RESEARCH ARTICLE

Five-Bit Axial Subspace Differential Uniformity: Exact Optima and a DDT-Support Realizability Gap

Jingchuan Ma[1], Yuqing Shao[1], Xin Wei[1], Qiaoyun Huang[1*]

[1*]Department of Computer Engineering, Fuzhou University Zhicheng College, No. 50 Yangqiao West Road, Gulou District, Fuzhou, 350002, Fujian, China.

ORCID: Jingchuan Ma: 0009-0001-3703-471X

*Corresponding author(s). E-mail(s): huangqy@fdzcxy.edu.cn;

Abstract

Subspace differential uniformity (SDU) measures the concentration of a differential distribution table (DDT) on affine subspaces. We determine the exact optima of the two axial affine-SDU coordinates in dimension five and establish a strict gap between locally admissible support designs and supports realizable by almost perfect nonlinear (APN) permutations. We first prove that every **16**-subset of $\mathbb{F}_2^5$ meets some affine **3**-flat in at least six points. Equality holds precisely for balanced quadratic indicators of polar rank four, forming a single affine orbit represented by $\mathbf{supp}(\mathbf{Tr}(x^3))$. This gives the relaxed axial optimum **12**. We then classify all **900** ordered monomial-trace candidates $M_{r,s}(u,v) = \mathbf{Tr}(u^r v^s)$: exactly **100** attain both relaxed axial optima, yielding **20** labelled arrays and three product-linear types up to transpose. Despite satisfying regularity and zero-vector-sum constraints, every optimal type violates a global necessary condition for DDT realizability: its two-dimensional character transform contains a negative coefficient where a vectorial Walsh square is required. Finally, the published exhaustive affine classification of five-bit APN permutations, together with direct recomputation of all five class representatives, gives realizable axial minimum **14**, attained simultaneously by the x^{15} class. Thus local incidence and vector-sum constraints permit score **12**, whereas genuine APN-permutation DDT supports require score **14**.

Keywords: Boolean functions, APN permutations, subspace differential uniformity, differential distribution tables, finite geometry, trace constructions

MSC 2020: 94A60 , 94D10 , 05B25 , 11T71

1 Introduction

Differential uniformity measures how frequently an input difference can produce a prescribed output difference. For a mapping $F : \mathbb{F}_2^n \to \mathbb{F}_2^n$, this information is recorded by the differential distribution table (DDT), and the almost perfect nonlinear (APN) condition requires every nonzero input-difference row to contain only the values zero and two [1, 2]. The usual differential uniformity is a pointwise maximum. Subspace differential uniformity (SDU), introduced by Rønjom, Sandrib, and Sunde, instead aggregates DDT entries over affine or linear subspaces on the two difference axes [3]. It therefore detects geometric concentration that is invisible to the largest individual DDT entry.

We study the two axial affine-SDU coordinates $(0, 3)$ and $(3, 0)$ in dimension five. Under the asymmetric zero convention of Rønjom et al. [3], the coordinate $(0, 3)$ is twice the largest intersection between a nonzero DDT-support row and an affine 3-flat of $\mathbb{F}_2^5$; the transposed coordinate $(3, 0)$ has the analogous column interpretation. The relaxed problem is consequently a finite-geometric optimization over 16-subsets of $\mathbb{F}_2^5$. Rønjom, Sandrib, and Sunde also proposed the scored bivariate family $2\,\mathrm{Tr}(u^r v^s)$ as candidate DDT indicators. We write

$$M_{r,s}(u, v) = \mathrm{Tr}(u^r v^s)$$

for the Boolean support and restore the factor two through the scoring convention in (2.11).

Dimension five combines several features that make an exact analysis possible without reducing the paper to an undifferentiated exhaustive search. A nonzero APN-permutation DDT row has a 16-point support, while the axial coordinates test intersections with 8-point affine 3-flats. Reed–Muller duality reduces an extremal halfset indicator to degree at most three; exterior duality leaves only two nonzero cubic top-part types; and a quadratic form in five variables has maximal polar rank four. In the multiplicative group of $\mathbb{F}_{32}$, the 30 unit exponents split into six binary cyclotomic cosets of size five. Finally, Theorem 3 of Brinkmann and Leander supplies a single external completeness input: five-bit APN permutations form exactly five ordinary affine-equivalence classes [4]. These geometric, Boolean-algebraic, and cryptographic structures permit a proof-led classification with bounded exact computations at explicitly identified steps.

Three nested support models separate local combinatorial constraints from global DDT realizability. Level R retains only the 16-regularity forced on every nonzero row and column of an APN-permutation DDT. Level S_1 additionally requires zero vector XOR in every nonzero row and column. Level P consists of supports of actual five-bit APN permutations. Thus

$$\mathrm{P} \subseteq \mathrm{S}_1 \subseteq \mathrm{R}.$$

The first two models have the same axial minima, and every simultaneous Level R optimizer automatically lies in Level S_1; the strict separation occurs between these locally admissible designs and Level P realizability.

Our first result classifies the equality case for arbitrary halfsets, not merely for sets already known to be quadratic. Every 16-subset of $\mathbb{F}_2^5$ meets an affine 3-flat in at least

six points. If the cap is six, then the vector XOR of the set is zero, its indicator lies in $\mathrm{RM}(3,5)$, and a fourth-incidence congruence excludes both nonzero cubic top-part types. Equality therefore forces a balanced quadratic of polar rank four. The classical affine classification of quadratic Boolean functions by polar rank and radical derivative then gives one $\mathrm{AGL}(5,2)$ orbit, represented by

$$q_\star(x) = x_1x_2 + x_3x_4 + x_5$$

and, after a field identification, by $\mathrm{Tr}(x^3)$. Every equality set has affine-3-flat spectrum

$$\{2 : 120,\ 4 : 380,\ 6 : 120\},$$

and the orbit contains 27,776 labelled supports. Reed–Muller duality and quadratic rank theory are the classical inputs [2, 5–7]; the passage from an arbitrary cap-six halfset to quadraticity is proved here.

Our second result classifies the complete five-bit monomial-trace family

$$M_{r,s}(u,v) = \mathrm{Tr}(u^r v^s), \qquad r,s \in \mathbb{Z}_{31}^*,$$

with zero axes adjoined. Among the 900 ordered exponent pairs, all 900 yield Level R matrices, exactly 625 yield Level S_1 matrices, and exactly 100 attain simultaneous axial score $(12,12)$. The optimal pairs represent 20 distinct labelled Boolean arrays and three product-linear equivalence classes after adjoining transpose, represented by $M_{3,3}$, $M_{3,5}$, and $M_{5,5}$. The proof first reduces the exponent domain by cyclotomic and algebraic-degree arguments and then evaluates a finite set of exact invariants; exponent pairs, literal matrices, and geometric equivalence classes are kept distinct throughout.

Our third result identifies the obstruction underlying the relaxed-versus-realizable gap. At score 12, every row and column is forced into the same rank-four quadratic orbit, so the local incidence and vector-sum constraints are mutually consistent. They are nevertheless insufficient for a DDT: the unnormalized two-dimensional character transform of a raw DDT must be the square of the vectorial Walsh transform [2, 3, 8]. A regular APN-support pattern forces a unique candidate raw DDT, and each optimal trace type has a negative transform coefficient; for example,

$$H_{M_{3,3}}(9,9) = -16.$$

Thus none of the optimal trace designs is the derivative support of any mapping. This is a global Fourier-positivity obstruction, not a failure of a particular reconstruction algorithm.

Finally, the finite-geometric lower bound and trace construction give relaxed and vector-sum-balanced minima 12 on both axes. Theorem 3 and Table 1 of Brinkmann and Leander give exactly five ordinary affine classes of five-bit APN permutations and one representative of each class [4]. Direct recomputation of the DDT profiles of those representatives gives realizable minimum 14 on each axis, attained simultaneously by the x^{15} class. Hence the Level R and Level S_1 minima are both 12 for each axial coordinate, whereas the corresponding Level P minimum is 14. The universal Level P

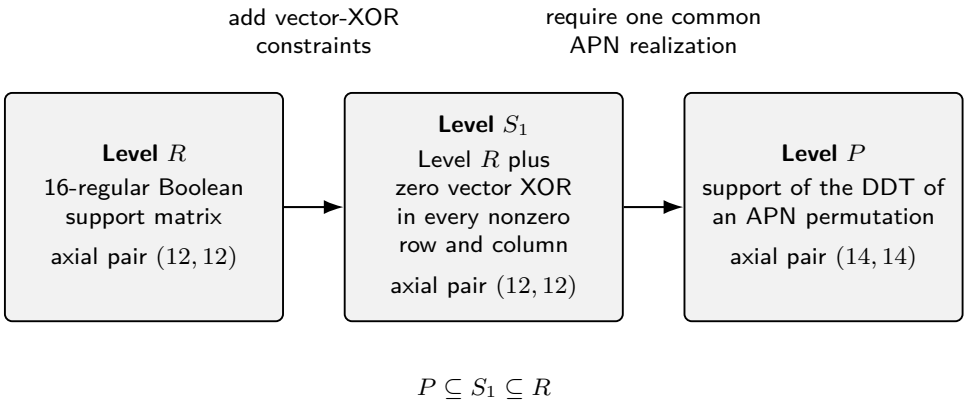

Fig. 1 Nested support models $P \subseteq S_1 \subseteq R$. Level R imposes 16-regularity, Level S_1 additionally imposes zero vector XOR in every nonzero row and column, and Level P consists of APN-permutation DDT supports. The simultaneous axial minima are $(12, 12)$, $(12, 12)$, and $(14, 14)$, respectively

statement imports the published completeness theorem; the five representative profiles are recomputed locally, and the remaining principal results do not depend on that external exhaustive quantifier.

Figure 1 summarizes the three models and the exact gap.

Section 2 fixes the SDU convention and the three support models. Sections 3 and 4 prove the incidence bound and classify equality. Section 5 classifies the monomial-trace family. Section 6 gives the Walsh-square obstruction, and Section 7 establishes the exact relaxed-versus-realizable gap. Section 8 records a bounded odd-dimensional extension, with full derivations in Appendix B. Section 9 states the computational and external trust boundaries.

2 Preliminaries and model hierarchy

2.1 Affine flats, halfsets, and caps

Let $V_n = \mathbb{F}_2^n$. We use the same notation for the underlying vector and affine space. A k-dimensional affine flat is a coset $a + U$ of a k-dimensional linear subspace $U \leq V_n$, and $\mathcal{A}_k(V_n)$ denotes the set of all such flats. For $S \subseteq V_n$, define

$$\mathrm{cap}_k(S) = \max_{A \in \mathcal{A}_k(V_n)} |S \cap A| . \tag{2.1}$$

A halfset is a subset of cardinality 2^{n-1}. Its Boolean indicator is denoted by $\mathbf{1}_S$. The vector XOR of S is

$$\mathrm{XOR}(S) = \bigoplus_{x \in S} x = \sum_{x \in S} x \in V_n.$$

2.2 Boolean, Walsh, quadratic, and Reed–Muller conventions

For a Boolean function $f : \mathbb{F}_2^k \to \mathbb{F}_2$, its unnormalized Walsh transform is

$$W_f(a) = \sum_{x \in \mathbb{F}_2^k} (-1)^{f(x)+a\cdot x}.$$

For a Boolean function $q : V_n \to \mathbb{F}_2$ of algebraic degree at most two, its polar form is

$$B_q(x, y) = q(x + y) + q(x) + q(y) + q(0). \tag{2.2}$$

The rank of q means the rank of B_q, and $\mathrm{rad}(B_q)$ denotes its radical. If $r \in \mathrm{rad}(B_q)$, then

$$\lambda_q(r) = D_r q = q(x + r) + q(x) = q(r) + q(0) \tag{2.3}$$

is independent of x and defines a linear functional on the radical. A quadratic function is balanced exactly when this radical derivative is nonzero [2, pp. 194–200]; see also Carlet [6], Charpin et al. [7].

We use the full-length, unpunctured binary evaluation code

$$\mathrm{RM}(r, 5) = \left\{ \big(f(x)\big)_{x \in \mathbb{F}_2^5} : \deg_{\mathrm{ANF}} f \le r \right\} \subseteq \mathbb{F}_2^{32}, \tag{2.4}$$

with coordinates indexed by all points of $\mathbb{F}_2^5$ and with the standard binary Euclidean inner product. Under this convention,

$$\mathrm{RM}(1, 5)^{\perp} = \mathrm{RM}(3, 5). \tag{2.5}$$

Indeed, the product of a degree-at-most-one function and a degree-at-most-three function has degree at most four, and every squarefree monomial of degree below five sums to zero over $\mathbb{F}_2^5$. Hence $\mathrm{RM}(3, 5) \subseteq \mathrm{RM}(1, 5)^{\perp}$. The dimensions are

$$\dim \mathrm{RM}(1, 5) = 1 + 5 = 6, \qquad \dim \mathrm{RM}(3, 5) = 1 + 5 + 10 + 10 = 26,$$

so equality follows. This is the only Reed–Muller convention used below [5].

2.3 DDTs, APN permutations, and affine SDU

For $F : V_n \to V_n$, the raw differential distribution table is

$$\delta_F(a, b) = \left| \{x \in V_n : F(x + a) + F(x) = b\} \right|. \tag{2.6}$$

The map F is APN if $\delta_F(a, b) \le 2$ for every $a \neq 0$ and every b. If F is a permutation, then $\delta_F(a, 0) = 0$ whenever $a \neq 0$.

For affine flats $A, B \subseteq V_n$, we follow the asymmetric zero convention of Rønjom et al. [3] and put

$$T_F(A, B) = \sum_{a \in A \setminus \{0\}} \sum_{b \in B} \delta_F(a, b). \tag{2.7}$$

The affine-SDU entry of dimensions (d_1, d_2) is

$$\Delta^{(d_1,d_2)}(F) = \max_{\substack{A\in\mathcal{A}_{d_1}(V_n)\\ B\in\mathcal{A}_{d_2}(V_n)}} T_F(A, B). \tag{2.8}$$

Linear SDU restricts A and B to linear subspaces; all results in this paper concern the affine version unless stated otherwise.

The deletion of input difference zero in (2.7) is important. In particular, the raw DDT has $\delta_F(0, 0) = 2^n$, whereas the effective matrix used by affine SDU has a zero input axis. We retain this convention throughout and reinsert the raw $(0, 0)$ entry only in the Walsh-square transform of Section 6.

2.4 Three nested support models

We now specialize to $V = V_5$ and write $V^* = V \setminus \{0\}$. A Boolean support matrix is a function

$$M : V \times V \to \{0, 1\}$$

whose zero row and zero column are identically zero. Its scored entries are $2M(a, b)$.

Definition 2.1 (Level R)**.** A support matrix is in Level R if every nonzero row and every nonzero column has Hamming weight 16.

Definition 2.2 (Level S_1)**.** A Level R matrix is in Level S_1 if

$$\bigoplus_{b:M(a,b)=1} b = 0 \quad (a \neq 0), \qquad \bigoplus_{a:M(a,b)=1} a = 0 \quad (b \neq 0). \tag{2.9}$$

Definition 2.3 (Level P)**.** A support matrix is in Level P if there is an APN permutation $F : V \to V$ such that, for all $a, b \neq 0$,

$$M(a, b) = 1 \quad \Longleftrightarrow \quad \delta_F(a, b) = 2.$$

The next proposition establishes that these are nested relaxations of the realizable domain.

Proposition 2.4. *Every five-bit APN-permutation support is in Level* S_1*, and every Level* S_1 *matrix is in Level* R*. Hence*

$$\mathrm{P} \subseteq \mathrm{S}_1 \subseteq \mathrm{R}.$$

Proof Fix an APN permutation F. For $a \neq 0$, the derivative $x \mapsto F(x + a) + F(x)$ has 32 inputs. Its solutions occur in pairs $x \leftrightarrow x + a$, and APN implies that every attained output difference has multiplicity two. Thus each nonzero row has exactly 16 support entries.

For a fixed $b \neq 0$ and each x, there is a unique $y = F^{-1}(F(x) + b)$, with $y \neq x$. Writing $a = x + y$ shows that

$$\sum_a \delta_F(a, b) = 32.$$

Again every positive entry equals two, so each nonzero column has weight 16. This proves Level R regularity.

For the row XOR, the 16 unordered pairs $\{x, x + a\}$ partition V. Their output differences are distinct, and therefore

$$\bigoplus_{b:M(a,b)=1} b = \bigoplus_{\{x,x+a\}} \big(F(x) + F(x + a)\big) = \bigoplus_{x \in V} F(x) = \bigoplus_{y \in V} y = 0.$$

The column identity follows from the same argument applied to F^{-1}, since

$$\delta_{F^{-1}}(a, b) = \delta_F(b, a). \tag{2.10}$$

Thus the support is in Level S_1. □

The row and column XOR constraints are necessary but not sufficient for realizability. In coding-theoretic language, every nonzero row and column incidence vector is a weight-16 codeword in the binary Hamming code whose parity-check columns are the 31 nonzero vectors of V. Level P additionally requires a common permutation of the 32 vertices that simultaneously realizes all derivative matchings.

2.5 Support-matrix SDU and the axial coordinates

For a support matrix M, define

$$T_M(A, B) = 2 \sum_{a \in A \setminus \{0\}} \sum_{b \in B} M(a, b) \tag{2.11}$$

and define $\Delta^{(d_1,d_2)}(M)$ by maximizing (2.11) over affine flats of dimensions d_1, d_2. If M is the support of a five-bit APN permutation, this agrees with (2.8).

For a nonzero row support $S_a = \{b : M(a, b) = 1\}$ and a nonzero column support $C_b = \{a : M(a, b) = 1\}$,

$$\Delta^{(0,3)}(M) = 2 \max_{a \neq 0} \mathrm{cap}_3(S_a), \qquad \Delta^{(3,0)}(M) = 2 \max_{b \neq 0} \mathrm{cap}_3(C_b). \tag{2.12}$$

Thus the axial optimization problem is controlled by affine-3-flat caps of 16-subsets of V.

2.6 Equivalence conventions

Ordinary affine equivalence of permutations is

$$G(x) = QF(Px + p) + q,$$

where $P, Q \in \mathrm{GL}(5, 2)$. Translation cancels in a derivative and

$$\delta_G(a, b) = \delta_F(Pa, Q^{-1}b). \tag{2.13}$$

Consequently every affine-SDU coordinate is invariant under ordinary affine equivalence. We do not assert invariance under unrestricted extended-affine equivalence, because an added nonzero linear term shears output differences in a row-dependent manner.

For abstract support matrices, product-linear equivalence means

$$M'(u, v) = M(Au, Bv), \qquad A, B \in \mathrm{GL}(5, 2). \tag{2.14}$$

Transpose is adjoined only when explicitly stated. Arbitrary permutations of rows or columns, nonlinear power maps, and exponent inversion are not admitted as geometric equivalences: they need not preserve affine flats or the labels in the XOR constraints.

3 An exact affine-incidence lower bound

The axial lower bound follows from a short incidence certificate. It applies to every halfset of $V = \mathbb{F}_2^5$ and does not assume algebraic structure. Problems of avoiding prescribed affine-flat intersection sizes have been studied more generally in finite affine spaces; see, for example, Kovács and Nagy [9]. The numerical equality used here is specific to 16-sets and affine 3-flats in V.

Lemma 3.1 (Incidence parameters)**.** *There are* 620 *affine* 3*-flats in* V*. The numbers of affine* 3*-flats containing a fixed set of one, two, or three distinct points are respectively*

$$155, \qquad 35, \qquad 7.$$

Proof The number of 3-dimensional directions is the Gaussian binomial coefficient

$$\begin{bmatrix} 5 \\ 3 \end{bmatrix}_2 = 155.$$

Each direction has $2^{5-3} = 4$ cosets, giving 620 affine 3-flats. Through a fixed point there is one coset for each direction. A flat through two distinct points must have a direction containing their nonzero difference, so the count is $\begin{bmatrix} 4 \\ 2 \end{bmatrix}_2 = 35$. Three distinct points in a binary affine space span an affine plane; a 3-direction containing its 2-direction can be chosen in $\begin{bmatrix} 3 \\ 1 \end{bmatrix}_2 = 7$ ways. □

For $S \subseteq V$ with $|S| = 16$, write

$$n_A = |S \cap A|, \qquad A \in \mathcal{A}_3(V).$$

Double counting one-, two-, and three-subsets of S in affine 3-flats gives

$$\sum_A 1 = 620, \tag{3.1}$$

$$\sum_A n_A = 16 \cdot 155 = 2480, \tag{3.2}$$

$$\sum_A \binom{n_A}{2} = \binom{16}{2} \cdot 35 = 4200, \tag{3.3}$$

$$\sum_A \binom{n_A}{3} = \binom{16}{3} \cdot 7 = 3920. \tag{3.4}$$

Theorem 3.2 (Halfset incidence bound)**.** *Every* 16*-subset* $S \subseteq \mathbb{F}_2^5$ *satisfies*

$$\mathrm{cap}_3(S) \geq 6.$$

Proof Assume that $n_A \leq 5$ for every affine 3-flat. Consider the polynomial in the binomial basis

$$p(t) = 5 - 5t + 5\binom{t}{2} - 3\binom{t}{3}. \tag{3.5}$$

Its values on the assumed range are

$$\big(p(0), p(1), p(2), p(3), p(4), p(5)\big) = (5, 0, 0, 2, 3, 0),$$

so $p(n_A) \geq 0$ for every A. On the other hand, (3.1)–(3.4) give

$$\begin{aligned} \sum_A p(n_A) &= 5(620) - 5(2480) + 5(4200) - 3(3920) \\ &= -60, \end{aligned}$$

which is impossible. □

Corollary 3.3 (Axial lower bound)**.** *Every Level* R *matrix satisfies*

$$\Delta^{(0,3)}(M) \geq 12, \qquad \Delta^{(3,0)}(M) \geq 12.$$

The same lower bounds hold at Levels S_1 *and* P*.*

Proof Every nonzero row and column of a Level R matrix is a 16-subset of V. Apply Theorem 3.2 and (2.12). The remaining statements follow from the model inclusions. □

The certificate is sharp: Section 4 classifies all halfsets attaining cap six, and Section 5 embeds such halfsets simultaneously as every nonzero row and column of explicit Level S_1 matrices.

4 Classification of cap-six halfsets

We now determine all equality cases in Theorem 3.2. The starting point is an arbitrary halfset; quadraticity is a conclusion rather than an assumption.

Theorem 4.1 (Complete cap-six classification)**.** *Let* $S \subseteq \mathbb{F}_2^5$ *with* $|S| = 16$*. The following are equivalent.*

(i) $\mathrm{cap}_3(S) \leq 6$*;*

(ii) $\mathrm{cap}_3(S) = 6$;

(iii) $\mathbf{1}_S$ *is a balanced quadratic Boolean function of polar rank four;*

(iv) S *is* $\mathrm{AGL}(5,2)$*-equivalent to the support of*

$$q_\star(x) = x_1x_2 + x_3x_4 + x_5;$$

(v) *after an identification* $\mathbb{F}_2^5 \cong \mathbb{F}_{32}$*, the set* S *is affine-equivalent to* $\mathrm{supp}(\mathrm{Tr}(x^3))$.

Every such set has affine-3-flat intersection spectrum

$$\{2 : 120,\ 4 : 380,\ 6 : 120\}. \tag{4.1}$$

The equality class contains exactly 27,776 *labelled supports and forms one* $\mathrm{AGL}(5,2)$ *orbit.*

The implication (i)⇒(ii) follows from Theorem 3.2. The remainder of this section proves the structural implication from (i) to (iii), then derives the required orbit, count, and spectrum consequences from polar rank and the radical derivative.

4.1 Cap six forces vector XOR zero

We first isolate a small affine-hyperplane lemma.

Lemma 4.2. *If* $T \subseteq \mathbb{F}_2^4$ *and* $|T| \le 5$*, then some affine hyperplane of* $\mathbb{F}_2^4$ *meets* T *in at most one point.*

Proof If T is empty there is nothing to prove. Translate one point of T to zero and let the remaining at most four vectors be $v_1, \dots, v_t$. If they span a proper subspace, choose a nonzero linear functional that vanishes on their span; its affine level set $\ell = 1$ misses all of T. If they span the whole four-space, then $t = 4$ and the v_i form a basis. A dual coordinate functional has value one on exactly one v_i, so its level set $\ell = 1$ meets T once. □

Lemma 4.3 (XOR bridge)**.** *If* $S \subseteq \mathbb{F}_2^5$*,* $|S| = 16$*, and* $\mathrm{cap}_3(S) \le 6$*, then*

$$\mathrm{XOR}(S) = 0. \tag{4.2}$$

Proof Suppose that $w = \mathrm{XOR}(S) \neq 0$. Choose a linear functional ℓ with $\ell(w) = 1$ and write

$$S_i = S \cap \{x : \ell(x) = i\}, \qquad i \in \mathbb{F}_2.$$

Because

$$|S_1| \equiv \sum_{x \in S} \ell(x) = \ell(w) = 1 \pmod 2$$

and $|S| = 16$, both slice sizes are odd.

Neither slice can have at most five points. Indeed, if one slice has size at most five, the complement of the other slice in its affine four-space also has size at most five. Lemma 4.2 gives an affine 3-flat in that four-space meeting the complement in at most one point, and

therefore meeting S in at least seven points. This contradicts the cap. The two odd slice sizes must consequently be 7 and 9.

Every 7-subset of $\mathbb{F}_2^4$ contains an affine plane. Its $\binom{7}{2} = 21$ unordered pair sums take only 15 possible nonzero values. Two distinct pairs with the same sum cannot share exactly one endpoint, so the repeated sum gives four distinct points with XOR zero, hence an affine plane P.

Let W be the two-dimensional direction of such a plane in the 7-slice. The other affine four-space is partitioned into four W-cosets. Its 9 selected points place at least three points in one coset. That coset together with P lies in an affine 3-flat containing at least seven points of S, again a contradiction. Thus (4.2) holds. □

The conclusion has an immediate Reed–Muller interpretation under the full-length evaluation-code convention of (2.4). The even cardinality of S makes $\mathbf{1}_S$ orthogonal to the constant function, and (4.2) makes it orthogonal to every linear function. Therefore (2.5) gives

$$\mathbf{1}_S \in \mathrm{RM}(1,5)^{\perp} = \mathrm{RM}(3,5). \tag{4.3}$$

4.2 Parallel-flat patterns

Fix a three-dimensional direction U. Its four cosets have intersection sizes n_1, n_2, n_3, n_4. The union of any two cosets is an affine hyperplane. Since $\mathbf{1}_S$ is orthogonal to every affine function, every affine-hyperplane intersection has even cardinality. Thus the four n_i have the same parity. They sum to 16 and are at most six, so their sorted pattern is one of

$$\begin{aligned} &E_1 = (0,4,6,6),\ E_2 = (2,2,6,6),\ E_3 = (2,4,4,6),\ E_4 = (4,4,4,4),\\ &O_1 = (1,5,5,5),\ O_2 = (3,3,5,5). \end{aligned} \tag{4.4}$$

Let a, b, c, d, e, f be the numbers of directions with these six patterns.

The incidence parameters also give power moments. Expanding $n^2 = n + 2\binom{n}{2}$ and $n^3 = n + 6\binom{n}{2} + 6\binom{n}{3}$ in (3.2)–(3.4),

$$\sum_A n_A^2 = 10880, \qquad \sum_A n_A^3 = 51200. \tag{4.5}$$

Subtracting the contribution of 155 copies of E_4 and scaling gives

$$6a + 4b + 2c + 3e + f = 240, \tag{4.6}$$

$$5a + 4b + 2c + \frac{5}{2}e + f = 240. \tag{4.7}$$

Therefore $a = e = 0$. In particular, no affine 3-flat meets S in zero or one point.

4.3 Excluding a nonzero cubic top part

Let h be the cubic homogeneous part of $\mathbf{1}_S$. In the Boolean ANF algebra, filtered by degree, the associated graded relation is $x_i^2 = 0$; hence the degree-three component

is the space of squarefree cubics, naturally $\Lambda^3(V^*)$. Affine translation changes only lower-degree terms, so the top part is acted on by the linear part of $\mathrm{AGL}(5,2)$. After choosing a volume form, exterior duality identifies $\Lambda^3(V^*)$ with alternating two-forms on V. Nonzero alternating forms in dimension five have rank two or four, so nonzero cubic top parts have two $\mathrm{GL}(5,2)$ types, represented by

$$h_2 = x_1x_2x_3, \qquad h_4 = x_5(x_1x_2 + x_3x_4). \tag{4.8}$$

Let t be the number of three-dimensional directions on which h restricts nontrivially. For h_2, this occurs exactly when the projection of the direction onto the first three coordinates is an isomorphism. Such a direction is the graph of a linear map $\mathbb{F}_2^3 \to \mathbb{F}_2^2$, giving

$$t = 2^{3\cdot 2} = 64. \tag{4.9}$$

For h_4, put $W = \{x_5 = 0\}$ and let $\omega = e_1^* \wedge e_2^* + e_3^* \wedge e_4^*$ be the nondegenerate alternating two-form on W represented by the polar form of $x_1x_2 + x_3x_4$. A direction with nonzero cubic restriction is not contained in W. Its intersection L with W has dimension two, and the restriction is nonzero exactly when $\omega|_L$ is nondegenerate. Of the 35 two-subspaces of this symplectic four-space, 15 are totally isotropic and 20 are nondegenerate. For each such L, the remaining generator is determined modulo L, giving four directions. Hence

$$t = 20 \cdot 4 = 80. \tag{4.10}$$

On a direction where the cubic part is nonzero, all four affine restrictions have degree three. A Boolean function in three variables has odd truth-table weight exactly when its cubic coefficient is nonzero. Since the pattern O_1 has been excluded, every such direction has pattern O_2. If N_j denotes the number of affine 3-flats with intersection size j, then

$$N_3 = N_5 = 2t. \tag{4.11}$$

Writing $N_2 = A$, $N_4 = B$, and $N_6 = C$, the zeroth, first, and second incidence moments yield

$$A = C = 120 - \frac{t}{2}, \qquad B = 380 - 3t. \tag{4.12}$$

Consequently

$$\begin{aligned} \sum_A \binom{|S \cap A|}{4} &= N_4 + 5N_5 + 15N_6 \\ &= 2180 - \frac{t}{2}. \end{aligned} \tag{4.13}$$

The right-hand side is 2148 when $t = 64$ and 2140 when $t = 80$.

A four-subset of V lies in seven affine 3-flats if it is an affine plane and in exactly one otherwise. If P is the number of affine-plane four-subsets contained in S, then independently

$$\sum_A \binom{|S \cap A|}{4} = \binom{16}{4} + 6P = 1820 + 6P \equiv 2 \pmod 6. \tag{4.14}$$

But $2148 \equiv 0 \pmod 6$ and $2140 \equiv 4 \pmod 6$. Both nonzero cubic types are impossible. Together with (4.3), this proves

$$\mathbf{1}_S \in \mathrm{RM}(2,5). \tag{4.15}$$

In particular, every affine 3-flat meets S in an even number of points. This is the central equality bridge.

4.4 Classical quadratic consequences

We now use the standard classification of quadratic Boolean functions by polar rank and the radical derivative [2, pp. 194–200]; see also [6, 7]. Up to affine equivalence, the balanced quadratics in five variables have the three forms shown in Table 1; the count and the rank-four orbit reduction are included below to close the argument explicitly.

Table 1 Balanced quadratic types in five variables and their affine-3-flat spectra.

Polar rank	Normal form	Spectrum {weight : flats}	Cap
0	x_5	$\{0:30,\ 4:560,\ 8:30\}$	8
2	$x_1x_2 + x_5$	$\{0:6,\ 2:96,\ 4:416,\ 6:96,\ 8:6\}$	8
4	$x_1x_2 + x_3x_4 + x_5$	$\{2:120,\ 4:380,\ 6:120\}$	6

For the rank-four form, every affine 3-flat restriction has degree at most two and therefore even weight. A constant restriction would require a three-dimensional totally isotropic direction. Such a direction contains the unique nonzero radical vector, but translation by that vector toggles the function because the radical derivative is one. Hence weights zero and eight are impossible, and the zeroth, first, and second incidence-moment equations force the spectrum (4.1). The rank-zero and rank-two types have cap eight; details of their spectra are recorded in Appendix A.

The vector space of alternating bilinear forms on V has dimension $\binom{5}{2} = 10$, hence contains $2^{10} = 1024$ forms. A rank-two form has a three-dimensional radical. Choosing that radical gives $\begin{bmatrix}5\\3\end{bmatrix}_2 = 155$ possibilities, and the two-dimensional quotient supports a unique nonzero alternating form over $\mathbb{F}_2$. Thus the rank strata have sizes

$$1, \qquad 155, \qquad 1024 - 1 - 155 = 868 \tag{4.16}$$

for ranks 0, 2, and 4.

Fix a rank-$2r$ polar form B, its radical R, and one quadratic refinement q_0. Every quadratic with polar form B is uniquely $q_0 + \ell + c$, with ℓ linear and $c \in \mathbb{F}_2$. The map from linear perturbations to the radical derivative $\lambda_{q_0+\ell}|_R$ is affine, and each fiber has size 2^{2r} because the annihilator of R has dimension $2r$. Consequently $2^5 - 2^{2r}$ linear perturbations give a nonzero radical derivative, which is exactly the balance

criterion, and either constant is allowed. The numbers of balanced Boolean functions of the three ranks are therefore

$$2(32-1)=62, \qquad 155\cdot 2(32-4)=8680, \qquad 868\cdot 2(32-16)=27776. \tag{4.17}$$

For completeness, the balanced rank-four functions form one affine orbit. Choose a symplectic basis on a four-dimensional complement of the radical and a radical vector e_5, so a fixed refinement is $q_0 = x_1x_2 + x_3x_4$. The balance criterion forces the radical derivative in the e_5 direction to be one. Translation in the nondegenerate four-space removes the remaining linear functional, because $t \mapsto B(t,\cdot)$ is an isomorphism there; translation by e_5 then toggles and removes the constant. The resulting normal form is $q_\star = x_1x_2 + x_3x_4 + x_5$. Thus every cap-six halfset belongs to this single $\mathrm{AGL}(5,2)$ orbit, whose 27,776 labelled members are exactly the balanced rank-four functions counted above. This completes the proof of Theorem 4.1, except for the explicit trace representative, which we give next.

4.5 Trace-cube representative and an optimal Level S_1 design

Work in

$$\mathbb{F}_{32} = \mathbb{F}_2[z]/(z^5+z^2+1) \tag{4.18}$$

and let $\mathrm{Tr}:\mathbb{F}_{32}\to\mathbb{F}_2$ be the absolute trace. For $c\neq 0$, define

$$q_c(x) = \mathrm{Tr}(cx^3).$$

Its polar form is

$$B_c(x,y) = \mathrm{Tr}\big(c(x^2y+xy^2)\big). \tag{4.19}$$

Using nondegeneracy of the trace pairing, the nonzero radical equation is $y^{14} = c^{-15}$. Since $\gcd(14,31)=1$, there is a unique nonzero solution, $y=c^{10}$, and

$$q_c(c^{10}) = \mathrm{Tr}(c^{31}) = \mathrm{Tr}(1) = 1.$$

Thus every q_c is balanced of polar rank four. By the classification above, its support lies in the cap-six orbit.

An explicit affine certificate is useful for reproducibility. In the polynomial basis $(1,z,z^2,z^3,z^4)$, let ϕ have offset 7 and linear columns

$$[2,4,12,18,1]. \tag{4.20}$$

Equivalently,

$$\begin{aligned}\phi(x) = {} & 1+z+z^2+x_1z+x_2z^2 \\ & + x_3(z^2+z^3)+x_4(z+z^4)+x_5.\end{aligned}$$

Reduction in the field model (4.18) gives, for all $x\in\mathbb{F}_2^5$,

$$\mathrm{Tr}(\phi(x)^3) = x_1x_2 + x_3x_4 + x_5. \tag{4.21}$$

Moreover $3 \cdot 21 \equiv 1 \pmod{31}$, so

$$\mathrm{Tr}(cx^3) = \mathrm{Tr}((c^{21}x)^3).$$

Every nonzero trace-cube component is therefore linearly equivalent to $\mathrm{Tr}(x^3)$.

Define the zero-extended matrix

$$M_{3,3}(u,v) = \begin{cases} \mathrm{Tr}(u^3v^3), & u,v \neq 0, \\ 0, & uv = 0. \end{cases} \tag{4.22}$$

Its nonzero rows and columns are precisely the nonzero q_c supports. It is symmetric, 16-regular, and every row and column has cap six. Theorem 4.1 also gives zero XOR for every row and column. Hence $M_{3,3}$ is Level S_1 and

$$\Delta^{(0,3)}(M_{3,3}) = \Delta^{(3,0)}(M_{3,3}) = 12. \tag{4.23}$$

Corollary 4.4 (Equality automatically enforces Level S_1)**.** *Every Level* R *matrix satisfying*

$$\Delta^{(0,3)}(M) = \Delta^{(3,0)}(M) = 12$$

is automatically in Level S_1*.*

Proof Every nonzero row and column is a 16-set of cap at most six. Theorem 4.1 gives zero vector XOR for each support. □

Combining Corollary 3.3 with (4.23) already determines the Level R and Level S_1 axial minima. Section 5 determines every monomial-trace matrix that attains them.

5 Complete monomial-trace classification

Rønjom, Sandrib, and Sunde proposed the scored bivariate family $2\,\mathrm{Tr}(u^r v^s)$ as a candidate DDT indicator [3]. We classify the complete five-bit unit-exponent domain of its Boolean support; the factor two is restored by the support-matrix scoring convention (2.11).

For $r, s \in \mathbb{Z}_{31}^*$, define

$$M_{r,s}(u,v) = \begin{cases} \mathrm{Tr}(u^r v^s), & u,v \neq 0, \\ 0, & uv = 0. \end{cases} \tag{5.1}$$

The domain contains $30^2 = 900$ ordered exponent pairs. A *literal matrix* below means an identical labelled 32×32 Boolean array in the fixed field labelling, before any row, column, or transpose equivalence is imposed. Let

$$C_1 = \{1, 2, 4, 8, 16\},$$
$$C_3 = \{3, 6, 12, 17, 24\},$$

$$
\begin{aligned}
C_5 &= \{5, 9, 10, 18, 20\},\\
C_7 &= \{7, 14, 19, 25, 28\},\\
C_{11} &= \{11, 13, 21, 22, 26\},\\
C_{15} &= \{15, 23, 27, 29, 30\} = -C_1
\end{aligned}
\tag{5.2}
$$

be the binary cyclotomic cosets modulo 31. Put

$$
H = C_1 = \langle 2 \rangle, \qquad G = C_3 \cup C_5. \tag{5.3}
$$

Theorem 5.1 (Trace-exponent classification)**.** *For the matrices* (5.1)*:*

(i) *for every ordered pair* (r, s)*, the matrix* $M_{r,s}$ *lies in Level* R*;*

(ii) $M_{r,s}$ *is Level* S_1 *if and only if* $r, s \notin C_{15}$*;*

(iii) $M_{r,s}$ *has simultaneous axial score* $(12, 12)$ *if and only if* $r, s \in G$*;*

(iv) *literal equality is characterized by*

$$
M_{r,s} = M_{r',s'} \quad \Longleftrightarrow \quad (r', s') = (2^k r, 2^k s) \pmod{31} \tag{5.4}
$$

for a common $k \in \mathbb{Z}/5\mathbb{Z}$*;*

(v) $M_{r,s}$ *is symmetric if and only if* $r = s$*;*

(vi) *the optimal matrices form exactly three product-linear equivalence classes after adjoining transpose, represented by*

$$
M_{3,3}, \qquad M_{3,5}, \qquad M_{5,5}.
$$

The exact counts are shown in Table 2.

Table 2 Counts in the complete 900-pair trace domain. Exponent pairs and literal matrices are different quotients.

Property	Ordered exponent pairs	Literal matrices
Level R	900	180
Level S_1	625	125
Simultaneous axial score $(12, 12)$	100	20
Symmetric	30	6
Symmetric and Level S_1	25	5
Symmetric and optimal	10	2

5.1 Regularity and the exact XOR criterion

For a unit exponent d, the map $x \mapsto x^d$ permutes $\mathbb{F}_{32}$. Hence every nonzero row and column of (5.1) is a permuted support of a nonzero trace functional and has weight 16. This proves Level R.

To determine Level S_1, for $c \neq 0$ define the support XOR

$$X_d(c) = \bigoplus_{\mathrm{Tr}(cx^d)=1} x. \tag{5.5}$$

Interpreting XOR as field addition and expanding the absolute trace,

$$\begin{aligned} X_d(c) &= \sum_{x \in \mathbb{F}_{32}} x \,\mathrm{Tr}(cx^d) \\ &= \sum_{i=0}^{4} c^{2^i} \sum_{x \in \mathbb{F}_{32}} x^{1+d2^i}. \end{aligned} \tag{5.6}$$

The inner power sum, taken in $\mathbb{F}_{32}$, equals one exactly when $d2^i \equiv -1 \pmod{31}$ and equals zero otherwise. Therefore

$$X_d(c) = \begin{cases} 0, & d \notin C_{15}, \\ c^{2^i} \neq 0, & d2^i \equiv -1 \pmod{31}. \end{cases} \tag{5.7}$$

Rows have component exponent s and columns have component exponent r. This proves the criterion in Theorem 5.1(ii). It also shows why the XOR constraint is genuinely stronger than regularity: exactly one of the six cyclotomic exponent classes violates it.

5.2 The optimal exponent set

For a fixed unit d, all nonzero functions $x \mapsto \mathrm{Tr}(cx^d)$ are linearly equivalent. Indeed, for $c, c' \neq 0$ there is an $\alpha \neq 0$ with $\alpha^d = c'/c$, and then $\mathrm{Tr}(c'x^d) = \mathrm{Tr}(c(\alpha x)^d)$. The standard monomial degree formula states that a nonzero trace function $\mathrm{Tr}(cx^d)$ has algebraic degree equal to the binary weight $\mathrm{wt}_2(d)$ [2, Prop. 7, p. 63]. Multiplication by 2 modulo 31 cyclically rotates the five-bit expansion, so this weight is constant on each binary cyclotomic coset. Therefore

coset	C_1	C_3, C_5	C_7, C_{11}	C_{15}
degree	1	2	3	4.

(5.8)

The C_1 components are nonconstant linear functions and have cap eight.

For $d = 2^k + 1$ with $k = 1$ or 2, the polar equation of $q_c(x) = \mathrm{Tr}(cx^d)$ has a unique nonzero radical vector. Indeed, if $r \neq 0$ is radical and $t = cr^{2^k+1}$, the equation reduces to $t^{2^k} = t$. Since $\gcd(k, 5) = 1$, the fixed field is $\mathbb{F}_2$, so $t = 1$. Moreover $q_c(r) = \mathrm{Tr}(1) = 1$. Thus the representative exponents 3 and 5 give balanced rank-four quadratics.

This conclusion covers their full cyclotomic cosets. If $e \equiv 2^j d \pmod{31}$, Frobenius invariance of the absolute trace gives

$$\mathrm{Tr}(cx^e) = \mathrm{Tr}\big((cx^{2^j d})^{2^{5-j}}\big) = \mathrm{Tr}\big(c^{2^{5-j}} x^d\big). \tag{5.9}$$

Hence the component families for e and d coincide, and every component with exponent in C_3 or C_5 is a balanced rank-four quadratic with cap six by Theorem 4.1.

Conversely, a component of cap six would have an indicator of algebraic degree at most two and polar rank four by Theorem 4.1. The degree table excludes C_7, C_{11}, and C_{15}, while C_1 has cap eight. Thus the good exponents are exactly

$$G = C_3 \cup C_5.$$

Rows use s and columns use r, proving Theorem 5.1(iii).

5.3 Literal equality and symmetry

Choose a primitive element $g \in \mathbb{F}_{32}$ and define the binary trace sequence

$$t_j = \mathrm{Tr}(g^j), \qquad j \in \mathbb{Z}_{31}.$$

On nonzero indices,

$$M_{r,s}(g^i, g^j) = t_{ri+sj}. \tag{5.10}$$

Suppose $M_{r,s} = M_{r',s'}$, and let $L, L' : \mathbb{Z}_{31}^2 \to \mathbb{Z}_{31}$ be the corresponding nonzero linear forms. If $\ker L \neq \ker L'$, choose $h \in \ker L \setminus \ker L'$. Along a translate by the 31 multiples of h, the value of L is fixed while L' runs through every element of $\mathbb{Z}_{31}$. Equation (5.10) would force the trace sequence to be constant, a contradiction. Hence $\ker L = \ker L'$ and $L' = \lambda L$ for some unit λ.

It remains to determine the multiplier stabilizer of t_j. Equality $\mathrm{Tr}(x^\lambda) = \mathrm{Tr}(x)$ for every nonzero x is equivalent to equality of the reduced polynomials

$$\sum_{i=0}^{4} X^{\lambda 2^i} \quad \text{and} \quad \sum_{i=0}^{4} X^{2^i} \pmod{X^{31} - 1}.$$

Their exponent supports agree exactly when $\lambda H = H$, namely when $\lambda \in H$. Frobenius invariance gives the converse. This proves (5.4).

Every common-Frobenius orbit has size five. Dividing the ordered-pair counts by five gives 180 literal matrices in total, 125 Level S_1 matrices, and 20 optimal matrices. A canonical list of optimal literal representatives is

$$\{(3, s) : s \in G\} \ \cup \ \{(5, s) : s \in G\}. \tag{5.11}$$

Transpose sends $M_{r,s}$ to $M_{s,r}$. If a matrix is symmetric, (5.4) gives $(s, r) = h(r, s)$ for some $h \in H$, and hence $h^2 = 1$. Since H has order five, $h = 1$ and $r = s$. The converse is immediate. This proves Theorem 5.1(v) and the symmetric counts in Table 2.

5.4 Exactly three optimal geometric classes

The maps $u \mapsto u^{2^i}$ and $v \mapsto v^{2^j}$ are independent $\mathbb{F}_2$-linear permutations. They multiply r and s independently by elements of H. Thus the 100 optimal pairs reduce under product-linear equivalence to the four ordered types

$$C_3 \times C_3, \quad C_3 \times C_5, \quad C_5 \times C_3, \quad C_5 \times C_5. \tag{5.12}$$

Transpose merges the middle two.

To show that the three remaining types do not merge, we use diagonal affine-SDU coordinates, which are invariant under product-linear maps and transpose. Exact enumeration gives Table 3.

Table 3 Diagonal affine-SDU invariants separating the three optimal trace types.

Representative	$\Delta^{(3,3)}$	$\Delta^{(4,4)}$
$M_{3,3}$	84	280
$M_{3,5}$	88	280
$M_{5,5}$	92	292

The distinct $(3, 3)$ values separate the three types and prove Theorem 5.1(vi). The values were obtained by two independent exact implementations; the finite enumeration supplies a complete invariant calculation, not a heuristic sample.

Figure 2 displays the quotient structure. In particular, within the monomial-trace family, $M_{3,3}$ and $M_{5,5}$ are inequivalent symmetric optimizers. This is a family-specific nonuniqueness theorem; no uniqueness or classification is asserted among all Level S_1 optimizers.

5.5 Component spaces and scope

For later reference, define

$$\mathcal{Q}_d = \left\{ x \mapsto \mathrm{Tr}(cx^d) : c \in \mathbb{F}_{32} \right\}. \tag{5.13}$$

Reduced-polynomial supports give

$$\mathcal{Q}_d = \mathcal{Q}_e \quad \Longleftrightarrow \quad e \in dH. \tag{5.14}$$

The nonzero rows of $M_{r,s}$ form $\mathcal{Q}_s \setminus \{0\}$, and its nonzero columns form $\mathcal{Q}_r \setminus \{0\}$. Every individual nonzero member of $\mathcal{Q}_3$ or $\mathcal{Q}_5$ belongs to the same balanced rank-four affine orbit, even though the two full component spaces are not identified by (5.14). This statement concerns component spaces and the product-linear equivalence

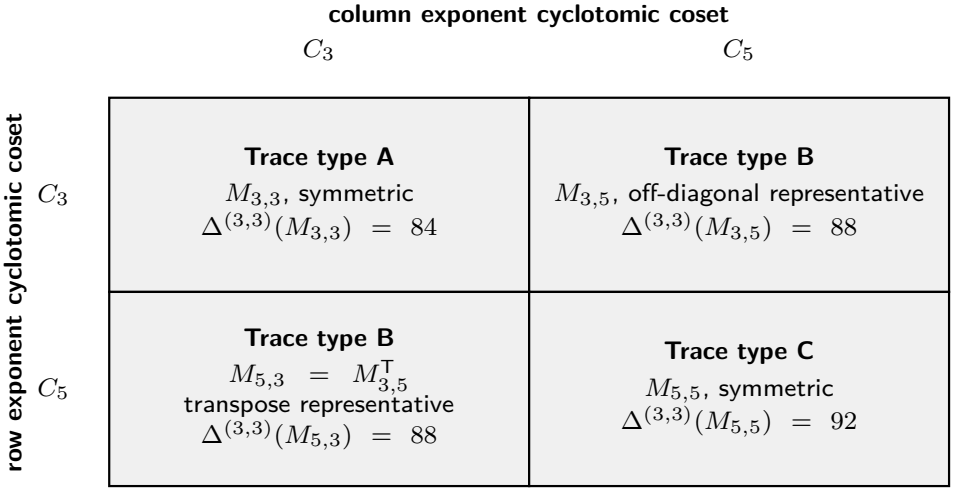

Fig. 2 Quotient of the 100 optimal exponent pairs under independent Frobenius actions and transpose. Each $C_i \times C_j$ block forms one product-linear type, and transpose identifies the two off-diagonal blocks. Each displayed value is $\Delta^{(3,3)}(M_{r,s})$; the values 84, 88, and 92 separate the three resulting optimal trace-matrix types

of bivariate support matrices; it does not invoke EA or CCZ equivalence of vectorial power mappings.

The complete theorem concerns exactly the 900 ordered unit-exponent pairs in (5.1). Simultaneous multiplication by arbitrary exponents, exponent inversion, or arbitrary row and column permutations would apply nonlinear relabellings and are not included in the geometric quotient.

6 A direct Walsh-square obstruction to DDT realizability

The trace optimizers satisfy the local regularity and XOR constraints, but they fail a global necessary condition for the DDT of any mapping. The argument is support-only and independent of the published exhaustive affine classification of five-bit APN permutations.

6.1 The DDT–Walsh-square identity

For $F : V \to V$, define the unnormalized vectorial Walsh transform

$$W_F(u, v) = \sum_{x \in V} (-1)^{u \cdot x + v \cdot F(x)}. \tag{6.1}$$

The raw DDT and Walsh transform satisfy the classical identity [2, 3, 8]

$$W_F(u,v)^2 = \sum_{a,b\in V} \delta_F(a,b)(-1)^{u\cdot a+v\cdot b}. \tag{6.2}$$

For completeness,

$$\begin{aligned} W_F(u,v)^2 &= \sum_{x,y\in V} (-1)^{u\cdot(x+y)+v\cdot(F(x)+F(y))} \\ &= \sum_{a,x\in V} (-1)^{u\cdot a+v\cdot(F(x)+F(x+a))}, \end{aligned}$$

and grouping by $b = F(x) + F(x+a)$ gives (6.2). No normalization factor occurs with convention (6.1).

6.2 A support matrix forces a unique raw candidate table

For $a \neq 0$, solutions counted by $\delta_F(a,b)$ occur in pairs $x \leftrightarrow x+a$, so every positive DDT entry is even. If a Level R support matrix M is the nonzero derivative support of a mapping F, then each nonzero row of M has exactly 16 positive positions. The corresponding positive even DDT entries sum to 32, so they must all equal two and no additional entry in that row can be positive. Thus M forces the raw candidate table

$$D_M(a,b) = \begin{cases} 32, & (a,b) = (0,0), \\ 2M(a,b), & a,b \neq 0, \\ 0, & \text{otherwise.} \end{cases} \tag{6.3}$$

Define its two-dimensional character transform by

$$H_M(u,v) = \sum_{a,b\in V} D_M(a,b)(-1)^{u\cdot a+v\cdot b}. \tag{6.4}$$

Proposition 6.1 (Walsh-square admissibility)**.** *Let M be a Level* R *support matrix. If a mapping $F : V \to V$ has M as its nonzero derivative support, then*

$$H_M(u,v) = W_F(u,v)^2$$

for all $u, v \in V$. In particular, every $H_M(u,v)$ is a nonnegative perfect square.

Proof The preceding row-sum argument shows that $D_M = \delta_F$. The identity now follows from (6.2). □

A single negative coefficient therefore rules out every mapping, not only APN permutations.

6.3 An exact four-count certificate for the trace cube

Use the field model (4.18) with low-to-high polynomial-basis bits. The absolute-trace values on the basis are

$$\big(\mathrm{Tr}(1), \mathrm{Tr}(z), \mathrm{Tr}(z^2), \mathrm{Tr}(z^3), \mathrm{Tr}(z^4)\big) = (1, 0, 0, 1, 0),$$

so the packed-bit mask for $x \mapsto \mathrm{Tr}(x)$ is 9.

For $M = M_{3,3}$, partition the 496 support cells by the two trace bits of a and b. The exact counts are

	$\mathrm{Tr}(b) = 0$	$\mathrm{Tr}(b) = 1$
$\mathrm{Tr}(a) = 0$	110	130
$\mathrm{Tr}(a) = 1$	130	126.

(6.5)

The support character sum at $(u, v) = (9, 9)$ is therefore

$$110 - 130 - 130 + 126 = -24.$$

After restoring the raw DDT origin from (6.3),

$$H_{M_{3,3}}(9, 9) = 32 + 2(-24) = -16. \tag{6.6}$$

This is incompatible with Proposition 6.1.

Theorem 6.2 (Direct nonrealizability of the trace cube)**.** *No mapping $F : \mathbb{F}_2^5 \to \mathbb{F}_2^5$ has $M_{3,3}$ as its nonzero derivative support. In particular, the trace-cube Level S_1 optimizer is not the DDT support of an APN permutation.*

Proof Equation (6.6) contradicts the necessary Walsh-square identity. □

Figure 3 records the complete certificate: one field model, one character mask, four cell counts, and one integer coefficient.

6.4 Excluding the complete optimal trace family

The three optimal trace classes have the full transform histograms shown in Table 4. Each histogram contains negative coefficients.

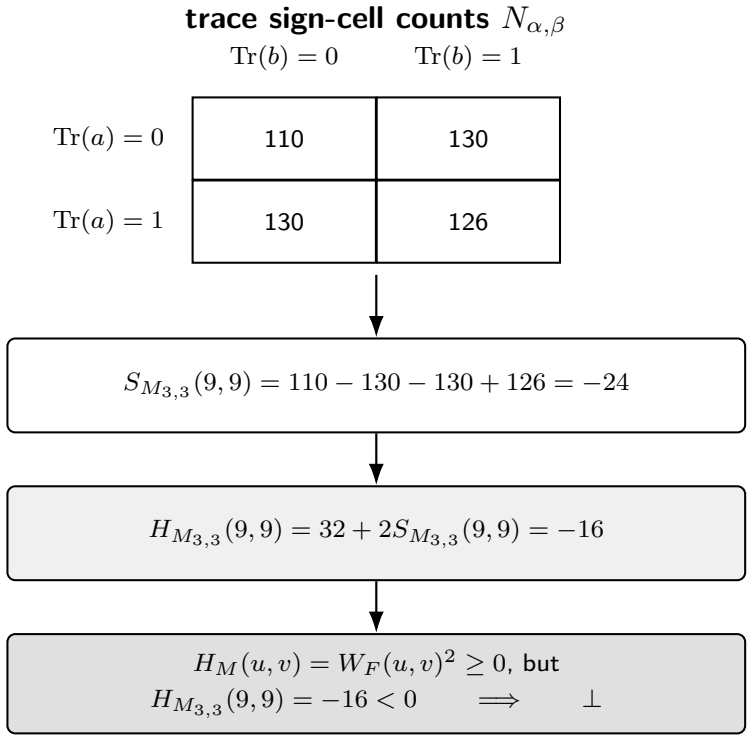

Fig. 3 Negative Walsh-square certificate for $M_{3,3}$. The displayed entries are the trace sign-cell counts $N_{\alpha,\beta} = \#\{(a,b) \in \mathrm{supp}(M_{3,3}) : \mathrm{Tr}(a) = \alpha, \mathrm{Tr}(b) = \beta\}$. They sum to 496 and give signed character sum -24. Here 9 is the packed polynomial-basis mask of the absolute-trace character. Restoring $D_M(0,0) = 32$ and doubling the Boolean support yields $H_{M_{3,3}}(9,9) = -16$, contradicting $H_M(u,v) = W_F(u,v)^2 \geq 0$

Table 4 Exact histograms of the 1024 coefficients $H_M(u,v)$ for the three optimal trace types.

Type	Histogram {coefficient : multiplicity}
$M_{3,3}$	$\{-16 : 31,\ 0 : 217,\ 16 : 310,\ 32 : 155,\ 48 : 155,\ 96 : 155,\ 1024 : 1\}$
$M_{3,5}$	$\{-96 : 31,\ 0 : 217,\ 32 : 465,\ 64 : 310,\ 1024 : 1\}$
$M_{5,5}$	$\{-16 : 155,\ 0 : 62,\ 16 : 155,\ 32 : 310,\ 48 : 155,\ 64 : 155,\ 144 : 31,\ 1024 : 1\}$

In the primary polynomial basis, convenient negative witnesses are

$$H_{3,3}(1,19) = -16, \qquad H_{3,5}(1,13) = -96, \qquad H_{5,5}(1,4) = -16. \tag{6.7}$$

A product-linear change of the two axes permutes the character coefficients by the dual linear maps, and transpose swaps the two character variables. Hence negativity is invariant under the equivalences used in Theorem 5.1.

Theorem 6.3 (Optimal trace-family exclusion)**.** *No matrix $M_{r,s}$ arising from any of the* 100 *optimal ordered exponent pairs is realizable as the nonzero derivative support of a mapping $\mathbb{F}_2^5 \to \mathbb{F}_2^5$. Consequently, none of the* 20 *optimal literal matrices or the three product-linear/transpose types in Theorem 5.1 is realizable.*

Proof Every optimal monomial-trace matrix belongs to one of the three equivalence types, and each type has a negative coefficient in (6.7). The exact checker also evaluates the matrices associated with all 100 optimal exponent pairs directly, providing a redundant complete-family verification. □

As positive controls, applying the same transform to the genuine raw DDTs of the five APN representatives used in Section 7 gives only nonnegative squares. Four representatives have histogram

$$\{0 : 527,\ 64 : 496,\ 1024 : 1\},$$

and the x^{15} representative has

$$\{0 : 217,\ 16 : 465,\ 64 : 310,\ 144 : 31,\ 1024 : 1\}.$$

Thus the negative trace coefficients are not a normalization artifact.

Remark 6.4 (What the direct obstruction does not prove). Theorem 6.3 excludes the completely classified optimal monomial-trace family without using any APN-permutation classification. It does not by itself determine the minimum over all realizable APN permutations; the latter universal quantifier is addressed separately in Section 7.

7 The exact five-bit axial gap

We now combine the finite-geometric lower bound, the trace construction, and the realizable-domain classification. The proof is deliberately factored so that the external completeness theorem enters only at the final Level P step.

7.1 Exact relaxed and XOR-balanced optima

Theorem 7.1 (Level R and Level S_1 optima)**.** *For the two axial affine-SDU coordinates,*

$$\min_{M \in \mathrm{R}} \Delta^{(0,3)}(M) = \min_{M \in \mathrm{S}_1} \Delta^{(0,3)}(M) = 12, \tag{7.1}$$

$$\min_{M \in \mathrm{R}} \Delta^{(3,0)}(M) = \min_{M \in \mathrm{S}_1} \Delta^{(3,0)}(M) = 12. \tag{7.2}$$

Moreover, every Level R *matrix attaining both minima is automatically in Level* S_1.

Proof Corollary 3.3 gives the lower bounds. The trace-cube matrix (4.22) is Level S_1 and attains both values by (4.23). The final statement is Corollary 4.4. □

For reference, the complete affine-SDU profile of the scored matrix $2M_{3,3}$, indexed by dimensions $0, \ldots, 5$ on each axis, is

$$\begin{pmatrix} 2 & 4 & 8 & 12 & 20 & 32 \\ 4 & 8 & 16 & 24 & 40 & 64 \\ 8 & 16 & 28 & 48 & 80 & 128 \\ 12 & 24 & 48 & 84 & 152 & 256 \\ 20 & 40 & 80 & 152 & 280 & 512 \\ 32 & 64 & 128 & 256 & 512 & 992 \end{pmatrix}. \tag{7.3}$$

This profile is an exact enumeration; only its axial entries are needed for Theorem 7.1.

7.2 The realizable minimum

Let $\mathcal{A}_5^{\mathrm{perm}}$ denote the set of all five-bit APN permutations.

Theorem 3 of Brinkmann and Leander [4] proves that there are exactly five ordinary affine-equivalence classes of five-bit APN permutations; Table 1 of their paper gives representatives labelled

$$x^{15}, \quad x^{11}, \quad x^7, \quad x^3, \quad x^5. \tag{7.4}$$

Equation (2.13) proves that every affine-SDU coordinate is invariant under precisely this equivalence.

Direct recomputation of the DDTs and all 620 affine 3-flats gives Table 5. The lookup tables in Table 1 of Brinkmann and Leander were checked literally; separately generated finite-field power maps are used only as transported cross-checks, not as a claim about an unspecified source field basis.

Table 5 Axial affine-SDU profiles of the five affine classes of five-bit APN permutations.

Representatives	$\Delta^{(0,3)}$	$\Delta^{(3,0)}$
x^3, x^5	16	14
x^7, x^{11}	14	16
x^{15}	14	14

The transpose relation is consistent with inversion. For every permutation,

$$\delta_{F^{-1}}(a,b) = \delta_F(b,a), \qquad \Delta^{(d_1,d_2)}(F^{-1}) = \Delta^{(d_2,d_1)}(F). \tag{7.5}$$

The asymmetric deletion of zero input difference does not introduce an unmatched boundary term because the raw zero row and, for a permutation, the raw zero column are both trivial away from the origin.

Theorem 7.2 (Classification-backed realizable optimum)**.** *Among five-bit APN permutations,*

$$\min_{F\in\mathcal{A}_5^{\mathrm{perm}}} \Delta^{(0,3)}(F) = 14, \qquad \min_{F\in\mathcal{A}_5^{\mathrm{perm}}} \Delta^{(3,0)}(F) = 14. \tag{7.6}$$

The x^{15} affine class attains both minima simultaneously.

Proof By Theorem 3 of Brinkmann and Leander [4], every five-bit APN permutation is affine-equivalent to one of the five representatives in (7.4). Affine invariance reduces the minimum to Table 5, whose entries were recomputed directly. □

Remark 7.3 (External trust boundary). Theorem 7.2 is classification-backed. The exhaustive enumeration underlying Theorem 3 of Brinkmann and Leander [4], which proves that the five affine classes represented in (7.4) are complete among five-bit APN permutations, is imported rather than replayed. This work independently checks the five source representatives, their APN and permutation properties, affine transport of SDU, inverse transport, and the displayed profiles. The halfset classification, trace classification, and direct nonrealizability theorems do not depend on the completeness claim.

7.3 Relaxed versus realizable

Theorem 7.4 (Exact axial realizability gap)**.** *For the two axial affine-SDU coordinates,*

$$\begin{aligned} \min_{M\in\mathrm{R}} \Delta^{(0,3)}(M) &= \min_{M\in\mathrm{S}_1} \Delta^{(0,3)}(M) = 12 < 14 = \min_{M\in\mathrm{P}} \Delta^{(0,3)}(M), \\ \min_{M\in\mathrm{R}} \Delta^{(3,0)}(M) &= \min_{M\in\mathrm{S}_1} \Delta^{(3,0)}(M) = 12 < 14 = \min_{M\in\mathrm{P}} \Delta^{(3,0)}(M). \end{aligned} \tag{7.7}$$

In particular, the strict relaxed-versus-realizable gap holds simultaneously on both axes.

Proof The Level R and Level S_1 equalities follow from Theorem 7.1. By Definition 2.3 and the agreement between the function and support-matrix SDU conventions established in Section 2.5, minimizing over $M \in \mathrm{P}$ is equivalent to minimizing over five-bit APN permutations. The Level P equalities therefore follow from Theorem 7.2. □

The two nonrealizability arguments have different logical roles. Theorem 6.3 directly excludes every optimal monomial-trace design, independently of the classification of APN permutations. Theorem 7.2 is stronger in a different direction: it proves that no five-bit APN permutation, trace-based or otherwise, attains axial score 12. Its universal quantifier is precisely where the published completeness theorem is required.

8 Odd-dimensional perspective

The five-bit trace cube is a member of a broader odd-dimensional construction, but the exact five-bit optimum and realizability gap do not automatically extend. We retain

only the construction-level consequence in the main text and move the restriction spectra, existence proof, probabilistic comparison, and bounded transform observations to Appendix B.

Let $n = 2m + 1 \geq 5$, let $1 \leq k < n$ with $\gcd(k, n) = 1$, and put $d = 2^k + 1$. The functions

$$q_c(x) = \mathrm{Tr}_{\mathbb{F}_{2^n}/\mathbb{F}_2}(cx^d), \qquad c \neq 0,$$

are the nonzero Boolean components $x \mapsto \mathrm{Tr}(cF(x))$ of the Gold power map $F(x) = x^d$. Each q_c is a balanced quadratic whose polar form has maximal rank $2m$ [2, 10]. Define the zero-extended symmetric support matrix

$$M_k(a, b) = \begin{cases} \mathrm{Tr}_{\mathbb{F}_{2^n}/\mathbb{F}_2}\big((ab)^d\big), & a, b \neq 0, \\ 0, & ab = 0. \end{cases} \tag{8.1}$$

Proposition 8.1 (Odd-dimensional Gold-trace construction)**.** *Every nonzero row and column of* (8.1) *has weight* 2^{n-1} *and vector XOR zero. At affine-flat dimension* $m + 1$*,*

$$\Delta^{(0,m+1)}(M_k) = \Delta^{(m+1,0)}(M_k) = 3 \cdot 2^m. \tag{8.2}$$

For $n = 3$*, the same support matrix is regular but its row and column XORs are nonzero.*

Proof Maximal-rank quadratic restriction theory gives row and column cap $3 \cdot 2^{m-1}$. Multiplying by the score factor two gives (8.2). For $n \geq 5$, each coordinate of a support XOR is the sum over $\mathbb{F}_2^n$ of a degree-at-most-three Boolean polynomial, hence vanishes because its degree is below n. The exceptional $n = 3$ power sum is nonzero. Appendix B supplies the complete derivation, including existence of every stated restriction weight. □

This proposition is a construction statement, not a classification or optimality theorem for arbitrary halfsets in general dimension. The present work does not determine the arbitrary-halfset optimum for $m = 3, \ldots, 7$. The appendix proves that maximal-rank quadratic halfsets are not universally optimal for $m \geq 8$, and the bounded negative transform computations at $n = 3, 5, 7, 9$ do not establish a nonrealizability theorem for all odd dimensions.

9 Computational verification and reproducibility

The argument separates human proofs, bounded exact computation, and external literature. The incidence lower bound and the arbitrary-halfset cap-six classification are proved mathematically in Sections 3 and 4; solver models and enumerators only corroborate those proofs. Exact computation is used for three finite tasks: classifying the stated 900 monomial-trace exponent pairs and their invariants, evaluating the character-transform certificates, and recomputing the DDT profiles of the five published APN representatives. These computations use integer or finite-field arithmetic over explicitly bounded domains rather than statistical sampling.

The principal finite results were checked through independent implementations where practical. The support classification and trace-family calculations have separate Python and C++ paths; the Walsh-square obstruction is evaluated by independent two-dimensional transforms; and genuine APN DDTs serve as positive controls. The source code, finite inputs, retained computational evidence, replay instructions, and integrity manifests supporting the reported computations are maintained in a companion reproducibility artifact, available from the corresponding author upon request.

The remaining trust boundary is Theorem 3 of Brinkmann and Leander, which states that the five APN-permutation representatives in their Table 1 exhaust the ordinary affine classes in dimension five. The companion reproducibility artifact recomputes those five representatives and their profiles but does not reproduce the historical exhaustive search proving completeness. Consequently, Theorem 7.2 and the universal gap theorem import that completeness result, whereas the incidence theorem, equality classification, trace-family classification, and direct Walsh-square exclusions do not. Optional solver replays are additional corroboration and are not treated as substitutes for either the human proofs or the external completeness theorem.

10 Conclusion

The five-bit axial SDU problem separates local extremality from global DDT realizability. A finite-geometric incidence certificate gives the common relaxed lower bound, and the equality analysis shows that every optimal row or column is a balanced rank-four quadratic halfset. This rigidity determines the relaxed axial score 12, makes the complete monomial-trace classification possible, and shows that the regular and vector-sum-balanced models attain the same optimum on both axes.

The strict gap is not merely a failed attempt to reconstruct a permutation from an attractive support. At score 12, the local row and column conditions are compatible, but the forced candidate DDTs of the optimal trace types violate global Fourier positivity: each has a negative coefficient where a vectorial Walsh square is required. Theorem 3 and Table 1 of the published Brinkmann–Leander classification then supply the universal realizable comparison: five-bit APN permutations form exactly five ordinary affine classes. Direct recomputation of one representative from each class gives minimum 14 on each axis, attained simultaneously by the x^{15} class, and hence

$$12 < 14.$$

The exhaustive completeness theorem is imported only for this universal Level P statement and its corollaries; the direct trace-family exclusions and the relaxed optimum do not depend on it.

The odd-dimensional Gold-trace construction shows that the five-bit trace cube belongs to a broader family of symmetric regular support matrices with zero row and column XOR and the classical maximal-rank quadratic axial score. It does not imply general-dimensional optimality or an infinite realizability gap. The arbitrary-halfset comparison remains unresolved here for $m = 3, \ldots, 7$, while Appendix B proves that maximal-rank quadratic halfsets are not universally optimal for $m \geq 8$. The most immediate next question is therefore structural: to identify support-level conditions

stronger than row and column XOR balance, but weaker than explicit permutation reconstruction, that capture the missing global DDT constraints.

A Supplementary proof details

A.1 Deriving the three balanced quadratic spectra

We record elementary derivations of Table 1.

For $q = x_5$, a three-dimensional direction contained in $\ker(x_5)$ gives two zero and two full cosets. There are $\begin{bmatrix}4\\3\end{bmatrix}_2 = 15$ such directions, hence 30 flats of weight zero and 30 of weight eight. The remaining 140 directions give four balanced restrictions, accounting for 560 flats of weight four.

For $q = x_1x_2 + x_5$, translation by the radical vector e_5 complements each restriction, so $N_j = N_{8-j}$. Each fiber over (x_1, x_2, x_3) is a full $\mathbb{F}_2^2$ in the remaining radical coordinates. The four-point graph of x_1x_2 in $\mathbb{F}_2^3$ is not an affine plane, because its four points have nonzero XOR. A constant affine 3-flat inside a fiber must therefore project to one of the six affine lines determined by two of these four points and contain the full two-dimensional kernel. This gives $N_0 = N_8 = 6$. The identities

$$\begin{aligned} N_0 + N_2 + N_4 + N_6 + N_8 &= 620, \\ 2N_2 + 4N_4 + 6N_6 + 8N_8 &= 2480, \\ N_2 + 6N_4 + 15N_6 + 28N_8 &= 4200 \end{aligned}$$

then give $N_2 = N_6 = 96$ and $N_4 = 416$.

For $q = q_\star$, the radical is $\langle e_5 \rangle$ and $D_{e_5}q = 1$. Any three-dimensional totally isotropic direction contains the radical, so no restriction can be constant. A degree-at-most-two Boolean function in three variables has even weight, leaving only weights two, four, and six. The zeroth, first, and second incidence-moment equations give $N_2 = N_6 = 120$ and $N_4 = 380$.

A.2 Counting balanced functions of each polar rank

The vector space of alternating bilinear forms on $\mathbb{F}_2^5$ has dimension $\binom{5}{2} = 10$ and therefore contains 1024 forms. Its rank strata have sizes

$$1, \qquad 155, \qquad 868$$

for ranks $0, 2, 4$, respectively. These numbers may be obtained from the standard orbit formula for alternating forms; they sum to 1024.

Fix a rank-$2r$ polar form B, its radical R, and a quadratic refinement q_0. Every quadratic with polar form B is $q_0 + \ell + c$. The relevant object is not the coordinate-dependent ANF linear part by itself but the radical derivative $r \mapsto D_r(q_0 + \ell)$. As ℓ varies, the perturbations giving zero radical derivative form an affine fiber of size 2^{2r}, because the annihilator of R has dimension $2r$. Hence $2^5 - 2^{2r}$ perturbations give a nonzero radical derivative and therefore a balanced function; either constant is allowed. Multiplying by the polar-form counts gives (4.17).

A.3 Details of the cubic fourth-incidence congruence

For completeness, a four-subset $T \subseteq V$ has one of two affine-span dimensions. If $\mathrm{XOR}(T) = 0$, then T is an affine plane. Its two-dimensional direction is contained in exactly

$$\begin{bmatrix} 3 \\ 1 \end{bmatrix}_2 = 7$$

three-dimensional directions, and each determines one affine 3-flat containing T. If $\mathrm{XOR}(T) \neq 0$, the four points are affinely independent and span a unique affine 3-flat. Summing over all four-subsets of S therefore gives (4.14).

The two cubic top-part counts can also be checked directly. For $h_2 = x_1x_2x_3$, restriction to a direction is nonzero exactly when the three coordinate forms remain independent; the direction is a graph over $\mathbb{F}_2^3$, giving 64. For $h_4 = x_5(x_1x_2 + x_3x_4)$, the intersection with $x_5 = 0$ must be one of the 20 nondegenerate planes of the symplectic four-space, with four graph extensions, giving 80.

A.4 The explicit trace affine certificate

In the field model (4.18), the packed integers in (4.20) denote polynomial-basis vectors. The five linear columns are

$$z, \quad z^2, \quad z^2 + z^3, \quad z + z^4, \quad 1.$$

They are linearly independent. Substituting their affine combination with offset $1+z+z^2$, reducing powers by $z^5 = z^2 + 1$, and applying the absolute trace gives (4.21). The companion checker evaluates the same identity on all 32 inputs; the finite evaluation is an exact certificate for the displayed map, while the orbit theorem independently establishes existence of an affine equivalence.

A.5 Literal trace-matrix representatives

The 20 optimal literal matrices may be represented by the exponent pairs

$$\{(3, s) : s \in G\} \cup \{(5, s) : s \in G\}.$$

Within each matrix, the remaining four ordered pairs are obtained by common Frobenius multiplication. The two symmetric literal matrices are represented by $(3, 3)$ and $(5, 5)$. Product-linear Frobenius maps collapse the 20 literal matrices to the four ordered coset rectangles in (5.12); transpose then gives the three classes in Theorem 5.1.

B Odd-dimensional details and bounded scope

This appendix gives the complete derivations underlying Section 8. The statements below describe what is proved in the present work; they do not assert that unresolved intermediate cases are absent from the literature.

B.1 Restriction weights of maximal-rank balanced quadratics

Let $n = 2m + 1$ and let $q : V_n \to \mathbb{F}_2$ be quadratic. Assume its polar form has rank $2m$, with unique nonzero radical vector r, and

$$D_r q = q(r) + q(0) = 1. \tag{B.1}$$

Equivalently, q is balanced and affinely equivalent to

$$q_m(z, x, y) = z + \sum_{i=1}^{m} x_i y_i. \tag{B.2}$$

Proposition B.1 (Classical restriction-weight consequence)**.** *As A ranges over the affine $(m+1)$-flats of V_n, the possible weights of $q|_A$ are exactly*

$$\{2^m\} \ \cup \ \left\{ 2^m \pm 2^{m-j} : 1 \le j \le \left\lfloor \frac{m+1}{2} \right\rfloor \right\}. \tag{B.3}$$

In particular,

$$\max_{A \in \mathcal{A}_{m+1}(V_n)} |A \cap \operatorname{supp}(q)| = 3 \cdot 2^{m-1}. \tag{B.4}$$

Proof Write $V_n = \langle r \rangle \oplus W$, where the induced polar form B on the $2m$-space W is symplectic. Let U be an $(m+1)$-dimensional direction.

If $r \in U$, translation by r pairs the points of every coset and toggles q, so every restriction is balanced and has weight 2^m.

Assume $r \notin U$. Projection onto W identifies U with an $(m+1)$-subspace $H \le W$. The restricted alternating form has rank $2j$. Since a symplectic $2m$-space has no totally isotropic subspace of dimension $m+1$,

$$1 \le j \le \left\lfloor \frac{m+1}{2} \right\rfloor.$$

The standard quadratic weight formula gives the balanced value 2^m when the restricted radical derivative is nonzero and the two values $2^m \pm 2^{m-j}$ when it is zero [2, 6, 7].

It remains to justify existence, not only necessity. Every displayed rank occurs by choosing j hyperbolic pairs and $m + 1 - 2j$ mutually orthogonal isotropic basis vectors. For such an H, let $R_H = \operatorname{rad}(B|_H)$. As the affine offset $w_0 \in W$ varies, the radical derivative of the restriction changes by

$$B(w_0, \cdot)|_{R_H}.$$

The map $W \to R_H^*$, $w_0 \mapsto B(w_0, \cdot)|_{R_H}$, is surjective because B is nondegenerate on W. Hence some affine translate has zero restricted radical derivative and realizes one of the two nonbalanced weights. Translating that affine flat by r complements the restriction, because $D_r q = 1$, and realizes the opposite sign. Thus every weight in (B.3) occurs. The largest is obtained for $j = 1$ and equals (B.4). □

For $m = 1, 2, 3$, exact enumeration of the normal form gives

n	flat dimension	full spectrum
3	2	$\{1 : 4,\ 2 : 6,\ 3 : 4\}$
5	3	$\{2 : 120,\ 4 : 380,\ 6 : 120\}$
7	4	$\{4 : 5040,\ 6 : 21504,\ 8 : 41400,\ 10 : 21504,\ 12 : 5040\}$.

(B.5)

The five-bit row agrees with Theorem 4.1.

B.2 Gold components and the zero-XOR construction

Let $n = 2m + 1$ be odd. Reduce k modulo n to k_0 with $1 \le k_0 < n$, assume $\gcd(k_0, n) = 1$, and put

$$d = 2^{k_0} + 1. \tag{B.6}$$

Then d is a unit modulo $2^n - 1$. The functions

$$q_c(x) = \mathrm{Tr}_{\mathbb{F}_{2^n}/\mathbb{F}_2}(cx^d), \qquad c \neq 0, \tag{B.7}$$

are the nonzero Boolean components of the Gold power map $F(x) = x^d$.

The polar radical equation is equivalent, for $r \neq 0$, to

$$t^{2^{k_0}} = t, \qquad t = cr^{2^{k_0}+1}.$$

The fixed field is $\mathbb{F}_2$, so $t = 1$. Since d is a unit, there is a unique nonzero radical vector, and

$$q_c(r) = \mathrm{Tr}(1) = 1$$

because n is odd. Thus every nonzero component is a balanced quadratic of polar rank $2m$. These maximal-rank properties are classical [2, 10].

Zero-extend the symmetric matrix

$$M_k(a, b) = \begin{cases} \mathrm{Tr}_{\mathbb{F}_{2^n}/\mathbb{F}_2}\big((ab)^d\big), & a, b \neq 0, \\ 0, & ab = 0. \end{cases} \tag{B.8}$$

Its nonzero rows and columns are supports of the q_c, so each has weight 2^{n-1}. Proposition B.1 gives cap $3 \cdot 2^{m-1}$ on affine $(m + 1)$-flats.

For $n \ge 5$, every coordinate of a support XOR is the parity of

$$\sum_{x \in V_n} \ell(x) q_c(x),$$

where ℓ is linear. The product has algebraic degree at most three, strictly below n, so its sum over V_n is zero. Hence every nonzero row and column has vector XOR zero, and the score factor two yields

$$\Delta^{(0,m+1)}(M_k) = \Delta^{(m+1,0)}(M_k) = 3 \cdot 2^m. \tag{B.9}$$

At $n = 3$, degree three is no longer below the ambient dimension; the exact power-sum calculation analogous to (5.6) gives a nonzero support XOR.

B.3 Why the quadratic cap is not universally optimal in large dimension

For an arbitrary halfset $S \subseteq V_{2m+1}$ and a uniformly random affine $(m+1)$-flat A, put $X = |S \cap A|$. The first three factorial moments agree with the corresponding hypergeometric moments. With $\mu = 2^m$,

$$\mathbb{E}(X-\mu)^3 = 0, \qquad \mathrm{Var}(X) = \frac{2^{2m}(2^m-1)}{2^{2m+1}-1}. \tag{B.10}$$

Let $Y = X - \mu$ and $M = \max Y$. Since $Y \le M$,

$$(M-Y)(Y+M)^2 \ge 0.$$

Using $\mathbb{E}Y = 0$ and $\mathbb{E}Y^3 = 0$, we obtain

$$0 \le \mathbb{E}[(M-Y)(Y+M)^2] = M\big(M^2 - \mathrm{Var}(X)\big).$$

Thus $M \ge \sqrt{\mathrm{Var}(X)}$. Taking the integer ceiling yields the universal lower bound

$$\max_A |S \cap A| \ge L_m = 2^m + \left\lceil 2^m \sqrt{\frac{2^m-1}{2^{2m+1}-1}} \right\rceil. \tag{B.11}$$

For $m = 1, 2$, this equals the quadratic cap $Q_m = 3 \cdot 2^{m-1}$; it is weaker from $m = 3$ onward.

More strongly, the quadratic cap is not the optimum among all halfsets for sufficiently large m. Let $K = 2^{m+1}$ be the size of a flat. For a uniformly random exact halfset and a fixed flat, a sampling-without-replacement Hoeffding bound gives

$$\Pr[X \ge 3K/4] \le \exp(-K/8) = \exp(-2^{m-2}). \tag{B.12}$$

The number of affine $(m+1)$-flats satisfies

$$N_m = 2^m \begin{bmatrix} 2m+1 \\ m+1 \end{bmatrix}_2 < 4 \cdot 2^{m^2+2m}. \tag{B.13}$$

The logarithm of the union bound is at most

$$B_m = (m^2 + 2m + 2)\log 2 - 2^{m-2}. \tag{B.14}$$

At $m = 8$, $B_8 = 82 \log 2 - 64 < 0$, and

$$B_{m+1} - B_m = (2m+3)\log 2 - 2^{m-2} < 0 \qquad (m \ge 8).$$

Hence for every $m \geq 8$ there is an exact halfset whose intersection with every affine $(m+1)$-flat is at most $Q_m - 1$.

Corollary B.2 (Failure of universal quadratic optimality)**.** *Maximal-rank quadratic supports are not universally optimal halfsets for any $m \geq 8$.*

The present work does not determine the arbitrary-halfset optimum for $m = 3, 4, 5, 6, 7$. It also does not establish an all-odd-dimensional DDT-support nonrealizability theorem. Exact computations at $n = 3, 5, 7, 9$ are bounded observations recorded in the companion reproducibility artifact; they are not used to assert a general nonrealizability theorem.

Reproducibility Artifact

A companion reproducibility artifact contains the finite-computation code, datasets, representative checks, computational certificates, solver records, replay instructions, and integrity manifests reported in this preprint. It is available from the corresponding author upon request.

Statements and Declarations

Funding. The authors did not receive support from any organization for the submitted work.

Competing interests. The authors have no relevant financial or non-financial interests to disclose.

Author contributions. Jingchuan Ma: Conceptualization, methodology, formal analysis, software, investigation, writing—original draft, and writing—review and editing. Yuqing Shao: Validation, data curation, visualization, and writing—review and editing. Xin Wei: Validation, data curation, and writing—review and editing. Qiaoyun Huang: Conceptualization, supervision, and writing—review and editing. All authors read and approved the final manuscript.

Data availability. The finite datasets, enumerated profiles, computational certificates, manifests, and result records supporting the findings of this study are maintained in the companion reproducibility artifact and are available from the corresponding author upon request.

Code availability. The generators, checkers, independent cross-checks, and bounded replay scripts used in this study are maintained in the companion reproducibility artifact and are available from the corresponding author upon request. The artifact contains environment information, replay instructions, expected results, and file-integrity manifests.

References

[1] Nyberg, K.: Differentially uniform mappings for cryptography. In: Helleseth, T. (ed.) Advances in Cryptology—EUROCRYPT '93. Lecture Notes in Computer Science, vol. 765, pp. 55–64. Springer, Berlin, Heidelberg (1994). https://doi.org/10.1007/3-540-48285-7_6

[2] Carlet, C.: Boolean Functions for Cryptography and Coding Theory. Cambridge University Press, Cambridge (2021). https://doi.org/10.1017/9781108606806

[3] Rønjom, S., Sandrib, A., Sunde, J.: Subspace differential uniformity. Cryptogr. Commun. (2026) https://doi.org/10.1007/s12095-026-00896-w

[4] Brinkmann, M., Leander, G.: On the classification of APN functions up to dimension five. Des. Codes Cryptogr. **49**(1–3), 273–288 (2008) https://doi.org/10.1007/s10623-008-9194-6

[5] MacWilliams, F.J., Sloane, N.J.A.: The Theory of Error-Correcting Codes. North-Holland, Amsterdam (1977)

[6] Carlet, C.: Partially bent functions. Des. Codes Cryptogr. **3**(2), 135–145 (1993) https://doi.org/10.1007/BF01388412

[7] Charpin, P., Pasalic, E., Tavernier, C.: On bent and semi-bent quadratic boolean functions. IEEE Trans. Inf. Theory **51**(12), 4286–4298 (2005) https://doi.org/10.1109/TIT.2005.858929

[8] Canteaut, A., Kölsch, L., Li, C., Li, C., Li, K., Qu, L., Wiemer, F.: Autocorrelations of vectorial boolean functions. In: Progress in Cryptology—LATINCRYPT 2021. Lecture Notes in Computer Science, vol. 12912, pp. 233–253. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-88238-9_12

[9] Kovács, B., Nagy, Z.L.: Avoiding intersections of given size in finite affine spaces AG$(n, 2)$. J. Comb. Theory Ser. A **209**, 105959 (2025) https://doi.org/10.1016/j.jcta.2024.105959

[10] Lahtonen, J., McGuire, G., Ward, H.N.: Gold and kasami–welch functions, quadratic forms, and bent functions. Adv. Math. Commun. **1**(2), 243–250 (2007) https://doi.org/10.3934/amc.2007.1.243