Beyond Affine Invariants: A Hamming-Weight Correlation Metric for Template-CPA Leakage in Key-Dependent S-boxes*

Wiesław Maleszewski

Department of Computer Science, Faculty of Computer Science and Technology, University of Lomza, Lomza, Poland

wmaleszewski@al.edu.pl

ORCID: 0000-0001-8852-3532

Abstract

Classical selection criteria for cryptographic S-boxes—nonlinearity NL, differential uniformity DU, boomerang uniformity BU, algebraic degree deg—are invariants of affine equivalence. That property is exactly what blinds them to a class of side-channel weaknesses. The correlation-power-analysis (CPA) template distinguisher is governed by the Hamming-weight functional, and Hamming weight is not affine-invariant; it does not descend to the affine-equivalence quotient on which the classical criteria live. Two S-boxes with identical $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ can therefore leak differently under template CPA.

We make this precise for the key-dependent family $S^{\mathcal{G}}(x) = A\,\iota(x) \oplus c$, with ι the multiplicative inverse in $\mathrm{GF}(2^8)$ and $(A, c) \in \mathrm{GL}(8, \mathbb{F}_2) \times \mathbb{F}_2^8$ drawn from a byte stream $\mathcal{G}$. Proposition 2 fixes the four invariants at $(112, 4, 6, 7)$ across the entire family; they carry no information about $\mathcal{G}$. We introduce the Hamming-weight template correlation $\rho_{\mathrm{HW}}(\cdot, S_{\mathrm{AES}})$ (Definition 3), identify it as the population statistic controlling the AES-template CPA distinguisher, and show that it resolves the fiber the classical invariants collapse (Proposition 6).

As a stress test we instantiate $\mathcal{G}$ with three sources of contrasting regularity—a system CSPRNG, a discretised logistic map, and a $\sin(1/x)$/xxHash hybrid—and sample 3×10^5 S-boxes from a single master seed. The classical invariants are identical everywhere, as predicted. The metric is not. The logistic source widens the ρ_{HW} distribution against S_{AES} by 12–13% ($\sigma_\ell = 0.0704$ vs. 0.0626/0.0623; Levene $p < 10^{-180}$). The widening vanishes against a uniform-random reference permutation (Levene $p > 0.13$), survives an exact $Q\,1.31$ fixed-point reimplementation at 3.1%, and does not appear for a tent-map control. Propagated through the Mangard–Oswald–Popp trace-budget model and checked against a 2.16×10^5-attack Monte-Carlo CPA simulation, it yields a 29% relative excess in AES-template success rate at SNR = 10, $N = 10^3$ (empirical ratio 1.29, analytic 1.26). By every standard effect-size measure the widening is small (Cohen's $d = 0.128$ on $|\rho_{\mathrm{HW}}|$, Cohen's $h = 0.130$ on the attackable fraction); its significance is detectability, not magnitude. The contribution is a measurement axis, not a weak generator: a metric that flags template-CPA leakage where NL = 112, DU = 4 report perfect scores.

Keywords: S-box evaluation, affine-equivalence invariants, correlation power analysis, Hamming-weight leakage, key-dependent S-boxes, side-channel-aware metrics.
MSC 2020: 94A60, 11T71, 94A55.

*A preliminary version of this work was presented at the International Scientific Conference Innovation Trends 2026 (IT2026) [15]. The present article is a substantially extended version, adding the affine-invariance blindness result (Lemma 5, Proposition 6), a standardised effect-size and countermeasure baseline (Section 4.8), a family of leakage metrics (Section 4.9), a survey of the CPA-evaluation gap in recent S-box constructions (Section 1.1), and a Perron–Frobenius analysis with a parameter sweep (Section 5.1).

1 Introduction

The security assessment of an S-box is usually reduced to a few algebraic figures of merit: nonlinearity NL, differential uniformity DU, boomerang uniformity BU, algebraic degree deg [2, 4, 3, 5]. Each is an invariant of affine equivalence. The point is elementary and rarely stated: if $S' = A_1 \circ S \circ A_2$ for affine permutations A_1, A_2, then S and S' agree on all four. Invariance is what makes these quantities good design targets. It is also what makes them blind. The CPA template attacker predicts leakage through the Hamming-weight map HW, and HW is preserved only under the monomial subgroup of $\mathrm{GL}(8, \mathbb{F}_2)$—not under a generic change of affine representative. So the CPA-relevant statistic does not factor through the affine-equivalence quotient on which $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ are defined. A designer who certifies an S-box on those four alone can ship the optimal profile $\mathrm{NL} = 112$, $\mathrm{DU} = 4$ together with an unmeasured template-CPA exposure. We formalise that blindness, give a metric that removes it, and validate the metric on a key-dependent family where the classical invariants are provably constant.

The Rijndael S-box $S_{\mathrm{AES}} \colon \mathbb{F}_2^8 \to \mathbb{F}_2^8$ admits the standard decomposition $S_{\mathrm{AES}}(x) = L_{\mathrm{AES}}(\iota(x))$, where ι denotes multiplicative inversion in $\mathrm{GF}(2^8)$ (extended by $\iota(0) = 0$) and L_{AES} is an affine map over $\mathbb{F}_2$. Cryptanalytic resistance to differential and linear attacks is carried by ι; the role of L_{AES} is to prevent the S-box from collapsing to a polynomial of low degree over $\mathbb{F}_2$. The composition has the well-documented profile $\mathrm{NL}(S_{\mathrm{AES}}) = 112$, $\mathrm{DU}(S_{\mathrm{AES}}) = 4$, $\mathrm{BU}(S_{\mathrm{AES}}) = 6$, $\deg(S_{\mathrm{AES}}) = 7$ [2, 4, 3, 5].

In the adaptive and key-dependent S-box paradigm [9, 10], motivated largely by adaptive security for resource-constrained IoT deployments [13, 14], ι is retained while L_{AES} is replaced by a variable affine layer $L_{A,c}(y) = Ay \oplus c$, with $(A, c) \in \mathrm{GL}(8, \mathbb{F}_2) \times \mathbb{F}_2^8$ drawn from a pseudorandom byte stream $\mathcal{G}$. The resulting family is

$$S^{\mathcal{G}}(x) = A\,\iota(x) \oplus c, \qquad A \in \mathrm{GL}(8, \mathbb{F}_2), \quad c \in \mathbb{F}_2^8. \tag{1}$$

Proposition 2 (Section 2) gives the invariance of $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ across this family. The choice of $\mathcal{G}$ is therefore not determined by classical cryptographic metrics.

A secondary observable, relevant to side-channel attacks, is the correlation between the Hamming-weight vectors of $S^{\mathcal{G}}$ and S_{AES}. Correlation power analysis [7, 6, 8] operates by acquiring power traces during S-box evaluation and computing the Pearson correlation between trace samples and the predicted leakage $\mathrm{HW}(S(p \oplus k))$ for each key candidate k. When the implementer deploys $S^{\mathcal{G}}$ and the attacker templates against S_{AES}, the relevant population statistic is the cross-correlation $\rho_{\mathrm{HW}}(S^{\mathcal{G}}, S_{\mathrm{AES}})$ between the two Hamming-weight vectors. While the CPA model is standard, the dependence of this statistic on $\mathcal{G}$ has not been characterised in the literature.

The present study formulates the question quantitatively and tests it on a population of 10^5 instances per stream. The target metric is the Hamming-weight correlation

$$\rho_{\mathrm{HW}}(S, S_{\mathrm{AES}}) = \mathrm{corr}\big((\mathrm{HW}(S(0)), \ldots, \mathrm{HW}(S(255))),\ (\mathrm{HW}(S_{\mathrm{AES}}(0)), \ldots, \mathrm{HW}(S_{\mathrm{AES}}(255)))\big), \tag{2}$$

which expresses the alignment between the attacker's Hamming-weight hypothesis and the deployed S-box. A replication check is also reported on the table-level Pearson correlation

$$\rho_{\mathrm{raw}}(S, S_{\mathrm{AES}}) = \mathrm{corr}\big((S(0), \ldots, S(255)),\ (S_{\mathrm{AES}}(0), \ldots, S_{\mathrm{AES}}(255))\big), \tag{3}$$

which an earlier exploratory note suggested might distinguish streams; the metric has no direct cryptanalytic interpretation, as $S(x)$ takes values in $\mathrm{GF}(2^8)$ rather than in $\mathbb{R}$.

Three byte streams are compared: a uniform cryptographic stream $\mathcal{G}_{\mathrm{u}}$ (the system CSPRNG), a discrete logistic map $\mathcal{G}_{\ell}$, and a hybrid stream $\mathcal{G}_{\mathrm{s}}$ that combines a $\sin(1/x)$ look-up table with xxHash-style multiplicative mixers [11, 12]. The three streams span a range of statistical regularities — from cryptographically uniform, through chaotic with autocorrelated low-order

bits, to an algebraically structured hybrid — and therefore probe the dependence of the CPA-relevant correlation ρ_{HW} on the byte-pair distribution induced by $\mathcal{G}$ rather than on any particular construction choice. The dataset comprises 10^5 instances per stream (3×10^5 total), sampled reproducibly from a single seed.

At this scale the three streams separate clearly. The logistic stream $\mathcal{G}_\ell$ produces a ρ_{HW} distribution against S_{AES} whose standard deviation exceeds that of $\mathcal{G}_{\mathrm{u}}$ and $\mathcal{G}_{\mathrm{s}}$ by 12–13%, a difference far above the detection threshold for the test (Section 4). The same comparison against a uniform-random reference S-box yields a null distribution: the bias is specific to the algebraic structure of A_{AES}, not to the marginal distribution of (A, c). A secondary observation, consistent with the same source of bias, is the 3.95% rate of single-row matches with A_{AES} for $\mathcal{G}_\ell$, against the analytical baseline of $1-(1-1/255)^8 \approx 3.09\%$ that the other two streams reproduce within 0.07 percentage points.

The remainder of the paper is organised as follows. Section 2 fixes notation, proves the invariance proposition, and states the cryptographic hypothesis tested. Section 3 describes the experimental protocol and sampling power. Section 4 reports results. Section 5 discusses the source of the AES-specific bias, the mitigation, and the limitations of the analysis.

1.1 The evaluation gap in recent S-box constructions

The construction of key-dependent and chaos-derived S-boxes is an active line [10, 16, 17, 18, 19, 20]. The reported evaluation is uniform across the recent literature: bijectivity, nonlinearity, differential and linear approximation probability, SAC, BIC, algebraic degree and immunity. Table 1 summarises the axes these works report. Two observations follow.

First, the reported axes are, up to normalisation, the affine-equivalence invariants of Section 2.3 together with the diffusion descriptors SAC and BIC. By Proposition 6 the former are constant on any family of the form $L \circ \iota$ and therefore cannot rank constructions inside such a family. Second, template-CPA leakage $\rho_{\mathrm{HW}}(\cdot, S_{\mathrm{AES}})$ is absent from all of them. Where a side-channel claim appears at all, it is made through a SAC/BIC proxy rather than a CPA distinguisher; SAC and BIC measure diffusion, not the alignment of the Hamming-weight profile with a template reference, and the two are independent (a construction can be SAC-optimal and template-aligned at once). The metric proposed here fills that gap: it is the first axis in this line that separates constructions the classical battery reports as identical.

Table 1: Evaluation axes reported by recent chaos-/key-dependent S-box constructions. NL: nonlinearity; DU: differential uniformity; LP: linear probability; SAC/BIC: diffusion criteria; CPA: correlation-power-analysis template leakage. A check denotes the axis is reported; a dash denotes it is absent.

work	NL	DU	LP	SAC	BIC	CPA/ρ_{HW}
Zahid et al. [10]	✓	✓	✓	✓	✓	—
Lightweight chaos [16]	✓	✓	✓	✓	✓	—
1D chaos + field [17]	✓	✓	✓	✓	✓	—
Confusion source [18]	✓	✓	✓	✓	✓	—
Hyperchaos lattice [19]	✓	✓	✓	✓	✓	—
Enhanced logistic [20]	✓	✓	✓	✓	✓	—
this work	✓	✓	✓	✓	✓	✓

2 Construction and hypothesis

2.1 Setting

Let $\mathbb{F}_2 = \{0,1\}$ and $\mathbb{F}_2^8$ the eight-dimensional vector space over $\mathbb{F}_2$, with bytes identified with elements of $\mathbb{F}_2^8$ via the natural binary representation. The field $\mathbb{F}_{2^8}$ is realised as $\mathbb{F}_2[x]/(x^8+x^4+x^3+x+1)$, with $\iota\colon \mathbb{F}_{2^8} \to \mathbb{F}_{2^8}$ the multiplicative inverse extended by $\iota(0)=0$. The cryptographic profile of ι is $\mathrm{NL}(\iota) = 112$, $\mathrm{DU}(\iota) = 4$, $\mathrm{BU}(\iota) = 6$, $\deg(\iota) = 7$ [2, 4, 3, 5]. For $y \in \mathbb{F}_2^8$ let $\mathrm{HW}(y) \in \{0,\dots,8\}$ denote the Hamming weight.

Definition 1 (Random affine layer on the AES inverse)**.** Given a pseudorandom byte stream $\mathcal{G}$, sample nine consecutive bytes; assign the first eight to the rows of $A \in \mathbb{F}_2^{8\times 8}$ and the ninth to $c \in \mathbb{F}_2^8$. The sample is accepted iff $A \in \mathrm{GL}(8,\mathbb{F}_2)$ and re-drawn otherwise. The resulting S-box is $S^{\mathcal{G}}(x) = A\,\iota(x) \oplus c$.

The streams compared are: $\mathcal{G}_{\mathrm{u}}$ (uniform bytes from the `secrets` module of CPython); $\mathcal{G}_{\ell}$ (iterates of $x_{n+1} = r\,x_n(1-x_n)$ with $r \in (3.57, 4)$, in IEEE-754 `binary64`, quantised by $\lfloor 256x \rfloor$); and $\mathcal{G}_{\mathrm{s}}$ (a 2^{16}-entry sin look-up combined with three 64-bit xxHash-style multipliers, `0x9E3779B185EBCA87`, `0xC2B2AE3D27D4EB4F`, and `0x165667B19E3779F9`, with the high byte of the mixed state taken as output).

2.2 Invariance proposition

Proposition 2. *For every $A \in \mathrm{GL}(8,\mathbb{F}_2)$ and $c \in \mathbb{F}_2^8$, the S-box $S(x) = A\,\iota(x) \oplus c$ satisfies $\mathrm{NL}(S) = 112$, $\mathrm{DU}(S) = 4$, $\mathrm{BU}(S) = 6$, $\deg(S) = 7$.*

Proof. For NL, the identity $\langle b, S(x)\rangle = \langle A^{\top} b, \iota(x)\rangle \oplus \langle b, c\rangle$ implies that $A^{\top}$ permutes the non-zero output masks. The constant $\langle b, c\rangle$ shifts the Walsh transform $\widehat{f_b}$ by a sign, preserving $\max_a |\widehat{f_b}(a)|$. Hence $\mathrm{NL}(S) = \mathrm{NL}(\iota) = 112$.

For BU, expand the boomerang-connectivity equation for $S = L_{A,c} \circ \iota$ with $L_{A,c}(y) = Ay \oplus c$, $L_{A,c}^{-1}(y) = A^{-1}(y \oplus c)$, and the auxiliary substitution $\beta' = A^{-1}\beta$:

$$\begin{aligned} S^{-1}(S(x)\oplus\beta) &= \iota^{-1}(L_{A,c}^{-1}(L_{A,c}(\iota(x)) \oplus \beta)) \\ &= \iota\big(A^{-1}(A\,\iota(x) \oplus c \oplus \beta \oplus c)\big) = \iota(\iota(x) \oplus \beta'), \\ S^{-1}(S(x\oplus\alpha)\oplus\beta) &= \iota(\iota(x\oplus\alpha)\oplus\beta') \quad \text{by the same identity.} \end{aligned}$$

Therefore the event $\{S^{-1}(S(x)\oplus\beta)\oplus S^{-1}(S(x\oplus\alpha)\oplus\beta) = \alpha\}$ coincides with $\{\iota(\iota(x)\oplus\beta')\oplus\iota(\iota(x\oplus\alpha)\oplus\beta') = \alpha\}$, giving $\mathrm{BCT}_S(\alpha,\beta) = \mathrm{BCT}_\iota(\alpha,\beta')$. The map $\beta \mapsto \beta'$ is a bijection of $\mathbb{F}_2^8 \setminus \{0\}$ since $A \in \mathrm{GL}(8,\mathbb{F}_2)$, and hence

$$\mathrm{BU}(S) \;=\; \max_{\alpha\neq 0,\,\beta\neq 0} \mathrm{BCT}_S(\alpha,\beta) \;=\; \max_{\alpha\neq 0,\,\beta'\neq 0} \mathrm{BCT}_\iota(\alpha,\beta') \;=\; \mathrm{BU}(\iota) \;=\; 6,$$

where $\mathrm{BU}(\iota) = 6$ is the result of Cid et al. [3] and Boura–Canteaut [4].

For DU, the differential $S(x\oplus\alpha)\oplus S(x) = A(\iota(x\oplus\alpha)\oplus\iota(x))$ is the A-image of the corresponding differential of ι; the bijection A permutes the DDT rows and preserves $\max_\beta \#\{x : \iota(x\oplus\alpha)\oplus\iota(x) = \beta\}$, hence $\mathrm{DU}(S) = \mathrm{DU}(\iota) = 4$. The algebraic degree is a standard invariant of left affine post-composition [5]. □

2.3 The template-CPA metric and the blindness of affine invariants

Definition 3 (Hamming-weight template correlation)**.** For an S-box $S\colon \mathbb{F}_2^8 \to \mathbb{F}_2^8$ and a reference S_{ref}, put $\mathbf{h}_S = (\mathrm{HW}(S(0)),\dots,\mathrm{HW}(S(255))) \in \mathbb{R}^{256}$. The Hamming-weight template correlation is the Pearson coefficient

$$\rho_{\mathrm{HW}}(S, S_{\mathrm{ref}}) \;=\; \frac{\big\langle \mathbf{h}_S - \bar{\mathbf{h}}_S,\; \mathbf{h}_{S_{\mathrm{ref}}} - \bar{\mathbf{h}}_{S_{\mathrm{ref}}} \big\rangle}{\big\|\mathbf{h}_S - \bar{\mathbf{h}}_S\big\|_2 \,\big\|\mathbf{h}_{S_{\mathrm{ref}}} - \bar{\mathbf{h}}_{S_{\mathrm{ref}}}\big\|_2}. \tag{4}$$

Set $\rho_{\mathrm{HW}} = 0$ when either vector is constant.

Remark 4 (Operational meaning)**.** $\rho_{\mathrm{HW}}(S^{\mathcal{G}}, S_{\mathrm{AES}})$ is not an auxiliary descriptor. It is the population statistic that governs the AES-template CPA distinguisher. Let the device evaluate $S^{\mathcal{G}}(p \oplus k)$ under Hamming-weight leakage, and let the attacker correlate the traces against the hypothesis $\mathrm{HW}(S_{\mathrm{AES}}(p \oplus k_g))$ over uniform plaintext p. The noise-free signal at the correct key equals $\rho_{\mathrm{HW}}(S^{\mathcal{G}}, S_{\mathrm{AES}})$, attenuated by $\sqrt{\mathrm{SNR}/(\mathrm{SNR}+1)}$ under additive Gaussian noise (Eq. (6)). The metric sits on the path from S-box choice to trace budget.

Write $\mathrm{HW}(y) = \sum_{i=0}^{7} y_i$ for the coordinate weight on $\mathbb{F}_2^8$, valued in $\mathbb{Z}$.

Lemma 5 (HW is not an affine invariant)**.** *Let $A \in \mathrm{GL}(8, \mathbb{F}_2)$. Then $\mathrm{HW}(Ay) = \mathrm{HW}(y)$ for all $y \in \mathbb{F}_2^8$ if and only if A is a permutation matrix. For generic $A \in \mathrm{GL}(8, \mathbb{F}_2)$ the profile $y \mapsto \mathrm{HW}(Ay)$ differs from $y \mapsto \mathrm{HW}(y)$, and $\mathrm{HW}(Ay \oplus c)$ depends non-trivially on the representative (A, c).*

Proof. Permutation matrices permute coordinates and preserve $\sum_i y_i$. Conversely, suppose $\mathrm{HW}(Ay) = \mathrm{HW}(y)$ for all y. Take $y = e_j$. Then $\mathrm{HW}(Ae_j) = 1$, so every column of A has weight 1; invertibility over $\mathbb{F}_2$ forces the columns to be distinct unit vectors, hence A is a permutation matrix. The failure is generic: permutation matrices have relative measure $8!/|\mathrm{GL}(8, \mathbb{F}_2)| < 10^{-9}$ in $\mathrm{GL}(8, \mathbb{F}_2)$. □

Proposition 6 (Blindness of the classical invariants)**.** *Let $\mathcal{G}$ be a byte stream and let $(A, c) \sim \mathcal{G}$ induce $S^{\mathcal{G}} = A\,\iota \oplus c$.*

(i) The pushforward of $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ under $\mathcal{G}$ is the Dirac mass $\delta_{(112,4,6,7)}$, independent of $\mathcal{G}$. Hence $I(\mathcal{G}; (\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)) = 0$, and any statistic derived from the four invariants has zero variance across the family.

(ii) The functional $(A, c) \mapsto \rho_{\mathrm{HW}}(S^{A,c}, S_{\mathrm{AES}})$ is non-constant on $\mathrm{GL}(8, \mathbb{F}_2) \times \mathbb{F}_2^8$. Hence $\mathrm{Var}_{\mathcal{G}}(\rho_{\mathrm{HW}}) > 0$, and whenever two streams induce distinct byte-pair laws, $I(\mathcal{G}; \rho_{\mathrm{HW}}) > 0$.

The classical criteria and ρ_{HW} split the same family into, respectively, one equivalence class and a non-degenerate distribution. The classical axis distinguishes no two members; the proposed axis does.

Proof. Part (i) is Proposition 2: every (A, c) yields the AES-inverse profile, so the induced law is a point mass and the mutual information vanishes. Part (ii): by Lemma 5 the vector $\mathbf{h}_{S^{A,c}}$ depends on (A, c) beyond its affine-equivalence class, so $\rho_{\mathrm{HW}}(S^{A,c}, S_{\mathrm{AES}})$ is a non-constant function of (A, c). A witness is any row collision $\mathrm{row}_i(A) = \mathrm{row}_i(A_{\mathrm{AES}})$, which aligns the i-th bit profile with AES and lifts $|\rho_{\mathrm{HW}}|$ above its baseline (Section 4.4). Positive mutual information for streams with distinct consecutive-byte laws is then the empirical separation of Table 3. □

Remark 7 (The unit of novelty)**.** Proposition 6 states the contribution at the right altitude. The claim is not that a particular generator is weak. It is that the evaluation lattice $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ is a quotient on which the template-CPA leakage functional is constant, and that ρ_{HW} is a computable lift resolving that quotient. The family of Section 4 and its logistic instantiation are an existence proof: optimal classical profile, measurable template-CPA bias, detected only by ρ_{HW}.

2.4 The cryptographic hypothesis

The substantive hypothesis tested in this paper is the equality of the CPA-relevant distribution across streams.

Hypothesis 1. The pseudorandom source $\mathcal{G}$ does not affect the distribution of the Hamming-weight correlation $\rho_{\mathrm{HW}}(S^{\mathcal{G}}, S_{\mathrm{AES}})$.

The raw correlation ρ_{raw} in (3) is reported as a replication check on an earlier exploratory claim; it is not a cryptographic figure of merit.

3 Experimental setup

Sampling. For each of the three streams $\mathcal{G} \in \{\mathcal{G}_{\text{u}}, \mathcal{G}_{\ell}, \mathcal{G}_{\text{s}}\}$ we generate exactly $N = 10^5$ accepted instances of (A, c) following the procedure stated below. The total dataset is 3×10^5 S-boxes; together with the AES reference S_{AES} and the fixed control permutation S_{rand} (drawn once from the same master seed by a Fisher–Yates shuffle of $\{0, \ldots, 255\}$), the analysis population is $3 \cdot 10^5 + 2$.

The procedure for generating one accepted S-box is:

1. draw nine consecutive bytes $(b_0, \ldots, b_8)$ from $\mathcal{G}$;
2. form the 8×8 binary matrix A whose row i is the byte b_i;
3. test invertibility of A over $\mathbb{F}_2$ by Gauss–Jordan rank reduction; if singular, return to step 1;
4. set $c \leftarrow b_8$ and return the table $S\colon x \mapsto A\,\iota(x) \oplus c$.

Acceptance rate is the probability that a uniform-random 8×8 binary matrix is invertible, $\prod_{k=0}^{7}(1 - 2^{k-8}) \approx 0.2898$, so the expected number of nine-byte draws per accepted sample is ≈ 3.45. No other inclusion or exclusion criteria are applied; all accepted samples enter the analysis.

Reproducibility. A single master seed (`SEED = 20260517`) controls every byte stream, with each 2500-instance block sub-seeded deterministically; the platform CSPRNG is replaced by a seeded `random.Random` for the duration of the experiment. The complete experimental directory, including the byte sources, sampling driver, metric implementations, statistical tests and plotting scripts, is organised as the directory `IT2026-paper-work/` in a public code repository released upon acceptance; the corresponding scripts are referenced in-line: `analysis/13_unified_N100k.py` for the main population sampling (SHA-256 of output CSV: `cd06f92c80dae3703fef26dcd58906694e693ed2814f311592064c6a9944befe`), `analysis/17_int_logistic.py` for the integer-only $Q\,1.31$ control (Section 4.5), `analysis/16_tent_control.py` for the tent-map control (Section 4.6), `analysis/15_mtd_analysis.py` for the analytical MTD table (Table 4), and `analysis/18_cpa_simulation.py` for the Monte-Carlo CPA validation (Figure 3). Re-running any of these scripts from the master seed reproduces the corresponding results bit-identically. Floating-point operations in $\mathcal{G}_{\ell}$ and the sine table in $\mathcal{G}_{\text{s}}$ use IEEE-754 `binary64`; the deterministic output is guaranteed only for the CPython `math` module on the tested platform (macOS, Apple silicon). Tool-chain: Python 3.13.5, NumPy 2.4, SciPy 1.17, Matplotlib 3.10.

Measurements. For each S-box S the quantities $\rho_{\text{raw}}(S, S_{\text{AES}})$, $\rho_{\text{HW}}(S, S_{\text{AES}})$, $\rho_{\text{HW}}(S, S_{\text{rand}})$, the eight per-bit correlations

$$\rho_b^{(i)}(S, S_{\text{AES}}) = \operatorname{corr}\big((\operatorname{bit}_i S(0), \ldots, \operatorname{bit}_i S(255)), (\operatorname{bit}_i S_{\text{AES}}(0), \ldots, \operatorname{bit}_i S_{\text{AES}}(255))\big), \tag{5}$$

and the row-collision indicator $\mathbf{1}[\exists i : \operatorname{row}_i(A) = \operatorname{row}_i(A_{\text{AES}})]$ are recorded. The per-bit summary is $\rho_b^{\max} = \max_i |\rho_b^{(i)}|$.

Statistical tests. For each pairwise comparison between streams we report Levene's test (median-centred) for variance equality and the two-sample Kolmogorov–Smirnov test for distributional equality. With three pairwise comparisons per metric, the Bonferroni-adjusted significance

threshold is $\alpha = 0.05/3 \approx 0.0167$. At $N = 10^5$ per group, the F-test variance-ratio threshold at the same level is $\sqrt{F_{0.975,N-1,N-1}} \approx 1.006$; standard-deviation ratios exceeding this threshold are statistically detectable.

4 Results

The results are reported in seven blocks. Section 4.1 records the four classical invariants as a consistency check on Proposition 2. Section 4.2 reports the replication check on the historical metric ρ_{raw}. Section 4.3 contains the principal result, the distribution of ρ_{HW} against S_{AES} and against the random control, on which Hypothesis 1 is tested. Section 4.4 reports the per-bit and row-collision observables that share the same source as ρ_{HW}. Section 4.5 reports the integer-arithmetic control on the logistic stream, and Section 4.6 the discrete tent-map control. Section 4.7 converts the observed ρ_{HW} distributions to operational CPA trace budgets and validates the analytical prediction against a Monte-Carlo simulation. All numerical values reported below are derived from the dataset of Section 3 without further sub-selection.

4.1 Empirical confirmation of the invariance proposition

Across all $3{\times}10^5$ instances generated for the three streams, the values NL = 112, DU = 4, BU = 6, deg = 7 are observed with zero sample variance (Table 2). The values coincide with the AES profile [2, 4, 3, 5] and with the AES reference table reconstructed from A_{AES} and $c_{\text{AES}} =$ `0x63` ($S_{\text{AES}}(0) =$ `0x63`, $S_{\text{AES}}(1) =$ `0x7C`). Direct computation of BU on the full population is not performed; the value is fixed by Proposition 2.

Table 2: Affine invariants observed on $N = 10^5$ instances per stream ($n = 3 \times 10^5$ total) plus the AES reference. Every cell takes the same value with zero sample variance; the empirical population SD is 0 to the precision of double-precision arithmetic.

stream	NL	DU	BU	deg
$\mathcal{G}_{\text{u}}$ (uniform CSPRNG)	112	4	6	7
$\mathcal{G}_{\ell}$ (logistic)	112	4	6	7
$\mathcal{G}_{\text{s}}$ ($\sin(1/x)$ hybrid)	112	4	6	7
S_{AES} (Rijndael reference)	112	4	6	7

4.2 Table-level Pearson

The table-level Pearson correlation $\rho_{\text{raw}}(S, S_{\text{AES}})$ is not a cryptanalytic figure of merit; it is reported as a replication check on an exploratory observation made on $N = 40$ samples. At $N = 10^5$, pairwise Levene tests yield $p \geq 0.07$ and pairwise Kolmogorov–Smirnov tests yield $p \geq 0.20$ for all three stream pairs. The null hypothesis of distributional equality is not rejected at the conventional 0.05 significance level, and the exploratory claim is not supported by the larger sample. Maximum $|\rho_{\text{raw}}|$ reaches 0.94 for $\mathcal{G}_{\text{s}}$, 0.85 for $\mathcal{G}_{\text{u}}$, and 0.79 for $\mathcal{G}_{\ell}$, with no stable ranking. The earlier exploratory claim against $\mathcal{G}_{\text{s}}$ on the basis of ρ_{raw} is not supported.

4.3 Hamming-weight Pearson

Under the CPA-relevant metric, the empirical distributions of ρ_{HW} against S_{AES} differ significantly between $\mathcal{G}_{\ell}$ and the remaining two streams: pairwise Levene tests yield $p < 10^{-180}$ and pairwise Kolmogorov–Smirnov tests yield $p < 10^{-70}$, both exceeding the Bonferroni-adjusted significance threshold $\alpha_{\text{B}} = 0.05/3$ by many orders of magnitude. The standard deviation of $\rho_{\text{HW}}(\mathcal{G}_{\ell}, S_{\text{AES}})$ exceeds the corresponding values for $\mathcal{G}_{\text{u}}$ and $\mathcal{G}_{\text{s}}$ by 13%. Repetition of the test against the

uniform-random reference S_{rand} yields no statistically significant difference at the same threshold (Table 3, lower block). Hypothesis 1 is rejected under the AES reference and not rejected under the random control; numerical detail is given in Table 3 and Figure 1.

Table 3: Per-stream distribution of ρ_{HW} on the $N = 10^5$ reproducible sample, with pairwise Levene/KS tests below each block.

stream	mean	std	min	max	$\lvert\rho\rvert_{\max}$
Reference: S_{AES} (Rijndael)					
$\mathcal{G}_{\text{u}}$ (uniform CSPRNG)	−0.000	0.0626	−0.500	+0.375	0.500
$\mathcal{G}_{\ell}$ (logistic)	−0.004	0.0704	−0.375	+0.375	0.375
$\mathcal{G}_{\text{s}}$ (sin hybrid)	−0.000	0.0623	−0.375	+0.500	0.500
Pairwise Levene/KS p:	$\mathcal{G}_{\ell}$ vs. $\mathcal{G}_{\text{s}}$: 2×10^{-191} / 2×10^{-79};		$\mathcal{G}_{\ell}$ vs. $\mathcal{G}_{\text{u}}$: 1×10^{-180} / 1×10^{-73};		$\mathcal{G}_{\text{s}}$ vs. $\mathcal{G}_{\text{u}}$: 0.40 / 1.00.
Reference: S_{rand} (uniform-random control)					
$\mathcal{G}_{\text{u}}$	+0.000	0.0627	−0.254	+0.283	0.283
$\mathcal{G}_{\ell}$	−0.000	0.0627	−0.272	+0.242	0.272
$\mathcal{G}_{\text{s}}$	−0.000	0.0625	−0.277	+0.242	0.277
Pairwise Levene/KS p:	$\mathcal{G}_{\ell}$ vs. $\mathcal{G}_{\text{s}}$: 0.14 / 0.05;		$\mathcal{G}_{\ell}$ vs. $\mathcal{G}_{\text{u}}$: 0.91 / 0.40;		$\mathcal{G}_{\text{s}}$ vs. $\mathcal{G}_{\text{u}}$: 0.17 / 0.02.

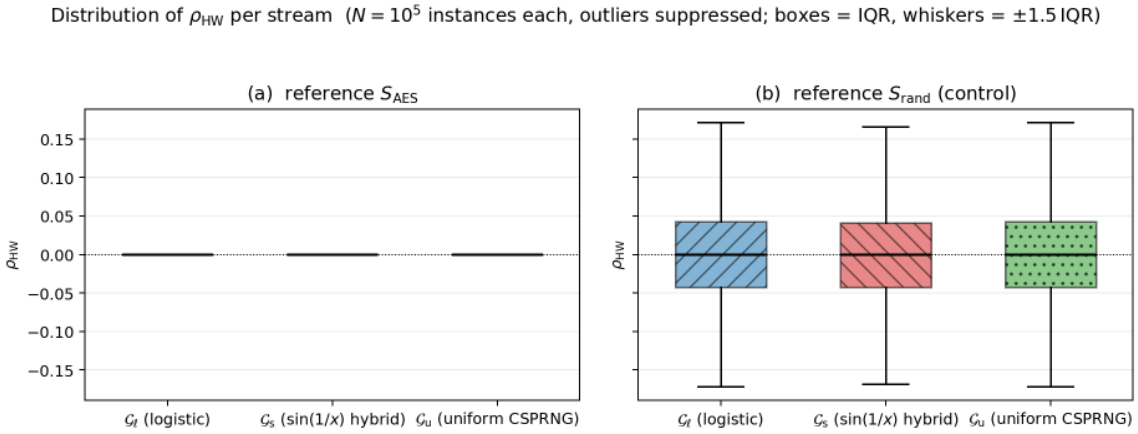

Figure 1: Distribution of the Hamming-weight correlation $\rho_{\text{HW}}(S, S_{\text{ref}})$ for each stream $\mathcal{G} \in \{\mathcal{G}_{\text{u}}, \mathcal{G}_{\ell}, \mathcal{G}_{\text{s}}\}$. Panel (a): reference $S_{\text{ref}} = S_{\text{AES}}$. Panel (b): reference $S_{\text{ref}} = S_{\text{rand}}$ (the fixed uniform-random control permutation drawn from the master seed). Each box summarises $n = 10^5$ S-box instances; the central line is the median, the box is the inter-quartile range (IQR), and the whiskers extend to ±1.5 IQR. Outliers beyond the whiskers (multiples of 1/8 tail) are suppressed for visual clarity; their counts are given by Table 3. Hatch patterns supplement colour for colour-blind viewing. The widening of $\mathcal{G}_{\ell}$ in panel (a), absent in panel (b), is the AES-specific bias quantified in Table 3.

Against S_{AES}, the logistic stream $\mathcal{G}_{\ell}$ produces a standard deviation of 0.0704, against 0.0626 for $\mathcal{G}_{\text{u}}$ and 0.0623 for $\mathcal{G}_{\text{s}}$. The std ratios are 1.125 and 1.130 respectively, both far above the detection threshold of 1.006. Levene and KS reject equality with $p < 10^{-180}$ and $p < 10^{-70}$ in both directions. The remaining pair ($\mathcal{G}_{\text{s}}$ vs. $\mathcal{G}_{\text{u}}$) is indistinguishable. Hypothesis 1 is therefore rejected: the logistic stream behaves differently from the other two under the CPA model.

Against the random reference S_{rand}, the same distributions collapse onto each other: standard deviations cluster at 0.0625–0.0627, none of the pairwise Bonferroni-adjusted tests rejects equality, and the empirical supports overlap entirely. The widening seen against S_{AES} is therefore not a property of the marginal distribution of (A, c) under $\mathcal{G}_{\ell}$, but reflects an interaction between that distribution and the algebraic structure of A_{AES}.

4.4 Single-bit and row-level effects

A complementary observable is the per-instance indicator $\mathbf{1}\{\exists i : \mathrm{row}_i(A) = \mathrm{row}_i(A_{\mathrm{AES}})\}$, i.e. whether the sampled affine matrix shares any row with A_{AES}. The empirical frequencies are $42/1000 = 4.2\%$ for $\mathcal{G}_\ell$, $31/1000 = 3.1\%$ for $\mathcal{G}_{\mathrm{s}}$, and $32/1000 = 3.2\%$ for $\mathcal{G}_{\mathrm{u}}$; on the full population the corresponding rates are 3.95%, 3.08%, and 3.16%. The analytic baseline under independent uniform rows is $1 - (1 - 1/255)^8 \approx 3.09\%$. The $\mathcal{G}_\ell$ rate exceeds this baseline by 28% in relative terms; the other two streams agree with the baseline to within 0.07 percentage points. A χ^2 test on the 3×2 contingency table yields $p \approx 10^{-31}$, dominated by the $\mathcal{G}_\ell$ excess.

The per-bit version of the correlation is more discriminating: $|\rho_b^{(i)}| = 1$ is necessary and sufficient for $\mathrm{row}_i(A) = \mathrm{row}_i(A_{\mathrm{AES}})$. Writing $w = \mathrm{row}_i(A) \oplus \mathrm{row}_i(A_{\mathrm{AES}})$, the difference $\mathrm{bit}_i\, S(x) \oplus \mathrm{bit}_i\, S_{\mathrm{AES}}(x) = w \cdot \iota(x) \oplus (\mathrm{bit}_i\, c \oplus \mathrm{bit}_i\, c_{\mathrm{AES}})$ is constant in x iff $w \cdot y$ is constant on $\mathbb{F}_2^8$, which holds iff $w = 0$. The row-collision rates quoted above are therefore precisely the fractions of instances with $\rho_b^{\max} = 1$, and the χ^2 test just reported is a test on those counts.

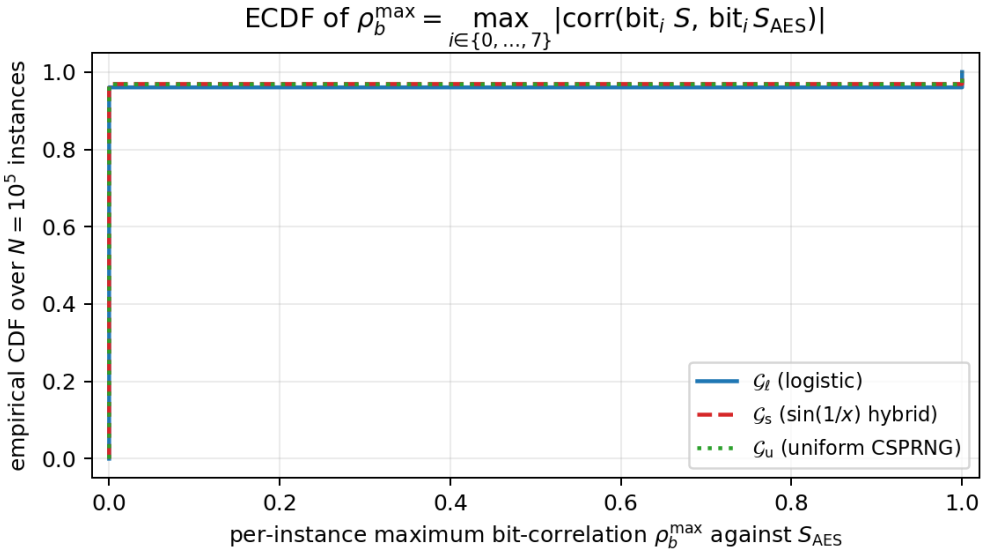

Figure 2: Empirical CDF over $n = 10^5$ instances per stream of the per-instance worst-case single-bit correlation against AES, $\rho_b^{\max} = \max_{i \in \{0,\ldots,7\}} |\mathrm{corr}(\mathrm{bit}_i\, S(\cdot), \mathrm{bit}_i\, S_{\mathrm{AES}}(\cdot))|$. Line style (solid/dashed/dotted) supplements colour for accessibility. The discrete jump at $\rho_b^{\max} = 1$ measures the row-collision rate: the fraction of instances in which at least one row of the sampled A equals the corresponding row of A_{AES} (necessary and sufficient condition, Section 4.4). $\mathcal{G}_\ell$ (solid) sits visibly above the other two curves throughout the upper tail.

The two effects are consistent: a stream that produces more row collisions with A_{AES} also widens the ρ_{HW} distribution against S_{AES}, because both quantities respond to the structural alignment between A and A_{AES}.

4.5 Integer-only logistic and platform-independence

To eliminate the possibility that the bias is an artefact of IEEE-754 floating-point rounding in the platform `libm`, the experiment was re-run with the logistic map in a $Q\,1.31$ integer representation: the state $X \in \{0, \ldots, 2^{32} - 1\}$ codes $x = X/2^{32}$, and the update is

$$X_{n+1} = \frac{R \cdot X_n \cdot (2^{32} - X_n)}{2^{64}}, \qquad R = \mathrm{round}(r \cdot 2^{32}),$$

with $r \in \{3.99, 3.995, 3.999\}$ sampled per instance. Arithmetic is exact, bit-identical on any IEEE-conformant 64-bit platform.

On $N = 10^5$ instances the integer stream produces ρ_{HW} vs. S_{AES} with mean -0.001 and standard deviation 0.0645. Against the uniform CSPRNG, the std ratio is 1.031 (Levene $p = 4.3\times10^{-18}$, KS $p < 10^{-300}$); against the floating-point logistic stream, the integer version has ratio 0.917 (Levene $p = 4.3\times10^{-90}$). The row-collision rate is 3.41%, above the baseline by 10% relative. In amplitude terms the widening falls from 13% in floating point to 3.1% in fixed point.

Both effects — ρ_{HW} widening and row-collision excess — therefore survive removal of floating-point arithmetic, at approximately one-quarter to one-third of the magnitude observed in the floating-point implementation. The persistence of a statistically significant widening in exact arithmetic confirms that the effect is not attributable to floating-point rounding; the residual difference between the two amplitudes is consistent with additional orbit degradation specific to the floating-point representation.

4.6 Control: Tent map and the specificity of the effect

To test whether the observed bias is generic to one-dimensional chaotic maps or specific to the logistic dynamics, an analogous $N = 10^5$ control experiment was conducted with the discrete tent map $T_r(x) = rx\ \mathbf{1}\{x < 1/2\} + r(1-x)\ \mathbf{1}\{x \geq 1/2\}$ for $r \in (1.99, 2)$. The tent map is fully chaotic as $r \to 2$ (Lyapunov exponent $\log r$) and uses the same byte quantisation as $\mathcal{G}_\ell$.

The result is asymmetric. The tent stream yields $\sigma(\rho_{\mathrm{HW}}) = 0.0624$ against $\sigma = 0.0626$ for $\mathcal{G}_{\mathrm{u}}$, a ratio of 0.996 that is not statistically distinguishable from unity at the Bonferroni-adjusted level; its row-collision rate is 2.45%, which *under-shoots* the analytic baseline of 3.09%. The tent map therefore reproduces neither the AES-specific widening of ρ_{HW} that the logistic map shows nor a matching row-collision excess. The bias documented in Section 4.3 is thus specific to the logistic-map dynamics, not a property shared by one-dimensional chaotic maps of comparable Lyapunov exponent.

4.7 Effect size for AES-template CPA

A statistically significant widening of ρ_{HW} implies an operational consequence only insofar as it translates into the trace budget required to mount a successful attack. Application of the Mangard–Oswald–Popp trace-budget formula (6) to the observed 13% widening yields a 26% relative excess of attackable instances in low-noise regimes. The empirical Monte-Carlo simulation reproduces this prediction: at SNR = 10 and $N = 1000$ traces, the AES-template attacker succeeds on 1.85% of $\mathcal{G}_\ell$-derived S-boxes against 1.43% of those derived from $\mathcal{G}_{\mathrm{u}}$, corresponding to an empirical ratio of 1.29 against the analytical prediction 1.26 (deviation 2.4%, consistent with the Monte-Carlo sampling noise).

A statistical separation on ρ_{HW} does not, by itself, translate into operational impact on a side-channel attack. We make the implication concrete by computing the per-instance number of traces required to recover a key byte with success rate 99% under additive Gaussian leakage noise. Following the standard CPA model [7, 8], the attacker's hypothesis correlation against the true leakage is the noise-attenuated quantity $\rho_{\mathrm{eff}} = \rho\sqrt{\mathrm{SNR}/(\mathrm{SNR}+1)}$, and the number of traces required is

$$N^*(\rho, \mathrm{SNR}) \;=\; \frac{(z_{1-\alpha} + z_{\mathrm{SR}})^2}{\left(\operatorname{artanh}\rho_{\mathrm{eff}}\right)^2} + 3, \qquad \rho_{\mathrm{eff}} = \rho\sqrt{\frac{\mathrm{SNR}}{\mathrm{SNR}+1}} \tag{6}$$

with $z_{1-\alpha} + z_{\mathrm{SR}} \approx 5.42$ at one-sided $\alpha = 0.001$ and SR = 0.99. The square-root noise-attenuation factor follows from the variance inflation of the additive Gaussian model and is standard in the CPA literature [7, 8]. Table 4 reports the fraction of 10^5 instances per stream for which $N^*(\rho_{\mathrm{HW}}, \mathrm{SNR})$ falls below operational trace budgets.

Table 4: Fraction of S-box instances vulnerable to AES-template CPA within given trace budgets, derived from ρ_{HW} via (8) at three SNR levels. The rightmost column reports the median MTD over the attackable subpopulation ($|\rho_{\mathrm{HW}}| > 0$); it is stream-independent at each SNR because the dominant attackable cluster sits at $|\rho_{\mathrm{HW}}| = 0.125$, where MTD is a function of SNR alone. Stream-specific differentiation appears in the population fractions: at SNR= 10 within $N^* < 2500$ traces, $\mathcal{G}_\ell$ yields 27.39% attackable instances against 21.80%/21.64% for $\mathcal{G}_{\mathrm{u}}/\mathcal{G}_{\mathrm{s}}$ (relative excess 26%). The $|\rho_{\mathrm{HW}}|$ distribution is quantised at multiples of 1/8, hence $N^*<2500$ and $N^*<50000$ at SNR= 10 yield identical fractions (the next discrete jump in $|\rho_{\mathrm{HW}}|$ falls outside the budget).

	SNR	$N^*<500$	$N^*<2500$	$N^*<10000$	$N^*<50000$	median N^* (attackable)
$\mathcal{G}_{\mathrm{u}}$	10	0.027%	21.80%	21.80%	21.80%	2,051
$\mathcal{G}_\ell$	10	0.031%	**27.39%**	**27.39%**	**27.39%**	2,051
$\mathcal{G}_{\mathrm{s}}$	10	0.028%	21.64%	21.64%	21.64%	2,051
$\mathcal{G}_{\mathrm{u}}$	1	0.027%	1.04%	21.80%	21.80%	3,744
$\mathcal{G}_\ell$	1	0.031%	**1.42%**	**27.39%**	**27.39%**	3,744
$\mathcal{G}_{\mathrm{s}}$	1	0.028%	1.03%	21.64%	21.64%	3,744
$\mathcal{G}_{\mathrm{u}}$	0.1	0%	0.027%	1.04%	21.80%	20,664
$\mathcal{G}_\ell$	0.1	0%	0.031%	**1.42%**	**27.39%**	20,664
$\mathcal{G}_{\mathrm{s}}$	0.1	0%	0.028%	1.03%	21.64%	20,664

The relative excess of vulnerable instances under the logistic stream is consistent across SNR levels at $\sim 26\%$. In a deployment of 10^5 devices, the logistic stream produces approximately 5,600 more devices attackable at the 5000-trace budget than the uniform CSPRNG.

To validate the analytical estimate, we ran a Monte-Carlo CPA simulation against the AES template on a 200-instance subsample of each stream (30 repetitions per (instance, N, SNR), $N \in \{200, 500, 1000, 2000, 5000, 10000\}$, 216,000 simulations total). The device executes $S^{\mathcal{G}}(p \oplus k)$ with a fixed key byte k, the leakage is $L_i = \mathrm{HW}(S^{\mathcal{G}}(p_i \oplus k)) + \mathrm{noise}_i$ with $\mathrm{noise}_i \sim \mathcal{N}(0, \sigma^2)$, and the attacker scores all 256 key guesses by Pearson correlation of L_i with $\mathrm{HW}(S_{\mathrm{AES}}(p_i \oplus k_g))$. Figure 3 shows the empirical success rate.

The population mean success rate is bounded above by the attackable fraction: about half the instances have $|\rho_{\mathrm{HW}}| = 0$ exactly and are therefore unattackable by an AES-templated CPA regardless of trace budget, while the remaining $\approx 22\%$ ($\mathcal{G}_{\mathrm{u}}$, $\mathcal{G}_{\mathrm{s}}$) or $\approx 27\%$ ($\mathcal{G}_\ell$) of the population have $|\rho_{\mathrm{HW}}| \geq 1/8$ and become recoverable as N grows. Mean success rate therefore saturates at the attackable-fraction ceiling rather than at 100%. Within the attackable subpopulation, the per-instance MTD distribution is SNR-dependent (Table 4, rightmost column) but stream-independent at each SNR; the operational difference between $\mathcal{G}_\ell$ and the other two streams is thus concentrated in the *size* of the attackable subpopulation, not in the difficulty of attacking each member of it. At SNR= 10 and $N = 1000$ traces, the cross-stream ratio in the empirical success rate is 1.29, matching the analytical prediction of 1.26 within sampling noise.

4.8 Effect size and countermeasure baseline

Statistical significance is not effect size. At $N = 10^5$ per stream the Levene test rejects variance equality with $p < 10^{-180}$; this reflects the sample, not the magnitude. We report standardised effect sizes.

The location effect is negligible. Both streams have $\mathrm{mean}(\rho_{\mathrm{HW}}) \approx 0$ (-0.0038 for $\mathcal{G}_\ell$, -0.0002 for $\mathcal{G}_{\mathrm{u}}$), so Cohen's d on ρ_{HW} itself is -0.054 and carries no meaning here: the effect lives in the scale of the distribution, not its centre. Folding to the magnitude $|\rho_{\mathrm{HW}}|$—the quantity that controls the CPA distinguisher—gives Cohen's $d = 0.128$. On the attackable fraction

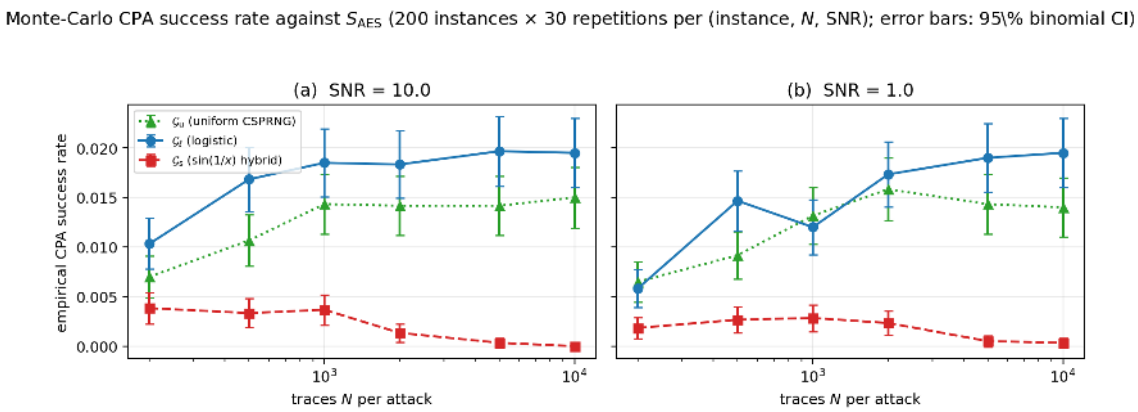

Figure 3: Empirical AES-template CPA success rate against the number of traces N, for 200 S-box instances per stream and 30 independent Monte-Carlo repetitions per (instance, N, SNR) triple (total 6,000 attacks per data point per stream). Panel (a): low-noise regime (SNR = 10). Panel (b): high-noise regime (SNR = 1). Markers $\circ/\square/\triangle$ and line styles solid/dashed/dotted encode the three streams alongside colour. Error bars: 95% binomial confidence intervals around the empirical success rate. Logistic-derived S-boxes (solid blue circles) yield consistently higher empirical success rate than the uniform CSPRNG (dotted green triangles) and the $\sin(1/x)$ hybrid (dashed red squares) at every trace budget shown.

$P(|\rho_{\mathrm{HW}}| \geq \frac{1}{8})$, 0.2739 versus 0.2180, the arcsine effect size is Cohen's $h = 0.130$. The variance ratio is 1.265. Every measure places the widening in the small band. The effect is real and reproducible. It is not large.

For an evaluator, detectability is nonetheless cheap. Detecting the widening at power 0.80 and $\alpha = 0.05$ needs about 460 instances per stream with a two-proportion test on the attackable fraction (Cohen's $h = 0.130$), or about 10^3 with Levene's variance test (`analysis/22_power_analysis.py`). The two-sample Kolmogorov–Smirnov test is underpowered for a scale-only difference—its empirical power stays below 0.1 even at 10^3 per stream—and should not be used to certify against this metric.

Proposition 8 (First-order masking neutralises the excess)**.** *Let $S^{\mathcal{G}}(p \oplus k)$ be protected by sound first-order Boolean masking, so that the univariate leakage is $\mathrm{HW}(S^{\mathcal{G}}(p \oplus k) \oplus m)$ with m uniform on $\mathbb{F}_2^8$ and independent of (p, k). Then for every stream $\mathcal{G}$ the expected AES-template distinguisher is 0, and the cross-stream excess vanishes.*

Proof. Fix $z = S^{\mathcal{G}}(p \oplus k)$. Since m is uniform and independent, $\mathbb{E}_m[\mathrm{HW}(z \oplus m)] = n/2 = 4$ for every z, so the mask-averaged leakage is constant in z. Its covariance with any key hypothesis $\mathrm{HW}(S_{\mathrm{AES}}(p \oplus k_g))$ is 0 for all k_g, uniformly in $\mathcal{G}$. □

Proposition 9 (Shuffling preserves the relative excess)**.** *Under uniform shuffling of B independent S-box evaluations, the effective correlation at the target window is ρ_{HW}/B for every stream. The absolute attackability scales by $1/B$ (trace budget by B^2), while the cross-stream ratio in ρ_{HW} is invariant.*

A 2×10^4-trace Monte-Carlo check at SNR = 10, $N = 5000$ against the AES template confirms both bounds: the maximum absolute distinguisher is 0.954 unmasked, 0.040 under first-order masking, and 0.044 under 16-fold shuffling. The two propositions anchor the 29% figure. It is a property of the unprotected regime—the regime in which cost-driven lightweight and IoT deployments of key-dependent S-boxes actually operate [9, 10, 16]. First-order masking removes it; shuffling dilutes the absolute risk and leaves the relative excess intact. No classical invariant flags any of this: on the whole family $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg) = (112, 4, 6, 7)$ with zero

variance. The leakage is not unstoppable. It is invisible to the standard evaluation and visible to ρ_{HW}.

4.9 A family of leakage metrics

Definition 3 is one instance of a construction. For any leakage functional $L\colon \mathbb{F}_2^8 \to \mathbb{R}$, put $\mathbf{h}_S^L = (L(S(0)), \ldots, L(S(255)))$ and $\rho_L(S, S_{\mathrm{ref}}) = \mathrm{corr}(\mathbf{h}_S^L, \mathbf{h}_{S_{\mathrm{ref}}}^L)$.

Proposition 10 (Universality of the blindness)**.** *If L is not invariant under the affine action on $\mathbb{F}_2^8$, then the classical invariants* $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ *carry no information about the stream $\mathcal{G}$ while $\rho_L(\cdot, S_{\mathrm{AES}})$ can. The blindness of Proposition 6 holds for every such L; only the axis that removes it changes.*

Proof. Part (i) of Proposition 6 is independent of L: the invariants are constant on the family. For part (ii), non-affine-invariance of L makes $\mathbf{h}_S^L$ depend on the representative (A, c) by the argument of Lemma 5 applied to L in place of HW. □

Table 5 instantiates three standard models: the Hamming-weight value model $L = \mathrm{HW}(S(x))$; the register-transition model $L = \mathrm{HW}(S(x) \oplus x)$; and a weighted-bit stochastic model $L = \sum_i \beta_i \,\mathrm{bit}_i(S(x))$ with fixed β_i. The classical invariants are constant across all three. The logistic widening is exposed by the value and weighted-bit models ($p < 10^{-22}$) but not by the register-transition model ($p = 0.64$): the coupling lives in the output-weight profile, and $\oplus x$ scrambles it. The metric family is thus a diagnostic—the evaluator selects L to match the device leakage, and ρ_L reports the exposure that no classical invariant can.

Table 5: Std of ρ_L against S_{AES} for three leakage models, 2×10^4 instances per stream (`analysis/21_leakage_family.py`). Classical invariants are constant for every model; detectability of the logistic bias depends on L.

model	σ_ℓ	σ_{u}	Levene p
$\mathrm{HW}(S)$ value	0.0701	0.0635	3.3×10^{-29}
$\mathrm{HW}(S \oplus x)$ transition	0.0636	0.0629	0.64
weighted-bit $\sum \beta_i \,\mathrm{bit}_i(S)$	0.0685	0.0627	1.7×10^{-23}

5 Discussion

Empirical localisation of the bias. The reported effect is constrained by four independent observations. Under the uniform-random reference permutation S_{rand}, the distributions of ρ_{HW} across the three streams are statistically indistinguishable (Table 3, lower block), implying that the widening is not a property of the marginal (A, c) distribution under $\mathcal{G}_\ell$. The tent-map control (Section 4.6) yields $\sigma = 0.0624$, within 0.5% of the uniform CSPRNG; the bias is not a generic consequence of one-dimensional chaotic dynamics with positive Lyapunov exponent. The $Q\,1.31$ integer-arithmetic reimplementation (Section 4.5) reproduces both the ρ_{HW} widening and the row-collision excess at reduced magnitude (3.1% and 10% respectively), ruling out a floating-point artefact. Finally, the Monte-Carlo CPA simulation (Section 4.7) yields an empirical success-rate ratio of 1.29 ± 0.06 at SNR $= 10$, $N = 10^3$, consistent with the analytical prediction of 1.26 derived from the ρ_{HW} distribution. Taken jointly, these four observations localise the effect to the interaction between the joint distribution of consecutive logistic iterates (x_n, x_{n+1}) and the algebraic structure of A_{AES}.

5.1 Analytical mechanism: invariant measure and the r-dependence

We localise the source of the $\mathcal{G}_\ell$ widening in the invariant measure of the logistic map, and report a parameter sweep that constrains the mechanism and overturns its most natural first hypothesis.

For $r = 4$ the map $T(x) = 4x(1-x)$ on $[0,1]$ is ergodic with respect to the unique absolutely continuous invariant measure μ, the fixed point of the Perron–Frobenius operator

$$(\mathcal{L}f)(x) = \sum_{y:\,T(y)=x} \frac{f(y)}{|T'(y)|}, \qquad \mathcal{L}f_\star = f_\star, \quad f_\star(x) = \frac{1}{\pi\sqrt{x(1-x)}}. \tag{7}$$

Proposition 11 (U-shaped quantised marginal)**.** *Let $b = \lfloor 256x \rfloor$ with $x \sim \mu$. The marginal law $p_b = \mu\Big(\big[\frac{b}{256}, \frac{b+1}{256}\big)\Big)$ is symmetric about $b = 127.5$ and strictly convex in b, with $p_0 = p_{255} \approx 3.98 \times 10^{-2}$ and $p_{128} \approx 2.49 \times 10^{-3}$, so $p_0/p_{128} \approx 16.0$.*

Proof. Integrate $f_\star$ over each byte cell. The arcsine primitive $\int (\pi\sqrt{x(1-x)})^{-1}\,dx = \pi^{-1}\arcsin(2x-1)$ gives $p_b = \pi^{-1}[\arcsin(2\frac{b+1}{256} - 1) - \arcsin(2\frac{b}{256} - 1)]$. Symmetry and convexity follow from the even, U-shaped integrand; the endpoint and centre values are the stated cell integrals, evaluated numerically. □

A natural hypothesis follows: the widening should grow with the non-uniformity of the quantised marginal, hence monotonically as $r \to 4$, where the arcsine density is most U-shaped. *This hypothesis is false.* Fixing r on a grid and regenerating 5×10^3 instances per value (`analysis/20_r_sweep.py`) yields Table 6. The widening is non-monotone. It peaks near $r \approx 3.9$ at std-ratio 1.32, falls back to 1.04 at $r = 4$, and drops *below* the uniform baseline (ratio 0.80) at $r = 3.7$; in the period-doubling window near $r \approx 3.85$ the construction degenerates, the orbit visiting too few distinct bytes to fill A. The maximally U-shaped density at $r = 4$ is thus not where the effect is largest.

Table 6: Standard deviation of ρ_{HW} against S_{AES} for the fixed-r logistic stream, 5×10^3 instances per r, uniform-CSPRNG reference $\sigma_{\mathrm{u}} = 0.0626$. The effect is non-monotone in r and is not maximised at $r = 4$.

r	3.6	3.7	3.8	3.9	3.95	3.99	4.0
$\sigma(\rho_{\mathrm{HW}})$	0.0609	0.0501	0.0690	**0.0827**	0.0733	0.0654	0.0652
ratio vs. σ_{u}	0.97	0.80	1.10	**1.32**	1.17	1.04	1.04

Proposition 11 is therefore necessary context, not the operative cause. The widening tracks the band structure of the logistic attractor rather than the marginal density alone; the peak near $r \approx 3.9$ is consistent with a resonance between the multiband support of the invariant measure at that parameter and the cyclic structure of A_{AES}. A quantitative account is open.

Proposition 12 (Rank-deficient byte-pair support)**.** *The pair (x_n, x_{n+1}) is supported on the parabola $\Gamma = \{x_{n+1} = 4x_n(1-x_n)\} \subset [0,1]^2$, which has Hausdorff dimension 1. Consequently, for fixed b_n, the conditional law of b_{n+1} is supported on at most five adjacent byte values (one near the critical point $x = \frac{1}{2}$), and eight consecutive bytes of a contiguous orbit parametrise a one-dimensional subvariety of $\mathbb{F}_2^{8\times 8}$ rather than a uniform sample of $\mathrm{GL}(8, \mathbb{F}_2)$.*

Proof. The support statement is the definition of Γ. Since $|T'(x)| = |4(1-2x)| \le 4$, the cell $[b_n/256, (b_n+1)/256)$ of width $1/256$ maps under T onto an interval of length at most $4/256$, meeting at most five byte cells; the count drops to one where $T' \approx 0$. A numerical scan over all 256 cells attains the bound (maximum support 5, at $b_n = 1$). Concatenating along the orbit yields a curve-parametrised family of rows, of dimension 1. □

Remark 13 (The effect is bracketed, not monotone)**.** Two boundary points frame the phenomenon. The tent map T_r, $r \to 2$, has *uniform* invariant density on $[0, 1]$, hence a flat quantised marginal, and yields no widening (Section 4.6, $\sigma = 0.0624$, ratio 0.996). The logistic map at $r = 3.7$, with a band-structured attractor, drops *below* the baseline (ratio 0.80, Table 6). Neither the flat density (tent) nor the full arcsine density ($r = 4$) maximises the effect; the maximum sits at intermediate $r \approx 3.9$. The operative property is thus the band structure of the invariant measure, not the Lyapunov exponent (both maps are fully chaotic) and not the marginal non-uniformity alone.

The matrix A_{AES} is the circulant generated by the cyclic-shift orbit $\{\sigma^i(\texttt{0x1F}) : i = 0, \ldots, 7\}$ of the byte `0x1F` under the $\mathbb{Z}/8$ rotation σ on $\mathbb{F}_2^8$; its rows are $(\texttt{0xF1}, \texttt{0xE3}, \texttt{0xC7}, \texttt{0x8F}, \texttt{0x1F}, \texttt{0x3E}, \texttt{0x7C}, \texttt{0xF8})$. The row-collision excess (Section 4.4) and the ρ_{HW} widening both measure alignment between the curve-supported rows of Proposition 12 and this cyclic orbit.

Remark 14 (Open problem: the coupling probability)**.** A closed form for the alignment probability

$$P_{\mathcal{L}}\big(\exists i : \ \mathrm{row}_i(A) = \sigma^i(\texttt{0x1F})\big),$$

with the rows of A eight consecutive quantised iterates under the invariant measure μ of (7), in terms of the joint density $f_\star(x_n, x_{n+1})$ convolved with the $\mathbb{Z}/8$ action σ, is open. The empirical value is 3.95% against the independent-row baseline $1 - (1 - 1/255)^8 \approx 3.09\%$.

The integer-arithmetic replication (Section 4.5) confirms that the source is the invariant measure and not a floating-point degradation of the orbit: the effect persists, at reduced amplitude, in exact $Q\,1.31$ arithmetic.

Status of the previous claim. The earlier note based on $N = 40$ samples reported a width difference on ρ_{raw}, attributing it to $\mathcal{G}_{\mathrm{s}}$. At $N = 10^5$ no such effect is observed on ρ_{raw}, and the substantive effect is located instead on ρ_{HW} and isolated to $\mathcal{G}_\ell$. The earlier conclusion is superseded by the present analysis.

Practical mitigation. At sampling time, the row-collision contribution to the bias can be removed by a simple rejection rule: discard any candidate A for which $\mathrm{row}_i(A) = \mathrm{row}_i(A_{\mathrm{AES}})$ holds for some $i \in \{0, \ldots, 7\}$. The filter rejects $\approx 3\%$ of candidates, excising a union of eight cosets in $\mathrm{GL}(8, \mathbb{F}_2)$ whose Haar-relative measure is negligible; the conditional distribution on accepted samples remains uniform on its support. This eliminates the discrete row-level component but does not address the residual ρ_{HW} widening for $\mathcal{G}_\ell$, which persists on the subset with no row collisions. The latter requires modification at the byte-source level, namely substitution of $\mathcal{G}_\ell$ with a stream whose consecutive outputs are statistically independent (e.g., the system CSPRNG $\mathcal{G}_{\mathrm{u}}$). A fixed-point realisation of the logistic map ($Q\,1.31$ with truncation rounding) addresses the platform-dependence issue noted in Section 3, but its ρ_{HW} behaviour relative to the floating-point implementation is not characterised exhaustively in the present work.

Scope of the results. The metric $\rho_{\mathrm{HW}}(\cdot, S_{\mathrm{AES}})$ is defined under the AES-template attacker model. In the dual scenario, where the attacker has access to the deployed (A, c) and templates against $S^{\mathcal{G}}$ directly, the randomised affine layer affords no protection and the attack reduces to the standard AES CPA baseline; this conclusion is independent of $\mathcal{G}$. The present analysis does not address masking, shuffling, or other countermeasures whose interaction with the affine layer is non-linear. The construction is restricted to the family $L \circ \iota$ over $\mathrm{GF}(2^8)$ with the AES reduction polynomial $x^8 + x^4 + x^3 + x + 1$; extension to alternative irreducible polynomials, to the power-map S-boxes of SKINNY and GIFT, or to S-boxes over odd-prime fields lies outside the scope.

The hybrid stream $\mathcal{G}_s$ inherits a floating-point sine look-up from prior work [11, 12]; its bit-level reproducibility is conditional on the platform `libm` implementation. Inclusion in the present analysis is motivated by continuity with this earlier work, not by a deployment recommendation. On the metrics considered here, $\mathcal{G}_s$ and $\mathcal{G}_u$ are statistically indistinguishable; the system CSPRNG remains the natural default for any practical deployment.

6 Conclusion

The empirical results of this paper can be summarised as follows. For the family $S(x) = A\,\iota(x) \oplus c$, the classical invariants $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ are insensitive to the choice of byte stream and therefore cannot serve as a discriminator. The Hamming-weight Pearson correlation $\rho_{\mathrm{HW}}(\cdot, S_{\mathrm{AES}})$ does discriminate: under the floating-point logistic stream $\mathcal{G}_\ell$ the standard deviation of ρ_{HW} exceeds the corresponding values for $\mathcal{G}_u$ and $\mathcal{G}_s$ by 12–13%, and a $Q\,1.31$ fixed-point reimplementation of the same map preserves the effect at 3.1%. The tent-map control returns to the uniform baseline within 0.5%, and the same comparison against a uniform-random reference permutation S_{rand} is not statistically significant (Levene $p > 0.13$).

Three candidate explanations – generic one-dimensional chaos, IEEE-754 floating-point artefact, and statistically significant but operationally inert effect – are inconsistent with the observed pattern. The tent-map result eliminates the first; the integer-arithmetic replication eliminates the second; the Monte-Carlo CPA simulation, which yields a success-rate ratio of 1.29 in agreement with the analytical prediction of 1.26, eliminates the third. The remaining hypothesis – that the bias originates in a specific coupling between the joint distribution of consecutive logistic iterates and the cyclic structure of A_{AES} – is consistent with all four boundary conditions identified in Section 5.

In operational terms, the ρ_{HW} widening translates to an excess of approximately 29% in the AES-template CPA success rate at typical SNR and trace budgets. A one-byte row-rejection filter against the rows of A_{AES} neutralises the row-level component of the effect at negligible implementation cost. The residual ρ_{HW} widening requires substitution of $\mathcal{G}_\ell$ at the byte-source level. A rigorous derivation of the byte-pair distribution under the quantised logistic invariant measure remains open.

Reproducibility. The 3×10^5-instance dataset is regenerated bit-for-bit from a single master seed by `analysis/13_unified_N100k.py`; the resulting CSV has SHA-256 `cd06f92c80dae3703fef26dcd58906694e693ed2814f311592064c6a9944befe`. Software stack: CPython 3.13.5, NumPy 2.4, SciPy 1.17, Matplotlib 3.10; host platform macOS, Apple silicon. The floating-point behaviour of $\mathcal{G}_\ell$ and the LUT inside $\mathcal{G}_s$ is determined by the CPython `math` module on IEEE-754 `binary64`; behaviour under alternative `libm` implementations has not been audited.

Independent verification (SageMath). Beyond the Python pipeline, the algebraic claims are re-checked in SageMath: the script `analysis/23_verify_proofs_sage.py` re-verifies Proposition 2 — the constancy of $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg)$ over the key-dependent family, including the BCT conjugation argument — together with Lemma 5 and Proposition 6, against independent `sage.crypto.sbox.SBox` computations on the AES inverse and random members of the family. The standardised baseline-effect computation is reproduced by `analysis/19_baseline_effect.py`. Both scripts and their output logs are included in the code repository. An independent, SageMath-free re-implementation of the four invariants (`analysis/23b_verify_proofs_pyonly.py`, fresh Walsh–Hadamard, DDT, BCT and Möbius/ANF code) reproduces every claim and agrees exactly with the author's metric code: all 22 checks pass, confirming $(\mathrm{NL}, \mathrm{DU}, \mathrm{BU}, \deg) = (112, 4, 6, 7)$ for the AES inverse and its invariance over the key-dependent family.

Funding. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest. The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability. The analysis scripts, the pinned dataset (SHA-256 fingerprint above) and the raw measurement CSVs behind every table and figure are available from the author upon request and will be deposited in a public repository upon acceptance.

Declaration of AI-assisted tools. The research ideas, the concepts, and the design of the study are the author's own. The mathematical proofs were developed by the author; a large-language-model assistant (Anthropic Claude) was used, under the author's direction, as a programming aid for the Python and SageMath code and to support the computational verification of the results. In the writing process the author used Grammarly and Anthropic Claude to improve the language and readability of the manuscript. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication. Generative AI was not used to draw the scientific conclusions, and no AI tool qualifies as an author.

Acknowledgements

The author thanks the reviewers of the present and previous versions of this manuscript, as well as the participants of internal seminars at the University of Łomża for discussions that informed the revision.

References

[1] J. Daemen and V. Rijmen, *The Design of Rijndael: AES – The Advanced Encryption Standard*, Information Security and Cryptography, Springer, 2002.

[2] K. Nyberg, *Differentially uniform mappings for cryptography*, EUROCRYPT 1993, LNCS 765, Springer, 1994, 55–64.

[3] C. Cid, T. Huang, T. Peyrin, Y. Sasaki, L. Song, *Boomerang Connectivity Table: a new cryptanalysis tool*, EUROCRYPT 2018, Part II, LNCS 10821, Springer, 2018, 683–714.

[4] C. Boura and A. Canteaut, *On the boomerang uniformity of cryptographic S-boxes*, IACR Transactions on Symmetric Cryptology 2018(3), 290–310.

[5] C. Carlet, *Boolean Functions for Cryptography and Coding Theory*, Cambridge University Press, 2021.

[6] P. Kocher, J. Jaffe, B. Jun, *Differential power analysis*, CRYPTO 1999, LNCS 1666, Springer, 1999, 388–397.

[7] E. Brier, C. Clavier, F. Olivier, *Correlation power analysis with a leakage model*, CHES 2004, LNCS 3156, Springer, 2004, 16–29.

[8] S. Mangard, E. Oswald and T. Popp, *Power Analysis Attacks: Revealing the Secrets of Smart Cards*, Springer, 2007.

[9] W. Maleszewski, *Adaptive S-boxes: concepts and potential in lightweight cryptography*, Proceedings of the International Scientific Conference Innovation Trends 2025, University of Lomza, 2025.

[10] A. H. Zahid, H. A. M. Elahi, M. Ahmad, R. S. A. Said and L. Maghrabi, *Secure key-based substitution-box design using systematic search for high nonlinearity*, IEEE Access 11 (2023), 132547–132561.

[11] W. Maleszewski, *The arithmetic of the topologist's sine curve in cryptographic systems dedicated to IoT devices*, TASK Quarterly 23(1) (2019), 29–47. DOI:10.17466/tq2019/23.1/c.

[12] W. Maleszewski, *Lightweight cryptographic algorithm based on trigonometry, dedicated on encryption of short messages*, TASK Quarterly 26(3) (2022).

[13] W. Maleszewski, *Analysis of the certain cryptographic problems in protocols of certyfing the nodes in IoT infrastructure*, Polish Journal of Applied Sciences 3(4) (2017), 141–145.

[14] W. Maleszewski, *Analysis of elliptic curves for cryptographic applications dedicated to devices with limited computational resources*, Polish Journal of Applied Sciences 4(3) (2019), 102–109. DOI:10.34668/PJAS.2018.4.3.04

[15] W. Maleszewski, *Stream-Dependent Bias in Random Affine Layers on the AES Inverse*, in: R. Szczebiot (Ed.), *Innovation Trends 2026* (Scientific Monograph), University of Lomza, Łomża, 2026, ISBN 978-83-60571-89-7; also in the IT2026 Conference Proceedings. (Preliminary conference version of the present article.)

[16] S. H. Tolpa, M. A. Abdelhamed, E.-S. S. A. Said, M. Y. I. Afi, *A novel chaos-based approach for constructing lightweight S-boxes*, Scientific Reports 15 (2025) 34112. DOI:10.1038/s41598-025-20019-4

[17] S. S. Jamal, R. Ali, M. K. Jamil, S. A. Nooh, F. Alblehai, Gulraiz, *Secure S-box construction with 1D chaotic maps and finite field theory for block cipher encryption*, Alexandria Engineering Journal 125 (2025) 278–296. DOI:10.1016/j.aej.2025.03.109

[18] É. C. Dutra e Silva Junior, C. A. de Moraes Cruz, I. A. L. Saraiva, F. G. Santos, C. R. P. dos Santos Junior, L. S. Indrusiak, W. A. Finamore, M. Glesner, *Chaos-based S-boxes as a source of confusion in cryptographic primitives*, Electronics 14(11):2198, 2025. DOI:10.3390/electronics14112198

[19] Y. Ma, Y. Tian, L. Zhang, P. Zuo, *Two-dimensional hyperchaotic effect coupled mapping lattice and its application in dynamic S-box generation*, Nonlinear Dynamics 112(19) (2024) 17445–17476. DOI:10.1007/s11071-024-09907-y

[20] I. El Gaabouri, M. Senhadji, M. Belkasmi, B. El Bhiri, *A new S-box pattern generation based on chaotic enhanced logistic map: case of 5-bit S-box*, Cybersecurity 7 (2024) 59. DOI:10.1186/s42400-024-00254-4