Generalized Integral Properties: Separable Structure and Structured Search

Dongchen Chai[1], Shenghu Hu[1] Thomas Peyrin[2], Zilong Wang[1] Trevor Yap[2], Hongyi Zhang[2], Liu Zhang[2✉], and Chunning Zhou[2]

[1] State Key Laboratory of Integrated Services Networks, School of Cryptology, Xidian University, Xi'an, China chaidc@foxmail.com, shenghu.hu@stu.xidian.edu.cn, zlwang@xidian.edu.cn

[2] School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore {thomas.peyrin,trevor.yap,liu.zhang,chunning.zhou}@ntu.edu.sg, hongyi003@e.ntu.edu.sg

Abstract. Generalized integral properties extend classical integral distinguishers by allowing general Boolean mappings over plaintext and ciphertext variables. This additional flexibility raises two natural cryptanalytic questions: whether generalized mappings can yield higher-round distinguishers, and whether the resulting relations can be exploited to improve key-recovery attacks. Addressing these questions requires understanding which generalized integral structures are relevant and how they can be searched systematically. In this work, we introduce a structural classification of key-independent generalized integral properties into separably generated and residual-cancellation properties. Although some previously reported generalized integral properties appear to involve cancellations among key-dependent residuals, we show through explicit decompositions that they are in fact separably generated. Motivated by this view, we search for balanced separable terms within explicit plaintext-side and ciphertext-side candidate spaces.

We develop two complementary automated search strategies: a unified Matrix–Bipartite Graph–Matrix framework for joint plaintext–ciphertext search and a specialized Gaussian-elimination-based solver for ciphertext-side mappings under a fixed plaintext structure. Across five representative block ciphers, the proposed methods extend previous integral distinguishers by one round for `SIMON` and for `SPECK` variants other than `SPECK` 32, while the other evaluated ciphers mainly yield additional balanced mappings at the same round depth. We further show that the underlying round-extension mechanisms for the `SPECK` family and `SIMON` arise from exploitable algebraic structures in their round functions, whereas no comparable simplification is observed for the remaining ciphers under the current search setting. Using `SKINNY-64-64` as a case study, we show that the additional nonlinear relations provide stronger filtering but destroy the computational separability exploited by existing partial-sum techniques, and therefore do not yield a better key-recovery attack under the current approach.

Keywords: Integral Cryptanalysis, Generalized Integral Properties, Separable Structure, Automated Search, Gaussian Elimination

1 Introduction

Integral cryptanalysis is a fundamental technique for analyzing symmetric-key primitives. Knudsen and Wagner introduced it as a generalization of the `SQUARE` attack based on zero-sum distinguishers [14]. Todo later formulated integral distinguishers through the division property with two subsets [17], which was subsequently refined by Todo and Morii into the bit-based division property with three subsets [18]. Boura and Canteaut provided an alternative formulation through parity sets, connecting the division property to the algebraic normal form of cipher components [4]. These developments enabled increasingly accurate automated propagation models. Xiang *et al.* [22] introduced division trails and an MILP framework for the two-subset setting, while Wang *et al.* [21] extended this approach to the three-subset setting. Hao *et al.* [7] developed an efficient method without unknown subsets, and Hu *et al.* [10] further improved the propagation model through monomial prediction.

Early integral distinguishers mainly focused on the balancedness of individual ciphertext bits. Todo and Morii [18] observed that this coordinate-wise view is incomplete: even when several output bits are individually classified as *unknown*, an appropriate combination of them may still be balanced. This observation motivated the study of integral properties involving mappings over multiple output bits. Lambin *et al.* [15] exploited the dependence of division-property propagation on linear representations and searched for distinguishers over $L_{\mathrm{out}} \circ E \circ L_{\mathrm{in}}$. Derbez and Fouque [5] subsequently developed more efficient procedures for input/output linear mappings together with more precise propagation techniques. Beyne and Verbauwhede [3] formalized this broader perspective through generalized integral properties, allowing a Boolean function to depend jointly on plaintext and ciphertext variables.

The increased flexibility raises two natural cryptanalytic questions. First, to what extent can generalized plaintext- and ciphertext-side mappings extend integral distinguishers beyond the round depths attainable by classical integral properties? Second, when additional generalized integral relations are obtained, to what extent can they be exploited to improve key-recovery attacks? Answering these questions requires understanding not only how such properties can be searched systematically, but also which structural features enable additional round propagation and whether the resulting relations are compatible with efficient attack procedures.

At the same time, the general definition permits arbitrary Boolean functions of the plaintext and ciphertext variables, making direct exploration of the full function space computationally infeasible. Practical automated methods must therefore exploit additional structure or restrict the admissible candidate space. Many classical and mapping-based integral constructions naturally admit a separable form, with one function on the plaintext side and another on the ciphertext side. However, a separable decomposition does not necessarily imply that the individual integral sum of each component is key-independent. More general formulations may, in principle, combine separable components whose individual integral sums depend on the key, while their key-dependent residual

sums cancel in the final combination. This raises a structural question that has not been made explicit in previous work: which generalized integral properties admit decompositions into individually key-independent separable terms, and which genuinely rely on cancellations among key-dependent residual sums? This distinction is relevant both for understanding the expressive power of existing approaches and for identifying a tractable search space.

Our Contributions Motivated by these observations, we introduce a classification of key-independent generalized integral properties into separably generated and residual-cancellation properties. Focusing on balanced separable terms, we construct structured plaintext-side and ciphertext-side candidate spaces and develop automated methods for their search. We further use the resulting framework to study when generalized mappings extend integral distinguishers and whether the resulting relations benefit key-recovery attacks. Our main contributions are summarized as follows.

- **Structural Classification of Generalized Integral Properties.** We characterize separably generated properties as those admitting a decomposition in which each separable term has a key-independent integral sum, and distinguish them from residual-cancellation properties, whose key-independent overall sum relies on cancellations among key-dependent residual sums of different separable components. Classical integral constructions and representative mapping-based distinguishers naturally fall into the separably generated class. More general formulations can, in principle, represent residual-cancellation properties. However, through explicit decompositions, we show that previously reported generalized integral properties that appear to involve such cancellations are still separably generated.
- **Structured Candidate Spaces and Automated Search.** We construct explicit plaintext-side and ciphertext-side candidate spaces for balanced separable terms, covering combinations of basic plaintext structures and degree-bounded Boolean output mappings. For joint plaintext–ciphertext search, we develop the MBM (Matrix–Bipartite Graph–Matrix) framework, which connects surviving plaintext-side sources to ciphertext-side monomials through a middle propagation graph and reduces the resulting search to an MILP problem. For a fixed plaintext structure, we further show that the ciphertext-side certification condition forms a homogeneous linear system in the mapping coefficients. This yields a specialized Gaussian-elimination solver that recovers the entire solution space certified by the induced criterion within the chosen degree bound, without enumerating mapping combinations.
- **Applications, Key-Recovery Impact, and Structural Insights.** We apply the proposed methods to `SPECK`, `SIMON`, `PRESENT`, `RECTANGLE`, and `SKINNY`[3]. The discovered mappings extend previous integral distinguishers

[3] A repository containing the code and experimental results is available at https://anonymous.4open.science/r/general-framework-for-mapping-based-integral-property-CA84/.

by one round for SIMON and for the SPECK variants except SPECK32. For the other evaluated ciphers, the methods recover or enlarge the set of balanced mappings at the same round depth. We further analyze this difference and show that the underlying round-extension mechanisms for the SPECK family and SIMON arise from algebraic relations in their round functions that can be exploited by the searched mappings, whereas no comparable simplification is observed for the other ciphers within the current candidate spaces and propagation models. Finally, using SKINNY-64-64 as a case study, we investigate whether the additional generalized integral relations can improve key-recovery attacks. The nonlinear relations provide stronger filtering but destroy the computational separability exploited by the existing partial-sum procedure, and therefore do not yield a better attack under current techniques. A detailed comparison with previous distinguisher results is provided in Table 1.

A very recent work by Wang, Hadipour, and Gerhalter [19] independently develops a closely related kernel-based ciphertext-side search. Our fixed-plaintext matrix-based solver had already been included in an earlier conference submission before this concurrent work became publicly available. The two works were therefore developed independently. Beyond the fixed-plaintext ciphertext-side search, the present work further develops a structural classification of key-independent generalized integral properties, a joint plaintext–ciphertext search framework, and an analysis of the impact of nonlinear mappings on key-recovery attacks.

Outline. Section 2 reviews the preliminaries. Section 3 presents the structural classification and candidate spaces. Section 4 develops the proposed search methods. Section 5 presents applications and analyzes distinguisher results, key-recovery implications, and round-extension mechanisms. Section 6 concludes the paper and discusses future work.

2 Preliminaries

Notions Let $\mathbb{F}_2 = \{0, 1\}$ denote the finite field with two elements, and let $\mathbb{F}_2^n$ denote the n-dimensional vector space over $\mathbb{F}_2$. For $u \in \mathbb{F}_2^n$, let $\mathrm{wt}(u)$ denote its Hamming weight. For $\kappa, \kappa' \in \mathbb{F}_2^n$, we write $\kappa \succeq \kappa'$ if $\kappa_i \geq \kappa'_i$ for all $i \in \{0, \ldots, n-1\}$, and let $e_i \in \mathbb{F}_2^n$ denote the i-th unit vector. For a fixed vector $a = (a_0, \ldots, a_{n-1}) \in \mathbb{F}_2^n$, define

$$\mathbb{X}_i = \{x \in \mathbb{F}_2^n : x_i = a_i\}, \quad i \in \{0, \ldots, n-1\},$$

where $a_i \in \mathbb{F}_2$ is the prescribed fixed value of the i-th plaintext bit. Thus, $\mathbb{X}_i$ is the basic plaintext structure obtained by fixing only the i-th bit and leaving the remaining $n-1$ bits unrestricted. For any $\mathbb{X} \subseteq \mathbb{F}_2^n$, we use $\mathbf{1}_{\mathbb{X}}$ to denote its indicator function.

Table 1. Detailed comparison of integral distinguishers

Cipher	Round	Data	(S, L, NL)	Time	Reference
SPECK32	7	2^{30}	(0, 1, 0)	–	[13]
	7	2^{29}	(0, 1, 0)	18.14 min	Section 5.1
SPECK48	6	2^{45}	(1, 0, 0)	2 s	[11]
	7	2^{44}	(0, 1, 0)	12.15 min	Section 5.1
SPECK64	6	2^{61}	(1, 0, 0)	3.45 min	[11]
	7	2^{52}	(0, 1, 0)	12.48 min	Section 5.1
SPECK96	6	2^{93}	(1, 0, 0)	18.6 min	[11]
	7	2^{68}	(0, 1, 0)	12.57 min	Section 5.1
SPECK128	6	2^{125}	(1, 0, 0)	16 h	[11]
	7	2^{84}	(0, 1, 0)	13.87 min	Section 5.1
SIMON32	14	2^{31}	(3, 0, 0)	2 min	[21]
	15	2^{31}	(0, 0, 48)	1.62 h	Section 5.2
SIMON48	15	2^{47}	(24, 0, 0)	–	[22]
	16	2^{47}	(0, 0, 576)	11.41 min	Section 5.2
SIMON64	17	2^{63}	(23, 0, 0)	6.7 min	[21]
	18	2^{63}	(0, 0, 704)	21.72 min	Section 5.2
SIMON96	21	2^{95}	(5, 0, 0)	–	[22]
	22	2^{95}	(0, 0, 240)	41.23 min	Section 5.2
SIMON128	25	2^{127}	(3, 0, 0)	–	[22]
	26	2^{127}	(0, 0, 192)	1.95 h	Section 5.2
PRESENT	9	2^{63}	(28, 0, 0)	10 min	[21]
	9	2^{63}	(130, 85 , 0)	28.7 h‡	[3]
	9	2^{63}	(138, 66, 0)	1.2 h	Section 5.2
RECTANGLE	9	2^{60}	(27, 0, 0)	10 min	[21]
	10	2^{63}	(1, 0, 0)	6.67 min	[15]
	10	2^{63}	(16, 0, 0)	11.8 min	Section 5.2
SKINNY-64-64	11	2^{60}	(0, 17, 0)	–	[23]
	11	2^{60}	(0, 21†, 0)	–	[13]
	11	2^{60}	(0, 21, 40)	5.46 h	Section 5.3
SKINNY-128-128	11	2^{120}	(0, 33*, 0)	–	[23]
	11	2^{120}	(0, 33, 0)	2.58 h	Section 5.3

(S, L, NL) denotes the numbers of generalized integral properties classified according to the type of the ciphertext-side mapping $g(y)$, namely single-bit, linear, and nonlinear mappings, respectively. Each pair $(h(x), g(y))$ is counted as one generalized integral property, while its category is determined solely by the type of $g(y)$. The plaintext-side mapping $h(x)$ specifies the choice of the input set and therefore does not affect the classification.
The standard one-round free extensions for SIMON [20] and SPECK are not included in the reported round numbers.
‡: Beyne and Verbauwhede [3] report 215 essentially distinct distinguishers over eight partitions in 28.7 h; the 1-7-1 partition contains 204 distinguishers and takes 2.3 h.
†: the original paper reports 30 mappings; after removing equivalent representations, 21 distinct mappings remain.

*: the value reported in [23] contains a typo; the correct number is 33.

Boolean Functions A function $f : \mathbb{F}_2^n \to \mathbb{F}_2$ is called a Boolean function. Every Boolean function admits a unique algebraic normal form (ANF),

$$f(x) = \bigoplus_{u \in \mathbb{F}_2^n} a_u \pi_u(x),$$

where $a_u \in \mathbb{F}_2$ and $\pi_u(x) = x^u = \prod_{i=0}^{n-1} x_i^{u_i}$, for $x = (x_0, \ldots, x_{n-1})$ and $u = (u_0, \ldots, u_{n-1})$. The algebraic degree of f, denoted by $\deg(f)$, is the maximum degree of the monomials with nonzero coefficients in its ANF.

We next recall the notions used to characterize monomial sums in integral cryptanalysis. The bit-based division property with two subsets distinguishes monomial sums that are guaranteed to be zero from those that remain unknown, while the three-subset setting additionally tracks sums equal to one. Monomial prediction provides a finer monomial-level characterization and connects these integral sums to the ANF representation.

Division Property with Two Subsets Todo introduced the division property at EUROCRYPT 2015 [17]. Todo and Morii later introduced its bit-based formulation, including the two-subset and three-subset settings [18]. In the two-subset setting, each monomial sum is classified as either guaranteed to be zero or *unknown*.

Definition 1 (Bit-based Division Property with Two Subsets). *A multiset $\mathbb{X}$ over $\mathbb{F}_2^n$ is said to have division property $D_{\mathbb{K}}^n$, where $\mathbb{K} \subseteq \mathbb{F}_2^n$, if for every $u \in \mathbb{F}_2^n$,*

$$\bigoplus_{x \in \mathbb{X}} \pi_u(x) = \begin{cases} \textit{unknown}, & \textit{if there exists } \kappa \in \mathbb{K} \textit{ such that } u \succeq \kappa, \\ 0, & \textit{otherwise.} \end{cases}$$

Thus, $\mathbb{K}$ specifies the monomials whose sums are unknown, while all remaining monomials are guaranteed to have zero sum.

Division Property with Three Subsets Todo and Morii [18] further introduced the bit-based division property with three subsets, which distinguishes monomial sums equal to one in addition to zero and unknown sums.

Definition 2 (Bit-Based Division Property with Three Subsets). *A multiset $\mathbb{X}$ over $\mathbb{F}_2^n$ is said to have division property $D_{\mathbb{K},\mathbb{L}}^{1^n}$, where $\mathbb{K}, \mathbb{L} \subseteq \mathbb{F}_2^n$, if for every $u \in \mathbb{F}_2^n$,*

$$\bigoplus_{x \in \mathbb{X}} \pi_u(x) = \begin{cases} \textit{unknown}, & \textit{if there exists } \kappa \in \mathbb{K} \textit{ such that } u \succeq \kappa, \\ 1, & \textit{if there exists } \lambda \in \mathbb{L} \textit{ such that } u = \lambda, \\ 0, & \textit{otherwise.} \end{cases}$$

Here, $\mathbb{K}$ characterizes monomials whose sums are unknown, whereas $\mathbb{L}$ identifies monomials whose sums are equal to one. All remaining monomials are guaranteed to have zero sum.

Monomial Prediction Hu *et al.* [10] introduced monomial prediction as an algebraic technique for tracing the contributions of ANF monomials to Boolean-function sums. For the purposes of this work, we use the following basic monomial-level characterization, which expresses an integral sum as the XOR of the corresponding monomial sums.

Definition 3 (Monomial-Level Characterization). *Let $\mathbb{X}$ be a multiset over $\mathbb{F}_2^n$, and let $f(x) = \bigoplus_{u \in \mathbb{F}_2^n} a_u x^u$ be the* ANF *of f. Then*

$$\bigoplus_{x \in \mathbb{X}} f(x) = \bigoplus_{u \in \mathbb{F}_2^n} a_u \left(\bigoplus_{x \in \mathbb{X}} x^u \right).$$

3 Separable Structure and Candidate Spaces

This section analyzes the structure of key-independent generalized integral properties and constructs the candidate spaces used in the subsequent search methods. We first distinguish separably generated properties, which admit decompositions into separable terms whose individual integral sums are key-independent, from residual-cancellation properties, whose key-independent overall sums necessarily rely on cancellations among key-dependent residual sums of different separable components. We then examine how representative existing approaches fit into this classification. Motivated by this structural characterization, our search focuses on balanced separable terms within structured candidate spaces and models their plaintext-side and ciphertext-side components separately. Specifically, we construct a degree-bounded Boolean candidate space on the ciphertext side and a structured indicator-based candidate space on the plaintext side.

3.1 Separable Structure of Generalized Integral Properties

Beyne and Verbauwhede introduced the notion of a generalized integral property, providing a unified formal framework for a broad class of integral properties [3].

Definition 4 (Generalized Integral Property [3]). *Let $E_k : \mathbb{F}_2^n \to \mathbb{F}_2^n$ be an n-bit block cipher parameterized by $k \in \mathbb{F}_2^m$. A Boolean function $G : \mathbb{F}_2^n \times \mathbb{F}_2^n \to \mathbb{F}_2$ defines a generalized integral property of E_k, whose evaluation is*

$$\bigoplus_{x \in \mathbb{F}_2^n} G(x, E_k(x)).$$

For cryptanalytic purposes, we primarily consider generalized integral properties with key-independent evaluations. Accordingly, for a keyed cipher E_k, we focus on Boolean functions G such that

$$\bigoplus_{x \in \mathbb{F}_2^n} G(x, E_k(x)) = c, \quad c \in \mathbb{F}_2,$$

for every key k, where c is key-independent. The special case $c = 0$ corresponds to a balanced, or zero-sum, generalized integral property.

To examine the structure of such properties, we distinguish separable contributions from more general joint plaintext–ciphertext Boolean functions. A Boolean function of the form

$$h(x)g(y),$$

where $h, g : \mathbb{F}_2^n \to \mathbb{F}_2$, is called a *separable term*. After substituting $y = E_k(x)$, its contribution to the generalized integral evaluation is

$$\bigoplus_{x \in \mathbb{F}_2^n} h(x)g(E_k(x)).$$

A separable term $h(x)g(E_k(x))$ is said to have a key-independent integral sum if

$$\bigoplus_{x \in \mathbb{F}_2^n} h(x)g(E_k(x)) = c, \quad c \in \mathbb{F}_2,$$

for every key k, where c is key-independent. The special case $c = 0$ is called a *balanced separable term*. We classify key-independent generalized integral properties according to whether they admit a decomposition into separable terms with individually key-independent integral sums.

Definition 5 (Two Classes of Key-Independent Generalized Integral Properties). *Let G define a generalized integral property of E_k whose evaluation is independent of the key. The property defined by G is called* separably generated *if there exist an integer $N \geq 1$ and Boolean functions*

$$h_t : \mathbb{F}_2^n \to \mathbb{F}_2, \quad g_t : \mathbb{F}_2^n \to \mathbb{F}_2, \quad 1 \leq t \leq N,$$

such that

$$G(x, y) = \bigoplus_{t=1}^{N} h_t(x)g_t(y),$$

and the integral sum of every separable term is key-independent, i.e.,

$$\bigoplus_{x \in \mathbb{F}_2^n} h_t(x)g_t(E_k(x)) = c_t, \quad c_t \in \mathbb{F}_2,$$

for every key k, where c_t is key-independent. We denote the set of all such properties by $\mathcal{G}_{\text{sep}}$. Otherwise, the property defined by G is called a residual-cancellation property, and we denote the set of all such properties by $\mathcal{G}_{\text{rc}}$. If $\mathcal{G}_{\text{ki}}$ denotes the set of generalized integral properties whose evaluations are independent of the key, then

$$\mathcal{G}_{\text{ki}} = \mathcal{G}_{\text{sep}} \cup \mathcal{G}_{\text{rc}}, \quad \mathcal{G}_{\text{sep}} \cap \mathcal{G}_{\text{rc}} = \varnothing.$$

The distinction is not whether algebraic cancellations occur within the evaluation of an individual separable term; such cancellations may occur in both

classes. Rather, it is whether cancellation among key-dependent residual evaluations of different separable components is necessary to obtain the final key-independent sum.

Since products of plaintext-side and ciphertext-side Boolean functions span the full Boolean function space on $\mathbb{F}_2^n \times \mathbb{F}_2^n$, every generalized Boolean function admits a decomposition into separable terms. The distinction is therefore not the existence of a separable decomposition, but whether there exists one in which the integral sum of every separable term is individually key-independent.

Classical word- and bit-based integral distinguishers, representative division-property-based distinguishers and their common refinements, and mapping-based distinguishers with explicit plaintext-side or ciphertext-side mappings are naturally expressible in terms of separable terms with key-independent integral sums and therefore fall within $\mathcal{G}_{\mathrm{sep}}$.

The algebraic transition matrix framework of Beyne and Verbauwhede [3] is more general. In this framework, the final key-independent relation may be obtained by combining several separable components whose individual integral sums still depend on the key. These key-dependent residual sums may cancel only after the components are combined. Therefore, the framework is, in principle, expressive enough to represent residual-cancellation properties in $\mathcal{G}_{\mathrm{rc}}$.

For the eight remaining generalized integral properties of `PRESENT` reported by Beyne and Verbauwhede, however, we find that residual cancellation among key-dependent residuals is not necessary. Each of these properties admits an explicit decomposition into separable terms with individually key-independent integral sums and therefore belongs to $\mathcal{G}_{\mathrm{sep}}$. The following example illustrates the decomposition for one of these properties.

Example 1 (A Generalized Integral Property for `PRESENT`*).* Following the notation of Beyne and Verbauwhede, write $F = E_k$ for the keyed `PRESENT` permutation, with F^v defined as in the algebraic transition matrix formulation. Following their reported `PRESENT` properties, consider

$$\mathbb{X}_5 = \{x \in \mathbb{F}_2^n : x_5 = 0\},\ \mathbb{X}_9 = \{x \in \mathbb{F}_2^n : x_9 = 0\},\ \mathbb{X}_{5,9} = \{x \in \mathbb{F}_2^n : x_5 + x_9 = 0\}.$$

Consider the key-independent generalized integral property

$$\bigoplus_{x \in \mathbb{X}_5} F^{e_5}(x) \oplus \bigoplus_{x \in \mathbb{X}_9} F^{e_{13}}(x) = c,$$

where $c \in \mathbb{F}_2$ is independent of the key. Its original form does not immediately exhibit a decomposition into separable terms with individually key-independent integral sums. By the indicator-function identity

$$\mathbf{1}_{\mathbb{X}_9} = \mathbf{1}_{\mathbb{X}_5} \oplus \mathbf{1}_{\mathbb{X}_{5,9}} \oplus 1,$$

we obtain

$$\begin{aligned} &\bigoplus_{x \in \mathbb{X}_5} F^{e_5}(x) \oplus \bigoplus_{x \in \mathbb{X}_9} F^{e_{13}}(x) \\ &= \bigoplus_{x \in \mathbb{X}_5} (F^{e_5}(x) \oplus F^{e_{13}}(x)) \oplus \bigoplus_{x \in \mathbb{X}_{5,9}} F^{e_{13}}(x) \oplus \bigoplus_{x \in \mathbb{F}_2^n} F^{e_{13}}(x). \end{aligned}$$

The first two terms are key-independent separable terms already contained in the PRESENT results of Beyne and Verbauwhede. The last term is a trivial zero-sum separable term, since F is a permutation and therefore $\bigoplus_{x\in\mathbb{F}_2^n} F^{e_{13}}(x) = 0$. Hence, the original generalized integral property belongs to $\mathcal{G}_{\mathrm{sep}}$.

The other seven PRESENT properties reported in [3] admit analogous decompositions, and therefore all eight belong to $\mathcal{G}_{\mathrm{sep}}$. Although the broader algebraic transition matrix framework is, in principle, expressive enough to represent residual-cancellation properties, whether such properties can actually be found remains an interesting direction for future work.

Searching for $\mathcal{G}_{\mathrm{rc}}$ is fundamentally different from searching for key-independent separable terms. For a separable term whose evaluation is not key-independent, the corresponding integral leaves a key-dependent residual

$$\bigoplus_{x\in\mathbb{F}_2^n} h_t(x)g_t(E_k(x)) = c_t(k).$$

A residual-cancellation property requires several such residual functions to cancel so that

$$\bigoplus_{t=1}^{N} c_t(k) = c, \quad c \in \mathbb{F}_2,$$

for every key k, where c is independent of the key. Hence, searching for $\mathcal{G}_{\mathrm{rc}}$ requires recovering the key-dependent residual functions themselves and detecting cancellations among them, rather than merely determining whether individual separable terms have key-independent evaluations. In contrast, properties in $\mathcal{G}_{\mathrm{sep}}$ can be constructed from individually key-independent separable terms.

The above classification is structural rather than algorithmic. Our search methods do not attempt to enumerate all properties in $\mathcal{G}_{\mathrm{sep}}$. Instead, we focus on key-independent separable terms within structured plaintext-side and ciphertext-side candidate spaces. Each recovered term therefore defines an element of $\mathcal{G}_{\mathrm{sep}}$, while XOR combinations of such terms may generate more general separably generated properties. Specifically, we search for balanced separable terms satisfying

$$\bigoplus_{x\in\mathbb{F}_2^n} h(x)g(E_k(x)) = 0$$

for every key k. Here, h determines the plaintext-side support over which the integral is evaluated, whereas g specifies the Boolean mapping tested on the ciphertext side. We next model these two components separately, starting with the ciphertext-side mapping g.

3.2 Modeling Ciphertext-Side Mappings

Todo and Morii [18] pointed out that coordinate-wise division-property analysis does not fully capture integral properties arising from combinations of ciphertext bits. In particular, even when the sums of several ciphertext bits are individually

classified as *unknown* over a plaintext structure $\mathbb{X}$, an appropriate combination of these bits may still be balanced. For example,

$$\bigoplus_{x\in\mathbb{X}} \big(E_k(x)_{j_1} \oplus E_k(x)_{j_2}\big) = 0$$

may hold although the two component sums are individually unknown. Equivalently, this property corresponds to the ciphertext-side mapping

$$g(y) = (e_{j_1} \cdot y) \oplus (e_{j_2} \cdot y).$$

More generally, such properties can arise from cancellations among unknown monomials in the induced ANF, which motivates searching over ciphertext-side Boolean mappings rather than individual ciphertext bits.

From Single Bit to Ciphertext-Side Boolean Mapping We formulate the ciphertext-side search problem directly at the level of the induced ANF representation. Let $E : \mathbb{F}_2^n \times \mathbb{F}_2^m \to \mathbb{F}_2^n$ denote the keyed encryption function, and write

$$E(x,k) = (f_0(x,k), \ldots, f_{n-1}(x,k)),$$

where each $f_i : \mathbb{F}_2^{n+m} \to \mathbb{F}_2$ represents the i-th ciphertext bit as a Boolean function of the joint variables (x,k). For a fixed key k, we write $E_k(x) = E(x,k)$. The cancellation analysis is carried out at the monomial level over the joint variables (x,k); in particular, x_1k_1 and x_1k_2 are distinct monomials and therefore cannot cancel each other.

Given a plaintext structure $\mathbb{X} \subseteq \mathbb{F}_2^n$, our goal is to find a nonzero ciphertext-side Boolean mapping $g : \mathbb{F}_2^n \to \mathbb{F}_2$ such that

$$\bigoplus_{x\in\mathbb{X}} g(f_0(x,k), \ldots, f_{n-1}(x,k)) = 0$$

for all keys $k \in \mathbb{F}_2^m$. Equivalently, this amounts to searching for a balanced separable term $\mathbf{1}_{\mathbb{X}}(x)g(y)$ within the ciphertext-side candidate space. To express this search problem algebraically, we write the induced expression in ANF form over the joint variables (x,k):

$$g(f_0(x,k), \ldots, f_{n-1}(x,k)) = \bigoplus_{(u,v)\in\mathbb{F}_2^n\times\mathbb{F}_2^m} a_{u,v} x^u k^v.$$

For the ciphertext-side search considered here, the chosen integral-analysis backend provides a conservative classification of these monomials as either *balanced* or *unknown* with respect to $\mathbb{X}$. Our task is to choose g so that no monomial classified as *unknown* survives in the induced ANF. If this holds, every remaining monomial is certified as balanced under the adopted classification, and hence the resulting ciphertext-side mapping is zero-sum over $\mathbb{X}$.

Degree-Bounded Candidate Space Since the full Boolean function space on the ciphertext side is too large for direct search, we restrict g to the span of low-degree ciphertext monomials. Let $y = (y_0, \ldots, y_{n-1})$ denote the ciphertext variables. For a degree bound d, define

$$\mathcal{B}_d = \{\beta \in \mathbb{F}_2^n : 1 \le \mathrm{wt}(\beta) \le d\}, \quad |\mathcal{B}_d| = \sum_{i=1}^{d} \binom{n}{i}.$$

We then write

$$g(y) = \bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta y^\beta, \quad y^\beta = \prod_{i=0}^{n-1} y_i^{\beta_i},$$

where $\gamma_\beta \in \mathbb{F}_2$. Thus, the candidate space is

$$\mathrm{span}_{\mathbb{F}_2}\{y^\beta : \beta \in \mathcal{B}_d\},$$

which parameterizes all nonconstant Boolean mappings of algebraic degree at most d.

Ciphertext-Side Mapping Matrix For each $\beta \in \mathcal{B}_d$, define

$$F_\beta(x, k) = \prod_{i=0}^{n-1} f_i(x, k)^{\beta_i},$$

which is obtained from the ciphertext monomial y^β by substituting $y_i = f_i(x, k)$. We expand F_β in ANF over the joint variables (x, k) and classify the resulting monomials with respect to the plaintext structure $\mathbb{X}$ using a chosen integral-analysis backend. For the present ciphertext-side model, the backend is required to conservatively distinguish monomials certified as balanced from those that must remain unknown. The resulting information is encoded in a matrix that records, for each candidate ciphertext monomial y^β, which monomials remain unknown.

Definition 6 (Ciphertext-Side Mapping Matrix). *For each $\beta \in \mathcal{B}_d$, let $\mathcal{U}_\beta$ denote the set of monomials that appear with coefficient 1 in the ANF of F_β and are classified as unknown with respect to $\mathbb{X}$. Let*

$$\mathcal{M} = \bigcup_{\beta \in \mathcal{B}_d} \mathcal{U}_\beta, \quad M_{\mathcal{B}_d} \in \mathbb{F}_2^{|\mathcal{B}_d| \times |\mathcal{M}|},$$

where the rows of $M_{\mathcal{B}_d}$ are indexed by $\beta \in \mathcal{B}_d$, the columns are indexed by $\mu \in \mathcal{M}$, and

$$(M_{\mathcal{B}_d})_{\beta,\mu} = \begin{cases} 1, & \text{if } \mu \in \mathcal{U}_\beta, \\ 0, & \text{otherwise.} \end{cases}$$

A ciphertext-side mapping in the degree-bounded candidate space is represented by

$$\gamma = (\gamma_\beta)_{\beta \in \mathcal{B}_d} \in \mathbb{F}_2^{|\mathcal{B}_d|}.$$

The product $\gamma M_{\mathcal{B}_d}$ records the parity with which each $\mu \in \mathcal{M}$ appears in

$$\bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta F_\beta(x, k).$$

Hence,

$$\gamma M_{\mathcal{B}_d} = 0$$

if and only if every recorded unknown monomial cancels in the resulting ANF.

Theorem 1 (Monomial-Cancellation Criterion for Ciphertext-Side Mappings). *Let $\gamma \in \mathbb{F}_2^{|\mathcal{B}_d|}$, and let g be the ciphertext-side mapping represented by γ. If*

$$\gamma M_{\mathcal{B}_d} = 0,$$

then all unknown monomials recorded in

$$g(f_0(x, k), \ldots, f_{n-1}(x, k)) = \bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta F_\beta(x, k)$$

cancel at the ANF level. Consequently, under the adopted monomial classification,

$$\bigoplus_{x \in \mathbb{X}} g(E_k(x)) = 0$$

for every key k. Hence, $\mathbf{1}_{\mathbb{X}}(x)g(y)$ defines a balanced separable term and therefore belongs to $\mathcal{G}_{\text{sep}}$.

A proof is provided in Appendix A.1. Therefore, the left kernel of $M_{\mathcal{B}_d}$ is exactly the solution space of the induced monomial-cancellation system and characterizes all ciphertext-side mappings certified by this criterion. Its dimension is

$$\dim(\ker_{\text{left}}(M_{\mathcal{B}_d})) = |\mathcal{B}_d| - \operatorname{rank}(M_{\mathcal{B}_d}).$$

Our ciphertext-side modeling is closely related to the matrix-based approaches of Hebborn *et al.* [8,9], as both encode monomial cancellation through linear-algebraic conditions. Hebborn *et al.* use rank-based criteria to derive degree lower bounds or to exclude certain integral distinguishers under the independent-round-key assumption. In contrast, $M_{\mathcal{B}_d}$ is used directly as a search representation: its left kernel characterizes the mappings in the chosen degree-bounded candidate space for which all recorded unknown monomials cancel. Moreover, whereas their setting considers linear output mappings, our candidate space includes degree-bounded nonlinear Boolean mappings.

3.3 Modeling Plaintext-Side Mappings

For plaintext-side mapping, we combine multiple plaintext structures. Since different structures may induce different propagation behaviors, plaintext-side mappings primarily act by modifying the input summation domain, rather than through the ciphertext-side unknown-monomial cancellation modeled. This viewpoint is related in spirit to dynamic cube attacks [6] and conditional cube attacks [12], which also manipulate public input variables to control algebraic behavior, although their objectives and constructions are different. To keep as many plaintext bits active as possible, we use as basic plaintext structures those obtained by fixing exactly one plaintext bit while leaving the other $n-1$ bits active. For an n-bit block cipher, this gives n basic structures, one for each bit position.

In the separable framework, plaintext-side mappings are in general Boolean functions $h : \mathbb{F}_2^n \to \mathbb{F}_2$. Here, we restrict h to a structured candidate space generated by the indicator functions of the basic affine plaintext structures $\mathbb{X}_0, \ldots, \mathbb{X}_{n-1}$ defined in the preliminaries.

Definition 7 (Plaintext-Side Mapping as a Linear Combination of Indicators). *For $\sigma = (\sigma_0, \ldots, \sigma_{n-1}) \in \mathbb{F}_2^n$, define*

$$h_\sigma(x) = \bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}(x).$$

Accordingly, the plaintext-side candidate space is

$$\mathrm{span}_{\mathbb{F}_2}\{\mathbf{1}_{\mathbb{X}_0}, \ldots, \mathbf{1}_{\mathbb{X}_{n-1}}\}.$$

Since $h_\sigma(x) \in \mathbb{F}_2$, its support is the set of points for which $h_\sigma(x) = 1$. Thus, the support of h_σ is exactly the symmetric difference of the selected plaintext structures. We denote this induced plaintext structure by $\mathbb{X}_\sigma$.

Definition 8 (Symmetric Difference of Plaintext Structures). *For plaintext structures $\mathbb{X}_0, \ldots, \mathbb{X}_{n-1} \subseteq \mathbb{F}_2^n$ and a selection vector $\sigma = (\sigma_0, \ldots, \sigma_{n-1}) \in \mathbb{F}_2^n$, define*

$$\mathbb{X}_\sigma = \triangle_{i:\sigma_i=1} \mathbb{X}_i, \quad \mathbf{1}_{\mathbb{X}_\sigma}(x) = \bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}(x).$$

Because each basic indicator $\mathbf{1}_{\mathbb{X}_i}$ is an affine Boolean function, any nontrivial XOR of these indicators remains affine. The following theorem makes the resulting plaintext structure explicit.

Theorem 2 (Symmetric Differences Induce Affine Linear Constraints). *Let $\mathbb{X}_i = \{x \in \mathbb{F}_2^n : x_i = a_i\}$ for $0 \le i < n$, and let $\sigma = (\sigma_0, \ldots, \sigma_{n-1}) \in \mathbb{F}_2^n$ be nonzero. Then*

$$\mathbb{X}_\sigma = \triangle_{i:\sigma_i=1} \mathbb{X}_i = \left\{ x \in \mathbb{F}_2^n : \bigoplus_{i=0}^{n-1} \sigma_i x_i = 1 \oplus \bigoplus_{i=0}^{n-1} \sigma_i a_i \oplus \bigoplus_{i=0}^{n-1} \sigma_i \right\}.$$

Hence, $\mathbb{X}_\sigma$ is an affine hyperplane of $\mathbb{F}_2^n$, and therefore $|\mathbb{X}_\sigma| = 2^{n-1}$.

A proof is provided in Appendix A.2. Thus, every nonzero plaintext-side combination in the chosen candidate space preserves the affine-hyperplane structure and the corresponding 2^{n-1}-point data complexity.

Linearity of Integral Sums under Plaintext-Side Combinations The indicator representation of $\mathbb{X}_\sigma$ also implies that integral sums over the combined plaintext structure are linear in the selected basic structures.

Theorem 3 (Linearity of Integral Sums under Symmetric Difference). *For any plaintext structures $\mathbb{X}_0, \ldots, \mathbb{X}_{n-1}$, any selection vector $\sigma \in \mathbb{F}_2^n$, and any Boolean function $F : \mathbb{F}_2^n \to \mathbb{F}_2$, we have*

$$\bigoplus_{x \in \mathbb{X}_\sigma} F(x) = \bigoplus_{i=0}^{n-1} \sigma_i \left(\bigoplus_{x \in \mathbb{X}_i} F(x) \right),$$

where

$$\mathbf{1}_{\mathbb{X}_\sigma} = \bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}.$$

In particular, the identity applies to $F(x) = g(E_k(x))$ for any ciphertext-side Boolean mapping g and any key k.

A proof is provided in Appendix A.3. Theorem 3 shows that plaintext-side combinations preserve linearity at the level of integral sums for arbitrary Boolean functions, including mappings of the form $g(E_k(x))$. Although restricting the plaintext domain to an affine hyperplane may change the ANF obtained by an explicit reparameterization, such a reparameterization is unnecessary here because the combination is handled directly through the indicator functions of the summation domains. In principle, the resulting behavior could be analyzed at the monomial level, but repeating such an analysis for many candidate selection vectors σ quickly becomes expensive. We therefore use the bit-based division property with three subsets to encode the effect of plaintext-side combinations more efficiently. Its deterministic-one contributions combine by parity and may cancel, whereas unknown contributions must be handled conservatively.

Three-Subset Combination Rule Suppose that each basic plaintext structure $\mathbb{X}_i$ is represented by a three-subset division-property description $(\mathbb{K}_i, \mathbb{L}_i)$. Here, $\mathbb{L}_i$ identifies monomials whose sums are deterministically equal to 1, while $\mathbb{K}_i$ characterizes unknown monomial sums: a monomial indexed by u is classified as *unknown* whenever $u \succeq \kappa$ for some $\kappa \in \mathbb{K}_i$. Under plaintext-side parity combinations, the $\mathbb{L}$-part combines exactly by XOR, whereas the $\mathbb{K}$-part is handled conservatively by union.

Definition 9 (Three-Subset Combination under Plaintext-Side Mappings). *Let $\sigma \in \mathbb{F}_2^n$ be a selection vector. We define the induced combined three-subset description by*

$$\mathbb{L}_\sigma = \triangle_{i:\sigma_i=1} \mathbb{L}_i, \quad \mathbb{K}_\sigma = \bigcup_{i:\sigma_i=1} \mathbb{K}_i.$$

Thus, deterministically-one contributions may cancel under parity, while unknown contributions are retained whenever they occur in any selected basic structure.

Plaintext-Side Combination Matrices To encode the above rule for subsequent search, we construct two matrices from the three-subset descriptions of the basic plaintext structures. Since the plaintext-side candidate space is generated by $\mathbb{X}_0, \ldots, \mathbb{X}_{n-1}$, the rows of both matrices are indexed by these n structures.

Let $\mathcal{L} = \{l_1, \ldots, l_{|\mathcal{L}|}\}$ be the set of all $\mathbb{L}$-type source vectors appearing among the basic structures, and define

$$M_{\mathcal{L}} \in \mathbb{F}_2^{n \times |\mathcal{L}|}, \quad (M_{\mathcal{L}})_{i,r} = \begin{cases} 1, & \text{if } l_r \in \mathbb{L}_i, \\ 0, & \text{otherwise.} \end{cases}$$

The i-th row records the deterministically-one source vectors associated with $\mathbb{X}_i$. For a selection vector σ, the combined $\mathbb{L}$-part is represented by

$$\sigma M_{\mathcal{L}}.$$

Similarly, let $\mathcal{K} = \{k_1, \ldots, k_{|\mathcal{K}|}\}$ be the set of all $\mathbb{K}$-type source vectors appearing among the basic structures, and define

$$M_{\mathcal{K}} \in \mathbb{F}_2^{n \times |\mathcal{K}|}, \quad (M_{\mathcal{K}})_{i,r} = \begin{cases} 1, & \text{if } k_r \in \mathbb{K}_i, \\ 0, & \text{otherwise.} \end{cases}$$

The i-th row records the unknown-source vectors associated with $\mathbb{X}_i$. Unlike the $\mathbb{L}$-part, the $\mathbb{K}$-part is not combined linearly. A vector k_r remains in $\mathbb{K}_\sigma$ whenever it appears in at least one selected row, i.e.,

$$\bigvee_{i:\sigma_i=1} (M_{\mathcal{K}})_{i,r} = 1.$$

The pair $(M_{\mathcal{L}}, M_{\mathcal{K}})$ therefore provides a structured encoding of plaintext-side combinations. The $\mathbb{L}$-part may cancel under XOR and suppress some input-side sources, whereas the $\mathbb{K}$-part conservatively retains sources that may remain active. Because the two parts obey different combination rules, the plaintext-side search does not reduce to a single linear system. Whether a surviving source is relevant further depends on whether it can propagate to an unknown ciphertext-side monomial. This interaction motivates the unified MBM framework introduced next.

4 Search Frameworks and Solving Strategies

This section presents two solving strategies for searching balanced separable terms within the candidate spaces constructed in Section 3. Both strategies avoid exhaustive enumeration of candidate mappings by representing the search

through structured matrix-based models. The first is a unified framework that jointly searches plaintext-side and ciphertext-side mappings and reduces the resulting constraints to an MILP formulation. The second is a specialized strategy for a fixed plaintext structure, where the ciphertext-side search is reduced to a homogeneous linear system and solved by Gaussian elimination.

4.1 A Unified MBM Search Framework

Sections 3.2 and 3.3 provide complementary models for the two sides of the separable search. The ciphertext side captures cancellation of unknown monomials, whereas the plaintext side captures suppression of their propagation sources. The sufficient condition used by the MBM model is that no surviving plaintext-side source is connected to any surviving ciphertext-side unknown monomial. This motivates the unified MBM (*Matrix–Bipartite Graph–Matrix*) framework, which links the two matrix components through a middle bipartite propagation graph. Fig. 1 illustrates the resulting structure.

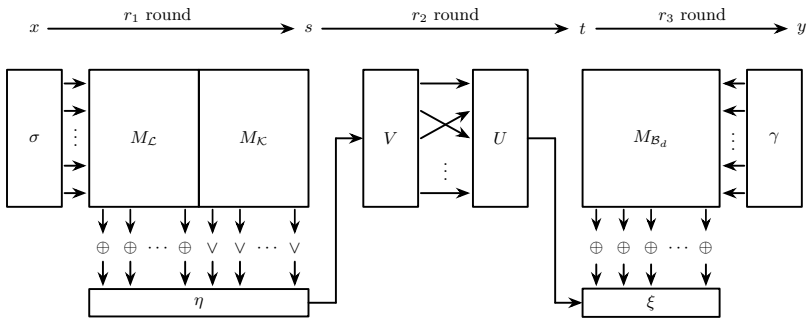


Fig. 1. Overview of the MBM search framework

The plaintext side is parameterized by a selection vector

$$\sigma = (\sigma_0, \ldots, \sigma_{n-1}) \in \{0,1\}^n,$$

where σ_i indicates whether the basic plaintext structure $\mathbb{X}_i$ is selected. The ciphertext side is parameterized by a coefficient vector

$$\gamma \in \{0,1\}^{|\mathcal{B}_d|},$$

which specifies a mapping in the degree-bounded candidate space introduced in Section 3.2. Thus, an MBM search instance is determined by the pair (σ, γ).

Static Components of the MBM Model To connect the plaintext-side and ciphertext-side components, we decompose the r-round keyed encryption E_k into

three consecutive parts:

$$E_k = E_k^{(r_3)} \circ E_k^{(r_2)} \circ E_k^{(r_1)}, \quad r_1 + r_2 + r_3 = r,$$

and write

$$s = E_k^{(r_1)}(x), \quad t = E_k^{(r_2)}(s), \quad y = E_k^{(r_3)}(t).$$

The first and third parts are used to construct the plaintext-side and ciphertext-side matrix components, respectively, while the middle part determines the propagation relation between them. Accordingly, the MBM model consists of three precomputed components: the plaintext-side membership matrices $M_{\mathcal{K}}$ and $M_{\mathcal{L}}$, the ciphertext-side mapping matrix $M_{\mathcal{B}_d}$, and the middle bipartite propagation graph $\mathcal{G}$. The decision variables are σ and γ; all other quantities introduced below are determined by this pair.

Plaintext-Side Component. Propagation of each basic plaintext structure $\mathbb{X}_i$ through the first r_1 rounds yields a three-subset description $(\mathbb{K}_i, \mathbb{L}_i)$. Collecting the corresponding source vectors as

$$\mathcal{K} = \bigcup_{i=0}^{n-1} \mathbb{K}_i = \{k_1, \ldots, k_{p_{\mathcal{K}}}\}, \quad \mathcal{L} = \bigcup_{i=0}^{n-1} \mathbb{L}_i = \{l_1, \ldots, l_{p_{\mathcal{L}}}\},$$

we obtain the plaintext-side membership matrices

$$M_{\mathcal{K}} \in \{0,1\}^{n \times p_{\mathcal{K}}}, \quad (M_{\mathcal{K}})_{i,a} = \begin{cases} 1, & \text{if } k_a \in \mathbb{K}_i, \\ 0, & \text{otherwise,} \end{cases}$$

$$M_{\mathcal{L}} \in \{0,1\}^{n \times p_{\mathcal{L}}}, \quad (M_{\mathcal{L}})_{i,a} = \begin{cases} 1, & \text{if } l_a \in \mathbb{L}_i, \\ 0, & \text{otherwise.} \end{cases}$$

Ciphertext-Side Component. We apply the ciphertext-side matrix construction of Section 3.2 to the degree-bounded candidate space $\mathcal{B}_d$ at the interface t. After expanding the ciphertext-side basis monomials backward through the last r_3 rounds, we obtain an interface-specific matrix

$$M_{\mathcal{B}_d} \in \{0,1\}^{|\mathcal{B}_d| \times q},$$

whose columns are indexed by the ciphertext-side monomial objects

$$U = \{u_1, \ldots, u_q\}.$$

Middle Propagation Component. The propagation relation between the plaintext-side and ciphertext-side components is represented by a bipartite graph

$$\mathcal{G} = (V \cup U, \mathcal{E}), \quad V = \mathcal{K} \cup \mathcal{L} = \{v_1, \ldots, v_p\}.$$

A vertex $v_a \in V$ is connected to $u_b \in U$ if the plaintext-side source represented by v_a may propagate through the middle r_2 rounds to a contribution that makes

the monomial u_b unknown. This edge relation is determined by a chosen propagation backend. The graph is constructed conservatively: an edge is omitted only when the backend proves that the corresponding propagation is impossible. Thus, the approximation may discard valid solutions, but it cannot certify an invalid distinguisher. For the subsequent formulation, we equivalently represent $\mathcal{G}$ by its adjacency matrix

$$P \in \{0,1\}^{p\times q}, \quad P_{a,b} = 1 \iff (v_a, u_b) \in \mathcal{E}.$$

Surviving Objects Given a plaintext-side selector $\sigma = (\sigma_0, \ldots, \sigma_{n-1}) \in \{0,1\}^n$, the surviving $\mathcal{K}$-type objects follow the union rule, whereas the surviving $\mathcal{L}$-type objects follow the XOR rule:

$$\rho_a(\sigma) = \bigvee_{i=0}^{n-1} \big(\sigma_i \wedge (M_{\mathcal{K}})_{i,a}\big), \quad 1 \le a \le p_{\mathcal{K}},$$
$$\lambda_a(\sigma) = \bigoplus_{i=0}^{n-1} \big(\sigma_i \wedge (M_{\mathcal{L}})_{i,a}\big), \quad 1 \le a \le p_{\mathcal{L}}.$$

These indicators are lifted to the common vertex set

$$V = \mathcal{K} \cup \mathcal{L} = \{v_1, \ldots, v_p\}$$

by

$$\eta_a(\sigma) = \left(\bigvee_{h:\, v_a = k_h \in \mathcal{K}} \rho_h(\sigma)\right) \vee \left(\bigvee_{h:\, v_a = l_h \in \mathcal{L}} \lambda_h(\sigma)\right),$$

where an empty disjunction is interpreted as 0.

On the ciphertext side, for $\gamma \in \{0,1\}^{|\mathcal{B}_d|}$, a monomial object u_b survives exactly when it appears with odd parity among the selected basis monomials:

$$\xi_b(\gamma) = \bigoplus_{\beta \in \mathcal{B}_d} \big(\gamma_\beta \wedge (M_{\mathcal{B}_d})_{\beta,b}\big), \quad 1 \le b \le q.$$

Thus, $\eta_a(\sigma) = 1$ indicates that the plaintext-side source v_a survives the combination induced by σ, whereas $\xi_b(\gamma) = 1$ indicates that the ciphertext-side monomial u_b remains uncanceled.

Danger Predicate and Unified Criterion The interaction between the surviving plaintext-side sources and ciphertext-side monomials is captured by the following predicate.

Definition 10 (Danger Predicate in the MBM Model). *For a plaintext-side selector $\sigma \in \{0,1\}^n$ and a ciphertext-side coefficient vector $\gamma \in \{0,1\}^{|\mathcal{B}_d|}$, define*

$$\Phi(\sigma, \gamma) = \bigvee_{a=1}^{p} \bigvee_{b=1}^{q} \big(\eta_a(\sigma) \wedge P_{a,b} \wedge \xi_b(\gamma)\big).$$

Thus, $\Phi(\sigma, \gamma) = 1$ if and only if the MBM graph $\mathcal{G}$ contains an edge between a surviving plaintext-side source and a surviving ciphertext-side monomial.

The following criterion gives a sufficient condition for a balanced separable term in the MBM model.

Theorem 4 (Unified Criterion under the MBM Model). *Let $\sigma \neq 0$ be a plaintext-side selector and let $\gamma \neq 0$ be a ciphertext-side coefficient vector over $\mathcal{B}_d$. If*

$$\Phi(\sigma, \gamma) = 0,$$

then

$$h_\sigma(x) g_\gamma(y)$$

defines a balanced separable term under the MBM criterion and therefore an element of $\mathcal{G}_{\text{sep}}$.

A proof is provided in Appendix A.4. A concrete `PRESENT` example illustrating how the plaintext-side L-part component affects the danger predicate and the resulting MBM certificate is given in Appendix B.

Therefore, the MBM search reduces to finding nonzero σ and γ such that

$$\Phi(\sigma, \gamma) = 0.$$

Together with the definitions of ρ, λ, η, and ξ, this gives a Boolean constraint system that can be encoded as an MILP instance. The complete joint-search procedure, including construction of the static components and the final constraint-solving step, is given in Algorithm 1 in Appendix C.

4.2 Gaussian Elimination for Ciphertext-Side Low-Degree Mappings

Under a fixed plaintext structure $\mathbb{X}$, the ciphertext side can be treated independently. This is particularly useful because the ciphertext-side model in Section 3.2 already supports a degree-bounded candidate space rather than only linear output combinations. We therefore consider how to recover all ciphertext-side mappings certified by the induced cancellation system for a given plaintext structure $\mathbb{X}$. Each such mapping g corresponds to a balanced separable term $\mathbf{1}_{\mathbb{X}}(x) g(y)$.

We decompose the r-round keyed encryption as

$$E_k = E_k^{(r_2)} \circ E_k^{(r_1)}, \quad r_1 + r_2 = r,$$

and write

$$t = E_k^{(r_1)}(x), \quad y = E_k(x) = E_k^{(r_2)}(t).$$

Using t as the analysis interface, the ciphertext is represented as

$$y = (y_0, \ldots, y_{n-1}) = (f_0(t, k), \ldots, f_{n-1}(t, k)),$$

where $f_j(t,k)$ denotes the j-th ciphertext coordinate function.

We work in the degree-bounded ciphertext-side candidate space $\mathcal{B}_d$ introduced in Section 3.2. For each $\beta \in \mathcal{B}_d$, the corresponding substituted monomial $F_\beta(t,k)$ is expanded over the interface variables (t,k). The resulting monomials are classified using the propagation information induced by $\mathbb{X}$ through the first r_1 rounds, yielding the ciphertext-side mapping matrix $M_{\mathcal{B}_d}$. By Theorem 1, a coefficient vector $\gamma \in \mathbb{F}_2^{|\mathcal{B}_d|}$ is certified by the induced cancellation system whenever

$$\gamma M_{\mathcal{B}_d} = 0.$$

Thus, the fixed-plaintext ciphertext-side search reduces to computing the left kernel of $M_{\mathcal{B}_d}$.

Theorem 5 (Kernel Characterization in the Fixed-Plaintext Setting). *Let $\delta = |\mathcal{B}_d| - \mathrm{rank}(M_{\mathcal{B}_d})$. The homogeneous system $\gamma M_{\mathcal{B}_d} = 0$ admits a nonzero solution if and only if $\mathrm{rank}(M_{\mathcal{B}_d}) < |\mathcal{B}_d|$, or equivalently, if and only if $\delta > 0$. In this case, the solution space has dimension δ. Let $\gamma^{(1)}, \ldots, \gamma^{(\delta)} \in \mathbb{F}_2^{|\mathcal{B}_d|}$ be a basis of the left kernel of $M_{\mathcal{B}_d}$. Then every nonzero solution can be written as*

$$\gamma = \bigoplus_{u=1}^{\delta} \alpha_u \gamma^{(u)}, \quad \alpha_u \in \{0,1\}, \quad (\alpha_1, \ldots, \alpha_\delta) \neq 0.$$

Accordingly, every ciphertext-side mapping certified by the induced monomial-cancellation system can be written as

$$g_\gamma(y) = \bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta y^\beta.$$

A proof is provided in Appendix A.5.

Theorem 5 shows that the fixed-plaintext search can be handled entirely by linear algebra. Gaussian elimination either shows that the left kernel is trivial or returns a basis that compactly represents the full solution space of the induced cancellation system. Each basis vector $\eta^{(u)}$ determines a ciphertext-side basis mapping, and all mappings certified by this criterion are obtained as nonzero $\mathbb{F}_2$-linear combinations of these basis mappings. Each resulting mapping g_γ yields a balanced separable term $\mathbf{1}_{\mathbb{X}}(x) g_\gamma(y)$ and therefore an element of $\mathcal{G}_{\mathrm{sep}}$.

Accordingly, the solving procedure consists of three steps. First, decompose the encryption into two parts and express the ciphertext coordinate functions over the intermediate state t. Second, construct $M_{\mathcal{B}_d}$ for the chosen degree-bounded candidate space using the propagation information induced by $\mathbb{X}$. Third, compute a basis of its left kernel by Gaussian elimination and recover the corresponding ciphertext-side mappings. The detailed procedure is given in Algorithm 2 in Appendix C.

At a high level, our approach is related to the algebraic transition matrix framework of Beyne and Verbauwhede [3] in that both reduce a fixed-plaintext search to a kernel computation. The distinction lies in the information encoded by the corresponding matrices. Beyne and Verbauwhede represent generalized

integral propagation through algebraic transition matrices over a chosen candidate space, whereas $M_{\mathcal{B}_d}$ records the incidence between candidate ciphertext monomials and monomials classified as unknown after substitution and propagation analysis. Its left kernel therefore characterizes the mappings in the chosen degree-bounded candidate space for which all recorded unknown monomials cancel.

5 Applications and Analysis

In this section, we evaluate the proposed search methods on `SPECK`, `SIMON`, `PRESENT`, `RECTANGLE`, and `SKINNY`. Our analysis has three objectives. First, we assess the practical effectiveness of the two search frameworks in identifying balanced plaintext- and ciphertext-side mappings across representative block ciphers. Second, using `SKINNY` as a case study, we investigate whether ciphertext-side nonlinear mappings can be exploited to improve key-recovery attacks. Third, we analyze why the underlying round-extension mechanisms arise for the `SPECK` family and for `SIMON`, but not for the other evaluated ciphers, with particular attention to the algebraic structures exposed by backward expansion.

All experiments were conducted on a platform equipped with Intel(R) Xeon(R) Gold 6226R processors running at 2.90 GHz and 256 GB RAM. Throughout this section, we use $y = (y_0, y_1, \ldots, y_{n-1})$ and $x = (x_0, x_1, \ldots, x_{n-1})$ to denote the ciphertext and plaintext, respectively. Unless stated otherwise, bit positions are indexed from left to right.

5.1 Application to `SPECK`

`SPECK` is an ARX block-cipher family [1] with a two-word state (L_i, R_i). One round updates the state as

$$L_{i+1} = ((L_i \ggg \alpha) + R_i) \oplus RK_i, \quad R_{i+1} = (R_i \lll \beta) \oplus L_{i+1},$$

where RK_i is the round key, $+$ denotes addition modulo 2^n, and $\ggg, \lll$ denote right and left rotations, respectively.

For `SPECK`, constructing the plaintext-side component of MBM is prohibitively expensive because of the modular-addition structure of its round function. In particular, the $E_k^{(r_1)}$ phase requires propagating the three-subset descriptions $(\mathbb{K}_i, \mathbb{L}_i)$ for all basic plaintext structures, and the resulting propagation sets grow rapidly even after the first round. We therefore fix the plaintext structure and restrict the search to the ciphertext side, using the specialized solver of Section 4.2.

Baseline Ciphertext-Side Search. We first consider `SPECK32`. To keep the search feasible, we restrict the ciphertext-side mapping to algebraic degree at most 1 and use one round of backward expansion ($r_2 = 1$). Under the plaintext structure `0xFFFFFF1F`, we obtain a 6-round integral distinguisher with data complexity

2^{29}. The searched linear candidate space yields six balanced mappings, all of the form $y_i \oplus y_{i+16}$, such as $y_8 \oplus y_{24}$; the complete list is given in Appendix D.2. These mappings follow from the relation

$$L_{i+1} \oplus R_{i+1} = R_i \lll \beta,$$

and are therefore closely related to balanced linear relations on the rotated right branch of the preceding round.

The direct ciphertext-side search faces two limitations. For SPECK32, extending the 6-round result by another round would require sufficiently strong balanced relations on the 6-round intermediate state. For the larger SPECK variants, fully expanding the carry terms of modular addition causes the backward ANF to grow rapidly, leading to prohibitive time and memory costs.

To address these limitations, we use three complementary refinements. Linear enhancement reuses previously certified balanced relations on intermediate states to enable an additional round of backward expansion. Abstraction enhancement replaces carry terms with independent formal symbols to control the cost of ANF expansion, while truncated enhancement restricts the candidate mapping space to selected output coordinates.

Linear Enhancement. For 7-round SPECK32, consider the ciphertext-side mapping

$$y_6 \oplus y_{13} \oplus y_{15} \oplus y_{22} \oplus y_{29}.$$

Rewriting it over the 6-round intermediate state and canceling common terms gives

$$y_6 \oplus y_{13} \oplus y_{15} \oplus y_{22} \oplus y_{29} = s_8 \oplus s_{24} \oplus k_{15}.$$

The complete derivation is given in Appendix D.3. Although the right-hand side is classified as unknown by the original backend, its nonconstant part $s_8 \oplus s_{24}$ is a previously certified 6-round balanced relation, while k_{15} is constant with respect to the plaintext summation. Hence, the 7-round mapping is certified by reducing its backward expansion to a known 6-round balanced relation.

More generally, a higher-round ciphertext-side mapping can be certified when its backward expansion reduces to a previously certified balanced relation plus a zero-sum remainder. This principle is formalized in Theorem 6 in Appendix D.1.

Abstraction Enhancement. For the larger SPECK variants, fully expanding carry terms in the backward ANF quickly becomes infeasible. We therefore replace carry terms with independent formal symbols and treat them as unknown variables in the subsequent determination step. This abstraction is conservative: it does not introduce spurious distinguishers, although it may miss cancellations that become visible only after full carry expansion.

Truncated Enhancement. We further reduce the search cost by restricting the ciphertext-side candidate space to the last 16 bits of each branch. This restriction only prunes the candidate masks and does not alter the algebraic expressions of the retained output bits. Hence, any integral property found in the truncated search space remains valid in the full ciphertext space.

Results for Larger Variants. For SPECK48, SPECK64, SPECK96, and SPECK128, abstraction enhancement is first used to obtain 6-round balanced linear relations. Linear enhancement then reuses these relations when analyzing the corresponding 7-round expressions, while truncated enhancement controls the candidate-space size. This combination yields 7-round integral distinguishers for all four variants, improving the previously reported 6-round results by one round. The detailed results are given in Appendix E, Table 3.

5.2 Applications to SIMON, PRESENT, and RECTANGLE

The SIMON, PRESENT, and RECTANGLE experiments use the unified MBM framework with $r_1 = r_3 = 1$. We set $d = 2$ for SIMON and $d = 1$ for PRESENT and RECTANGLE. For the long-round searches, we use basic plaintext structures with one fixed bit and $n - 1$ active bits. The middle verification uses the two-subset division-property model, except for SIMON32, where the three-subset model remains feasible. We report essentially distinct properties after tensor-based deduplication. The deduplication procedure is described in Appendix F.

Application to SIMON For every SIMON variant, the proposed search extends the previously reported integral distinguisher by one round, and all such improvements are obtained through nonlinear ciphertext-side mappings. No linear distinguisher is found in the searched subclass. Specifically, the round numbers increase from 14 to 15, 15 to 16, 17 to 18, 21 to 22, and 25 to 26 for SIMON32, SIMON48, SIMON64, SIMON96, and SIMON128, respectively. The search yields at least 48, 576, 704, 240, and 192 essentially distinct nonlinear distinguishers, with data complexities 2^{31}, 2^{47}, 2^{63}, 2^{95}, and 2^{127}, respectively. Representative examples are given in Appendix E, Table 4. The reported round numbers do not include the additional free extension of [20].

Application to PRESENT For PRESENT, the proposed search recovers a large set of certifiable balanced mappings at the same round depth as previous results. With $r_2 = 7$, we obtain at least 204 essentially distinct 9-round properties, all with data complexity 2^{63}. For comparison, we apply the same deduplication procedure to the results of Beyne and Verbauwhede [3]. Under the same partition $(r_1, r_2, r_3) = (1, 7, 1)$, their results also yield 204 essentially distinct properties. When combining their search results from all eight partitions, the total increases to 215 essentially distinct properties. Representative examples are given in Appendix E, Table 5.

Application to RECTANGLE For RECTANGLE, the proposed search reproduces the previously known 10-round plaintext-side mappings obtained via linearly equivalent S-boxes. It also identifies additional balanced plaintext-side mappings at the same round depth. With $r_2 = 8$, we obtain 16 such mappings, all yielding 10-round integral properties with data complexity 2^{63}. For example, $x_{16} \oplus x_{32}$

corresponds to the balanced ciphertext bit y_5, while $x_{17} \oplus x_{33}$ corresponds to y_6. Representative results are given in Appendix E, Table 6.

5.3 Application to SKINNY

SKINNY is a family of tweakable block ciphers [2]. Its internal state is a 4×4 array, and each round applies SubCells, AddConstants, AddRoundTweakey, ShiftRows, and MixColumns.

Plaintext-side search is substantially more expensive for SKINNY than for the ciphers considered above because MixColumns causes the propagated division-property sets to grow rapidly. In particular, propagating the three-subset division property from the plaintext side results in a prohibitively large search space. We therefore perform only ciphertext-side search using the specialized solver from Section 4.2.

Integral Distinguisher Results For SKINNY-64-64, we set $(r_1, r_2) = (9, 2)$ and $d = 2$, with plaintext structure 0x0FFFFFFFFFFFFFFF. The search yields 11-round integral properties with data complexity 2^{60}, including 21 balanced linear and 40 balanced nonlinear ciphertext-side mappings. Compared with previous results, the round number is unchanged; the main improvement is the identification of additional certifiable ciphertext-side mappings at the same round depth.

For SKINNY-128-128, we set $(r_1, r_2) = (10, 1)$, restrict the search to linear mappings with $d = 1$, and use plaintext structure 0x00FFFFFFFFFFFFFFFFFFFFFFFFFFFFFF. We recover 33 balanced linear ciphertext-side mappings for 11 rounds with data complexity 2^{120}, matching the result of Zhang *et al.* [23]. Representative mappings are reported in Appendix E, Table 7.

Impact on Key-Recovery Attacks Motivated by the key-recovery attack of Zhang *et al.* [23], which exploits key-independent linear relations between two output cells and processes the two corresponding backward branches independently through separate partial-sum computations, we investigate whether the additional nonlinear mappings found by our search can provide a similar benefit.

We apply the Gaussian-elimination-based ciphertext-side solver of Section 4.2 to SKINNY-64-64 with a plaintext structure containing a single active cell. Under this setting, we obtain the following 7-round integral distinguisher.

Distinguisher 1 (Extended 7-Round Integral distinguisher for SKINNY-64-64) *If plaintext cell* 15 *is active, the four linear mappings*

$$b_4 \oplus b_{52}, \quad b_5 \oplus b_{53}, \quad b_6 \oplus b_{54}, \quad b_7 \oplus b_{55}$$

are balanced after seven rounds, as established in [23]. In addition, our search identifies the following four balanced quadratic mappings:

$$b_4b_6 \oplus b_4b_{54} \oplus b_6b_{52} \oplus b_{52}b_{54}, \quad b_5b_6 \oplus b_5b_{54} \oplus b_6b_{53} \oplus b_{53}b_{54},$$
$$b_5b_7 \oplus b_5b_{55} \oplus b_7b_{53} \oplus b_{53}b_{55}, \quad b_6b_7 \oplus b_6b_{55} \oplus b_7b_{54} \oplus b_{54}b_{55}.$$

The four quadratic relations provide additional filtering information beyond the linear relations used in the baseline attack. However, this stronger filtering comes at a substantial computational cost. With only the linear relations, the two backward branches can be processed separately. The quadratic mappings introduce cross terms between the branches and destroy this branch-wise separability. Under the analyzed construction, the nonlinear relations reduce the ideal data complexity from 2^{50} to 2^{49}, but increase the time complexity from approximately 2^{50} to $2^{67.97}$ equivalent encryptions. Hence, they do not improve the overall key-recovery attack under the currently available techniques. The complete complexity analysis is provided in Appendix G.

This comparison shows that the cryptanalytic value of a key-independent generalized integral mapping depends not only on its filtering power, but also on whether its algebraic structure is compatible with efficient partial decryption and partial-sum evaluation.

5.4 Structural Analysis of Distinguisher Results

The experimental results show that the observed round extensions can be explained by a common structural mechanism: during backward expansion, some searched ciphertext-side mappings admit sufficient algebraic simplification to reduce to relations that are already certifiable at the preceding round depth. For `SPECK` and `SIMON`, this mechanism leads to one-round extensions, whereas for `PRESENT`, `RECTANGLE`, and `SKINNY`, no comparable simplification is identified within the current candidate spaces and propagation models, and the main gain is therefore an enlarged set of certifiable mappings at the same round depth.

For `SPECK`, one round satisfies

$$L_{i+1} = ((L_i \ggg \alpha) + R_i) \oplus RK_i, \quad R_{i+1} = (R_i \lll \beta) \oplus L_{i+1},$$

and hence

$$L_{i+1} \oplus R_{i+1} = R_i \lll \beta.$$

Therefore, after one round of backward expansion, certain ciphertext-side linear mappings reduce to simple linear expressions, up to key-dependent constants that are independent of the plaintext. This allows previously certified balanced relations over fewer rounds to be extended by one round.

For `SIMON`, the relevant structure comes from the low-degree round function

$$f(x) = (x \lll a \,\&\, x \lll b) \oplus (x \lll c).$$

The nonlinear mappings found by our search can cancel quadratic terms arising during backward expansion, allowing degree-two ciphertext-side mappings to remain certifiable for one additional round.

6 Conclusion

In this work, we introduced a basis-independent structural classification of key-independent generalized integral properties into separably generated and residual-cancellation properties, and showed that the eight additional `PRESENT` properties

of Beyne and Verbauwhede [3] admit explicit key-independent separable decompositions. Motivated by this structure, we developed the unified MBM framework and a specialized Gaussian-elimination-based solver to search for balanced separable terms within structured plaintext-side and ciphertext-side candidate spaces. Our experiments identify additional integral properties on several representative ciphers, including one-round extensions for `SIMON` and for the `SPECK` variants except `SPECK` 32, and additional certifiable mappings at the same round depth for the other evaluated ciphers. The `SKINNY-64-64` key-recovery analysis further shows that nonlinear integral relations do not necessarily yield better attacks: additional quadratic relations reduce the ideal data complexity from 2^{50} to 2^{49}, but their cross-branch coupling increases the direct evaluation cost to approximately $2^{67.97}$ equivalent encryptions. These results highlight that the cryptanalytic value of generalized integral mappings depends not only on their filtering power, but also on whether their algebraic structure can be exploited efficiently.

An important direction for future work is to understand the role of residual-cancellation properties in $\mathcal{G}_{\mathrm{rc}}$. Among the concrete generalized integral properties examined in this work, we have not identified an example whose key-independent evaluation genuinely relies on cancellation among key-dependent residuals. `PRESENT` provides a natural open case: the longest currently known integral properties reach 9 rounds, whereas Peng *et al.* [16] exclude 11-round integral distinguishers but provide only partial nonexistence results for 10 rounds. Thus, one round remains unresolved between the known distinguishers and the 11-round resistance result. If a nontrivial 10-round generalized integral property exists, it may be separably generated but lie outside the structured search spaces considered here, or it may exhibit a residual-cancellation structure in $\mathcal{G}_{\mathrm{rc}}$. Determining whether residual-cancellation properties can occur in such unresolved high-round regimes, and whether they can provide a genuine round-depth advantage over separably generated properties, remains an interesting open problem; a more formal discussion is given in Appendix H.

Acknowledgements

Generative AI tools were used to assist with language editing and with preparing initial drafts of selected portions of this manuscript. All scientific content was reviewed, verified, and revised by the authors, who take full responsibility for the final manuscript.

References

1. Beaulieu, R., Shors, D., Smith, J., Treatman-Clark, S., Weeks, B., Wingers, L.: The SIMON and SPECK Families of Lightweight Block Ciphers. IACR Cryptol. ePrint Arch. p. 404 (2013)
2. Beierle, C., Jean, J., Kölbl, S., Leander, G., Moradi, A., Peyrin, T., Sasaki, Y., Sasdrich, P., Sim, S.M.: The SKINNY Family of Block Ciphers and Its Low-Latency

Variant MANTIS. In: CRYPTO (2). Lecture Notes in Computer Science, vol. 9815, pp. 123–153. Springer (2016)
3. Beyne, T., Verbauwhede, M.: Integral Cryptanalysis Using Algebraic Transition Matrices. IACR Trans. Symmetric Cryptol. **2023**(4), 244–269 (2023)
4. Boura, C., Canteaut, A.: Another View of the Division Property. In: CRYPTO (1). Lecture Notes in Computer Science, vol. 9814, pp. 654–682. Springer (2016)
5. Derbez, P., Fouque, P.A.: Increasing precision of division property. IACR Transactions on Symmetric Cryptology pp. 173–194 (2020)
6. Dinur, I., Shamir, A.: Breaking grain-128 with dynamic cube attacks. In: International Workshop on Fast Software Encryption. pp. 167–187. Springer (2011)
7. Hao, Y., Leander, G., Meier, W., Todo, Y., Wang, Q.: Modeling for Three-Subset Division Property Without Unknown Subset - Improved Cube Attacks Against Trivium and Grain-128AEAD. In: EUROCRYPT (1). Lecture Notes in Computer Science, vol. 12105, pp. 466–495. Springer (2020)
8. Hebborn, P., Lambin, B., Leander, G., Todo, Y.: Lower Bounds on the Degree of Block Ciphers. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 12491, pp. 537–566. Springer (2020)
9. Hebborn, P., Lambin, B., Leander, G., Todo, Y.: Strong and tight security guarantees against integral distinguishers. In: ASIACRYPT (1). pp. 362–391. Lecture Notes in Computer Science, Springer (2021)
10. Hu, K., Sun, S., Wang, M., Wang, Q.: An Algebraic Formulation of the Division Property: Revisiting Degree Evaluations, Cube Attacks, and Key-Independent Sums. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 12491, pp. 446–476. Springer (2020)
11. Hu, K., Yap, T.: Perfect Monomial Prediction for Modular Addition. IACR Trans. Symmetric Cryptol. **2024**(3), 177–199 (2024)
12. Huang, S., Wang, X., Xu, G., Wang, M., Zhao, J.: Conditional cube attack on reduced-round keccak sponge function. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 259–288. Springer (2017)
13. Hwang, Y., Kim, I., Kim, S., Lee, M., Shin, H., Hong, D., Hong, S., Lee, D., Sung, J., Seok, B.: Improving neural-inspired integral distinguishers via a linear-algebraic approach. Cryptology ePrint Archive (2026)
14. Knudsen, L.R., Wagner, D.A.: Integral Cryptanalysis. In: FSE. Lecture Notes in Computer Science, vol. 2365, pp. 112–127. Springer (2002)
15. Lambin, B., Derbez, P., Fouque, P.: Linearly equivalent S-boxes and the division property. Des. Codes Cryptogr. **88**(10), 2207–2231 (2020)
16. Peng, S., He, J., Hu, K., Wang, M.: Delving deep into security guarantees against integral distinguishers with applications to PRESENT, TWINE and LBLOCK. Designs, Codes and Cryptography **94**(6) (2026). https://doi.org/10.1007/s10623-026-01871-5
17. Todo, Y.: Structural Evaluation by Generalized Integral Property. In: EUROCRYPT (1). Lecture Notes in Computer Science, vol. 9056, pp. 287–314. Springer (2015)
18. Todo, Y., Morii, M.: Bit-Based Division Property and Application to Simon Family. In: Peyrin, T. (ed.) Fast Software Encryption - 23rd International Conference, FSE 2016, Bochum, Germany, March 20-23, 2016, Revised Selected Papers. pp. 357–377. Lecture Notes in Computer Science (2016)
19. Wang, D., Hadipour, H., Gerhalter, S.: On extending integral distinguishers. IACR Cryptol. ePrint Arch. **2026**, 1402 (2026)

20. Wang, Q., Liu, Z., Varici, K., Sasaki, Y., Rijmen, V., Todo, Y.: Cryptanalysis of reduced-round SIMON32 and SIMON48. In: INDOCRYPT. pp. 143–160. Lecture Notes in Computer Science, Springer (2014)
21. Wang, S., Hu, B., Guan, J., Zhang, K., Shi, T.: MILP-aided method of searching division property using three subsets and applications. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 398–427. Springer (2019)
22. Xiang, Z., Zhang, W., Bao, Z., Lin, D.: Applying MILP Method to Searching Integral Distinguishers Based on Division Property for 6 Lightweight Block Ciphers. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 10031, pp. 648–678 (2016)
23. Zhang, L., Yao, Y., Shi, D., Chai, D., Guo, J., Wang, Z.: Neural-inspired advances in integral cryptanalysis. In: EUROCRYPT. pp. 451–481. Lecture Notes in Computer Science, Springer (2026)

A Deferred Proofs

A.1 Proof of Theorem 1

Proof. By definition of $M_{\mathcal{B}_d}$, for each $\beta \in \mathcal{B}_d$, the row indexed by β records the unknown monomials appearing in the ANF of $F_\beta(x, k)$. Therefore, for a coefficient vector $\gamma \in \mathbb{F}_2^{|\mathcal{B}_d|}$, the product $\gamma M_{\mathcal{B}_d}$ records over $\mathbb{F}_2$ the parity with which each unknown monomial in $\mathcal{M}$ appears in the ANF of

$$\bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta F_\beta(x, k).$$

Hence, $\gamma M_{\mathcal{B}_d} = 0$ is equivalent to requiring every recorded unknown monomial to occur with even parity and therefore cancel.

Under the adopted balanced/unknown monomial classification, every remaining monomial is certified to have zero sum with respect to $\mathbb{X}$. For a remaining monomial of the form $x^{u_1}k^{u_2}$, the factor k^{u_2} is constant once the key is fixed and therefore does not affect the zero-sum property over $\mathbb{X}$. Consequently,

$$\bigoplus_{x \in \mathbb{X}} \bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta F_\beta(x, k) = 0$$

for every key k, or equivalently,

$$\bigoplus_{x \in \mathbb{X}} g(E_k(x)) = 0.$$

Thus, $\mathbf{1}_{\mathbb{X}}(x)g(y)$ defines a balanced separable term under the adopted monomial classification.

Finally,

$$\left\{\gamma \in \mathbb{F}_2^{|\mathcal{B}_d|} : \gamma M_{\mathcal{B}_d} = 0\right\}$$

is exactly the left kernel of $M_{\mathcal{B}_d}$, i.e., the solution space of the induced cancellation system.

A.2 Proof of Theorem 2

Proof. For each i, a plaintext x belongs to $\mathbb{X}_i$ if and only if $x_i = a_i$, or equivalently $x_i \oplus a_i = 0$. Hence,

$$\mathbf{1}_{\mathbb{X}_i}(x) = 1 \oplus x_i \oplus a_i.$$

Since symmetric difference corresponds to XOR of indicator functions,

$$\mathbf{1}_{\mathbb{X}_\sigma}(x) = \bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}(x) = \bigoplus_{i=0}^{n-1} \sigma_i(1 \oplus x_i \oplus a_i).$$

Therefore, $x \in \mathbb{X}_\sigma$ if and only if

$$\bigoplus_{i=0}^{n-1} \sigma_i(1 \oplus x_i \oplus a_i) = 1,$$

which is equivalent to

$$\bigoplus_{i=0}^{n-1} \sigma_i x_i = 1 \oplus \bigoplus_{i=0}^{n-1} \sigma_i a_i \oplus \bigoplus_{i=0}^{n-1} \sigma_i.$$

Since $\sigma \neq 0$, this is a nontrivial affine linear constraint. Hence, $\mathbb{X}_\sigma$ is an affine hyperplane of $\mathbb{F}_2^n$ and therefore has size 2^{n-1}.

A.3 Proof of Theorem 3

Proof. Using the indicator-function representation of $\mathbb{X}_\sigma$, we have

$$\mathbf{1}_{\mathbb{X}_\sigma}(x) = \bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}(x).$$

Therefore,

$$\begin{aligned}
\bigoplus_{x \in \mathbb{X}_\sigma} F(x) &= \bigoplus_{x \in \mathbb{F}_2^n} \mathbf{1}_{\mathbb{X}_\sigma}(x) F(x) \\
&= \bigoplus_{x \in \mathbb{F}_2^n} \left(\bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}(x) \right) F(x) \\
&= \bigoplus_{i=0}^{n-1} \sigma_i \left(\bigoplus_{x \in \mathbb{F}_2^n} \mathbf{1}_{\mathbb{X}_i}(x) F(x) \right) \\
&= \bigoplus_{i=0}^{n-1} \sigma_i \left(\bigoplus_{x \in \mathbb{X}_i} F(x) \right).
\end{aligned}$$

This proves the claimed identity. Taking $F(x) = g(E_k(x))$ gives the stated special case for any ciphertext-side Boolean mapping g and any key k.

A.4 Proof of Theorem 4

Proof. Let

$$\mathbb{X}_\sigma = \triangle_{i:\sigma_i=1} \mathbb{X}_i, \quad V_\sigma = \{v_a \in V : \eta_a(\sigma) = 1\}, \quad U_\gamma = \{u_b \in U : \xi_b(\gamma) = 1\}.$$

By construction of $M_{\mathcal{B}_d}$, after backward expansion through the last r_3 rounds, the ciphertext-side mapping can be written as

$$g_\gamma(E_k(x)) = \bigoplus_{u_b \in U_\gamma} u_b(x, k),$$

where U_γ contains exactly the ciphertext-side monomials that remain after cancellation among the selected basis monomials. The condition $\Phi(\sigma, \gamma) = 0$ implies that no vertex in V_σ is connected to any vertex in U_γ. By construction of the middle propagation graph, an edge is omitted only when the propagation backend proves that the corresponding plaintext-side source cannot contribute to the ciphertext-side monomial. Hence, for every $u_b \in U_\gamma$, none of the plaintext-side sources surviving under σ can contribute an unknown term to u_b. Under the combined three-subset description induced by σ, this certifies

$$\bigoplus_{x \in \mathbb{X}_\sigma} u_b(x, k) = 0$$

for every key k and every $u_b \in U_\gamma$. Therefore,

$$\begin{aligned}\bigoplus_{x \in \mathbb{X}_\sigma} g_\gamma(E_k(x)) &= \bigoplus_{x \in \mathbb{X}_\sigma} \left(\bigoplus_{u_b \in U_\gamma} u_b(x, k) \right) \\ &= \bigoplus_{u_b \in U_\gamma} \left(\bigoplus_{x \in \mathbb{X}_\sigma} u_b(x, k) \right) = 0.\end{aligned}$$

Since $h_\sigma = \mathbf{1}_{\mathbb{X}_\sigma}$, the Boolean function $h_\sigma(x) g_\gamma(y)$ therefore defines a balanced separable term under the MBM criterion. Hence, the resulting generalized integral property belongs to $\mathcal{G}_{\text{sep}}$.

A.5 Proof of Theorem 5

Proof. By Theorem 1, the ciphertext-side mappings certified by the induced cancellation system in the chosen degree-bounded candidate space are exactly the coefficient vectors $\gamma \in \mathbb{F}_2^{|\mathcal{B}_d|}$ satisfying $\gamma M_{\mathcal{B}_d} = 0$, i.e., the left kernel of $M_{\mathcal{B}_d}$. Hence, the homogeneous system admits a nonzero solution if and only if this left kernel is nontrivial, or equivalently,

$$\operatorname{rank}(M_{\mathcal{B}_d}) < |\mathcal{B}_d|.$$

By the rank–nullity theorem, the dimension of the left kernel is

$$\delta = |\mathcal{B}_d| - \operatorname{rank}(M_{\mathcal{B}_d}).$$

Let $\gamma^{(1)}, \ldots, \gamma^{(\delta)} \in \mathbb{F}_2^{|\mathcal{B}_d|}$ be a basis of this solution space. Then every solution γ can be written uniquely as

$$\gamma = \bigoplus_{u=1}^{\delta} \alpha_u \gamma^{(u)}, \quad \alpha_u \in \{0, 1\}.$$

Whenever the coefficients are not all zero, $\gamma \neq 0$ and determines the ciphertext-side mapping $g_\gamma(y) = \bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta y^\beta$. Conversely, every nonzero vector in the left kernel defines a mapping certified by this criterion, completing the proof.

B Illustration of the Plaintext-Side L-Part in MBM Search

This appendix illustrates how the plaintext-side L-part affects the MBM certificate. In the three-subset representation, L-type sources are combined by parity, whereas K-type sources are retained conservatively by union. The surviving L-type sources may propagate through the middle rounds to ciphertext-side monomials that remain uncertified by the propagation backend. The following `PRESENT` example illustrates how cancellation of such sources can affect the certification of a ciphertext-side mapping.

Example for ***`PRESENT`****.* Consider the partition

$$(r_1, r_2, r_3) = (1, 7, 1)$$

and the plaintext structures $\mathbb{X}_{13}$ and $\mathbb{X}_{14}$. We examine the ciphertext-side mapping

$$g(y) = y_7 \oplus y_{39}.$$

After propagation through the first round, the two plaintext structures induce the same L-part:

$$\mathbb{L}_{13} = \mathbb{L}_{14} = \{\ell_0, \ell_1, \ell_2\},$$

where

$$\ell_0 = \texttt{0xEFFFEFFFFFFFEFFF},$$
$$\ell_1 = \texttt{0xEFFFFFFFEFFFEFFF},$$
$$\ell_2 = \texttt{0xFFFFEFFFEFFFEFFF}.$$

For either $\mathbb{X}_{13}$ or $\mathbb{X}_{14}$ individually, these L-type sources survive in the plaintext-side representation.

Backward expansion of $g(y)$ through the last round, followed by the seven-round middle propagation, leaves the following three monomials at the middle interface classified as unknown:

$$t_{29}t_{30}, \qquad t_{28}t_{30}, \qquad t_{28}t_{29}.$$

The surviving L-type sources ℓ_0, ℓ_1, ℓ_2 are precisely the plaintext-side sources connected to these remaining unknown monomials. Consequently, the MBM danger condition is triggered for either $\mathbb{X}_{13}$ or $\mathbb{X}_{14}$, and the method cannot certify that g is balanced over either structure individually.

Now consider their symmetric difference

$$\mathbb{X}_{13} \triangle \mathbb{X}_{14}.$$

Since the L-part is combined by parity,

$$\mathbb{L}_{13} \triangle \mathbb{L}_{14} = \varnothing.$$

Hence, the three L-type sources above disappear from the combined plaintext-side representation. The corresponding K-part, which is retained conservatively by union, does not introduce a dangerous connection to any of the three target monomials. The MBM criterion therefore certifies

$$\bigoplus_{x \in \mathbb{X}_{13} \triangle \mathbb{X}_{14}} g(E_k(x)) = 0$$

for every key k. Equivalently,

$$\mathbf{1}_{\mathbb{X}_{13} \triangle \mathbb{X}_{14}}(x)\, g(y)$$

is a balanced separable term.

This example illustrates the operational role of the plaintext-side L-part in MBM. For the individual structures, the surviving L-type sources remain connected to the relevant unknown middle-interface monomials and therefore prevent certification. For their combination, the same sources cancel according to the standard three-subset parity rule, removing the corresponding dangerous connections and allowing the target mapping to be certified.

C Deferred Algorithms

Algorithm 1 MBM-Based Joint Search Procedure

Input: $E_k = E_k^{(r_3)} \circ E_k^{(r_2)} \circ E_k^{(r_1)}$; $\mathcal{B}_d$; $a \in \mathbb{F}_2^n$; middle-check backend $\mathcal{C}$
Output: A nonzero pair (σ, γ) satisfying $\Phi(\sigma, \gamma) = 0$, if one exists

Step 1: Plaintext-side matrices
for $i = 0, \ldots, n-1$ **do**
    $\mathbb{X}_i \leftarrow \{x \in \mathbb{F}_2^n : x_i = a_i\}$, $(\mathbb{K}_i, \mathbb{L}_i) \leftarrow$ **ThreeSubsetPropagate**$(\mathbb{X}_i, E_k^{(r_1)})$
end for
$\mathcal{K} \leftarrow \bigcup_i \mathbb{K}_i$, $\mathcal{L} \leftarrow \bigcup_i \mathbb{L}_i$, $V \leftarrow \mathcal{K} \cup \mathcal{L}$; build $M_{\mathcal{K}}$ and $M_{\mathcal{L}}$
Step 2: Ciphertext-side matrix
Expand all basis monomials in $\mathcal{B}_d$ backward through $E_k^{(r_3)}$, collect U, and build $M_{\mathcal{B}_d}$
Step 3: Middle propagation graph
for each $v_a \in V$ and $u_b \in U$ **do**
    $P_{a,b} \leftarrow$ **AutomatedCheck**$(v_a, u_b, E_k^{(r_2)}, \mathcal{C})$
    if $P_{a,b} = 1$ **then**
        Add (v_a, u_b) to $\mathcal{E}$
    end if
end for
$\mathcal{G} \leftarrow (V \cup U, \mathcal{E})$
Step 4: Solve
Introduce $\sigma \in \{0,1\}^n$ and $\gamma \in \{0,1\}^{|\mathcal{B}_d|}$; construct $\Phi(\sigma, \gamma)$ and solve $\Phi(\sigma, \gamma) = 0$ with $\sigma \neq 0$ and $\gamma \neq 0$
    return (σ, γ)

Algorithm 2 Search Ciphertext-Side Low-Degree Mappings

Input: Plaintext structure $\mathbb{X}$; $\mathcal{B}_d$; $E_k = E_k^{(r_2)} \circ E_k^{(r_1)}$; integral-analysis backend
Output: A basis $\gamma^{(1)}, \ldots, \gamma^{(\delta)}$ of the left kernel of $M_{\mathcal{B}_d}$

Step 1: Fix the interface
$t \leftarrow E_k^{(r_1)}(x)$, $y = (f_0(t,k), \ldots, f_{n-1}(t,k))$
Step 2: Construct $M_{\mathcal{B}_d}$
for each $\beta \in \mathcal{B}_d$ **do**
    Compute $F_\beta(t,k) = \prod_{i=0}^{n-1} f_i(t,k)^{\beta_i}$ and expand it in ANF over (t,k)
    Classify the resulting monomials with respect to $\mathbb{X}$ and record those classified as unknown
end for
Build $M_{\mathcal{B}_d}$ as in Section 3.2
Step 3: Gaussian elimination
Compute a basis $\gamma^{(1)}, \ldots, \gamma^{(\delta)}$ of the left kernel of $M_{\mathcal{B}_d}$
    return $\gamma^{(1)}, \ldots, \gamma^{(\delta)}$

D Additional Details for SPECK

D.1 Formalization of Linear Enhancement

Let $\mathbf{s}_k^{(r)}(x) = (s_0, \ldots, s_{n-1})$ denote the intermediate state after Round r under key k. We write

$$\mathcal{L}^{(r)} = \left\{ \ell(\mathbf{s}^{(r)}) = \bigoplus_{i=0}^{n-1} \lambda_i s_i \;\middle|\; \lambda_i \in \mathbb{F}_2 \right\}$$

for the space of linear Boolean functions on the intermediate state, and define

$$\mathcal{S}_{\mathbb{X}}^{(r)} = \left\{ \ell \in \mathcal{L}^{(r)} \;\middle|\; \bigoplus_{x \in \mathbb{X}} \ell(\mathbf{s}_k^{(r)}(x)) = 0 \text{ for every key } k \right\}$$

as the subspace of linear relations that are balanced over $\mathbb{X}$ for every key.

Theorem 6 (Linear Enhancement via Intermediate-State Relations). *Let $\mathbb{X} \subseteq \mathbb{F}_2^n$ be a plaintext structure, and let $\mathbf{s}_k^{(r)}(x)$ be the intermediate state after Round r. Suppose that a ciphertext-side linear mapping g_μ can be rewritten as*

$$g_\mu(E_k(x)) = \ell(\mathbf{s}_k^{(r)}(x)) \oplus q(x, k),$$

where $\ell \in \mathcal{S}_{\mathbb{X}}^{(r)}$ and

$$\bigoplus_{x \in \mathbb{X}} q(x, k) = 0$$

for every key k. Then

$$\bigoplus_{x \in \mathbb{X}} g_\mu(E_k(x)) = 0$$

for every key k.

Proof. By assumption,

$$g_\mu(E_k(x)) = \ell(\mathbf{s}_k^{(r)}(x)) \oplus q(x, k).$$

Therefore,

$$\begin{aligned} \bigoplus_{x \in \mathbb{X}} g_\mu(E_k(x)) &= \bigoplus_{x \in \mathbb{X}} \ell(\mathbf{s}_k^{(r)}(x)) \oplus \bigoplus_{x \in \mathbb{X}} q(x, k) \\ &= 0, \end{aligned}$$

where the first term is zero because $\ell \in \mathcal{S}_{\mathbb{X}}^{(r)}$, and the second is zero by assumption. Hence, g_μ is balanced over $\mathbb{X}$ for every key k.

The theorem formalizes the linear-enhancement strategy used for SPECK: a higher-round ciphertext-side mapping can be certified whenever its backward expansion reduces to a previously certified balanced linear relation plus a remainder that is itself zero-sum over the plaintext structure.

D.2 Balanced Mappings for the 6-Round SPECK32 Distinguisher

Table 2. Six balanced ciphertext-side mappings for the 6-round SPECK32 distinguisher under the plaintext structure 0xFFFFFF1F

Ciphertext-Side Mapping	Ciphertext-Side Mapping	Ciphertext-Side Mapping
$y_8 \oplus y_{24}$	$y_9 \oplus y_{25}$	$y_{10} \oplus y_{26}$
$y_{11} \oplus y_{27}$	$y_{12} \oplus y_{28}$	$y_{13} \oplus y_{29}$

D.3 Complete ANF Derivation for the 7-Round SPECK32 Example

We illustrate linear enhancement on the 7-round SPECK32 mapping

$$y_6 \oplus y_{13} \oplus y_{15} \oplus y_{22} \oplus y_{29}.$$

Rather than analyzing this mapping directly at the ciphertext level, we rewrite it over the intermediate state after 6 rounds and reduce it to a previously certified balanced linear relation.

Let $\mathbf{s}^{(6)} = (s_0, \ldots, s_{31})$, $y = (y_0, \ldots, y_{31})$, and $k = (k_0, \ldots, k_{15})$ denote the state after 6 rounds, the ciphertext, and the key bits involved in the last round, respectively. To keep the expressions readable, let Δ denote the common terms appearing in the ANFs of y_6 and y_{22}. Then

$$\begin{aligned}
y_6 &= s_{15} \oplus s_{22} \oplus \Delta \oplus k_6, \\
y_{13} &= s_6 \oplus s_{29} \oplus s_7 s_{30} \oplus s_7 s_8 s_{31} \oplus s_8 s_{30} s_{31} \oplus k_{13}, \\
y_{15} &= s_8 \oplus s_{31} \oplus k_{15}, \\
y_{22} &= s_{15} \oplus s_{22} \oplus s_{24} \oplus \Delta \oplus k_6, \\
y_{29} &= s_6 \oplus s_{29} \oplus s_7 s_{30} \oplus s_{31} \oplus s_7 s_8 s_{31} \oplus s_8 s_{30} s_{31} \oplus k_{13}.
\end{aligned} \tag{1}$$

Under the original backend, these expressions are classified as unknown when considered separately. After XORing them, however, the common terms cancel pairwise, yielding

$$y_6 \oplus y_{13} \oplus y_{15} \oplus y_{22} \oplus y_{29} = s_8 \oplus s_{24} \oplus k_{15}.$$

The nonconstant part $s_8 \oplus s_{24}$ is exactly the previously certified 6-round balanced relation reported in Table 2. Moreover, k_{15} is independent of the plaintext. Since the plaintext structure contains 2^{29} elements,

$$\bigoplus_{x \in \mathbb{X}} k_{15} = 0.$$

Thus,

$$s_8 \oplus s_{24} \oplus k_{15}$$

is balanced over $\mathbb{X}$, and hence so is the 7-round ciphertext-side mapping. This is precisely the linear-enhancement mechanism formalized in Theorem 6.

E Representative Experimental Results

To keep the appendix concise, we report only representative distinguishers here.

Table 3. Representative 7-round integral distinguishers for the SPECK family

Variant	Data	Plaintext Structure	Ciphertext-Side Mapping
SPECK32	2^{29}	0xFFFFFF1F	$y_6 \oplus y_{13} \oplus y_{15} \oplus y_{22} \oplus y_{29}$
SPECK48	2^{44}	0xFFFFFFFFFE1F	$y_{12} \oplus y_{20} \oplus y_{23} \oplus y_{36} \oplus y_{44}$
SPECK64	2^{52}	0xFFFFFFFFFFFE001F	$y_{20} \oplus y_{28} \oplus y_{31} \oplus y_{52} \oplus y_{60}$
SPECK96	2^{68}	0xFFFFFFFFFFFF_ FFFE0000001F	$y_{36} \oplus y_{44} \oplus y_{47} \oplus y_{84} \oplus y_{92}$
SPECK128	2^{84}	0xFFFFFFFFFFFFFFFF_ FFFE00000000001F	$y_{52} \oplus y_{60} \oplus y_{63} \oplus y_{116} \oplus y_{124}$

Table 4. Representative integral distinguishers for the SIMON family

Variant	Round	Data	Plaintext-Side Mapping	Ciphertext-Side Mapping
SIMON32	15	2^{31}	x_2	$y_{10} \oplus y_{18}y_{27} \oplus y_{28}$
SIMON48	16	2^{47}	x_{13}	$y_{21} \oplus y_{29}y_{46} \oplus y_{47}$
SIMON64	18	2^{63}	x_{15}	$y_9 \oplus y_{42}y_{49} \oplus y_{43}$
SIMON96	22	2^{95}	x_{37}	$y_{30} \oplus y_{79}y_{86} \oplus y_{80}$
SIMON128	26	2^{127}	x_3	$y_1 \oplus y_{66}y_{73} \oplus y_{67}$

Table 5. Two representative 9-round integral distinguishers for PRESENT

Data	Plaintext-Side Mapping	Ciphertext-Side Mapping
2^{63}	x_{12}	$y_{15} \oplus y_{47}$
	$x_{13} \oplus x_{14}$	y_{51}

Table 6. Representative 10-round integral distinguishers for RECTANGLE

Data	Plaintext-Side Mapping	Ciphertext-Side Mapping
2^{63}	$x_{16} \oplus x_{32}$	y_5
	$x_{17} \oplus x_{33}$	y_6

Table 7. Representative 11-round ciphertext-side mappings for SKINNY

Variant	Round	Data	Ciphertext-Side Mapping
SKINNY-64-64	11	2^{60}	$y_{12} \oplus y_{28} \oplus y_{44}$
			$y_{20}y_{23} \oplus y_{20}y_{55} \oplus y_{23}y_{52} \oplus y_{36} \oplus y_{38} \oplus y_{52}y_{55}$
SKINNY-128-128	11	2^{120}	$y_{32} \oplus y_{64} \oplus y_{96}$

F Tensor-Based Deduplication of Mapping-Based Integral Properties

The MBM search may return multiple balanced separable terms that are linearly redundant as generalized Boolean functions. For example, consider

$$(x_1 \oplus x_2, y_1 \oplus y_2), \quad (x_1 \oplus x_2, y_1 \oplus y_3),$$
$$(x_1 \oplus x_3, y_2 \oplus y_3), \quad (x_2 \oplus x_3, y_2 \oplus y_3).$$

Under the tensor encoding introduced below, the fourth property is the XOR of the first three and is therefore linearly redundant. For joint plaintext- and ciphertext-side mappings, this redundancy can be detected directly in the tensor product of the two candidate spaces.

Theorem 7 (Tensor-Based Deduplication). *Let*

$$\sigma = (\sigma_0, \ldots, \sigma_{n-1}) \in \mathbb{F}_2^n$$

and

$$\gamma = (\gamma_\beta)_{\beta \in \mathcal{B}_d} \in \mathbb{F}_2^{|\mathcal{B}_d|}$$

encode the plaintext-side and ciphertext-side mappings

$$h_\sigma(x) = \bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}(x)$$

and

$$g_\gamma(y) = \bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta y^\beta,$$

respectively. We encode the mapping pair (σ, γ) *by*

$$\sigma \otimes \gamma = (\sigma_i \gamma_\beta)_{0 \le i < n,\ \beta \in \mathcal{B}_d} \in \mathbb{F}_2^{n|\mathcal{B}_d|}.$$

Any linear dependence among these tensor encodings induces the same linear dependence among the corresponding generalized Boolean functions $h_\sigma(x) g_\gamma(y)$*. Therefore, retaining a basis of the tensor encodings is sufficient to remove linear redundancies among the recovered mapping-based integral properties.*

Proof. By the definitions of h_σ and g_γ,

$$\begin{aligned} h_\sigma(x) g_\gamma(y) &= \left(\bigoplus_{i=0}^{n-1} \sigma_i \mathbf{1}_{\mathbb{X}_i}(x) \right) \left(\bigoplus_{\beta \in \mathcal{B}_d} \gamma_\beta y^\beta \right) \\ &= \bigoplus_{i=0}^{n-1} \bigoplus_{\beta \in \mathcal{B}_d} \sigma_i \gamma_\beta \mathbf{1}_{\mathbb{X}_i}(x) y^\beta. \end{aligned}$$

Thus, the tensor $\sigma \otimes \gamma$ records exactly the coefficients of the corresponding generalized Boolean function in the product candidate space generated by

$$\left\{ \mathbf{1}_{\mathbb{X}_i}(x) y^{\beta} : 0 \le i < n,\ \beta \in \mathcal{B}_d \right\}.$$

Suppose that a collection of recovered solutions satisfies

$$\bigoplus_k \lambda_k \left(\sigma^{(k)} \otimes \gamma^{(k)} \right) = 0, \quad \lambda_k \in \mathbb{F}_2.$$

Then the coefficient of every product term $\mathbf{1}_{\mathbb{X}_i}(x) y^{\beta}$ in the corresponding XOR is zero. Hence,

$$\bigoplus_k \lambda_k h_{\sigma^{(k)}}(x) g_{\gamma^{(k)}}(y) = 0.$$

Therefore, any linear dependence among the tensor encodings induces the corresponding linear dependence among the recovered mapping-based integral properties. It is consequently sufficient to retain a basis of the tensor encodings when removing linear redundancies.

With this tensor encoding, the deduplication step reduces to computing a basis of the recovered tensor vectors over $\mathbb{F}_2$, which can be performed efficiently by Gaussian elimination.

G Key-Recovery Analysis of Nonlinear Integral Relations for SKINNY-64-64

We analyze whether the additional quadratic relations in Distinguisher 1 can improve the 16-round $4 + 7 + 5$ key-recovery attack of Zhang *et al.* [23]. Their attack uses four key-independent linear relations between output cells 1 and 13, whose two backward branches can be processed independently by partial sums. Our search provides four additional quadratic relations over the same two cells.

Key-Recovery Setting and Loss of Branch Separability We follow the 16-round $4 + 7 + 5$ framework of Zhang *et al.* [23]. The target cipher is divided into three parts: rounds 0–3, the 7-round distinguisher over rounds 4–10, and the final five rounds 11–15.

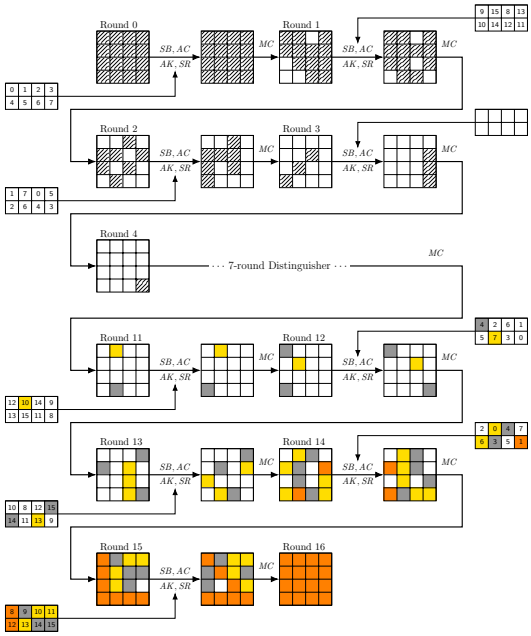


Fig. 2. The 16-round $4 + 7 + 5$ key-recovery framework for SKINNY-64-64. Yellow and gray indicate the two backward branches, while orange marks their shared dependencies. With only the linear relations, the two branches can be processed independently; the quadratic relations introduce cross-branch terms and require joint evaluation.

Let

$$u = x_{11}[1] = (u_0, u_1, u_2, u_3), \qquad v = x_{11}[13] = (v_0, v_1, v_2, v_3), \qquad d = u \oplus v.$$

The four linear relations are

$$\bigoplus_t d_i(t) = 0, \qquad 0 \le i < 4,$$

whereas the four additional quadratic relations involve

$$d_0d_2, \quad d_1d_2, \quad d_1d_3, \quad d_2d_3.$$

For a quadratic term d_id_j,

$$\bigoplus_t d_i(t)d_j(t) = \bigoplus_t u_i(t)u_j(t) \oplus \bigoplus_t v_i(t)v_j(t) \oplus \bigoplus_t u_i(t)v_j(t) \oplus \bigoplus_t u_j(t)v_i(t).$$

The first two terms are branch-local, whereas the last two are cross-branch terms. Hence, unlike the linear relations, the quadratic relations cannot in general be evaluated from the marginal summaries produced by two independent branch-wise partial sums. The values of the two backward branches must remain aligned for each text, which destroys the branch-wise separability exploited by the baseline attack.

Importantly, the nonlinear relations do not require guessing additional key material. Both the linear and nonlinear procedures involve the same 14 master-key cells. The increase in complexity is instead caused by the need to evaluate the two branches jointly.

Complexity Comparison Without the quadratic relations, the two backward branches can be processed independently using the original partial-sum procedure. Under the same wrong-key randomization model as in Zhang *et al.* [23], four large structures of 2^{48} chosen plaintexts are required, giving

$$D_{\text{linear}} = 4 \cdot 2^{48} = 2^{50}.$$

The branch-wise partial-sum computation remains below the data-collection cost, so the total time complexity is approximately

$$T_{\text{linear}} \approx 2^{50}.$$

When the four quadratic relations are included, the additional filtering allows two large structures to suffice, reducing the data complexity to

$$D_{\text{nonlinear}} = 2 \cdot 2^{48} = 2^{49}.$$

However, the cross-branch terms force a joint partial-sum computation. We construct an explicit joint evaluation procedure over the same 14 involved master-key cells and optimize the key-cell ordering within the considered cell-wise model. For one large structure, the joint computation costs approximately

$$2^{66.9688}$$

equivalent 16-round encryptions. Using two structures therefore gives

$$T_{\text{nonlinear}} \approx 2^{67.9688}.$$

Table 8. Linear-only and nonlinear joint key-recovery complexities for `SKINNY-64-64`

Method	Large Structures	Data	Total Time
Four linear tests only	4	2^{50}	2^{50}
Eight linear/quadratic tests	2	2^{49}	$2^{67.9688}$

Thus, the quadratic relations save one bit in the data exponent but increase the direct evaluation cost from approximately 2^{50} to $2^{67.9688}$. The increase is not caused by a larger key-guess space, but by the loss of branch-wise separability and the resulting joint computation.

The value $2^{67.9688}$ is the complexity of the explicit joint construction analyzed here, rather than a lower bound on every possible method for exploiting the quadratic relations. Alternative structured MITM, transform-based, low-rank, or other specialized techniques are not excluded by this analysis.

This comparison shows that the additional filtering power of nonlinear integral relations does not by itself imply a better key-recovery attack. Their usefulness also depends on whether they can be evaluated without destroying the computational separability exploited by the baseline attack.

H High-Round Structure of Generalized Integral Properties

Definition 5 partitions key-independent generalized integral properties into the separably generated class $\mathcal{G}_{\text{sep}}$ and the residual-cancellation class $\mathcal{G}_{\text{rc}}$. The search methods developed in this work focus on structured balanced separable generators of $\mathcal{G}_{\text{sep}}$, rather than the full class itself. This leaves open a broader question: whether residual-cancellation properties can occur at round depths beyond those attainable by separably generated properties.

The formal nonemptiness of $\mathcal{G}_{\text{rc}}$ can be established by construction.[4] From a cryptanalytic perspective, however, the relevant question is whether such residual cancellation can arise naturally from cipher propagation and provide an additional round-depth advantage over separably generated properties. Among the properties examined in this work, we have not identified an example whose key-independent integral sum genuinely requires residual cancellation. As shown in Section 3.1, the eight additional `PRESENT` properties of Beyne and Verbauwhede [3] all belong to $\mathcal{G}_{\text{sep}}$.

To formalize the possible round-depth advantage of residual-cancellation properties, for an r-round cipher, let $\mathcal{G}_{\text{sep}}^{(r)}$ and $\mathcal{G}_{\text{ki}}^{(r)}$ denote the sets of nontrivial separably generated properties and nontrivial key-independent generalized integral properties, respectively. Define

$$R_{\text{sep}} = \max\left\{r : \mathcal{G}_{\text{sep}}^{(r)} \neq \varnothing\right\}, \quad R_{\text{ki}} = \max\left\{r : \mathcal{G}_{\text{ki}}^{(r)} \neq \varnothing\right\}.$$

Since $\mathcal{G}_{\text{sep}} \subseteq \mathcal{G}_{\text{ki}}$, we necessarily have

$$R_{\text{sep}} \leq R_{\text{ki}}.$$

The central question is whether this inequality can be strict for practical block ciphers, i.e., whether residual-cancellation properties can occupy an additional high-round range after separably generated properties cease to exist.

Conjecture 1 (Separable-Dominance Conjecture). For practical block ciphers, residual-cancellation generalized integral properties do not increase the maximal attainable round depth beyond that of separably generated properties, i.e.,

$$R_{\text{ki}} = R_{\text{sep}}.$$

Existing results are consistent with this conjecture but do not establish it in general. For several ciphers, including `SKINNY-64`, `CRAFT`, the `SIMON` family, and

[4] For one ordinary round of `PRESENT`, write $E_K(x) = T(x \oplus K)$, where $T = P \circ S$ denotes the composition of the S-box layer and the permutation layer. Take $A = \{x : x_0 = 0\} \cup \{e_0\}$ and $B = A \oplus e_0$, and define $h_1(x) = \mathbf{1}_A(x)$, $h_2(x) = \mathbf{1}_B(x)$, $g_1(y) = \mathbf{1}_B(T^{-1}(y))$, and $g_2(y) = \mathbf{1}_A(T^{-1}(y))$. Since $T^{-1}(E_K(x)) = x \oplus K$ and $B = A \oplus e_0$, the two separable terms leave identical key-dependent residual sums, which cancel in their XOR. Thus, the resulting property belongs to $\mathcal{G}_{\text{rc}}$, showing that this class is formally nonempty.

the SIMECK family, known integral distinguishers already meet the corresponding integral-resistance upper bounds of Hebborn *et al.* [9]. Within the settings covered by these bounds, this leaves no room for further round-depth improvements.

PRESENT provides a natural open case. The longest currently known integral properties reach 9 rounds, whereas Peng *et al.* [16] establish resistance for 11 rounds and provide only partial nonexistence results for 10 rounds. Thus, the existence of a nontrivial 10-round property remains unresolved.

If no 10-round key-independent generalized integral property exists, then the currently known 9-round depth may already be maximal. If a 10-round property exists and is separably generated, it must lie outside the structured search spaces or certification criteria considered in this work. In contrast, if a 10-round property exists only through residual cancellation, then

$$R_{\mathrm{sep}} < R_{\mathrm{ki}},$$

providing a counterexample to Conjecture 1.

Consequently, if a 10-round key-independent generalized integral property for PRESENT exists, at least two possibilities remain. It may be separably generated but not captured by the structured search spaces and certification criteria considered in this work, or it may exhibit a residual-cancellation structure in $\mathcal{G}_{\mathrm{rc}}$.

The latter case would provide evidence that residual-cancellation properties can extend the attainable round depth beyond that of separably generated properties. Determining whether such a property exists therefore provides a concrete test case for Conjecture 1 and motivates techniques that track and exploit key-dependent residual functions directly.

If these residual functions retain exploitable algebraic structure at high round depths, explicitly tracking them may allow integral analysis to move from merely eliminating key dependence to exploiting its underlying structure. This may provide a new avenue for overcoming current round-depth limitations of integral distinguishers, while also yielding additional information for key-recovery attacks.