finds an alternative candidate with an estimated correlation magnitude of $2^{-14.03}$, close to the full-state estimate of $2^{-14.02}$.
- **Failure to recover a candidate.** For SPECK32 with the $2+9+0$ split, the reference mask includes coordinate 29, which ranks 31st by HW1 correlation magnitude. The mask lies outside all tested subspaces with $k < 32$, and the search produces no complete candidate in these configurations.

The two SPECK32 configurations represent different round partitions of related 11-round distinguishers. Although the restricted subspaces fail to recover a candidate for the $2 + 9 + 0$ split, the $2 + 5 + 4$ split yields a comparable 11-round result even with $k = 8$. This observation highlights the importance of evaluating search outcomes across different round partitions, rather than focusing solely on the recovery of a particular middle output mask.

These controlled experiments support HW1-guided subspace search when full-state spectral evaluation is infeasible. A selected subspace can retain the reference mask or yield an alternative candidate with a comparable three-part correlation estimate.

5.3 Mixed-Key Discovery versus Fixed-Key Validation

We compare mixed-key discovery estimates with independent fixed-key measurements for a representative DL candidate. As discussed in Section 2.3, mixed-key sampling estimates $r_{\text{mix}} = \mathbb{E}_K[r_K]$, whereas fixed-key validation considers $\mathbb{E}_K[|r_K|]$. Differences in correlation signs across keys may cause cancellation in the mixed-key correlation, potentially obscuring candidates with substantial fixed-key correlations.

We consider a middle difference–mask pair discovered for CHAM-128/128 with the $10 + 39 + 11$ round split:

$$\Delta_m = (\texttt{00000000}, \texttt{00000000}, \texttt{00000000}, \texttt{00004000}),$$
$$\Gamma_m = (\texttt{00000000}, \texttt{00000000}, \texttt{00000018}, \texttt{00000000}).$$

During discovery, we use $N = 2^{34}$ random-key-per-pair samples. For independent fixed-key validation, we evaluate the same middle candidate under $T = 50$ master keys, using 2^{34} fresh input pairs per key. The resulting correlation estimates are

$$|\widehat{r}_{\text{mix}}| = 5.67 \times 10^{-5} \approx 2^{-14.11}, \quad \widehat{r}_{\text{abs}} = 5.22 \times 10^{-5} \approx 2^{-14.22}.$$

Among the 50 fixed keys, 47 yield negative empirical correlations and three yield positive ones; all three positive estimates are close to zero. The mixed-key discovery magnitude is approximately $2^{0.12}$ times the mean absolute fixed-key correlation. Most other retained candidates with per-key records also show consistent signs. The modest excess of the mixed-key estimate may reflect fluctuations in the independent finite samples and variation arising from the sampled validation keys. In subsequent experiments, we report the two estimates separately and use independent fixed-key validation to evaluate the retained differential–linear candidates.

5.4 Single-Trail Differential–Linear Results

We report DL results for SIMON, SIMECK, SPECK, and CHAM. For each candidate, p and q are the probability and correlation of the best single trails found for E_1 and E_2, respectively. The tables report the discovery magnitude $|\widehat{r}_{\mathrm{mix}}|$ and the independent fixed-key estimate $\widehat{r}_{\mathrm{abs}}$. The resulting three-part correlation magnitude is estimated as $p\widehat{r}_{\mathrm{abs}}q^2$. Appendix E lists the external endpoints and single-trail factors.

Results for the SIMON Family. SIMON is a lightweight Feistel block cipher family with nonlinear function $f(x) = ((x \lll 1) \,\&\, (x \lll 8)) \oplus (x \lll 2)$. Table 5 reports the results for the five tested variants.

Table 5: Summary of three-part differential–linear search results for the SIMON family.

Cipher	Split	Δ_m	Γ_m	p	q^2	$\lvert\widehat{r}_{\mathrm{mix}}\rvert$	$\widehat{r}_{\mathrm{abs}}$	$p\widehat{r}_{\mathrm{abs}}q^2$
SIMON32	$3+11+0$	[26]	$[11,12,17,29]$	2^{-4}	1	$2^{-10.02}$	$2^{-10.11}$	$2^{-14.11}$
SIMON48	$2+12+4$	[23]	[4]	2^{-4}	2^{-10}	$2^{-8.66}$	$2^{-8.69}$	$2^{-22.69}$
SIMON64	$5+14+3$	[32]	$[6,40]$	2^{-10}	2^{-8}	$2^{-11.85}$	$2^{-11.83}$	$2^{-29.83}$
SIMON96	$4+17+8$	[0]	$[2,48]$	2^{-10}	2^{-22}	$2^{-13.95}$	$2^{-14.16}$	$2^{-46.16}$
SIMON128	$8+22+6$	[29]	[30]	2^{-28}	2^{-20}	$2^{-13.53}$	$2^{-13.91}$	$2^{-61.91}$

Results for the SIMECK Family. SIMECK is a lightweight Feistel block cipher family with nonlinear function $f(x) = (x \,\&\, (x \lll 5)) \oplus (x \lll 1)$. Table 6 reports the results for the three tested variants.

Table 6: Summary of three-part differential–linear search results for the SIMECK family.

Cipher	Split	Δ_m	Γ_m	p	q^2	$\lvert\widehat{r}_{\mathrm{mix}}\rvert$	$\widehat{r}_{\mathrm{abs}}$	$p\widehat{r}_{\mathrm{abs}}q^2$
SIMECK32	$2+12+0$	[0]	$[13,14,15,30]$	2^{-4}	1	$2^{-9.87}$	$2^{-9.93}$	$2^{-13.93}$
SIMECK48	$3+15+4$	[44]	$[19,42]$	2^{-4}	2^{-6}	$2^{-13.71}$	$2^{-13.94}$	$2^{-23.94}$
SIMECK64	$4+20+4$	[29]	[28]	2^{-10}	2^{-10}	$2^{-11.82}$	$2^{-11.82}$	$2^{-31.82}$

Results for the SPECK Family. SPECK is an ARX block cipher family with a two-word state (x_i, y_i). Given the round subkey k_i, the state update is

$$x_{i+1} = \big((x_i \ggg \alpha) + y_i\big) \oplus k_i, \qquad y_{i+1} = (y_i \lll \beta) \oplus x_{i+1}.$$

For w-bit words, addition is modulo 2^w. The variants in Table 7 use $(\alpha, \beta) = (8, 3)$.

Table 7: Summary of three-part differential–linear search results for the SPECK family.

Cipher	Split	Δ_m	Γ_m	p	q^2	$\lvert\widehat{r}_{\mathrm{mix}}\rvert$	$\widehat{r}_{\mathrm{abs}}$	$p\widehat{r}_{\mathrm{abs}}q^2$
SPECK48	$2+8+2$	[39]	$[8,32]$	2^{-5}	2^{-6}	$2^{-10.90}$	$2^{-10.98}$	$2^{-21.98}$
SPECK64	$2+9+3$	[39]	[32]	2^{-6}	2^{-12}	$2^{-13.57}$	$2^{-13.51}$	$2^{-31.51}$
SPECK96	$3+9+4$	[79]	$[5,53]$	2^{-11}	2^{-24}	$2^{-12.48}$	$2^{-12.46}$	$2^{-47.46}$
SPECK128	$5+9+4$	[125]	[77]	2^{-30}	2^{-18}	$2^{-6.53}$	$2^{-6.47}$	$2^{-54.47}$

Results for the CHAM Family. CHAM is a lightweight ARX block cipher family with a four-branch generalized Feistel structure over w-bit words. In round i, the first three words follow $X_{i+1}[j] = X_i[j+1]$ for $0 \le j \le 2$, and the last word is

$$X_{i+1}[3] = \mathrm{ROL}_{b_i}\big((X_i[0] \oplus i) \boxplus \big(\mathrm{ROL}_{a_i}(X_i[1]) \oplus RK_{i \bmod 2m}\big)\big).$$

Here $(a_i, b_i) = (1, 8)$ for even i and $(8, 1)$ for odd i, $\boxplus$ denotes addition modulo 2^w, and m is the number of words in the master key. The round constant and rotations depend on the absolute round index, which we preserve in middle-spectrum evaluation and outer-trail completion. Table 8 reports the results for CHAM64 and CHAM128.

Table 8: Summary of three-part differential–linear search results for the CHAM family.

Cipher	Split	Δ_m	Γ_m	p	q^2	$\lvert\widehat{r}_{\mathrm{mix}}\rvert$	$\widehat{r}_{\mathrm{abs}}$	$p\widehat{r}_{\mathrm{abs}}q^2$
CHAM64	$7+33+7$	[11]	[49]	2^{-13}	2^{-6}	$2^{-10.71}$	$2^{-10.91}$	$2^{-29.91}$
CHAM128	$10+39+11$	[110]	$[67,68]$	2^{-21}	2^{-28}	$2^{-14.11}$	$2^{-14.22}$	$2^{-63.22}$

5.5 Outer-Path Aggregation Results

We refine selected SIMON and SIMECK candidates by aggregating differential-prefix contributions and empirically evaluating nonempty linear suffixes under fixed keys. For the prefix, we combine the dynamic-window transition DP of Niu et al. [NLZW26] with SAT-based trail enumeration. The aggregated probability $P_{E_1}^{\mathrm{agg}}$ sums the DP contribution within the window and the probability mass of enumerated trails outside it. Appendix C illustrates the implementation for SIMON128.

For outer-path aggregation, the single-trail stopping and candidate-retention rules of Section 4.4 are not applicable, because the best single-trail contribution does not upper-bound the corresponding aggregate. We therefore retain the middle lengths identified by the general search, extend the total number of rounds assigned to E_1 and E_2, and separately examine the resulting round splits to identify candidates suitable for outer-path aggregation.

For each candidate, we evaluate E_m and E_2 under the same set of $T = 50$ fixed master keys. For each key, we derive the round subkeys for the two components at zero-based offsets R_1 and $R_1 + R_m$, respectively. We use 2^{34} input pairs to estimate the middle correlation $\widehat{r}_{K_j}(\Delta_m, \Gamma_m)$. For $R_2 > 0$, we use 2^{34} independent inputs to estimate the suffix correlation $\widehat{q}_{K_j}$. For $R_2 = 0$, we set $\widehat{q}_{K_j} = 1$.

We estimate the three-part correlation magnitude by averaging the products paired by master key:

$$\left|\widehat{\mathrm{Cor}}_{\mathrm{agg}}\right| = \frac{P_{E_1}^{\mathrm{agg}}}{T} \sum_{j=1}^{T} \left|\widehat{r}_{K_j}(\Delta_m, \Gamma_m)\right| \widehat{q}_{K_j}^2.$$

Table 9 also reports the marginal means $\widehat{r}_{\mathrm{abs}}$ and $\widehat{Q}^{\mathrm{emp}}_{E_2}$ of the middle correlation magnitudes and squared suffix correlations, respectively.

Table 9: Three-part correlation estimates using aggregated differential prefixes and paired fixed-key middle and suffix measurements for the `SIMON` and `SIMECK` families.

Cipher	Split	Δ_m	Γ_m	$P^{\mathrm{agg}}_{E_1}$	$\widehat{Q}^{\mathrm{emp}}_{E_2}$	$\lvert\widehat{r}_{\mathrm{mix}}\rvert$	$\widehat{r}_{\mathrm{abs}}$	$\lvert\widehat{\mathrm{Cor}}_{\mathrm{agg}}\rvert$
`SIMON48`	$2+12+4$	[0]	[5]	$2^{-4.00}$	$2^{-9.21}$	$2^{-8.68}$	$2^{-8.73}$	$2^{-21.91}$
`SIMON64`	$4+13+5$	[0]	[40]	$2^{-9.30}$	$2^{-9.08}$	$2^{-9.86}$	$2^{-9.86}$	$2^{-28.27}$
`SIMON96`	$9+17+4$	[48]	[4]	$2^{-24.17}$	$2^{-9.14}$	$2^{-12.49}$	$2^{-12.50}$	$2^{-45.83}$
`SIMON128`	$18+21+0$	[29]	[27, 31, 93]	$2^{-51.62}$	1	$2^{-10.27}$	$2^{-10.28}$	$2^{-61.90}$
`SIMECK64`	$2+20+6$	[5]	[4, 35]	$2^{-4.00}$	$2^{-11.39}$	$2^{-13.98}$	$2^{-14.52}$	$2^{-29.93}$

The largest gain is obtained for the 39-round `SIMON128` candidate with the $18+21+0$ split. Prefix aggregation increases its differential contribution from 2^{-64} to $2^{-51.62}$. Since $R_2 = 0$, this raises the three-part correlation estimate by a factor of $2^{12.38}$ to $2^{-61.90}$. Appendix D reports the prefix gain, marginal suffix gain, and total gain for each candidate relative to its single-trail evaluation under the same round split and boundary values.

5.6 Hourglass Structures and Middle-Output Boundary Selection

We examine how the evaluated middle-output boundaries relate to the hourglass and hourglass-like structures described by Gong et al. [GWH+25]. Following Section 3.2, we distinguish three cases according to whether the evaluated boundary coincides with, precedes, or follows a reference single-bit output-mask position.

Case 1: At a Low-Weight Output Mask. The evaluated middle output mask is single-bit for `SIMON48`, `SPECK64`, `SPECK128`, and `CHAM-64/128`, as well as for `SIMON64` $(4+13+5)$, `SIMON96` $(9+17+4)$, `SIMON128` $(8+22+6)$, and `SIMECK64` $(4+20+4)$. The `CHAM-128/128` candidate with the $10+39+11$ split has the adjacent two-bit output mask $[67, 68]$, corresponding to an hourglass-like structure.

For Cases 2 and 3, Table 10 gives the single-bit reference configurations obtained by reallocating rounds between E_m and E_2.

Case 2: Before a Single-Bit Output Mask. For `SIMECK48`, the single-bit reference configuration uses the $3+16+3$ split with $\Gamma_m = [43]$. An independent evaluation yields a middle correlation magnitude of approximately $2^{-14.69}$, below the screening threshold $\tau = 2^{-14.29}$. This correlation would therefore fail the middle-mask screening criterion. In contrast, the $3+15+4$ split with the two-bit output mask $[19, 42]$ passes screening and yields a retained DL candidate. This example illustrates that evaluating a shorter middle component can reveal candidates that would be missed when screening only at the corresponding single-bit mask position.

Case 3: After a Single-Bit Output Mask. For `SIMON32`, the $3+8+3$ reference configuration combines an empirically evaluated 8-round middle correlation with a subsequent single linear-trail contribution, giving an estimate of approximately $2^{-10.95}$. By extending the middle component to 11 rounds, the $3+11+0$ configuration directly evaluates the additional propagation and obtains an average absolute fixed-key correlation of approximately $2^{-10.11}$. This comparison illustrates how extending E_m can reduce reliance on single-trail approximations and yield a stronger empirical correlation estimate when the longer middle component remains reliably evaluable.

Table 10: Comparison of evaluated middle-output configurations and their corresponding single-bit reference configurations.

Case	Cipher	Δ_m	Evaluated boundary		Single-bit boundary	
			Split	Γ_m	Split	Γ_m
2	SIMON96	[0]	$4+17+8$	[2, 48]	$4+18+7$	[50]
	SIMECK48	[44]	$3+15+4$	[19, 42]	$3+16+3$	[43]
	SIMECK64	[5]	$2+20+6$	[4, 35]	$2+21+5$	[36]
3	SIMON32	[26]	$3+11+0$	[11, 12, 17, 29]	$3+8+3$	[1]
	SIMON64	[32]	$5+14+3$	[6, 40]	$5+13+4$	[8]
	SIMON128	[29]	$18+21+0$	[27, 31, 93]	$18+19+2$	[31]
	SIMECK32	[0]	$2+12+0$	[13, 14, 15, 30]	$2+10+2$	[15]
	SPECK48	[39]	$2+8+2$	[8, 32]	$2+7+3$	[5]
	SPECK96	[79]	$3+9+4$	[5, 53]	$3+8+5$	[2]

Discussion. These results support the motivation for flexible middle-output boundary selection. Although the reported distinguishers admit hourglass or hourglass-like representations, the corresponding low-weight mask positions need not provide the most favorable boundaries for empirical evaluation. A shorter middle component may preserve more readily detectable correlations, while a longer one can incorporate additional propagation into the empirical evaluation. Together, these observations illustrate the value of efficiently evaluating multi-bit middle output masks at different round partitions.

6 Conclusion

We presented a spectral search framework for differential–linear cryptanalysis that efficiently evaluates multiple middle output masks using the Walsh–Hadamard transform over HW1-guided low-dimensional subspaces. By combining spectral evaluation with automated outer-trail search, the framework enables flexible exploration of three-part DL distinguishers. Experiments on several ARX and AND–RX ciphers yield improved distinguishers in terms of round coverage or estimated correlation. Our results also show that, even when a distinguisher admits an hourglass or hourglass-like representation, the corresponding sparse-mask position need not be the most favorable boundary for empirical middle-correlation evaluation. These findings highlight the value of efficient multi-bit mask evaluation and flexible middle boundary selection in automated DL cryptanalysis.

References

[BDKW19] Achiya Bar-On, Orr Dunkelman, Nathan Keller, and Ariel Weizman. DLCT: A new tool for differential-linear cryptanalysis. In *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 313–342. Springer, 2019. doi:10.1007/978-3-030-17653-2_11.

[BGG+23] Emanuele Bellini, David Gérault, Juan Grados, Rusydi H. Makarim, and Thomas Peyrin. Fully automated differential-linear attacks against ARX ciphers. In *Topics in Cryptology – CT-RSA 2023*, volume 13871 of *Lec-*

ture Notes in Computer Science, pages 252–276. Springer, 2023. doi:10.1007/978-3-031-30872-7_10.

[CBY23] Yi Chen, Zhenzhen Bao, and Hongbo Yu. Differential-linear approximation semi-unconstrained searching and partition tree: Application to LEA and Speck. In *Advances in Cryptology – ASIACRYPT 2023, Part III*, volume 14440 of *Lecture Notes in Computer Science*, pages 223–255. Springer, 2023. doi:10.1007/978-981-99-8727-6_8.

[CXZQ24] Siwei Chen, Zejun Xiang, Xiangyong Zeng, and Guangxue Qin. Enhancing the MILP/MIQCP-based automatic search for differential-linear distinguishers of Simon-like ciphers. arXiv preprint, 2024. arXiv:2408.01052, doi:10.48550/arXiv.2408.01052.

[FA76] Bernard J. Fino and V. Ralph Algazi. Unified matrix treatment of the fast Walsh–Hadamard transform. *IEEE Transactions on Computers*, C-25(11):1142–1146, November 1976. doi:10.1109/TC.1976.1674569.

[GWH+25] Xinxin Gong, Qingju Wang, Yonglin Hao, Lin Jiao, and Xichao Hu. Persistence of hourglass(-like) structure: Improved differential-linear distinguishers for several ARX ciphers. In *Advances in Cryptology – ASIACRYPT 2025, Part I*, volume 16245 of *Lecture Notes in Computer Science*, pages 317–351. Springer, 2025. doi:10.1007/978-981-95-5018-0_11.

[HDE24] Hosein Hadipour, Patrick Derbez, and Maria Eichlseder. Revisiting differential-linear attacks via a boomerang perspective with application to AES, Ascon, CLEFIA, SKINNY, PRESENT, KNOT, TWINE, WARP, LBlock, Simeck, and SERPENT. In *Advances in Cryptology – CRYPTO 2024, Part IV*, volume 14923 of *Lecture Notes in Computer Science*, pages 38–72. Springer, 2024. doi:10.1007/978-3-031-68385-5_2.

[HHN+26] Jiahui He, Kai Hu, Zhongfeng Niu, Bart Preneel, and Meiqin Wang. Revisiting differential-linear cryptanalysis via a Walsh-transform perspective. Cryptology ePrint Archive, Paper 2026/2013, 2026. URL: https://eprint.iacr.org/2026/2013.

[HMAV26] Hosein Hadipour, Saleh Khalaj Monfared, Jens Alich, and Jan Vorloeper. LibFWHT: From exact Walsh spectra to key dependence in differential-linear correlations. Cryptology ePrint Archive, Paper 2026/1997, 2026. URL: https://eprint.iacr.org/2026/1997.

[HNW26] Kai Hu, Zhongfeng Niu, and Meiqin Wang. Round-based approximation of (higher-order) differential-linear correlation: A geometric approach perspective. In *Advances in Cryptology – EUROCRYPT 2026, Part VI*, volume 16546 of *Lecture Notes in Computer Science*, pages 152–180. Springer, 2026. doi:10.1007/978-3-032-25333-0_6.

[Leu16] Gaëtan Leurent. Improved differential-linear cryptanalysis of 7-round chaskey with partitioning. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 344–371. Springer, 2016.

[LH94] Susan K. Langford and Martin E. Hellman. Differential-linear cryptanalysis. In *Advances in Cryptology – CRYPTO '94*, volume 839 of *Lecture Notes in Computer Science*, pages 17–25. Springer, 1994. doi:10.1007/3-540-48658-5_3.

[LXZZ25] Ying Liu, Zejun Xiang, Shun Zhang, and Xiangyong Zeng. MILP-based automatic search algorithms for differential-linear distinguishers. *The Computer Journal*, 68(11):1772–1784, 2025. doi:10.1093/comjnl/bxaf073.

[LZ26] Guangqiu Lv and Yi Zhang. Automated differential–linear cryptanalysis of SIMON-like ciphers via truncated differential. *Designs, Codes and Cryptography*, 94:196, 2026. doi:10.1007/s10623-026-01931-w.

[NLZW26] Chao Niu, Muzhou Li, Jifu Zhang, and Meiqin Wang. Improved differential and linear cryptanalysis on round-reduced SIMON. *Designs, Codes and Cryptography*, 94:107, 2026. doi:10.1007/s10623-026-01827-9.

[Ope] Open-CP. OCP: Open cryptanalysis platform. Accessed: 2026-09-08. URL: https://github.com/Open-CP/OCP.

[PZWD24] Ting Peng, Wentao Zhang, Jingsui Weng, and Tianyou Ding. New approaches for estimating the bias of differential-linear distinguishers. In *Advances in Cryptology – CRYPTO 2024, Part IV*, volume 14923 of *Lecture Notes in Computer Science*, pages 174–205. Springer, 2024. doi:10.1007/978-3-031-68385-5_6.

[Roh24] Dongyoung Roh. Differential–linear approximations of CHAM. *Electronics*, 13(16):3141, 2024. doi:10.3390/electronics13163141.

[ZWH24] Yanyan Zhou, Senpeng Wang, and Bin Hu. MILP/MIQCP-based fully automatic method of searching for differential-linear distinguishers for SIMON-like ciphers. *IET Information Security*, 2024(1):8315115, 2024. doi:10.1049/2024/8315115.

A Overall Search Algorithm

Given the candidate round splits $\mathcal{R}$, Algorithm 1 describes the main search for three-part DL candidates. It uses HW1-guided subspace selection, FWHT-based middle-correlation evaluation, mask filtering, and SAT-based outer-trail completion to construct and rank candidates.

B Extension to SPN-Based Designs

We adapt the three-part search framework of Section 4.3 to SPN designs and report differential–linear results for `PRESENT` and `KNOT`.

B.1 Three-Part Search Framework for S-Box-Based Designs

We retain the three-part procedure used for AND–RX and ARX ciphers and adapt the generation of candidate middle input differences. For a state with c S-box positions of s bits each, we generate $c(2^s - 1)$ candidates by assigning a nonzero difference to a single S-box position. This uses the local S-box structure to limit the search space. For large candidate sets, we can rank and filter candidates by their differential propagation probabilities in E_1.

When a linear diffusion layer follows the S-box layer, we can generate differences with a single active unit at the S-box output. Let δ_{mid} denote a candidate at this position and L_{rem} the linear transformation to the input boundary of E_m. The middle input difference is $\Delta_m = L_{\mathrm{rem}}(\delta_{\mathrm{mid}})$. We use the mapped Δ_m for both differential-prefix search and middle-spectrum evaluation within the given three-part round split.

Algorithm 1 Three-Part Differential–Linear Search

Require: Primitive E, candidate round splits $\mathcal{R}$, subspace dimension k, sample size N, threshold τ, HW bound $h_{\max}$, candidate budgets B_{HW}, B_{cand}, and B_{keep}

Ensure: Ranked discovery candidates

$\mathcal{L}_{\mathrm{search}} \leftarrow \varnothing$
for each $(R_1, R_m, R_2) \in \mathcal{R}$ **do**
  Split E into (E_1, E_m, E_2)
  Construct the candidate middle-difference set $\mathcal{D}_m$
  $\mathcal{L}_{\mathrm{split}} \leftarrow \varnothing$
  for each $\Delta_m \in \mathcal{D}_m$ **do**
    Find the best single differential trail in E_1 ending at Δ_m, obtaining $(\Delta_{\mathrm{in}}, p)$
    if $k = n$ **then**
      $I \leftarrow \{0, \ldots, n-1\}$
    else
      Estimate all HW1 correlations using N input pairs
      Select I according to Eq. (4)
    end if
    Generate a fresh batch of N input pairs under Δ_m
    Evaluate $\widehat{r}(\Delta_m, \Gamma_m)$ for all $\Gamma_m \in U_I$ by FWHT
    Discard the zero mask and masks with $|\widehat{r}| \le \tau$
    if $R_2 > 0$ and the HW bound applies **then**
      Discard masks with $\mathrm{HW}(\Gamma_m) > h_{\max}$
    end if
    Retain up to B_{HW} masks with the largest $|\widehat{r}|$ at each permitted Hamming weight
    From their union, retain the top B_{cand} masks by $|\widehat{r}|$
    for each retained $(\Gamma_m, \widehat{r})$ **do**
      if $R_2 = 0$ **then**
        $\Gamma_{\mathrm{out}} \leftarrow \Gamma_m$, $q \leftarrow 1$
      else
        if $|\mathcal{L}_{\mathrm{split}}| \ge B_{\mathrm{keep}}$ **then**
          Set $\mathrm{Cor}_{\mathrm{cut}}$ to the smallest retained $|\widehat{\mathrm{Cor}}|$
          $q_{\min} \leftarrow \sqrt{\dfrac{\mathrm{Cor}_{\mathrm{cut}}}{p|\widehat{r}|}}$
          Find the best single linear trail in E_2 from Γ_m subject to $|q| > q_{\min}$
        else
          Find the best single linear trail in E_2 from Γ_m
        end if
      end if
      if $R_2 = 0$ or a qualifying linear trail is found **then**
        $\Pi \leftarrow (R_1, R_m, R_2, \Delta_{\mathrm{in}}, \Delta_m, \Gamma_m, \Gamma_{\mathrm{out}})$
        Compute $|\widehat{\mathrm{Cor}}(\Pi)| \leftarrow p|\widehat{r}|q^2$
        Insert Π into $\mathcal{L}_{\mathrm{split}}$, retaining the B_{keep} largest correlation magnitudes
      end if
    end for
  end for
  $\mathcal{L}_{\mathrm{search}} \leftarrow \mathcal{L}_{\mathrm{search}} \cup \mathcal{L}_{\mathrm{split}}$
end for
Rank $\mathcal{L}_{\mathrm{search}}$ by total rounds and then by $|\widehat{\mathrm{Cor}}|$
return $\mathcal{L}_{\mathrm{search}}$

B.2 SPN Experimental Results

For PRESENT, the three-part estimate uses $\widehat{r}_{\mathrm{abs}}$ from independent fixed-key validation. We sample 50 master keys independently and evaluate 2^{34} fresh input pairs under each key. For the unkeyed KNOT permutation, we reevaluate the retained candidates with 50 independent batches of 2^{34} input pairs.

Results for PRESENT. PRESENT is a lightweight SPN block cipher with a 64-bit block and an 80- or 128-bit key. Given the input state X_i and round subkey K_i, one round updates the state as

$$X_{i+1} = P(S(X_i \oplus K_i)),$$

where S applies 16 parallel 4-bit S-boxes and P is the fixed bit permutation.

For PRESENT-64/80, we use a $2 + 9 + 4$ split and generate $16 \times 15 = 240$ differences with one active S-box at the input to P. Their optimal two-round differential probabilities occur with the multiplicities

$$2^{-4} : 112, \qquad 2^{-5} : 64, \qquad 2^{-6} : 16, \qquad 2^{-7} : 48.$$

We retain the 112 candidates with probability 2^{-4} and map them through P to the E_m input boundary for middle-spectrum evaluation. Table 11 reports the resulting DL distinguisher.

Table 11: Main differential–linear result for PRESENT-64/80.

Variant	Split	Δ_m	Γ_m	p	q^2	$\lvert\widehat{r}_{\mathrm{mix}}\rvert$	$\widehat{r}_{\mathrm{abs}}$	$p\widehat{r}_{\mathrm{abs}}q^2$
64/80	$2+9+4$	00000002 00000002	00800000 00000000	2^{-4}	2^{-14}	$2^{-13.69}$	$2^{-13.90}$	$2^{-31.90}$

Hadipour et al. reported a 13-round DL distinguisher for PRESENT [HDE24]. Liu et al. used pre-pruning and MILP search to obtain a 14-round distinguisher with a $1+8+5$ split and correlation $2^{-29.11}$ [LXZZ25]. Our $2 + 9 + 4$ split extends the three-part result of Liu et al. by one round and yields an estimated correlation of $2^{-31.90}$.

Hu et al. subsequently obtained a 17-round result using a round-based approximation. Our PRESENT experiment therefore serves to demonstrate the applicability of the proposed search framework to SPN designs, rather than to establish the longest PRESENT distinguisher.

Results for KNOT. KNOT is a family of lightweight permutation-based authenticated-encryption schemes. The KNOT-256 permutation considered here acts on four 64-bit words. KNOT-AEAD$(128, 256, 64)$ uses this permutation with a 128-bit key, a 128-bit nonce, and a 64-bit rate.

In the *AEAD-constrained setting*, we compare two initializations with a common fixed key and distinct nonces. The input difference is zero in the key portion, and the output mask is supported on the rate portion. Under our word ordering, these constraints take the form

$$\Delta_{\mathrm{in}} = [0, 0, *, *], \qquad \Gamma_{\mathrm{out}} = [*, 0, 0, 0].$$

Each tuple lists four ordered 64-bit words, with $*$ denoting an arbitrary 64-bit word. In the *unconstrained permutation setting*, we search over general difference–mask pairs for the same core permutation. Table 12 lists the four boundary values of each complete path.

In both settings, we use uniformly sampled states to estimate the middle correlations. The estimates $\widehat{r}_{\mathrm{search}}$ and $\widehat{r}_{\mathrm{val}}$ use discovery samples and fresh independent validation samples, respectively. Table 13 reports the resulting distinguishers.

Table 12: Endpoints of the complete differential–linear paths in the two KNOT-256 analysis settings.

Setting	Δ_{in}	Δ_m	Γ_m	Γ_{out}
AEAD	0000000000000000	0000000000000000	0020000000000000	0040000000000000
	0000000000000000	0000000000002000	0040000000000000	0000000000000000
	0000000000000008	0000000000000000	0020000000000000	0000000000000000
	0000000000000008	0000002000000000	0040000000000000	0000000000000000
Perm.	0004000000000000	0000000000000000	0000000000000030	0000000000000080
	0004000200000000	0000000000000010	0000000000000020	0000000100000000
	0004000000000000	0000000000000000	0000000000000000	0000000000000000
	0000000200000000	0000000010000000	0000000000000000	0100000100000000

Table 13: Main differential–linear results for KNOT-256 in the two analysis settings.

Setting	Split	p	q^2	$\lvert\widehat{r}_{\text{search}}\rvert$	$\lvert\widehat{r}_{\text{val}}\rvert$	$p\lvert\widehat{r}_{\text{val}}\rvert q^2$
AEAD	$4+9+4$	2^{-10}	2^{-18}	$2^{-3.77}$	$2^{-3.77}$	$2^{-31.77}$
Perm.	$5+10+9$	2^{-14}	2^{-42}	$2^{-7.04}$	$2^{-7.04}$	$2^{-63.04}$

In the AEAD-constrained setting, the longest previous result covers 16 rounds with a $4+8+4$ split and an estimated correlation of $2^{-29.52}$ after conversion to our correlation measure [PZWD24]. Our $4+9+4$ split extends the middle component to 9 rounds and yields a 17-round distinguisher with an estimated correlation of $2^{-31.77}$.

In the unconstrained permutation setting, the longest previous result covers 23 rounds with a $6+9+8$ split and an estimated correlation of $2^{-58.88}$ [HDE24]. Our $5+10+9$ split extends the middle component to 10 rounds and yields a 24-round distinguisher with an estimated correlation of $2^{-63.04}$.

C SAT-Guided Window Aggregation for SIMON128

We illustrate differential-trail aggregation on an 18-round SIMON128 prefix. Using the bit-index convention from the main text, we fix the input and output differences as

$$\Delta_{\text{in}} = [29, 33, 95] = (\texttt{0x0000000080000000}, \texttt{0x0000000220000000}),$$
$$\Delta_m = [29] = (\texttt{0x0000000000000000}, \texttt{0x0000000020000000}).$$

The best single-trail weight between these endpoints is $W_{\text{best}} = 64$.

C.1 SAT-Guided Dynamic-Window Aggregation

To restrict the state-difference space, we assign width-w bit windows W_i^L and W_i^R to the left and right branches at each round boundary, following Niu et al. [NLZW26]. Setting the remaining difference bits to zero leaves 2^{2w} permitted state differences per boundary. The window positions can vary across boundaries and must accommodate the fixed differences at the two endpoints.

For a window family $\mathcal{W}$, the permitted state differences at boundary i are

$$\mathcal{S}_i(\mathcal{W}) = \left\{(L_i, R_i) : \text{supp}(L_i) \subseteq W_i^L, \quad \text{supp}(R_i) \subseteq W_i^R\right\}.$$

Let $T(\alpha, \gamma)$ denote the one-round differential transition probability from α to γ. We compute backward from the fixed output difference using the recursion

$$B_{18}(\gamma) = \mathbf{1}[\gamma = \Delta_m],$$

$$B_i(\alpha) = \sum_{\gamma \in \mathcal{S}_{i+1}(\mathcal{W})} T(\alpha, \gamma) B_{i+1}(\gamma), \qquad \alpha \in \mathcal{S}_i(\mathcal{W}).$$

The DP result is

$$P_{\mathrm{DW}}(\mathcal{W}) = B_0(\Delta_{\mathrm{in}}) = \sum_{\pi \in \mathcal{T}_{\mathrm{DW}}(\mathcal{W})} 2^{-W(\pi)},$$

where $\mathcal{T}_{\mathrm{DW}}(\mathcal{W})$ contains all trails satisfying the window constraints and $W(\pi)$ is the differential weight of trail π under the adopted Markov model. The recursion sums their probabilities implicitly, giving a lower bound on the full fixed-endpoint differential-hull probability in this model.

We initialize the windows from the per-round active-position rankings obtained by propagating the fixed output difference backward. With both endpoints fixed, projected AllSAT enumerates trails by total weight using the transition constraints of the single-trail search. Each state-difference sequence $(\alpha_0, \ldots, \alpha_{18})$ identifies a trail. Blocking clauses use only state-difference variables, so a trail counts once across auxiliary CNF assignments. A fully enumerated weight layer W containing N_W trails contributes $N_W 2^{-W}$. Full enumeration of weights $64, \ldots, 75$ yields 829,042 distinct trails.

Let $\mathcal{T}_{\mathrm{SAT}}$ denote the enumerated trails. For each trail outside the current windows, we record the first exit boundary r and the set S of active positions outside the windows at that boundary. The corresponding probability mass is

$$M(r, S) = \sum_{\substack{\pi \in \mathcal{T}_{\mathrm{SAT}} \\ \text{first exit at } r \text{ through } S}} 2^{-W(\pi)}.$$

We rank the pairs (r, S) by $M(r, S)$ and prioritize positions in higher-ranked sets for insertion. To preserve the window width, we remove the lowest-ranked existing positions outside the support of the protected trails. We protect the supports of both trails in the fully enumerated weight-64 layer throughout.

We evaluate each candidate replacement using transition DP and retain the window family with the largest fixed-endpoint probability. This selects a single DP result because the window families may cover overlapping trail sets. We reuse the enumerated SAT trails across iterations.

For the selected window family $\mathcal{W}$, we add the probability mass of the enumerated trails outside the windows to the DP contribution

$$P_{\mathrm{union}} = P_{\mathrm{DW}}(\mathcal{W}) + \sum_{\pi \in \mathcal{T}_{\mathrm{SAT}} \setminus \mathcal{T}_{\mathrm{DW}}(\mathcal{W})} 2^{-W(\pi)}.$$

The two terms sum probabilities over disjoint trail sets, so P_{union} is a lower bound on the full fixed-endpoint differential-hull probability under the model. We use it as the aggregated prefix contribution.

C.2 Aggregation Results

Table 14 classifies the fully enumerated trails of weights $68 \le W \le 75$ by membership in the final SAT-guided width-17 windows. The 256 trails of weights $64 \le W \le 67$ all lie within the final windows.

The final windows contain 94.37% of the trails in the eight displayed layers. Outside-window trails account for 3.99% of their probability mass and for 3.84% of the mass across

Table 14: SAT trail audit for the final width-17 windows. The total-row exponents are computed after summing probabilities across the displayed layers.

W	N_W	N_W^{in}	**In (%)**	N_W^{out}	$-\log_2(N_W^{\text{in}}2^{-W})$	$-\log_2(N_W^{\text{out}}2^{-W})$
68	663	659	99.40	4	58.64	66.00
69	1,994	1,972	98.90	22	58.05	64.54
70	5,648	5,551	98.28	97	57.56	63.40
71	14,842	14,496	97.67	346	57.18	62.57
72	37,301	36,107	96.80	1,194	56.86	61.78
73	89,574	86,025	96.04	3,549	56.61	61.21
74	208,480	197,378	94.67	11,102	56.41	60.56
75	470,284	439,941	93.55	30,343	56.25	60.11
Total	828,786	782,129	94.37	46,657	54.01	58.59

the fully enumerated range $64 \le W \le 75$, including the 256 in-window trails from the lower weight layers. The outside-window contribution is approximately $2^{-58.59}$; adding it to $P_{\text{DW}} \approx 2^{-51.63}$ yields $P_{\text{union}} \approx 2^{-51.62}$, an improvement by a factor of approximately $2^{0.01}$. By comparison, SAT-guided window adjustment improves the width-17 DP result from $2^{-52.65}$ to $2^{-51.63}$, corresponding to a factor of $2^{1.02}$, by bringing additional trails into the DP, including trails beyond the enumerated weight range.

Table 15: Dynamic-window DP probabilities at different window widths.

Window width	**Initial P_{DW}**	**SAT-guided P_{DW}**	**Gain**
14	$2^{-52.82}$	$2^{-51.82}$	$2^{1.00}$
15	$2^{-52.77}$	$2^{-51.76}$	$2^{1.01}$
16	$2^{-52.68}$	$2^{-51.66}$	$2^{1.02}$
17	$2^{-52.65}$	$2^{-51.63}$	$2^{1.02}$

As shown in Table 15, SAT guidance improves the DP probability by factors between $2^{1.00}$ and $2^{1.02}$ at the tested widths. Increasing the window width from 16 to 17 further improves the SAT-guided result from $2^{-51.66}$ to $2^{-51.63}$, corresponding to a factor of approximately $2^{0.03}$. Convergence to the full differential-hull probability remains undetermined. The width-17 DP contains 2^{34} states per layer and takes approximately 29.80 hours.

At width 17, the DP and final aggregated prefix probabilities are

$$P_{E_1}^{\text{DW}} = P_{\text{DW}}(\mathcal{W}) \approx 2^{-51.63}, \qquad P_{E_1}^{\text{agg}} = P_{\text{union}} \approx 2^{-51.62}.$$

Relative to the best single-trail probability of 2^{-64}, the aggregated prefix probability is improved by a factor of approximately $2^{12.38}$.

D Detailed Analysis of Outer-Path Aggregation Gains

For each candidate in Section 5.5, we compare the refined three-part estimate with a single-trail baseline under the same round split and boundary values. Section 5.4 reports the best results of the single-trail search.

Let p and q denote the probability and correlation of the best single trails for the differential prefix and linear suffix, respectively. The total gain over the baseline $p\widehat{r}_{\mathrm{abs}}q^2$ is

$$G_{\mathrm{agg}} = \frac{\widehat{|\mathrm{Cor}_{\mathrm{agg}}|}}{p\widehat{r}_{\mathrm{abs}}q^2}.$$

The prefix gain is $P_{E_1}^{\mathrm{agg}}/p$. For the suffix, we compute the mean squared empirical correlation across the fixed keys,

$$\widehat{Q}_{E_2}^{\mathrm{emp}} = \frac{1}{T}\sum_{j=1}^{T}\widehat{q}_{K_j}^2,$$

and report $\widehat{Q}_{E_2}^{\mathrm{emp}}/q^2$ as the marginal suffix gain. Table 16 lists the prefix, marginal suffix, and total gains.

Table 16: Differential-prefix and linear-suffix contributions to the total correlation gain relative to single-trail evaluation of the same fixed-endpoint candidates.

Cipher	p	$P_{E_1}^{\mathrm{agg}}$	$P_{E_1}^{\mathrm{agg}}/p$	q^2	$\widehat{Q}_{E_2}^{\mathrm{emp}}$	$\widehat{Q}_{E_2}^{\mathrm{emp}}/q^2$	G_{agg}
SIMON48	$2^{-4.00}$	$2^{-4.00}$	1	$2^{-10.00}$	$2^{-9.21}$	$2^{0.79}$	$2^{0.82}$
SIMON64	$2^{-10.00}$	$2^{-9.30}$	$2^{0.70}$	$2^{-10.00}$	$2^{-9.08}$	$2^{0.92}$	$2^{1.59}$
SIMON96	$2^{-28.00}$	$2^{-24.17}$	$2^{3.83}$	$2^{-10.00}$	$2^{-9.14}$	$2^{0.86}$	$2^{4.67}$
SIMON128	$2^{-64.00}$	$2^{-51.62}$	$2^{12.38}$	1	1	1	$2^{12.38}$
SIMECK64	$2^{-4.00}$	$2^{-4.00}$	1	$2^{-12.00}$	$2^{-11.39}$	$2^{0.61}$	$2^{0.59}$

For SIMON48 and SIMECK64, the gain comes entirely from empirical suffix evaluation. Both prefix aggregation and empirical suffix evaluation improve the estimates for SIMON64 and SIMON96. The $2^{12.38}$ gain for SIMON128 comes entirely from prefix aggregation because the suffix has zero rounds.

The refined estimate uses the mean of the paired fixed-key products

$$\frac{1}{T}\sum_{j=1}^{T}|\widehat{r}_{K_j}|\widehat{q}_{K_j}^2,$$

and thus retains the joint variation of the middle and suffix contributions across keys. This variation can make the total gain G_{agg} differ from the product of the prefix and marginal suffix gain ratios.

E Complete Distinguisher Parameters

Table 17 lists the round splits, boundary differences and masks, and outer single-trail factors for the distinguishers in Section 5.4. Here p is the probability of the differential prefix trail, and q^2 is the squared correlation of the linear suffix trail. For a zero-round suffix, $\Gamma_{\mathrm{out}} = \Gamma_m$ and $q^2 = 1$.

Table 17: Parameters of the single-trail DL distinguishers. Brackets list active bit indices.

Cipher / split	Δ_{in}	Δ_m	Γ_m	Γ_{out}	p	q^2
SIMON32 $3+11+0$	[10, 14, 28]	[26]	[11, 12, 17, 29]	[11, 12, 17, 29]	2^{-4}	1
SIMON48 $2+12+4$	[3, 23, 25]	[23]	[4]	[0, 4, 20, 46]	2^{-4}	2^{-10}
SIMON64 $5+14+3$	[0, 4, 8, 38]	[32]	[6, 40]	[0, 4, 8, 34]	2^{-10}	2^{-8}
SIMON96 $4+17+8$	[0, 4, 8, 54]	[0]	[2, 48]	[2, 38, 42, 48, 88]	2^{-10}	2^{-22}
SIMON128 $8+22+6$	[29, 37, 41, 45, 107]	[29]	[30]	[18, 22, 30, 84, 92]	2^{-28}	2^{-20}
SIMECK32 $2+12+0$	[0, 2, 17]	[0]	[13, 14, 15, 30]	[13, 14, 15, 30]	2^{-4}	1
SIMECK48 $3+15+4$	[20, 22, 45]	[44]	[19, 42]	[17, 19, 42]	2^{-4}	2^{-6}
SIMECK64 $4+20+4$	[1, 29, 31, 32]	[29]	[28]	[24, 26, 28, 57]	2^{-10}	2^{-10}
SPECK48 $2+8+2$	[9, 17, 20, 25, 31, 41]	[39]	[8, 32]	[14, 19, 22, 26, 27, 38, 39, 40, 42, 46]	2^{-5}	2^{-6}
SPECK64 $2+9+3$	[1, 9, 12, 41, 49, 55]	[39]	[32]	[1, 11, 17, 22, 33, 39, 40, 43, 46, 49]	2^{-6}	2^{-12}
SPECK96 $3+9+4$	[22, 33, 38, 44, 49, 52, 55, 78, 94]	[79]	[5, 53]	[0, 6, 8, 17, 27, 33, 37, 38, 48, 53, 56, 65, 72, 75, 77, 78, 81, 93, 94]	2^{-11}	2^{-24}
SPECK128 $5+9+4$	[4, 9, 14, 15, 26, 46, 57, 73, 76, 86, 89, 98, 101, 118]	[125]	[77]	[1, 5, 6, 9, 16, 17, 54, 56, 57, 65, 69, 73, 80, 81, 114, 115, 120, 121]	2^{-30}	2^{-18}
CHAM-64/128 $7+33+7$	[7, 22, 28, 36, 43, 46, 58, 61]	[11]	[49]	[2, 15, 17, 41, 56]	2^{-13}	2^{-6}
CHAM-128/128 $10+39+11$	[5, 12, 18, 24, 30, 43, 49, 62, 73, 104, 110]	[110]	[67, 68]	[4, 5, 12, 18, 24, 29, 30, 36, 42, 73, 97, 103, 116]	2^{-21}	2^{-28}