Multiple-of Property for Related-Differential Distinguishers on 5-Round AES

Hanbeom Shin[1], Donggeun Kwon[2], Byoungjin Seok[3], Deukjo Hong[4], Jaechul Sung[5], Seokhie Hong[6] and Dongjae Lee[7]

[1] Korea University, Seoul, Republic of Korea
[2] Kunsan National University, Gunsan, Republic of Korea
[3] Hansung University, Seoul, Republic of Korea
[4] Jeonbuk National University, Jeonju, Republic of Korea
[5] University of Seoul, Seoul, Republic of Korea
[6] SmartM2M, Seoul, Republic of Korea
[7] Kangwon National University, Chuncheon, Republic of Korea

Abstract

A *key-independent secret-key distinguisher* identifies inherent structural deviations of a block cipher from an ideal random permutation without recovering any information about the secret key. For the Advanced Encryption Standard (AES), various key-independent secret-key distinguishers have been proposed on reduced-round versions. In this paper, we study related-differential distinguishers for 5-round AES that combine a 1-round related differential trail with the 4-round generalized zero-difference property. We prove that the number of valid quartets satisfying the underlying property takes the multiple-of form, $N_q = 2^{2-n_z}(2^w)^{n_z}A + 8B$, where A, B are non-negative integers, w is the cell size, and n_z is the number of inactive bytes of ΔX_2 (the difference at the second-round input) under the chosen *pairing* of the 4-plaintext quartet into two pairs. The pairing fixes the bundle size $2^{2-n_z}(2^w)^{n_z}$, which sets the variance of the valid-quartet count. We compare the three pairings $n_z \in \{0, 1, 2\}$. In the chosen-plaintext setting, where a single structure is examined as a whole, a smaller bundle keeps the count closer to its mean, so $n_z = 0$ is the most reliable single-structure distinguisher, $n_z = 1$ is close behind, and $n_z = 2$ almost fails. The pairing $n_z = 1$ is that of Yan et al. Its $2 \cdot 2^w$ bundle raises the variance, so it reaches 63% only at $2^{27.2}$, rather than at the 2^{27} where one valid quartet is expected. At the same $2^{27.2}$, the smaller bundle of $n_z = 0$ reaches 65%. In the adaptively chosen-plaintext setting the bundle does not form across the separate base collisions, so $n_z = 1$, which produces more valid quartets, is the better pairing. All claims are verified experimentally on both small-AES ($w = 4$) and the standard AES ($w = 8$).

Keywords: AES, Structural Property, Related Difference, Mixture Differential, Distinguisher

1 Introduction

The Advanced Encryption Standard (AES) [DR02] is a block cipher with a 128-bit block size, supporting key sizes of 128, 192, and 256 bits, which correspond to 10, 12, and 14 rounds, respectively. Designed by Daemen and Rijmen and standardized by the National

Emails: newonetiger@korea.ac.kr (Hanbeom Shin), dgkwon@kunsan.ac.kr (Donggeun Kwon), bjseok@hansung.ac.kr (Byoungjin Seok), deukjo.hong@jbnu.ac.kr (Deukjo Hong), jcsung@uos.ac.kr (Jaechul Sung), shhong@smartm2m.co.kr (Seokhie Hong), dongjae.lee@kangwon.ac.kr (Dongjae Lee)

This work is licensed under a "CC BY 4.0" license.
Date of this document: 2026-06-19.

Institute of Standards and Technology (NIST) in 2001, AES remains the most widely used block cipher to date. Its security has been extensively evaluated against a wide range of cryptanalytic techniques over the past two decades. Moreover, several block ciphers have adopted structures similar to AES, such as SKINNY [BJK+16] and Midori [BBI+15], and many tweakable block ciphers reuse the AES round function in their designs.

The security analysis of reduced-round AES has been studied extensively. Differential cryptanalysis [BS91], which exploits non-uniform propagation of input differences, and its numerous extensions, including truncated differentials [Knu94], impossible differentials [Knu98, BBS99], and boomerang attacks [Wag99, DKRS20], have been fundamental tools in this effort. For 5-round AES in particular, a number of structural and distinguishing properties have been discovered in recent years.

Evaluating the security of a block cipher requires analyzing both its resistance to key recovery and its indistinguishability from an ideal random permutation. While a successful key recovery directly implies distinguishability, the converse does not hold, since a cipher may exhibit detectable structural deviations from a random permutation without revealing any information about the secret key. A *secret-key distinguisher* formalizes this notion. Under the standard black-box model, the adversary is given oracle access to either the cipher with an unknown random key or a truly random permutation, and its goal is to decide which one the oracle implements. Grassi et al. [GRR17] further classified secret-key distinguishers into two categories: (1) *key-independent* distinguishers, which verify a structural property that holds for every key without discovering any key information, and (2) *key-dependent* distinguishers, which are derived from key-recovery attacks and rely on the existence of a (partial) key satisfying a certain condition. A *key-independent secret-key distinguisher* is particularly significant because it exposes inherent weaknesses of the round function itself, independent of the key schedule, and remains valid even in the independent subkey model where key-recovery attacks may be invalidated.

The multiple-of-8 property, introduced by Grassi et al. at EUROCRYPT 2017 [GRR17], was the first key-independent secret-key distinguisher for 5-round AES. It states that the number of plaintext pairs from a diagonal space whose ciphertexts collide on an inverse diagonal space is always a multiple of 8. This property was later generalized to arbitrary Substitution-Permutation Network (SPN) block ciphers by Boura et al. [BCC19]. Grassi subsequently proposed mixture-differential cryptanalysis [Gra18], which reformulates the multiple-of-8 property to facilitate key-recovery attacks on 5-round AES. Bar-On et al. [BODK+20] further improved the key-recovery complexity. In [SKK+24], the multiple-of property for SKINNY was revisited, and it was shown that the number of right pairs is not merely a multiple of a constant but a fixed value determined by the choice of input and output subspaces. More recently, the multiple-of-8 property was directly utilized for the first time in a key-recovery attack on 5-round AES [SKS+25b].

Other notable 5-round distinguishers include the yoyo trick by Rønjom et al. [RBH17], later refined by Mondal et al. [MRSA23], and the exchange attack by Bardeh and Rønjom [BR19]. The truncated boomerang [BL23] and boomerang chain [YTXQ24] techniques have also been applied to construct distinguishers for reduced-round AES. Recently, the exchanged boomerang framework was used to construct the massive and multiple exchanged boomerang distinguishers for 5-round AES [SKS+25a], where the multiple exchanged boomerang achieves a data and time complexity of $2^{27.1}$ in the adaptive setting, which represents the current best-known result for 5-round AES. Shin et al. [SKS+26] recently extended the generalized zero-difference property to 7-round AES. Yan et al. [YTQ26] recently presented related-differential distinguishers for 5, 6, and 7-round AES, including a 5-round chosen-plaintext distinguisher with data complexity $2^{27.2}$.

Our Contribution

In this paper, we study related-differential distinguishers for 5-round AES, based on a 5-round property obtained from the first round of the related differential trail of Bardeh and Rijmen [BR22] combined with the 4-round generalized zero-difference property. Independently of Yan et al. [YTQ26], we state this property, which holds on a 4-plaintext quartet with probability 2^{-22}, and prove that its valid-quartet count takes a multiple-of form whose bundle size depends on the chosen pairing of the quartet, controlling the variance of the count and the efficiency of the resulting distinguisher. Our contributions are summarized as follows.

1. **Multiple-of property and pairing analysis.** We prove that the valid-quartet count takes the multiple-of form $N_q = 2^{2-n_z}(2^w)^{n_z}A + 8B$ (Theorem 5), where w is the cell size, A, B are non-negative integers, and n_z is the number of inactive bytes of ΔX_2 under the chosen pairing of the quartet. A 4-plaintext quartet can be partitioned into two pairs in three ways, and each pairing fixes the value of n_z; for the related differential of Bardeh and Rijmen [BR22], the three pairings give $n_z \in \{0, 1, 2\}$, with multiple-of forms $N_q = 4A + 8B$, $N_q = 2 \cdot 2^w A + 8B$, and $N_q = (2^w)^2 A + 8B$ and bundle sizes 4, $2 \cdot 2^w$, and $(2^w)^2$. The bundle therefore fixes the variance of the valid-quartet count and the efficiency of the resulting distinguisher.

2. **Multiple-of explanation of Yan et al.'s success rate.** The distinguisher of Yan et al. [YTQ26] reports a 63% success rate at $2^{27.2}$ data, although a Poisson distribution already gives 63% at 2^{27}, where one valid quartet is expected on average. We show that this gap is a direct consequence of the $2 \cdot 2^w$-bundle from the $n_z = 1$ pairing, and that the structure size $s = 2^{13.6}$ chosen in their work matches the prediction of our framework.

3. **Comparison of the three pairings.** We compare the three pairings experimentally. At 2^{27} data the success rate is 57% for $n_z = 0$, 53% for $n_z = 1$, and 1% for $n_z = 2$; at $2^{27.2}$ data the $n_z = 1$ pairing of Yan et al. [YTQ26] reaches 64% (close to the 63% they report) and the smaller-bundle $n_z = 0$ reaches 65%. In the adaptively chosen-plaintext setting the bundle does not form, so the higher-mean $n_z = 1$ pairing is the better choice. A comparison with existing key-independent secret-key distinguishers for 5-round AES is given in Table 1.

4. **Experimental verification.** The proposed properties and bundle structures are verified experimentally on both small-AES ($w = 4$) and the standard AES ($w = 8$). Source code is available at `https://anonymous.4open.science/r/multiple-of-related-differential`.

Organization

The remainder of this paper is organized as follows. Section 2 provides a description of AES and reviews the related difference, generalized zero-difference property, mixture, and multiple-of-8 property. Section 3 presents the 5-round AES property from related differentials and the multiple-of property of the quartet count, and provides experimental verification. Section 4 constructs chosen-plaintext distinguishers from the three pairings and compares them experimentally, and Section 5 does the same in the adaptively chosen-plaintext setting. Section 6 concludes the paper.

Table 1: Comparison of key-independent secret-key distinguishers for 5-round AES. Data complexity is measured in chosen plaintexts (CP), adaptively chosen plaintexts (ACP), adaptive chosen plaintexts and ciphertexts (ACPC), or adaptive chosen ciphertexts (ACC). Time complexity is measured in 5-round AES encryptions (E) or memory accesses (M).

Property	Data	Time	Setting	Succ.	Ref.
Multiple-of-8	2^{32}	$2^{35.6}$ M	CP	99%	[GRR17]
Exchange	2^{30}	2^{30} E	CP	63%	[BR19]
Yoyo	$2^{29.95}$	$2^{29.95}$ M	ACPC	55%	[MRSA23]
Yoyo	$2^{30.65}$	$2^{30.65}$ M	ACPC	81%	[MRSA23]
Massive ex. boomerang	2^{31}	2^{31} M	ACC	70%	[SKS+25a]
Multiple ex. boomerang	$2^{27.1}$	$2^{27.1}$ M	ACC	80%	[SKS+25a]
Related diff. ($n_z = 1$ pairing)	$2^{27.2}$	$2^{28.05}$ E	CP	63%	[YTQ26]
Related diff. ($n_z = 0$ pairing)	$2^{27.2}$	$2^{28.05}$ E	CP	65%	Section 4
Related diff. ($n_z = 1$ pairing)	$2^{23.3}$	$2^{23.3}$ M	ACP	63%	[YTQ26]

2 Preliminaries

2.1 Description of AES

The 128-bit internal state of AES is represented as a 4×4 matrix of bytes. The bytes of a state S are indexed from 0 to 15 in a column-major order, denoted by $S[i]$ for $0 \leq i \leq 15$. To denote a set of bytes, we use $S[I]$ where I is a set of indices.

The encryption process of AES begins with a 0-round key addition operation, where the whitening key is XORed with the initial plaintext. This is followed by the repeated application of the round function. Let R denote the round function, which consists of the following four operations:

- **SubBytes** (SB): A non-linear byte substitution applied independently to each byte of the state using an S-box, denoted by S.
- **ShiftRows** (SR): A linear cyclic shift operation. The i-th row of the state is cyclically shifted to the left by i bytes ($0 \leq i \leq 3$).
- **MixColumns** (MC): A linear mixing operation. Each column of the state is multiplied by a constant invertible 4×4 matrix M over $GF(2^8)$.
- **AddRoundKey** (AK): A bitwise XOR operation between the state and the round key.

The overall structure of the round function R is illustrated in Figure 1. In the final round of AES, the MC operation is omitted. We denote the composition of i consecutive round functions as R^i. We also write $R' := SR \circ SB$ for the first-half sublayer of a round, consisting of SB followed by SR. We also write $R^{0.5} := AK \circ SR \circ SB$ for a round with the MC operation omitted.

In this paper, we use the following notations to represent the intermediate states and differences:

- P and C denote the plaintext and ciphertext, respectively. Multiple plaintexts are distinguished by subscripts, such as P^1 and P^2.
- X_r denotes the input state of the r-th round. Due to the 0-round key addition, the input of the 1st round is $X_1 = P \oplus K_0$, where K_0 is the whitening key.
- Y_r, Z_r, and W_r denote the intermediate states after applying SB, SR, and MC in the r-th round, respectively. Thus, $X_{r+1} = W_r \oplus K_r$ holds.

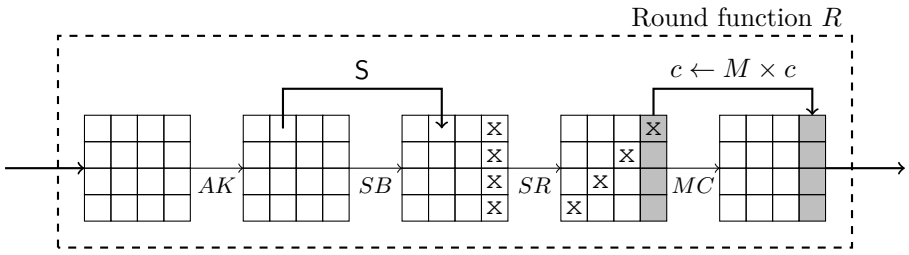


Figure 1: The overall structure of the AES round function R.

- Δ denotes the XOR difference between two states. For instance, the difference of the input states at the r-th round is denoted by $\Delta X_r = X_r^1 \oplus X_r^2$.

To describe the structural properties of AES, we define the index sets for the j-th column, diagonal, and inverse diagonal ($0 \leq j \leq 3$) as follows:

- $Col(j) = \{4j, 4j+1, 4j+2, 4j+3\}$
- $Dia(j) = \{4j, 4((j+1) \bmod 4) + 1, 4((j+2) \bmod 4) + 2, 4((j+3) \bmod 4) + 3\}$
- $InvDia(j) = \{4j, 4((j+3) \bmod 4) + 1, 4((j+2) \bmod 4) + 2, 4((j+1) \bmod 4) + 3\}$

For example, the bytes belonging to the j-th diagonal of a plaintext P are denoted by $P[Dia(j)]$.

Based on these index sets, we define the corresponding subspaces of $\mathbb{F}_{2^8}^{16}$: the *column space* $\mathcal{C}_j$, *diagonal space* $\mathcal{D}_j$, *inverse diagonal space* $\mathcal{ID}_j$, and *mixed space* $\mathcal{M}_j$ are the subspaces whose non-zero bytes are restricted to the indices $Col(j)$, $Dia(j)$, $InvDia(j)$, and $MC(InvDia(j))$, respectively [GRR16]. A *coset* of a subspace $\mathcal{V}$ is a set of the form $\mathcal{V} \oplus a = \{v \oplus a \mid v \in \mathcal{V}\}$ for some fixed state a. By the subspace trail property [GRR16], one round of AES maps a coset of $\mathcal{D}_j$ to a coset of $\mathcal{C}_j$ and a coset of $\mathcal{C}_j$ to a coset of $\mathcal{M}_j$, each with probability 1. Consequently, two rounds of AES map a coset of $\mathcal{D}_j$ to a coset of $\mathcal{M}_j$ with probability 1.

Since the structural property and the distinguisher proposed in this paper do not depend on the key schedule, we describe only the encryption round function and omit details of the key schedule.

2.2 Related Difference and Generalized Zero-Difference Property

In this section, we define the concepts of related difference and related differential, which form the basis of our 5-round distinguisher. Based on these concepts, we explain the characteristics of related differentials and the Generalized Zero-Difference Property.

Definition 1 (Related difference [DR09])**.** Two differences $\Delta x, \Delta x' \in (\mathbb{F}_{2^c})^n$ are called related differences if they satisfy the following condition for all $0 \leq i \leq n-1$:

$$\Delta x_i \Delta x_i' (\Delta x_i \oplus \Delta x_i') = 0.$$

According to this definition, at any word position i, at least one of Δx_i, $\Delta x_i'$, or $\Delta x_i \oplus \Delta x_i'$ must be zero. Using such related differences Δx and $\Delta x'$, we can construct an input quartet consisting of four states: $(\alpha, \alpha \oplus \Delta x, \alpha \oplus \Delta x', \alpha \oplus \Delta x \oplus \Delta x')$.

Definition 2 (Related differential [DR09])**.** Let $(\Delta x, \Delta y)$ and $(\Delta x', \Delta y')$ be two differentials for a linear operation M such that $\Delta y = M(\Delta x)$ and $\Delta y' = M(\Delta x')$. If the input differences $\Delta x, \Delta x'$ are related differences and the output differences $\Delta y, \Delta y'$ also satisfy the related difference condition, this is called a related differential.

The MC operation of AES allows specific forms of related differentials. The four basic types of byte-related differentials defined over $\mathbb{F}_{2^8}^4$ are shown in Table 2. Other related differentials can be derived by rotating or multiplying these four basic types by a scalar value.

Table 2: The sets of byte-related differentials over the AES MC operation [DR09], where $\Delta y = MC(\Delta x)$. The remaining related differentials are derived by rotation and scalar multiplication.

Δx	Δy	$\Delta x'$	$\Delta y'$	$\Delta x \oplus \Delta x'$	$\Delta y \oplus \Delta y'$
$[0,1,4,7]$	$[0,9,0,B]$	$[5,1,0,7]$	$[E,0,D,0]$	$[5,0,4,0]$	$[E,9,D,B]$
$[2,1,1,3]$	$[5,0,4,0]$	$[0,1,0,3]$	$[0,1,4,7]$	$[2,0,1,0]$	$[5,1,0,7]$
$[7,0,7,7]$	$[E,9,0,0]$	$[7,7,7,0]$	$[0,0,E,9]$	$[0,7,0,7]$	$[E,9,E,9]$
$[2,3,2,3]$	$[0,1,0,1]$	$[0,3,2,0]$	$[7,0,7,1]$	$[2,0,0,3]$	$[7,1,7,0]$

Using this, we can construct a 3-round related differential for reduced-round AES, originally introduced in [BR22]. Suppose a state α and related differences $\Delta x, \Delta x'$ have non-zero values only in a specific diagonal, e.g., $Dia(0)$. Here, $\Delta x'$ has non-zero values only at two bytes (the 0-th and 10-th bytes) in the diagonal, while Δx has non-zero values at all four bytes in the diagonal. Let the non-zero byte differences be $\Delta x_0, \Delta x_5, \Delta x_{10}, \Delta x_{15} \in \mathbb{F}_{2^8} \setminus \{0\}$. The differences Δx and $\Delta x'$ can be represented as the following matrices:

$$\Delta x = \begin{bmatrix} \Delta x_0 & 0 & 0 & 0 \\ 0 & \Delta x_5 & 0 & 0 \\ 0 & 0 & \Delta x_{10} & 0 \\ 0 & 0 & 0 & \Delta x_{15} \end{bmatrix}, \qquad \Delta x' = \begin{bmatrix} \Delta x_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \Delta x_{10} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

Using these differences, we construct an input quartet (P^0, P^1, P^2, P^3) such that $P^0 \oplus P^1 = \Delta x$, as follows:

$$P^0 = \alpha, P^1 = \alpha \oplus \Delta x, P^2 = \alpha \oplus \Delta x', P^3 = \alpha \oplus \Delta x \oplus \Delta x'. \tag{1}$$

A quartet constructed from related differences maintains an identical zero-difference pattern after passing through a 2-round SPN structure. To formally define this, we first introduce two functions that represent the zero-difference pattern.

Definition 3 (Zero-Difference Pattern [RBH17])**.** For any state $\alpha \in (\mathbb{F}_{2^c})^n$, the zero-difference pattern $\nu(\alpha) \in \mathbb{F}_2^n$ is defined as:

$$\nu(\alpha)_i = \begin{cases} 1 & \text{if } \alpha_i = 0, \\ 0 & \text{otherwise,} \end{cases}$$

for $0 \le i \le n-1$. Furthermore, for a linear layer L, we define $\mu(\alpha) = \nu(L^{-1}(\alpha))$, which reveals the zero-difference pattern of the state before the final linear transformation L.

With these functions defined, we can now formally state that a quartet constructed from two related differences produces two pairs whose output differences share the same zero-difference pattern after passing through a 2-round SPN.

Theorem 1 (Generalized Zero-Difference Property [BR22])**.** *Let $\alpha \in (\mathbb{F}_{2^c})^n$ and $\Delta x, \Delta x' \in (\mathbb{F}_{2^c})^n$ be two related differences. Let $F = L \circ S \circ L \circ S$ be a 2-round SPN function, where S is a non-linear layer applying n parallel bijective S-boxes and L is a linear layer. Then, the following equation holds:*

$$\mu(F(\alpha) \oplus F(\alpha \oplus \Delta x)) = \mu(F(\alpha \oplus \Delta x') \oplus F(\alpha \oplus \Delta x \oplus \Delta x')).$$

Proof. Let $\beta = \alpha \oplus \Delta x$, $\gamma = \alpha \oplus \Delta x'$, and $\delta = \alpha \oplus \Delta x \oplus \Delta x'$. Since Δx and $\Delta x'$ are related differences, at every word index i at least one of Δx_i, $\Delta x'_i$, or $\Delta x_i \oplus \Delta x'_i$ is zero. We examine the three cases separately to show that the set $\{\alpha[i], \beta[i], \gamma[i], \delta[i]\}$ contains at most two distinct values.

- If $\Delta x_i = 0$, then $\beta[i] = \alpha[i]$ and $\delta[i] = \gamma[i] \oplus \Delta x_i = \gamma[i]$, so $\{\alpha[i], \beta[i], \gamma[i], \delta[i]\} = \{\alpha[i], \gamma[i]\}$.
- If $\Delta x'_i = 0$, then $\gamma[i] = \alpha[i]$ and $\delta[i] = \beta[i] \oplus \Delta x'_i = \beta[i]$, so $\{\alpha[i], \beta[i], \gamma[i], \delta[i]\} = \{\alpha[i], \beta[i]\}$.
- If $\Delta x_i \oplus \Delta x'_i = 0$ (equivalently, $\Delta x_i = \Delta x'_i$), then $\beta[i] = \alpha[i] \oplus \Delta x_i = \alpha[i] \oplus \Delta x'_i = \gamma[i]$ and $\delta[i] = \alpha[i] \oplus (\Delta x_i \oplus \Delta x'_i) = \alpha[i]$, so $\{\alpha[i], \beta[i], \gamma[i], \delta[i]\} = \{\alpha[i], \beta[i]\}$.

In every case, the four values consist of two pairs of equal values. Because each S-box is bijective and operates independently on each word, the outputs at index i also consist of two identical pairs, so their XOR sum is zero. This gives:

$$S(\alpha) \oplus S(\beta) \oplus S(\gamma) \oplus S(\delta) = 0.$$

Since L is linear, applying L preserves the zero sum:

$$L(S(\alpha)) \oplus L(S(\beta)) \oplus L(S(\gamma)) \oplus L(S(\delta)) = 0.$$

Rearranging, the input differences to the second S-box layer are identical:

$$L(S(\alpha)) \oplus L(S(\beta)) = L(S(\gamma)) \oplus L(S(\delta)).$$

Because each S-box is a permutation, it maps a non-zero input difference to a non-zero output difference and a zero input difference to a zero output difference. Since S is a bijection, $S(a)_i = S(b)_i$ holds if and only if $a_i = b_i$, so the zero-difference pattern at each word index is preserved. Therefore, the equality of the input differences implies that the zero-difference patterns of the output differences are identical:

$$\nu(S(L(S(\alpha))) \oplus S(L(S(\beta)))) = \nu(S(L(S(\gamma))) \oplus S(L(S(\delta)))).$$

Finally, by the definition $\mu(\cdot) = \nu(L^{-1}(\cdot))$, we obtain:

$$\mu(F(\alpha) \oplus F(\beta)) = \mu(F(\gamma) \oplus F(\delta)).$$

□

This property can be directly applied to AES. It is well known that 4-round AES can be represented as a 2-round SPN by utilizing the Super-Box representation [DR09, Gil14, RBH17]. By defining a 32-bit superbox operation, 4-round AES can be viewed as an instance of the structure $L \circ S \circ L \circ S$. Consequently, the result of Theorem 1 holds for 4-round AES. Throughout this paper, AES states, plaintexts, ciphertexts, and their differences are written in the byte-level representation $\mathbb{F}_{2^8}^{16}$. In the AES case, the zero-difference pattern μ takes values in $\mathbb{F}_2^4$, where $\mu(\alpha)_i = 1$ if and only if the i-th inverse diagonal of $MC^{-1}(\alpha)$ is zero, for $i \in \{0, 1, 2, 3\}$.

Corollary 1 (Generalized Zero-Difference Property on AES [BR22])**.** *Let $\alpha \in \mathbb{F}_{2^8}^{16}$ and $\Delta x, \Delta x' \in \mathbb{F}_{2^8}^{16}$ be two related differences. Then, the following holds for the 4-round AES function R^4:*

$$\mu(R^4(\alpha) \oplus R^4(\alpha \oplus \Delta x)) = \mu(R^4(\alpha \oplus \Delta x') \oplus R^4(\alpha \oplus \Delta x \oplus \Delta x')).$$

Proof. Four-round AES has the decomposition $R^4 = AK \circ MC \circ SR \circ \mathcal{S} \circ P \circ \mathcal{S} \circ SR \circ AK$, where $\mathcal{S}$ denotes the Super-Box layer consisting of four parallel 32-bit superboxes (each operating independently on one column after SR), and $P = SR \circ AK \circ MC \circ SR$ is a linear permutation [DR09, Gil14]. This matches the 2-round SPN structure $F = L \circ S \circ L \circ S$ of Theorem 1, with $S = \mathcal{S}$ and $L = P$. The outer key additions AK and SR do not affect the related-difference condition nor the zero-difference pattern, since AK cancels in any XOR difference and SR is a fixed byte permutation. The result then follows directly from Theorem 1. □

The equality in Corollary 1 is stated for a specific pair partition of the quartet $(\alpha, \alpha \oplus \Delta x, \alpha \oplus \Delta x', \alpha \oplus \Delta x \oplus \Delta x')$. In fact, the same equality holds for every pair partition, since the proof of Theorem 1 establishes $S(\alpha) \oplus S(\beta) \oplus S(\gamma) \oplus S(\delta) = 0$ at every byte position, which is symmetric in the four states.

Corollary 2 (Joint Generalized Zero-Difference Property on AES [SKS+26])**.** *Under the assumptions of Corollary 1, let $\beta = \alpha \oplus \Delta x$, $\gamma = \alpha \oplus \Delta x'$, and $\delta = \alpha \oplus \Delta x \oplus \Delta x'$. Then, for any partition of the quartet $(\alpha, \beta, \gamma, \delta)$ into two pairs (P, P') and (P'', P'''),*

$$\mu(R^4(P) \oplus R^4(P')) = \mu(R^4(P'') \oplus R^4(P''')).$$

2.3 Mixture and multiple-of-8 property

We use the subspace framework introduced in Section 2.1 to describe the multiple-of-8 property.

The concept of mixture, introduced by Grassi et al. in 2018 [Gra18], refers to a structural relationship established among a quartet of states by combining state values according to specific conditions.

Definition 4 (Mixture [Gra18])**.** Let x and y be two states that differ only in a specific column. For two distinct states z and w, if for all byte indices j $(0 \le j \le 15)$, the values $(x[j], y[j], z[j], w[j])$ always consist of two pairs of equal values, then (z, w) is defined as the mixture counterpart of (x, y), and these four states form a mixture quartet.

Specifically, for each byte j, one of the following two conditions must hold:

1. If $x[j] \neq y[j]$: $\{z[j], w[j]\} = \{x[j], y[j]\}$ holds (i.e., the original values are either maintained or swapped).

2. If $x[j] = y[j]$: $z[j] = w[j]$ holds. In this case, the value does not need to be equal to $x[j]$; as long as the two values are equal to each other, they can freely take any arbitrary value.

Based on this definition of mixture, the input quartet (P^0, P^1, P^2, P^3) constructed from related differences in Section 2.2 satisfies the mixture condition by construction. By examining the byte-wise properties of the two pairs (P^0, P^1) and (P^2, P^3), it can be verified that the active bytes and the bytes with a zero difference fully satisfy the conditions of mixture defined above. We note that related differences and mixture describe different aspects of the same structural relationship: related differences characterize a condition on two difference vectors, while mixture characterizes a relationship between two pairs of states. A quartet constructed from related differences naturally forms a mixture quartet.

Such mixture relationships between states provide the mathematical basis for the multiple-of-8 property of AES, presented at EUROCRYPT 2017 [GRR17].

Theorem 2 (Multiple-of-8 property [GRR17])**.** *Let $E = R^{4.5}$ be the 4.5-round AES encryption function. For any $j \in \{0, 1, 2, 3\}$ and any state a, the number of valid plaintext*

pairs $(P, P') \in (\mathcal{D}_j \oplus a) \times (\mathcal{D}_j \oplus a)$ *whose ciphertexts collide on an inverse diagonal space satisfies the following congruence:*

$$|\{(P, P') \in (\mathcal{D}_j \oplus a) \times (\mathcal{D}_j \oplus a) \mid E(P) \oplus E(P') \in \mathcal{ID}_i\}| \equiv 0 \pmod 8$$

for any $i \in \{0, 1, 2, 3\}$.

Proof. The proof follows [GRR17]. By the subspace trail, R^2 maps a coset of $\mathcal{D}_j$ to a coset of $\mathcal{M}_j$ with probability 1, and symmetrically, $R^{-1.5}$ maps a coset of $\mathcal{ID}_i$ back to a coset of $\mathcal{D}_i$ with probability 1. Therefore, the 4.5-round collision count reduces to a 1-round problem: given all 2^{32} elements in a coset of $\mathcal{M}_j$ after round 2, count the number of pairs whose outputs after a single round R belong to the same coset of $\mathcal{D}_i$.

Each element in the coset of $\mathcal{M}_j$ is parameterized by four independent byte variables. Without loss of generality, assume $j = 0$; then the elements of $\mathcal{M}_0 \oplus a$ are parameterized by $(x, y, z, w) \in (\mathbb{F}_{2^8})^4$ via the MC matrix applied to the 0-th inverse diagonal. Consider two elements Z_2^0 (parameterized by (x, y, z, w)) and Z_2^1 (parameterized by (x', y', z', w')) in this coset such that $R(Z_2^0) \oplus R(Z_2^1) \in \mathcal{D}_i$. By the definition of $\mathcal{M}_0 = MC(\mathcal{ID}_0)$, each column of $\mathcal{ID}_0$ contains exactly one active byte, and since MC operates column-wise, each column of $Z_2 \in \mathcal{M}_0 \oplus a$ depends on exactly one of the four variables (x, y, z, w). Consequently, each byte of the difference $R(Z_2^0) \oplus R(Z_2^1)$ depends only on the variables that differ between Z_2^0 and Z_2^1, since equal variables produce identical byte values that cancel in the XOR. Thus, swapping variables between Z_2^0 and Z_2^1 preserves the difference $R(Z_2^0) \oplus R(Z_2^1)$, and the equal variables can take any value freely without affecting this difference. This leads to a case analysis based on the number of differing variables:

1. When all 4 variables differ ($x \neq x'$, $y \neq y'$, $z \neq z'$, $w \neq w'$): swapping any subset of variables between Z_2^0 and Z_2^1 preserves the output difference, giving $2^{4-1} = 8$ equivalent pairs. The contribution is a multiple of 8.

2. When 3 variables differ and 1 is equal (e.g., $w = w'$): swapping the 3 differing variables gives $2^{3-1} = 4$ equivalent pairs. The equal variable w can take any of 2^8 values without affecting the difference, giving a total contribution of $4 \cdot 2^8$, which is a multiple of 8.

3. When 2 variables differ and 2 are equal (e.g., $z = z'$, $w = w'$): swapping gives $2^{2-1} = 2$ equivalent pairs with 2^{16} free choices for the equal variables, giving $2 \cdot 2^{16}$, a multiple of 8.

4. When only 1 variable differs: by the branch number 5 of MC, a single active input byte produces 4 active output bytes per column, making it impossible to satisfy $R(Z_2^0) \oplus R(Z_2^1) \in \mathcal{D}_i$. No collision occurs.

Since the coset of $\mathcal{M}_0$ contains all 2^{32} elements, the existence of all equivalent pairs is guaranteed. The total number of collisions is the sum of contributions from cases 1–3, each independently a multiple of 8, and the result follows. □

The above proof can be equivalently understood through the mixture framework [Gra18]. The variable-swapping argument in each case corresponds to constructing mixture counterparts at the intermediate state Z_2: the active bytes are either maintained or swapped, while the bytes with a zero difference can freely take any value. A mixture quartet constructed in this way naturally satisfies the related difference condition, since at each byte position at most two distinct values appear among the four states. The generalized zero-difference property (Theorem 1) then guarantees that all pairs generated through this mixture construction share the same zero-difference pattern μ after R^4. This connection between the variable swapping and the mixture relation motivates the framework used throughout this paper.

This variable-swapping view was recently extended by Shin et al. [SKS+25b] for 5-round AES. They observed that a variable swap preserves not only the collision but the entire intermediate difference, so that the pairs in a single multiple-of-8 bundle all follow one common difference through the middle rounds. In the minimal case, this gives the following property.

Theorem 3 (Same-difference property [SKS+25b])**.** *For 5-round AES, suppose exactly* 8 *plaintext pairs in a coset of $\mathcal{D}_j$ have ciphertexts that collide on a fixed inverse diagonal space. Then these* 8 *pairs share the same difference at every state from Y_1 (immediately after the first-round SB) to X_4 (immediately before the fourth-round SB).*

This same-difference structure is the mechanism behind the valid-quartet count established in Section 3.

3 5-Round AES Property from Related Differentials and Its Multiple-of Property

3.1 5-Round AES Property from Related Differentials

$$\begin{bmatrix} \Delta x_0 & 0 & 0 & 0 \\ 0 & \Delta x_5 & 0 & 0 \\ 0 & 0 & \Delta x_{10} & 0 \\ 0 & 0 & 0 & \Delta x_{15} \end{bmatrix} \xrightarrow{R'} \begin{bmatrix} 2\lambda_0 & 0 & 0 & 0 \\ \lambda_0 & 0 & 0 & 0 \\ \lambda_0 & 0 & 0 & 0 \\ 3\lambda_0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{MC} \begin{bmatrix} 5\lambda_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 4\lambda_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} \Delta x_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \Delta x_{10} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{R'} \begin{bmatrix} 2\lambda_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \lambda_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{MC} \begin{bmatrix} 5\lambda_0 & 0 & 0 & 0 \\ \lambda_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 7\lambda_0 & 0 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & \Delta x_5 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \Delta x_{15} \end{bmatrix} \xrightarrow{R'} \begin{bmatrix} 0 & 0 & 0 & 0 \\ \lambda_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 3\lambda_0 & 0 & 0 & 0 \end{bmatrix} \xrightarrow{MC} \begin{bmatrix} 0 & 0 & 0 & 0 \\ \lambda_0 & 0 & 0 & 0 \\ 4\lambda_0 & 0 & 0 & 0 \\ 7\lambda_0 & 0 & 0 & 0 \end{bmatrix}$$

Figure 2: 1-round related differential trails starting with Δx, $\Delta x'$ and $\Delta x \oplus \Delta x'$ respectively. The function R' is $SR \circ SB$.

In this section, we utilize only the first round of the 3-round related differential presented in Section 2.2, and combine it with the 4-round generalized zero-difference property to derive a 5-round AES property based on related differentials. An equivalent property was independently derived by Yan et al. [YTQ26].

We begin by describing the 1-round related differential trail. For a quartet whose differences lie on the diagonal $Dia(0)$, the first round sends the active differences into a single column at Z_1, namely ΔZ_1 and $\Delta Z_1'$; when these match the $(\Delta x, \Delta x')$ of a byte-related differential of Table 2 (up to rotation and scaling), the corresponding MC outputs $\Delta X_2, \Delta X_2'$ are the $(\Delta y, \Delta y')$ of that family, and the four second-round inputs $X_2^0, X_2^1, X_2^2, X_2^3$ form a mixture quartet. Figure 2 shows one such trail, and the remaining families of Table 2 give analogous trails. The probability that a quartet follows such a trail is computed in the proof of the following theorem.

Using this 1-round related differential trail, the core property of 5-round AES can be formally stated as follows.

Theorem 4 (5-Round Property from Related Differentials)**.** *Let $\alpha \in \mathbb{F}_{2^8}^{16}$, and let $\Delta x, \Delta x' \in \mathbb{F}_{2^8}^{16}$ be related differences supported on the diagonal $Dia(0)$ whose active bytes are those of*

a 1-round related differential trail of Table 2 (up to rotation and scaling); Figure 2 shows one such trail. Then, for the 5-round AES encryption function R^5, the following equation holds with a probability of 2^{-22}:

$$\mu(R^5(\alpha) \oplus R^5(\alpha \oplus \Delta x)) = \mu(R^5(\alpha \oplus \Delta x') \oplus R^5(\alpha \oplus \Delta x \oplus \Delta x')). \tag{2}$$

Proof. As in Equation (1) of Section 2.2, let the plaintext quartet formed from the related differences Δx and $\Delta x'$ be $P^0 = \alpha$, $P^1 = \alpha \oplus \Delta x$, $P^2 = \alpha \oplus \Delta x'$, $P^3 = \alpha \oplus \Delta x \oplus \Delta x'$.

When this plaintext quartet undergoes 1-round encryption, the probability that it follows a 1-round related differential trail of Table 2 is 2^{-22}: of the 2^{32} possible column differences at ΔZ_1, only the $4(2^8 - 1)$ valid byte-related forms admit the trail (the four cyclic rotations and the $2^8 - 1$ scalings of a Table 2 family), giving $4 \cdot 2^{-24} = 2^{-22}$; this count is the same for every family of Table 2, so each trail holds with probability 2^{-22}. If the quartet successfully follows this 1-round trail with the probability of 2^{-22}, the internal states $(X_2^0, X_2^1, X_2^2, X_2^3)$ immediately after the first round form a mixture quartet. This state fully fulfills the input condition for the 4-round generalized zero-difference property required in Corollary 1.

Therefore, for the remaining 4-round operation (denoted as R^4) after the first round, the following holds with a probability of 1 according to Corollary 1:

$$\mu(R^4(X_2^0) \oplus R^4(X_2^1)) = \mu(R^4(X_2^2) \oplus R^4(X_2^3)).$$

Consequently, if the 1-round related differential trail is satisfied with a probability of 2^{-22}, the property is maintained with a probability of 1 over the remaining 4 rounds. Thus, from the perspective of the entire 5-round function R^5, the equality in (2) holds with the exact same probability of 2^{-22}. □

3.2 Multiple-of Property of the Quartet Count

Based on the structural property of Theorem 4, we formalize the quartet count as a multiple-of property in Theorem 5, which holds for every pairing of the quartet.

A bundle is the group of valid quartets that necessarily appear together whenever a single base pair satisfies the trail. Its size depends only on n_z, and Theorem 5 below shows that the valid-quartet count decomposes as a non-negative integer combination of the bundle size and 8.

Given the property of Theorem 4, the four plaintexts of the quartet can be partitioned into two pairs in different ways. We call each such partition a *pairing* and identify it by its base pair. In every pairing the two pairs share one difference. For the pairing $(P^0, P^1) \mid (P^2, P^3)$, this common difference is Δx, and similarly $\Delta x'$ for $(P^0, P^2) \mid (P^1, P^3)$ and $\Delta x \oplus \Delta x'$ for $(P^0, P^3) \mid (P^1, P^2)$. We call this common value the *base-pair difference* of the pairing. Writing $\mathrm{wt}(\cdot)$ for the number of active bytes of a difference in its column, the base-pair difference at X_2 has $n_z := 4 - \mathrm{wt}(\Delta X_2)$ inactive bytes.

Table 3 summarizes, for the three pairings of a quartet, the weight wt of the base-pair difference at Z_1 and at X_2 in the two quartet families of Table 2. Family A consists of rows 2 and 4, and Family B of rows 1 and 3. The three pairings of a single quartet thus realize $n_z \in \{0, 1, 2\}$, with a Family-A quartet giving $n_z = 2, 1, 1$ and a Family-B quartet giving $n_z = 2, 2, 0$. The multiple-of count established below depends on the pairing only through n_z.

For every pairing, the following lemma shows that the partner pair is the base pair translated by a fixed offset at the second-round input X_2, and that this translation sends mixture counterparts of the base pair to mixture counterparts of the partner pair; it underlies the count below for all three values of n_z.

Lemma 1 (Translation action at X_2)**.** *Let P^a, P^b, P^c, P^d be the four plaintexts of a quartet of Theorem 4, with second-round inputs $X_2^a, X_2^b, X_2^c, X_2^d$. For a chosen pairing, let*

Table 3: The three pairings of a related-differential quartet (P^0, P^1, P^2, P^3) and their inactive-byte structure, for the two families defined in the text. The base-pair differences of (P^0, P^1), (P^0, P^2), (P^0, P^3) are Δx, $\Delta x'$, $\Delta x \oplus \Delta x'$; for each, the table gives its active-byte weight wt at Z_1 (equal to that at the plaintext diagonal, since *SB* and *SR* preserve it) and at X_2, with $n_z = 4 - \mathrm{wt}(\Delta X_2)$.

Family	base pair	difference	wt at P/Z_1	wt at X_2	n_z
A	(P^0, P^1)	Δx	4	2	2
A	(P^0, P^2)	$\Delta x'$	2	3	1
A	(P^0, P^3)	$\Delta x \oplus \Delta x'$	2	3	1
B	(P^0, P^1)	Δx	3	2	2
B	(P^0, P^2)	$\Delta x'$	3	2	2
B	(P^0, P^3)	$\Delta x \oplus \Delta x'$	2	4	0

the base pair be (X_2^a, X_2^b) and the partner pair (X_2^c, X_2^d). We write $\Delta X_2 = X_2^a \oplus X_2^b$ for the base-pair difference at X_2, whose inactive bytes define n_z, and $\sigma = X_2^a \oplus X_2^c$ for the offset to the partner pair, also a difference at X_2. Then:

(i) $X_2^a \oplus X_2^b \oplus X_2^c \oplus X_2^d = 0$, whether or not the quartet follows a 1-round related differential trail;

(ii) σ is fixed (one of the three pairwise differences of the quartet at X_2) and the partner pair is the base pair translated by it, $(X_2^c, X_2^d) = (X_2^a \oplus \sigma, X_2^b \oplus \sigma)$, so it has the same difference ΔX_2;

(iii) for every mixture counterpart $(\overline{X}_2^a, \overline{X}_2^b)$ of the base pair, the translate $(\overline{X}_2^a \oplus \sigma, \overline{X}_2^b \oplus \sigma)$ is a mixture counterpart of the partner pair;

(iv) the map $x \mapsto x \oplus \sigma$ is a fixed-point-free involution on the mixture counterparts of (X_2^a, X_2^b), since $\sigma \neq 0$ and $\sigma \neq \Delta X_2$.

Proof. (i) The four plaintexts form a mixture at the input: since the quartet has related differences (Definition 1), at every byte j the four bytes $P^a[j], P^b[j], P^c[j], P^d[j]$ take two values, each appearing twice. As *SB* acts on each byte by the same bijection, the four bytes $Y_1^a[j], Y_1^b[j], Y_1^c[j], Y_1^d[j]$ are likewise two equal pairs at every byte, so $Y_1^a \oplus Y_1^b \oplus Y_1^c \oplus Y_1^d = 0$. The map from Y_1 to X_2 is affine ($MC \circ SR$ followed by the round-key addition, whose four copies cancel), so $X_2^a \oplus X_2^b \oplus X_2^c \oplus X_2^d = 0$. This argument uses only the byte-wise structure of the plaintexts, hence holds whether or not the trail is followed.
(ii) The offset $\sigma = X_2^a \oplus X_2^c$ is a difference between the base and partner pairs, hence one of the three pairwise differences of the quartet, and is fixed. By (i), $X_2^d = X_2^a \oplus X_2^b \oplus X_2^c = X_2^b \oplus \sigma$, so $(X_2^c, X_2^d) = (X_2^a \oplus \sigma, X_2^b \oplus \sigma)$ and $X_2^c \oplus X_2^d = X_2^a \oplus X_2^b = \Delta X_2$.
(iii) Let $(\overline{X}_2^a, \overline{X}_2^b)$ be a mixture counterpart of (X_2^a, X_2^b), and set $(\overline{X}_2^c, \overline{X}_2^d) := (\overline{X}_2^a \oplus \sigma, \overline{X}_2^b \oplus \sigma)$. The counterpart satisfies, at every byte j, $\{\overline{X}_2^a[j], \overline{X}_2^b[j]\} = \{X_2^a[j], X_2^b[j]\}$ at the active bytes of ΔX_2 and $\overline{X}_2^a[j] = \overline{X}_2^b[j]$ at its inactive bytes. Adding $\sigma[j]$ and using (ii), at every byte the four values $X_2^c[j], X_2^d[j], \overline{X}_2^c[j], \overline{X}_2^d[j]$ form two equal pairs, so $(\overline{X}_2^c, \overline{X}_2^d)$ is a mixture counterpart of (X_2^c, X_2^d). The four states $\overline{X}_2^a, \overline{X}_2^b, \overline{X}_2^c, \overline{X}_2^d$ carry the same pairwise differences as $X_2^a, X_2^b, X_2^c, X_2^d$ (namely ΔX_2, σ, $\Delta X_2 \oplus \sigma$), hence the same differences at Y_1 since *SR* and *MC* are linear. Thus the plaintexts $\overline{P}^a, \overline{P}^b, \overline{P}^c, \overline{P}^d$ that encrypt to $\overline{X}_2^a, \overline{X}_2^b, \overline{X}_2^c, \overline{X}_2^d$ form a related-differential quartet on $Dia(0)$, that is, an input quartet of Theorem 4.

(iv) Translating twice by σ is the identity. A pair (p, q) with $p \oplus q = \Delta X_2$ is fixed as an unordered pair only if $\sigma = 0$ or $\sigma = \Delta X_2$; but $\sigma \neq 0$ as the four states at X_2 are distinct, and $\sigma \neq \Delta X_2$ as σ and ΔX_2 are distinct pairwise differences. Hence no counterpart is fixed. □

Theorem 5 (Multiple-of Property of the Quartet Count)**.** *Let N_q be the number of valid quartets satisfying the property of Theorem 4 over 5-round AES with cell size w, counted under a fixed pairing whose base-pair difference, when the four second-round states form a mixture quartet at X_2, has n_z inactive bytes ($n_z \in \{0, 1, 2\}$ by Table 3). Then*

$$N_q = 2^{2-n_z}(2^w)^{n_z} A + 8B, \tag{3}$$

where A and B are non-negative integers. In particular, N_q is a multiple of 8 for $n_z \in \{1, 2\}$, and a multiple of 4 for $n_z = 0$.

Proof. By Lemma 1 the four second-round inputs are $X_2^a, X_2^b, X_2^c, X_2^d$ with base pair (X_2^a, X_2^b) and partner pair (X_2^c, X_2^d), base-pair difference at X_2 given by $\Delta X_2 = X_2^a \oplus X_2^b$, and offset $\sigma = X_2^a \oplus X_2^c$; the partner pair has the same difference ΔX_2, and $\sigma \neq 0$, $\sigma \neq \Delta X_2$. We partition the valid quartets according to whether their four states form a mixture quartet at X_2.

Case 1: the four states form a mixture quartet at X_2. Then ΔX_2 has n_z inactive bytes in its column (Table 3), hence $n_a := 4 - n_z$ active bytes, and the partner pair is itself a mixture counterpart of the base pair: at an inactive byte $X_2^a[j] = X_2^b[j]$ forces $X_2^c[j] = X_2^d[j]$, while at an active byte the two-equal-pairs condition gives $\{X_2^c[j], X_2^d[j]\} = \{X_2^a[j], X_2^b[j]\}$. Both the base pair and the partner pair therefore lie in the *same* family $\mathcal{F}$ of mixture counterparts of the base pair, namely the pairs (p, q) with $\{p[j], q[j]\} = \{X_2^a[j], X_2^b[j]\}$ at the active bytes of ΔX_2 and $p[j] = q[j] \in \mathbb{F}_{2^w}$ at its inactive bytes. This family has size $|\mathcal{F}| = 2^{n_a-1}(2^w)^{n_z}$, from 2^{n_a-1} unordered value-pair assignments over the n_a active bytes and $(2^w)^{n_z}$ free common values over the inactive bytes. Since both pairs lie in $\mathcal{F}$ and, by Lemma 1(iv), the translation $x \mapsto x \oplus \sigma$ is a fixed-point-free involution on $\mathcal{F}$, it partitions the $|\mathcal{F}|$ pairs into $|\mathcal{F}|/2 = 2^{2-n_z}(2^w)^{n_z}$ orbits of size two, each orbit $\{(\overline{X}_2^a, \overline{X}_2^b), (\overline{X}_2^c, \overline{X}_2^d)\}$ forming one quartet. By Lemma 1(iii) the plaintexts of each such quartet satisfy the input condition of Theorem 4; and since the base pair collides, every mixture counterpart of it in $\mathcal{F}$ collides on the same inverse diagonal, by Corollary 1. Both pairs of each generated quartet therefore collide, so the quartet is valid. This case contributes valid quartets in bundles of $|\mathcal{F}|/2 = 2^{2-n_z}(2^w)^{n_z}$, giving the term $2^{2-n_z}(2^w)^{n_z} A$.

Case 2: the four states do not form a mixture quartet at X_2. Then the partner pair is not a mixture counterpart of the base pair. Let ΔX_2 be active on m bytes. The mixture counterparts of the base pair, $\mathcal{F}_0$, and of the partner pair, $\mathcal{F}_1$, both consist of pairs of difference ΔX_2, so

$$|\mathcal{F}_0| = |\mathcal{F}_1| = 2^{m-1}(2^w)^{4-m},$$

a multiple of 8 for every $m \geq 1$. By Lemma 1(iii) the translation σ maps $\mathcal{F}_0$ into $\mathcal{F}_1$, and being an injective involution between sets of equal size it maps $\mathcal{F}_0$ bijectively onto $\mathcal{F}_1$. The two sets are disjoint: a common pair would equal $\{X_2^a[j], X_2^b[j]\}$ and $\{X_2^c[j], X_2^d[j]\}$ at every active byte and be constant at every inactive byte, making the four states a mixture quartet, contrary to assumption. Since the original quartet is valid, the base and partner pairs both collide, so each pair of $\mathcal{F}_0$ and of $\mathcal{F}_1$ collides as well, by Corollary 1. Pairing each member of $\mathcal{F}_0$ with its σ-image in $\mathcal{F}_1$ produces a valid quartet; distinct members give distinct quartets, none repeated since $\mathcal{F}_0 \cap \mathcal{F}_1 = \varnothing$. This case contributes $|\mathcal{F}_0| = 2^{m-1}(2^w)^{4-m}$ valid quartets per bundle, a multiple of 8, giving the term $8B$.

Cases 1 and 2 partition all valid quartets, so $N_q = 2^{2-n_z}(2^w)^{n_z} A + 8B$. □

3.3 Experimental Verification

To validate the structural properties established in Theorems 4 and 5, we conducted experimental analyses on both small-AES [CMR05] and the standard AES block cipher. As a proxy for a random permutation, we use the full 10-round variant of (small-)AES, which we treat as behaving like a random permutation in view of the decades-long security record of full-round AES. The 5-round and 10-round (small-)AES used in our experiments omit the MC operation in the final round, following the standard AES convention; consequently a zero difference on an inverse diagonal of the ciphertext directly realizes the zero-difference pattern μ, and no MC^{-1} needs to be applied when testing the property. The count of Theorem 4 is verified below using the Family-A pairing with base pair (P^0, P^1) of Table 3 (difference Δx, 4-active in Figure 2; $n_z = 2$); the same figure also depicts the Family-A pairing (P^0, P^2), with difference $\Delta x'$ ($n_z = 1$). The multiple-of property of Theorem 5 is verified for all three cases of Table 3.

3.3.1 Verification on Small-AES ($w = 4$)

Small-AES [CMR05] is a scaled-down variant of AES that operates on a 4×4 matrix of 4-bit nibbles, resulting in a 64-bit block size. Its round function follows the same structure as the standard AES, consisting of SB, SR, MC, and AK. In particular, the MC operation uses the same MDS matrix $[2, 1, 1, 3]$ over $GF(2^4)$, which ensures that the related differentials described in Table 2 also exist for small-AES. Moreover, 4-round small-AES admits a Super-Box representation analogous to the standard AES, and can therefore be viewed as a 2-round SPN structure. Consequently, the generalized zero-difference property (Theorem 1) applies to 4-round small-AES as well. The 1-round related differential trail in Figure 2 also applies to small-AES, but with a different probability: the total number of possible column differences is 2^{16}, and the number of valid ones is 2^4 (with 4 independent cases from byte rotation), giving a probability of $4 \cdot 2^{-12} = 2^{-10}$.

We utilized a structure of 2^{16} plaintexts (a full diagonal space) and performed 1,000 independent experiments for both 5-round small-AES and a random permutation (10-round small-AES). In each experiment, we counted the number of valid quartets (P^0, P^1, P^2, P^3), where (P^0, P^1) is a pair whose plaintext difference satisfies the input condition Δx (i.e., 4 active bytes in a specific diagonal) and whose ciphertexts (C^0, C^1) have a zero difference in a specific inverse diagonal, and (P^2, P^3) is the pair generated by swapping the 0-th and 10-th bytes between P^0 and P^1, whose ciphertexts (C^2, C^3) also have a zero difference in the same inverse diagonal. By construction, the quartet (P^0, P^1, P^2, P^3) automatically realizes the three input differences Δx, $\Delta x'$, and $\Delta x \oplus \Delta x'$ of Figure 2.

Verification of Theorem 4. We first verify the property of Theorem 4 by comparing the observed average number of valid quartets with the theoretical expected value. Within a structure of 2^{16} plaintexts, there are $\binom{2^{16}}{2} \approx 2^{31}$ plaintext pairs. For 5-round small-AES, the probability that a pair has a zero difference in one of the 4 inverse diagonals is $4 \cdot 2^{-16} = 2^{-14}$. For each such pair, the swapped pair satisfies the same inverse diagonal condition either through the structural property of Theorem 4 (with probability 2^{-10}) or, when the structural path does not apply, by a random collision (with probability 2^{-16}); the combined probability is approximately 2^{-10}. A valid quartet is counted twice. It is counted once when (P^0, P^1) serves as the base pair, and once when (P^2, P^3) does, since swapping the 0-th and 10-th bytes of (P^2, P^3) recovers (P^0, P^1). We therefore divide by 2 to obtain the expected number of unique valid quartets:

$$2^{31} \cdot 4 \cdot 2^{-16} \cdot 2^{-10}/2 = 2^6 = 64.$$

For a random permutation, the swapped pair collides solely by chance with probability 2^{-16}. Applying the same factor of $1/2$ to account for the double-counting of quartets, the

expected quartet count is:

$$2^{31} \cdot 2^{-16} \cdot 2^{-14}/2 = 1.$$

The experimentally observed averages over 1,000 trials are 73.34 for 5-round small-AES and 0.73 for the random permutation. These values are consistent with the theoretical expectations of 64 and 1, respectively, confirming that the property of Theorem 4 holds for small-AES. The slight overshoot above 64 for 5-round small-AES is attributable to sample-mean noise from the high variance imposed by the 256-bundle (see Table 8).

Verification of Theorem 5. By Theorem 5, the valid quartet count takes the form $N_q = 2^{2-n_z}(2^w)^{n_z}A + 8B$, whose bundle $2^{2-n_z}(2^w)^{n_z}$ equals 256, 32, and 4 for $n_z = 2, 1, 0$ at $w = 4$. We verify the three cases of Table 3: the Family-A pairings with base pairs (P^0, P^1) ($n_z = 2$) and (P^0, P^2) ($n_z = 1$), and the Family-B pairing with base pair (P^0, P^3) ($n_z = 0$), each over 1,000 structures of 2^{16} plaintexts. Table 4 summarizes the result, and Appendix A gives the full frequency tables (Tables 8, 9, and 10). In each case the 5-round counts are multiples of the bundle: for $n_z = 2$ and 1 they appear at the bundle multiples $256, 512, 768$ and $32, 64, 96$ (the A term), and the $8B$ term adds residues such as 264 and 40; for $n_z = 0$ they include $12, 28, 44$, which are multiples of 4 but not of 8, so the bundle is genuinely 4 rather than the standard multiple of 8. The detection rate $\Pr[N_q \geq 1]$, the fraction of structures with at least one valid quartet, grows from 0.30 to 0.86 to 1.00 as n_z decreases from 2 to 1 to 0, so a smaller bundle gives a more reliable distinguisher; this drives the data-complexity analysis of Section 4. The mean count moves the other way, increasing with the bundle (about 46, 60, and 73 for $n_z = 0, 1, 2$), so the smaller bundle is the more reliable even though it produces fewer valid quartets on average. The random permutation gives at most 6 valid quartets per structure, and these counts are not multiples of the bundle (for instance $1, 2, 3$), unlike the 5-round counts.

Table 4: Small-AES verification of the multiple-of property (Theorem 5) for the three cases of Table 3, with 1,000 structures of 2^{16} plaintexts per case. Parenthesized numbers are frequencies over the 1,000 structures, and $\Pr[N_q \geq 1]$ is the fraction with at least one valid quartet. The 5-round counts are multiples of the bundle in $N_q = 2^{2-n_z}(2^w)^{n_z}A + 8B$, whereas the random counts (≤ 6) are not; see the text for discussion.

case (Table 3)	n_z	bundle	observed 5-round counts (frequency)	$\Pr[N_q \geq 1]$	random
Family A, (P^0, P^1)	2	256	256 (198), 512 (20), 768 (8), 264 (13)	0.30	≤ 6
Family A, (P^0, P^2)	1	32	32 (283), 64 (246), 96 (158), 40 (14)	0.86	≤ 5
Family B, (P^0, P^3)	0	4	44 (148), 36 (96), 28 (54), 12 (1)	1.00	≤ 5

3.3.2 Verification on Standard AES ($w = 8$)

To validate the property on the standard AES block cipher ($w = 8$) for the $n_z = 2$ pairing, we performed experiments using a structure of 2^{32} plaintexts over 200 independent trials each for 5-round AES and a random permutation (10-round AES), following the same counting setup as Section 3.3.1.

Verification of Theorem 4. We verify the property of Theorem 4 by comparing the observed average number of valid quartets with the theoretical expected value. Within a structure of 2^{32} plaintexts, there are $\binom{2^{32}}{2} \approx 2^{63}$ plaintext pairs. For 5-round AES, the probability that a pair has a zero difference in one of the 4 inverse diagonals is $4 \cdot 2^{-32} = 2^{-30}$. For each such pair, the swapped pair satisfies the same inverse diagonal condition either through the structural property of Theorem 4 (with probability 2^{-22}) or, when the structural path does not apply, by a random collision (with probability 2^{-32});

the combined probability is approximately 2^{-22}. Applying the same 1/2 factor as in Section 3.3.1:

$$2^{63} \cdot 4 \cdot 2^{-32} \cdot 2^{-22}/2 = 2^{10} = 1024.$$

For a random permutation, the swapped pair collides solely by chance with probability 2^{-32}. Applying the same factor of 1/2 to account for the double-counting of quartets, the expected quartet count is:

$$2^{63} \cdot 2^{-32} \cdot 2^{-30}/2 = 1.$$

The experimentally observed averages are 328.12 for 5-round AES and 0.99 for the random permutation (each over 200 trials). The average for the random permutation closely matches the theoretical expectation of 1. For 5-round AES, the observed average is lower than the theoretical expectation of 1024; this gap arises from the structural form of N_q established in Theorem 5, which states that valid quartets occur in bundles of size $(2^w)^2 = 65536$ or of size 8. Since each 65536-bundle contributes 65536 quartets at once, the observed mean within only 200 trials is determined almost entirely by how many 65536-bundles happen to occur. The observed 328.12 decomposes as one 65536-bundle that occurred across the 200 trials (contributing $65536/200 \approx 327.68$) together with 11 small bundles of size 8 (contributing $11 \cdot 8/200 = 0.44$). The mean is therefore not a reliable indicator at this sample size, and we instead assess the property of Theorem 4 through the observed frequency of 65536-bundles, as analyzed below.

We quantify the expected frequency of 65536-bundles as follows. Since each occurrence of a 65536-bundle corresponds to a structural instance contributing $(2^w)^2 = 65536$ quartets, and the total expected number of quartets per structure is 1024, the expected number of 65536-bundles per structure is $1024/65536 = 2^{-6} \approx 0.0156$. Let X denote the number of 65536-bundles observed across the 200 trials. Then X follows the binomial distribution $B(200, 2^{-6})$ with $\mathbb{E}[X] = 200 \cdot 2^{-6} \approx 3.12$, and

$$\Pr[X = 1] = \binom{200}{1} \cdot 2^{-6} \cdot (1 - 2^{-6})^{199} \approx 0.136.$$

Our experiment observes $X = 1$, which has $\Pr[X = 1] \approx 0.136$ under this distribution and is therefore consistent with the property of Theorem 4 for the standard AES.

Verification of Theorem 5. Based on Theorem 5, the valid quartet count N_q for the standard AES should follow $N_q = (2^8)^2 A + 8B = 65536A + 8B$. The frequencies of the observed valid quartet counts are summarized in Table 5.

Table 5: Frequency of valid quartet counts in AES ($w = 8$, Structure size 2^{32}, 200 trials each)

5-round AES		Random Permutation (10-round)	
Count (N_q)	Frequency	Count (N_q)	Frequency
0	188	0	79
8	11	1	66
65536	1	2	35
		3	18
		4	2

For 5-round AES, the observed non-zero quartet counts are 8 and 65536, both of which are consistent with $N_q = 65536A + 8B$: the count 8 corresponds to the $8B$ term (with $A = 0$, $B = 1$), while the count 65536 corresponds to the $65536A$ term (with $A = 1$, $B = 0$). In particular, the single occurrence of the count 65536 observed across 200 trials is consistent with the $65536A$ component of the theorem, which is otherwise difficult

to observe within a limited number of trials. Conversely, the results for the random permutation exhibit a completely unstructured distribution, with observed counts such as 1, 2, and 3 that are not multiples of 8, confirming that the multiple-of-8 phenomenon is a unique structural characteristic intrinsic to 5-round AES. This standard-AES experiment realizes the Family-A (P^0, P^1) case of Table 3 ($n_z = 2$), whose 65536-bundle is the most demanding to observe within a limited number of trials. The Family-A (P^0, P^2) ($n_z = 1$) and Family-B (P^0, P^3) ($n_z = 0$) cases are confirmed on small-AES in Appendix A and hold for $w = 8$ by the cell-size-independent proof of the theorem.

4 Chosen-Plaintext Distinguishers for 5-Round AES

4.1 Distinguishing Strategy

A *valid quartet* is a base pair together with a partner pair that both collide on the same inverse diagonal of the ciphertext. Such quartets arise for both 5-round AES and a random permutation, but at different rates. For a random permutation the two collisions are independent, so a candidate quartet is valid with probability about $2^{-30} \cdot 2^{-32} = 2^{-62}$, where 2^{-30} is the chance that the base pair collides on one of the four inverse diagonals and 2^{-32} is the chance that the partner collides on the same one. For 5-round AES, Theorem 4 raises the per-partner success probability from 2^{-32} to $\lambda = 2^{-22}$ once the base pair follows the trail, so the chance of a valid quartet becomes $2^{-30} \cdot 2^{-22} = 2^{-52}$, larger by a factor of 2^{10}. Over a full structure of 2^{32} plaintexts this amounts to about 2^{10} expected valid quartets for 5-round AES against 1 for a random permutation (Section 3.3.2), and the distinguisher detects this difference. The multiple-of property of Theorem 5 is exact only over this full structure; with fewer plaintexts the count is no longer a multiple of the bundle, so the distinguisher relies on the presence of a valid quartet, which the 2^{10} rate gap keeps far more likely for 5-round AES than for a random permutation.

In the chosen-plaintext setting we exploit the excess as follows. We build a structure, encrypt it, and search it for one valid quartet under a fixed pairing of Table 3. Sorting the ciphertexts on each inverse diagonal exposes the base pairs that collide; the partner of each base pair is formed by the offset of the pairing and already lies in the structure, so its ciphertexts are read off without further queries. If at least one valid quartet is found, the target is classified as 5-round AES. If none is found, it is classified as a random permutation. Algorithm 1 states the search for a general pairing; its cost is $O(N_s \log N_s)$ memory accesses per inverse diagonal, comparable to N_s encryptions [GRR17], and the three pairings differ only in the offset that forms the partner.

How few plaintexts are needed depends on the pairing. The distinguisher succeeds when a structure contains at least one valid quartet, so what matters is the detection rate $\Pr[N_q \geq 1]$ rather than the mean count. The three pairings share the per-partner success probability $\lambda = 2^{-22}$ (Theorem 4) and differ only in how those quartets are grouped. By Theorem 5 a base pair that follows the trail makes all of its mixture counterparts valid together, forming a bundle of size $2^{2-n_z}(2^w)^{n_z}$, and this bundle reshapes the distribution of N_q. A large bundle (large n_z) makes valid quartets rare but many at once, so a single structure is almost always empty and detection needs many repeated structures. A small bundle (small n_z) makes them frequent but few at a time, keeping the count close to its mean, so a single structure usually contains one. The detection rate therefore rises as the bundle becomes smaller, which the small-AES experiment confirms, with $\Pr[N_q \geq 1] = 0.30, 0.86, 1.00$ for $n_z = 2, 1, 0$ (Section 3.3.1). Hence $n_z = 0$ and $n_z = 1$, with small bundles, are distinguished within a single structure, while $n_z = 2$, with the largest bundle, needs many.

Section 4.2 compares the three pairings, $n_z = 2$, $n_z = 1$ which is the pairing of Yan et al. [YTQ26], and $n_z = 0$, on a single structure.

Algorithm 1 Distinguishing 5-Round AES under a chosen pairing in the CP setting

Input: Oracle access to target cipher E, Structure size N_s
Output: 5-round AES or Random Permutation
Construct a structure S of N_s plaintexts that vary only the four diagonal bytes $\{0, 5, 10, 15\}$, keeping the other 12 bytes constant, arranged so that the partner of every candidate pair under the chosen pairing also lies in S.
for each plaintext $P^i \in S$ **do**
    $C^i \leftarrow E(P^i)$
end for
for each inverse diagonal index $d \in \{0, 1, 2, 3\}$ **do**
    Sort S by $C^i[InvDia(d)]$
    for each candidate pair (P^0, P^1) of consecutive entries with equal $C^i[InvDia(d)]$ that is active as the chosen pairing requires **do**
        Form the partner pair (P^2, P^3) from (P^0, P^1) by the offset of the pairing
        Read C^2, C^3 from S
        if $C^2 \oplus C^3$ has a zero difference in $InvDia(d)$ **then**
            return 5-round AES
        end if
    end for
end for
return Random Permutation

4.2 Comparison of the Three Pairings

We run the distinguisher of Algorithm 1 for the three pairings of Table 3 on a single structure of about 2^{27} plaintexts and compare the per-structure success probability P_{find}. The three differ only in how the partner of a base pair is formed: for $n_z = 2$ the base is active on all four diagonal bytes and the partner swaps bytes $\{0, 10\}$; for $n_z = 1$ [YTQ26] the base is active on $\{0, 10\}$ and the partner shifts $\{5, 15\}$; for $n_z = 0$ the base is active on $\{5, 15\}$ and the partner swaps byte 15 while moving to another $\{0, 10\}$ value.

Table 6: Per-structure success probability of the three pairings on a single structure (5-round AES). Each entry uses 300 trials.

Pairing	bundle size	P_{find} at $\sim 2^{27}$	P_{find} at $\sim 2^{27.2}$
$n_z = 2$	2^{16}	0.01	0.02
$n_z = 1$ [YTQ26]	512	0.53	0.64
$n_z = 0$	4	0.57	0.65

In theory the three pairings reach the same expected count in a structure of about 2^{27} plaintexts, so the single-structure success is set by the spread of the count, which the bundle controls. With the smallest bundle, $n_z = 0$ keeps the count closest to its mean and succeeds most often, with P_{find} rising from 0.57 at 2^{27} to 0.65 at $2^{27.2}$. The pairing of Yan et al. ($n_z = 1$) follows closely, 0.53 and 0.64, the last close to the 63% they report at $2^{27.2}$; its bundle of 512 concentrates the count into fewer structures and leaves slightly more of them empty. The largest bundle, $n_z = 2$, concentrates the count so strongly that a single structure almost never contains a valid quartet, so P_{find} stays at only a few percent and many repeated structures are needed.

5 Adaptive Chosen-Plaintext Distinguishers for 5-Round AES

5.1 Distinguishing Strategy

In the adaptively chosen-plaintext (ACP) setting the attacker may choose each query after seeing earlier ciphertexts. This changes what sets the data cost. In the chosen-plaintext setting of Section 4 every partner of a base pair has to be present in one fixed structure, so a large bundle forces a large structure; in the adaptive setting the distinguisher first finds the base pairs that collide and then queries only their partners.

For a chosen pairing the distinguisher proceeds as in Algorithm 2. It builds chosen-plaintext structures over the base-active bytes, encrypts them, and collects every base pair whose ciphertexts collide on an inverse diagonal. For each such base pair it forms the partners prescribed by the offset of the pairing, queries them adaptively, and tests whether any of them collides on the same inverse diagonal. If a valid quartet is found, the target is classified as 5-round AES; otherwise it is classified as a random permutation.

By Theorem 4, once a base pair is found to collide on an inverse diagonal, each partner pair formed from it collides on the same inverse diagonal with probability $\lambda = 2^{-22}$. About 2^{22} partner pairs, gathered across the base collisions found, then give one valid quartet on average. When the offset of the pairing ranges over a two-byte space, as for $n_z = 1$ and $n_z = 0$, each base collision produces up to 2^{16} partners, so the 2^{22} partner pairs come from only about 2^6 base collisions. A base structure varies only the two base-active bytes, so it holds at most 2^{16} plaintexts and gives about 2 base collisions; taking 2^5 such structures (2^{21} chosen plaintexts) reaches 2^6, and with the roughly 2^{23} adaptive partner queries the data complexity is about $2^{23.3}$ [YTQ26]. The $n_z = 2$ pairing has no such offset: its partner is a single swap of bytes $\{0, 10\}$, so each base collision produces only one partner pair. Reaching 2^{22} partner pairs would then require about 2^{22} base collisions, hence a structure of about $2^{26.5}$ plaintexts, far larger than the 2^{21} used by $n_z = 0$ and $n_z = 1$. Section 5.2 compares the three pairings experimentally.

Algorithm 2 Distinguishing 5-Round AES under a chosen pairing in the ACP setting

Input: Oracle access to target cipher E, total base-structure size N_b
Output: 5-round AES or Random Permutation
Construct chosen-plaintext structures of N_b plaintexts in total that vary only the base-active bytes of the chosen pairing, keeping the other bytes constant.
for each plaintext P in the structures **do**
    $C \leftarrow E(P)$
end for
for each inverse diagonal index $d \in \{0, 1, 2, 3\}$ **do**
    Sort the structures by $C[InvDia(d)]$
    for each base pair (P^0, P^1) with equal $C[InvDia(d)]$ that is active as the chosen pairing requires **do**
        for each partner offset of the chosen pairing **do**
            Form the partner pair (P^2, P^3) by applying the offset to (P^0, P^1)
            Query $C^2 \leftarrow E(P^2)$ and $C^3 \leftarrow E(P^3)$ adaptively
            if $C^2 \oplus C^3$ has a zero difference in $InvDia(d)$ **then**
                return 5-round AES
            end if
        end for
    end for
end for
return Random Permutation

5.2 Comparison of the Three Pairings

We evaluate the adaptive distinguisher of Algorithm 2 for the three pairings of Table 3 at the data complexity $2^{23.32}$, the point at which $n_z = 0$ and $n_z = 1$ reach one expected valid quartet (2^{21} base plaintexts together with 2^{23} adaptive queries), and report the success probability P_{find} over 300 trials in Table 7.

Table 7: Success probability of the adaptive distinguisher of Algorithm 2 for the three pairings at data complexity $2^{23.32}$ (5-round AES, 300 trials).

Pairing	P_{find} at $2^{23.32}$
$n_z = 2$	0.01
$n_z = 1$ [YTQ26]	0.59
$n_z = 0$	0.50

At this complexity the pairing of Yan et al. ($n_z = 1$) succeeds with probability 0.59, in line with the 63% they report [YTQ26], and $n_z = 0$ reaches 0.50. Both attain it because their two-byte offset turns each base collision into about 2^{16} partner pairs, so a few base collisions already supply the 2^{22} partner pairs needed for one valid quartet. The swap of $n_z = 2$ supplies only one partner pair per base collision, so at the same $2^{23.32}$ it almost never succeeds (about 0.01); matching the others would require about $2^{26.5}$ plaintexts.

The small-AES experiment shows that the mean increases with the bundle (Section 3.3.1), so $n_z = 1$ produces more valid quartets than $n_z = 0$ at the same data, even though both pairings share the same per-partner success probability $\lambda = 2^{-22}$. In a single structure this difference is hidden, because the larger bundle of $n_z = 1$ makes its valid quartets rare but many at once, so one structure is empty more often than for $n_z = 0$ despite the higher mean, leaving $n_z = 0$ the more reliable (Section 4.2). In the adaptive setting the partner pairs come from many separate base collisions, so the bundle does not form and the count stays close to its mean for both pairings. The detection rate then tracks the mean, and the higher mean of $n_z = 1$ makes it the better pairing here. The pairing $n_z = 2$ stays far behind both. Its offset forms only one partner per base collision, too few to reach the 2^{22} partner pairs that one valid quartet needs, so it fails for a reason unrelated to the bundle.

6 Conclusion

In this paper, we studied related-differential distinguishers for 5-round AES based on the composition of a 1-round related differential trail with the 4-round generalized zero-difference property. We proved that the valid-quartet count N_q takes the multiple-of form $N_q = 2^{2-n_z}(2^w)^{n_z}A + 8B$, where w is the cell size and n_z is the number of inactive bytes of ΔX_2 under the chosen pairing of the 4-plaintext quartet into two pairs. The value of n_z controls the bundle size, and a smaller n_z reduces the variance of the valid-quartet count.

We applied the distinguisher to the three pairings. The $n_z = 2$ pairing, with the largest bundle, almost never succeeds in a single structure, whereas the smaller-bundle pairings distinguish within a single structure of about 2^{27} plaintexts. In a single structure $n_z = 0$, with the smallest bundle, is the most reliable, while $n_z = 1$ is the concurrent distinguisher of Yan et al. [YTQ26] at $2^{27.2}$. In the adaptive setting the bundle does not form, since the partner pairs are gathered across separate base collisions, so $n_z = 1$, which produces more valid quartets, becomes the better pairing, reversing the single-structure ordering. All claims are verified experimentally on both small-AES ($w = 4$) and the standard AES ($w = 8$).

The choice of pairing directly affects the variance of related-differential distinguishers. Future work includes generalizing the multiple-of bundle analysis to other SPN ciphers and

investigating whether pairing choice has a similar effect in other cryptanalytic techniques.

References

[BBI+15] Subhadeep Banik, Andrey Bogdanov, Takanori Isobe, Kyoji Shibutani, Harunaga Hiwatari, Toru Akishita, and Francesco Regazzoni. Midori: A block cipher for low energy. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 411–436. Springer, 2015.

[BBS99] Eli Biham, Alex Biryukov, and Adi Shamir. Cryptanalysis of skipjack reduced to 31 rounds using impossible differentials. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 12–23. Springer, 1999.

[BCC19] Christina Boura, Anne Canteaut, and Daniel Coggia. A general proof framework for recent aes distinguishers. *IACR Transactions on Symmetric Cryptology*, 2019(1):170–191, 2019.

[BJK+16] Christof Beierle, Jérémy Jean, Stefan Kölbl, Gregor Leander, Amir Moradi, Thomas Peyrin, Yu Sasaki, Pascal Sasdrich, and Siang Meng Sim. The skinny family of block ciphers and its low-latency variant mantis. In *Annual international cryptology conference*, pages 123–153. Springer, 2016.

[BL23] Augustin Bariant and Gaëtan Leurent. Truncated boomerang attacks and application to aes-based ciphers. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 3–35. Springer, 2023.

[BODK+20] Achiya Bar-On, Orr Dunkelman, Nathan Keller, Eyal Ronen, and Adi Shamir. Improved key recovery attacks on reduced-round aes with practical data and memory complexities: A. bar-on et al. *Journal of Cryptology*, 33(3):1003–1043, 2020.

[BR19] Navid Ghaedi Bardeh and Sondre Rønjom. The exchange attack: How to distinguish six rounds of aes with chosen plaintexts. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 347–370. Springer, 2019.

[BR22] Navid Ghaedi Bardeh and Vincent Rijmen. New key-recovery attack on reduced-round aes. *IACR Transactions on Symmetric Cryptology*, pages 43–62, 2022.

[BS91] Eli Biham and Adi Shamir. Differential cryptanalysis of des-like cryptosystems. *Journal of CRYPTOLOGY*, 4(1):3–72, 1991.

[CMR05] Carlos Cid, Sean Murphy, and Matthew JB Robshaw. Small scale variants of the aes. In *International Workshop on Fast Software Encryption*, pages 145–162. Springer, 2005.

[DKRS20] Orr Dunkelman, Nathan Keller, Eyal Ronen, and Adi Shamir. The retracing boomerang attack. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 280–309. Springer, 2020.

[DR02] Joan Daemen and Vincent Rijmen. *The design of Rijndael*, volume 2. Springer, 2002.

[DR09] Joan Daemen and Vincent Rijmen. New criteria for linear maps in aes-like ciphers. *Cryptography and Communications*, 1(1):47–69, 2009.

[Gil14] Henri Gilbert. A simplified representation of aes. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 200–222. Springer, 2014.

[Gra18] Lorenzo Grassi. Mixture differential cryptanalysis: a new approach to distinguishers and attacks on round-reduced aes. *IACR Transactions on Symmetric Cryptology*, pages 133–160, 2018.

[GRR16] Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. Subspace trail cryptanalysis and its applications to aes. *Cryptology ePrint Archive*, 2016.

[GRR17] Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. A new structural-differential property of 5-round aes. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 289–317. Springer, 2017.

[Knu94] Lars R Knudsen. Truncated and higher order differentials. In *International Workshop on Fast Software Encryption*, pages 196–211. Springer, 1994.

[Knu98] Lars Knudsen. Deal-a 128-bit block cipher. 1998.

[MRSA23] Sandip Kumar Mondal, Mostafizar Rahman, Santanu Sarkar, and Avishek Adhikari. Revisiting yoyo tricks on aes. *IACR Transactions on Symmetric Cryptology*, 2023(4):28–57, 2023.

[RBH17] Sondre Rønjom, Navid Ghaedi Bardeh, and Tor Helleseth. Yoyo tricks with aes. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 217–243. Springer, 2017.

[SKK+24] Hanbeom Shin, Insung Kim, Sunyeop Kim, Seonggyeom Kim, Deukjo Hong, Jaechul Sung, and Seokhie Hong. Revisiting the multiple-of property for skinny: The exact computation of the number of right pairs. *IEEE Access*, 12:35502–35511, 2024.

[SKS+25a] Hanbeom Shin, Seonkyu Kim, Byoungjin Seok, Dongjae Lee, Deukjo Hong, Jaechul Sung, and Seokhie Hong. New exchanged boomerang distinguishers for 5-round aes. *Cryptology ePrint Archive*, 2025.

[SKS+25b] Hanbeom Shin, Sunyeop Kim, Byoungjin Seok, Dongjae Lee, Deukjo Hong, Jaechul Sung, and Seokhie Hong. Key-recovery attack on 5-round aes with multiple-of-8 property. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 2025.

[SKS+26] Hanbeom Shin, Sunyeop Kim, Byoungjin Seok, Deukjo Hong, Jaechul Sung, Seokhie Hong, Sangjin Lee, and Dongjae Lee. Key-independent secret-key distinguisher for 7-round AES based on the joint generalized zero-difference property. Cryptology ePrint Archive, Paper 2026/980, 2026. URL: https://eprint.iacr.org/2026/980.

[Wag99] David Wagner. The boomerang attack. In *International Workshop on Fast Software Encryption*, pages 156–170. Springer, 1999.

[YTQ26] Xueping Yan, Lin Tan, and Wenfeng Qi. Related-differential distinguishers on up to 7 rounds of AES. Cryptology ePrint Archive, Paper 2026/1039, 2026. URL: https://eprint.iacr.org/2026/1039.

[YTXQ24] Xueping Yan, Lin Tan, Hong Xu, and Wenfeng Qi. The boomerang chain distinguishers: New record for 6-round aes. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 301–329. Springer, 2024.

A Detailed Small-AES Frequency Tables

This appendix lists the full valid-quartet frequency tables behind Table 4, for the three pairings $n_z = 2, 1, 0$ of small-AES ($w = 4$), each measured over 1,000 structures of 2^{16} plaintexts. The 5-round counts are multiples of the bundle $2^{2-n_z}(2^w)^{n_z}$, namely 256, 32, and 4; the random permutation (10-round) gives unstructured counts.

Table 8: Frequency of valid quartet counts in small-AES for the Family-A pairing (P^0, P^1) of Table 3 ($n_z = 2$, bundle 256; $w = 4$, 1,000 trials).

5-round Small-AES		Random Permutation (10-round)	
Count (N_q)	Frequency	Count (N_q)	Frequency
0	704	0	494
8	50	1	336
24	1	2	124
64	1	3	40
256	198	4	5
264	13	6	1
272	1		
512	20		
520	4		
768	8		

Table 9: Frequency of valid quartet counts in small-AES for the Family-A pairing (P^0, P^2) of Table 3 ($n_z = 1$, bundle 32; $w = 4$, 1,000 trials).

5-round Small-AES		Random Permutation (10-round)	
Count (N_q)	Frequency	Count (N_q)	Frequency
0	138	0	456
8	10	1	352
32	283	2	144
40	14	3	35
64	246	4	11
72	30	5	2
80	1		
96	158		
104	18		
128	60		
136	5		
160	24		
168	3		
192	6		
232	2		
256	2		

Table 10: Frequency of valid quartet counts in small-AES for the Family-B pairing (P^0, P^3) of Table 3 ($n_z = 0$, bundle 4; $w = 4$, 1,000 trials). All 5-round counts (shown in two column pairs) are multiples of 4.

5-round Small-AES				Random (10-round)	
N_q	Freq	N_q	Freq	N_q	Freq
8	1	60	53	0	435
12	1	64	37	1	353
16	3	68	28	2	147
20	10	72	13	3	49
24	39	76	14	4	13
28	54	80	16	5	3
32	69	84	4		
36	96	88	2		
40	118	92	1		
44	148	164	2		
48	110	172	1		
52	96	176	1		
56	83				