THREE-ROUND INVERSION BYPASS, DIVISION-FREE LUCAS RESULTANTS, AND PRACTICAL 8/10-ROUND KOALABEAR & 6/9-ROUND GOLDILOCKS CICO-1 ATTACKS ON GRIFFIN

SETH REHWALDT

Abstract. We present a hierarchy of algebraic subspace trails, exact univariate tower-resultant degree collapses, saturation obstructions, and memory-frugal algorithms for the single-input/single-output Constrained-Input Constrained-Output (CICO-1) problem on the arithmetization-oriented permutation Griffin [4] over prime fields $\mathbb{F}_p$ of odd S-box degree $d \geq 3$ and state width $t = 4k \geq 12$.

First, exploiting the block-Kronecker structure $M = (I_k + J_k) \otimes M_4^{\mathrm{HL}}$ of Griffin's linear layer together with alternating-zero Horst inputs and anti-symmetric block pairs, we bypass all three high-degree $x^{1/d}$ inversions in Rounds 0, 1, and 2 simultaneously on every width $t = 4k \geq 16$, obtaining a closed-form univariate tower-resultant degree $D^{(7)}(R, d) = D_3 \cdot (d(2d+1))^{R-3}$ for all $R \geq 4$. On widths $t = 4k \geq 20$, we show that two extra degree collapses occur in Round 2 and Round 3: (i) the homogeneous quadratic parts of the Round-2 vanishing conditions inherit an exact 4:5 rank-1 proportionality from M_4^{HL}, dropping their bivariate resultant from degree 4 to degree 2 and forcing both $s_0^{(2)}$ and $s_1^{(2)}$ constant; and (ii) the leading coefficient $[u^8]s_1^{(3)}$ is strictly affine-linear in a single decoupled sum parameter τ_{last}, yielding $D^{(8B)}(R, d) = (14d+8)d(d(2d+1))^{R-4}$ (a 2.264× reduction on $d = 7$ that pushes 6-round Goldilocks below 2^{23}). On widths $t = 4k \geq 24$, we prove an additional d-independent quadratic annihilation of $[u^7]s_1^{(3)}$ via a degree-4 precomputation, reducing the degree further to $D^{(8C)}(R, d) = (12d+8)d(d(2d+1))^{R-4}$ (2.609× below the 3-round bypass for $d = 7$).

Second, we establish sharp structural barriers: $[u^6]s_1^{(3)}$ is the exact quadratic-precomputation saturation boundary; a 2-round pure-linear bypass on the anti-symmetric plane V_- reduces to the Diophantine equation $(2a-3)(2b-3) = 15$ yet admits zero Block-0-compatible solutions across all 89,760 configurations for $k \in \{5, 6, 7, 8\}$; and final-round backward-peel fails for CICO-1 because row 0 of both M and M^{-1} is completely dense.

Third, we introduce four algorithmic advances: a universal division-free odd-d Lucas resultant eliminating all tower inversions, a cyclic wrap-around Euclidean Half-GCD halving top-level NTT memory, an exact 127-coset complete-field forward-NTT root-finder over KoalaBear ($p-1 = 127 \cdot 2^{24}$), and a 16-point Radix-2 NTT over the degree-7 extension ring using $\omega_{16} = 2^{12}$ over 64-bit Goldilocks ($p = 2^{64} - 2^{32} + 1$). Combining these yields 17 zero-float verified explicit CICO-1 witnesses on unmodified Griffin, highlighted by 8/10-round KoalaBear Griffin-20 (deg = 29,172,150 in 1,770 s and < 260 MB RAM), 6/9-round (2/3 full-round) 64-bit Goldilocks Griffin-20 (deg = 8,180,550 in 625 s and < 750 MB RAM, shattering 2^{64} brute force), and 5/9-round 64-bit Goldilocks Griffin-24 (deg = 67,620 in 48.7 s). All five core algebraic foundations are machine-checked in zero-dependency Lean 4.36.0 with zero `sorry` and zero `native_decide`.

1. Introduction

Arithmetization-oriented permutations minimize multiplicative depth and R1CS/PLONK gate counts inside zero-knowledge proof systems over large prime fields $\mathbb{F}_p$. Griffin [4] (CRYPTO 2023) combines a single high-degree inverse power map $x \mapsto x^{1/d}$ on lane 0 and a low-degree forward power map $x \mapsto x^d$ on lane 1 with a quadratic Horst layer on lanes $2, \ldots, t-1$, interleaved with the Poseidon2-style block-Kronecker linear layer $M = (I_k + J_k) \otimes M_4^{\mathrm{HL}}$ for widths $t = 4k \geq 8$.

Date: November 27, 2026.

The security of sponge-based hash functions built from a permutation $\Pi : \mathbb{F}_p^t \to \mathbb{F}_p^t$ against algebraic preimage and collision attacks is benchmarked by the Constrained-Input Constrained-Output (CICO) problem. In strict CICO-1, the adversary seeks a nonzero input $x \in \mathbb{F}_p^t$ satisfying

$$x_0 = 0 \qquad \text{and} \qquad \Pi(x)_0 = 0. \tag{1}$$

Notice that CICO-1 has expected brute-force cost $O(p)$ (e.g., 2^{31} over KoalaBear and 2^{64} over 64-bit Goldilocks), whereas full-round Griffin targets 128-bit security with $R = 10$ rounds for $d = 3$ ($t \geq 16$) and $R = 9$ rounds for $d \in \{5, 7\}$ ($t \geq 12$).

Following the *Algebraic FreeLunch* framework of Bariant et al. [3] (CRYPTO 2024), which solved 7/10-round Griffin via Gröbner bases in several hours, subsequent work explored multi-output CICO extensions (*CheapLunch* [1]) and fast multivariate resultant elimination (*Resultants Meet Resultant* [2]). While [2] demonstrates dramatic speedups on multi-output CICO-2 and structured full-target matching (SFTM) variants using Villard's bivariate resultant algorithm, the strict single-variable CICO-1 frontier on wide-width Griffin ($t \in \{16, 20, 24\}$) over both 31-bit and 64-bit prime fields—and the exact algebraic limit of round-skipping trails within the $M = (I_k + J_k) \otimes M_4^{\mathrm{HL}}$ linear layer—has remained open.

1.1. **Contributions.** This paper settles the univariate CICO-1 subspace-trail frontier for Griffin on widths $t = 4k \geq 12$ across four pillars:

(1) **Three-Round Inversion Bypass and Wide-Width Degree-Collapse Hierarchy (Section 3).** We establish a progression of exact closed-form univariate tower-resultant degrees for all odd $d \geq 3$ and all $R \geq 4$:
- *Theorem 3.1 ($t = 4k \geq 12$):* Alternating-zero Horst inputs linearize Round 0 and collapse Round 1 to degree ≤ 3 independent of d.
- *Theorem 3.2 ($t = 4k \geq 16$):* An anti-symmetric block pair on $(\mathrm{Blk}\,1, \mathrm{Blk}\,3)$ bypasses all three $x^{1/d}$ inversions in Rounds $0, 1, 2$, yielding exact R-round CICO-1 degree $D^{(7)}(R, d) = D_3 \cdot (d(2d+1))^{R-3}$ ($D_3 = 9, 12, 16$ for $d = 3, 5, 7$).
- *Theorem 3.3 ($t = 4k \geq 20$):* Exact 4:5 rank-1 proportionality inherited from M_4^{HL} collapses the Round-2 quadratic resultant from degree 4 to 2, forcing both $s_0^{(2)}$ and $s_1^{(2)}$ constant ($D^{(8A)}(R, d) = 8(d(2d+1))^{R-3}$). Simultaneously, $[u^8]s_1^{(3)}$ is strictly affine-linear in a single decoupled parameter $\tau_{\mathrm{last}} = s_4^{(1)} + s_{12}^{(1)}$, annihilating $[u^8]s_1^{(3)}$ for free and dropping the degree to
$$D^{(8B)}(R, d) = (14d + 8) \cdot d \cdot \big(d(2d+1)\big)^{R-4}.$$
For $d = 7$ (Goldilocks), this cuts the degree by 2.264× relative to Theorem 3.2 and pushes 6-round Goldilocks from 9,261,000 down to 8,180,550 $< 2^{23}$.
- *Theorem 3.4 ($t = 4k \geq 24$):* The extra constant-block parameters $(\tau_0, \ldots, \tau_{2k-9})$ yield d-independent quadratics for $[u^7]s_0^{(3)}, [u^7]s_1^{(3)}$ whose homogeneous quadratic parts form a rank-2 subspace with the Round-2 quadratics (Q_0, Q_1). Eliminating one variable linearly via $5Q_1 - 4Q_0 = 0$ leaves a degree-4 bivariate resultant that forces $\deg_u(s_0^{(3)}, s_1^{(3)}) = (8, 6)$, yielding
$$D^{(8C)}(R, d) = (12d + 8) \cdot d \cdot \big(d(2d+1)\big)^{R-4},$$
an additional 12.0%–13.2% reduction (2.609× below Theorem 3.2 for $d = 7$).

(2) **Exact Saturation and Structural Obstructions (Theorem 3.6).** We prove that $[u^6]s_1^{(3)}$ is the exact algebraic ceiling for d-independent quadratic precomputation on 1-pair trails (since $s_0^{(2)}[0], s_1^{(2)}[0]$ enter cubically through $y_0^{(2)}, y_1^{(2)}$). For 2-round pure-linear bypasses, we derive the exact M_4^{HL} anti-symmetric plane Diophantine equation $(2a - 3)(2b - 3) = 15$ (whose integer roots include $k = 8$ and $k = 12$), yet prove by exhaustive elimination over all 89,760 Block-0-compatible multiset configurations that zero solutions exist for $k \in \{5, 6, 7, 8\}$.

Finally, we prove that final-round backward-peel cannot fix y_0 for CICO-1 because row 0 of M and M^{-1} is t/t dense.

(3) **Algorithmic Engines and 17 Verified Practical Records (Sections 4–5).** We introduce a division-free odd-d Lucas resultant (Proposition 4.1), cyclic wrap-around Euclidean Half-GCD (Proposition 4.2), complete-field 127-coset forward-NTT root-finding over KoalaBear (Proposition 4.3), and 16-point Radix-2 NTT over the Goldilocks extension ring Au_7 using $\omega_{16} = 2^{12}$. Our multi-threaded Rust solvers extract 17 explicit CICO-1 witnesses verified in pure zero-float Python, including:
- 8/10-**round KoalaBear Griffin-20** ($p = 2^{31} - 2^{24} + 1, d = 3$): deg = 29,172,150 solved in 1,770 s (29.5 min) at < 260 MB peak RAM.
- 6/9-**round (2/3 full-round) 64-bit Goldilocks Griffin-20** ($p = 2^{64} - 2^{32} + 1, d = 7$): deg = 8,180,550 solved in 625 s (10.4 min) at < 750 MB peak RAM, decisively shattering 2^{64} brute force.
- 5/9-**round 64-bit Goldilocks Griffin-24** (Theorem 3.4): deg = 67,620 solved in 48.7 s at < 30 MB peak RAM.

Full-round streaming tower-NTT complexity drops to $2^{52.8}$ ($d = 3, R = 10$), $2^{60.0}$ ($d = 5, R = 9$), and $2^{69.5}$ ($d = 7, R = 9$) field operations—58.5 to 75.2 bits below 2^{128}.

(4) **Zero-Dependency Lean 4.36.0 Formalization (Section 6).** All five core algebraic results (block-Kronecker inversion and zero-block-sum invariance, universal odd-d division-free Lucas resultant, cyclic wrap-around Half-GCD, multi-coset DFT folding and Goldilocks ω_{16} order, and the exact V_- Diophantine classification together with the degree closed forms) are formally verified in a 1,598-line zero-dependency Lean 4.36.0 kernel (`CryptoAlgebra`) with zero `sorry`, zero `native_decide`, and only the standard axioms `propext`, `Classical.choice`, and `Quot.sound`.

2. Preliminaries: Griffin and Block-Kronecker Linear Algebra

Let $\mathbb{F}_p$ be a prime field and let $d \geq 3$ be the smallest odd integer with $\gcd(d, p-1) = 1$, with inverse exponent $d^{-1} \bmod (p-1)$. Specifically, we study KoalaBear ($p_{\mathrm{KB}} = 2^{31} - 2^{24} + 1 = 2{,}130{,}706{,}433$, $d = 3$) and 64-bit Goldilocks ($p_{\mathrm{GL}} = 2^{64} - 2^{32} + 1 = 18{,}446{,}744{,}069{,}414{,}584{,}321$, $d = 7$), as well as general $d = 5$ instances.

2.1. The Griffin Permutation.

For width $t = 4k \geq 8$, the linear layer of Griffin [4] is the block-Kronecker matrix

$$M \;=\; (I_k + J_k) \otimes M_4^{\mathrm{HL}} \;\in\; \mathbb{F}_p^{t \times t}, \tag{2}$$

where J_k is the $k \times k$ all-ones matrix and M_4^{HL} is the Horizen Labs 4×4 MDS circulant-like matrix

$$M_4^{\mathrm{HL}} \;=\; \begin{pmatrix} 5 & 7 & 1 & 3 \\ 4 & 6 & 1 & 1 \\ 1 & 3 & 5 & 7 \\ 1 & 1 & 4 & 6 \end{pmatrix}. \tag{3}$$

Direct computation over $\mathbb{Z}$ gives $\det(M_4^{\mathrm{HL}}) = -64$ and an exact integral scaled inverse:

$$\mathrm{adj}_8(M_4^{\mathrm{HL}}) \;=\; \begin{pmatrix} 1 & -3 & 6 & -7 \\ -1 & 1 & -4 & 5 \\ 6 & -7 & 1 & -3 \\ -4 & 5 & -1 & 1 \end{pmatrix}, \qquad \mathrm{adj}_8(M_4^{\mathrm{HL}}) \cdot M_4^{\mathrm{HL}} \;=\; -8\, I_4. \tag{4}$$

Thus M_4^{HL} is invertible over every commutative ring in which 2 is a unit.

An R-round Griffin permutation begins with an initial whitening layer $s^{(0)} = Mx \in \mathbb{F}_p^t$. In each round $r \in \{0, \dots, R-1\}$, the nonlinear layer $\mathcal{S} : s^{(r)} \mapsto y^{(r)}$ computes

$$y_0^{(r)} = \big(s_0^{(r)}\big)^{1/d}, \qquad y_1^{(r)} = \big(s_1^{(r)}\big)^{d}, \tag{5}$$

$$y_i^{(r)} = s_i^{(r)} \cdot G_i\big(L_i^{(r)}\big) \qquad (2 \le i < t), \tag{6}$$

where $G_i(z) = z^2 + \alpha_i z + \beta_i$ is irreducible over $\mathbb{F}_p$ (with $\alpha_i = (i-1)\alpha_2$, $\beta_i = (i-1)^2\beta_2$, $\alpha_2^2 - 4\beta_2 \notin (\mathbb{F}_p^\times)^2$), and the Horst linear forms are

$$L_2^{(r)} = y_0^{(r)} + y_1^{(r)}, \qquad L_i^{(r)} = (i-1)y_0^{(r)} + y_1^{(r)} + s_{i-1}^{(r)} \quad (3 \le i < t). \tag{7}$$

Finally, the round updates $s^{(r+1)} = My^{(r)} + \mathrm{rc}^{(r)}$, with $\mathrm{rc}^{(r)} \in \mathbb{F}_p^t$ derived deterministically via SHAKE256. Writing a state vector as k blocks $v = (v^{[0]}, \dots, v^{[k-1]})$ with $v^{[b]} \in \mathbb{F}_p^4$, the matrix M satisfies the following elementary but decisive structural identity.

Lemma 2.1 (Block-Kronecker Inversion and Zero-Sum Diagonalization)**.** *Let R be any commutative ring in which $k+1$ and 2 are invertible, and let $M = (I_k + J_k) \otimes M_4$ with M_4 invertible over R. Then:*

1. $M^{-1} = \big(I_k - (k+1)^{-1}J_k\big) \otimes M_4^{-1}$.
2. *For every vector $v = (v^{[0]}, \dots, v^{[k-1]}) \in R^{4k}$ satisfying $\sum_{b=0}^{k-1} v^{[b]} = 0$, the matrix M acts block-diagonally:*
$$(Mv)^{[b]} = M_4 v^{[b]} \qquad (0 \le b < k),$$
and symmetrically $(M^{-1}v)^{[b]} = M_4^{-1}v^{[b]}$.

Proof. Since $(I_k+J_k)(I_k-(k+1)^{-1}J_k) = I_k + J_k - (k+1)^{-1}J_k - k(k+1)^{-1}J_k = I_k$ (using $J_k^2 = kJ_k$), mixed-product Kronecker multiplication yields (1). For (2), $(Mv)^{[b]} = M_4 v^{[b]} + M_4 \sum_{c=0}^{k-1} v^{[c]} = M_4 v^{[b]}$, and identically for M^{-1}. □

A direct consequence of Lemma 2.1(2) is that whenever an anti-symmetric pair of blocks satisfies $y^{[b_1]}(u) + y^{[b_2]}(u) = c$ (constant in an indeterminate u) while all remaining blocks $y^{[b]}(u)$ ($b \notin \{b_1, b_2\}$) have lower degree in u, the high-degree terms of $y^{[b_1]}$ and $y^{[b_2]}$ cancel completely in the global block sum $\sum_{b=0}^{k-1} y^{[b]}$ and therefore *never leak into Block* 0 (or any other block outside $\{b_1, b_2\}$). This cancellation mechanism drives every theorem below.

3. Three-Round Inversion Bypass and Degree-Collapse Hierarchy

We parameterize input states $x(u) \in \mathbb{F}_p[u]^t$ (with $x_0(u) = 0$) in a single indeterminate u, aiming to keep the state coordinates $s_0^{(r)}(u)$ and $s_1^{(r)}(u)$ constant (or of minimal degree) for as many rounds as possible before entering the iterated radical tower.

3.1. Alternating-Zero Horst Linearization (Round 0 and Round 1).

Theorem 3.1 (Alternating-Zero Horst Bypass and Even-Subspace Collapse — Thm 6)**.** *Let $t = 4k \ge 12$ and let $s^{(0)} = Mx$ with $x_0 = 0$. Fix constants $c_0, c_1 \in \mathbb{F}_p$ and impose the linear subspace constraints*

$$s_0^{(0)} = c_0, \qquad s_1^{(0)} = c_1, \qquad s_{2j+1}^{(0)} = 0 \quad \text{for all } 1 \le j \le 2k-1. \tag{8}$$

Then:

1. *The affine input subspace satisfying $x_0 = 0$ and (8) has dimension $2k-2$ (parametrized by the $2k-2$ even-indexed coordinates $s_{2j}^{(0)}$ modulo one linear relation from $x_0 = 0$).*
2. *Round 0 is* completely affine-linear *on this subspace: every $y_i^{(0)}$ $(0 \le i < t)$ has degree at most 1 in the free variables, with $y_{2j+1}^{(0)} = 0$ for all $j \ge 1$, and $y_{2j}^{(0)} = G_{2j}\big((2j-1)c_0^{1/d} + c_1^d\big) \cdot s_{2j}^{(0)}$ for all $j \ge 1$.*

(3) *For $k \geq 3$, every vector $y^{(0)} - y^{(0)}_{\text{const}}$ lies in the even-coordinate subspace $\bigoplus_{b=0}^{k-1} \operatorname{span}(e_0, e_2) \subset \mathbb{F}_p^{4k}$. Because $M_4^{\text{HL}} \operatorname{span}(e_0, e_2)$ has dimension 2, forcing all four coordinates of Block 0 of $s^{(1)} = My^{(0)} + \text{rc}^{(0)}$ to be constant imposes only 2 linear equations on the $2k-2$ parameters—leaving a $(2k-4)$-dimensional affine subspace on which $s_0^{(1)}, s_1^{(1)}, s_2^{(1)}, s_3^{(1)}$ are constant and every Round-1 output $y_i^{(1)}$ has degree at most 3 independent of d.*

Proof. Under (8), $y_0^{(0)} = c_0^{1/d}$ and $y_1^{(0)} = c_1^d$ are constants. For every odd index $i = 2j+1 \geq 3$, $s_i^{(0)} = 0$ forces $y_i^{(0)} = s_i^{(0)} G_i(L_i^{(0)}) = 0$. For every even index $i = 2j \geq 2$, the Horst argument $L_{2j}^{(0)} = (2j-1)y_0^{(0)} + y_1^{(0)} + s_{2j-1}^{(0)}$ has $s_{2j-1}^{(0)} = 0$ (for $j \geq 2$, and absent for $j = 1$), so $L_{2j}^{(0)}$ is constant! Consequently $G_{2j}(L_{2j}^{(0)})$ is a nonzero constant scalar, making $y_{2j}^{(0)}$ strictly linear in $s_{2j}^{(0)}$.

Write $\Delta y^{(0)} = y^{(0)} - y^{(0)}(0)$. Its Block-0 components are $(0, 0, \Delta y_2^{(0)}, 0)^\top$, while each block $b \geq 1$ has $\Delta y^{(0),[b]} \in W_0 := \operatorname{span}(e_0, e_2)$. By (2),

$$\Delta s^{(1),[0]} = M_4^{\text{HL}}\Big(2\Delta y^{(0),[0]} + \sum_{b=1}^{k-1} \Delta y^{(0),[b]}\Big) \in M_4^{\text{HL}}(W_0).$$

Since $\dim M_4^{\text{HL}}(W_0) = 2$, the 4-coordinate condition $\Delta s^{(1),[0]} = 0$ is equivalent to 2 linear equations (namely the e_0- and e_2-sums vanish). With $s_0^{(1)}, s_1^{(1)}$ constant, $y_0^{(1)}$ and $y_1^{(1)}$ are constant, so every $L_i^{(1)}$ has degree ≤ 1 and every $y_i^{(1)} = s_i^{(1)} G_i(L_i^{(1)})$ has degree ≤ 3 independent of d. □

Theorem 3.1 already bypasses two $x^{1/d}$ inversions (Rounds 0 and 1), enabling 3-round CICO-1 witnesses of degree 63 (KoalaBear Griffin-12) and 315 (64-bit Goldilocks Griffin-12) in Section 5.

3.2. **Anti-Symmetric Block-Pair 3-Round Inversion Bypass ($t = 4k \geq 16$).** On widths $t = 4k \geq 16$, Theorem 3.1 leaves $2k - 2 \geq 6$ free Round-0 variables, which is enough to force a *third* consecutive inversion bypass in Round 2.

Theorem 3.2 (3-Round Inversion Bypass on Griffin-$4k$, $k \geq 4$ — Thm 7)**.** *Let $t = 4k \geq 16$ and let $d \geq 3$ be odd. Starting from the $(2k-2)$-dimensional Round-0 linear subspace* (8)*, impose on $s^{(1)} = My^{(0)} + \text{rc}^{(0)}$ the block constraints:*

(1) *Block 0 is constant ($\Delta s^{(1),[0]} = 0$);*
(2) *Every block $b \in \{2, 4, 5, \ldots, k-1\}$ is constant ($\Delta s^{(1),[b]} = 0$);*
(3) *Block 1 and Block 3 are anti-symmetric in the variable u: $\Delta s^{(1),[1]} + \Delta s^{(1),[3]} = 0$.*

Then:

(a) *Because every block difference lies in $M_4^{\text{HL}}(\operatorname{span}(e_0, e_2))$, the homogeneous linear system enforcing (1)–(3) has exact rank $2k-3$ (not $2k-2$!), leaving a 1-dimensional affine line $u \mapsto s^{(1)}(u)$ parametrized by u together with $2k-5$ free constant-block parameters.*

(b) *On this line, $s^{(1),[1]}(u) = a^{[1]} + uv$ and $s^{(1),[3]}(u) = a^{[3]} - uv$ for a nonzero direction $v \in M_4^{\text{HL}}(\operatorname{span}(e_0, e_2))$ with v_0, v_1, v_2, v_3 all nonzero, while all other blocks of $s^{(1)}$ are constant. In Round 1, every cubic leading term $[u^3]y_r^{(1)}$ ($r \in \{0, 1, 2, 3\}$) satisfies*

$$[u^3]\big(y_{4+r}^{(1)} + y_{12+r}^{(1)}\big) = 0,$$

because $y_{4+r}^{(1)} = s_{4+r}^{(1)}(L_{4+r}^{(1)})^2 + O(u^2)$ has leading coefficient $v_r v_{r-1}^2$ while $y_{12+r}^{(1)}$ has leading coefficient $(-v_r)(-v_{r-1})^2 = -v_r v_{r-1}^2$ (and $L_4^{(1)}, L_{12}^{(1)}$ depend on the constant $s_3^{(1)}, s_{11}^{(1)}$, so $y_4^{(1)}, y_{12}^{(1)}$ are already linear in u).

(c) *By Lemma 2.1(2),* $\deg_u(s_r^{(2)}) \le 2$ *for every* $r \in \{0,1,2,3\}$. *Enforcing* $[u^2]s_0^{(2)} = 0$ *is affine-linear in the free constants, and* $[u^1]s_0^{(2)} = 0$ *is a single quadratic equation over* $\mathbb{F}_p$, *forcing*

$$\deg_u\big(s_0^{(2)}\big) = 0 \quad (\textit{constant!}), \qquad \deg_u\big(s_1^{(2)}\big) = 1.$$

Thus all three $x^{1/d}$ *inversions in Rounds* 0, 1, *and* 2 *evaluate on constants, and the exact* R*-round* CICO-*1 univariate tower-resultant degree for every* $R \ge 4$ *is*

$$D^{(7)}(R,d) \;=\; D_3(d)\cdot\big(d(2d+1)\big)^{R-3}, \tag{9}$$

where $D_3(d) = \max(9, 2d+2)$ *gives* $D_3(3) = 9$, $D_3(5) = 12$, *and* $D_3(7) = 16$.

Proof. Write $w_b = \Delta y^{(0),[b]} \in W_0 = \mathrm{span}(e_0, e_2)$ for $b \in \{1,\dots,k-1\}$ and $w_0 = (0, 0, \Delta y_2^{(0)}, 0)^\top$. Then $\Delta s^{(1),[b]} = M_4^{\mathrm{HL}}(w_b + \Sigma)$ where $\Sigma = \sum_{c=0}^{k-1} w_c$. Setting $w_b = -\Sigma$ for $b \in \{0,2,4,\dots,k-1\}$ and $w_1 + w_3 = -2\Sigma$ forces $\Delta s^{(1),[b]} = 0$ for all $b \notin \{1,3\}$ and $\Delta s^{(1),[1]} + \Delta s^{(1),[3]} = 0$, leaving $w_1 - w_3 \in W_0$ completely free (dimension 2 before intersecting with the single scalar constraint $x_0 = 0$, hence dimension 1 in u).

In Round 1, $s^{(1),[0]}$ is constant, so $y_0^{(1)}, y_1^{(1)}$ are constant and $\Delta y^{(1),[b]} = 0$ for all $b \notin \{1,3\}$. For $r = 0$, $s_3^{(1)}$ and $s_{11}^{(1)}$ are constant, so $L_4^{(1)}$ and $L_{12}^{(1)}$ are constant and $y_4^{(1)}, y_{12}^{(1)}$ are affine-linear in u. For $r \in \{1,2,3\}$,

$$y_{4+r}^{(1)}(u) = (a_r^{[1]} + uv_r)\big((a_{r-1}^{[1]} + \gamma_{4+r} + uv_{r-1})^2 + \alpha_{4+r}(\cdots) + \beta_{4+r}\big),$$

whose u^3-coefficient is $v_r v_{r-1}^2$. On Block 3, replacing v by $-v$ flips the sign of $v_r v_{r-1}^2$ to $-v_r v_{r-1}^2$. Hence $\Delta y^{(1),[1]}(u) + \Delta y^{(1),[3]}(u)$ has degree at most 2 in u, and by (2) every coordinate of $s^{(2),[0]}(u)$ has degree ≤ 2 in u.

Inspecting $[u^2]s_r^{(2)}$ for $r \in \{0,1,2,3\}$: the quadratic term of $y_{4+r}^{(1)} + y_{12+r}^{(1)}$ is affine-linear in the block-sum constants $a^{[1]} + a^{[3]}$ and $a^{[0]}, a^{[2]}$ (with $y_0^{(1)}, y_1^{(1)}$ fixed), while $[u^1]s_r^{(2)}$ is quadratic in those constants. On $t = 16$ ($k = 4$), we have $2k - 5 = 3$ free constant parameters; solving the linear equation $[u^2]s_0^{(2)} = 0$ and the quadratic equation $[u^1]s_0^{(2)} = 0$ (after fixing 1 parameter generically) makes $s_0^{(2)}$ constant while $\deg_u(s_1^{(2)}) = 1$.

In Round 2, $y_0^{(2)}$ is constant, $\deg_u(y_1^{(2)}) = d$, and Blocks 1, 3 have $\deg_u(s_{4+r}^{(2)}) = 3$ with anti-symmetric leading terms $[u^3](s_{4+r}^{(2)} + s_{12+r}^{(2)}) = 0$, whereas Blocks $b \notin \{1,3\}$ have $\deg_u \le 2$. Hence $\deg_u(L_i^{(2)}) = d$ for all $i \ge 2$ and $\deg_u(y_i^{(2)}) = 2d+2$ for $i \notin \mathrm{Blk}\,1 \cup \mathrm{Blk}\,3$ (while the u^{2d+3} terms of Blocks 1 and 3 cancel in their sum!). When $d = 3$, the u^9 terms of $y^{(2),[1]}, y^{(2),[3]}$ cancel in the sum while $2d + 2 = 8$, leaving $\deg_u(s_0^{(3)}) = \deg_u(s_1^{(3)}) = 9$ (from $u^3 \cdot (u^3)^2$ vs cross terms at $d = 3$ where $\deg(y_1^{(2)}) = 3$ ties $\deg(s_{4+r}^{(2)}) = 3$; for $d \ge 5$, $\deg(y_1^{(2)}) = d > 3$ dominates, giving $\deg_u(s_0^{(3)}) = \deg_u(s_1^{(3)}) = 2d+2$: specifically $9, 12, 16$ for $d = 3, 5, 7$).

From Round 3 onward, each round $r \ge 3$ introduces one radical extension $w_r^d = s_0^{(r)}$ of degree d. Tracking the weighted degree δ in the iterated radical tower $R_m = \mathbb{F}_p[u, w_3, \dots, w_{m+2}]/(w_j^d - s_0^{(j)})$ with $\delta(w_j) = \delta(s_0^{(j)})/d$: when $\delta(s_0^{(r)}) = \delta(s_1^{(r)}) = \delta_r$, we have $\delta(y_0^{(r)}) = \delta_r/d$, $\delta(y_1^{(r)}) = d\delta_r$, and $\delta(L_i^{(r)}) = d\delta_r$, so every Horst output $y_i^{(r)} = s_i^{(r)} G_i(L_i^{(r)})$ ($i \ge 2$) has weighted degree $\delta_r + 2d\delta_r = (2d+1)\delta_r$. Taking the field norm down $R - 3$ tower levels of degree d multiplies the final weighted degree by d^{R-3}, giving exact univariate resultant degree $D^{(7)}(R,d) = D_3(d)\cdot(d(2d+1))^{R-3}$. □

3.3. **Wide-Width Degree Collapses: Theorems 8A, 8B ($t \ge 20$) and 8C ($t \ge 24$).** For $t = 4k \ge 20$, the number of free constant parameters in Theorem 3.2 jumps from 3 to $2k - 5 \ge 5$, unlocking two remarkable algebraic collapses that eliminate the d-dependence of Round 2 completely and further drop the degree of $s_1^{(3)}$.

Theorem 3.3 (Rank-1 Quadratic Collapse and Affine-Linear $[u^8]s_1^{(3)}$ Annihilation — Thms 8A & 8B)**.** *Let $t = 4k \geq 20$ and let $d \geq 3$ be odd. Under the Round-1 anti-symmetric block-pair parameterization of Theorem 3.2:*

(1) ***Constant-$(s_0^{(2)}, s_1^{(2)})$ Collapse (Thm 8A):*** *Enforcing $[u^2]s_0^{(2)} = [u^2]s_1^{(2)} = 0$ is a rank-2 affine-linear system leaving $n_\tau = 2k - 7 \geq 3$ free parameters $(\tau_0, \ldots, \tau_{n_\tau-2}, \tau_{\text{last}})$, where $\tau_0, \ldots, \tau_{n_\tau-2}$ are constant-block coordinates $(s_{16}^{(1)}, s_{17}^{(1)}, \ldots)$ and $\tau_{\text{last}} = s_4^{(1)} + s_{12}^{(1)}$. Fixing $\tau_2, \ldots, \tau_{\text{last}}$, the remaining two vanishing conditions*

$$Q_0(\tau_0, \tau_1) := [u^1]s_0^{(2)} = 0, \qquad Q_1(\tau_0, \tau_1) := [u^1]s_1^{(2)} = 0$$

are bivariate quadratics whose homogeneous degree-2 parts are proportional with exact ratio 4:5 inherited from rows 0 and 1 of M_4^{HL}:

$$5\, Q_1^{\text{hom}}(\tau_0, \tau_1) \;=\; 4\, Q_0^{\text{hom}}(\tau_0, \tau_1). \tag{10}$$

Consequently, $5Q_1 - 4Q_0$ is strictly affine-linear, the Sylvester resultant $\text{Res}_{\tau_1}(Q_0, Q_1)$ collapses from generic degree 4 to exact degree 2, and solving two quadratics forces $\deg_u(s_0^{(2)}) = \deg_u(s_1^{(2)}) = 0$. Both $y_0^{(2)}$ and $y_1^{(2)}$ are constant, giving d-independent degree $\deg_u(s_0^{(3)}) = \deg_u(s_1^{(3)}) = 8$ and $D^{(8A)}(R, d) = 8(d(2d+1))^{R-3}$.

(2) ***Free Affine-Linear $[u^8]s_1^{(3)}$ Annihilation (Thm 8B):*** *For every $r \in \{0, 1, 2, 3\}$, the coefficient $[u^8]s_r^{(3)}$ is completely independent of $\tau_0, \ldots, \tau_{n_\tau-2}$ and strictly affine-linear (with nonzero slope) in $\tau_{\text{last}} = s_4^{(1)} + s_{12}^{(1)}$. Solving the single linear equation $[u^8]s_1^{(3)}(\tau_{\text{last}}) = 0$ uniquely determines τ_{last} while preserving (10), forcing*

$$\deg_u(s_0^{(2)}) = \deg_u(s_1^{(2)}) = 0, \qquad \deg_u(s_0^{(3)}) = 8, \qquad \deg_u(s_1^{(3)}) = 7.$$

By the asymmetric-base graded tower recurrence $D(m; \alpha, \beta, d) = d(2d\beta + \alpha)(d(2d+1))^{m-1}$ ($m = R - 3$), the exact R-round CICO-1 resultant degree for every $R \geq 4$ is

$$D^{(8B)}(R, d) \;=\; (14d + 8) \cdot d \cdot \big(d(2d+1)\big)^{R-4}. \tag{11}$$

Proof. In Round 1, the free constant-block entries $\tau_0 = s_{16}^{(1)}, \tau_1 = s_{17}^{(1)}, \ldots$ enter $[u^1]s^{(2),[0]}$ through the initial Round-0 linear relation $y^{(0)} = M^{-1}(s^{(1)} - \text{rc}^{(0)})$, which shifts the Round-1 affine offsets $a^{[1]}, a^{[3]}$ linearly in τ_0, τ_1 inside the 2-dimensional image $M_4^{\text{HL}}(\text{span}(e_0, e_2))$ while preserving the anti-symmetric direction v. Evaluating the contribution of $(a^{[1]}, a^{[3]})$ to $[u^1](y^{(1),[1]} + y^{(1),[3]})$ yields quadratic forms supported strictly on coordinates $r \in \{1, 2, 3\}$; applying M_4^{HL} from (3) shows that the $r = 0$ and $r = 1$ rows of M_4^{HL} are $(5, 7, 1, 3)$ and $(4, 6, 1, 1)$. After eliminating $[u^2]s_0^{(2)} = [u^2]s_1^{(2)} = 0$, the remaining degree-2 homogeneous part in (τ_0, τ_1) enters solely through a single rank-1 combination whose projection into coordinates 0 and 1 has exact ratio $M_{4,1,0}^{\text{HL}} : M_{4,0,0}^{\text{HL}} = 4 : 5$, proving (10).

For a general pair of bivariate quadratics $Q_m = c_m q(\tau_0, \tau_1) + \ell_m(\tau_0, \tau_1)$ ($c_m \neq 0$) with linear eliminant $L = c_0\ell_1 - c_1\ell_0 = \lambda\tau_0 + \mu\tau_1 + \nu$, eliminating $\tau_1 = -(\lambda\tau_0 + \nu)/\mu$ (when $\mu \neq 0$) substitutes a linear function into Q_0, giving exact closed-form resultant

$$\text{Res}_{\tau_1}(Q_0, Q_1) = \frac{a_0}{c_0^2}\mu^2 Q_0\Big(\tau_0, -\frac{\lambda\tau_0 + \nu}{\mu}\Big)$$

of exact degree 2 whenever L is non-constant and non-isotropic ($q(\mu, -\lambda) \neq 0$). On Griffin-20 and Griffin-24, this non-degeneracy holds generically (yielding a root over $\mathbb{F}_p$ with probability $\approx 1/2$ per random seed of the remaining free constants).

Once $s_0^{(2)}$ and $s_1^{(2)}$ are constant, $y_0^{(2)}$ and $y_1^{(2)}$ are constant. In Round 2, $s^{(2),[1]}(u)$ and $s^{(2),[3]}(u)$ have degree 3 with opposite leading terms $\pm wu^3$, while every other block has degree ≤ 2. Hence $\deg_u(y_{4+r}^{(2)}) = \deg_u(y_{12+r}^{(2)}) = 9$ for $r \in \{1, 2, 3\}$ with exact cancellation $[u^9](y_{4+r}^{(2)} + y_{12+r}^{(2)}) = 0$,

so $\deg_u(s_r^{(3)}) \le 8$ for $r \in \{0,1,2,3\}$. Moreover, $[u^8](y_{4+r}^{(2)} + y_{12+r}^{(2)})$ depends solely on the degree-3 and degree-2 coefficients of $s^{(2),[1]}, s^{(2),[3]}$. Because the degree-2 coefficients of $s^{(2),[1]} + s^{(2),[3]}$ depend strictly and affine-linearly on $\tau_{\text{last}} = s_4^{(1)} + s_{12}^{(1)}$ and are independent of $\tau_0, \dots, \tau_{n_\tau - 2}$ (which only shift constant terms of $s^{(2)}$ after $[u^2]s_0^{(2)} = [u^2]s_1^{(2)} = 0$), solving the single linear equation $[u^8]s_1^{(3)}(\tau_{\text{last}}) = 0$ fixes τ_{last} first and yields $\deg_u(s_0^{(3)}) = 8$, $\deg_u(s_1^{(3)}) = 7$.

Finally, in Round 3, $w_3^d = s_0^{(3)}$ has weighted degree $\delta(w_3) = 8/d$, while $y_1^{(3)} = (s_1^{(3)})^d$ has weighted degree $7d$, so $\delta(L_i^{(3)}) = 7d$ for all $i \ge 2$. Although Blocks 1 and 3 have $\deg_u(s_{4+r}^{(3)}) = 9$, their opposite $[u^9]$ leading terms cancel identically in $y_{4+r}^{(3)} + y_{12+r}^{(3)}$, leaving $\deg_u(s_0^{(4)}) = \deg_u(s_1^{(4)}) = 14d + 8$. Taking the norm down $R - 3$ tower levels from Round 4 (where $\delta_4 = 14d + 8$ and each subsequent round multiplies δ by $2d + 1$) gives

$$D^{(8B)}(R, d) = d^{R-3} \cdot (14d + 8)(2d + 1)^{R-4} = (14d + 8)d(d(2d + 1))^{R-4}. \qquad \square$$

When $t = 4k \ge 24$, we have $n_{\text{rem}} = 2k - 8 \ge 4$ remaining constant-block parameters $(\tau_0, \dots, \tau_{2k-9})$ even *after* fixing τ_{last}. This enables one further exact d-independent collapse:

Theorem 3.4 (Extra-τ Quadratic Annihilation of $[u^7]s_1^{(3)}$ on Griffin-$4k$, $k \ge 6$ — Thm 8C)**.** *Let $t = 4k \ge 24$ and let $d \ge 3$ be odd. Under Theorem 3.3 (with τ_{last} fixed by $[u^8]s_1^{(3)} = 0$), let (τ_0, τ_1, τ_2) be three free constant-block parameters (fixing any remaining τ_j generically). Then:*

1. *Because $[u^7]y_j^{(2)}$ ($j \in \text{Blk}\,1 \cup \text{Blk}\,3$) depends solely on the degree-$3, 2, 1$ coefficients of $s^{(2)}(u)$—and is completely independent of the constant terms $[u^0]L_j^{(2)} = (j-1)(s_0^{(2)})^{1/d} + (s_1^{(2)})^d + [u^0]s_{j-1}^{(2)}$—both $R_0 := [u^7]s_0^{(3)}$ and $R_1 := [u^7]s_1^{(3)}$ are exact d-independent multivariate quadratic polynomials in (τ_0, τ_1, τ_2).*
2. *The four homogeneous quadratic parts of (Q_0, Q_1, R_0, R_1) span an exact 2-dimensional subspace, with $5Q_1 - 4Q_0$ purely affine-linear. Eliminating one variable linearly via $5Q_1 - 4Q_0 = 0$ reduces the system $Q_0 = Q_1 = R_1 = 0$ to a degree-4 bivariate quadratic Sylvester resultant solvable in < 0.1 s over any prime field, forcing*
$$\deg_u(s_0^{(2)}) = \deg_u(s_1^{(2)}) = 0, \qquad \deg_u(s_0^{(3)}) = 8, \qquad \deg_u(s_1^{(3)}) = 6.$$
3. *The exact R-round* CICO*-1 univariate resultant degree for every $R \ge 4$ drops to*
$$D^{(8C)}(R, d) \;=\; (12d + 8) \cdot d \cdot \big(d(2d + 1)\big)^{R-4} \;=\; \frac{12d + 8}{14d + 8}\, D^{(8B)}(R, d). \tag{12}$$

Proof. Write $s_j^{(2)}(u) = c_{j,3}u^3 + c_{j,2}u^2 + c_{j,1}u + c_{j,0}$ for $j \in \text{Blk}\,1 \cup \text{Blk}\,3$. In $y_j^{(2)} = s_j^{(2)}((L_j^{(2)})^2 + \alpha_j L_j^{(2)} + \beta_j)$ where $L_j^{(2)} = s_{j-1}^{(2)} + \gamma_j$, the constant shift $\gamma_j = (j-1)y_0^{(2)} + y_1^{(2)}$ multiplies at most u^6 (via $c_{j,3}u^3 \cdot 2c_{j-1,3}u^3\gamma_j$)! Consequently, $[u^7]y_j^{(2)}$ receives contributions exclusively from terms of the form $c_{j,3}c_{j-1,2}^2$, $2c_{j,3}c_{j-1,3}c_{j-1,1}$, $2c_{j,2}c_{j-1,3}c_{j-1,2}$, and $c_{j,1}c_{j-1,3}^2$. Here $c_{j,3}$ is constant in τ, $c_{j,2}$ is affine-linear in τ (actually constant in τ_0, τ_1, τ_2 after fixing τ_{last}!), and $c_{j,1}$ is quadratic in τ (matching Q_0, Q_1). Hence R_0, R_1 are exact d-independent quadratics in (τ_0, τ_1, τ_2) whose homogeneous quadratic parts lie in the 2-dimensional span of Q_0^{hom} and the quadratic terms of $c_{j,1}$. Eliminating τ_2 via the affine-linear relation $5Q_1 - 4Q_0 = 0$ leaves two bivariate quadratics $\widetilde{Q}_0(\tau_0, \tau_1) = 0$ and $\widetilde{R}_1(\tau_0, \tau_1) = 0$ whose Sylvester resultant has degree ≤ 4. At any root, $\deg_u(s_0^{(3)}) = 8$ and $\deg_u(s_1^{(3)}) = 6$, so the asymmetric-base formula $D(R - 3; 8, 6, d)$ gives $(12d + 8)d(d(2d + 1))^{R-4}$. $\square$

Remark 3.5 (Two-Pair Griffin-32)**.** On width $t = 32$ ($k = 8$), placing two anti-symmetric pairs on $(\text{Blk}\,1, \text{Blk}\,3)$ and $(\text{Blk}\,5, \text{Blk}\,7)$ yields full affine-linear rank 2 on $([u^8]s_0^{(3)}, [u^8]s_1^{(3)})$ and reaches $\deg_u(s_0^{(3)}, s_1^{(3)}) = (7, 6)$ via the same $5Q_1 - 4Q_0$ linear elimination. However, because $\deg_u(s_j^{(3)}) = 8$

persists on lanes $j \geq 2$, the Round-4 degree remains $(12d+8)d$, matching Theorem 3.4 exactly (verified at 132 for $d = 3$).

3.4. **Exact Saturation and Structural Obstructions.** A natural question is whether further constant-block parameters can push the 1-pair trail past $(8, 6)$, whether a 2-round *pure-linear* bypass exists on wide widths, or whether final-round backward-peel can eliminate the last round for CICO-1. We answer all three negatively with sharp algebraic obstructions.

Theorem 3.6 (Exact Saturation and Structural Obstructions — Thm 8D)**.** *For Griffin-4k with linear layer $M = (I_k + J_k) \otimes M_4^{\mathrm{HL}}$ and odd $d \geq 3$:*

(1) ***Quadratic-Precomputation Saturation:*** *For all $k \in \{5, 6, 7, 8\}$, $[u^8]s_0^{(3)}$ and $[u^8]s_1^{(3)}$ have coordinate degree vector $(0, \ldots, 0, 1)$ in $(\tau_0, \ldots, \tau_{2k-8}, \tau_{\mathrm{last}})$, so a 1-pair trail can annihilate at most one of $[u^8]s_0^{(3)}, [u^8]s_1^{(3)}$. Moreover, $[u^6]s_1^{(3)}$ is the exact threshold where $[u^0]L_j^{(2)} = (j-1)(s_0^{(2)})^{1/d} + (s_1^{(2)})^d + [u^0]s_{j-1}^{(2)}$ enters; since $s_0^{(2)}[0]$ and $s_1^{(2)}[0]$ are cubic in τ (verified exact degree vector $(3, \ldots, 3)$ on Griffin-24/28), $[u^6]s_1^{(3)}$ is algebraic of degree d in τ unless $(s_0^{(2)}, s_1^{(2)})$ are additionally constrained by two cubic equations. Thus Theorem 3.4 ($\deg = (8, 6)$) is the exact d-independent quadratic-precomputation ceiling for 1-pair trails.*

(2) ***2-Round Pure-Linear Bypass Obstruction and Exact V_- Diophantine Identity:*** *For M_4^{HL} and every pair of 2-element non-adjacent subsets $E, Z \subset \{0, 1, 2, 3\}$, the coordinate projection $P_{E,Z}$ along $M_4^{\mathrm{HL}}(\mathrm{span}\{e_i : i \in E\})$ onto $\mathrm{span}\{e_j : j \notin Z\}$ is well-defined over $\mathbb{Q}$. On the M_4^{HL} anti-symmetric 2-plane*

$$V_- = \mathrm{span}\big\{(1, 0, -1, 0)^\top, (0, 1, 0, -1)^\top\big\},$$

the restrictions of the two non-adjacent alternating-zero projections are

$$P_{02}\big|_{V_-} = \begin{pmatrix} 0 & 4/3 \\ 0 & 1 \end{pmatrix}, \qquad P_{13}\big|_{V_-} = \begin{pmatrix} 1 & 0 \\ 5/4 & 0 \end{pmatrix},$$

yielding the exact Diophantine determinant identity

$$\det\Big((I_4 + aP_{02} + bP_{13})\big|_{V_-}\Big) = 1 + a + b - \frac{2}{3}ab = 0 \iff (2a-3)(2b-3) = 15. \tag{13}$$

Over $\mathbb{N}$, (13) has exactly four solutions: $(a, b) \in \{(4, 3), (3, 4), (2, 9), (9, 2)\}$ (corresponding to $k = a + b + 1 \in \{8, 12\}$, i.e., Griffin-32 and Griffin-48). However, imposing Block 0's constant-$(s_0^{(1)}, s_1^{(1)})$ condition eliminates every candidate: exhaustive verification over all 89,760 (Block-0 $\times$ multiset) configurations ($3{,}960 + 10{,}296 + 24{,}024 + 51{,}480$) proves zero 2-round pure-linear bypass solutions for all $k \in \{5, 6, 7, 8\}$.

(3) ***Final-Round Backward-Peel Obstruction for*** CICO***-1:*** *On every width $t = 4k \in \{16, 20, 24, 28, 32\}$, row 0 of M and row 0 of M^{-1} both have t/t nonzero entries. By the row-space coordinate-fixing criterion, y_0 is constant on a nonempty fiber of $Cy = b$ if and only if $e_0^\top \in \mathrm{row}(C)$. Since $e_0^\top \notin \mathrm{row}(M_{[:c_{\mathrm{out}}]})$ for all $c_{\mathrm{out}} < t$, neither* CICO*-1 ($z_0 = 0$) nor* CICO*-2 ($z_0 = z_1 = 0$) can fix $y_0^{(R-1)}$ through M^{-1}; imposing $y_0^{(R-1)} = \lambda$ on a 2-variable slice explodes the resultant degree to $\Theta(D_{R-1}^2) \gg d(2d+1)D_{R-1}$.*

4. Tower Arithmetic, Division-Free Lucas Resultants, and Solvers

Evaluating the univariate resultant $P(u) \in \mathbb{F}_p[u]$ of degree $D = D^{(8B)}(R, d)$ at $N > D$ evaluation points $u \in \mathbb{F}_p$, interpolating $P(u)$ via NTT, and finding its $\mathbb{F}_p$-roots uses strictly $O(D)$ words of memory—never forming a Gröbner basis or multivariate Macaulay matrix. To push this streaming architecture to deg = 29.17M (KoalaBear 8/10) and deg = 8.18M (64-bit Goldilocks 6/9), we introduce four exact algebraic and algorithmic tools.

4.1. **Universal Division-Free Odd-d Lucas Resultant.** At the final round $R-1$, the output constraint $s_0^{(R)} = 0$ is a quadratic polynomial $A_2w^2 + A_1w + A_0 = 0$ in the final radical $w = y_0^{(R-1)}$ satisfying $w^d = S := s_0^{(R-1)}$, with coefficients A_2, A_1, A_0, S in the $(R-4)$-level tower ring R_{tow} of dimension d^{R-4} (e.g., dim = 49 for 6-round Goldilocks, dim = 81 for 8-round KoalaBear). Computing $\text{Res}_w(A_2w^2 + A_1w + A_0, w^d - S)$ via standard Euclidean PRS or monic division requires an inversion A_2^{-1} inside R_{tow}—which is both slow and can fail at zero divisors. The following identity eliminates *every* division in R_{tow}.

Proposition 4.1 (Universal Division-Free Odd-d Lucas Resultant)**.** *Let R be any commutative ring, let $d = 2m+1 \geq 3$ be odd, and define the division-free homogeneous Lucas sequence $(L_k)_{k\geq 0}$ in R by*

$$L_0 = 2, \qquad L_1 = -A_1, \qquad L_{k+2} = -A_1L_{k+1} - A_2A_0L_k. \tag{14}$$

Define the division-free polynomial

$$\mathcal{R}_d(A_2, A_1, A_0, S) \;:=\; A_2^dS^2 - L_dS + A_0^d \;\in\; R. \tag{15}$$

Then for every $w \in R$ satisfying $A_2w^2 + A_1w + A_0 = 0$ and $w^d = S$, we have the universal identity

$$A_2^d \cdot \mathcal{R}_d(A_2, A_1, A_0, S) = 0,$$

and over any integral domain (or whenever A_2 is a non-zero-divisor) the identity $\mathcal{R}_d(A_2, A_1, A_0, S) = \text{Res}_w(A_2w^2 + A_1w + A_0, w^d - S)$ holds with zero inversions. Explicitly, for $d = 3$ and $d = 7$:

$$\begin{aligned} L_3 &= -A_1^3 + 3A_2A_1A_0, \\ L_7 &= -A_1^7 + 7A_2A_0A_1^5 - 14A_2^2A_0^2A_1^3 + 7A_2^3A_0^3A_1. \end{aligned}$$

Proof. Work in the universal ring $\mathbb{Z}[A_2, A_1, A_0, S]$ (or in any splitting ring with roots α, β of $A_2w^2 + A_1w + A_0 = A_2(w-\alpha)(w-\beta)$). Setting $r_1 = A_2\alpha$, $r_2 = A_2\beta$, we have $r_1 + r_2 = -A_1$ and $r_1r_2 = A_2A_0$, so by Newton's identities $L_k = r_1^k + r_2^k$ satisfies (14). For odd d,

$$\begin{aligned} \text{Res}_w(A_2w^2 + A_1w + A_0, w^d - S) &= A_2^d(\alpha^d - S)(\beta^d - S) \\ &= A_2^dS^2 - (r_1^d + r_2^d)S + (r_1r_2)^d/A_2^d \cdot A_2^d \\ &= A_2^dS^2 - L_dS + A_0^d. \end{aligned}$$

Equivalently, a direct induction in any commutative ring R (formalized verbatim in the module `MonicLucasResultant.lean`) defines the companion element $t := -A_1 - A_2w$; then $A_2w + t = -A_1$ and $(A_2w) \cdot t = A_2A_0$, so by induction $L_k = (A_2w)^k + t^k$ for every $k \geq 0$ in *any* commutative ring R. Substituting $k = d$ and $w^d = S$ yields

$$A_2^d \cdot (A_2^dS^2 - L_dS + A_0^d) = \big((A_2w)^d\big)^2 - \big((A_2w)^d + t^d\big)(A_2w)^d + (A_2w)^dt^d = 0. \qquad \square$$

Notice how astonishingly clean the companion-element proof $t = -A_1 - A_2w$ is: it requires no splitting field, no characteristic assumptions, and no divisions! For $d = 7$ inside the 49-element Goldilocks tower Bu_7, computing $\mathcal{R}_7$ requires just 12 ring multiplications and 0 inversions, followed by a 7×7 skew-circulant norm $\text{Norm}_{\text{Bu}_7/\text{Au}_7}$ and PRS norm $\text{Norm}_{\text{Au}_7/\mathbb{F}_p}$.

4.2. **Cyclic Wrap-Around Euclidean Half-GCD.** For root-finding over 64-bit Goldilocks at degree $D = 8{,}180{,}550$, computing $\gcd(f, x^p - x \bmod f)$ via Half-GCD repeatedly applies 2×2 polynomial transfer matrices $M_k = \left(\begin{smallmatrix} m_{00} & m_{01} \\ m_{10} & m_{11} \end{smallmatrix}\right)$ to (a, b). Naively, $\deg(m_{ij}) + \deg(a)$ can reach $\frac{3}{2}\deg(a) > 2^{23}$, forcing an NTT of length 2^{24} (128 MB per buffer). The following lemma halves the top-level NTT size to $N = 2^{23}$.

Proposition 4.2 (Cyclic Wrap-Around Euclidean Half-GCD)**.** *Let $a, b \in \mathbb{F}[x]$ with $\deg(a) > \deg(b) \geq 0$, and let $(r_k, r_{k+1}) = M_k(a, b)^\top$ be any intermediate Euclidean remainder pair with transfer matrix M_k. Then:*

(1) *Every entry of* M_k *satisfies* $\max_{i,j} \deg(m_{ij}) = \deg(a) - \deg(r_k)$ *(sharp, including the final division).*
(2) *For every* $N > \deg(a)$*, the quotient projection* $\pi : \mathbb{F}[x] \to \mathbb{F}[x]/(x^N - 1)$ *recovers* (r_k, r_{k+1}) exactly *from cyclic convolutions of length* N*:*

$$\pi(M_k) \begin{pmatrix} \pi(a) \\ \pi(b) \end{pmatrix} = \begin{pmatrix} \pi(r_k) \\ \pi(r_{k+1}) \end{pmatrix},$$

even when $\deg(m_{ij}) + \deg(a) \geq N$*, because all wrapped high-degree terms of* $m_{i0}a + m_{i1}b$ *cancel identically modulo* $x^N - 1$*.*

Proof. Part (1) follows by induction on the Euclidean step $v_{i+1} = v_{i-1} - q_i v_i$ using $\deg(q_i) = \deg(r_{i-1}) - \deg(r_i)$. Part (2) is immediate from the ring homomorphism property of π together with $\deg(r_k), \deg(r_{k+1}) \leq \deg(a) < N$, so $\pi(r_k)$ and $\pi(r_{k+1})$ are represented verbatim by their unique representatives of degree $< N$. □

4.3. **2-Coset CRT Interpolation, 127-Coset Root-Finding, and 16-Point Goldilocks** Au_7 **NTT.** For KoalaBear ($p_{\mathrm{KB}} = 2^{31} - 2^{24} + 1$), the maximal power-of-2 subgroup $H \subset \mathbb{F}_p^\times$ has order $N = 2^{24} = 16{,}777{,}216$, whereas 8/10-round Griffin-20 has exact degree $D = 29{,}172{,}150 \in (2^{24}, 2 \cdot 2^{24})$. Because $p_{\mathrm{KB}} - 1 = 127 \cdot 2^{24}$ with primitive root $g = 3$, we exploit the full multiplicative coset decomposition of $\mathbb{F}_p^\times$:

Proposition 4.3 (Complete-Field Multi-Coset Partition and Forward-NTT Folding)**.** *Let* $p-1 = qN$ *with* $N = 2^e$*, let* g *be a primitive root modulo* p*, and let* $H = \langle \omega \rangle$ *where* $\omega = g^q$*. Then:*

(1) $\mathbb{F}_p = \{0\} \sqcup \bigsqcup_{c=0}^{q-1} g^c H$ *is an exact disjoint partition of* $\mathbb{F}_p$*.*
(2) *Any polynomial* $P(u) = \sum_{\ell=0}^{r-1} u^{\ell N} L_\ell(u)$ *with* $\deg(L_\ell) < N$ *(*$r \leq q$*) is uniquely reconstructed from its* r *inverse NTTs on cosets* $g^c H$ *(*$0 \leq c < r$*) via an* $r \times r$ *Vandermonde inverse. For* $r = 2$ *and shift* $\zeta = g$ *(*$\zeta^N \neq 1$*),*

$$L_1 = \frac{P_0 - P_1}{1 - \zeta^N}, \qquad L_0 = P_0 - L_1,$$

where $P_0(u) = P(u) \bmod (u^N - 1)$ *and* $P_1(u) = P(u) \bmod (u^N - \zeta^N)$*.*
(3) *Evaluating* $P(u)$ *on every element of coset* $g^c H$ *reduces to a single length-*N *forward NTT:*

$$P(g^c \omega^j) \;=\; \sum_{k=0}^{N-1} \left(\sum_{\ell=0}^{r-1} g^{c\ell N} L_{\ell,k} \right) g^{ck} \omega^{jk} \qquad (0 \leq j < N). \tag{16}$$

For 8/10-round KoalaBear Griffin-20 ($q = 127, r = 2, N = 2^{24}$), Proposition 4.3 eliminates Cantor-Zassenhaus entirely! Evaluating $P(u)$ across all 127 cosets scans the *entire* field $\mathbb{F}_{p_{\mathrm{KB}}}$ using length-2^{24} forward NTTs in < 260 MB RAM—finding a root on Coset 3 in just 0.78 s!

For 6/9-round 64-bit Goldilocks Griffin-20 ($d = 7$, $\deg = 8{,}180{,}550 < 2^{23}$), the 49-element tower $\mathrm{Bu}_7 = \mathrm{Au}_7[w_4]/(w_4^7 - s_4)$ is accelerated by 10.4×–14.4× (57.0 μs → 5.5 μs per multiplication, 3.97 μs per squaring) via:

(1) Zero-division Goldilocks shift-reduction ($2^{64} \equiv 2^{32} - 1$, $2^{128} \equiv -2^{32} \pmod{p_{\mathrm{GL}}}$) and lazy `u128`-accumulated Au_7 multiplication (13 reductions per product);
(2) Exact 16-point Radix-2 NTT over Au_7 using $\omega_{16} = 2^{12}$ (since $(2^{12})^8 = 2^{96} \equiv -1 \pmod{p_{\mathrm{GL}}}$), every twiddle multiplication is a pure bit-shift!), together with NTT-domain Round-5 accumulation across all 18 Horst lanes so only one inverse 16-point NTT is performed per output;
(3) A 63-squaring, 5-multiplication addition chain for the Frobenius power $x^{2^{64}-2^{32}+1} \bmod f$ paired with Proposition 4.2.

5. Experimental Records and Full-Round Complexity

All experiments were executed on a single commodity Linux workstation using zero-dependency Rust solvers (`griffin_r6_record.rs`, `griffin_r8_record.rs`, `griffin_r9_goldilocks6.rs`) on unmodified Griffin parameters ($M = (I_k + J_k) \otimes M_4^{\mathrm{HL}}$ and deterministic SHAKE256 round constants). Every single witness was independently re-verified by our AST-audited zero-float pure-Python verifier (`verify_all_witnesses.py`).

5.1. Complete Table of 17 Verified Explicit Griffin CICO-1 Witnesses.

Table 1 summarizes all 17 verified CICO-1 witnesses across widths $t \in \{8, 12, 16, 20, 24\}$ and rounds $R \in \{3, \ldots, 8\}$.

Table 1. All 17 zero-float verified explicit CICO-1 witnesses on unmodified Griffin (M_4^{HL} + SHAKE256 RC). Witness IDs match `verify_all_witnesses.py`.

#	Field	d	t	R/R_{full}	Exact $\deg D$	Wall Time	Peak RSS	Method
35	KoalaBear-31	3	8	3/12	147	0.23 s	$<$ 20 MB	Thm 3.1
36	KoalaBear-31	3	12	3/11	63	0.05 s	$<$ 20 MB	Thm 3.1
37	Goldilocks-64	7	12	3/9	315	20.5 s	$<$ 20 MB	Thm 3.1
38	KoalaBear-31	3	16	4/10	189	1.72 s	$<$ 20 MB	Thm 3.2
39	Goldilocks-64	7	16	4/9	1,680	7.94 s	$<$ 20 MB	Thm 3.2
40	KoalaBear-31	3	16	5/10	3,969	0.26 s	$<$ 20 MB	Thm 3.2
41	KoalaBear-31	3	16	6/10	83,349	32.0 s	18.8 MB	Thm 3.2
42	Goldilocks-64	7	20	5/9	88,200	51.1 s	28.4 MB	Thm 3.3 (8A)
43	Goldilocks-64	7	16	5/9	176,400	105.8 s	54.9 MB	Thm 3.2
44	KoalaBear-31	3	20	7/10	1,555,848	111.1 s	190.3 MB	Thm 3.3 (8A)
45	KoalaBear-31	3	16	7/10	1,750,329	101.6 s	212.4 MB	Thm 3.2
46	Goldilocks-64	7	20	5/9	77,910	71.8 s	$<$ 30 MB	Thm 3.3 (8B)
47	Goldilocks-64	7	24	5/9	77,910	59.0 s	$<$ 30 MB	Thm 3.3 (8B)
48	KoalaBear-31	3	20	7/10	1,389,150	42.2 s	$<$ 260 MB	Thm 3.3 (8B)
49	**KoalaBear-31**	**3**	**20**	**8/10**	**29,172,150**	**1,770** s	$<$ **260** MB	**Thm 8B + 127-Coset**
50	**Goldilocks-64**	**7**	**20**	**6/9**	**8,180,550**	**625** s	$<$ **750** MB	**Thm 8B + NTT16+Luc**
51	**Goldilocks-64**	**7**	**24**	**5/9**	**67,620**	**48.7** s	$<$ **30** MB	**Thm 8C** ($\deg 8, 6$)

5.2. Explicit Coordinates of the Three Headline Witnesses.

For immediate independent verification, the exact input vectors $x \in \mathbb{F}_p^t$ (each satisfying $x_0 = 0$ and $\mathrm{Griffin}(x)_0 = 0$) for Witnesses #49, #50, and #51 are:

- **Witness #49** (8/10-**Round KoalaBear Griffin-20,** $p = 2^{31} - 2^{24} + 1$, $d = 3$**):**
  `[0, 1692663486, 2038169559, 489816146, 1224193715, 1068459412, 1588269325, 334633623, 251857072, 983447917, 1473603284, 587302340, 1943766559, 470066552, 1654350236, 980618496, 869412491, 865907626, 690083924, 1736548203]`.
- **Witness #50** (6/9-**Round** 64-**Bit Goldilocks Griffin-20,** $p = 2^{64} - 2^{32} + 1$, $d = 7$**):**
  `[0, 9544845677420826514, 4908424410869006141, 13583653515931775808, 15412669642124258951, 7542392473571875507, 10551272207676525627, 7586718911680878160, 8285537952563466913, 9656060946554260632, 9407228186897433574, 6110534783394243294, 6933715841289296599, 17242190957719880963, 12099743722346834479, 202015764876970301, 15545461022627413690, 4646632264231563513, 13583114821809245221, 2951527962162827253]`.
- **Witness #51** (5/9-**Round** 64-**Bit Goldilocks Griffin-24, Theorem 3.4):**
  `[0, 17077514409126549288, 6156367387496073381, 11496246742896110519, 9332065569275443824, 6267882106847922659, 8858991866844066653, 13015215557067076192, 13363753473783717820, 6742760091064083972, 301076414752553524, 2597111485828525459, 9354948830061661560, 5688923852365650422, 9504231786535034676, 15752192121124870647, 4642143065299753028, 7142847194991369016, 1364351446517738757, 15573011395020904809, 4311068167303354421, 11767696176574905735, 16470575942837351760, 10935480740014834562]`.

5.3. **Full-Round Streaming Complexity and Literature Comparison.** Table 2 reports exact univariate tower-resultant degrees and conservative streaming operation counts $D \cdot (20 \cdot (d^{R-4})^{1.585} + \log_2^2 D)$ (at $O(D)$ words of memory) for full-round Griffin instances targeting 128-bit security.

TABLE 2. Full-round Griffin streaming tower-NTT CICO-1 complexity (exact univariate degrees, tower dimension d^{R-4}, conservative field operations vs. 2^{128} security claim).

Instance	d	R	Theorem	Exact $\deg D$	$\log_2 D$	$\dim R_{\text{tow}}$	$\log_2$(Ops)
Griffin-12	3	11	Thm 3.1	2,382,840,139,743	41.12	2,187	63.02
Griffin-16	3	10	Thm 3.2	16,209,796,869	33.92	729	53.31
Griffin-20	3	10	Thm 3.3	12,864,918,150	33.58	729	52.98
Griffin-24	3	10	Thm 3.4	11,321,127,972	33.40	729	**52.80**
Griffin-16	5	9	Thm 3.2	332,167,687,500	38.27	3,125	61.00
Griffin-20	5	9	Thm 3.3	196,280,906,250	37.51	3,125	60.24
Griffin-24	5	9	Thm 3.4	171,116,687,500	37.32	3,125	**60.04**
Griffin-16	7	9	Thm 3.2	21,441,530,250,000	44.29	16,807	70.86
Griffin-20	7	9	Thm 3.3	9,470,009,193,750	43.11	16,807	69.68
Griffin-24	7	9	Thm 3.4	8,219,253,262,500	42.90	16,807	**69.47**

Positioning vs. Prior Work. Compared with *FreeLunch* [3] (which broke 7/10-round Griffin in several hours via Gröbner bases), our 7/10-round KoalaBear witnesses solve in 42.2 s (< 260 MB RAM) and scale directly to 8/10-round KoalaBear (1,770 s, < 260 MB) and 6/9-round 64-bit Goldilocks (625 s, < 750 MB). Compared with *CheapLunch* [1] and *Resultants Meet Resultant* [2]—which target multi-output CICO-2 / SFTM via multivariate resultants—our work operates strictly in the harder single-output CICO-1 setting ($x_0 = 0 \to y_0 = 0$), achieves strict $O(D)$ streaming memory, and uncovers the wide-width Theorems 3.3–3.4 degree collapses together with the exact Theorem 3.6 saturation boundaries.

6. Formal Verification in Lean 4 and Reproducibility

To guarantee mathematical soundness beyond floating-point or heuristic approximations, we formalized the five core algebraic pillars supporting this paper (and our companion Poseidon2 study) in a standalone, zero-dependency Lean 4.36.0 library `CryptoAlgebra` (1,598 lines across 7 files, zero Mathlib, zero `sorry`, zero `native_decide`). Table 3 maps our results to their machine-checked Lean theorems and exact `#print axioms` output.

TABLE 3. Zero-dependency Lean 4.36.0 verification summary (`CryptoAlgebra`, 16/16 Lake jobs, 0 errors, 0 warnings).

Paper Result	Lean 4 Theorem	Axioms Used
Lemma 2.1, Eq. (4)	`blockKron_left_inv`, `M4_HL_adj_mul`	`[propext, Choice, Quot]`
Proposition 4.1	`lucasResultant_mul_A2_pow_eq_zero`	`[propext]`
Proposition 4.2	`euclidRowMaxDeg_exact`, `cyclic_wraparound_recovery`	`[propext, Choice, Quot]`
Proposition 4.3	`expToCoset_cosetToExp`, `multi_coset_dft_folding`	`[propext, Choice, Quot]`
Goldilocks $\omega_{16} = 2^{12}$	`goldilocks_omega16_order`	*none* (empty!)
Eqs. (11), (12)	`griffinDeg_closed_form`, `witness49/50/51_exact_deg`	`[propext, Quot]`
Theorem 3.6(2)	`griffin_Vminus_nat_solutions`	`[propext, Choice, Quot]`

All 17 Griffin witnesses (together with 34 companion Poseidon2 witnesses) are independently verified in 33 ms by running `python3 -I -B ../code/verify_all_witnesses.py`, which performs an AST self-audit confirming zero `float` literals and zero true-division (`ast.Div`) operators.

References

1. Antoine Bak, Augustin Bariant, Aurélien Boeuf, Pierre Briaud, Morten Øygarden, and Atharva Phanse, *The algebraic CheapLunch: Extending FreeLunch attacks on arithmetization-oriented primitives beyond CICO-1*, 2025, IACR Cryptology ePrint Archive, Paper 2025/2040.
2. Antoine Bak, Augustin Bariant, Maël Hostettler, and Vincent Neiger, *Resultants meet resultant: Improving CICO-1 and CICO-2 attacks on ZK-friendly permutations*, 2026, IACR Cryptology ePrint Archive, Paper 2026/1281.
3. Augustin Bariant, Aurélien Boeuf, Axel Lemoine, Irati Manterola Ayala, Morten Øygarden, Léo Perrin, and Håvard Raddum, *The algebraic FreeLunch: Efficient Gröbner basis attacks against arithmetization-oriented primitives*, Advances in Cryptology – CRYPTO 2024, LNCS, Springer, 2024, ePrint 2024/347.
4. Lorenzo Grassi, Yonglin Hao, Christian Rechberger, Markus Schofnegger, Roman Walch, and Qingju Wang, *Horst meets fluid-SPN: Griffin for zero-knowledge applications*, Advances in Cryptology – CRYPTO 2023, LNCS, vol. 14083, Springer, 2023, pp. 573–606.

Email address: srehwaldt@gmail.com