

Settling Conjectures on Linear Structures of Inverse ChiChi Generalizations

Four Proofs and a Counterexample

Hui Wang¹, Ricardo Rodriguez Reveco²  and Kai Hu¹

¹ Shandong University, China

² Ruhr University Bochum, Germany

Abstract

Belkheyar et al. introduced ChiChi as the even-dimensional, low-latency nonlinear core of ChiLow (Eurocrypt 2025), and Andreoli et al. generalized it into the SWAP and 4-cycle families (ToSC 2025/3), conjecturing five bounds on the linear height and component linear nullity of their inverses, supported by experiments in small dimensions. We settle all five conjectures: Conjectures 1–4 hold for every admissible branch length, and Conjecture 5 is false. In addition, we derive consequences for ChiLow’s design and correct one value in the data of Andreoli et al.

Of independent interest, we settle the single-cycle map χ_n^{-1} itself. For every odd $n \geq 3$ it has linear height at most one in every nonzero direction and linear nullity at most one in every nonzero component, with both equality loci classified.

Keywords: ChiChi, χ function, low-latency permutations, linear structures, linear height, linear nullity, affine hulls

1 Introduction

Daemen’s χ mapping is popular for its short hardware path and efficient bitsliced implementation [Dae95]. Each output coordinate depends on only three adjacent input coordinates, and the nonlinear part costs one layer of AND gates followed by XORs. Several permutation-based primitives build on these features, though each uses χ differently. The Keccak permutation underlying SHA-3 applies the standard five-bit χ map independently to its rows [NIS15]. The round function of Subterranean 2.0 contains a standard χ step on its 257-bit state [DMMR20]. Ascon instead applies 64 parallel instances of a five-bit S-box, which combines a χ -related quadratic core with lightweight affine transformations [DEMS21, NIS25].

However, χ_m is a permutation exactly when m is odd [Dae95]. This rules out its direct use as an even-dimensional nonlinear permutation, despite the prevalence of even and power-of-two state sizes. Applying two odd χ maps in parallel restores even dimension, but the two branches stay independent. The inverse then splits into two smaller inverses, one per branch, so no output coordinate depends on both halves of the state.

At Eurocrypt 2025, Belkheyar et al. introduced ChiLow, a family of tweakable block ciphers and a related pseudorandom function for low-latency embedded code encryption [BDG⁺25]. Its nonlinear core ChiChi (also written as $\chi\chi$) attains even dimension while

To help verify the proofs in this paper, we provide Lean 4 formalizations of the proofs, included in the auxiliary material; the README there describes what each development covers.

Emails: hui.wang2025@mail.sdu.edu.cn (Hui Wang), ricardo.rodriguezreveco@rub.de (Ricardo Rodriguez Reveco), kai.hu@sdu.edu.cn (Kai Hu)




Figure 1: Partial decryption through a last-round generalized ChiChi layer after guessing part of the last-round key, adapted from [ALPS25].

retaining the latency profile of χ . It achieves this by coupling two odd χ cycles through a small number of boundary coordinates, and so avoids the complete separation of the parallel construction. The two cycles are χ_m and χ_{m+2} for odd m , so the dimension is always divisible by four. Andreoli et al. later placed ChiChi in a broader framework [ALPS25]. Their SWAP and 4-cycle families pair χ_m with $\chi_{m'}$ for any odd m, m' , which reaches every even dimension rather than only the multiples of four. We refer to them collectively as the *generalized ChiChi families*. They introduced two criteria, component linear nullity and directional linear height, for the inverse dependencies relevant to partial key recovery.

Figure 1 illustrates why these two criteria matter in partial key recovery. Suppose a distinguisher exploits a non-random property of the state y immediately before a final generalized ChiChi layer. Starting from a ciphertext c , an attacker guesses part of the N -bit last-round key, partially inverts the layer, and tries to evaluate $\langle \lambda, y \rangle$, the linear combination of state bits selected by a binary mask λ . Write H for the inverse of the layer and $\Delta_\alpha f(x) = f(x + \alpha) + f(x)$ for the derivative along a difference α . For a fixed mask λ , the directions along which $\Delta_\alpha \langle \lambda, H \rangle$ is constant form the space E_λ , whose dimension ν is the *linear nullity* of the component. Operationally, key guesses separated by any of the 2^ν directions in E_λ change its value predictably and need not be tested independently. The 2^N candidates therefore fall into $2^{N-\nu}$ relevant classes, leaving only $N - \nu$ independent bits of key information to guess. A large nullity can thus make a chosen component accessible through a comparatively cheap partial-key guess, whereas a small nullity forces the attacker to resolve nearly the entire last-round key.

The complementary viewpoint fixes a nonzero direction α corresponding to an unresolved key relation and asks how many state combinations remain usable. The masks λ for which $\Delta_\alpha \langle \lambda, H \rangle$ is constant form $K_H(\alpha)$, whose dimension h is the *linear height* in that direction. A height h yields 2^h such masks; excluding the zero mask, these give $2^h - 1$ nontrivial linear combinations whose change along α is predictable without resolving the corresponding linear combination of key bits. A large height therefore offers the attacker many usable components, whereas a small height leaves very little choice. Since $\alpha \in E_\lambda$ exactly when $\lambda \in K_H(\alpha)$, the two criteria describe complementary sides of the same relation. Bounding both controls both the cost of evaluating one chosen component and the number of components that can benefit from an unguessed key relation.

The five conjectures of [ALPS25] bound the linear height and the linear nullity of the inverses of both the SWAP and 4-cycle families. None of them is proved there. Their evidence is a table of small branch lengths, computed exhaustively, while the statements range over every odd pair m, m' . Many primitives reach for χ when they want a nonlinear layer that is cheap and shallow, and the generalized ChiChi families offer that profile at every even dimension, so designers can be expected to adopt them for the same purpose. Settling the five conjectures therefore serves everyone who wants to build ciphers on the generalized ChiChi families.

Our contributions. In this paper we settle all five conjectures of [ALPS25]. Our results have five main components.

1. We prove Conjectures 1–4 of [ALPS25] for all admissible odd branch lengths $m, m' \geq 3$. For the inverse SWAP family, the linear height in every nonzero direction and the linear nullity of every nonzero component are at most $\max\{m, m'\} + 1$. This proves Conjectures 2 and 1. For the inverse 4-cycle family, the linear height is at most $(\max\{m, m'\} + 1)/2$, which proves Conjecture 4. A complementary fixed-mask analysis gives $\dim E_\lambda \leq 2$ for every nonzero λ . That proves Conjecture 3, and shows that every nonzero inverse component has at most four linear structures. The exceptional (3, 3) SWAP case and the finite induction bases for the 4-cycle classifications come from exact exhaustive verification. Symbolic two-coordinate recurrences carry them to all lengths.
2. We disprove Conjecture 5 by exact exhaustive evaluation. The parameter pairs $(m, m') = (5, 7), (7, 9), (9, 7)$ have linear height greater than the proposed bound of 3, whereas (7, 5) attains but does not exceed it. We also correct the reported data of [ALPS25]: the value listed there for (5, 7) is 3, whereas the correct value is 4, as the authors have confirmed (Remark 2). More generally, Corollary 2 gives the sharp all-length bound for $H = (\chi_{m, m'}^{(0,1)})^{-1}$: for every $\alpha \neq 0$, $\dim K_H(\alpha) \leq (m' + 1)/2$ when $m' \geq 5$, and $\dim K_H(\alpha) \leq (m + 1)/2$ when $m' = 3$; equality is attained in each case. Thus no absolute constant can replace the conjectured constant, and the governing branch is orientation-dependent.
3. We derive design consequences for the deployed ChiChi instances. For the (0, 1) variant the exact maximum $(m' + 1)/2$ depends only on m' , so it distinguishes the two branch orderings of a given even dimension, and the unbalanced splits with $m' = 5$ pin the maximum at 3 for every even $N \geq 8$; an exhaustive sweep at $N = 32$ confirms the resulting predictions (Section 9).
4. Two reductions organize all four proofs. Fix an input difference α and let S_α be the set of attainable output differences. Write D_α for its affine hull and $\text{dir}(D_\alpha)$ for the direction space of that hull. We show that $K_H(\alpha)$ is the orthogonal complement $\text{dir}(D_\alpha)^\perp$, so an upper bound on linear height is a lower bound on $\dim D_\alpha$. Each family is two ordinary χ cycles altered at a few coordinates, two of them in the SWAP case and four in the 4-cycle case, and fixing those coordinates leaves two independent ordinary- χ systems. Everything then rests on which differences a single cycle attains, decided by a local test on the runs of zeros and ones. The bounds come from constructing such differences. With a difference fixed, enough of them make D_α large; with a mask fixed, each one is a direction the mask has to annihilate, and enough of them leave no nonzero mask standing.
5. Of independent interest, we settle the single-cycle benchmark exactly (Section 10). For every odd $n \geq 3$, the inverse map χ_n^{-1} has linear height at most one in every nonzero direction, and linear nullity at most one in every nonzero component. Both equality loci are classified. This is the base case against which [ALPS25] measures both families. There it is recorded numerically, for small dimensions only.

Outline. The remainder of the paper is organized as follows. Section 2 fixes the notation, defines the two ChiChi families, and states the five conjectures. Section 3 sets out the two reductions the proofs rest on and supplies the local analysis of one ordinary χ cycle that they lead to. Sections 4 and 5 prove Conjectures 2 and 1 for the SWAP family, and Sections 6 and 7 prove Conjectures 4 and 3 for the 4-cycle family. Section 8 gives the counterexamples to Conjecture 5 and identifies Conjecture 4 as the corrected length-dependent statement, and Section 9 draws the design consequences for the deployed

ChiLow instances. Section 10 settles the single ordinary cycle itself. Section 11 concludes. The appendices contain the finite-instance verification, the zero-internal fixed-boundary attainable sets, the two-coordinate recurrences, the boundary classifications used in the proofs, and the direction constructions for the SWAP family.

2 Preliminaries

This section fixes the notation used throughout the paper, defines the two generalized ChiChi families, and states the five conjectures of [ALPS25] in that notation.

2.1 Notation

The state. All algebraic operations are over $\mathbb{F}_2$, whose elements are denoted by 0 and 1; vector addition is componentwise. Let $m, m' \geq 3$ be odd and set $N = m + m'$. We write $x = (x_0, \dots, x_{m-1}) \in \mathbb{F}_2^m$ and $\hat{x} = (\hat{x}_0, \dots, \hat{x}_{m'-1}) \in \mathbb{F}_2^{m'}$, with unhatted and hatted indices reduced modulo m and m' , respectively. For $a, b \in \mathbb{F}_2^r$, their standard inner product is $\langle a, b \rangle = \sum_{i=0}^{r-1} a_i b_i \in \mathbb{F}_2$.

Differences. Let F be a member of either of the two generalized ChiChi families, defined in Section 2.2, and put $H = F^{-1}$. An input difference of H is

$$\alpha = (\alpha_L, \alpha_R) = (\alpha_0, \dots, \alpha_{m-1}, \hat{\alpha}_0, \dots, \hat{\alpha}_{m'-1}) \in \mathbb{F}_2^N.$$

Its internal and truncated internal parts are

$$\begin{aligned} \alpha_L^{\text{int}} &= (\alpha_1, \dots, \alpha_{m-1}), & \alpha_R^{\text{int}} &= (\hat{\alpha}_1, \dots, \hat{\alpha}_{m'-1}), \\ \eta_L &= (\alpha_2, \dots, \alpha_{m-1}), & \eta_R &= (\hat{\alpha}_2, \dots, \hat{\alpha}_{m'-1}). \end{aligned}$$

For example, $\alpha_L^{\text{int}} \neq 0$ means that some α_i , $1 \leq i < m$, is nonzero. Two truncations are needed because the two families couple their branches at different numbers of coordinates: the SWAP family touches one coordinate per branch, so what is left free there is α_L^{int} , while the 4-cycle family touches two, leaving η_L .

For $y \in \mathbb{F}_2^N$, let $d = H(y + \alpha) + H(y)$ denote the corresponding output difference, written as $d = (d_0, \dots, d_{m-1}, \hat{d}_0, \dots, \hat{d}_{m'-1})$. Throughout the paper, $\alpha_i, \hat{\alpha}_i$ denote input-difference coordinates of H , while $d_i, \hat{d}_i$ denote its output-difference coordinates. When the same constraint is read on a forward branch, these roles are reversed:

quantity	viewpoint of H	viewpoint of a forward branch
α	input difference	output difference
d	output difference	input difference

The four spaces. An output mask of H is $\lambda = (\lambda_0, \dots, \lambda_{m-1}, \hat{\lambda}_0, \dots, \hat{\lambda}_{m'-1}) \in \mathbb{F}_2^N$, and the inverse differential-linear bias is $\text{Bias}(\alpha, \lambda) = 2^{-N} \sum_{y \in \mathbb{F}_2^N} (-1)^{\langle \lambda, H(y + \alpha) + H(y) \rangle}$. For fixed α , let $S_\alpha = \{H(y + \alpha) + H(y) : y \in \mathbb{F}_2^N\}$ and $D_\alpha = \text{Aff}(S_\alpha)$, where $\text{Aff}(S_\alpha)$ is called the affine hull of S_α , that is, the smallest affine subspace of $\mathbb{F}_2^N$ containing it. For fixed λ , let $E_\lambda = \{\alpha \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\}$; this is the set of input differences for which $H_\lambda = \langle \lambda, H \rangle$ has a constant derivative. For fixed α , let $K_H(\alpha) = \{\lambda \in \mathbb{F}_2^N : \langle \lambda, d \rangle \text{ is constant for } d \in S_\alpha\}$. The notation E_λ fixes λ and varies α , whereas $K_H(\alpha)$ fixes α and varies λ . Both are needed because the conjectures come in pairs: one of each pair bounds $\dim E_\lambda$ over all masks, the other bounds $\dim K_H(\alpha)$ over all differences, and neither bound follows from the other. For an affine space $A = v + V$ with V a linear

subspace, its *direction space* is $\text{dir}(A) = V$; concretely, $\text{dir}(\text{Aff}(T)) = \text{span}\{t + t^{(0)} : t \in T\}$ for any fixed $t^{(0)} \in T$, independently of that choice. $\text{span}(T)$, $T^\perp$, and the standard unit vectors e_i have their usual meanings. A mask λ is *full-bias* for α when $\text{Bias}(\alpha, \lambda) = \pm 1$.

Projections. The branch projections are $\pi_L : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^m$, $\pi_L(d) = (d_0, \dots, d_{m-1})$, and $\pi_R : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^{m'}$, $\pi_R(d) = (\hat{d}_0, \dots, \hat{d}_{m'-1})$. Their internal counterparts are $\pi_L^{\text{int}}(d) = (d_1, \dots, d_{m-1})$ and $\pi_R^{\text{int}}(d) = (\hat{d}_1, \dots, \hat{d}_{m'-1})$. For $A \subseteq \mathbb{F}_2^N$, projections act pointwise, for example $\pi_L(A) = \{\pi_L(v) : v \in A\}$. They are linear, so $\text{dir}(\pi_L(D_\alpha)) = \pi_L(\text{dir}(D_\alpha))$, and analogously for π_R , π_L^{int} , and π_R^{int} .

2.2 The SWAP and 4-Cycle Families

The χ -function [Dae95] is

$$(\chi_r(z))_i = z_i + (z_{i+1} + 1)z_{i+2} \quad (0 \leq i < r),$$

where all indices are taken modulo r .

The SWAP family [ALPS25] consists of the mappings $\mathbb{X}_{m,m'}^{(\text{swap})} : \mathbb{F}_2^m \times \mathbb{F}_2^{m'} \rightarrow \mathbb{F}_2^m \times \mathbb{F}_2^{m'}$ defined by

$$\begin{aligned} \mathbb{X}_{m,m'}^{(\text{swap})}(x, \hat{x}) &= (y, \hat{y}), \\ y_i &= \begin{cases} \hat{x}_0 + (x_1 + 1)x_2, & i = 0, \\ x_i + (x_{i+1} + 1)x_{i+2}, & i \neq 0, \end{cases} \\ \hat{y}_i &= \begin{cases} x_0 + (\hat{x}_1 + 1)\hat{x}_2, & i = 0, \\ \hat{x}_i + (\hat{x}_{i+1} + 1)\hat{x}_{i+2}, & i \neq 0. \end{cases} \end{aligned}$$

The construction consists of two χ -functions whose zeroth input coordinates are exchanged.

We next define the 4-cycle family. Let $b, \hat{b} \in \mathbb{F}_2$. The family [ALPS25] consists of the mappings $\mathbb{X}_{m,m'}^{(b,\hat{b})} : \mathbb{F}_2^m \times \mathbb{F}_2^{m'} \rightarrow \mathbb{F}_2^m \times \mathbb{F}_2^{m'}$ defined by

$$\begin{aligned} \mathbb{X}_{m,m'}^{(b,\hat{b})}(x, \hat{x}) &= (y, \hat{y}), \\ y_i &= \begin{cases} \hat{x}_1 + (x_1 + 1)(x_2 + b), & i = 0, \\ \hat{x}_0 + (x_2 + 1)x_3, & i = 1, \\ x_i + (x_{i+1} + 1)x_{i+2}, & \text{otherwise,} \end{cases} \\ \hat{y}_i &= \begin{cases} x_0 + (\hat{x}_1 + 1)(\hat{x}_2 + \hat{b}), & i = 0, \\ x_1 + (\hat{x}_2 + 1)\hat{x}_3, & i = 1, \\ \hat{x}_i + (\hat{x}_{i+1} + 1)\hat{x}_{i+2}, & \text{otherwise.} \end{cases} \end{aligned}$$

Equivalently, this construction is obtained from two χ -functions by the coordinate 4-cycle $(x_0, \hat{x}_1, x_1, \hat{x}_0)$. The parameters b and $\hat{b}$ are applied at the two modified coordinates.

Both families consist of permutations: the SWAP map for all odd $m, m' \geq 3$ [ALPS25, Theorem 1], and the 4-cycle map for all odd $m, m' \geq 3$ with $b\hat{b} = 0$ [ALPS25, Theorem 3]. The excluded pair $b = \hat{b} = 1$ does not give a permutation, which is why the admissible parameters are restricted throughout. The inverse $H = F^{-1}$ used in this paper is therefore well defined.

2.3 Linear Structures and Related Conjectures

For a Boolean function $f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ and a direction $a \in \mathbb{F}_2^n$, define

$$\Delta_a f(x) = f(x + a) + f(x).$$

The direction a is a *linear structure* of f when $\Delta_a f$ is constant. The linear structures form the set

$$\text{LS}(f) = \{a \in \mathbb{F}_2^n : \Delta_a f \text{ is constant}\},$$

and $\dim \text{LS}(f)$ is the *linear nullity* of f . It is classical that $\text{LS}(f)$ is a vector subspace of $\mathbb{F}_2^n$ [Lai95, Car21], so that $\#\text{LS}(f) = 2^{\dim \text{LS}(f)}$.

Let $G : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^s$ be vectorial and let $G_v = \langle v, G \rangle$ be its component selected by $v \in \mathbb{F}_2^s$. For a fixed direction a , the partial linear kernel is

$$K_G(a) = \{v \in \mathbb{F}_2^s : \Delta_a G_v \text{ is constant}\}.$$

Its dimension is the *linear height* of G in direction a . For $G = H$ these two spaces are exactly the E_λ and $K_H(\alpha)$ defined above, so linear nullity and linear height are the dimensions of those. The terminology is that of [ALPS25].

We now state the five conjectures of [ALPS25] in the notation of Section 2.1.

Conjecture 1 (Linear nullity of the inverse SWAP family). *Let $H = (\mathbb{X}_{m,m'}^{(\text{swap})})^{-1}$. For every nonzero mask λ ,*

$$\dim E_\lambda \leq \max\{m, m'\} + 1.$$

Conjecture 2 (Linear height of the inverse SWAP family). *Let $H = (\mathbb{X}_{m,m'}^{(\text{swap})})^{-1}$. For every nonzero input difference α ,*

$$\dim K_H(\alpha) \leq \max\{m, m'\} + 1.$$

Conjecture 3 (Linear nullity of the inverse 4-cycle family). *Let $b, \hat{b} \in \mathbb{F}_2$ satisfy $b\hat{b} = 0$, and let $H = (\mathbb{X}_{m,m'}^{(b,\hat{b})})^{-1}$. For every nonzero mask λ ,*

$$\dim E_\lambda \leq 2.$$

Conjecture 4 (Linear height of the inverse 4-cycle family). *Let $b, \hat{b} \in \mathbb{F}_2$ satisfy $b\hat{b} = 0$, and let $H = (\mathbb{X}_{m,m'}^{(b,\hat{b})})^{-1}$. For every nonzero input difference α ,*

$$\dim K_H(\alpha) \leq \frac{\max\{m, m'\} + 1}{2}.$$

Conjecture 5 (The $(0, 1)$ 4-cycle family). *Let $m, m' \geq 3$ be odd and satisfy $|m - m'| \leq 2$, and put $H = (\mathbb{X}_{m,m'}^{(0,1)})^{-1}$. For every nonzero input difference α ,*

$$\dim K_H(\alpha) \leq 3.$$

3 Proof Strategy and Single-Cycle Properties

Each of the five conjectures bounds $\dim K_H(\alpha)$ or $\dim E_\lambda$ for $H = F^{-1}$. Attacking H directly is unattractive, because H is neither local nor of low degree. We avoid it with two reductions.

The first is a duality, and it is general. For a fixed input difference α , the full-bias masks are the annihilator of the direction space of $D_\alpha = \text{Aff}(S_\alpha)$, so an upper bound on linear height is the same thing as a lower bound on an affine rank, and a counting problem becomes a rank problem. We prove it once, in Subsection 3.1, for both families.

The second replaces H by the forward map, and it is specific to each family. A difference d lies in S_α exactly when F carries the input difference d to the output difference α , and F is two ordinary χ cycles altered at two coordinates in the SWAP case and at four in the 4-cycle case. Fixing those few coordinates leaves two independent ordinary- χ

difference systems, each with its first coordinates prescribed. Carrying that out needs the construction itself in hand, so we do it where each family is treated, in Lemma 8 for the SWAP family and Lemma 19 for the 4-cycle family.

Throughout the section, $r \geq 3$ is odd and $H_r = \chi_r^{-1}$, which is well defined since χ_r is a bijection for odd r . The spaces S_α , $D_\alpha = \text{Aff}(S_\alpha)$ and $K_{H_r}(\alpha)$ are those of Section 2.1 with H_r in place of H and cycle length r in place of N ; the proof of Lemma 1 uses only the bijectivity of the underlying map, so

$$K_{H_r}(\alpha) = \text{dir}(D_\alpha)^\perp \quad (1)$$

holds verbatim. All indices are reduced modulo r . A *cyclic interval* $[p, p']$ is the set of ℓ consecutive indices $p, p+1, \dots, p+\ell-1$ read modulo r , where ℓ is the residue of the integer $p' - p + 1$ in $\{0, 1, \dots, r-1\}$; it is empty when $\ell = 0$. Reading the length off the endpoint expressions, rather than off their residues, is what distinguishes an interval such as $[q, q-1]$, which is empty, from one such as $[j+3, j-1]$, which carries $r-3$ indices. The vector $\mathbf{1} \in \mathbb{F}_2^r$ is the all-ones vector. The weight $\text{wt}(\alpha)$ is the number of nonzero coordinates of α .

3.1 The Affine-Hull Duality

The masks counted by $K_H(\alpha)$ are exactly those on which the functional $\langle \lambda, \cdot \rangle$ is constant over S_α , and that is a statement about the affine hull of S_α .

Lemma 1 (Full-bias masks and affine hulls). *For every fixed α , the full-bias masks form the vector space*

$$K_H(\alpha) = \text{dir}(D_\alpha)^\perp, \quad \dim K_H(\alpha) = N - \dim D_\alpha. \quad (2)$$

In particular,

$$\#\{\lambda \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\} = 2^{N - \dim D_\alpha}.$$

Proof. For each $y \in \mathbb{F}_2^N$, put $\varepsilon_y = (-1)^{\langle \lambda, H(y+\alpha) + H(y) \rangle}$. The normalized sum of these signs has absolute value one exactly when all signs agree. Therefore $\text{Bias}(\alpha, \lambda) = \pm 1$ exactly when $\langle \lambda, d \rangle$ is constant on S_α .

Choose $d^{(0)} \in S_\alpha$. Constancy is equivalent to $\langle \lambda, d + d^{(0)} \rangle = 0$ for every $d \in S_\alpha$, while

$$\text{dir}(D_\alpha) = \text{span}\{d + d^{(0)} : d \in S_\alpha\}.$$

The full-bias masks are the orthogonal complement in (2). The dimension and cardinality statements follow from orthogonal-complement dimension over $\mathbb{F}_2$. $\square$

In every proof of this paper we take the same first step, splitting ChiChi into two independent ordinary χ cycles, and then work entirely with a single cycle. The subsections below collect the properties of one cycle that our proofs will frequently use.

3.2 The Local Membership Test

The set S_α is defined by ranging over all inputs, so deciding whether a particular d belongs to it means producing one. We obtain every bound by exhibiting elements of S_α , so that decision has to be cheap. Lemma 2 moves the question to the forward map χ_r , whose equations are local, and Theorem 1 then removes the input altogether: whether d lies in S_α is read off the runs of zeros and ones in d , together with α .

Lemma 2. *For every $\alpha \in \mathbb{F}_2^r$,*

$$S_\alpha = \{d \in \mathbb{F}_2^r : \chi_r(y) + \chi_r(y + d) = \alpha \text{ for some } y \in \mathbb{F}_2^r\}.$$

Proof. If $d = H_r(x) + H_r(x + \alpha)$, put $y = H_r(x)$; then $y + d = H_r(x + \alpha)$ and $\chi_r(y) + \chi_r(y + d) = x + (x + \alpha) = \alpha$. Conversely, if $\chi_r(y) + \chi_r(y + d) = \alpha$, put $x = \chi_r(y)$; then $H_r(x) = y$, $H_r(x + \alpha) = y + d$, and $d \in S_\alpha$. $\square$

Theorem 1 (Local test). *Let $\alpha \in \mathbb{F}_2^r$. Then $0 \in S_\alpha$ if and only if $\alpha = 0$, and $\mathbf{1} \in S_\alpha$ if and only if $\text{wt}(\alpha)$ is odd. For $d \notin \{0, \mathbf{1}\}$, the membership $d \in S_\alpha$ holds if and only if both of the following are satisfied:*

- (i) $d_i = \alpha_i$ whenever $d_{i+1} = d_{i+2} = 0$;
- (ii) $d_{q-2} = \alpha_{q-2} + \alpha_{q-1}$ at every isolated zero q of d .

Proof. By Lemma 2, $d \in S_\alpha$ means that $\chi_r(y) + \chi_r(y + d) = \alpha$ for some y . Expand $\chi_r(y + d)_i = (y_i + d_i) + (y_{i+1} + d_{i+1} + 1)(y_{i+2} + d_{i+2})$ and add $\chi_r(y)_i$. The terms quadratic in y cancel, so the condition is the solvability of the linear system $L_d(y) = w$, where

$$L_d(y)_i = d_{i+2}y_{i+1} + d_{i+1}y_{i+2}, \quad c_{d,i} = d_i + d_{i+2} + d_{i+1}d_{i+2}, \quad w_i = \alpha_i + c_{d,i}.$$

Re-indexing the i -th equation by $k = i + 1$, it reads

$$d_{k+1}y_k + d_k y_{k+1} = w_{k-1}, \quad k \in \mathbb{Z}_r, \quad (3)$$

and its shape at k depends only on (d_k, d_{k+1}) : the pair $(0, 0)$ yields the condition $w_{k-1} = 0$; the pair $(0, 1)$ pins $y_k = w_{k-1}$; the pair $(1, 0)$ pins $y_{k+1} = w_{k-1}$; the pair $(1, 1)$ chains $y_k + y_{k+1} = w_{k-1}$.

Let $d \notin \{0, \mathbf{1}\}$, so that d has both maximal one-runs and maximal zero-runs. Inside a one-run $[p, p + \ell - 1]$, the $(1, 1)$ equations chain $y_p, \dots, y_{p+\ell-1}$, leaving one free parameter and no condition. At $k = p - 1$ the pair is $(0, 1)$ and pins y_{p-1} ; at $k = p + \ell - 1$ the pair is $(1, 0)$ and pins $y_{p+\ell}$. Hence each zero-run has its first coordinate pinned by the preceding one-run and its last coordinate pinned by the following one-run; these coincide exactly for a zero-run of length one. The system is therefore solvable if and only if every $(0, 0)$ condition holds and the two pinnings agree at every isolated zero.

The $(0, 0)$ equations occur at the k with $d_k = d_{k+1} = 0$; writing $i = k - 1$ and using $d_{i+1} = d_{i+2} = 0$ in c_d gives $w_i = \alpha_i + d_i$, so the condition is $d_i = \alpha_i$: this is the first assertion. At an isolated zero q , the pinnings are $y_q = w_{q-2}$ and $y_q = w_{q-1}$; here $d_q = 0$ and $d_{q-1} = d_{q+1} = 1$ give $w_{q-2} = \alpha_{q-2} + d_{q-2}$ and $w_{q-1} = \alpha_{q-1}$, so agreement is $d_{q-2} = \alpha_{q-2} + \alpha_{q-1}$: this is the second assertion.

For $d = 0$, every equation reads $0 = \alpha_{k-1}$, so $\alpha = 0$; and $0 \in S_0$ since $H_r(y) + H_r(y) = 0$. For $d = \mathbf{1}$, every pair is $(1, 1)$ and (3) is a single cyclic chain, solvable exactly when $\sum_k w_{k-1} = 0$; here $w_i = \alpha_i + 1$, so the condition is $\text{wt}(\alpha) + r \equiv 0 \pmod{2}$, that is, $\text{wt}(\alpha)$ is odd, since r is odd. $\square$

3.3 Differences with a Single Zero-Run

The test recognises attainable differences, while the constructions later on have to produce them. The differences whose zero set is a single cyclic interval are the ones that can be produced, because for them the two conditions collapse into a statement about α alone.

For $q \in \mathbb{Z}_r$ and $0 \leq g \leq r - 1$, let $I_{q,g} \in \mathbb{F}_2^r$ denote the indicator of the cyclic interval $[q, q + g - 1]$, which has g indices, and consider the attainable differences of the form $d = \mathbf{1} + I_{q,g}$, the vectors whose zero set is one cyclic interval.

Lemma 3. $\mathbf{1} + I_{q,g} \in S_\alpha$ if and only if one of the following holds:

1. $g = 0$ and $\text{wt}(\alpha)$ is odd;
2. $g = 1$ and $\alpha_{q-2} \neq \alpha_{q-1}$;

3. $2 \leq g \leq r-1$, $\alpha_{q-1} = 1$, and $\alpha_i = 0$ for all $i \in [q, q+g-3]$.

Proof. Immediate from Theorem 1, since such a d has the single zero-run $[q, q+g-1]$. For $g = 1$, that zero is isolated and condition (ii) applies with $d_{q-2} = 1$. For $g \geq 2$, condition (i) applies at the $i \in [q-1, q+g-3]$; at $i = q-1$ one has $d_{q-1} = 1$, giving $\alpha_{q-1} = 1$, and at $i \in [q, q+g-3]$ one has $d_i = 0$, giving $\alpha_i = 0$. $\square$

3.4 Boundary Slices of Attainable Differences

Fixing the coordinates where the two branches interact prescribes the first two coordinates of each ordinary cycle. The corresponding subset of attainable differences is a *fixed-boundary attainable set*. Two things about such a set matter downstream, which boundary values are possible and how much internal rank it carries, and both answers split on whether the branch has a nonzero internal input difference. A rigid branch, one without such a difference, admits only eight of the sixteen boundary quadruples, which constrains the global boundary combinations, but its sets span all $r-2$ internal directions. A flexible branch admits every quadruple and so constrains nothing, and its sets still span at least $r-3$ directions. The four statements are Lemma 4 and (6) for the rigid case, and Lemma 6 and Lemma 5 for the flexible one.

Let $r \geq 3$ be odd. Written out coordinatewise, the difference equation behind Theorem 1 reads

$$\alpha_i = d_i + d_{i+2} + d_{i+1}d_{i+2} + d_{i+1}x_{i+2} + d_{i+2}x_{i+1}, \quad 0 \leq i < r, \quad (4)$$

where x is a preimage and α the output difference that the input difference d produces at it. Theorem 1 asked which d solve this for a given α . The family analyses need the same question asked with the first two coordinates of d , and the first two of α , held fixed. For fixed $(d_0, d_1, \alpha_0, \dots, \alpha_{r-1}) \in \mathbb{F}_2^{r+2}$, the corresponding fixed-boundary attainable set is

$$\begin{aligned} & \mathcal{C}_r(\alpha_2, \dots, \alpha_{r-1}; d_0, d_1, \alpha_0, \alpha_1) \\ &= \left\{ d \in \mathbb{F}_2^r : \begin{array}{l} \text{there exists } x \in \mathbb{F}_2^r \text{ such that} \\ \Delta_d \chi_r(x) = (\alpha_0, \dots, \alpha_{r-1}), \\ \text{the first two coordinates of } d \text{ equal the fixed values } d_0, d_1 \end{array} \right\}. \end{aligned} \quad (5)$$

Four pieces of notation are needed to state the results, and we collect them here. All differences in one such set share their first two coordinates, so their pairwise sums lie in

$$I_r = \{v \in \mathbb{F}_2^r : v_0 = v_1 = 0\} = \langle e_2, \dots, e_{r-1} \rangle,$$

and every direction space below is a subspace of I_r . Second, such a set need not be linear, so we record an anchor for its affine hull and the corresponding direction space. Call a boundary quadruple $q = (d_0, d_1, \alpha_0, \alpha_1) \in \mathbb{F}_2^4$ *supported* when $\mathcal{C}_r(0, \dots, 0; q) \neq \emptyset$, and for supported q write

$$c_r(q) \in \text{Aff } \mathcal{C}_r(0, \dots, 0; q), \quad U_r(q) = \text{dir Aff } \mathcal{C}_r(0, \dots, 0; q).$$

Any element of the hull serves as an anchor, so $c_r(q)$ is a choice; $U_r(q)$ is not, since translating a set does not change its directions. Third, boundary quadruples are indexed by

$$\text{enc}(d_0, d_1, \alpha_0, \alpha_1) = d_0 + 2d_1 + 4\alpha_0 + 8\alpha_1.$$

Fourth, for a set $\mathcal{B} \subseteq \mathbb{F}_2^4$ of quadruples and fixed internal coordinates $\alpha_2, \dots, \alpha_{r-1}$, write

$$W_r(\alpha_2, \dots, \alpha_{r-1}; \mathcal{B}) = \sum_{(d_0, d_1, \alpha_0, \alpha_1) \in \mathcal{B}} \text{dir Aff } \mathcal{C}_r(\alpha_2, \dots, \alpha_{r-1}; d_0, d_1, \alpha_0, \alpha_1)$$

for the directions that $\mathcal{B}$ contributes together. A single set need not span much on its own, so the rank estimate below asks that $\mathcal{B}$ offer several sets over one boundary output: say

that $\mathcal{B}$ has the *three-input-difference property* if, for some fixed output difference (α_0, α_1) , at least three of the four input differences $(d_0, d_1) \in \mathbb{F}_2^2$ give quadruples in $\mathcal{B}$. Three of the four suffice, and three is what the boundary maps of both families deliver.

We now treat the rigid case, in which the internal input difference is zero.

Lemma 4 (Exact attainable sets for zero internal difference). *For every odd $r \geq 5$, the nonempty sets*

$$\mathcal{C}_r(0, \dots, 0; d_0, d_1, \alpha_0, \alpha_1)$$

are exactly the eight sets in Table 4, and the table gives their exact anchor–direction decompositions. When $r = 3$, the last row is absent and the direction space in every other nonzero row is $\langle e_2 \rangle$. Writing B_r for the set of encodings of the supported quadruples,

$$B_3 = \{0, 5, 7, 10, 11, 13, 14\}, \quad B_r = B_3 \cup \{15\} \quad (r \geq 5).$$

Proof. For fixed d , the cyclic difference equations form a linear system in x . Backward elimination gives the stated affine upper bounds and excludes the eight impossible boundary types, whereas explicit unit upper-triangular witness families attain every bound. The complete derivation is given in Appendix B. $\square$

So B_r is the precise sense in which a branch with no internal input difference is rigid: those are the only boundary types it admits, and the family analyses use exactly this to cut down the feasible boundary profiles. Both B_3 and B_r have affine dimension four in $\mathbb{F}_2^4$, which is what Section 6 needs of them. One entry already carries the whole internal rank: the type $\text{enc}(1, 0, 1, 0) = 5$, the row 10/10 of Table 4, has direction space I_r , the maximum possible internal direction space for such a set. Hence

$$\dim \sum_{\text{enc}(d_0, d_1, \alpha_0, \alpha_1) \in B_r} \text{dir Aff } \mathcal{C}_r(0, \dots, 0; d_0, d_1, \alpha_0, \alpha_1) = r - 2. \quad (6)$$

In the flexible case the internal input difference is nonzero. The boundary is then unconstrained, so what has to be measured is rank.

Lemma 5 (Rank supplied by a nonzero internal input difference). *Let $r \geq 3$ be odd. If $(\alpha_2, \dots, \alpha_{r-1}) \neq 0$ and $\mathcal{B}$ has the three-input-difference property, then*

$$\dim W_r(\alpha_2, \dots, \alpha_{r-1}; \mathcal{B}) \geq r - 3. \quad (7)$$

Lemma 6 (Full boundary support for nonzero internal differences). *Let $r \geq 3$ be odd and let $\eta = (\alpha_2, \dots, \alpha_{r-1}) \neq 0$. Then*

$$\mathcal{C}_r(\eta; d_0, d_1, \alpha_0, \alpha_1) \neq \emptyset \quad \text{for every } (d_0, d_1, \alpha_0, \alpha_1) \in \mathbb{F}_2^4.$$

Both proofs are in the appendices. That of Lemma 5 is in Appendix C, which resolves the three retained sets state by state; the bound above is the crude consequence, and Table 6 there records the exact annihilator for each boundary profile. That of Lemma 6 is a construction, one explicit difference for each of the four prefixes (d_0, d_1) checked against the two conditions of Theorem 1, and is in Appendix B.

4 Proof of Conjecture 2

The four proofs of Conjectures 1–4 come in two pairs, one per family, and within each pair the height bound comes first. Note that it is not the order of the conjectures, but it is the order of the argument. By Lemma 1 a height bound asks us to exhibit enough attainable differences, whereas a nullity bound fixes a mask and has to rule directions out, which needs a supply of attainable differences already in hand.

As the proof strategy in Section 3 shows, we prove Conjecture 2 by translating linear height into an affine-rank problem. After reducing the attainable differences to the two ordinary χ branches, we construct enough independent directions on whichever branch carries a nonzero internal difference.

Let $F = \mathbb{X}_{m,m'}^{(\text{swap})}$ and $H = F^{-1}$. Fix a nonzero input difference α of H . Conjecture 2 asks for an upper bound on the dimension of the mask space $K_H(\alpha)$. By Lemma 1, its cardinality is

$$\#\{\lambda \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\} = 2^{\dim K_H(\alpha)} = 2^{N - \dim D_\alpha}. \quad (8)$$

The required linear-height bound is equivalent to

$$\dim D_\alpha \geq N - \max\{m, m'\} - 1 = \min\{m, m'\} - 1. \quad (9)$$

So it suffices to exhibit $\min\{m, m'\} - 1$ independent directions in D_α , and one branch is enough to supply them. The argument runs in three steps, one per subsection below. The two branches interact only through the pair $(d_0, \hat{d}_0)$, so fixing that pair decouples them. Whenever $\alpha_L \neq 0$ the coupling bit $\hat{d}_0$ can be forced to 1. Once it is, the right branch ranges over the attainable set of a single ordinary cycle, whose internal projection is all of $\mathbb{F}_2^{m'-1}$; this gives $\dim D_\alpha \geq m' - 1$. If $\alpha_L = 0$ then $\alpha_R \neq 0$ and the mirror argument gives $\dim D_\alpha \geq m - 1$. Either estimate implies (9).

4.1 Decoupling the Two Branches

We use the defining equations of the SWAP family. Write

$$x = (x_0, \dots, x_{m-1}, \hat{x}_0, \dots, \hat{x}_{m'-1}) \in \mathbb{F}_2^N, \quad y = (y_0, \dots, y_{m-1}, \hat{y}_0, \dots, \hat{y}_{m'-1}),$$

and let $F(x) = y$, where

$$y_0 = \hat{x}_0 + (x_1 + 1)x_2, \quad y_i = x_i + (x_{i+1} + 1)x_{i+2} \quad (1 \leq i \leq m-1), \quad (10)$$

and

$$\hat{y}_j = x_0 + (\hat{x}_1 + 1)\hat{x}_2, \quad \hat{y}_j = \hat{x}_j + (\hat{x}_{j+1} + 1)\hat{x}_{j+2} \quad (1 \leq j \leq m'-1). \quad (11)$$

The indices in (10) and (11) are reduced modulo m and m' on the left and right cycles, respectively. Set $H = F^{-1}$.

The local difference equations mention both preimages at once, so both need names. For $y \in \mathbb{F}_2^N$ put $x = H(y)$ and $x_d = H(y + \alpha) = x + d$, with coordinates

$$x_d = (x_{d,0}, \dots, x_{d,m-1}, \hat{x}_{d,0}, \dots, \hat{x}_{d,m'-1}),$$

so that $d_i = x_{d,i} + x_i$ and $\hat{d}_i = \hat{x}_{d,i} + \hat{x}_i$. The subscript d marks the translated preimage and is not a coordinate index.

Lemma 7. *Let α be an input difference of H , and let $d \in S_\alpha$ be an attainable output difference. On either ordinary χ -cycle, x_i and $x_{d,i}$, or $\hat{x}_i$ and $\hat{x}_{d,i}$ on the right branch, are the two preimage coordinates occurring in the local equations; they are not the external coordinates y_i and $y_i + \alpha_i$. Then*

$$\alpha_i = d_i + d_{i+2} + d_{i+1}d_{i+2} + d_{i+1}x_{i+2} + d_{i+2}x_{i+1}. \quad (12)$$

At the left branch connection,

$$\hat{d}_0 = \alpha_0 + d_2x_1 + d_1x_2 + d_1d_2 + d_2. \quad (13)$$

Proof. Following the standard XOR-difference formalism for Boolean mappings [Dae95], add the two local χ -equations at x and $x_d = x + d$ over $\mathbb{F}_2$. This gives (12). Adding the two instances of the first equation of the SWAP representative gives (13). $\square$

Lemma 8 (Branch decoupling). *Let $c, e \in \mathbb{F}_2$. There is an attainable difference $d \in S_\alpha$ with $d_0 = c$ and $\hat{d}_0 = e$ if and only if both*

- *some $d_L \in \mathbb{F}_2^m$ with $(d_L)_0 = c$ is an attainable difference of the ordinary cycle χ_m for the output difference $\alpha_L + (c + e)e_0$, and*
- *some $d_R \in \mathbb{F}_2^{m'}$ with $(d_R)_0 = e$ is an attainable difference of the ordinary cycle $\chi_{m'}$ for the output difference $\alpha_R + (c + e)e_0$.*

Proof. The SWAP representative differs from two independent ordinary cycles only in the two coordinates y_0 and $\hat{y}_0$, so the difference equations of index $1 \leq i \leq m - 1$ are exactly the ordinary χ_m equations in (x_L, d_L) , and likewise on the right. At index 0, adding the two instances of $y_0 = \hat{x}_0 + (x_1 + 1)x_2$ gives (13), that is $\alpha_0 = \hat{d}_0 + g$ with $g = d_2x_1 + d_1x_2 + d_1d_2 + d_2$, whereas the ordinary cycle χ_m evaluated on the same pair (x_L, d_L) produces the 0-coordinate $d_0 + g$. Hence the ordinary output difference produced by (x_L, d_L) is

$$(d_0 + g, \alpha_1, \dots, \alpha_{m-1}) = \alpha_L + (d_0 + \hat{d}_0)e_0,$$

since $g = \alpha_0 + \hat{d}_0$. The same computation on the right branch gives $\alpha_R + (\hat{d}_0 + d_0)e_0$. Conversely, a pair of ordinary differentials with the stated output differences and 0-coordinates c and e satisfies every SWAP difference equation when concatenated, because the two displayed identities are the only ones in which the branches interact. $\square$

4.2 Forcing $\hat{d}_0 = 1$

Decoupling leaves the coupling pair $(d_0, \hat{d}_0)$ to be chosen. What values it can take is decided branch by branch, so we first record which 0-coordinates a single cycle admits.

Lemma 9 (Attainable 0-coordinates on one cycle). *Let $r \geq 3$ be odd and $\beta \in \mathbb{F}_2^r$, and let $P(\beta) \subseteq \mathbb{F}_2$ be the set of values taken by d_0 over the attainable differences d of χ_r for the output difference β . Then*

$$P(\beta) = \begin{cases} \{0\}, & \beta = 0, \\ \{1\}, & \beta = e_0, \\ \mathbb{F}_2, & \text{otherwise.} \end{cases}$$

Proof. If the internal part $(\beta_2, \dots, \beta_{r-1})$ is nonzero, then Lemma 6 gives a nonempty fixed-boundary attainable set for every boundary quadruple, and in particular for both values of d_0 ; so $P(\beta) = \mathbb{F}_2$. There remain the four differences with zero internal part, which are covered by Lemma 4: with the encoding $\text{enc}(d_0, d_1, \beta_0, \beta_1) = d_0 + 2d_1 + 4\beta_0 + 8\beta_1$, the supported types are B_r . For $\beta = 0$ the admissible codes among $\{0, 1, 2, 3\}$ are $\{0\}$, forcing $d_0 = 0$; for $\beta = e_0$ the admissible codes among $\{4, 5, 6, 7\}$ are $\{5, 7\}$, both with $d_0 = 1$; for $\beta = e_1$ they are $\{10, 11\}$ and for $\beta = e_0 + e_1$ they are $\{13, 14\}$ (together with 15 when $r \geq 5$), and in these last two cases both values of d_0 occur. $\square$

Lemma 10. *If $\alpha_L \neq 0$, then there exists $d \in S_\alpha$ with $\hat{d}_0 = 1$. Symmetrically, if $\alpha_R \neq 0$, then there exists $d \in S_\alpha$ with $d_0 = 1$.*

Proof. Assume $\alpha_L \neq 0$ and apply Lemma 8 with $e = 1$; the shift is then $c + e = c + 1$, and it suffices to exhibit one $c \in \mathbb{F}_2$ with

$$c \in P(\alpha_L + (c + 1)e_0) \quad \text{and} \quad 1 \in P(\alpha_R + (c + 1)e_0).$$

By Lemma 9, the membership $0 \in P(\beta)$ fails only for $\beta = e_0$, and $1 \in P(\beta)$ fails only for $\beta = 0$.

If $\alpha_R \neq e_0$, take $c = 0$. The two conditions read $0 \in P(\alpha_L + e_0)$, which fails only if $\alpha_L + e_0 = e_0$, that is $\alpha_L = 0$; and $1 \in P(\alpha_R + e_0)$, which fails only if $\alpha_R + e_0 = 0$, that is $\alpha_R = e_0$. Both are excluded.

If $\alpha_R = e_0$, take $c = 1$. The shift vanishes and the two conditions read $1 \in P(\alpha_L)$, which fails only if $\alpha_L = 0$, and $1 \in P(e_0) = \{1\}$, which holds. Again both are satisfied.

In either case Lemma 8 produces $d \in S_\alpha$ with $\hat{d}_0 = 1$. The right-branch assertion follows by exchanging the two cycles. $\square$

4.3 Filling the Free Branch

With $\hat{d}_0 = 1$ secured, the right branch is free, and we have only to see that a free branch supplies the whole internal projection.

Lemma 11 (Fixed-endpoint slices). *Let $r \geq 3$ be odd. For $\gamma \in \mathbb{F}_2^r$, put $T_\gamma = \{d \in S_\gamma : d_0 = 1\}$. If $\beta \in \mathbb{F}_2^r$ and π^{int} deletes coordinate 0, then*

$$\pi^{\text{int}}(\text{dir Aff}(T_\beta \cup T_{\beta+e_0})) = \mathbb{F}_2^{r-1}.$$

Proof. Write $a_i = \beta_i$ for $1 \leq i < r$, and put

$$U = T_\beta \cup T_{\beta+e_0}, \quad W = \text{dir Aff}(U).$$

Since β and $\beta + e_0$ differ only at coordinate 0, U consists of the vectors d with $d_0 = 1$ that belong to $S_{(\varepsilon, a_1, \dots, a_{r-1})}$ for some $\varepsilon \in \mathbb{F}_2$. Thus

$$W \subseteq \langle e_1, \dots, e_{r-1} \rangle.$$

Choose $\varepsilon = 1 + \sum_{i=1}^{r-1} a_i$. Then $(\varepsilon, a_1, \dots, a_{r-1})$ has odd weight, so Theorem 1 gives $\mathbf{1} \in U$. We use the interval indicators $I_{q,g}$ from Section 3.3. All intervals below lie in $\{1, \dots, r-1\}$, so the corresponding differences have zeroth coordinate 1. For $\gamma = (1 + a_{r-1}, a_1, \dots, a_{r-1})$, Lemma 3 with $(q, g) = (1, 1)$ gives $\mathbf{1} + I_{1,1} \in T_\gamma \subseteq U$. Hence $e_1 = I_{1,1} \in W$.

For $2 \leq k < r$, let

$$j = \max(\{i : 1 \leq i \leq k-2, a_i = 1\} \cup \{0\}), \quad q = j+1, \quad g = k-j.$$

If $j = 0$, then $a_1 = \dots = a_{k-2} = 0$; applying Lemma 3 to $(1, a_1, \dots, a_{r-1})$ gives $\mathbf{1} + I_{q,g} \in U$. If $j > 0$, then $a_j = 1$ and $a_{j+1} = \dots = a_{k-2} = 0$, so the same lemma applied to β gives the same conclusion. Thus

$$I_{q,g} = e_q + e_{q+1} + \dots + e_k \in W.$$

Induction on k , starting from $e_1 \in W$, yields

$$e_k = I_{q,g} + \sum_{i=q}^{k-1} e_i \in W \quad (2 \leq k < r).$$

Hence $W = \langle e_1, \dots, e_{r-1} \rangle$, and the assertion follows. $\square$

Lemma 12 (Full internal projection). *If some $d \in S_\alpha$ satisfies $\hat{d}_0 = 1$, then $\dim \pi_R^{\text{int}}(D_\alpha) = m' - 1$. Symmetrically, if some $d \in S_\alpha$ satisfies $d_0 = 1$, then $\dim \pi_L^{\text{int}}(D_\alpha) = m - 1$.*

Proof. We prove the first assertion; the second follows by symmetry. For a fixed endpoint pair (c, e) , Lemma 8 shows that, whenever the pair is attainable, the right-branch data range over the entire e -slice of $S_{\alpha_R + (c+e)e_0}$. In particular, $(c, 1)$ is attainable if and only if

$$c \in P(\alpha_L + (c+1)e_0) \quad \text{and} \quad 1 \in P(\alpha_R + (c+1)e_0).$$

Lemma 9 therefore gives

$$\begin{aligned} (1, 1) \text{ is attainable} &\iff \alpha_L \neq 0 \text{ and } \alpha_R \neq 0, \\ (0, 1) \text{ is attainable} &\iff \alpha_L \neq 0 \text{ and } \alpha_R \neq e_0. \end{aligned}$$

The hypothesis implies that one of these pairs occurs, hence $\alpha_L \neq 0$. Moreover, T_{α_R} is empty when $\alpha_R = 0$, and otherwise the pair $(1, 1)$ realizes its whole right-branch slice. Likewise, $T_{\alpha_R + e_0}$ is empty when $\alpha_R = e_0$, and otherwise $(0, 1)$ realizes its whole slice. Consequently,

$$T_{\alpha_R} \cup T_{\alpha_R + e_0} \subseteq \pi_R(S_\alpha).$$

Taking affine hulls and direction spaces, and then projecting onto the internal coordinates, Lemma 11 gives

$$\pi_R^{\text{int}}(\text{dir}(D_\alpha)) = \mathbb{F}_2^{m'-1},$$

which proves the first assertion. $\square$

Theorem 2 (Conjecture 2 for the SWAP family). *For every nonzero $\alpha \in \mathbb{F}_2^N$, $\#\{\lambda \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\} \leq 2^{\max\{m, m'\}+1}$. Equivalently, $\dim K_H(\alpha) \leq \max\{m, m'\} + 1$ for every nonzero direction α .*

Proof. If $\alpha_L \neq 0$, Lemma 10 and Lemma 12 give $\dim D_\alpha \geq \dim \pi_R^{\text{int}}(D_\alpha) = m' - 1$. If $\alpha_L = 0$, then $\alpha_R \neq 0$, and the symmetric assertions give $\dim D_\alpha \geq \dim \pi_L^{\text{int}}(D_\alpha) = m - 1$. In both cases, $\dim D_\alpha \geq \min\{m, m'\} - 1$. Equation (8) proves the stated cardinality bound. The equivalent linear-height bound follows from (2). $\square$

5 Proof of Conjecture 1

In Section 4 we fixed an input difference α and exhibited enough attainable differences to make D_α large, which by Lemma 1 makes $K_H(\alpha)$ small. Here we fix an output mask λ instead, and the same constructions turn into an obstruction. Any α with $\text{Bias}(\alpha, \lambda) = \pm 1$ satisfies $\text{dir}(D_\alpha) \subseteq \lambda^\perp$, so each direction constructed inside $\text{dir}(D_\alpha)$ is one more direction that λ is required to annihilate. Construct enough of them and no nonzero λ survives. The resulting contradiction forces coordinates of α to vanish, and the number of coordinates it kills is what bounds $\dim E_\lambda$.

5.1 Statement and Proof Strategy

Fix a nonzero mask λ . By definition,

$$E_\lambda = \{\alpha \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\} = \{\alpha \in \mathbb{F}_2^N : \Delta_\alpha H_\lambda \text{ is constant}\}. \quad (14)$$

Thus $E_\lambda = \text{LS}(H_\lambda)$, a vector space by the classical linear-structure property recalled in Section 2.3, and

$$\#E_\lambda = 2^{\dim E_\lambda}. \quad (15)$$

Conjecture 1 is equivalent to

$$\dim E_\lambda \leq \max\{m, m'\} + 1 \quad (\lambda \neq 0).$$

The mask decides which construction is needed, so it decides the case split. A one-sided mask is treated in Section 5.2, a two-sided mask with both branches of length at least five in Section 5.3, and a two-sided mask with a three-cycle branch in Section 5.4. The one remaining case, $(m, m') = (3, 3)$, is settled by the finite verification of Appendix A.

We use branch-adapted unit-vector notation throughout this section: e_i has its unique nonzero coordinate at the left-branch position d_i , whereas $\hat{e}_j$ has its unique nonzero coordinate at the right-branch position $\hat{d}_j$. The right branch is indexed by $\hat{e}_0, \dots, \hat{e}_{m'-1}$. For two solutions (x, d) and (x', d') of one difference system with the same input difference α , we write $\delta d_i = d_i + d'_i$ and $\delta \hat{d}_i = \hat{d}_i + \hat{d}'_i$ for the coordinates of the resulting direction $d + d' \in \text{dir}(D_\alpha)$.

Fix $\alpha \in \mathbb{F}_2^N$, and put

$$S_\alpha = \{H(y + \alpha) + H(y) : y \in \mathbb{F}_2^N\}, \quad \text{dir}(D_\alpha) = \text{span}\{d + d^{(0)} : d \in S_\alpha\},$$

where $d^{(0)}$ is any fixed element of S_α . By (14),

$$\alpha \in E_\lambda \iff \text{dir}(D_\alpha) \subseteq \lambda^\perp. \quad (16)$$

Everything that follows consists of constructing elements of $\text{dir}(D_\alpha)$ until (16) can no longer hold.

5.2 One-Sided Masks

A one-sided mask sees only one branch. If the mask is supported on the left, a nonzero internal difference on the right is already enough to make the left projection $\pi_L(D_\alpha)$ equal $\mathbb{F}_2^m$, which no nonzero λ_L can tolerate.

Lemma 13. *If $\lambda_L \neq 0$ and $\lambda_R = 0$, then $\dim E_\lambda \leq m + 1$. Symmetrically, if $\lambda_L = 0$ and $\lambda_R \neq 0$, then $\dim E_\lambda \leq m' + 1$.*

Proof. Assume $\lambda_L \neq 0$ and $\lambda_R = 0$. For a fixed input difference α , the notation of Section 2.1 gives

$$\text{dir}(\pi_L(D_\alpha)) = \pi_L(\text{dir}(D_\alpha)) = \text{span}\{\pi_L(d) + \pi_L(d') : d, d' \in S_\alpha\}.$$

We first prove that

$$\alpha_R^{\text{int}} \neq 0 \implies \pi_L(D_\alpha) = \mathbb{F}_2^m. \quad (17)$$

Switching the right endpoint. Keep $\hat{d}_0$ fixed and move d_0 . Fix an arbitrary $\varepsilon \in \mathbb{F}_2$. We claim that the right-branch difference equations admit two solutions satisfying

$$\hat{d}_0 = \hat{d}'_0 = \varepsilon, \quad d_0 = 0, \quad d'_0 = 1.$$

On an ordinary right-branch coordinate, the local difference equation can be written in the backward form

$$\hat{d}_i = \hat{\alpha}_i + \hat{d}_{i+2}\hat{x}_{i+1} + \hat{d}_{i+1}\hat{x}_{i+2} + \hat{d}_{i+1}\hat{d}_{i+2} + \hat{d}_{i+2},$$

where right-branch indices are reduced modulo m' . The right connection equation is

$$d_0 = \hat{\alpha}_0 + \hat{d}_2\hat{x}_1 + \hat{d}_1\hat{x}_2 + \hat{d}_1\hat{d}_2 + \hat{d}_2.$$

Set $\hat{d}_0 = \varepsilon$ and choose $\hat{x}_0 = \varepsilon + 1$. The last two backward equations become

$$\begin{aligned} \hat{d}_{m'-1} &= \hat{\alpha}_{m'-1} + \varepsilon\hat{x}_1, \\ \hat{d}_{m'-2} &= \hat{\alpha}_{m'-2} + \hat{d}_{m'-1} + \varepsilon(\hat{x}_{m'-1} + 1). \end{aligned}$$

Let k be the largest index in $\{1, \dots, m' - 1\}$ for which $\hat{\alpha}_k = 1$. Such an index exists because $\alpha_R^{\text{int}} \neq 0$. Choose the local variables with index larger than k so that

$$\hat{d}_{k+1} = \dots = \hat{d}_{m'-1} = 0.$$

The backward equation at index k then gives $\hat{d}_k = \hat{\alpha}_k = 1$. By Lemma 8, an attainable difference with $d_0 = c$ and $\hat{d}_0 = \varepsilon$ exists precisely when the right branch admits an attainable difference of $\chi_{m'}$ with 0-coordinate ε for the output difference $\beta_c = \alpha_R + (c + \varepsilon)e_0$, and correspondingly on the left. Since e_0 has zero internal part, β_c and α_R have the same internal part; as $\alpha_R^{\text{int}} \neq 0$, it follows that $\beta_c \notin \{0, e_0\}$ for both values of c , so Lemma 9 gives $P(\beta_c) = \mathbb{F}_2$ and in particular $\varepsilon \in P(\beta_c)$. Both values of c are therefore realized with $\hat{d}_0 = \varepsilon$, which is the claimed pair of right-branch solutions.

Synchronizing the left branch. Construct two left-branch solutions with

$$d_0 = 0, \quad d'_0 = 1,$$

but with identical internal differences and identical right endpoint. Use base points x, x' and impose

$$x_0 = 1, \quad x'_0 = 0, \quad x_1 = x'_1 = 0, \quad x_{m-1} = x'_{m-1} = 1,$$

together with

$$x_i = x'_i \quad (2 \leq i \leq m-2).$$

The backward equation at coordinate $m-1$ gives

$$\begin{aligned} d_{m-1} &= \alpha_{m-1} + d_1 x_0 + d_0 x_1 + d_0 d_1 + d_1 = \alpha_{m-1}, \\ d'_{m-1} &= \alpha_{m-1} + d'_1 x'_0 + d'_0 x'_1 + d'_0 d'_1 + d'_1 = \alpha_{m-1}. \end{aligned}$$

The two values d_{m-1} and d'_{m-1} are equal. At coordinate $m-2$, the two equations reduce respectively to

$$d_{m-2} = \alpha_{m-2} + d_{m-1}, \quad d'_{m-2} = \alpha_{m-2} + d'_{m-1},$$

so $d_{m-2} = d'_{m-2}$. Proceeding backwards through $i = m-3, m-4, \dots, 1$, every equation involves only already synchronized successor differences and synchronized local variables. Induction gives

$$d_i = d'_i \quad (1 \leq i \leq m-1).$$

Finally, the left connection equation, together with $x_1 = x'_1 = 0$, $x_2 = x'_2$, $d_1 = d'_1$, and $d_2 = d'_2$, gives

$$\hat{d}_0 = \alpha_0 + d_1 x_2 + d_1 d_2 + d_2 = \hat{d}'_0.$$

Denote this common endpoint value by ε .

Combining the two halves. The switching construction above supplies compatible right-branch solutions for this same ε , with endpoint values $d_0 = 0$ and $d'_0 = 1$. The left and right halves combine into two complete attainable differences $d, d' \in S_\alpha$ satisfying

$$\pi_L(d) + \pi_L(d') = (1, 0, \dots, 0) = e_0.$$

This pair shows that $e_0 \in \text{dir}(\pi_L(D_\alpha))$. Since $\alpha_R^{\text{int}} \neq 0$, we have $\alpha_R \neq 0$. The symmetric assertion of Lemma 10 gives an attainable difference with $d_0 = 1$. The symmetric assertion of Lemma 12 therefore implies

$$\pi_L^{\text{int}}(\text{dir}(D_\alpha)) = \mathbb{F}_2^{m-1}.$$

For each $1 \leq i < m$, choose $v_i \in \text{dir}(\pi_L(D_\alpha))$ whose internal projection is the i -th standard unit vector. If the zeroth coordinate of v_i is one, add e_0 . The resulting vector is e_i , and hence

$$e_1, e_2, \dots, e_{m-1} \in \text{dir}(\pi_L(D_\alpha)).$$

These directions span $\mathbb{F}_2^m$, so $\text{dir}(\pi_L(D_\alpha)) = \mathbb{F}_2^m$ and the affine space $\pi_L(D_\alpha)$ itself equals $\mathbb{F}_2^m$. This proves (17).

The consequence for E_λ . Now let $\alpha \in E_\lambda$. Equation (16) and $\lambda_R = 0$ imply

$$\text{dir}(\pi_L(D_\alpha)) \subseteq \lambda_L^\perp.$$

If $\alpha_R^{\text{int}} \neq 0$, then (17) makes the left-hand direction space equal to $\mathbb{F}_2^m$, so $\mathbb{F}_2^m \subseteq \lambda_L^\perp$, which forces $\lambda_L = 0$, a contradiction. Every $\alpha \in E_\lambda$ must satisfy

$$\alpha_R^{\text{int}} = 0.$$

The $m' - 1$ right internal coordinates vanish identically on E_λ , and

$$\dim E_\lambda \leq N - (m' - 1) = m + 1.$$

The second assertion follows by exchanging the two branches. $\square$

5.3 Two-Sided Masks

When both halves of the mask are nonzero, neither branch can absorb a full projection, and the bound drops from $m + 1$ to 2. Getting there needs the internal directions of the opposite branch (Lemma 14) and its endpoint direction (Lemma 15); together they span the whole branch (Lemma 16), which leaves room for nothing but the two endpoint coordinates of α .

Lemma 14. *If $m' \geq 5$ and $\alpha_L^{\text{int}} \neq 0$, then*

$$\hat{e}_1, \dots, \hat{e}_{m'-1} \in \text{dir}(D_\alpha).$$

Symmetrically, if $m \geq 5$ and $\alpha_R^{\text{int}} \neq 0$, then $e_1, \dots, e_{m-1} \in \text{dir}(D_\alpha)$.

Proof. The construction is carried out in Appendix F. $\square$

Lemma 15. *Assume $m, m' \geq 5$, $\alpha_0 = \hat{\alpha}_0 = 0$, and $\alpha_L^{\text{int}} \neq 0$. Then there is a $\delta d \in \text{dir}(D_\alpha)$ such that*

$$\delta d_0 = \dots = \delta d_{m-1} = 0, \quad \delta \hat{d}_0 = 1. \quad (18)$$

The symmetric assertion holds with the two sides exchanged.

Proof. The construction is carried out in Appendix F. $\square$

Lemma 16. *Assume $m, m' \geq 5$. If $\alpha_0 = \hat{\alpha}_0 = 0$ and $\alpha_L^{\text{int}} \neq 0$, then*

$$\langle \hat{e}_0, \hat{e}_1, \dots, \hat{e}_{m'-1} \rangle \subseteq \text{dir}(D_\alpha). \quad (19)$$

The symmetric statement holds when $\alpha_R^{\text{int}} \neq 0$.

Proof. By Lemma 14, $\hat{e}_1, \dots, \hat{e}_{m'-1} \in \text{dir}(D_\alpha)$. Lemma 15 provides an element of $\text{dir}(D_\alpha)$ equal to $\hat{e}_0 + v$, where $v \in \langle \hat{e}_1, \dots, \hat{e}_{m'-1} \rangle$. Subtracting v gives $\hat{e}_0 \in \text{dir}(D_\alpha)$, proving (19). The symmetric assertion follows by exchanging the two cycles. $\square$

Lemma 17. *If $m, m' \geq 5$ and $\lambda_L \neq 0$, $\lambda_R \neq 0$, then $\dim E_\lambda \leq 2$.*

Proof. Let $\alpha \in E_\lambda$ satisfy $\alpha_0 = \hat{\alpha}_0 = 0$. If $\alpha \neq 0$, at least one of α_L^{int} and α_R^{int} is nonzero. Lemma 16 and (16) then force, respectively, $\lambda_R = 0$ or $\lambda_L = 0$, a contradiction. The map

$$\rho : E_\lambda \longrightarrow \mathbb{F}_2^2, \quad \rho(\alpha) = (\alpha_0, \hat{\alpha}_0),$$

has trivial kernel. It follows that $\dim E_\lambda \leq 2$. $\square$

5.4 A Three-Cycle Branch

A branch of length three is too short for the constructions of Section 5.3, which need $m, m' \geq 5$. It is also small enough for us to handle by direct computation.

Lemma 18. *If $m \geq 5$, $m' = 3$, and $\lambda_L \neq 0$, $\lambda_R \neq 0$, then $\dim E_\lambda \leq m+1$. The symmetric bound holds for $m = 3$, $m' \geq 5$.*

Proof. Suppose $m \geq 5$ and $m' = 3$, and consider the kernel of the coordinate map $\pi_0 : \alpha \mapsto \alpha_0$ on E_λ , and let $K = \ker \pi_0$. We show that $\alpha \in K$ implies $\alpha_R^{\text{int}} = 0$. Otherwise, $\alpha_R^{\text{int}} \neq 0$, so $\alpha_R \notin \{0, e_0\}$ and Lemma 9 gives $0 \in P(\alpha_R)$. Since $\alpha \in K$, we have $\alpha_0 = 0$, hence $\alpha_L \neq e_0$, and the same lemma gives $0 \in P(\alpha_L)$. Lemma 8, applied with $(d_0, \hat{d}_0) = (0, 0)$, now supplies an attainable difference with $\hat{d}_0 = 0$. The two internal equations of the right three-cycle are then

$$\hat{d}_1 = \hat{\alpha}_1 + \hat{\alpha}_2 \hat{x}_0, \quad \hat{d}_2 = \hat{\alpha}_2 + \hat{\alpha}_1 + \hat{\alpha}_1 \hat{x}_0.$$

Because $(\hat{\alpha}_1, \hat{\alpha}_2) \neq (0, 0)$, the pair $(\hat{d}_1, \hat{d}_2)$ is nonzero. Fix $\hat{x}_0$, choose two differences with the same values of $\hat{d}_0, \hat{d}_1, \hat{d}_2$, and subtract their endpoint equations. This gives

$$\delta d_0 = \hat{d}_1(\hat{x}_2 + \hat{x}_2') + \hat{d}_2(\hat{x}_1 + \hat{x}_1').$$

One of the two coefficients is nonzero, so the corresponding free variable can be chosen to make $\delta d_0 = 1$. This gives $e_0 + v \in \text{dir}(D_\alpha)$ for some $v \in \langle e_1, \dots, e_{m-1} \rangle$. By Lemma 14, all the directions in v also belong to $\text{dir}(D_\alpha)$, so $e_0 \in \text{dir}(D_\alpha)$. Equation (16) then forces $\lambda_L = 0$, a contradiction.

The coordinates $\alpha_0, \hat{\alpha}_1, \hat{\alpha}_2$ vanish on K , and hence $\dim K \leq N - 3 = m$. Since $\pi_0(E_\lambda) \subseteq \mathbb{F}_2$,

$$\dim E_\lambda = \dim K + \dim \pi_0(E_\lambda) \leq m + 1.$$

The second assertion is symmetric. □

Theorem 3 (Conjecture 1 for the SWAP family). *For every nonzero $\lambda \in \mathbb{F}_2^N$,*

$$\#\{\alpha \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\} \leq 2^{\max\{m, m'\} + 1}.$$

Equivalently, every nonzero component of H has linear nullity at most $\max\{m, m'\} + 1$.

Proof. For a one-sided mask, the result follows from Lemma 13. For a two-sided mask with $m, m' \geq 5$, it follows from Lemma 17; with exactly one three-cycle, it follows from Lemma 18. The exceptional case $(m, m') = (3, 3)$ is the proof-critical finite verification recorded in Lemma 29 of Appendix A. In every case, $\dim E_\lambda \leq \max\{m, m'\} + 1$. Equation (15) now proves the cardinality bound and the equivalent linear-nullity statement. □

6 Proof of Conjecture 4

We proved the height bound for the SWAP family in Section 4 by cutting the two branches apart. We could cut because the SWAP construction couples the branches at a single pair of coordinates: by Lemma 8, fixing $(d_0, \hat{d}_0)$ turns the global system into two independent ordinary χ systems, one per branch. The 4-cycle construction modifies two coordinates on each branch rather than one, and no pair of values separates them.

We replace it with a coarser cut. Conditioning on all four boundary differences at once does split the remaining constraints into two ordinary χ systems, but it leaves those four coordinates outside the split, which we handle separately. A branch whose internal difference is nonzero admits every boundary quadruple, so it constrains the boundary

not at all. When both internal differences vanish, both branches are rigid and the entire difference is confined to those four coordinates.

Let $m, m' \geq 3$ be odd, $N = m + m'$, $(b, \hat{b}) \in \{(0, 0), (0, 1), (1, 0)\}$, $F = \mathcal{X}_{m, m'}^{(b, \hat{b})}$, and $H = F^{-1}$. Fix a nonzero input difference $\alpha = (\alpha_0, \dots, \alpha_{m-1}, \hat{\alpha}_0, \dots, \hat{\alpha}_{m'-1})$ of H . Recall that $K_H(\alpha) = \{\lambda \in \mathbb{F}_2^N : \Delta_\alpha \langle \lambda, H \rangle \text{ is constant}\}$. The required linear-height bound is $\dim K_H(\alpha) \leq (\max\{m, m'\} + 1)/2$.

We use the convention of Section 2: α is an input difference of H and d is an output difference of H , with the forward-branch reading of both given by the viewpoint table there.

Let S_α be the attainable output-difference set and $D_\alpha = \text{Aff}(S_\alpha)$. By Lemma 1, $K_H(\alpha) = \text{dir}(D_\alpha)^\perp$ and hence $\dim K_H(\alpha) = N - \dim D_\alpha$. It is therefore enough to prove $\dim D_\alpha \geq N - (\max\{m, m'\} + 1)/2$. We obtain this lower bound by conditioning on the four boundary differences, reducing the remaining constraints to two ordinary- χ systems, and treating separately the boundary-only and nonzero-internal cases.

6.1 Conditioning on the Boundary

The two branches of the 4-cycle construction interact only through four boundary coordinates. We condition on those coordinates, reduce the remaining constraints to two ordinary χ cycles, and establish the local affine-rank estimates used in the two case analyses below. To formalize this reduction, let

$$d_{\text{bd}} = (d_0, d_1, \hat{d}_0, \hat{d}_1) \in \mathbb{F}_2^4$$

denote the boundary output-difference profile. Comparison with the boundary equations of two ordinary χ cycles gives

$$\begin{aligned} \alpha_0^{\chi, L} &= \alpha_0 + d_0 + \hat{d}_1 + b d_1, & \alpha_1^{\chi, L} &= \alpha_1 + d_1 + \hat{d}_0, \\ \hat{\alpha}_0^{\chi, R} &= \hat{\alpha}_0 + \hat{d}_0 + d_0 + \hat{b} \hat{d}_1, & \hat{\alpha}_1^{\chi, R} &= \hat{\alpha}_1 + \hat{d}_1 + d_1. \end{aligned} \quad (20)$$

These identities define the boundary maps

$$\begin{aligned} \phi_L(d_{\text{bd}}) &= (d_0, d_1, \alpha_0^{\chi, L}, \alpha_1^{\chi, L}), \\ \phi_R(d_{\text{bd}}) &= (\hat{d}_0, \hat{d}_1, \hat{\alpha}_0^{\chi, R}, \hat{\alpha}_1^{\chi, R}). \end{aligned} \quad (21)$$

Lemma 19 (Reduction to two ordinary χ cycles). *For each fixed boundary output-difference profile d_{bd} , the differential constraints of the 4-cycle construction separate into two independent ordinary- χ systems. Their boundary data are $\phi_L(d_{\text{bd}})$ and $\phi_R(d_{\text{bd}})$, respectively, while their internal input differences are η_L and η_R . Moreover, both ϕ_L and ϕ_R are affine bijections of $\mathbb{F}_2^4$.*

Proof. All cross-cycle terms occur in the four boundary equations. For fixed d_{bd} , equation (20) absorbs these terms into the boundary data of the two ordinary cycles. Every remaining coordinate equation then involves variables from only one branch and is exactly an ordinary- χ difference equation. Finally, the linear parts of the maps in (21) are nonsingular: either boundary quadruple uniquely determines the four components of d_{bd} . Both ϕ_L and ϕ_R are affine bijections. $\square$

The fixed-boundary attainable sets $\mathcal{C}_r$ of Subsection 3.4, and the ordinary- χ difference equation (4) they are built from, are exactly the local objects this reduction produces.

Lemma 19 now yields the product decomposition

$$S_\alpha = \bigcup_{d_{\text{bd}} \in \mathbb{F}_2^4} \mathcal{C}_m(\eta_L; \phi_L(d_{\text{bd}})) \times \mathcal{C}_{m'}(\eta_R; \phi_R(d_{\text{bd}})), \quad (22)$$

where the abbreviated arguments have the meaning specified in (5); a product with an empty factor is omitted. Accordingly, the feasible boundary output-difference profiles are exactly

$$T_\alpha = \{d_{\text{bd}} \in \mathbb{F}_2^4 : \mathcal{C}_m(\eta_L; \phi_L(d_{\text{bd}})) \neq \emptyset, \mathcal{C}_{m'}(\eta_R; \phi_R(d_{\text{bd}})) \neq \emptyset\}. \quad (23)$$

Here T_α is the set of feasible values of d_{bd} ; it controls only the four boundary output-difference bits, whereas the two sets control the internal coordinates of the left and right branches.

The two local estimates supplied by Subsection 3.4 are Lemma 4 for a branch with zero internal input difference and Lemma 5 for one with a nonzero internal input difference; Lemma 6 records which boundary types the latter supports.

To connect these local contributions to the global affine rank, we need two additional pieces of boundary information. For a set $T \subseteq \mathbb{F}_2^4$, the expressions $\phi_L(T)$ and $\phi_R(T)$ refer to the corresponding sets of ordinary- χ boundary quadruples. Direct substitution in (20) gives the multisets of admissible input-difference counts in Table 1. In each row, the four numbers count the possible values of the boundary input pair (d_0, d_1) for the four fixed values of the boundary output pair (α_0, α_1) .

Table 1: Numbers of admissible active-branch input differences induced by a zero branch.

length of the zero branch	multiset of admissible-input counts
3	$\{3, 2, 2, 0\}$ or $\{4, 1, 1, 1\}$
at least 5	$\{3, 2, 2, 1\}$ or $\{4, 2, 1, 1\}$

The two possibilities in each row cover all three admissible values of $(b, \hat{b})$ and both choices of the active branch. Every induced active-branch boundary set satisfies the three-input-difference hypothesis of Lemma 5.

Let

$$\pi(d, \hat{d}) = (d_0, d_1, \hat{d}_0, \hat{d}_1).$$

Directions within a product set in (22) lie in $\ker \pi$. Directions from the two branches have disjoint coordinate supports, whereas projection under π supplies the direction space of $\text{Aff}(T_\alpha)$. Applying the rank-nullity theorem [Axl24, Theorem 3.21] to the restricted projection $\pi|_{\text{dir}(D_\alpha)}$ gives the uniform lower bound

$$\begin{aligned} \dim \text{dir}(D_\alpha) &\geq \dim W_m(\eta_L; \phi_L(T_\alpha)) + \dim W_{m'}(\eta_R; \phi_R(T_\alpha)) \\ &\quad + \dim \text{Aff}(T_\alpha). \end{aligned} \quad (24)$$

In this decomposition, a zero internal difference contributes through (6), a nonzero internal difference contributes through Lemma 5, and the four boundary coordinates contribute $\dim \text{Aff}(T_\alpha)$. Equation (24) is the interface between the local fixed-boundary estimates and the two global cases.

6.2 Boundary-Only Directions

A branch with a nonzero internal difference imposes no boundary constraint, so the case that has to be worked out is the one where both internal words η_L and η_R vanish. There the entire input difference sits on the four boundary coordinates, these sets reduce to a single linear system, and Lemma 20 writes that system down. Section 6.3 then solves it.

Suppose first that the internal input-difference words of both ordinary- χ branches are zero, that is,

$$\eta_L = 0, \quad \eta_R = 0. \quad (25)$$

Under this assumption, the fixed input difference α of H can be nonzero only in the four boundary coordinates. We denote this fixed boundary input-difference profile by

$$\alpha_{\text{bd}} = (\alpha_0, \alpha_1, \hat{\alpha}_0, \hat{\alpha}_1) \in \mathbb{F}_2^4 \setminus \{0\},$$

whose coordinates correspond to $y_0, y_1, \hat{y}_0, \hat{y}_1$. Here and in the tables below, $e_{y_i}, e_{\hat{y}_i}, e_{x_i}$, and $e_{\hat{x}_i}$ denote the standard unit vectors of $\mathbb{F}_2^N$ supported at the named coordinate: input differences α of H are written in the y -coordinates, whereas output differences d and masks λ are written in the x -coordinates. In contrast, we write

$$d_{\text{bd}} = (d_0, d_1, \hat{d}_0, \hat{d}_1) \in \mathbb{F}_2^4$$

for a variable boundary output-difference profile of H . The roles of α and d are reversed when we pass from the inverse map H to a forward ordinary- χ branch: an input difference of H is an output difference of the forward branch, and an output difference of H is an input difference of that branch.

The boundary-only analysis follows the chain

$$d_{\text{bd}} \mapsto (\phi_L(d_{\text{bd}}), \phi_R(d_{\text{bd}})) \mapsto A_{d_{\text{bd}}} \mapsto R_\alpha \mapsto K_H(\alpha).$$

The boundary maps select the two sets. Their affine hulls determine the within-profile and between-profile directions, whose span determines the full-bias kernel.

For a candidate profile d_{bd} , the maps ϕ_L and ϕ_R from (21) produce the ordinary- χ boundary quadruples for the left and right branches, respectively. The sets B_m and $B_{m'}$ record the boundary types supported by those branches when their internal input differences are zero. A global profile is feasible only when both branch systems have a solution. Accordingly, (23) becomes

$$T_\alpha = \{d_{\text{bd}} \in \mathbb{F}_2^4 : \text{enc}(\phi_L(d_{\text{bd}})) \in B_m, \text{enc}(\phi_R(d_{\text{bd}})) \in B_{m'}\}.$$

Thus, T_α is exactly the set of variable boundary output-difference profiles compatible with the fixed input difference α on both branches simultaneously.

For each feasible profile, Table 4 describes the affine hull contributed by either branch. For a supported ordinary- χ boundary quadruple $q = (d_0, d_1, \alpha_0, \alpha_1)$, the table gives the affine hull, rather than necessarily the attainable set itself:

$$\text{Aff } \mathcal{C}_r(0, \dots, 0; d_0, d_1, \alpha_0, \alpha_1) = c_r(q) + U_r(q).$$

Here $c_r(q) \in \mathbb{F}_2^r$ is one chosen representative point, or anchor, and $U_r(q) \subseteq \mathbb{F}_2^r$ is the direction space of the affine hull. Explicitly,

$$c_r(q) + U_r(q) = \{c_r(q) + u : u \in U_r(q)\}.$$

The anchor is not unique, but the direction space is unique. For example, the row $(\alpha_0, \alpha_1) = 10, (d_0, d_1) = 10$ states that the relevant affine hull is $e_0 + I_r$: the boundary coordinates are represented by e_0 , while every direction in I_r may be added.

We now combine the two branch descriptions. For each $d_{\text{bd}} \in T_\alpha$, define the left and right anchors by

$$c_L(d_{\text{bd}}) = c_m(\phi_L(d_{\text{bd}})), \quad c_R(d_{\text{bd}}) = c_{m'}(\phi_R(d_{\text{bd}})),$$

and concatenate them to obtain the global anchor

$$c_{d_{\text{bd}}} = (c_L(d_{\text{bd}}), c_R(d_{\text{bd}})) \in \mathbb{F}_2^{m+m'}.$$

The parentheses concatenate the two branch vectors into an element of $\mathbb{F}_2^N$.

Similarly, define the branch direction spaces

$$U_L(d_{\text{bd}}) = U_m(\phi_L(d_{\text{bd}})), \quad U_R(d_{\text{bd}}) = U_{m'}(\phi_R(d_{\text{bd}})),$$

and their global direct sum

$$\begin{aligned} U_{d_{\text{bd}}} &= U_L(d_{\text{bd}}) \oplus U_R(d_{\text{bd}}) \\ &= \{(u_L, u_R) : u_L \in U_L(d_{\text{bd}}), u_R \in U_R(d_{\text{bd}})\}. \end{aligned}$$

The two components occupy disjoint coordinate blocks, so their directions vary independently and

$$\dim U_{d_{\text{bd}}} = \dim U_L(d_{\text{bd}}) + \dim U_R(d_{\text{bd}}).$$

Here $c_{d_{\text{bd}}}$ is one concatenated anchor, whereas $U_{d_{\text{bd}}}$ contains every concatenated pair of branch directions.

For each feasible boundary profile, the anchor and its direction space form the associated global affine set

$$\begin{aligned} A_{d_{\text{bd}}} &:= c_{d_{\text{bd}}} + U_{d_{\text{bd}}} \\ &= (c_L(d_{\text{bd}}) + U_L(d_{\text{bd}})) \times (c_R(d_{\text{bd}}) + U_R(d_{\text{bd}})). \end{aligned}$$

This affine set is the joint affine hull of the attainable branch differences when the global boundary profile is fixed to d_{bd} . The global affine hull is

$$D_\alpha = \text{Aff} \bigcup_{d_{\text{bd}} \in T_\alpha} A_{d_{\text{bd}}}.$$

Each profile affine set accounts for the internal variation under one fixed boundary profile; the outer affine hull also accounts for directions between sets with different feasible boundary profiles.

The generators of $\text{dir}(D_\alpha)$ therefore come in two distinct forms. For a fixed profile d_{bd} , the space $U_{d_{\text{bd}}} \subseteq \mathbb{F}_2^N$ contains the within-profile directions: they vary an attainable difference without changing its boundary profile. After a reference profile $d_{\text{bd}}^{(0)}$ is chosen, the vector $c_{d_{\text{bd}}} + c_{d_{\text{bd}}^{(0)}} \in \mathbb{F}_2^N$ compares the anchor points associated with two feasible profiles and supplies a between-profile direction. The anchor $c_{d_{\text{bd}}}$ itself depends on the representative chosen for that profile, but the span of all within-profile and between-profile directions does not.

A mask $\lambda \in \mathbb{F}_2^N$ is full-bias precisely when the linear functional $d \mapsto \langle \lambda, d \rangle$ is constant on the union of all feasible profile affine sets. It must consequently annihilate every $U_{d_{\text{bd}}}$ and every anchor difference. The following lemma records these requirements as a homogeneous linear system whose unknown is λ : its first family of equations comes from variation within one profile, its second from comparison between profiles, and its nullspace is exactly $K_H(\alpha)$.

Lemma 20 (Full-bias masks from the profile affine sets). *Fix any reference profile $d_{\text{bd}}^{(0)} \in T_\alpha$, and define*

$$R_\alpha = \text{span} \left(\bigcup_{d_{\text{bd}} \in T_\alpha} U_{d_{\text{bd}}} \cup \left\{ c_{d_{\text{bd}}} + c_{d_{\text{bd}}^{(0)}} : d_{\text{bd}} \in T_\alpha \right\} \right). \quad (26)$$

Then

$$\text{dir}(D_\alpha) = R_\alpha, \quad K_H(\alpha) = R_\alpha^\perp.$$

Equivalently, a mask $\lambda \in \mathbb{F}_2^N$ belongs to $K_H(\alpha)$ if and only if it satisfies both the within-profile and between-profile conditions

$$\begin{aligned} \langle \lambda, u \rangle &= 0 \quad (d_{\text{bd}} \in T_\alpha, \quad u \in U_{d_{\text{bd}}}), \\ \langle \lambda, c_{d_{\text{bd}}} + c_{d_{\text{bd}}^{(0)}} \rangle &= 0 \quad (d_{\text{bd}} \in T_\alpha). \end{aligned} \quad (27)$$

Proof. All additions below are in $\mathbb{F}_2$, so subtraction and bitwise XOR are the same operation. Fix a feasible profile d_{bd} . Every $d \in A_{d_{\text{bd}}}$ has the form

$$d = c_{d_{\text{bd}}} + u, \quad u \in U_{d_{\text{bd}}},$$

so

$$\langle \lambda, d \rangle = \langle \lambda, c_{d_{\text{bd}}} \rangle + \langle \lambda, u \rangle.$$

The inner product is constant on this affine set if and only if $\lambda \perp U_{d_{\text{bd}}}$. These are exactly the first-line, within-profile constraints in (27). When they hold, the constant value on the set is $\langle \lambda, c_{d_{\text{bd}}} \rangle$.

It remains to make these constants agree across different profiles. Using the reference profile, this requirement is

$$\langle \lambda, c_{d_{\text{bd}}} \rangle = \langle \lambda, c_{d_{\text{bd}}^{(0)}} \rangle \quad (d_{\text{bd}} \in T_\alpha),$$

which over $\mathbb{F}_2$ is equivalent to the second-line, between-profile constraints in (27). A full-bias mask annihilates precisely the generators in (26). It follows that $\text{dir}(D_\alpha) = R_\alpha$ and $K_H(\alpha) = R_\alpha^\perp$.

These spaces do not depend on the auxiliary choices. Replacing an anchor $c_{d_{\text{bd}}}$ by another anchor changes it by an element of $U_{d_{\text{bd}}}$, which is already included among the generators of R_α . Replacing the reference profile only replaces the anchor differences by another generating set of the same span. $\square$

The lemma also gives a direct computation. For a fixed α_{bd} , evaluate the sixteen pairs $\phi_L(d_{\text{bd}}), \phi_R(d_{\text{bd}})$ and retain the feasible profiles in T_α . For every retained profile, read direction bases and anchors from Table 4. Form a binary matrix M_α whose rows consist of all these direction-basis vectors and the anchor differences $c_{d_{\text{bd}}} + c_{d_{\text{bd}}^{(0)}}$. Then

$$K_H(\alpha) = \ker M_\alpha, \quad \dim K_H(\alpha) = N - \text{rank } M_\alpha.$$

The classifications below are obtained by exact row reduction over $\mathbb{F}_2$.

Put

$$\ell = \frac{m+1}{2}, \quad \hat{\ell} = \frac{m'+1}{2}. \quad (28)$$

6.3 Kernel Dimensions

Once the fixed-boundary attainable sets are known, the boundary-kernel system reduces by binary row operations. Its solution depends on the branch lengths only through ℓ and $\hat{\ell}$, so one reduction covers every admissible pair, and no boundary-only direction has kernel dimension above $\max\{\ell, \hat{\ell}\}$.

Assume first that $m, m' \geq 5$. Substitute the all-length anchor–direction descriptions from Table 4 into (27) and reduce the resulting binary systems. The resulting exact dimensions for all fifteen nonzero boundary-only directions are recorded in Table 9 of Appendix D.1. The length-dependent entries are formulas in ℓ and $\hat{\ell}$, so the classification covers every admissible pair in the stated range rather than a finite sample of lengths.

The entries larger than two have simple bases. Define

$$\begin{aligned} \mathcal{L}_m &= \langle e_{x_0}, e_{x_3}, e_{x_5}, \dots, e_{x_{m-2}}, e_{\hat{x}_0} \rangle, \\ \mathcal{R}_{m'} &= \langle e_{x_1}, e_{\hat{x}_0}, e_{\hat{x}_3}, e_{\hat{x}_5}, \dots, e_{\hat{x}_{m'-2}} \rangle. \end{aligned} \quad (29)$$

The rows that carry a length-dependent kernel, together with a worked row reduction showing how one of them is obtained, are recorded in Appendix D.2.

It remains to account for a branch of length three. Appendix D.3 lists the kernel dimension for each of the fifteen nonzero boundary-only directions, each of the three admissible parameter pairs, and the three short-cycle geometries $m = 3 < m'$, $m' = 3 < m$, and $(m, m') = (3, 3)$. These exact binary row reductions use B_3 and the direction spaces from the length-three case of Lemma 4 and exhaust $15 \cdot 3 \cdot 3 = 135$ systems. This finite calculation is part of the proof: its largest nullities establish the required bound for every short geometry. It does not supply the long-cycle conclusion, which follows from the preceding symbolic classification.

The length-dependent kernels, including the short cases, are exactly

$(b, \hat{b})$	condition and direction	$K_H(\alpha)$
$(0, 0)$	$\alpha = e_{y_1}$	$\mathcal{L}_m$
$(0, 0)$	$\alpha = e_{\hat{y}_1}$	$\mathcal{R}_{m'}$
$(0, 1)$	$\alpha = e_{\hat{y}_1}$	$\mathcal{R}_{m'}$
$(0, 1)$	$m' = 3, \alpha = e_{y_1} + e_{\hat{y}_1}$	$\mathcal{L}_m$
$(1, 0)$	$\alpha = e_{y_1}$	$\mathcal{L}_m$
$(1, 0)$	$\alpha = e_{y_0} + e_{\hat{y}_1}$	$\mathcal{R}_{m'}$

(30)

An empty alternating list in (29) is omitted. Counting the displayed basis vectors gives

$$\dim \mathcal{L}_m = \frac{m+1}{2}, \quad \dim \mathcal{R}_{m'} = \frac{m'+1}{2}. \quad (31)$$

Equations (30) and (31) give, for every boundary-only direction,

$$\dim K_H(\alpha) \leq \max\{\ell, \hat{\ell}\}.$$

By (28), the boundary-only case satisfies the required bound.

6.4 Directions with a Nonzero Internal Part

We now combine the global rank decomposition (24) with Lemma 5 to treat every configuration in which at least one of η_L, η_R is nonzero.

Lemma 21 (Rank with a nonzero internal difference). *If $\eta_L \neq 0$ or $\eta_R \neq 0$, then*

$$\dim \operatorname{dir}(D_\alpha) \geq N - 2, \quad \dim K_H(\alpha) \leq 2. \quad (32)$$

More precisely, the stronger estimate $\dim \operatorname{dir}(D_\alpha) \geq N - 1$ holds when exactly one of η_L, η_R is nonzero.

Proof. First suppose that $\eta_L \neq 0$ and $\eta_R = 0$. By Lemma 19,

$$T_\alpha = \phi_R^{-1}(B_{m'}).$$

Indeed, by Lemma 6 the active left branch supports every boundary quadruple, so feasibility of a profile is decided by the zero right branch alone. The set $B_{m'}$ has affine dimension four and ϕ_R is an affine bijection, so $\dim \operatorname{Aff}(T_\alpha) = 4$. Table 1 shows that $\phi_L(T_\alpha)$ has the three-input-difference property. Together, Lemma 5 and (6) give, respectively,

$$\dim W_m(\eta_L; \phi_L(T_\alpha)) \geq m - 3, \quad \dim W_{m'}(0; \phi_R(T_\alpha)) = m' - 2.$$

Substitution in (24) yields

$$\dim \operatorname{dir}(D_\alpha) \geq (m - 3) + (m' - 2) + 4 = N - 1.$$

Next suppose that $\eta_L = 0$ and $\eta_R \neq 0$. Exchanging the left and right cycles gives, again by Lemma 6, $T_\alpha = \phi_L^{-1}(B_m)$, $\dim \text{Aff}(T_\alpha) = 4$, and the three-input-difference property for $\phi_R(T_\alpha)$. Equation (6) and Lemma 5 give local contributions $m - 2$ and at least $m' - 3$, respectively, and

$$\dim \text{dir}(D_\alpha) \geq (m - 2) + (m' - 3) + 4 = N - 1.$$

Finally, suppose that $\eta_L \neq 0$ and $\eta_R \neq 0$. Every boundary value is feasible on both ordinary- χ cycles by Lemma 6, so $T_\alpha = \mathbb{F}_2^4$ and $\dim \text{Aff}(T_\alpha) = 4$. Both projected boundary sets have the three-input-difference property, so the two local contributions are at least $m - 3$ and $m' - 3$. Substitution in (24) gives

$$\dim \text{dir}(D_\alpha) \geq (m - 3) + (m' - 3) + 4 = N - 2.$$

These inequalities show that $\dim D_\alpha = \dim \text{dir}(D_\alpha)$ is at least $N - 1$ when exactly one internal difference is nonzero and at least $N - 2$ when both are nonzero. Lemma 1 then shows that all three cases satisfy

$$\dim K_H(\alpha) \leq 2.$$

This proves (32). $\square$

6.5 The Height Theorem and Its Consequences

Combining the preceding lemma with the boundary-only classification gives the claimed height bound for every nonzero input difference.

Theorem 4 (Conjecture 4 for the 4-cycle family). *Let $m, m' \geq 3$ be odd, $(b, \hat{b}) \in \{(0, 0), (0, 1), (1, 0)\}$, and*

$$H = \left(\mathcal{X}_{m, m'}^{(b, \hat{b})} \right)^{-1}.$$

For every $\alpha \neq 0$,

$$\dim K_H(\alpha) \leq \frac{\max\{m, m'\} + 1}{2}. \quad (33)$$

Equivalently,

$$\dim D_\alpha \geq N - \frac{\max\{m, m'\} + 1}{2}. \quad (34)$$

Proof. If $\eta_L \neq 0$ or $\eta_R \neq 0$, Lemma 21 gives $\dim K_H(\alpha) \leq 2$. Since $2 \leq (\max\{m, m'\} + 1)/2$, (33) holds in every case with a nonzero internal difference.

Under the boundary-only condition (25), every value in Table 9 of Appendix D.1 is at most $\max\{\ell, \hat{\ell}\}$ when $m, m' \geq 5$, and Appendix D.3 gives the same conclusion when at least one cycle has length three. Since $\max\{\ell, \hat{\ell}\} = (\max\{m, m'\} + 1)/2$ by (28), this proves (33) in the remaining cases. Finally, the duality $K_H(\alpha) = \text{dir}(D_\alpha)^\perp$ from Lemma 1 yields the equivalent affine form (34). $\square$

Remark 1 (Symplectic origin of the half-length). The factor $1/2$ in (33) is not an artifact of the row reductions behind Table 9. For odd n , let

$$C_n(a, b) = \sum_{i \in \mathbb{Z}_n} (a_i b_{i+1} + a_{i+1} b_i), \quad a, b \in \mathbb{F}_2^n, \quad (35)$$

be the cyclic nearest-neighbour alternating form. Its radical is exactly $\{0, \mathbf{1}\}$: a vector orthogonal to every e_i satisfies $a_{i-1} = a_{i+1}$ for all i , and for odd n the step-two shift is a single cycle on $\mathbb{Z}_n$, so a is constant. Hence C_n descends to a nondegenerate symplectic form on $\mathbb{F}_2^n / \langle \mathbf{1} \rangle$, of dimension $n - 1$. A totally isotropic $V \subseteq \mathbb{F}_2^n$ with $\mathbf{1} \notin V$ meets $\langle \mathbf{1} \rangle$

trivially, so it injects into that quotient with totally isotropic image, and the standard alternating-form bound [Bal15, Theorem 3.3] gives

$$\dim V \leq \frac{n-1}{2}. \quad (36)$$

The two length-dependent kernels of (29) attain this ceiling on their own branch. Since $C_n(e_i, e_j) = 1$ exactly when i and j are cyclically adjacent, a coordinate subspace is totally isotropic as soon as no two of its indices are neighbours. The left-branch part of $\mathcal{L}_m$ is supported on $\{0, 3, 5, \dots, m-2\} \subseteq \mathbb{Z}_m$ (just $\{0\}$ when $m = 3$), which contains no two cyclically adjacent indices and is a proper subset of $\mathbb{Z}_m$; so that part is totally isotropic, avoids $\mathbf{1}$, and has dimension $(m-1)/2$, the largest allowed by (36). Its one remaining generator $e_{\hat{x}_0}$ lies on the opposite branch. Therefore

$$\dim \mathcal{L}_m = \frac{m-1}{2} + 1 = \ell, \quad \dim \mathcal{R}_{m'} = \frac{m'-1}{2} + 1 = \hat{\ell},$$

The second equality follows by the same argument with the branches exchanged. These are the two length-dependent entries of Table 9, so (33) reads: one branch at its symplectic ceiling, plus the seam coordinate.

Corollary 1. *Under the assumptions of Theorem 4,*

$$\#\{\lambda \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\} \leq 2^{(\max\{m, m'\}+1)/2}.$$

Proof. The full-bias masks are exactly the elements of $K_H(\alpha)$. Their number is $2^{\dim K_H(\alpha)}$, so the claim follows from Theorem 4. $\square$

The kernel tables computed for the boundary-only case determine not only the upper bound but the exact maximum of the linear height, together with its maximizing directions.

Corollary 2 (Exact maximum and tightness). *Under the assumptions of Theorem 4,*

$$\max_{\alpha \neq 0} \dim K_H(\alpha) = \begin{cases} \max\{\ell, \hat{\ell}\}, & (b, \hat{b}) \in \{(0, 0), (1, 0)\}, \\ \hat{\ell}, & (b, \hat{b}) = (0, 1) \text{ and } m' \geq 5, \\ \ell, & (b, \hat{b}) = (0, 1) \text{ and } m' = 3. \end{cases}$$

In particular, the bound of Theorem 4 is attained for every admissible pair at the variants (0, 0) and (1, 0), and at (0, 1) exactly when $m \leq m'$ or $m' = 3$. For the variant (0, 1) with $\max\{m, m'\} \geq 5$, the maximizing direction is unique: $\alpha = e_{\hat{y}_1}$ when $m' \geq 5$, and the straddle $\alpha = e_{y_1} + e_{\hat{y}_1}$ when $m' = 3$.

Proof. The upper bounds are Theorem 4 together with the observation that at (0, 1) every entry of the corresponding columns of Table 9 in Appendix D.1 and Tables 10–12 other than the displayed value is at most two, while Lemma 21 bounds every direction with a nonzero internal difference by two. For attainment, the table entries are exact kernel dimensions, so each large entry names a witness: $\alpha = e_{y_1}$ attains ℓ at (0, 0) and (1, 0); $\alpha = e_{\hat{y}_1}$ attains $\hat{\ell}$ at (0, 0) and (0, 1); $\alpha = e_{y_0} + e_{\hat{y}_1}$ attains $\hat{\ell}$ at (1, 0); and for $m' = 3$, $\alpha = e_{y_1} + e_{\hat{y}_1}$ attains ℓ at (0, 1) by Table 11. When $\max\{m, m'\} \geq 5$ and $(b, \hat{b}) = (0, 1)$, the displayed value is at least three and occurs in exactly one row, so the maximizer is unique. At $(m, m') = (3, 3)$ every kernel dimension is at most two and the value two occurs. $\square$

7 Proof of Conjecture 3

In Section 6 we fixed an input difference α and used the fixed-boundary attainable sets to make D_α large. Here we fix an output mask λ and read the same sets as constraints: a

direction stays in E_λ only if λ annihilates every difference the sets make attainable. Thus larger direction spaces impose more constraints on E_λ .

The directions split by how many branches carry a nonzero internal difference. With both active no nonzero mask survives at all; with one active the survivors admit an exact classification; with neither, λ is confined to a three-dimensional space.

7.1 Fixing the Mask

Let $m, m' \geq 3$ be odd, put $N = m + m'$, and assume

$$(b, \hat{b}) \in \{(0, 0), (0, 1), (1, 0)\}.$$

Put $F = \mathbb{X}_{m, m'}^{(b, \hat{b})}$ and $H = F^{-1}$. We use the convention of Section 2: α is an input difference of H , d is an output difference of H , and λ is an output mask. The forward-branch interpretation is given by the viewpoint table in that section.

Conjecture 3 of [ALPS25] states that every nonzero component of the inverse of the 4-cycle family has linear nullity at most two. In the notation of Section 6, the claim is

$$\dim E_\lambda \leq 2 \quad \text{for every } \lambda \in \mathbb{F}_2^N \setminus \{0\}. \quad (37)$$

This assertion and Theorem 4 concern the two dual slices of the same incidence relation

$$\lambda \in K_H(\alpha) \iff \alpha \in E_\lambda.$$

Theorem 4 fixes α and bounds the dimension of the admissible-mask space $K_H(\alpha)$. Conjecture 3 instead fixes λ and bounds the dimension of the admissible-direction space. The linear-height estimate of the preceding section does not by itself imply (37); a compatibility argument between different directions α is required. For a fixed mask λ , define

$$E_\lambda = \{\alpha \in \mathbb{F}_2^N : \lambda \in K_H(\alpha)\}. \quad (38)$$

Lemma 22. *For every λ , the set E_λ is a vector subspace of $\mathbb{F}_2^N$.*

Proof. By (38), E_λ is the set of linear structures of $\langle \lambda, H \rangle$, so the assertion is the classical linear-structure property recalled in Section 2.3 [Lai95, Car21]. $\square$

Lemma 1 gives $K_H(\alpha) = \text{dir}(D_\alpha)^\perp$.

7.2 Reducing to Seven Directions

A three-dimensional binary subspace has exactly seven nonzero directions. This observation yields the following seven-direction formulation of the nullity problem.

Proposition 1 (Seven-direction formulation). *The following assertions are equivalent: (i) $\dim E_\lambda \leq 2$ for every $\lambda \neq 0$; (ii) for every three-dimensional subspace $A \leq \mathbb{F}_2^N$,*

$$\bigcap_{\alpha \in A \setminus \{0\}} K_H(\alpha) = \{0\};$$

and (iii) for every such A ,

$$\sum_{\alpha \in A \setminus \{0\}} \text{dir}(D_\alpha) = \mathbb{F}_2^N.$$

Proof. By Lemma 22, the inequality $\dim E_\lambda \leq 2$ fails precisely when E_λ contains a three-dimensional subspace A . By (38), the inclusion $A \subseteq E_\lambda$ is equivalent to

$$\lambda \in \bigcap_{\alpha \in A \setminus \{0\}} K_H(\alpha).$$

This proves the equivalence of the first two assertions. Moreover,

$$\bigcap_{\alpha \in A \setminus \{0\}} K_H(\alpha) = \bigcap_{\alpha \in A \setminus \{0\}} \text{dir}(D_\alpha)^\perp = \left(\sum_{\alpha \in A \setminus \{0\}} \text{dir}(D_\alpha) \right)^\perp.$$

The last orthogonal complement is zero if and only if the sum is $\mathbb{F}_2^N$. $\square$

The analysis of fixed-boundary attainable sets and the boundary-kernel tools required to evaluate these intersections were established in Section 6. It remains to compare their kernels across the seven directions of a fixed three-dimensional subspace. The next subsection treats active-branch directions, and the final subsection treats directions supported on the four boundary coordinates.

7.3 Both Branches Active

Write the two internal input differences of H as

$$\eta_L = (\alpha_2, \dots, \alpha_{m-1}), \quad \eta_R = (\hat{\alpha}_2, \dots, \hat{\alpha}_{m'-1}).$$

The term *one-active-branch direction* means that exactly one of η_L, η_R is nonzero; a *two-active-branch direction* has both parts nonzero. The first local result shows that a full set of boundary values eliminates every internal mask coefficient on an active branch. An active branch admits every boundary quadruple, so with both branches active the boundary imposes nothing and no nonzero mask survives. A single active branch leaves more room, and its survivors have to be classified one by one.

Lemma 23. *Let $r \geq 5$ be odd and fix $(\alpha_2, \dots, \alpha_{r-1}) \neq 0$. If all sixteen boundary quadruples*

$$(d_0, d_1, \alpha_0, \alpha_1) \in \mathbb{F}_2^4$$

are admitted, then the direction spaces of the corresponding sets span every internal coordinate:

$$\sum_{(d_0, d_1, \alpha_0, \alpha_1) \in \mathbb{F}_2^4} \text{dir Aff } \mathcal{C}_r(\alpha_2, \dots, \alpha_{r-1}; d_0, d_1, \alpha_0, \alpha_1) = I_r.$$

For $r = 3$, the only nonzero internal input-difference coordinate is $\alpha_2 = 1$, each such set is a singleton, and

$$d_2 = 1 + d_0\alpha_0 + d_1\alpha_1. \tag{39}$$

Proof. Fix $\beta = (\alpha_0, \alpha_1) = 11$ and exclude $\bar{u} = 00$. Table 6 lists no nonzero generator for this profile. The exact classification in Proposition 3 therefore gives

$$V_r(\eta; 11, 00)^{\perp_{I_r}} = \{0\}, \quad V_r(\eta; 11, 00) = I_r.$$

These three sets occur in the full collection of sixteen sets, so the full collection spans I_r as well. For $r = 3$, formula (39) follows by solving (4) for d_2 . $\square$

Lemma 24. *If both internal input differences are nonzero, then*

$$K_H(\alpha) = \{0\}.$$

Proof. When both branches are active, every value of $(d_0, d_1, \hat{d}_0, \hat{d}_1) \in \mathbb{F}_2^4$ is feasible in the product decomposition (22), by Lemma 6 applied to each branch; since ϕ_L and ϕ_R are affine bijections of $\mathbb{F}_2^4$ (Lemma 19), the sixteen boundary quadruples of each branch are then all admitted.

Suppose first that $m, m' \geq 5$. Applying Lemma 23 on each branch shows that a mask in $K_H(\alpha)$ has

$$\lambda_i = 0 \quad (2 \leq i \leq m-1), \quad \hat{\lambda}_j = 0 \quad (2 \leq j \leq m'-1).$$

The projection onto $(d_0, d_1, \hat{d}_0, \hat{d}_1)$ is all of $\mathbb{F}_2^4$, so constancy of the remaining linear phase also gives

$$\lambda_0 = \lambda_1 = \hat{\lambda}_0 = \hat{\lambda}_1 = 0.$$

It remains to treat the short-cycle cases, in which $m = 3$ or $m' = 3$. Dispose first of any branch of length at least five: if $m' \geq 5$, then Lemma 23 applied to the right branch gives $\hat{\lambda}_j = 0$ for $2 \leq j \leq m' - 1$, and symmetrically if $m \geq 5$, then $\lambda_i = 0$ for $2 \leq i \leq m - 1$. Both invocations are legitimate: the internal input difference of the branch in question is nonzero by hypothesis, and all sixteen boundary quadruples are admitted as recorded above. Hence the only internal mask coefficients that can survive are λ_2 , when $m = 3$, and $\hat{\lambda}_2$, when $m' = 3$, so the phase reduces to $\lambda_2 d_2 + \hat{\lambda}_2 \hat{d}_2$ — each term present only for the corresponding short branch — plus a linear expression in $d_0, d_1, \hat{d}_0, \hat{d}_1$.

Substituting the boundary conversion (20) in (39), the quadratic terms of d_2 , as a polynomial in $(d_0, d_1, \hat{d}_0, \hat{d}_1)$, are

$$d_0 \hat{d}_1 + d_1 \hat{d}_0 + b d_0 d_1, \tag{40}$$

and the corresponding calculation for a right branch of length three gives, for $\hat{d}_2$, the quadratic terms

$$\hat{d}_0 d_0 + \hat{d}_1 d_1 + \hat{b} \hat{d}_0 \hat{d}_1. \tag{41}$$

No linear expression in $d_0, d_1, \hat{d}_0, \hat{d}_1$ contains a quadratic monomial; moreover $d_0 \hat{d}_1$ occurs in (40) but not in (41), and $d_0 \hat{d}_0$ occurs in (41) but not in (40). Neither quadratic can therefore be cancelled, against the linear part or against the other, and $\lambda_2 = \hat{\lambda}_2 = 0$. All internal mask coefficients vanish in every short-cycle case. The full boundary projection then forces the four boundary mask coefficients to vanish. $\square$

7.4 One Branch Active

The two-active-branch case leaves no nonzero mask. For a one-active-branch direction, the following classification supplies the additional phase-consistency information that is not implied by a dimension bound on an individual kernel. In each row, all input-difference and mask coordinates not listed are zero.

Lemma 25 (One-active-branch survivor classification). *Suppose the left internal input difference is nonzero and the right internal input difference is zero.*

If $m \geq 5$, then $K_H(\alpha) \neq \{0\}$ exactly in the rows of Table 7 in Appendix C. In every listed case, $K_H(\alpha)$ is the line generated by the mask specified in the last column.

If $m = 3$, then necessarily $\alpha_2 = 1$. For every choice of the four boundary input-difference coordinates, the nonzero kernel is the line whose generator satisfies

$$\begin{aligned} \lambda_2 &= 1, \\ \lambda_0 &= \alpha_0 + b\hat{\alpha}_1, \quad \lambda_1 = \alpha_1 + b\hat{\alpha}_0, \\ \hat{\lambda}_0 &= \hat{\alpha}_1, \quad \hat{\lambda}_1 = \hat{\alpha}_0, \end{aligned} \tag{42}$$

with all right internal mask coefficients equal to zero.

Suppose instead that the left internal input difference is zero and the right internal input difference is nonzero. If $m' \geq 5$, the complete list is Table 8 in Appendix C; again each nonzero kernel is the line generated by the mask in the last column.

If $m' = 3$, then necessarily $\hat{\alpha}_2 = 1$, and the generator of the nonzero kernel satisfies

$$\begin{aligned}\hat{\lambda}_2 &= 1, \\ \lambda_0 &= \alpha_1, & \lambda_1 &= \alpha_0, \\ \hat{\lambda}_0 &= \hat{\alpha}_0 + \hat{b}\alpha_0, & \hat{\lambda}_1 &= \hat{\alpha}_1 + \hat{b}\alpha_1,\end{aligned}\tag{43}$$

with all left internal mask coefficients equal to zero. There are no other one-active-branch directions with a nonzero kernel.

Proof. Suppose first that the active branch has length at least five. The one-active estimate in Lemma 21 gives $\dim K_H(\alpha) \leq 1$. The zero-branch supports and the two boundary maps (21) then select the applicable local profiles from Table 6; the left and right maps are treated separately in Subsection C.2, without an assumed branch symmetry. The length-five and length-seven instances are established there by exact base calculations. From length seven onward, Proposition 5 supplies every two-coordinate update and the boundary coefficients recorded in Tables 7 and 8. Its no-revival and unique-predecessor parts exclude every unlisted case. Thus each displayed nonzero mask spans the whole kernel, and every other boundary case has zero kernel.

For an active branch of length three, (39) replaces the length-five base case: each such set is a singleton, so the mask condition reduces to constancy of $\lambda_2 d_2$ plus a linear expression in d_{bd} — on T_α , not on all of $\mathbb{F}_2^4$. That restriction is what makes a mask with $\lambda_2 = 1$ possible. On $\mathbb{F}_2^4$ the quadratic part (40) of d_2 is not affine, whereas on T_α it satisfies

$$d_0 \hat{d}_1 + d_1 \hat{d}_0 + b d_0 d_1 = (1 + b\hat{\alpha}_1)d_0 + (1 + b\hat{\alpha}_0)d_1 + \hat{\alpha}_1 \hat{d}_0 + \hat{\alpha}_0 \hat{d}_1 + c \tag{44}$$

for a constant c , for both admissible values of b and each of the sixteen boundary input-difference tuples $(\alpha_0, \alpha_1, \hat{\alpha}_0, \hat{\alpha}_1)$. Substituting (20) in (39) and reducing by (44) gives (42); the mirror computation gives (43). The right-hand side of (44) depends on b and on $\hat{\alpha}_0, \hat{\alpha}_1$, while T_α depends on $\hat{b}$ and on $\hat{\alpha}_0, \hat{\alpha}_1$; the reduction is therefore uniform in m' . The four boundary mask coefficients are determined by the four boundary input-difference coordinates, and the rank estimate above shows that each displayed line is the entire kernel. $\square$

Corollary 3. *Let $\lambda \neq 0$.*

1. *No element of E_λ is a two-active-branch direction.*
2. *If E_λ contains one-active-branch directions, then all of them use the same branch and the same nonzero internal input difference.*
3. *If E_λ contains an active direction, then*

$$\dim(E_\lambda \cap \mathcal{B}_{\text{out}}) \leq 1, \tag{45}$$

where

$$\mathcal{B}_{\text{out}} = \{\alpha : \alpha_i = 0 \ (2 \leq i \leq m-1), \ \hat{\alpha}_j = 0 \ (2 \leq j \leq m'-1)\}.$$

Proof. The first assertion is Lemma 24. The left- and right-active classifications have nonzero internal mask parts on disjoint branches, so a fixed nonzero mask cannot occur on both sides. For a branch of length three, (42) or (43) determines the boundary coordinates injectively. For a long branch, the only repeated masks in Tables 7 and 8 arise from boundary tuple 0000 and its listed repeated tuple; their internal input difference is the

same, while the endpoint mask coefficients remove all other apparent coincidences. This proves the second assertion.

Let $\alpha \in E_\lambda$ be active and let $\alpha' \in E_\lambda \cap \mathcal{B}_{\text{out}}$. By Lemma 22, $\alpha + \alpha' \in E_\lambda$, and it has the same internal input difference as α . The survivor tables permit at most two boundary tuples for that internal difference: 0000 and one fixed repeated tuple. All possible values of α' lie in one binary line, proving (45). $\square$

7.5 Boundary-Only Directions

It remains to compare the exact kernels for directions in $\mathcal{B}_{\text{out}}$. This comparison uses the boundary linear system of Lemma 20; no new calculation of these sets is needed.

Identify $\mathcal{B}_{\text{out}}$ with $\mathbb{F}_2^4$ through its four free coordinates $(\alpha_0, \alpha_1, \hat{\alpha}_0, \hat{\alpha}_1)$, and let $c = (c_0, c_1, \hat{c}_0, \hat{c}_1) \in \mathbb{F}_2^4 \setminus \{0\}$. The vector c is not an input difference, an output difference, a mask, or a parameter of the 4-cycle construction; it is the coefficient vector of a nonzero linear equation on the boundary coordinates. Write A_c for the solution space of

$$c_0\alpha_0 + c_1\alpha_1 + \hat{c}_0\hat{\alpha}_0 + \hat{c}_1\hat{\alpha}_1 = 0. \quad (46)$$

Since the equation is nonzero, A_c has dimension three. Conversely, every three-dimensional subspace of $\mathcal{B}_{\text{out}}$ is the kernel of a nonzero linear equation, whose coefficient vector is unique up to nonzero scalar multiplication. Since the only nonzero scalar in $\mathbb{F}_2$ is one, the $2^4 - 1 = 15$ nonzero choices of c give exactly the fifteen three-dimensional boundary subspaces. For example, $c = 1000$ gives $\alpha_0 = 0$, whereas $c = 1100$ gives $\alpha_0 + \alpha_1 = 0$.

The complete witness classification is given in Tables 13–16 of Appendix E. The four tables cover the long–long, short–long, long–short, and short–short geometries, respectively; they are obtained by exact binary row reduction rather than finite-instance experimentation.

Lemma 26 (Boundary spanning). *For every three-dimensional subspace $A \leq \mathcal{B}_{\text{out}}$,*

$$\bigcap_{\alpha \in A \setminus \{0\}} K_H(\alpha) = \{0\}. \quad (47)$$

It follows that, for every nonzero λ ,

$$\dim(E_\lambda \cap \mathcal{B}_{\text{out}}) \leq 2. \quad (48)$$

Proof. Let $A \leq \mathcal{B}_{\text{out}}$ have dimension three. By the preceding parametrization, $A = A_c$ for a unique nonzero coefficient vector c . For the fixed parameter pair and branch-length geometry, insert the one or two directions printed in the row indexed by c of Tables 13–16 in Appendix E into the exact system (27). The anchors and direction spaces are those of Table 4. Exact binary row reduction gives

$$\bigcap_{\text{directions printed in the row}} K_H(\alpha) = \{0\}.$$

Every printed direction is nonzero and satisfies (46), so it belongs to A_c . The intersection over all nonzero directions in A_c is contained in the displayed zero intersection. This proves (47).

For each direction, the reduction contains at most sixteen product sets. Across the four tables it covers all fifteen coefficient tuples, all three admissible parameter pairs, and all four short/long geometries. For odd branch lengths at least five, the required affine hulls follow from the all-length elimination and upper-triangular witnesses in Appendix B; branch length three is treated separately. The reductions apply to every odd pair (m, m') .

It remains to derive the stated dimension bound. Suppose, to the contrary, that $\lambda \neq 0$ and $\dim(E_\lambda \cap \mathcal{B}_{\text{out}}) \geq 3$. By Lemma 22, this intersection is a vector space and contains a three-dimensional subspace A_c . For every $\alpha \in A_c$, the inclusion $\alpha \in E_\lambda$ and the definition of E_λ imply $\lambda \in K_H(\alpha)$. Hence λ belongs to the intersection in (47), which is $\{0\}$. This contradicts $\lambda \neq 0$ and proves (48). $\square$

7.6 The Nullity Theorem

Combining the boundary estimate with the compatibility of active directions now gives the required nullity bound.

Theorem 5 (Conjecture 3). *Conjecture 3 of [ALPS25] holds. More precisely, let $m, m' \geq 3$ be odd and $(b, \hat{b}) \in \{(0, 0), (0, 1), (1, 0)\}$. For*

$$H = \left(\mathbb{X}_{m, m'}^{(b, \hat{b})} \right)^{-1}$$

and every nonzero mask $\lambda \in \mathbb{F}_2^N$,

$$\dim E_\lambda \leq 2. \quad (49)$$

Equivalently, every nonzero component of H has at most four linear structures.

Proof. Fix $\lambda \neq 0$ and consider the subspace E_λ from (38). If it contains no direction with a nonzero internal input-difference coordinate, then

$$E_\lambda \subseteq \mathcal{B}_{\text{out}}.$$

Lemma 26 gives $\dim E_\lambda \leq 2$.

If E_λ contains an active direction, then by Corollary 3, every active element has the same nonzero internal input difference. Hence the image of E_λ under the projection

$$\alpha \mapsto (\alpha_2, \dots, \alpha_{m-1}, \hat{\alpha}_2, \dots, \hat{\alpha}_{m'-1})$$

has dimension one. The kernel of this projection on E_λ is $E_\lambda \cap \mathcal{B}_{\text{out}}$, whose dimension is at most one by (45). Rank-nullity now yields

$$\dim E_\lambda = 1 + \dim(E_\lambda \cap \mathcal{B}_{\text{out}}) \leq 2.$$

This proves (49); since E_λ is a binary vector space, it contains at most $2^2 = 4$ directions. $\square$

Corollary 4. *Under the assumptions of Theorem 5, every nonzero $\lambda \in \mathbb{F}_2^N$ satisfies*

$$\#\{\alpha \in \mathbb{F}_2^N : \text{Bias}(\alpha, \lambda) = \pm 1\} \leq 4.$$

Moreover, every three-dimensional subspace $A \leq \mathbb{F}_2^N$ satisfies

$$\sum_{\alpha \in A \setminus \{0\}} \text{dir}(D_\alpha) = \mathbb{F}_2^N.$$

Proof. The full-bias condition is equivalent to $\lambda \in K_H(\alpha)$, so the set being counted is precisely E_λ . Its size is $2^{\dim E_\lambda} \leq 4$. The final assertion follows from Proposition 1. $\square$

8 Counterexamples to Conjecture 5

Conjecture 5 asserts $\dim K_H(\alpha) \leq 3$ for the $(0, 1)$ variant of the 4-cycle family whenever $|m - m'| \leq 2$. At $(m, m') = (5, 7)$ there is a direction of height 4, and at $(7, 9)$ one of height 5, both inside that range.

We now examine the 4-cycle instance $\mathbb{X}_{m, m'}^{(0, 1)}$. In the notation used throughout this paper, its parameters are $(b, \hat{b}) = (0, 1)$. For each pair (m, m') below, we construct the complete table of the forward permutation, invert it, and enumerate every nonzero direction α . For each such α , we form

$$S_\alpha = \{H(y + \alpha) + H(y) : y \in \mathbb{F}_2^N\}, \quad H = (\mathbb{X}_{m, m'}^{(0, 1)})^{-1},$$

Table 2: Exhaustive linear-height computation for the 4-cycle family with $(b, \hat{b}) = (0, 1)$. The displayed direction attains the maximum over all nonzero directions.

(m, m')	N	maximizing direction α	$\max_{\alpha \neq 0} \dim K_H(\alpha)$
$(5, 7)$	12	e_6	4
$(7, 5)$	12	e_8	3
$(7, 9)$	16	e_8	5
$(9, 7)$	16	e_{10}	4

and compute $\dim K_H(\alpha) = N - \dim D_\alpha$ using Lemma 1. Concretely, after choosing one anchor in S_α , the vectors obtained by adding that anchor to every element of S_α are placed in a binary matrix. Gaussian elimination over $\mathbb{F}_2$ gives $\dim D_\alpha$. The calculation covers every $y \in \mathbb{F}_2^N$ and every $\alpha \neq 0$. Here e_i denotes the i -th standard unit direction.

The rows $(5, 7)$, $(7, 9)$, and $(9, 7)$ have linear height greater than 3, so they are counterexamples to Conjecture 5. The asymmetric case $(7, 5)$ is included as a control: it attains, but does not exceed, the proposed bound 3. In particular, Conjecture 5 cannot hold for all odd pairs with $|m - m'| \leq 2$.

The failure is not repaired by replacing 3 with another absolute constant. Theorem 4 of Section 6 establishes the length-dependent bound

$$\dim K_H(\alpha) \leq \frac{\max\{m, m'\} + 1}{2}$$

for all odd $m, m' \geq 3$ and every admissible $(b, \hat{b}) \in \{(0, 0), (0, 1), (1, 0)\}$. Its scope is therefore strictly broader than that of Conjecture 5, which concerns only $(b, \hat{b}) = (0, 1)$ and $|m - m'| \leq 2$. Thus the bound of Conjecture 4 is the correct length-dependent replacement for the refuted constant bound.

Moreover, Corollary 2 determines the exact maximum. For the $(0, 1)$ variant it is

$$\max_{\alpha \neq 0} \dim K_H(\alpha) = \begin{cases} (m' + 1)/2, & m' \geq 5, \\ (m + 1)/2, & m' = 3. \end{cases}$$

Consequently, the maximum grows with the relevant branch length—even within the parameter range of Conjecture 5—so no length-independent constant bound can hold. The values 4 and 3 at $(m, m') = (5, 7)$ and $(7, 5)$, respectively, also illustrate the orientation dependence of the $(0, 1)$ construction.

Remark 2 (Erratum to [ALPS25, Table 1]). Table 1 of [ALPS25] lists experimental data for $m \leq m'$ only, on the stated grounds that the properties agree with those of the reversed ordering over the parameter range considered; in particular the maximum linear height 3 is reported for $(m, m') = (5, 7)$ and, implicitly, for $(7, 5)$. The exhaustive computation of Table 2 shows that the two orderings in fact differ: the maximum is 4 at $(5, 7)$ and 3 at $(7, 5)$. The authors of [ALPS25] have confirmed that the value reported there for $(5, 7)$ is incorrect (personal communication, August 2026). The discrepancy is exactly the orientation dependence discussed above, made precise by Corollary 2.

9 Design Implications for ChiLow

We refuted a constant bound in Section 8 because the maximum linear height grows with the branch lengths. Read the other way, that growth is something a designer controls. If the maximum depends on one branch length and not the other, then the split itself is a free parameter, and we can choose it to make the maximum small.

The ChiChi permutation deployed in ChiLow [BDG⁺25] on $N = 2M$ bits is, up to permutation equivalence, the instance $(m, m', b, \hat{b}) = (M - 1, M + 1, 0, 1)$ of the 4-cycle family [ALPS25]. Linear height is invariant under composition with bit permutations, so Corollary 2 applies to the deployed instances verbatim. Two consequences for the design choice follow, together with an assessment of the effect of imbalance.

Branch ordering. For $(b, \hat{b}) = (0, 1)$ and $m' \geq 5$ the exact maximum $(m' + 1)/2$ depends only on m' . Reversing the deployed ordering, so that the longer branch comes first as $(m, m') = (M + 1, M - 1)$, lowers the exact maximum linear height from $(M + 2)/2$ to $M/2$ for every even $M \geq 6$: from 9 to 8 at $N = 32$, and from 11 to 10 at $N = 40$. Corollary 2 gives the exact ordering-dependent maximum and therefore proves this asymmetry, despite the symmetric upper bound in Theorem 4. The fact that the two orderings of $\chi_{m, m'}^{(0, 1)}$ are in general not EA-equivalent [ALPS25] is a separate structural distinction; Table 2 shows that their linear heights already differ at $N = 12$. An exhaustive sweep of all nonzero directions confirms both values at $N = 32$, with the maximum attained by the single direction $e_{\hat{y}_1}$ in each case.

Branch balance. More is gained by dropping balance altogether. Since the maximum is independent of m when $m' \geq 5$, minimizing it over the odd splits of a fixed N means minimizing m' . The choice $m' = 5$ achieves maximum linear height 3 for every even $N \geq 8$, for instance $(27, 5)$ at $N = 32$ and $(35, 5)$ at $N = 40$, against 9 and 11 for the deployed balanced splits. The boundary case $m' = 3$ must be excluded, since there the maximum flips to $(m + 1)/2$, and $(29, 3)$ attains 15 at $N = 32$. The cycle lengths of ChiChi were chosen “with a constraint to split as equally as possible” [BDG⁺25]; under that constraint the maximum grows linearly with the block size and is approximately $N/4$, whereas the unbalanced split $m' = 5$ pins it at the constant 3. An exhaustive computation at $N = 32$ confirms the optimum, with the maximum over all nonzero directions of $(27, 5)$ equal to exactly 3, and explicit witnesses attain the value of Corollary 2 at every odd split of 32.

The cost of imbalance. The algebraic degree of the inverse is split-independent. Every nonzero component of $(\chi_{m, m'}^{(b, \hat{b})})^{-1}$ has degree $(m + m')/2$ [ALPS25, Theorem 4], so a short branch does not lower the degree of the inverse of the full map. Note that we do not claim a complete security assessment of strongly unbalanced splits: beyond the full-bias masks determined by the kernel classification (Corollary 1), differential and correlation properties, and the behavior of the short branch under structural attacks, require a separate analysis.

10 The Inverse of a Single Chi Cycle

The two families are settled. We now record what the same local analysis says about a single ordinary cycle, which is a question of independent interest. It is the benchmark against which [ALPS25] measures both families, recorded there numerically for small dimensions and attributed to the analysis of χ_n^{-1} in [LSMI22]. We settle it exactly, for every odd cycle length, with both equality loci classified. Everything below uses only Subsections 3.2 and 3.3; the notation is that of Section 3.

10.1 The Classification

For $r \geq 5$ only two families of directions give anything but the full space, and they are the only directions along which χ_r^{-1} has a linear structure at all. At $r = 3$ the cycle is too short for that dichotomy and every direction is exceptional.

Theorem 6 (Affine hulls of the inverse single cycle). *Let $r \geq 3$ be odd and let $\alpha \in \mathbb{F}_2^r$ be nonzero.*

1. *If $r \geq 5$, then*

$$\text{dir}(D_\alpha) = \begin{cases} \{e_j\}^\perp, & \alpha = e_j, \\ \{e_j + e_{j+2}\}^\perp, & \alpha = e_{j-1} + e_j, \\ \mathbb{F}_2^r & \text{otherwise.} \end{cases}$$

2. *If $r = 3$, then $\text{dir}(D_\alpha) = \{\alpha\}^\perp$, so $K_{H_3}(\alpha) = \{0, \alpha\}$ and dually $E_v = \{0, v\}$ for every nonzero v .*

The two parts agree where they overlap. At $r = 3$ the index $j + 2$ equals $j - 1$, so $e_j + e_{j+2} = e_{j-1} + e_j$, and both exceptional families of part (1) already give $\{\alpha\}^\perp$. The short cycle departs from the generic pattern at the single direction $\alpha = \mathbf{1}$, where part (1) would predict $\mathbb{F}_2^3$ and the annihilator is instead $\{0, \mathbf{1}\}$.

The proof identifies S_α with a column of the differential distribution table of the forward map χ_r , solves the resulting linear systems by a run analysis of the difference pattern, and pins the possible annihilating masks by evaluating one explicit family of attainable differences.

10.2 Pinning the Annihilating Masks

The next proposition shows that $\text{dir}(D_\alpha)^\perp$ is contained in the line predicted by Theorem 6. Its proof only evaluates the admissible intervals of Lemma 3 against a putative annihilating mask.

Proposition 2. *Let $r \geq 5$ be odd and $\alpha \neq 0$. If $c \perp \text{dir}(D_\alpha)$, then*

$$c \in \begin{cases} \{0, e_j\}, & \alpha = e_j, \\ \{0, e_j + e_{j+2}\}, & \alpha = e_{j-1} + e_j, \\ \{0\} & \text{otherwise.} \end{cases}$$

Proof. Since c annihilates the direction space of the affine hull of S_α , the value $\langle c, d \rangle$ is a constant ϵ over $d \in S_\alpha$. In particular $\langle c, I \rangle = \epsilon + \langle c, \mathbf{1} \rangle =: \epsilon'$ is constant over all admissible interval indicators I .

Suppose first $\alpha = \mathbf{1}$. By Lemma 3, the admissible intervals are $I_{q,0} = 0$ (as $\text{wt}(\mathbf{1}) = r$ is odd) and $I_{q,2}$ for every q . From $\langle c, 0 \rangle = 0$ we get $\epsilon' = 0$, so $c_q + c_{q+1} = 0$ for every q , forcing $c \in \{0, \mathbf{1}\}$. To exclude $c = \mathbf{1}$, note that $d^* = \mathbf{1} + e_0 + e_1 + e_3 \in S_\mathbf{1}$ by Theorem 1. Its zero set is the run $[0, 1]$ together with the isolated zero 3. Condition (i) applies only at $i = r - 1$, where $d_{r-1}^* = 1 = \alpha_{r-1}$ since $r \geq 5$, and condition (ii) at $q = 3$ reads $d_1^* = 0 = \alpha_1 + \alpha_2$. Then $\langle \mathbf{1}, d^* \rangle = r - 3$ is even while $\langle \mathbf{1}, \mathbf{1} \rangle = r$ is odd, so $\langle \mathbf{1}, \cdot \rangle$ is not constant on S_α . Hence $c = 0$.

Now let $\alpha \neq \mathbf{1}$, so that α has at least one zero coordinate. Decompose the support of α into maximal one-runs $[s_t, e_t]$, indexed cyclically by t , separated by zero-gaps of lengths $z_t = s_{t+1} - e_t - 1 \geq 1$. Lemma 3 provides the following admissible intervals and hence the following constraints on c :

(P1) $g = 1$ at $q = s_t + 1$, since $\alpha_{s_t-1} = 0 \neq 1 = \alpha_{s_t}$; hence $c_{s_t+1} = \epsilon'$.

(P2) $g = 1$ at $q = e_t + 2$, since $\alpha_{e_t} = 1 \neq 0 = \alpha_{e_t+1}$; hence $c_{e_t+2} = \epsilon'$.

(P3) $g = 2$ at $q = j + 1$ for every j in the support of α ; hence $c_{j+1} + c_{j+2} = \epsilon'$.

(P4) g with $2 \leq g \leq \min(z_t + 2, r - 1)$ at $q = e_t + 1$. Consecutive values of g differ by the single coordinate e_{e_t+g} , so $c_i = 0$ for $i \in [e_t + 3, e_t + \min(z_t + 2, r - 1)]$.

Suppose $\epsilon' = 0$. By (P1) and the chain (P3) along each run, $c_i = 0$ for $i \in [s_t + 1, e_t + 2]$; by (P4), $c_i = 0$ on $[e_t + 3, e_t + z_t + 2]$ when the cap does not interfere, and in the capped cases below the same conclusion is checked directly. When the cap is inactive, $e_t + z_t + 2 = s_{t+1} + 1$, so the intervals $[s_t + 1, e_t + 2]$ and $[e_t + 3, s_{t+1} + 1]$ together cover $\mathbb{Z}_r$ and $c = 0$. The cap is active only when $z_t \geq r - 2$, which forces a single one-run of length at most two; for $\alpha = e_j$ the constraints (P1)–(P3) read $c_{j+1} = c_{j+2} = 0$ and (P4) gives $c_i = 0$ for $i \in [j + 3, j - 1]$, so $c \in \{0, e_j\}$; for $\alpha = e_{j-1} + e_j$, (P1) and (P3) give $c_j = c_{j+1} = c_{j+2} = 0$ and (P4) gives $c_i = 0$ for $i \in [j + 3, j - 1]$, so $c = 0$.

Suppose $\epsilon' = 1$. If α has two or more one-runs, then every gap satisfies $z_t \leq r - 3$, the cap in (P4) is inactive, and (P4) gives $c_{s_{t+1}+1} = 0$, contradicting (P1) at run $t + 1$. If α is a single one-run of length $\ell \geq 3$, then $z = r - \ell \leq r - 3$, the cap is again inactive, and the same contradiction applies. If $\alpha = e_j$, then (P1) and (P2) give $c_{j+1} = c_{j+2} = 1$ while (P3) at j gives $c_{j+1} + c_{j+2} = 1$, a contradiction. The only remaining case is the single run of length two, $\alpha = e_{j-1} + e_j$; there (P1), (P2), and (P3) give $c_j = 1$, $c_{j+2} = 1$, $c_{j+1} = 0$, and (P4) with $z = r - 2$ gives $c_i = 0$ for $i \in [j + 3, j - 1]$, so $c = e_j + e_{j+2}$. $\square$

It remains to show that the two lines permitted by Proposition 2 do annihilate the respective direction spaces.

Lemma 27. $d_j = 1$ for every $d \in S_{e_j}$.

Proof. Put $\alpha = e_j$. Suppose $d \in S_{e_j}$ with $d_j = 0$; then $d \neq 0$ by Theorem 1. A zero-run of d of length at least two starts, by condition (i) at $i = q - 1$, at an index q with $\alpha_{q-1} = 1$, hence at $q = j + 1$; such a run cannot contain j without being all of $\mathbb{Z}_r$, which is the excluded $d = 0$. Hence j is an isolated zero of d , and condition (ii) gives $d_{j-2} = \alpha_{j-2} + \alpha_{j-1} = 0$. The same argument applies at $j - 2$, then at $j - 4$, and so on; since r is odd, the orbit of -2 is all of $\mathbb{Z}_r$, so every coordinate of d vanishes, a contradiction. $\square$

Lemma 28. $d_j + d_{j+2} = 1$ for every $d \in S_{e_{j-1}+e_j}$.

Proof. Put $\alpha = e_{j-1} + e_j$. Then $d \neq 0$, and $d \neq \mathbf{1}$ since $\text{wt}(\alpha) = 2$ is even (Theorem 1). The local test has two consequences. The right-hand side $\alpha_{q-2} + \alpha_{q-1}$ of condition (ii) equals one exactly for $q \in \{j, j + 2\}$, and a zero-run of length at least two must start at $q = j$ — where $\alpha_j = 1$ at its second index forces length exactly two — or at $q = j + 1$.

Suppose $d_j = d_{j+2} = 0$. A zero-run of length at least two containing j must start at j , since a start at $j + 1$ reaches j only after the full circle, giving the excluded $d = 0$. A run starting at j has length exactly two, so it is the maximal run $[j, j + 1]$ and forces $d_{j+2} = 1$, contradicting the assumption. Hence j is an isolated zero, in particular $d_{j+1} = 1$, and condition (ii) gives $d_{j-2} = 1$. The zero-run containing $j + 2$ can start neither at j , whose run ends at $j + 1$, nor at $j + 1$, where $d_{j+1} = 1$; so $j + 2$ is isolated as well, and condition (ii) gives $d_j = 1$, a contradiction.

Suppose $d_j = d_{j+2} = 1$. No zero-run of length at least two exists, since a start at j is excluded by $d_j = 1$ and a run starting at $j + 1$ contains $j + 2$, excluded by $d_{j+2} = 1$. So every zero of d is isolated, and condition (ii) at any isolated zero $q \notin \{j, j + 2\}$ gives $d_{q-2} = 0$. Since $d \neq \mathbf{1}$, some zero exists, and iterating along the -2 orbit — all of $\mathbb{Z}_r$, as r is odd — eventually asserts $d_{j+2} = 0$ or $d_j = 0$, a contradiction. $\square$

10.3 Proof of the Classification

Proof of Theorem 6. Fix $\alpha \neq 0$ and put $V = \text{dir}(D_\alpha)$. If $\alpha \notin \{e_j\}_j \cup \{e_{j-1} + e_j\}_j$, then Proposition 2 gives $V^\perp = \{0\}$, that is, $V = \mathbb{F}_2^r$. If $\alpha = e_j$, then Lemma 27 shows that $\langle e_j, \cdot \rangle \equiv 1$ on S_α , so $e_j \in V^\perp$. Together with Proposition 2, this gives $V^\perp = \{0, e_j\}$ and $V = \{e_j\}^\perp$. If $\alpha = e_{j-1} + e_j$, then Lemma 28 shows that $\langle e_j + e_{j+2}, \cdot \rangle \equiv 1$ on S_α , so $e_j + e_{j+2} \in V^\perp$. Together with Proposition 2, this gives $V^\perp = \{0, e_j + e_{j+2}\}$ and $V = \{e_j + e_{j+2}\}^\perp$.

For part (2), let $r = 3$, so that indices are read modulo 3 and $q - 2 = q + 1$ for every q . Here Theorem 1 determines S_α outright. The zero difference is excluded because $\alpha \neq 0$, and $\mathbf{1} \in S_\alpha$ exactly when $\text{wt}(\alpha)$ is odd. A difference of weight one, say e_p , has the single zero-run $\{p + 1, p + 2\}$ of length two and no isolated zero, so only condition (i) applies, at $i = p$, where it reads $\alpha_p = 1$. A difference of weight two, say $\mathbf{1} + e_q$, has the isolated zero q and no zero-run of length two, so only condition (ii) applies, and it reads $\alpha_{q+1} + \alpha_{q+2} = 1$.

Every difference that survives satisfies $\langle \alpha, d \rangle = 1$. For $d = e_p$ this is the condition $\alpha_p = 1$; for $d = \mathbf{1} + e_q$ it is $\text{wt}(\alpha) + \alpha_q = \alpha_{q+1} + \alpha_{q+2} = 1$; and for $d = \mathbf{1}$ it is $\text{wt}(\alpha) = 1$. So S_α lies in the affine hyperplane $\{d : \langle \alpha, d \rangle = 1\}$, which has four elements. It exhausts that hyperplane. If $\text{wt}(\alpha) = 1$ the members are $\mathbf{1}$, the one e_p with $\alpha_p = 1$, and the two $\mathbf{1} + e_q$ with q outside the support; if $\text{wt}(\alpha) = 2$ they are the two e_p in the support and the two $\mathbf{1} + e_q$ with q in the support; if $\text{wt}(\alpha) = 3$ they are $\mathbf{1}$ and the three e_p . Each count is four, so $S_\alpha = \{d : \langle \alpha, d \rangle = 1\}$ and $\text{dir}(D_\alpha) = \{\alpha\}^\perp$. $\square$

Corollary 5 (Height and nullity of χ_r^{-1}). *Let $r \geq 5$ be odd. For every nonzero direction α and every nonzero mask v ,*

$$\dim K_{H_r}(\alpha) \leq 1, \quad \dim E_v \leq 1.$$

Exactly, $K_{H_r}(e_j) = \{0, e_j\}$ and $K_{H_r}(e_{j-1} + e_j) = \{0, e_j + e_{j+2}\}$, while $K_{H_r}(\alpha) = \{0\}$ for all other $\alpha \neq 0$; dually, $E_{e_j} = \{0, e_j\}$ and $E_{e_j + e_{j+2}} = \{0, e_{j-1} + e_j\}$, while $E_v = \{0\}$ for all other $v \neq 0$. Both the directions' linear-height distribution and the components' linear-nullity distribution of χ_r^{-1} equal $\{1^{2r}, 0^{2^r-1-2r}\}$.

Proof. The height statements follow from Theorem 6 through (1), since the orthogonal complement of a coordinate hyperplane $\{c\}^\perp$ is the line $\{0, c\}$ and the complement of $\mathbb{F}_2^r$ is $\{0\}$. For the nullity statements, fix $v \neq 0$; by (1), a nonzero α lies in E_v exactly when $v \in K_{H_r}(\alpha)$, which by the height classification forces $v = e_j$ with $\alpha = e_j$, or $v = e_j + e_{j+2}$ with $\alpha = e_{j-1} + e_j$. The $2r$ listed masks are pairwise distinct. The two families have weights one and two, the first is injectively indexed by j , and $e_j + e_{j+2} = e_{j'} + e_{j'+2}$ with $j \neq j'$ would force $j' = j + 2$ and $j = j' + 2$, hence $4 \equiv 0 \pmod{r}$, impossible for odd r . Each E_v is therefore one of the listed lines or $\{0\}$, and the $2r$ directions of height one are likewise pairwise distinct. $\square$

The case $r = 3$ is part (2) of Theorem 6, where all seven nonzero directions have height one and all seven nonzero components have nullity one, so the bounds are unchanged and only the equality loci differ.

Remark 3. The corollary rests on no finite verification. Every step is an exact statement about the cyclic system (3), valid uniformly in r . It therefore settles, with classification, the base case that [ALPS25] records numerically on the strength of [LSMI22], and against which Theorems 4 and 5 are measured. Coupling two cycles through the 4-cycle boundary at most doubles the nullity of the inverse, one per branch, while the SWAP coupling allows the inverse invariants to grow with the branch length.

11 Conclusion

In this paper, we settle all five conjectures of [ALPS25]. Conjectures 1–4 hold at every admissible branch length, and Conjecture 5 already fails at $(m, m') = (5, 7)$. We show that no absolute constant repairs it. The exact maximum grows with the branch lengths, so the length-dependent bound of Conjecture 4 is the correct statement. Corollary 2 gives, for each variant, the maximal linear height and the directions attaining it. For $(0, 0)$ and $(1, 0)$ the bound is reached at every admissible pair; for $(0, 1)$, exactly when $m \leq m'$ or $m' = 3$. We turn the same corollary into a criterion for choosing branch lengths (Section 9).

We have treated the inverse nonlinear layer on its own. Natural next steps are sharper bounds for individual input-difference directions; an analysis of how these single-layer properties survive linear layers, round constants, and multiple rounds; and an assessment of the differential and correlation properties of strongly unbalanced splits.

Acknowledgments

We thank Enrico Piccione for his useful comments in an earlier version of this draft.

References

- [ALPS25] Samuele Andreoli, Gregor Leander, Enrico Piccione, and Lukas Stennes. Generalizations of ChiChi: Families of low-latency permutations in any even dimension. *IACR Transactions on Symmetric Cryptology*, 2025(3):800–826, 2025. doi:[10.46586/tosc.v2025.i3.800-826](https://doi.org/10.46586/tosc.v2025.i3.800-826).
- [Axl24] Sheldon Axler. *Linear Algebra Done Right*. Undergraduate Texts in Mathematics. Springer, 4th edition, 2024. doi:[10.1007/978-3-031-41026-0](https://doi.org/10.1007/978-3-031-41026-0).
- [Bal15] Simeon Ball. *Finite Geometry and Combinatorial Applications*, volume 82 of *London Mathematical Society Student Texts*. Cambridge University Press, 2015. doi:[10.1017/CB09781316257449](https://doi.org/10.1017/CB09781316257449).
- [BDG⁺25] Yanis Belkheyar, Patrick Derbez, Shibam Ghosh, Gregor Leander, Silvia Mella, Léo Perrin, Shahram Rasoolzadeh, Lukas Stennes, Siwei Sun, Gilles Van Assche, and Damian Vizár. Chilow and chichi: New constructions for code encryption. In *Advances in Cryptology – EUROCRYPT 2025*, Lecture Notes in Computer Science. Springer, 2025. doi:[10.1007/978-3-031-91107-1_8](https://doi.org/10.1007/978-3-031-91107-1_8).
- [Car21] Claude Carlet. *Boolean Functions for Cryptography and Coding Theory*. Cambridge University Press, 2021. doi:[10.1017/9781108606806](https://doi.org/10.1017/9781108606806).
- [Dae95] Joan Daemen. *Cipher and Hash Function Design: Strategies Based on Linear and Differential Cryptanalysis*. PhD thesis, Katholieke Universiteit Leuven, March 1995. URL: https://cs.ru.nl/~joan/papers/JDA_Thesis_1995.pdf.
- [DEMS21] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl  ffer. Ascon v1.2: Lightweight authenticated encryption and hashing. *Journal of Cryptology*, 34(3):33, 2021. doi:[10.1007/s00145-021-09398-9](https://doi.org/10.1007/s00145-021-09398-9).
- [DMMR20] Joan Daemen, Pedro Maat Costa Massolino, Alireza Mehrdad, and Yann Rotella. The subterranean 2.0 cipher suite. *IACR Transactions on Symmetric Cryptology*, 2020(S1):262–294, 2020. doi:[10.13154/tosc.v2020.iS1.262-294](https://doi.org/10.13154/tosc.v2020.iS1.262-294).
- [Lai95] Xuejia Lai. Additive and linear structures of cryptographic functions. In *Fast Software Encryption (FSE) 1994*, volume 1008 of *Lecture Notes in Computer Science*, pages 75–85. Springer, 1995. doi:[10.1007/3-540-60590-8_6](https://doi.org/10.1007/3-540-60590-8_6).
- [LSMI22] Fukang Liu, Santanu Sarkar, Willi Meier, and Takanori Isobe. The inverse of χ and its applications to Rasta-like ciphers. *Journal of Cryptology*, 35(4):28, 2022. doi:[10.1007/s00145-022-09439-x](https://doi.org/10.1007/s00145-022-09439-x).

-
- [NIS15] SHA-3 standard: Permutation-based hash and extendable-output functions. Federal Information Processing Standards Publication 202, 2015. [doi:10.6028/NIST.FIPS.202](#).
- [NIS25] Ascon-based lightweight cryptography standards for constrained devices. NIST Special Publication 800-232, 2025. [doi:10.6028/NIST.SP.800-232](#).

A Finite-Instance Verification

This appendix records the proof-critical finite verification for the shortest branch-length pair $(m, m') = (3, 3)$. It establishes the exceptional six-dimensional case in the proof of Theorem 3; it is not an independent consistency check.

Lemma 29. *For $(m, m') = (3, 3)$, Table 3 gives the complete enumeration of the 63 nonzero masks, grouped by $\dim E_\lambda$. In particular, $\dim E_\lambda \leq 2$ when $\lambda_L \neq 0$ and $\lambda_R \neq 0$, and $\dim E_\lambda \leq 4$ for every nonzero mask.*

Proof. For each of the 63 nonzero masks λ , we evaluate $\Delta_\alpha H_\lambda(y)$ for all 64 directions $\alpha \in \mathbb{F}_2^6$ and all 64 inputs $y \in \mathbb{F}_2^6$. A direction is retained exactly when the derivative is constant over all inputs. Binary Gaussian elimination on the retained directions gives $\dim E_\lambda$. No sampling or floating-point arithmetic is used. Table 3 groups the resulting 63 mask classifications. The last column separates the masks with $\lambda_L \neq 0$ and $\lambda_R \neq 0$ from the one-sided ones; it totals $(2^3 - 1)^2 = 49$, as it must. The two masks with $\dim E_\lambda = 4$ are $\lambda = e_0$ and $\lambda = \hat{e}_0$, both one-sided, so the two-sided maximum is 2. Both bounds follow. $\square$

Table 3: Complete enumeration for $(m, m') = (3, 3)$, grouped over all 63 nonzero masks. The difference count includes the zero difference.

$\dim E_\lambda$	# masks	# differences per mask	Bias = 1	Bias = -1	# with $\lambda_L \neq 0$ and $\lambda_R \neq 0$
1	36	2	2	0	36
2	25	4	2	2	13
4	2	16	8	8	0
Total	63				49

This exhaustive calculation supplies the finite $(3, 3)$ case and is not used for any larger branch length.

B Fixed-Boundary Attainable Sets

This appendix supplies the two proofs deferred from Subsection 3.4. The first two parts derive the zero-internal fixed-boundary attainable sets of Lemma 4: for every odd $r \geq 5$ they prove the affine hulls of Table 4 directly, without a base-length enumeration or an $r \mapsto r + 2$ lifting, and the exceptional length $r = 3$ is treated by direct substitution. The last part proves Lemma 6, that a nonzero internal input difference leaves every boundary quadruple attainable.

Throughout, the directions of such a set lie in $I_r = \{v \in \mathbb{F}_2^r : v_0 = v_1 = 0\}$. For odd $r \geq 5$, put

$$\begin{aligned} E_r &= \langle e_2, e_4, \dots, e_{r-1} \rangle, \\ o_r &= e_3 + e_5 + \dots + e_{r-2}, \\ J_r &= \{v \in I_r : v_3 = 0\}. \end{aligned}$$

Table 4 records the anchor and direction space of each of the eight sets; $c_r(q)$ and $U_r(q)$ are as defined in Subsection 3.4.

Table 4: Anchors and direction spaces of the zero-internal fixed-boundary attainable sets for odd $r \geq 5$.

(α_0, α_1)	(d_0, d_1)	anchor $c_r(q)$	direction space $U_r(q)$
00	00	0	$\{0\}$
10	10	e_0	I_r
10	11	$e_0 + e_1 + o_r$	E_r
01	01	e_1	E_r
01	11	$e_0 + e_1$	I_r
11	01	e_1	E_r
11	10	$e_0 + o_r$	E_r
11	11	$e_0 + e_1$	J_r

B.1 Linear Elimination and Affine Upper Bounds

Put

$$q = (\alpha_0, \alpha_1, 0, \dots, 0), \quad d_0 = u_0, \quad d_1 = u_1,$$

and let all operations be over $\mathbb{F}_2$. The coordinate equations for $\Delta_d \chi_r(x) = q$ are

$$q_i = d_i + d_{i+2} + d_{i+1}d_{i+2} + d_{i+1}x_{i+2} + d_{i+2}x_{i+1}, \quad 0 \leq i < r. \quad (50)$$

Once d is fixed, (50) is an affine linear system in $x_0, \dots, x_{r-1}$:

$$d_{i+1}x_{i+2} + d_{i+2}x_{i+1} = q_i + d_i + d_{i+2} + d_{i+1}d_{i+2}.$$

For later substitutions, its four local forms are

(d_{i+1}, d_{i+2})	resulting equation
(0, 0)	$q_i = d_i,$
(1, 0)	$x_{i+2} = q_i + d_i,$
(0, 1)	$x_{i+1} = q_i + d_i + 1,$
(1, 1)	$x_{i+1} + x_{i+2} = q_i + d_i.$

Ordinary Gaussian elimination over $\mathbb{F}_2$ decides whether a prescribed difference is feasible.

For the internal coordinates, $q_i = 0$, and the same equations can be read backwards as

$$d_i = d_{i+2} + d_{i+1}d_{i+2} + d_{i+1}x_{i+2} + d_{i+2}x_{i+1}.$$

Starting with the two cyclic equations adjacent to d_0, d_1 and eliminating two coordinates at a time gives the following necessary conditions.

For example, take $(d_0, d_1) = (0, 1)$. The equations at coordinates $r-1$ and $r-2$ first give

$$d_{r-1} = 1 + x_0, \quad d_{r-2} = d_{r-1}x_0 = (1 + x_0)x_0 = 0.$$

Continuing this two-coordinate elimination around the cycle yields

$$d_{r-4} = d_{r-6} = \dots = d_3 = 0.$$

Substitution into the coordinate-1 boundary equation forces $\alpha_1 = 1$, whereas α_0 remains free. Precisely the two rows (01, 01) and (01, 11) occur, and both are contained in $e_1 + E_r$. The other three choices of (d_0, d_1) are reduced in the same manner.

This elimination also lists the contradictions for all omitted boundary types without an exhaustive search:

(d_0, d_1)	forced boundary condition	excluded (α_0, α_1)
00	$\alpha_0 = \alpha_1 = 0$	10, 01, 11
10	$\alpha_0 = 1$	00, 01
01	$\alpha_1 = 1$	00, 10
11	$(\alpha_0, \alpha_1) \neq (0, 0)$	00

The contradiction is explicit in each row. For $d_0d_1 = 00$, a nonzero α_0 or α_1 makes the corresponding boundary equation read $0 = 1$, excluding 10, 01, 11. For $d_0d_1 = 10$, the coordinate-0 equation becomes $0 = 1$ when $\alpha_0 = 0$, which excludes 00, 01. For $d_0d_1 = 01$, the coordinate-1 equation becomes $0 = 1$ when $\alpha_1 = 0$, excluding 00, 10. Finally, for $d_0d_1 = 11$, completing the backward substitution with $(\alpha_0, \alpha_1) = 00$ makes one of the two cyclic boundary equations read $0 = 1$. These eight contradictions and the eight attainable types below exhaust all sixteen boundary types.

Table 5: Affine upper bounds obtained from the backward equations.

(d_0, d_1)	(α_0, α_1)	necessary internal condition	affine upper bound
00	00	$d_2 = \dots = d_{r-1} = 0$	$\{0\}$
10	10	no additional linear condition	$e_0 + I_r$
11	10	$d_3 = d_5 = \dots = d_{r-2} = 1$	$e_0 + e_1 + o_r + E_r$
01	01	$d_3 = d_5 = \dots = d_{r-2} = 0$	$e_1 + E_r$
11	01	no additional linear condition	$e_0 + e_1 + I_r$
01	11	$d_3 = d_5 = \dots = d_{r-2} = 0$	$e_1 + E_r$
10	11	$d_3 = d_5 = \dots = d_{r-2} = 1$	$e_0 + o_r + E_r$
11	11	$d_3 = 0$	$e_0 + e_1 + J_r$

The necessary conditions in the third column yield the affine upper bounds in the last column.

B.2 Explicit Realization of the Affine Hulls

To attain every upper bound in Table 5, define

$$S_k = e_k + e_{k+1} + \dots + e_{r-1}, \quad 2 \leq k \leq r-1,$$

and, for even k , define

$$T_k = e_k + e_{k+2} + \dots + e_{r-1}.$$

An out-of-range T_k is understood to be zero. In the ordered internal coordinates $2, \dots, r-1$, the vectors S_k are the rows

$$(1, 1, \dots, 1), (0, 1, \dots, 1), \dots, (0, \dots, 0, 1).$$

Their coefficient matrix is unit upper triangular. The same is true of $T_2, T_4, \dots, T_{r-1}$ after restriction to the even internal coordinates. These triangular families give

$$\langle S_2, S_3, \dots, S_{r-1} \rangle = I_r, \quad \langle T_2, T_4, \dots, T_{r-1} \rangle = E_r. \quad (51)$$

We begin with boundary type 00/00, whose attainable set has affine hull $\{0\}$. The difference $d = 0$ is realized, for example, at $x = 0$. The corresponding upper bound in Table 5 has no nonzero direction, so the affine hull is exactly $\{0\}$.

The two sets with direction space I_r are treated next. For $(d_0, d_1)/(\alpha_0, \alpha_1) = 10/10$, the baseline $d^{(0)} = e_0$ is feasible because

$$\Delta_{e_0} \chi_r(e_{r-1}) = e_0.$$

For every $2 \leq k \leq r-1$, set $d^{(k)} = e_0 + S_k$ and

$$A_k = \begin{cases} e_1 + T_4, & k = 2, \\ e_1 + e_{k-1} + e_k + T_{k+1}, & k \geq 3 \text{ odd}, \\ e_1 + e_{k-1} + T_{k+2}, & k \geq 4 \text{ even}. \end{cases}$$

Direct substitution in (50) gives

$$\Delta_{e_0+S_k} \chi_r(A_k) = e_0.$$

The difference $d^{(k)} + d^{(0)} = S_k$ lies in the direction space of the set, and (51) proves that its affine hull is $e_0 + I_r$.

For boundary type 11/01, take the baseline $d^{(0)} = e_0 + e_1$ and $x = e_2 + e_{r-1}$. For $d^{(k)} = e_0 + e_1 + S_k$, put

$$B_k = \begin{cases} A_k, & k = 2, 3, \\ A_k + e_2, & k \geq 4. \end{cases}$$

Then

$$\Delta_{e_0+e_1} \chi_r(e_2 + e_{r-1}) = e_1, \quad \Delta_{e_0+e_1+S_k} \chi_r(B_k) = e_1.$$

The point-difference argument used for the preceding case gives $e_0 + e_1 + I_r$.

For the four sets with direction space E_r , first let $\varepsilon \in \mathbb{F}_2$, take $d = e_1 + T_k$, and prescribe $q = \varepsilon e_0 + e_1$. The baseline $d = e_1$ is realized by $x = e_0 + \varepsilon e_2$. For even k , let

$$x = \begin{cases} \varepsilon e_2, & k = 2, \\ e_{k-1} + \varepsilon e_2, & k \geq 4. \end{cases}$$

Substitution gives

$$\Delta_{e_1+T_k} \chi_r(x) = \varepsilon e_0 + e_1.$$

The choices $\varepsilon = 0$ and 1 prove, respectively, the rows 01/01 and 01/11, both with affine hull $e_1 + E_r$.

For the remaining two rows, let

$$b_\delta = e_0 + \delta e_1 + o_r, \quad \delta \in \mathbb{F}_2.$$

The baseline is realized at $x = 0$:

$$\Delta_{b_\delta} \chi_r(0) = e_0 + (1 + \delta)e_1.$$

For even k , prescribe $d = b_\delta + T_k$ and choose

$$x = \begin{cases} e_1 + T_2, & k = 2, \\ e_1 + e_{k-1} + T_k, & k \geq 4. \end{cases}$$

Then

$$\Delta_{b_\delta+T_k} \chi_r(x) = e_0 + (1 + \delta)e_1.$$

Taking $\delta = 1$ yields $e_0 + e_1 + o_r + E_r$ for the row 11/10, while $\delta = 0$ yields $e_0 + o_r + E_r$ for the row 10/11.

It remains to treat boundary type 11/11, whose direction space is J_r . The baseline $d^{(0)} = e_0 + e_1$ is realized by $x = e_{r-1}$. Consider the direction family

$$\mathcal{G}_J = \{T_k : 2 \leq k \leq r-1, k \text{ even}\} \cup \{S_k : 5 \leq k \leq r-2, k \text{ odd}\}.$$

For an even k , the difference $e_0 + e_1 + T_k$ is realized by

$$x = \begin{cases} e_1 + e_2, & k = 2, \\ e_1 + e_{k-1}, & k \geq 4, \end{cases}$$

whereas, for odd $k \geq 5$, the difference $e_0 + e_1 + S_k$ is realized by

$$x = e_1 + e_{k-1} + e_k + T_{k+1}.$$

Every substitution gives output difference $e_0 + e_1$. Order the vectors in $\mathcal{G}_J$ by the pivot coordinates

$$e_2, e_4, e_5, \dots, e_{r-1}.$$

Their coefficient matrix is again unit upper triangular, so

$$\langle \mathcal{G}_J \rangle = \langle e_2, e_4, e_5, \dots, e_{r-1} \rangle = J_r,$$

and the affine hull is exactly $e_0 + e_1 + J_r$.

The upper bounds and witness families coincide row by row, proving Table 4 for every odd $r \geq 5$. For $r = 3$, direct substitution in the three cyclic equations gives

$$B_3 = \{0, 5, 7, 10, 11, 13, 14\}.$$

The 11/11 row is absent; the attainable set for 00/00 is $\{000\}$, and the direction space of every other nonempty set is $\langle e_2 \rangle$, as stated in Lemma 4. Together, the all-length witnesses and the separate length-three substitution provide precisely the zero-internal local structure used in the global analysis.

B.3 Full Boundary Support for Nonzero Internal Differences

We now prove Lemma 6.

Proof. Put $\alpha = (\alpha_0, \alpha_1, \alpha_2, \dots, \alpha_{r-1})$. By Lemma 2, the set consists of the $d \in S_\alpha$ with the prescribed first two coordinates, where S_α is the attainable-difference set of χ_r^{-1} from Section 3, and membership in S_α is decided by conditions (i) and (ii) of Theorem 1. For $r = 3$, put

$$d_2 = 1 + d_0\alpha_0 + d_1\alpha_1.$$

The four choices of (d_0, d_1) admit, respectively, the following base points, listed in the order 00, 10, 01, 11:

$$\begin{aligned} &(\alpha_1, \alpha_0 + 1, 0), \quad (0, \alpha_0, \alpha_0 + \alpha_1), \\ &(1 + \alpha_1, 0, \alpha_0), \quad (\alpha_0 + \alpha_1, 0, \alpha_0 + 1). \end{aligned}$$

Substitution in (4) gives $(\alpha_0, \alpha_1, 1)$ in every case, so every such set with $\alpha_2 = 1$ is nonempty. Let $r \geq 5$. In each case below, d equals $\mathbf{1}$ off the listed zero set, and the local-test conditions are verified directly; $\eta \neq 0$ means α has a nonzero coordinate outside $\{0, 1\}$.

Prefix $(d_0, d_1) = (1, 1)$. If $\text{wt}(\alpha)$ is odd, take $d = \mathbf{1}$. If $\text{wt}(\alpha)$ is even, then $\alpha \notin \{0, \mathbf{1}\}$, so some cyclic index t has $\alpha_t \neq \alpha_{t+1}$. If such a $t \leq r-3$ exists, take the single isolated zero $t+2 \notin \{0, 1\}$; condition (ii) there reads $d_t = \alpha_t + \alpha_{t+1} = 1$ and holds. Otherwise all such indices lie in $\{r-2, r-1\}$, forcing $\alpha = \mathbf{1} + e_{r-1}$; take the zero set $\{2, 3\}$, whose only condition is (i) at $i = 1$, namely $d_1 = \alpha_1 = 1$.

Prefix (0,1). If $\alpha_{r-2} \neq \alpha_{r-1}$, take the single isolated zero 0; condition (ii) reads $d_{r-2} = \alpha_{r-2} + \alpha_{r-1} = 1$ and holds. Otherwise let $\ell = \max\{i \in [1, r-2] : \alpha_i = 1\}$; this set is nonempty, since a nonzero coordinate outside $\{0,1\}$ located at $r-1$ would force $\alpha_{r-2} = \alpha_{r-1} = 1$. Take the zero set $[\ell+1, 0]$: condition (i) applies at $i \in [\ell, r-2]$ and reads $d_\ell = \alpha_\ell = 1$ at $i = \ell$ and $0 = \alpha_i$ at $i \in [\ell+1, r-2]$, which holds by maximality of ℓ .

Prefix (1,0). If $\alpha_{r-1} \neq \alpha_0$, take the single isolated zero 1; condition (ii) reads $d_{r-1} = \alpha_{r-1} + \alpha_0 = 1$. If $\alpha_{r-1} = \alpha_0 = 1$, let $f = \min\{i \in [1, r-1] : \alpha_i = 1\}$ and take the zero set $[1, \min(f+1, r-1)]$: condition (i) applies at $i = 0$, where $d_0 = \alpha_0 = 1$, and at the interior indices, where $0 = \alpha_i$ holds by minimality of f . If $\alpha_{r-1} = \alpha_0 = 0$, let $\ell = \max\{i \in [1, r-2] : \alpha_i = 1\}$, which exists since the nonzero coordinate outside $\{0,1\}$ cannot sit at $r-1$. For $\ell \leq r-3$, take the zero set $\{1\} \cup [\ell+1, r-1]$: condition (ii) at the isolated zero 1 reads $d_{r-1} = \alpha_{r-1} + \alpha_0 = 0$, which holds because $r-1$ lies in the zero run; condition (i) on the run reads $d_\ell = \alpha_\ell = 1$ and $0 = \alpha_i$ on $[\ell+1, r-3]$. For $\ell = r-2$ and $\alpha_{r-3} = 0$, take the zero set $\{1, r-1\}$: condition (ii) at 1 is as before, and at $r-1$ it reads $d_{r-3} = \alpha_{r-3} + \alpha_{r-2} = 1$, which holds. For $\ell = r-2$ and $\alpha_{r-3} = 1$, take the zero set $\{1, r-2, r-1\}$: condition (ii) at 1 is as before, and the run $[r-2, r-1]$ imposes only condition (i) at $i = r-3$, namely $d_{r-3} = \alpha_{r-3} = 1$.

Prefix (0,0). If $\alpha_{r-1} = 1$, take the zero set $\{0,1\}$, whose only condition is (i) at $i = r-1$, namely $d_{r-1} = \alpha_{r-1} = 1$. Otherwise let $\ell = \max\{i : \alpha_i = 1\} \in [2, r-2]$ and take the zero set $[\ell+1, 1]$: condition (i) applies at $i \in [\ell, r-1]$ and reads $d_\ell = \alpha_\ell = 1$ at $i = \ell$, $0 = \alpha_i$ on $[\ell+1, r-2]$ by maximality, and $0 = \alpha_{r-1}$ at $i = r-1$, which is the case assumption. $\square$

C Two-Coordinate Recurrences

This appendix supplies the active-direction analysis used in Lemmas 5 and 25. The first subsection gives a computer-assisted finite-state proof of Table 6; the second uses that classification in the one-active-branch analysis.

C.1 Exact Annihilator Classification for Three Retained Sets

For $2 \leq j \leq r-1$, let $\mathbf{q}_j$ denote the internal input-difference word whose only nonzero coordinate is α_j ; this is distinct from the coordinate vector $e_j \in I_r$. In Table 6 the symbol $*$ denotes an arbitrary excluded boundary input pair, and every source word not displayed for a fixed profile has zero annihilator.

Table 6: Exact annihilator classification for three fixed-boundary attainable sets with nonzero internal input difference and odd $r \geq 7$.

boundary output β	excluded input $\bar{u}$	internal input difference η	annihilator generator
00	*	$\mathbf{q}_j, 2 \leq j \leq r-1$ $\mathbf{q}_j + \mathbf{q}_{j+1}, 2 \leq j \leq r-4$ $\mathbf{q}_{r-3} + \mathbf{q}_{r-2}$ $\mathbf{q}_{r-2} + \mathbf{q}_{r-1}$	e_j $e_{j+1} + e_{j+3}$ e_{r-2} e_{r-1}
10	00, 01, 11	$\mathbf{q}_{r-1}$	e_2
10	10	$\mathbf{q}_{r-1}$ $\mathbf{q}_j, 3 \leq j \leq r-2, j \text{ odd}$	e_2 e_j
01	00, 10, 01	$\mathbf{q}_2$	$e_2 + e_4$
01	11	$\mathbf{q}_2$ $\mathbf{q}_j, 3 \leq j \leq r-2, j \text{ odd}$ $\mathbf{q}_j + \mathbf{q}_{j+1}, 2 \leq j \leq r-5, j \text{ even}$ $\mathbf{q}_{r-3} + \mathbf{q}_{r-2}$ $\mathbf{q}_{r-1} \text{ or } \mathbf{q}_{r-2} + \mathbf{q}_{r-1}$	$e_2 + e_4$ e_j $e_{j+1} + e_{j+3}$ e_{r-2} $e_{r-2} + e_{r-1}$
11	00, 01	$\eta \neq 0$	$\{0\}$
11	10	$\mathbf{q}_2 + \mathbf{q}_3$	$e_3 + e_5$
11	11	$\mathbf{q}_j, 3 \leq j \leq r-2, j \text{ odd}$	e_j

All vector spaces and relations in this subsection are over $\mathbb{F}_2$. Fix an odd $r \geq 7$, a nonzero internal source word $\eta = (\alpha_2, \dots, \alpha_{r-1})$, a boundary output $\beta = (\alpha_0, \alpha_1)$, and an excluded boundary input $\bar{u} \in \mathbb{F}_2^2$. The pair $(\beta, \bar{u})$ will be called a *profile*. We abbreviate the fixed-boundary attainable set in (5) by $\mathcal{C}_r(\eta; u, \beta)$, and put

$$V_r(\eta; \beta, \bar{u}) = \sum_{u \in \mathbb{F}_2^2 \setminus \{\bar{u}\}} \text{dir Aff } \mathcal{C}_r(\eta; u, \beta) \subseteq I_r, \quad I_r = \langle e_2, \dots, e_{r-1} \rangle. \quad (52)$$

The symbol $\mathbf{q}_j$ denotes the source word supported at j , whereas e_j denotes the coordinate vector of I_r supported at the same index.

Proposition 3 (Exact classification for three retained sets). *For every odd $r \geq 7$, Table 6 is exact. If its row for $(\eta, \beta, \bar{u})$ displays g , then*

$$V_r(\eta; \beta, \bar{u})^{\perp_{I_r}} = \langle g \rangle = \{0, g\}.$$

Every row marked $\{0\}$, and every nonzero source word not displayed for its fixed $(\beta, \bar{u})$, has annihilator $\{0\}$.

The proof passes from these sets to weighted transition relations and then to a finite-state quotient that is invariant under a two-coordinate extension. For a finite set $S \subseteq I_r$, write $\text{dir Aff } S = \langle s + s' : s, s' \in S \rangle$, with value $\{0\}$ if S is empty.

Lemma 30. *For $\lambda \in I_r$,*

$$\lambda \in V_r(\eta; \beta, \bar{u})^{\perp_{I_r}} \iff d \mapsto \langle \lambda, d \rangle \text{ is constant on each } \mathcal{C}_r(\eta; u, \beta), \quad u \neq \bar{u}. \quad (53)$$

The three constants in (53) need not agree.

Proof. For any S , orthogonality to $\langle s + s' : s, s' \in S \rangle$ is equivalent to $\langle \lambda, s \rangle = \langle \lambda, s' \rangle$ for every $s, s' \in S$. Taking the intersection of these orthogonal complements over the three retained sets proves the equivalence without relating the three constants. $\square$

Transition representation. Equation (4) can be solved backward at coordinate i . Immediately before that coordinate, retain

$$s_i = (x_{i+1}, x_{i+2}, d_{i+1}, d_{i+2}) = (X, Y, P, Q).$$

Given the source bit $a = \alpha_i$ and $w = x_i$, the emitted difference and successor are uniquely determined by

$$D = a + Q + PQ + PY + QX = d_i, \quad \tau_{a,w}(X, Y, P, Q) = (w, X, D, P). \quad (54)$$

The initial and terminal states are

$$s_{\text{in}} = (x_0, x_1, d_0, d_1), \quad s_{\text{out}} = (x_2, x_3, d_2, d_3).$$

The two equations not processed in the order $r-1, r-2, \dots, 2$ are imposed by

$$\begin{aligned} B_0(s_{\text{in}}, s_{\text{out}}) &= d_0 + d_2 + d_1 d_2 + d_1 x_2 + d_2 x_1, \\ B_1(s_{\text{in}}, s_{\text{out}}) &= d_1 + d_3 + d_2 d_3 + d_2 x_3 + d_3 x_2. \end{aligned} \quad (55)$$

An accepting run starts with difference boundary u and processes the source bits from α_{r-1} down to α_2 , using (54) at every step. Its endpoints must satisfy $B(s_{\text{in}}, s_{\text{out}}) = \beta$.

Lemma 31. *For fixed (r, η, u, β) , accepting runs are in bijection with*

$$\{(x, d) : \Delta_d \chi_r(x) = (\beta, \eta), (d_0, d_1) = u\}.$$

Their labels are precisely the points of $\mathcal{C}_r(\eta; u, \beta)$. If two copies are run synchronously with the same source and boundary data, while their choices of x_i are independent, the complete difference-label set is

$$\{d + d' : d, d' \in \mathcal{C}_r(\eta; u, \beta)\}.$$

The common boundary input makes the first two coordinates of every such label zero, so the labels lie in I_r . Consequently,

$$\text{span}\{d + d' : d, d' \in \mathcal{C}_r(\eta; u, \beta)\} = \text{dir Aff } \mathcal{C}_r(\eta; u, \beta), \quad (56)$$

and the three retained boundary inputs give exactly $V_r(\eta; \beta, \bar{u})$.

Proof. A differential solution fixes every state and, by (54), every emitted bit. After coordinate 2, equations (55) are exactly the two remaining cyclic differential equations. Conversely, the initial state and the successive choices of w recover x , while the emitted bits recover d . The copied entries in consecutive states enforce consistency on all overlapping coordinates. Thus the two constructions are inverse and the run labels are exactly the elements of that set.

Applying this bijection to both synchronized projections gives every ordered pair of differential solutions. Identifying witnesses with the same label pair removes only the possible multiplicity in x , leaving all pairs (d, d') of elements of the same set. Their differences span the affine direction space. Summing over the retained inputs gives (52). $\square$

Weighted boundary relation. The induction retains the full boundary relation, rather than only its annihilator. For a test mask $\lambda \in I_r$, augment the sixteen path states by

$$h = \sum_{j \text{ processed}} \lambda_j d_j,$$

giving the fixed set $\mathbb{F}_2^4 \times \mathbb{F}_2$ of 32 augmented states. For a source-mask pair $(a, \ell) \in \mathbb{F}_2^2$, define the Boolean relation

$$(s, h) G_{a,\ell}(\tau_{a,w}(s), h + \ell D) \quad (w \in \mathbb{F}_2), \quad (57)$$

where D is the unique bit in (54). We use the standard convention for relational composition: in $B \circ A$, the relation A is traversed first. The complete internal relation is therefore

$$\mathcal{R}(\eta, \lambda) = G_{\alpha_2, \lambda_2} \circ \cdots \circ G_{\alpha_{r-1}, \lambda_{r-1}}. \quad (58)$$

Thus, if R_k is the relation for coordinates $2, \dots, k$, adjoining coordinate $k+1$ gives

$$R_{k+1} = R_k \circ G_{\alpha_{k+1}, \lambda_{k+1}}. \quad (59)$$

The path bijection and Lemma 30 yield

$$\lambda \in V_r(\eta; \beta, \bar{u})^{\perp_{I_r}} \iff \begin{array}{c} \text{for every } u \neq \bar{u}, \\ \left| \left\{ h \in \mathbb{F}_2 : \begin{array}{l} \exists s_{\text{in}}, s_{\text{out}}, (s_{\text{in}}, 0) \mathcal{R}(\eta, \lambda)(s_{\text{out}}, h), \\ (d_0, d_1)(s_{\text{in}}) = u, \quad B(s_{\text{in}}, s_{\text{out}}) = \beta \end{array} \right\} \right| \leq 1. \end{array} \quad (60)$$

Thus $\mathcal{R}(\eta, \lambda)$ determines the annihilator membership value for each of the sixteen profiles $(\beta, \bar{u})$.

Table specification and quotient. For a profile $p = (\beta, \bar{u})$, write $P_p^{(r)}(\eta, \lambda)$ when $\lambda \neq 0$ and the row of Table 6 for p contains the source η and displays λ as its generator. Thus the predicate records the assertion of the table in a form that can be compared with the constancy test.

Lemma 32. *For fixed r , $\eta \neq 0$, and a profile p , at most one nonzero λ satisfies $P_p^{(r)}(\eta, \lambda)$. Moreover, the sixteen truth values of $P_p^{(r)}(\eta, \lambda)$ are determined by a finite-state scanner of the pairs*

$$(\alpha_2, \lambda_2), \dots, (\alpha_{r-1}, \lambda_{r-1}), \quad (61)$$

whose state set and update rule are independent of r . In particular, the state after appending one further pair is determined by the preceding state and that pair.

Proof. For a fixed profile, the rows of Table 6 separate first by the size and shape of the source support: a singleton, an adjacent pair, or one of the stated endpoint cases. The index ranges separate interior and terminal supports, and the remaining overlaps are separated by parity. Thus two rows for the same profile cannot assign different masks to the same source.

A scanner reads the pairs in (61) from left to right. It records the parity of the current coordinate, the bounded initial segment needed for the left-end cases, a rolling terminal segment for the right-end cases, and whether the nonzero pairs already read can still form one of the singleton or adjacent-support rows. Once a support violates every row, it enters a rejecting state. These records take only finitely many values: all variable-length parts of a row consist of zero pairs, parity uses one bit, and every nonzero local pattern and endpoint condition has bounded length. Taking the product of the row scanners determines all sixteen profile values and gives the asserted length-independent update. $\square$

Denote the resulting specification state after reading (61) by $\rho(\eta, \lambda)$, and define the quotient state

$$S(\eta, \lambda) = (\mathcal{R}(\eta, \lambda), \rho(\eta, \lambda), \mathbf{1}_{\eta \neq 0}, \mathbf{1}_{\lambda=0}), \quad (62)$$

where $\mathbf{1}_P$ is the Iverson indicator of the proposition P . The relation component determines the sixteen membership values in (60); the other components determine whether the source is nonzero and whether the zero mask or a Table 6 generator is present. The recurrence in (59), the scanner update, and the indicators are deterministic functions of the next pair (α_i, λ_i) . Consequently, equality of quotient states is preserved by one-coordinate, and hence by two-coordinate, extension.

Lemma 33. *For internal length $n = r - 2$, let $\mathcal{Q}_n$ be the set of quotient states (62) obtained from all 4^n sequences of source-mask pairs. The following properties hold.*

1. *Every pair sequence belongs to one quotient state, and the accumulated multiplicities account for all 4^n sequences.*
2. *At $n = 5, 7, 9, 11$, the membership values obtained from (60) agree, for every nonzero source, with the condition $\lambda = 0$ or $P_{\beta, \bar{u}}^{(r)}(\eta, \lambda)$.*
3. *$\mathcal{Q}_{11} = \mathcal{Q}_{13}$. Extending every state of $\mathcal{Q}_{11}$ by each of the sixteen two-pair tails gives exactly $\mathcal{Q}_{13}$, with multiplicities accounting for every extended sequence once.*

Proof. At length zero, take the identity relation, the initial recognizer state, and the two initial indicators. Appending a source-mask pair applies (57) by (59), advances the recognizer, and updates the indicators. Sequences with equal resulting quotient states are identified, and their multiplicities add. This finite recurrence defines each set $\mathcal{Q}_n$, each multiplicity, and every two-pair transition in the statement.

Source code for the finite-state quotient computation in this lemma is available in the [same repository](#). The program reconstructs the four local relations directly from (4). At each run it forms the quotient layers and their multiplicities in memory, compares the actual and specified profile values at every required odd layer, checks the two stable-set equalities, and tests all two-coordinate extensions. No precomputed transition data are used. These computations establish only the finite identities stated here; the all-length conclusion is deduced below from extension congruence and $\mathcal{Q}_{11} = \mathcal{Q}_{13}$. $\square$

Proposition 4. *For every odd $r \geq 7$, every nonzero η , every $\lambda \in I_r$, and every profile, the test (60) holds if and only if $\lambda = 0$ or $P_{\beta, \bar{u}}^{(r)}(\eta, \lambda)$.*

Proof. The internal length $n = r - 2$ is odd. Lemma 33 proves the assertion for $n = 5, 7, 9, 11$. Two-coordinate extension depends only on the current quotient state and the two appended source-mask pairs. Since $\mathcal{Q}_{13} = \mathcal{Q}_{11}$, the same extension gives $\mathcal{Q}_{15} = \mathcal{Q}_{13}$; induction yields $\mathcal{Q}_n = \mathcal{Q}_{11}$ for every odd $n \geq 11$. Both the constancy test and the Table 6 predicate are functions of the quotient state, so their equality persists at every later odd length. Together with the four initial odd layers, this covers all odd $r \geq 7$. $\square$

Proof of Proposition 3. Fix a nonzero source η , a profile $(\beta, \bar{u})$, and let λ range over all of I_r . Lemma 30, equation (60), and Proposition 4 give

$$V_r(\eta; \beta, \bar{u})^{\perp_{I_r}} = \{\lambda \in I_r : \lambda = 0 \text{ or } P_{\beta, \bar{u}}^{(r)}(\eta, \lambda)\}.$$

If Table 6 lists g , the definition of its predicate and Lemma 32 identify the right-hand side with $\{0, g\} = \langle g \rangle$. Since $\dim I_r = r - 2$, finite-dimensional orthogonality then gives

$$\dim V_r = r - 3, \quad V_r = \langle g \rangle^{\perp_{I_r}} = g^{\perp} \cap I_r.$$

If the table gives $\{0\}$, or does not list the source for the fixed profile, its predicate is false for every nonzero λ . Consequently the annihilator is $\{0\}$ and $V_r = I_r$. Because λ was arbitrary, this conclusion excludes every nonzero mask, not only the masks occurring in the table. $\square$

Corollary 6. *For every odd $r \geq 7$, every $\eta \neq 0$, and every $\beta, \bar{u} \in \mathbb{F}_2^2$,*

$$\dim V_r(\eta; \beta, \bar{u})^{\perp_{I_r}} \leq 1.$$

Proof. This is the listed/unlisted dichotomy of Proposition 3. $\square$

Lemma 5 follows.

Proof. If $r = 3$, then $\dim W_r(\alpha_2; \mathcal{B}) \geq 0 = r - 3$, so suppose $r \geq 5$. For $r = 5$, compose the three explicit weighted relations (57) and apply the constancy test (60). Among the $7 \cdot 4 \cdot 4 = 112$ profiles, 36 have a one-dimensional annihilator and 76 have zero annihilator; none has larger dimension. Thus the required rank is at least $2 = r - 3$. This is only the short rank statement and does not extend Table 6 to length five. Hence assume $r \geq 7$. Choose a boundary output pair $\beta = (\alpha_0, \alpha_1)$ for which three distinct boundary input pairs are admitted, and let $\bar{u}$ be the unique pair in $\mathbb{F}_2^2$ not selected. With $\eta = (\alpha_2, \dots, \alpha_{r-1})$, put

$$V_r(\eta; \beta, \bar{u}) = \sum_{u \in \mathbb{F}_2^2 \setminus \{\bar{u}\}} \text{dir Aff } \mathcal{C}_r(\eta; u, \beta).$$

Then

$$V_r(\eta; \beta, \bar{u}) \subseteq W_r(\eta; \mathcal{B}) \subseteq I_r.$$

Corollary 6 gives

$$\dim V_r(\eta; \beta, \bar{u})^{\perp_{I_r}} \leq 1.$$

Since $\dim I_r = r - 2$, rank-nullity yields

$$\dim W_r(\eta; \mathcal{B}) \geq \dim V_r(\eta; \beta, \bar{u}) \geq (r - 2) - 1 = r - 3,$$

which proves (7). $\square$

Residual transitions. For the application in Subsection C.2, it remains to record how the generators transform under a two-coordinate extension. This result is an interface for that application and is not used in the proof of Proposition 3. The transformation is obtained by projecting synchronized direction labels onto the old generator and the two appended coordinates; the same projection also treats sources whose old annihilator is zero. Write $R = r + 2$, $\eta^+ = (\eta, t)$, $t = (\alpha_r, \alpha_{r+1})$. If the old annihilator is $\langle g \rangle$, every mask whose old-coordinate part lies in this annihilator has the form

$$\lambda^+ = (\zeta g, \mu, \nu), \quad \zeta, \mu, \nu \in \mathbb{F}_2.$$

For two points in one new set, orthogonality is exactly

$$\zeta \langle g, d^- + d'^- \rangle + \mu(d_r + d'_r) + \nu(d_{r+1} + d'_{r+1}) = 0. \quad (63)$$

Lemma 31 shows that the rows in (63) form the full projected direction space.

Proposition 5 (Residual transition classification). *Fix a profile and one of the four source tails t .*

1. *Suppose $\eta \neq 0$ is listed at length r , with generator g . The kernel of the residual system (63) is nonzero exactly when Table 6 lists $\eta^+ = (\eta, t)$ at length R . If its generator is g^+ , then the old-coordinate part of g^+ is either 0 or g , and the unique triple satisfying*

$$g^+ = (\zeta g, \mu, \nu)$$

generates the residual kernel.

2. *If $\eta \neq 0$ is unlisted for the profile, the residual system with zero old-coordinate part has full rank, and no listed length- R source has η as its prefix.*
3. *If $\eta = 0$ and $t \neq 0$, the three resulting sources are listed exactly in the profiles specified by Table 6, with the displayed generators; every other such profile has zero annihilator.*

Every listed length- R source has a unique predecessor of the first or third type. All source–tail transitions not covered above have zero annihilator.

Proof. For a quotient state, the weighted relation determines all synchronized labels admitted by each retained set. Projecting a label to

$$(\langle g, d^- + d'^- \rangle, d_r + d'_r, d_{r+1} + d'_{r+1})$$

gives every row of (63), by Lemma 31. The same program independently row-reduces these projected systems for every stable quotient state, profile, and source tail. It also treats zero annihilators and zero source prefixes, and compares every result directly with Table 6. The verified transition is a function of the quotient state, so the extension congruence and $\mathcal{Q}_{11} = \mathcal{Q}_{13}$ propagate it to every later odd length.

The source decomposition

$$\eta_R = \eta_r + \alpha_r \mathbf{q}_r + \alpha_{r+1} \mathbf{q}_{r+1}$$

is unique, where η_r is extended by two zero coordinates. The residual triple (ζ, μ, ν) therefore identifies the target generator uniquely. Full rank in the second case excludes a mask supported only on the new coordinates, while Proposition 4 excludes a target with any other old-coordinate part. The zero-prefix calculation supplies the new unit and adjacent endpoint sources. The direct Table 6 comparison and the unique prefix–tail decomposition prove both coverage and uniqueness. $\square$

C.2 One-Active-Branch Kernel Classification

Only the left branch is called active when $\eta_L \neq 0$ and $\eta_R = 0$; only the right branch is active when $\eta_L = 0$ and $\eta_R \neq 0$. Throughout this subsection, a boundary tuple means the boundary *input* difference $(\alpha_0, \alpha_1, \hat{\alpha}_0, \hat{\alpha}_1)$. The two tables below give the exact $m, m' \geq 5$ cases; the length-three formulas are handled separately in the main text. Every undisplayed coordinate is zero, and a boundary case satisfying no row has zero kernel. Here “all parameters” means $(b, \hat{b}) \in \{(0, 0), (0, 1), (1, 0)\}$, and each local-family name refers to Table 6; the last two columns give the complete global data.

Table 7: Nonzero kernels for one-active-branch directions on the left and $m \geq 5$. The boundary tuple is $(\alpha_0, \alpha_1, \hat{\alpha}_0, \hat{\alpha}_1)$.

boundary case	local family and range	nonzero internal input-difference coordinates	generator of $K_H(\alpha)$: nonzero coordinates
0000	unit, $2 \leq j \leq m-1$	$\alpha_j = 1$	$\lambda_j = 1$
0000	adjacent, $2 \leq j \leq m-2$	$\alpha_j = \alpha_{j+1} = 1$	$\lambda_{j+1} = \lambda_{j+3} = 1$ if $j \leq m-4$; $\lambda_0 = \lambda_{m-2} = 1$ if $j = m-3$; $\lambda_1 = \lambda_{m-1} = 1$ if $j = m-2$
0100	left-end unit; all parameters	$\alpha_2 = 1$	$\lambda_2 = \lambda_4 = 1$
0100/0101	odd-position unit, $3 \leq j \leq m-2$	$\alpha_j = 1$	$\lambda_j = 1$
0100/0101	even adjacent, $2 \leq j \leq m-3$	$\alpha_j = \alpha_{j+1} = 1$	$\lambda_{j+1} = \lambda_{j+3} = 1$ if $j \leq m-4$; $\lambda_0 = \lambda_{m-2} = 1$ if $j = m-3$
0010	terminal unit; all parameters	$\alpha_{m-1} = 1$	$\lambda_2 = \hat{\lambda}_1 = 1, \lambda_1 = b$

In the two rows marked 0100/0101, the first boundary case applies when $\hat{b} = 0$; the second applies only when $\hat{b} = 1$ and $m' = 3$.

Table 8: Nonzero kernels for one-active-branch directions on the right and $m' \geq 5$. The boundary tuple is $(\alpha_0, \alpha_1, \hat{\alpha}_0, \hat{\alpha}_1)$.

boundary case	local family and range	nonzero internal input-difference coordinates	generator of $K_H(\alpha)$: nonzero coordinates
0000	unit, $2 \leq j \leq m' - 1$	$\hat{\alpha}_j = 1$	$\hat{\lambda}_j = 1$
0000	adjacent, $2 \leq j \leq m' - 2$	$\hat{\alpha}_j = \hat{\alpha}_{j+1} = 1$	$\hat{\lambda}_{j+1} = \hat{\lambda}_{j+3} = 1$ if $j \leq m' - 4$; $\hat{\lambda}_0 = \hat{\lambda}_{m'-2} = 1$ if $j = m' - 3$; $\hat{\lambda}_1 = \hat{\lambda}_{m'-1} = 1$ if $j = m' - 2$
0001	left-end unit; all parameters	$\hat{\alpha}_2 = 1$	$\hat{\lambda}_2 = \hat{\lambda}_4 = 1$
$(b, 0, 0, 1)$	odd-position unit, $3 \leq j \leq m' - 2$	$\hat{\alpha}_j = 1$	$\hat{\lambda}_j = 1$
$(b, 0, 0, 1)$	even adjacent, $2 \leq j \leq m' - 3$	$\hat{\alpha}_j = \hat{\alpha}_{j+1} = 1$	$\hat{\lambda}_{j+1} = \hat{\lambda}_{j+3} = 1$ if $j \leq m' - 4$; $\hat{\lambda}_0 = \hat{\lambda}_{m'-2} = 1$ if $j = m' - 3$
0100	terminal unit; all parameters	$\hat{\alpha}_{m'-1} = 1$	$\lambda_0 = \hat{\lambda}_2 = 1, \hat{\lambda}_1 = \hat{b}$

For the two-coordinate propagation, let $r \geq 7$ be the current odd length and put $R = r + 2$. As in Subsection C.1, $\mathbf{q}_j$ denotes the unit internal input-difference word at coordinate j , while e_i and $\hat{e}_i$ denote the unit mask vectors on the left and right branches. The notation $+ab$ appends $(\alpha_r, \alpha_{r+1}) = (a, b)$, so every length- R word has the unique form $\eta_R = (\eta_r, \alpha_r, \alpha_{r+1})$.

Lemma 34 (Two-coordinate propagation). *Retain, for each of the sixteen ordinary- χ boundary types, the complete weighted relation (58). Under the extension $r \mapsto R$, the following and only the following updates can have a nonzero compatible mask.*

1. Under tail 00, the unit and adjacent families for boundary case 0000, the odd-unit and even-adjacent families for 0100 (and for 0101 when $\hat{b} = 1, m' = 3$), and the left-end unit for 0100 retain the formulas in Table 7, with m replaced by R .
2. The remaining nonzero updates are the right-end transitions below. The generator is the one in the corresponding row of Table 7; hence its coordinate formula is not repeated here.

old family	tail	new family	boundary case
zero	10	unit at r	0000
zero	01	unit at $r + 1$	0000
zero	11	adjacent at r	0000
zero	10	odd unit at r	0100/0101
unit at $r - 1$	10	even adjacent at $r - 1$	0000 or 0100/0101
zero	01	terminal unit at $r + 1$	0010

Here 0101 has the same exceptional meaning as in the left-active table.

The corresponding right-branch local updates are obtained by placing hats on the active coordinates. Their global boundary tuples and boundary mask coefficients are determined independently below. Every source–tail pair absent from this list has zero compatible kernel.

Proof. Append both source–mask pairs by applying (59) twice. The boundary maps in (20) select the applicable $(\beta, \bar{u})$ -profiles, while the two parity values in the complete weighted relation determine the compatible boundary-mask coefficients. Proposition 5 gives all internal prefix–tail residuals and excludes revival from an old zero annihilator. Substitution of the selected boundary coefficients leaves precisely the two update classes above. In each

displayed case the stated vector spans the one-dimensional residual kernel. For every other source–tail pair, the corresponding system has full column rank.

The complete relation is a two-coordinate congruence. Since each length- R source has a unique prefix–tail decomposition, the four tails and all boundary types are covered, with no transition outside the displayed classes. $\square$

Derivation of the state tables. The required one-dimensional bound is independent of the tables. When exactly one branch is active, Lemma 21 gives $\dim \text{dir}(D_\alpha) \geq N - 1$, so $K_H(\alpha) = \text{dir}(D_\alpha)^\perp$ satisfies $\dim K_H(\alpha) \leq 1$. Once a displayed nonzero mask has been verified to lie in the kernel, it generates the whole kernel.

The active lengths five and seven form the two finite base layers. For both possible zero-branch geometries, direct construction of the product sets covers every nonzero internal source, boundary case, and admissible parameter pair. Exact row reduction gives the corresponding instances of Tables 7 and 8; every omitted case has full column rank.

For a left active branch, substituting the zero-right-branch supports of Lemma 4 into (20) selects exactly the following global systems: 0000 for the unit and adjacent-support class; 0100 for the $\alpha_2 = 1$ class; 0010 for the top-unit class; and, for the parity-restricted class, 0100 when $\hat{b} = 0$, 0101 when $\hat{b} = 1, m' = 3$, and no boundary tuple when $\hat{b} = 1, m' \geq 5$. Substituting these supports also fixes the endpoint generator as $e_2 + be_1 + \hat{e}_1$. At length seven, the verified global base agrees with the local classification in Proposition 3. Induction via the propagation lemma then gives every row of Table 7. The candidate in each row annihilates the exact difference-direction matrix, while the independent dimension bound makes its span the full kernel. If no row applies, the corresponding propagation system has full column rank, so the kernel is zero.

The right-active table is derived independently because the boundary coupling need not be invariant under a branch swap. Substitute the zero-left-branch support into the right boundary map ϕ_R of (21) and apply the same weighted-relation extension to $(\hat{x}_i, \hat{d}_i, \hat{\alpha}_i, \hat{\lambda}_i)$. It selects 0000 for the unit and adjacent-support class, 0001 for $\hat{\alpha}_2 = 1$, $(b, 0, 0, 1)$ for the parity-restricted class, and 0100 for the top unit. The endpoint equations give $e_0 + \hat{e}_2 + \hat{b}\hat{e}_1$. These are exactly the boundary tuples and generators of Table 8; all other right-active systems have full column rank. This proves the stated if-and-only-if classification for every odd active length at least seven, including the $\hat{b} = 1, m' = 3$ exception. Together with the two independently verified base layers, the tables therefore cover every odd active length at least five.

D Exact Kernel Dimensions for Boundary-Only Directions

This appendix completes the boundary-only part of Theorem 4. For each nonzero boundary-only input difference α , the tables give $\dim K_H(\alpha)$ in the three admissible parameter cases $(b, \hat{b}) = (0, 0), (0, 1), (1, 0)$. All internal input-difference coordinates are zero.

D.1 The Long–Long Case

Assume $m, m' \geq 5$, and retain $\ell = (m + 1)/2$ and $\hat{\ell} = (m' + 1)/2$ from (28). Substitution of the all-length anchor–direction descriptions in Table 4 into (27) gives the following exact dimensions for all odd branch lengths in the stated range.

Table 9: Dimensions of $K_H(\alpha)$ for $m, m' \geq 5$.

α	$(b, \hat{b}) = (0, 0)$	$(b, \hat{b}) = (0, 1)$	$(b, \hat{b}) = (1, 0)$
e_{y_0}	2	2	2
e_{y_1}	ℓ	2	ℓ
$e_{y_0} + e_{y_1}$	0	0	0
$e_{\hat{y}_0}$	2	2	2
$e_{y_0} + e_{\hat{y}_0}$	2	2	2
$e_{y_1} + e_{\hat{y}_0}$	1	1	1
$e_{y_0} + e_{y_1} + e_{\hat{y}_0}$	0	0	0
$e_{\hat{y}_1}$	$\hat{\ell}$	$\hat{\ell}$	2
$e_{y_0} + e_{\hat{y}_1}$	2	2	$\hat{\ell}$
$e_{y_1} + e_{\hat{y}_1}$	1	1	1
$e_{y_0} + e_{y_1} + e_{\hat{y}_1}$	1	0	1
$e_{\hat{y}_0} + e_{\hat{y}_1}$	0	0	1
$e_{y_0} + e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	0
$e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	0	0	0
$e_{y_0} + e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	0	0	0

D.2 Exact Forms of the Length-Dependent Kernels

Among the entries of Table 9, the following are exactly the rows whose kernels depend on a branch length. The spaces $\mathcal{L}_m$ and $\mathcal{R}_{m'}$ are defined in (29).

$(b, \hat{b})$	α	$K_H(\alpha)$
$(0, 0)$	e_{y_1}	$\mathcal{L}_m$
$(0, 0)$	$e_{\hat{y}_1}$	$\mathcal{R}_{m'}$
$(0, 1)$	$e_{\hat{y}_1}$	$\mathcal{R}_{m'}$
$(1, 0)$	e_{y_1}	$\mathcal{L}_m$
$(1, 0)$	$e_{y_0} + e_{\hat{y}_1}$	$\mathcal{R}_{m'}$

(64)

For example, for $(b, \hat{b}) = (0, 0)$ and $\alpha = e_{y_1}$, row reduction gives

$$\text{Aff}(S_{e_{y_1}}) = \left\{ (d, \hat{d}) : d_0 = 0, d_3 = d_5 = \dots = d_{m-2} = 0, \hat{d}_0 = 1 \right\}.$$

There is no further affine equation, so the annihilator of its direction space is $\mathcal{L}_m$. The other rows follow by substituting the displayed α and $(b, \hat{b})$ into (27) and applying the same binary row reduction.

D.3 The Short–Long, Long–Short, and Short–Short Cases

It remains to consider $m = 3$ or $m' = 3$. The three tables below use the length-three sets in Lemma 4 and cover all $15 \cdot 3 \cdot 3 = 135$ boundary-only cases.

Table 10: Dimensions of $K_H(\alpha)$ for $m = 3$ and $m' \geq 5$.

α	(0, 0)	(0, 1)	(1, 0)
e_{y_0}	2	2	2
e_{y_1}	2	2	2
$e_{y_0} + e_{y_1}$	0	1	0
$e_{\hat{y}_0}$	2	2	2
$e_{y_0} + e_{\hat{y}_0}$	2	2	2
$e_{y_1} + e_{\hat{y}_0}$	2	2	2
$e_{y_0} + e_{y_1} + e_{\hat{y}_0}$	1	0	1
$e_{\hat{y}_1}$	$\hat{\ell}$	$\hat{\ell}$	2
$e_{y_0} + e_{\hat{y}_1}$	2	2	$\hat{\ell}$
$e_{y_1} + e_{\hat{y}_1}$	1	1	1
$e_{y_0} + e_{y_1} + e_{\hat{y}_1}$	1	1	1
$e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	2
$e_{y_0} + e_{\hat{y}_0} + e_{\hat{y}_1}$	2	2	1
$e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	1
$e_{y_0} + e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	1

Table 11: Dimensions of $K_H(\alpha)$ for $m \geq 5$ and $m' = 3$.

α	(0, 0)	(0, 1)	(1, 0)
e_{y_0}	2	2	2
e_{y_1}	ℓ	2	ℓ
$e_{y_0} + e_{y_1}$	1	2	1
$e_{\hat{y}_0}$	2	2	2
$e_{y_0} + e_{\hat{y}_0}$	2	2	2
$e_{y_1} + e_{\hat{y}_0}$	1	1	1
$e_{y_0} + e_{y_1} + e_{\hat{y}_0}$	1	1	1
$e_{\hat{y}_1}$	2	2	2
$e_{y_0} + e_{\hat{y}_1}$	2	2	2
$e_{y_1} + e_{\hat{y}_1}$	2	ℓ	2
$e_{y_0} + e_{y_1} + e_{\hat{y}_1}$	2	1	2
$e_{\hat{y}_0} + e_{\hat{y}_1}$	0	0	1
$e_{y_0} + e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	0
$e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	1
$e_{y_0} + e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	1

Table 12: Dimensions of $K_H(\alpha)$ for $(m, m') = (3, 3)$.

α	(0, 0)	(0, 1)	(1, 0)
e_{y_0}	2	2	2
e_{y_1}	2	2	2
$e_{y_0} + e_{y_1}$	1	2	1
$e_{\hat{y}_0}$	2	2	2
$e_{y_0} + e_{\hat{y}_0}$	2	2	2
$e_{y_1} + e_{\hat{y}_0}$	2	2	2
$e_{y_0} + e_{y_1} + e_{\hat{y}_0}$	2	1	2
$e_{\hat{y}_1}$	2	2	2
$e_{y_0} + e_{\hat{y}_1}$	2	2	2
$e_{y_1} + e_{\hat{y}_1}$	2	2	2
$e_{y_0} + e_{y_1} + e_{\hat{y}_1}$	2	1	2
$e_{\hat{y}_0} + e_{\hat{y}_1}$	1	1	2
$e_{y_0} + e_{\hat{y}_0} + e_{\hat{y}_1}$	2	2	1
$e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	2	2	1
$e_{y_0} + e_{y_1} + e_{\hat{y}_0} + e_{\hat{y}_1}$	1	2	2

Together with Table 9, these tables cover all four long/short geometries. The only entries above two are the length-dependent values ℓ and $\hat{\ell}$ identified in (30); hence they give exactly the boundary-only estimates used in Theorem 4.

E Trivial Kernel Intersections in Three-Dimensional Boundary Subspaces

This appendix supplies the boundary-direction classification used in Lemma 26. For each three-dimensional boundary subspace, the tables select one or two input differences whose kernels have trivial intersection.

E.1 Structure and Interpretation of the Tables

The four-bit row label $c = (c_0, c_1, \hat{c}_0, \hat{c}_1) \neq 0$ is the coefficient vector of the hyperplane A_c defined by (46); it is not an input difference. A four-bit vector in an entry denotes $\alpha_{\text{bd}} = (\alpha_0, \alpha_1, \hat{\alpha}_0, \hat{\alpha}_1)$, with all internal input-difference coordinates zero. Every displayed α_{bd} lies in the hyperplane indexing its row.

An entry $\{\alpha_{\text{bd}}\}$ asserts $K_H(\alpha_{\text{bd}}) = \{0\}$. An entry $\{\alpha_{\text{bd}}, \alpha'_{\text{bd}}\}$ asserts

$$K_H(\alpha_{\text{bd}}) \cap K_H(\alpha'_{\text{bd}}) = \{0\}.$$

Thus a two-element entry records two witness directions, not generators of one kernel or of A_c . The columns correspond to $(b, \hat{b}) = (0, 0), (0, 1), (1, 0)$. All entries are obtained from the exact system (27), using the affine hulls in Table 4.

E.2 The Long–Long Case

Table 13: Kernel-intersection witnesses in A_c for $m, m' \geq 5$.

hyperplane coefficient c	$(0, 0)$	$(0, 1)$	$(1, 0)$
1000	$\{0011\}$	$\{0011\}$	$\{0111\}$
0100	$\{0011\}$	$\{0011\}$	$\{1011\}$
1100	$\{1100\}$	$\{1100\}$	$\{1100\}$
0010	$\{1100\}$	$\{1100\}$	$\{1100\}$
1010	$\{1110\}$	$\{1110\}$	$\{1110\}$
0110	$\{1110\}$	$\{1110\}$	$\{1110\}$
1110	$\{1100\}$	$\{1100\}$	$\{1100\}$
0001	$\{1100\}$	$\{1100\}$	$\{1100\}$
1001	$\{1111\}$	$\{1101\}$	$\{1011\}$
0101	$\{0111\}$	$\{1101\}$	$\{0111\}$
1101	$\{1100\}$	$\{1100\}$	$\{1100\}$
0011	$\{1100\}$	$\{1100\}$	$\{1100\}$
1011	$\{1110\}$	$\{1110\}$	$\{1110\}$
0111	$\{1110\}$	$\{1110\}$	$\{1110\}$
1111	$\{1100\}$	$\{1100\}$	$\{1100\}$

E.3 The Short–Long and Long–Short Cases

Table 14: Kernel-intersection witnesses in A_c for $m = 3$ and $m' \geq 5$.

hyperplane coefficient c	(0, 0)	(0, 1)	(1, 0)
1000	{0100, 0011}	{0100, 0011}	{0100, 0011}
0100	{1000, 0011}	{1000, 0011}	{1000, 0011}
1100	{1100}	{1110}	{1100}
0010	{1100}	{1000, 0100}	{1100}
1010	{0100, 1010}	{1110}	{0100, 1010}
0110	{1000, 0111}	{1110}	{1000, 0111}
1110	{1100}	{1100, 0001}	{1100}
0001	{1100}	{1110}	{1100}
1001	{0100, 1011}	{0100, 1001}	{0100, 1011}
0101	{1000, 0101}	{1000, 0101}	{1000, 0101}
1101	{1100}	{1110}	{1100}
0011	{1100}	{1000, 0100}	{1100}
1011	{0100, 1010}	{1110}	{0100, 1010}
0111	{1000, 0101}	{1110}	{1000, 0101}
1111	{1100}	{1100, 1001}	{1100}

Table 15: Kernel-intersection witnesses in A_c for $m \geq 5$ and $m' = 3$.

hyperplane coefficient c	(0, 0)	(0, 1)	(1, 0)
1000	{0011}	{0011}	{0100, 0011}
0100	{0011}	{0011}	{1011}
1100	{0011}	{0011}	{1100, 0010}
0010	{1000, 0100}	{1000, 0100}	{1000, 0100}
1010	{0100, 1010}	{0100, 1010}	{1011}
0110	{1000, 0110}	{1000, 0110}	{1000, 0110}
1110	{1100, 1010}	{1100, 1010}	{1011}
0001	{1000, 0100}	{1000, 0100}	{1000, 0100}
1001	{0100, 1011}	{0100, 1001}	{1011}
0101	{1000, 0101}	{1000, 0101}	{1000, 0101}
1101	{1100, 0010}	{1100, 0010}	{1011}
0011	{0011}	{0011}	{1011}
1011	{0011}	{0011}	{0100, 1010}
0111	{0011}	{0011}	{1011}
1111	{0011}	{0011}	{1100, 1010}

E.4 The Short–Short Case

Table 16: Kernel-intersection witnesses in A_c for $m = m' = 3$.

hyperplane coefficient c	(0, 0)	(0, 1)	(1, 0)
1000	{0100, 0011}	{0100, 0011}	{0100, 0011}
0100	{1000, 0011}	{1000, 0011}	{1000, 0011}
1100	{1100, 0010}	{1100, 0010}	{1100, 0010}
0010	{1000, 0100}	{1000, 0100}	{1000, 0100}
1010	{0100, 1010}	{0100, 1010}	{0100, 1010}
0110	{1000, 0111}	{1000, 0110}	{1000, 0111}
1110	{1100, 1010}	{1100, 0001}	{1100, 1010}
0001	{1000, 0100}	{1000, 0100}	{1000, 0100}
1001	{0100, 1011}	{0100, 1001}	{0100, 1011}
0101	{1000, 0101}	{1000, 0101}	{1000, 0101}
1101	{1100, 0010}	{1100, 0010}	{1100, 0010}
0011	{1000, 0100}	{1000, 0100}	{1000, 0100}
1011	{0100, 1010}	{0100, 1010}	{0100, 1010}
0111	{1000, 0101}	{1000, 0110}	{1000, 0101}
1111	{1100, 1010}	{1100, 0011}	{1100, 1010}

The four tables cover all fifteen nonzero coefficient vectors, all three parameter pairs, and all four long/short geometries. Their displayed intersections are exactly those required in Lemma 26.

F Direction Constructions for the SWAP Family

We give the constructions used in Lemmas 14 and 15. The single-cycle mask constraint is already proved in Proposition 2; here we use it to obtain directions within fixed-endpoint sets and then combine compatible local differences. Write

$$S_\gamma^\chi = \{d : \Delta_d \chi_r(x) = \gamma \text{ for some } x \in \mathbb{F}_2^r\}$$

for an ordinary-cycle attainable set, to distinguish it from the global SWAP set S_α .

Proof of Lemma 14. Put $r = m' \geq 5$. For $\gamma \in \mathbb{F}_2^r$, define

$$\begin{aligned} T_{\gamma, \varepsilon} &= \{d \in S_\gamma^\chi : d_0 = \varepsilon\}, \\ \mathcal{B}_r(\gamma) &= \sum_{\varepsilon \in \mathbb{F}_2} \text{span}\{d + d' : d, d' \in T_{\gamma, \varepsilon}\}, \\ M_r(\gamma) &= (\text{dir Aff } S_\gamma^\chi)^\perp. \end{aligned}$$

An empty endpoint set contributes the zero subspace. A linear functional constant on each nonempty endpoint set has the form $\langle v, d \rangle = a + bd_0$ on S_γ^χ , including when only one endpoint is attainable. Hence

$$\mathcal{B}_r(\gamma)^\perp = M_r(\gamma) + \langle e_0 \rangle.$$

We claim that for every $\beta \in \mathbb{F}_2^r$,

$$\mathcal{B}_r(\beta) + \mathcal{B}_r(\beta + e_0) = \langle e_1, \dots, e_{r-1} \rangle.$$

The orthogonal complement of the left-hand side is

$$(M_r(\beta) + \langle e_0 \rangle) \cap (M_r(\beta + e_0) + \langle e_0 \rangle),$$

which always contains $\langle e_0 \rangle$. If $\{\beta, \beta + e_0\} = \{0, e_0\}$, Proposition 2 applied to e_0 gives the reverse inclusion. The same is true if either nonzero output has $M_r = \{0\}$. In every remaining case, the same proposition restricts each output to a unit vector or an adjacent pair. Toggling coordinate zero leaves only the following possibilities, with the indicated upper bounds on the intersection:

$$\begin{aligned} \{e_1, e_0 + e_1\} : \quad & \langle e_0, e_1 \rangle \cap \langle e_0, e_1 + e_3 \rangle = \langle e_0 \rangle, \\ \{e_{r-1}, e_{r-1} + e_0\} : \quad & \langle e_0, e_{r-1} \rangle \cap \langle e_0, e_2 \rangle = \langle e_0 \rangle. \end{aligned}$$

The internal coordinates appearing here are distinct because $r \geq 5$. Taking orthogonal complements proves the claim.

Now fix $\sigma, \varepsilon \in \mathbb{F}_2$ and put $c = \sigma + \varepsilon$. Since $\alpha_L^{\text{int}} \neq 0$, Lemma 9 supplies

$$a \in S_{\alpha_L + \sigma e_0}^\chi, \quad a_0 = c.$$

For any $b, b' \in T_{\alpha_R + \sigma e_0, \varepsilon}$, Lemma 8 gives $(a, b), (a, b') \in S_\alpha$, since $c + \varepsilon = \sigma$. Thus

$$(0, b + b') = (a, b) + (a, b') \in \text{dir}(D_\alpha).$$

Taking spans within each endpoint set and summing over the four choices of (σ, ε) , the claim gives

$$\{0\} \times \langle e_1, \dots, e_{m'-1} \rangle \subseteq \text{dir}(D_\alpha).$$

Different endpoint sets may use different left differences a : each pair cancels its own left branch before the directions are added. This proves the first assertion. The other follows by exchanging the branches. $\square$

Proof of Lemma 15. Let

$$k = \max\{i \in \{1, \dots, m-1\} : \alpha_i = 1\}, \quad a = \begin{cases} e_1, & k = 1, \\ \sum_{i=2}^k e_i, & k \geq 2. \end{cases}$$

Then $a_0 = 0$, and we claim that

$$a \in S_{\alpha_L}^X \cap S_{\alpha_L + e_0}^X.$$

For $k = 1$, its zero set is the cyclic interval $[2, 0]$ of length $m - 1$. Lemma 3 requires only $\alpha_1 = 1$ and $\alpha_i = 0$ for $2 \leq i \leq m - 2$. For $k \geq 2$, the zero set is $[k + 1, 1]$ of length $m - k + 1$; the same lemma requires only $\alpha_k = 1$ and $\alpha_i = 0$ for $k < i \leq m - 1$. These conditions follow from maximality of k , and none uses α_0 , proving the claim.

Since $\hat{\alpha}_0 = 0$, we have $\alpha_R \neq e_0$ and $\alpha_R + e_0 \neq 0$. Lemma 9 therefore allows choices

$$b \in S_{\alpha_R}^X, \quad b_0 = 0, \quad b' \in S_{\alpha_R + e_0}^X, \quad b'_0 = 1.$$

Lemma 8 at the endpoint pairs $(0, 0)$ and $(0, 1)$ gives $(a, b), (a, b') \in S_\alpha$. Their sum

$$\delta d = (a, b) + (a, b') = (0, b + b') \in \text{dir}(D_\alpha)$$

has every left coordinate zero and right endpoint $b_0 + b'_0 = 1$, as required. Exchanging the branches proves the symmetric assertion. $\square$