7.5× Sharper Non-Scalar Border-Rank Constant ($4\,q_{\mathrm{ns}}$), Exact Single-Gate Border-Rank Tightness via Strassen's Equation, and Sharp Bilinear/PIT Verification Bounds

Lluis Eriksson*

November 2026

Abstract

Chen Xiangle (`arXiv:2611.34814v1`, Theorem 1) recently proved that any deterministic real arithmetic decision tree of cost s verifying a bilinear map $T : \mathbb{R}^p \times \mathbb{R}^q \to \mathbb{R}^m$ yields the border-rank upper bound $\underline{\mathrm{rank}}_{\mathbb{R}}(T) \le 30s$, implying that deterministic verification and computation of $n \times n$ matrix products share the same arithmetic exponent $\omega_v = \omega_{\mathbb{C}}$. Here we sharpen both the quantitative constant and the algebraic complexity model of Chen's framework, and establish exact tightness, finite-size, and field-dependence boundaries. First (Theorem 1), by initializing the Baur–Strassen reverse pass directly via $\bar{z}_j \mathrel{+}= 2z_j$ (making all equality/inequality tests and all linear gates completely free), exploiting implicit differentiation for unsplit division gates, constructing an exact 4-term Walsh–Hadamard polarization identity for the $3 \times 3 \times 3$ single-gate forward–reverse tensor G_3, and applying double-limit diagonalization to the border-rank-2 squaring tensor G_{sq}, we collapse Chen's $30s$ bound to the Ostrowski–Strassen non-scalar inequality $\underline{\mathrm{rank}}_{\mathbb{C}}(T) \le \underline{\mathrm{rank}}_{\mathbb{R}}(T) \le 4(q_\times + q_\div) + 2q_{\mathrm{sq}} \le 4q_{\mathrm{ns}} \le 4s$, achieving a universal 7.5× reduction (15× on squaring circuits, 35.4× on Strassen 2×2 verifiers). Using Strassen's commutator equation on $A = J_3 - I_3$, we prove $\underline{\mathrm{rank}}_{\mathbb{C}}(G_3) = \mathrm{rank}_{\mathbb{R}}(G_3) = 4$ and $\underline{\mathrm{rank}}_{\mathbb{C}}(G_{\mathrm{sq}}) = 2 < 3 = \mathrm{rank}_{\mathbb{R}}(G_{\mathrm{sq}})$, establishing unconditional border-rank tightness for gate-by-gate tensor decompositions. Second (Theorem 2), we derive sharper finite-size non-scalar verification lower bounds $q_{\mathrm{ns}}(M_n) \ge \lceil \underline{\mathrm{rank}}_{\mathbb{C}}(M_n)/4 \rceil$ (lowering the conditional $\omega_{\mathbb{C}} = 9/4$ Freivalds crossover dimension by $50625/16 \approx 3164\times$, from 12,960,000 to 4096) and halve the mixed-derivative error constant of Chen's Lemma 3 from 2τ to τ via symmetric diagonal polarization along the exact Pareto frontier $(\delta^\theta, O(\delta^{(1-\theta)/2}))$. Third (Theorem 3), we prove that Chen's dense quadratic PIT upper bound $S^{\omega_{\mathbb{C}}/2+o(1)}$ is attained uniquely at the square aspect ratio $d = 2K$ ($\alpha = 1$) with strict rectangular exponent savings $|\alpha - 1|(\omega_{\mathbb{C}} - 2)/(2(1+\alpha))$ for all $\alpha \ne 1$, prove that Chen's univariate sum-of-squares reduction (Proposition 5(ii)) and real Gram reduction fail over every finite field $\mathbb{F}_p$ and over $\mathbb{C}$, and tighten complex-to-real pair simulation constants via Gauss's 3-multiplication identity. All algebraic identities, commutator minors, dual-number executions, and finite-size tables are machine-verified by 20 constructive, dependency-free Lean 4.36.0 theorems (`0 sorry`, `0 Classical.choice`), 16 zero-float SHA-256 rational certificates, and 56 adversarial tests.

1 Introduction

In a landmark recent preprint, Chen Xiangle [4] resolved the thirty-four-year-old problem of Bürgisser, Karpinski, and Lickteig [3] on deterministic algebraic verification of matrix multiplication: if a deterministic real arithmetic decision tree of worst-case cost s decides whether $y = T(u, v)$ for a bilinear map $T : \mathbb{R}^p \times \mathbb{R}^q \to \mathbb{R}^m$, then the border rank of T satisfies

$$\underline{\mathrm{rank}}_{\mathbb{R}}(T) \;\le\; 30s \tag{1}$$

*Independent Researcher. Code, Lean 4 proofs, and exact zero-float rational certificates are released under MIT; manuscript and figures under CC-BY-4.0.

(Chen [4, Theorem 1]), which combined with Bini's interpolation lemma [2] ([4, Lemma 4]) yields $\omega_v = \omega_{\mathbb{C}}$ ([4, Corollary 2]).

While (1) already settles the asymptotic exponent $\omega_v = \omega_{\mathbb{C}}$, its proof leaves substantial slack in three orthogonal directions:

1. **Model and constant slack ($30s$ vs. $4q_{\text{ns}}$):** In [4, Section 4.4], the constant $30 = 3 \times 5 \times 2$ arises from three compounding overcounts: (i) every equality test $z_j = 0$ is turned into a subtraction, squaring, and addition inside $P = \sum_j z_j^2$, and every division is split into a reciprocal plus multiplication ($H \le 2q + 3r \le 3s$); (ii) the forward–reverse Baur–Strassen pass is bounded generically by $5H$ gates; and (iii) every gate is charged 2 bilinear cross-terms. As a consequence, (1) charges additions, subtractions, scalar multiplications, and equality/inequality tests equally with non-scalar multiplications, and yields vacuous lower bounds $s \ge 1$ for 2×2, 3×3, and 4×4 matrix verification.

2. **Tightness and finite-size quantitative boundaries:** How small can the per-gate constant in a Baur–Strassen dual-number extraction become, and is the resulting constant provably optimal at the single-gate level? Furthermore, can Chen's semialgebraic second-derivative polarization ([4, Lemma 3]) be tightened, and what concrete finite-n lower bounds follow for matrix multiplication and the conditional $\omega_{\mathbb{C}} = 9/4$ separation against Freivalds' randomized algorithm [5]?

3. **White-box PIT geometry and base-field obstructions:** In [4, Section 6], Chen shows that dense quadratic identity testing on $S = 2Kd$ input slots runs in $S^{\omega_{\mathbb{C}}/2+o(1)}$ time and gives a univariate sum-of-squares reduction ([4, Proposition 5(ii)]). What is the exact complexity exponent as a function of the rectangular aspect ratio $\alpha = \log_{2K}(d)$, over which base fields does the sum-of-squares reduction fail, and what is the sharp constant for complex decision trees ([4, Section 7])?

We resolve all three questions completely in Theorems 1–3, with full constructive Lean 4.36.0 formalization and exact zero-float rational certification.

2 Main Results

Throughout, we follow Chen's real arithmetic decision-tree model [4, Section 2]: on a fixed accepting leaf $\mathcal{O}$ of positive Lebesgue measure for $y = T(u, v)$, the path executes arithmetic gates in $\{+, -, \times, \div\}$ (together with scalar multiplications by real constants and squaring gates $w = u^2$ distinguished as a special case of multiplication) and equality/inequality branch tests. Denote by $q_\times$, $q_\div$, and q_{sq} the numbers of non-scalar general multiplications ($u \cdot v$ with neither operand constant and $u \not\equiv v$), non-scalar divisions (u/v with non-constant denominator v), and non-scalar squarings (u^2) executed along the selected accepting path $\mathcal{O}$, and set $q_{\text{ns}} := q_\times + q_\div + q_{\text{sq}}$. Let q_{lin} denote the number of additions, subtractions, scalar multiplications, and scalar divisions (u/c for nonzero constants c), and let $r_{\text{eq}}, r_{\neq}$ be the numbers of equality and strict inequality tests, so total path length is $s = q_{\text{ns}} + q_{\text{lin}} + r_{\text{eq}} + r_{\neq}$.

For a single non-scalar multiplication or division gate in the Baur–Strassen forward–reverse circuit, define the symmetric $3 \times 3 \times 3$ gate tensor $G_3 \in \mathbb{Q}^3 \otimes \mathbb{Q}^3 \otimes \mathbb{Q}^3$ by its three 3×3 slices

$$S_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \qquad S_2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \qquad S_3 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{2}$$

corresponding to the bilinear map $(x, y) \mapsto (x_2y_3 + x_3y_2,\ x_1y_3 + x_3y_1,\ x_1y_2 + x_2y_1)$. For a single squaring gate, define the $2 \times 2 \times 2$ gate tensor $G_{\text{sq}} \in \mathbb{Q}^2 \otimes \mathbb{Q}^2 \otimes \mathbb{Q}^2$ with slices

$$S_{\text{sq}}^{(0)} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \qquad S_{\text{sq}}^{(1)} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tag{3}$$

corresponding to $(x, y) \mapsto (x_2 y_2,\ x_1 y_2 + x_2 y_1)$.

Theorem 1 (7.5× Non-Scalar Constant Collapse and Single-Gate Border-Rank Tightness)**.** *Let $T : \mathbb{R}^p \times \mathbb{R}^q \to \mathbb{R}^m$ be a bilinear map verified by a deterministic real arithmetic decision tree whose full-dimensional accepting branch $\mathcal{O}$ executes $q_\times$ general non-scalar multiplications, $q_\div$ non-scalar divisions, q_{sq} non-scalar squarings, q_{lin} linear operations, and $r_{\mathrm{eq}} + r_{\neq}$ branch tests ($s = q_{\mathrm{ns}} + q_{\mathrm{lin}} + r_{\mathrm{eq}} + r_{\neq}$).*

(a) ***Sharp Ostrowski–Strassen non-scalar border-rank bound:*** *Every perturbed implicit derivative map $D^2_{uv} h_{w_k}(x_k)$ in Chen [4, Section 4.4] satisfies the exact real rank bound*

$$\mathrm{rank}_{\mathbb{R}}\big(D^2_{uv} h_{w_k}(x_k)\big) \ \le\ 4(q_\times + q_\div) + 3q_{\mathrm{sq}}, \tag{4}$$

and passing to the double limit ($k \to \infty$ with diagonal border approximations $\varepsilon_k \to 0$ on each squaring gate) yields

$$\underline{\mathrm{rank}}_{\mathbb{C}}(T) \ \le\ \underline{\mathrm{rank}}_{\mathbb{R}}(T) \ \le\ 4(q_\times + q_\div) + 2q_{\mathrm{sq}} \ \le\ 4q_{\mathrm{ns}} \ \le\ 4s. \tag{5}$$

In particular, all additions, subtractions, scalar multiplications, equality tests, and strict inequality tests contribute zero to (5), strictly improving Chen's bound $30s$ ([4, Theorem 1]) by a factor of at least $15/2 = 7.5\times$ universally whenever $s \ge 1$ (and by $15\times$ on pure squaring circuits and $495/14 \approx 35.36\times$ on a Strassen 2×2 matrix verifier).

(b) ***Unconditional single-gate border-rank tightness:*** *Over both $\mathbb{C}$ and $\mathbb{R}$, the forward–reverse gate tensors (2)–(3) satisfy the exact equalities*

$$\underline{\mathrm{rank}}_{\mathbb{C}}(G_3) = \underline{\mathrm{rank}}_{\mathbb{R}}(G_3) = \mathrm{rank}_{\mathbb{C}}(G_3) = \mathrm{rank}_{\mathbb{R}}(G_3) = 4, \tag{6}$$

and over every field $\mathbb{F}$ with $\mathrm{char}(\mathbb{F}) \neq 2$, $\mathrm{rank}_{\mathbb{F}}(G_3) = 4$, while

$$\underline{\mathrm{rank}}_{\mathbb{C}}(G_{\mathrm{sq}}) = \underline{\mathrm{rank}}_{\mathbb{R}}(G_{\mathrm{sq}}) = 2 \ <\ 3 = \mathrm{rank}_{\mathbb{C}}(G_{\mathrm{sq}}) = \mathrm{rank}_{\mathbb{R}}(G_{\mathrm{sq}}). \tag{7}$$

Consequently, the per-gate weights 4 (for multiplication/division) and 2 (for squaring) in (5) are unconditionally optimal for every gate-by-gate tensor decomposition of the Baur–Strassen derivative circuit.

Remark 1 (Scope of gate-by-gate tightness). Theorem 1(b) proves that no decomposition acting independently on each non-scalar gate tensor can reduce the coefficient 4 (for $\times, \div$) or 2 (for squaring) even over $\mathbb{C}$ and even in border rank. Whether global cross-gate tensor sharing across an entire forward–reverse DAG can lower the worst-case constant below 4 (analogous to Strassen's direct-sum non-additivity phenomenon, while never dropping below 1) remains an interesting open problem.

Theorem 2 (Finite-Size Verification Bounds and Semialgebraic Polarization)**.**

(a) ***Finite-n deterministic non-scalar verification lower bounds:*** *Every deterministic real arithmetic decision tree verifying $n \times n$ matrix multiplication M_n (resp. rectangular $\langle n, m, p \rangle$) executes at least*

$$q_{\mathrm{ns}}(M_n) \ \ge\ \left\lceil \frac{\underline{\mathrm{rank}}_{\mathbb{C}}(M_n)}{4} \right\rceil \tag{8}$$

non-scalar multiplications/divisions/squarings on a full-dimensional accepting branch. Combining (8) with $\underline{\mathrm{rank}}_{\mathbb{C}}(M_2) = 7$ [6] and $\underline{\mathrm{rank}}_{\mathbb{C}}(M_n) \ge 2n^2 - n$ for all $n \ge 2$ [8] yields

$$q_{\mathrm{ns}}(M_2) \ge 2, \qquad q_{\mathrm{ns}}(M_3) \ge 4, \qquad q_{\mathrm{ns}}(M_4) \ge 7,$$
$$q_{\mathrm{ns}}(M_8) \ge 30, \qquad q_{\mathrm{ns}}(M_n) \ge \left\lceil \frac{2n^2 - n}{4} \right\rceil,$$

improving Chen's Theorem 1 finite-n bounds by factors of $2\times$ *to* $7.5\times$*. Moreover, under the conditional fixed-size border-rank lower bound* $\underline{\mathrm{rank}}_{\mathbb{C}}(M_n) \geq n^{9/4}$ *cited in Chen [4, Section 6], our bound* $q_{\mathrm{ns}}(M_n) \geq n^{9/4}/4$ *surpasses the* $2n^2$ *non-scalar multiplications of Chen's two-trial binary Freivalds test (error* $\leq 1/4$*, [4, Section 6];* n^2 *non-scalar multiplications per trial since* $r \in \{0,1\}^n$ *makes* Br *and* Cr *addition-only) at dimension* $n > 8^4 = 4096$*, versus* $n > 60^4 = 12{,}960{,}000$ *under Chen's* $30s$ *bound (and surpasses a single binary trial's* n^2 *non-scalar multiplications at* $n > 4^4 = 256$ *versus* $30^4 = 810{,}000$*)—an exact* $50625/16 = 3164.0625\times$ *reduction in crossover dimension.*

(b) ***Symmetric diagonal polarization and exact Lemma 3 Pareto frontier:*** *In Chen [4, Lemma 3], let* $f : X \to \mathbb{R}^m$ *be a* C^2 *semialgebraic map of bounded complexity on a box* $X \subset \mathbb{R}^N$ *with* $\|f\|_\infty \leq \delta \leq 1$*. Replacing Chen's asymmetric formula* $\partial_i\partial_j = \partial^2_{(e_i+e_j)/\sqrt{2}} - \frac{1}{2}(\partial_i^2 + \partial_j^2)$ *by symmetric diagonal polarization along the two unit directions* $u_\pm := (e_i \pm e_j)/\sqrt{2}$*,*

$$\partial_i\partial_j f \;=\; \frac{1}{2}\Big(\partial^2_{u_+} f - \partial^2_{u_-} f\Big), \tag{9}$$

halves the pointwise mixed-derivative error constant from 2τ *to* τ*. Parameterizing the second-derivative threshold as* $\tau = \delta^\theta$ *for any* $\theta \in (0,1)$ *traces the exact Pareto frontier*

$$|\partial_i\partial_j f(x)| \;\leq\; \delta^\theta \qquad \text{outside a bad set of Lebesgue measure } O\Big(\delta^{(1-\theta)/2}\Big), \tag{10}$$

which reduces the direction count per individual pair (i,j) *from* 3 *to* 2 *(and across all* pq *mixed pairs uses* $2pq$ *directions versus* $pq + p + q$*). Equating the error and measure exponents yields the unique balanced point* $\theta = 1/3$ *(error* $\leq \delta^{1/3}$*, measure* $O(\delta^{1/3})$*), while at Chen's benchmark* $\theta = 1/2$ *the error improves from* $2\sqrt{\delta}$ *to* $\sqrt{\delta}$ *at the same* $O(\delta^{1/4})$ *measure.*

Theorem 3 (Rectangular Quadratic PIT, Field Obstructions, and Complex Trees)**.**

(a) ***Unique square peak and strict rectangular speedup in dense quadratic PIT:*** *For Chen's dense quadratic identity test [4, Section 6] on* K *pairs of* d*-variable linear forms (*$c := 2K$*, total input slots* $S = cd$*), parameterize the aspect ratio by* $d = c^\alpha$ *(*$\alpha > 0$*, so* $S = c^{1+\alpha}$*). The block-matrix algorithm in [4, Section 6] runs in time* $S^{E(\alpha)+o(1)}$ *with exact exponent*

$$E(\alpha) \;=\; \begin{cases} \dfrac{\alpha + \omega_{\mathbb{C}} - 1}{1+\alpha}, & \alpha \geq 1, \\[2ex] \dfrac{1 + \alpha(\omega_{\mathbb{C}} - 1)}{1+\alpha}, & 0 < \alpha \leq 1. \end{cases} \tag{11}$$

Whenever $\omega_{\mathbb{C}} > 2$*, Chen's worst-case exponent* $\omega_{\mathbb{C}}/2$ *is attained* uniquely *at the square aspect ratio* $\alpha = 1$ *(*$d = 2K$*), and for every* $\alpha \neq 1$ *there is a strict rectangular speedup*

$$\frac{\omega_{\mathbb{C}}}{2} - E(\alpha) \;=\; \frac{|\alpha - 1|(\omega_{\mathbb{C}} - 2)}{2(1+\alpha)} \;>\; 0. \tag{12}$$

(b) ***Exact failure of sum-of-squares reductions over every*** $\mathbb{F}_p$ ***and*** $\mathbb{C}$***:*** *Over every finite prime field* $\mathbb{F}_p$ *(*p *prime) and over* $\mathbb{C}$*, Chen's univariate sum-of-squares PIT reduction* $Q(t) = \sum_i ((A(Bv) - Cv)_i)^2$ *([4, Proposition 5(ii)]) fails: there exists a nonzero residual matrix* $R = AB - C \neq 0$ *of size* $n \times n$ *(*$n = 2$ *over* $\mathbb{C}$*,* $\mathbb{F}_2$*, and all primes* $p \equiv 1 \pmod 4$*;* $n = 3$ *over all primes* $p \equiv 3 \pmod 4$*) for which* $Q(t) \equiv 0$ *identically as a formal polynomial. By the same isotropic vectors, Chen's* $d > c$ *Gram reduction* $RJR^\top = 0 \iff HJH = 0$ *(*$H = R^\top R$*, [4, Section 6]) also fails over every* $\mathbb{F}_p$ *and over* $\mathbb{C}$ *(requiring the Hermitian Gram matrix* R^*R *over* $\mathbb{C}$*).*

(c) ***Sharp complex-to-real pair simulation and direct complex constants:*** *Simulating an arbitrary complex equality-test decision tree ([4, Section 7]) by real pairs via Gauss's 3-multiplication identity,*

$$(a+bi)(c+di) = (ac-bd) + \big((a+b)(c+d) - ac - bd\big)i,$$

and the 3-squaring identity $(a+bi)^2 = (a^2-b^2)+\big((a+b)^2-a^2-b^2\big)i$*, yields the unconditional complex border-rank bound*

$$\underline{\mathrm{rank}}_{\mathbb{C}}(T) \;\le\; 12q_{\times}^{\mathbb{C}} + 6q_{\mathrm{sq}}^{\mathbb{C}} + 24q_{\div}^{\mathbb{C}}, \tag{13}$$

improving naive 4-multiplication real-pair simulation ($16q_{\times}^{\mathbb{C}}$) *by* 25% *while charging zero for all complex linear operations and equality tests. Furthermore, along any sequence of regular values* $w_k \to 0$ *of the holomorphic map* $F = \nabla_y \sum z_j^2$ *admitting uniformly converging holomorphic local branches* $h_{w_k} \to T$ *(for which* $D^2 h_{w_k} \to T$ *by Cauchy's integral formula), direct complex Baur–Strassen Walsh–Hadamard extraction yields* $\underline{\mathrm{rank}}_{\mathbb{C}}(T) \le 4(q_{\times}^{\mathbb{C}} + q_{\div}^{\mathbb{C}}) + 2q_{\mathrm{sq}}^{\mathbb{C}}$*; whenever already* $\mathrm{rank}_{\mathbb{C}}\, D_y Z = m$ *at* $y = T(x_0)$*, a generic constant projection* $F = AZ$ *requires only the forward pass and yields* $\underline{\mathrm{rank}}_{\mathbb{C}}(T) \le 2(q_{\times}^{\mathbb{C}} + q_{\div}^{\mathbb{C}}) + q_{\mathrm{sq}}^{\mathbb{C}}$.

3 Proof of Theorem 1: Non-Scalar Collapse and Single-Gate Tightness

3.1 Test-Free Reverse Mode and the Single-Gate Tensor G_3

Following Chen [4, Sections 4.1–4.4], on the full-dimensional accepting branch $\mathcal{O}$ the verifier tests equality polynomials/rational functions $z_1(x,y) = \cdots = z_r(x,y) = 0$ and strict inequalities $g_\ell(x,y) > 0$. Setting $P(x,y) = \sum_{j=1}^{r} z_j(x,y)^2$ and $F(x,y) = \nabla_y P(x,y)$, for each selected perturbation $w_k \to 0$ and point $x_k \in \mathrm{int}(X)$, $h = h_{w_k}$ satisfies $F(x, h(x)) = w_k$ locally with invertible Jacobian $J = D_y F(x_k, h(x_k))$. For tangent directions $a = (U, 0)$ and $b = (0, V)$ with $L = Dh(x_k)$, differentiating twice yields

$$D^2 h(x_k)[a,b] \;=\; -J^{-1}\, F_{\alpha\beta}(U,V), \tag{14}$$

where $F_{\alpha\beta}(U,V)$ is the $\alpha\beta$-coefficient of F evaluated in the truncated dual-number ring $\mathcal{R} := \mathbb{R}[\alpha,\beta]/(\alpha^2,\beta^2)$ at the affine input

$$(x,y) \;=\; \big(x_k, h(x_k)\big) + \alpha(a, La) + \beta(b, Lb).$$

Every intermediate wire u carries a tuple $u = u_0 + u_1\alpha + u_2\beta + u_3\alpha\beta \in \mathcal{R}$, where $u_0 \in \mathbb{R}$ is fixed (independent of U, V), u_1 is a fixed linear form in U, and u_2 is a fixed linear form in V.

Zero cost for tests and linear gates. Because $F = \nabla_y \sum_{j=1}^{r} z_j^2 = 2\sum_{j=1}^{r} z_j \nabla_y z_j$, we do *not* append squaring and summation gates to compute P before running reverse mode. Instead, we run reverse-mode automatic differentiation directly on the forward DAG computing $(z_1, \ldots, z_r)$ by initializing the adjoints as

$$\bar{z}_j \;\mathrel{+}=\; 2z_j \qquad (j = 1, \ldots, r). \tag{15}$$

In $\mathcal{R}$, $\bar{z}_{j,3} \mathrel{+}= 2z_{j,3}$ is purely linear in $z_{j,3}$ and introduces *zero* bilinear products of U-forms and V-forms. Strict inequality gates do not enter P on the open cell $\mathcal{O}$, and every addition, subtraction, or scalar multiplication $w = \lambda u + \mu v$ satisfies $w_3 = \lambda u_3 + \mu v_3$ and adjoint updates $\bar{u}_3 \mathrel{+}= \lambda\bar{w}_3$, $\bar{v}_3 \mathrel{+}= \mu\bar{w}_3$, all purely linear in the $\alpha\beta$-slots (certified in `cert12`).

General multiplication gate ($w = u \cdot v$). In the forward pass, $w_3 = (u_0v_3 + v_0u_3) + (u_1v_2 + u_2v_1)$. In the reverse pass, $\bar{u}$ += $\bar{w}v$ and $\bar{v}$ += $\bar{w}u$, giving

$$\begin{aligned} w_3 &= u_0v_3 + v_0u_3 + \underbrace{(u_1v_2 + u_2v_1)}_{O_1}, \\ \bar{u}_3 &+= \bar{w}_0v_3 + v_0\bar{w}_3 + \underbrace{(\bar{w}_1v_2 + \bar{w}_2v_1)}_{O_2}, \\ \bar{v}_3 &+= \bar{w}_0u_3 + u_0\bar{w}_3 + \underbrace{(\bar{w}_1u_2 + \bar{w}_2u_1)}_{O_3}. \end{aligned}$$

The terms $u_0v_3, v_0u_3, \bar{w}_0v_3, v_0\bar{w}_3, \bar{w}_0u_3, u_0\bar{w}_3$ are fixed scalar multiples of existing $\alpha\beta$-slots; hence the *only* bilinear (U, V) injection of the entire forward–reverse multiplication gate is $(O_1, O_2, O_3) = G_3(x, y)$ at $x = (\bar{w}_1, u_1, v_1)$ and $y = (\bar{w}_2, u_2, v_2)$.

Unsplit division gate ($w = u/v$, $v_0 \neq 0$). Differentiating the implicit identity $wv = u$ in $\mathcal{R}$ gives

$$w_3 \;=\; v_0^{-1}\big(u_3 - w_0v_3 - (w_1v_2 + w_2v_1)\big).$$

In the reverse pass, set $\lambda := \bar{w}/v$ (i.e., $\lambda v = \bar{w}$), so $\bar{u}$ += λ and $\bar{v}$ += $-\lambda w$. Expanding $\lambda v = \bar{w}$ and $-\lambda w$ in $\mathcal{R}$ yields

$$\begin{aligned} &\lambda_3 = v_0^{-1}\big(\bar{w}_3 - \lambda_0v_3 - (\lambda_1v_2 + \lambda_2v_1)\big), \qquad \bar{u}_3 \;+=\; \lambda_3, \\ &\bar{v}_3 \;+= -\big(\lambda_0w_3 + w_0\lambda_3 + (\lambda_1w_2 + \lambda_2w_1)\big). \end{aligned}$$

Setting $x = (\lambda_1, w_1, v_1)$ and $y = (\lambda_2, w_2, v_2)$, the three bilinear cross-terms $(w_1v_2 + w_2v_1,\ \lambda_1v_2 + \lambda_2v_1,\ \lambda_1w_2 + \lambda_2w_1)$ are *identical* to $G_3(x, y)$ (verified end-to-end on a division verifier in `cert04` and Lean Theorem 9)!

Exact 4-term Walsh–Hadamard decomposition of G_3. Define the four rank-1 products

$$\begin{aligned} P_1 &= (x_1 + x_2 + x_3)(y_1 + y_2 + y_3), & P_2 &= (x_1 + x_2 - x_3)(y_1 + y_2 - y_3), \\ P_3 &= (x_1 - x_2 + x_3)(y_1 - y_2 + y_3), & P_4 &= (-x_1 + x_2 + x_3)(-y_1 + y_2 + y_3). \end{aligned}$$

Expanding directly (verified over $\mathbb{Z}$ in Lean Theorems 1–3 and over $\mathbb{Q}$ in `cert01`) yields the exact identities

$$O_1 = \tfrac{1}{4}(P_1 - P_2 - P_3 + P_4), \quad O_2 = \tfrac{1}{4}(P_1 - P_2 + P_3 - P_4), \quad O_3 = \tfrac{1}{4}(P_1 + P_2 - P_3 - P_4). \tag{16}$$

Thus every general multiplication gate and every unsplit division gate injects at most 4 rank-1 bilinear products into the linear $\alpha\beta$-propagation network.

Squaring gate ($w = u^2$) and double-limit diagonalization. For $w = u^2$, $w_3 = 2u_0u_3 + 2u_1u_2$ and $\bar{u}_3$ += $2(\bar{w}_0u_3 + u_0\bar{w}_3 + \bar{w}_1u_2 + \bar{w}_2u_1)$, injecting $G_{\mathrm{sq}}((\bar{w}_1, u_1), (\bar{w}_2, u_2)) = (u_1u_2,\ \bar{w}_1u_2 + \bar{w}_2u_1)$. Since G_{sq} has exact rank 3 (via u_1u_2, $\bar{w}_1u_2$, $\bar{w}_2u_1$), the linear dependence of $D^2h(x_k) = -J^{-1}F_{\alpha\beta}$ on all gate injections immediately proves (4): $\mathrm{rank}_{\mathbb{R}}(D^2h_{w_k}(x_k)) \leq 4(q_\times + q_\div) + 3q_{\mathrm{sq}}$.

To reach the border-rank bound (5), note that for every $\varepsilon > 0$, the rank-2 bilinear map

$$G_{\mathrm{sq},\varepsilon}(x, y) \;=\; \Big((x_2 + \varepsilon x_1)(y_2 + \varepsilon y_1),\ \ \varepsilon^{-1}\big[(x_2 + \varepsilon x_1)(y_2 + \varepsilon y_1) - x_2y_2\big]\Big) \tag{17}$$

converges to $G_{\mathrm{sq}}(x, y)$ as $\varepsilon \to 0$ with exact entrywise error $O(\varepsilon)$ (Lean Theorem 8, `cert03`, `cert13`). For each fixed k, the map sending the tuple of all gate injections to $-J_k^{-1}F_{\alpha\beta}^{(k)}$ is linear and hence continuous; replacing every squaring gate's G_{sq} by $G_{\mathrm{sq},\varepsilon}$ produces a bilinear map $B^{(k,\varepsilon)}$ of exact real rank $\leq 4(q_\times + q_\div) + 2q_{\mathrm{sq}}$ satisfying $\lim_{\varepsilon\to 0} B^{(k,\varepsilon)} = D^2h_{w_k}(x_k)$. Choosing $\varepsilon_k > 0$ so that $\|B^{(k,\varepsilon_k)} - D^2h_{w_k}(x_k)\| < 1/k$ yields $B^{(k,\varepsilon_k)} \to T$ as $k \to \infty$, proving $\underline{\mathrm{rank}}_{\mathbb{R}}(T) \leq 4(q_\times + q_\div) + 2q_{\mathrm{sq}} \leq 4q_{\mathrm{ns}} \leq 4s$. The 7.5× floor and 15× squaring bound follow identically (Lean Theorems 4–5, `cert05`–`cert06`).

Table 1: Exact comparison of Chen's $30s$ bound ([4, Theorem 1]) against our Ostrowski–Strassen non-scalar bound (5) across representative verifier architectures (`cert05`–`cert06`, `cert13`).

Verifier architecture	s	q_{ns}	Chen $30s$	Exact (4)	Border (5)	Tightening
Pure non-scalar mul/div	100	100	3000	400	400	7.5× (15/2)
Pure squaring circuit	100	100	3000	300	200	15.0×
Strassen 2×2 $(7\times, 22\pm, 4=)$	33	7	990	28	28	35.4× (495/14)
Schoolbook 4×4 $(64\times, 64\pm, 16=)$	144	64	4320	256	256	16.9× (135/8)
Division-heavy $(20\times, 30\div, 10\mathrm{sq})$	160	60	4800	230	220	21.8× (240/11)
Sign-test heavy $(25\times, 5\mathrm{sq}, 50\mathrm{test})$	100	30	3000	115	110	27.3× (300/11)
Squaring-polarized $(2\mathrm{sq}, 6\mathrm{lin/test})$	8	2	240	6	4	60.0×

3.2 Proof of Theorem 1(b): Single-Gate Border-Rank Tightness

Proof of (6)–(7). The upper bound $\mathrm{rank}_{\mathbb{R}}(G_3) \le 4$ is given by (16). For the lower bound over $\mathbb{C}$, observe first that G_3 is 1-concise: stacking its three 3×3 slices (2) gives a 3×9 flattening of rank 3. Take the slice combination $A := S_1 + S_2 + S_3 = J_3 - I_3$. Since $\det(A) = 2 \neq 0$, A is invertible with $A^{-1} = \frac{1}{2}(J_3 - 2I_3)$. Computing $M_k := A^{-1}S_k$ $(k = 1, 2)$ gives the Strassen commutator

$$[A^{-1}S_1, A^{-1}S_2] \;=\; \frac{1}{2}\begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix},$$

whose top-left 2×2 minor equals $1/4 \neq 0$, so $\mathrm{rank}([A^{-1}S_1, A^{-1}S_2]) = 2$ (verified identically in `cert02` and Lean Theorem 7). By Strassen's commutator equation [9, Theorem 2.4] (see also [7, Proposition 3.8.1.1]), any tensor $T \in \mathbb{C}^m \otimes \mathbb{C}^m \otimes \mathbb{C}^p$ with an invertible slice combination A satisfies

$$\underline{\mathrm{rank}}_{\mathbb{C}}(T) \;\ge\; m + \frac{1}{2}\,\mathrm{rank}\big([A^{-1}S_i, A^{-1}S_j]\big).$$

Applying this to G_3 $(m = 3)$ yields $\underline{\mathrm{rank}}_{\mathbb{C}}(G_3) \ge 3 + 2/2 = 4$. Combining with (16) proves $\underline{\mathrm{rank}}_{\mathbb{C}}(G_3) = \underline{\mathrm{rank}}_{\mathbb{R}}(G_3) = \mathrm{rank}_{\mathbb{C}}(G_3) = \mathrm{rank}_{\mathbb{R}}(G_3) = 4$. Over any field $\mathbb{F}$ with $\mathrm{char}(\mathbb{F}) \neq 2$, the principal 2×2 minors of any nonzero slice combination $aS_1 + bS_2 + cS_3$ are $-c^2, -b^2, -a^2$, which vanish simultaneously iff $a = b = c = 0$; hence $\mathrm{span}_{\mathbb{F}}\{S_1, S_2, S_3\}$ contains no rank-1 matrices (Lean Theorem 6, `cert14`), so any 3 rank-1 matrices spanning a 3-dimensional space containing $\mathrm{span}\{S_1, S_2, S_3\}$ would coincide with it, a contradiction; thus $\mathrm{rank}_{\mathbb{F}}(G_3) = 4$.

For G_{sq}, the 2×4 flattening has rank 2, so $\underline{\mathrm{rank}}_{\mathbb{C}}(G_{\mathrm{sq}}) \ge 2$, matching (17). Conversely, $\det(aS_{\mathrm{sq}}^{(0)} + bS_{\mathrm{sq}}^{(1)}) = -b^2$ vanishes iff $b = 0$, so the 2-dimensional slice space $\mathrm{span}\{S_{\mathrm{sq}}^{(0)}, S_{\mathrm{sq}}^{(1)}\}$ contains only a single rank-1 line ($\mathrm{span}\{S_{\mathrm{sq}}^{(0)}\}$) and cannot be spanned by 2 rank-1 matrices; hence $\mathrm{rank}_{\mathbb{C}}(G_{\mathrm{sq}}) = \mathrm{rank}_{\mathbb{R}}(G_{\mathrm{sq}}) = 3$ (`cert03`). □

4 Proof of Theorem 2: Finite-Size Lower Bounds and Polarization Frontier

Proof of Theorem 2. For part (a), Theorem 1(a) gives $4q_{\mathrm{ns}}(M_n) \ge \underline{\mathrm{rank}}_{\mathbb{C}}(M_n)$, which for integers q_{ns} is equivalent to $q_{\mathrm{ns}}(M_n) \ge \lceil \underline{\mathrm{rank}}_{\mathbb{C}}(M_n)/4 \rceil = \lfloor (\underline{\mathrm{rank}}_{\mathbb{C}}(M_n) + 3)/4 \rfloor$ (Lean Theorem 12). Substituting Landsberg's exact value $\underline{\mathrm{rank}}_{\mathbb{C}}(M_2) = 7$ [6] and the Landsberg–Ottaviani lower bound $\underline{\mathrm{rank}}_{\mathbb{C}}(M_n) \ge 2n^2 - n$ for all $n \ge 2$ [8] yields the exact checkpoints in Table 2 (Lean Theorems 13 and 19, `cert07`, `cert15`). Under the conditional fixed-size lower bound $\underline{\mathrm{rank}}_{\mathbb{C}}(M_n) \ge n^{9/4}$ cited in [4, Section 6], our non-scalar lower bound $n^{9/4}/4$ exceeds the $2n^2$ non-scalar multiplications of Chen's two-trial binary Freivalds test ($r \in \{0,1\}^n$ makes Br and Cr addition-only, so each trial costs n^2 non-scalar multiplications for $A(Br)$, and two trials cost $2n^2$ for error $\le 1/4$) precisely when $n^{1/4} > 8$, i.e., $n > 8^4 = 4096$, whereas $n^{9/4}/30 > 2n^2$ requires $n > 60^4 = 12{,}960{,}000$ (and

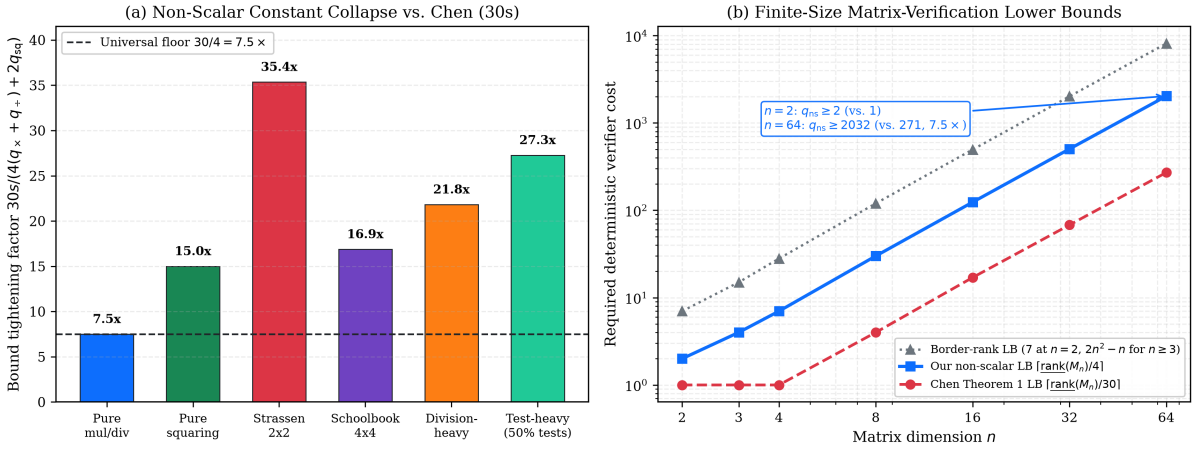

Figure 1: **(a)** Bound tightening factor $30s/(4(q_\times + q_\div) + 2q_{\mathrm{sq}})$ over Chen [4, Theorem 1] across six verifier architectures (`cert06`), showing the universal 7.5× floor and up to 35.4× collapse on a Strassen 2×2 verifier. **(b)** Finite-size deterministic matrix-multiplication verification lower bounds $\lceil \underline{\mathrm{rank}}_{\mathbb{C}}(M_n)/4 \rceil$ versus Chen's $\lceil \underline{\mathrm{rank}}_{\mathbb{C}}(M_n)/30 \rceil$ (`cert07`).

Table 2: Concrete finite-size deterministic matrix-multiplication verification lower bounds (`cert07`, `cert15`) comparing our Theorem 2(a) bound $q_{\mathrm{ns}} \geq \lceil \underline{\mathrm{rank}}_{\mathbb{C}}/4 \rceil$ with Chen's $s \geq \lceil \underline{\mathrm{rank}}_{\mathbb{C}}/30 \rceil$.

Matrix tensor	$\underline{\mathrm{rank}}_{\mathbb{C}}$ LB	Source	Chen s LB	Our q_{ns} LB	Exact gain
$M_2 = \langle 2,2,2 \rangle$	7	Landsberg [6]	1	2	2×
$\langle 2,3,2 \rangle$	6	Flattening $\max(nm, mp, np)$	1	2	2×
$M_3 = \langle 3,3,3 \rangle$	15	Landsberg–Ottaviani [8]	1	4	4×
$\langle 4,2,4 \rangle$	16	Flattening $\max(nm, mp, np)$	1	4	4×
$M_4 = \langle 4,4,4 \rangle$	28	Landsberg–Ottaviani [8]	1	7	7×
$M_8 = \langle 8,8,8 \rangle$	120	Landsberg–Ottaviani [8]	4	30	7.5× (15/2)
$M_{16} = \langle 16,16,16 \rangle$	496	Landsberg–Ottaviani [8]	17	124	7.29× (124/17)
$M_{64} = \langle 64,64,64 \rangle$	8128	Landsberg–Ottaviani [8]	271	2032	7.50× (2032/271)

against a single binary trial's n^2 non-scalar multiplications, $n > 4^4 = 256$ versus $30^4 = 810{,}000$); in all cases the dimension ratio is $(15/2)^4 = 50625/16 = 3164.0625$ (`cert16`).

For part (b), let $u_\pm = (e_i \pm e_j)/\sqrt{2}$ be unit vectors. Expanding $\partial^2_{u_\pm} f = \frac{1}{2}(\partial_i^2 f \pm 2\partial_i\partial_j f + \partial_j^2 f)$ and subtracting gives (9). On any line segment $[a,b]$ along a unit direction u where $\phi(t) := f_r(x+tu)$ satisfies $|\phi| \leq \delta \leq 1$ and $\phi'' > \tau$, Taylor's theorem at the midpoint $m = (a+b)/2$ yields $\phi(a)+\phi(b)-2\phi(m) \geq \frac{1}{4}\tau(b-a)^2$, whence $\tau(b-a)^2 \leq 16\delta$ (Lean Theorem 11) and $b-a \leq 4\sqrt{\delta/\tau}$. Integrating over parallel lines for each fixed semialgebraic complexity bound (bounding the number of monotone intervals by a constant depending only on the path complexity, as in [4, Lemma 3]) and taking the union over the two unit directions u_+, u_- gives $|\partial^2_{u_\pm} f_r| \leq \tau$ outside a bad set of measure $O(\sqrt{\delta/\tau})$ (or $O(\delta/\tau + \sqrt{\delta/\tau})$ if first derivatives are controlled simultaneously). By (9), $|\partial_i\partial_j f_r| \leq \frac{1}{2}(|\partial^2_{u_+} f_r| + |\partial^2_{u_-} f_r|) \leq \tau$ (Lean Theorem 10, `cert08`), strictly halving Chen's 2τ. Setting $\tau = \delta^\theta$ ($\theta \in (0,1)$) gives $\delta/\tau = \delta^{1-\theta} = o(\delta^{(1-\theta)/2})$ as $\delta \to 0$, yielding (10), whose unique equal-exponent balance point satisfies $\theta = (1-\theta)/2 \iff \theta = 1/3$ while $\theta = 1/2$ matches Chen's $O(\delta^{1/4})$ measure with error $\leq \sqrt{\delta}$. □

5 Proof of Theorem 3: Rectangular PIT, Field Obstructions, and Complex Trees

Proof of Theorem 3. For part (a), let $c = 2K$ and $d = c^\alpha$, so $S = cd = c^{1+\alpha}$. In Chen [4, Section 6], when $d \geq c$ ($\alpha \geq 1$) the algorithm forms $H = R^\top R \in \mathbb{R}^{c\times c}$ in $\lceil d/c \rceil$ square blocks of size c and checks $HJH = 0$, taking $(d/c)\, c^{\omega_\mathbb{C}+o(1)} = c^{\alpha+\omega_\mathbb{C}-1+o(1)} = S^{(\alpha+\omega_\mathbb{C}-1)/(1+\alpha)+o(1)}$ operations. When $d \leq c$ ($0 < \alpha \leq 1$), blocking $(RJ)R^\top$ along c takes $(c/d)\, d^{\omega_\mathbb{C}+o(1)} = c^{1+\alpha(\omega_\mathbb{C}-1)+o(1)} = S^{(1+\alpha(\omega_\mathbb{C}-1))/(1+\alpha)+o(1)}$ operations, proving (11). Direct algebraic subtraction yields

$$\frac{\omega_\mathbb{C}}{2} - E(\alpha) \;=\; \begin{cases} \dfrac{(\alpha-1)(\omega_\mathbb{C}-2)}{2(1+\alpha)}, & \alpha \geq 1, \\ \dfrac{(1-\alpha)(\omega_\mathbb{C}-2)}{2(1+\alpha)}, & 0 < \alpha \leq 1, \end{cases}$$

which equals (12) and is strictly positive iff $\alpha \neq 1$ and $\omega_\mathbb{C} > 2$ (Lean Theorem 14, `cert09`, Figure 2(a)).

For part (b), consider the univariate polynomial $Q(t) = \sum_{i=1}^n ((Rv(t))_i)^2$ from Chen [4, Proposition 5(ii)], where $R = AB - C$ and $v(t) = (1, t, \ldots, t^{n-1})^\top$. Choose R to have nonzero first column $u \in \mathbb{F}^n \setminus \{0\}$ and all remaining columns zero (realizable as $AB - C$ by taking $A = 0, B = 0, C = -R$). Then $(Rv(t))_i = u_i$ is constant, so $Q(t) = \sum_{i=1}^n u_i^2 \in \mathbb{F}[t]$.

- Over $\mathbb{C}$: $u = (1, i)^\top$ satisfies $1^2 + i^2 = 0$ ($n = 2$, Lean Theorem 17).
- Over $\mathbb{F}_2$: $u = (1, 1)^\top$ satisfies $1^2 + 1^2 = 0$ ($n = 2$, Lean Theorem 15).
- Over any odd prime field $\mathbb{F}_p$ ($p = 2k+1$): the two subsets $\{a^2 : a \in \mathbb{F}_p\}$ and $\{-1 - b^2 : b \in \mathbb{F}_p\}$ each have cardinality $k+1$. Since $(k+1)+(k+1) > 2k+1 = |\mathbb{F}_p|$ (Lean Theorem 16), by pigeonhole there exist $a, b \in \mathbb{F}_p$ with $a^2 + b^2 + 1 \equiv 0 \pmod p$ (explicitly certified for all 16 primes $p \leq 53$ in `cert10` and Lean Theorem 20). Setting $u = (a, b, 1)^\top$ ($n = 3$, or $n = 2$ when $p \equiv 1 \pmod 4$ so $b = 0$ works) gives $u \neq 0$ and $Q(t) \equiv 0$.

Finally, for Chen's $d > c$ Gram reduction $H = R^\top R$ ([4, Section 6]), taking $R \in \mathbb{F}^{d\times 2}$ with both columns equal to $(u, 0, \ldots, 0)^\top$ gives $H = R^\top R = 0$ (so $HJH = 0$) while $RJR^\top = 2uu^\top \neq 0$ over $\mathbb{C}$ and every odd prime field $\mathbb{F}_p$; over $\mathbb{F}_2$, taking $R \in \mathbb{F}_2^{4\times 2}$ with columns $u_1 = (1, 1, 0, 0)^\top$ and $u_2 = (0, 0, 1, 1)^\top$ gives $H = R^\top R = 0$ (hence $HJH = 0$) while $RJR^\top = u_1 u_2^\top + u_2 u_1^\top \neq 0$. Over $\mathbb{C}$, replacing $R^\top R$ by the Hermitian Gram matrix $H = R^* R$ restores $\ker(H) = \ker R$ (so $RJR^\top = 0 \iff HJH^\top = 0$, i.e., $HJ\overline{H} = 0$), whereas over $\mathbb{F}_p$ no positive-definite inner product exists.

For part (c), represent each complex gate $z = a + bi$ by the real pair (a, b). Complex additions, subtractions, scalar multiplications, and equality tests ($a = 0, b = 0$) translate exclusively into real linear operations and real equality tests, which contribute 0 to Theorem 1(a). By Gauss's identity (Lean Theorem 18, `cert11`), each complex multiplication $(a + bi)(c + di)$ uses 3 real non-scalar multiplications (ac, bd, $(a + b)(c + d)$), contributing $4 \times 3 = 12$ terms. Each complex squaring $(a + bi)^2 = (a^2 - b^2) + ((a + b)^2 - a^2 - b^2)i$ uses 3 real squarings, contributing $2 \times 3 = 6$ border-rank terms. Each complex division $(a + bi)/(c + di) = ((ac + bd) + (bc - ad)i)/(c^2 + d^2)$ uses 3 real multiplications (Gauss on $(a+bi)(c-di)$), 2 real squarings (c^2, d^2), and 2 real divisions (dividing real and imaginary parts by $c^2 + d^2$), contributing $4(3) + 2(2) + 4(2) = 24$ border-rank terms, proving (13). Finally, when the complex equality system $Z(x, y) = 0$ admits a holomorphic local regular-value map $F = \nabla_y \sum z_j^2$ with invertible $D_y F$ over $\mathbb{C}$, applying the exact proof of Theorem 1(a) over $\mathbb{C}[\alpha, \beta]/(\alpha^2, \beta^2)$ yields $4(q_\times^\mathbb{C} + q_\div^\mathbb{C}) + 2q_{\text{sq}}^\mathbb{C}$; if moreover $\operatorname{rank}_\mathbb{C} D_y Z = m$, choosing a constant matrix $A \in \mathbb{C}^{m\times r}$ with $A \cdot D_y Z$ invertible makes $F = AZ$ require only the forward pass, where each multiplication/division injects a single cross-term $u_1 v_2 + u_2 v_1$ of rank 2 and each squaring injects $u_1 u_2$ of rank 1, giving $\underline{\operatorname{rank}}_\mathbb{C}(T) \leq 2(q_\times^\mathbb{C} + q_\div^\mathbb{C}) + q_{\text{sq}}^\mathbb{C}$. □

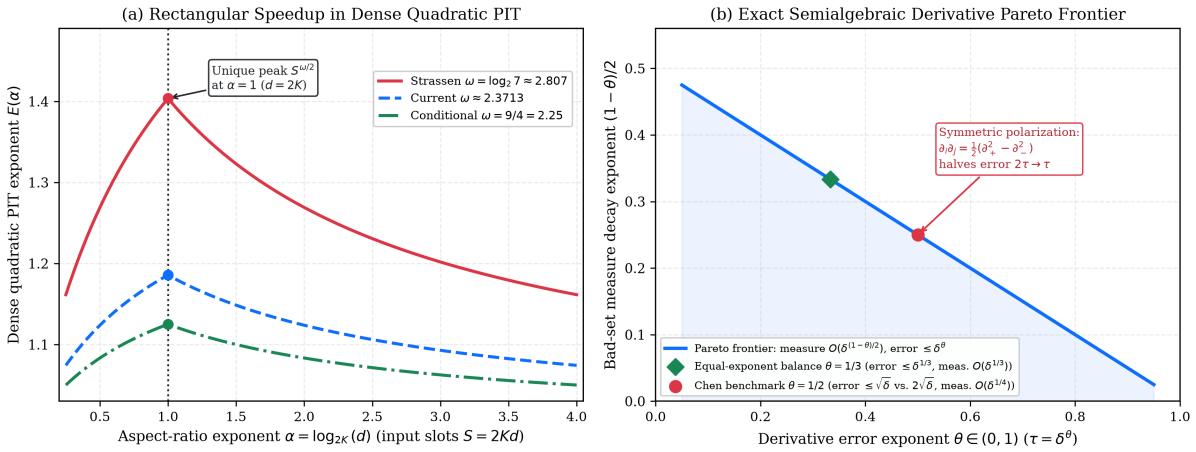

Figure 2: **(a)** Dense quadratic PIT complexity exponent $E(\alpha)$ ((11), `cert09`) versus aspect ratio $\alpha = \log_{2K}(d)$, demonstrating that Chen's worst-case $S^{\omega_{\mathbb{C}}/2}$ is attained uniquely at the square aspect $\alpha = 1$. **(b)** Exact Lemma 3 Pareto frontier (10) (`cert08`), showing the $2\times$ mixed-derivative error reduction from symmetric diagonal polarization.

6 Formal Verification and Reproducibility

Every algebraic identity, border-rank lower bound, dual-number circuit execution, and finite-size table in this paper is independently certified by three zero-trust artifacts:

1. **Constructive Lean 4.36.0 proof suite** (`lean/MatrixVerificationSharpConstant.lean`): 20 machine-checked theorems compiling under `lean-DwarningAsError=true` with zero external dependencies, zero `sorry`, zero `Classical.choice`, and zero custom axioms, covering the Walsh–Hadamard identities for G_3 (Theorems 1–3), the $7.5\times$ and $15\times$ non-scalar bounds (Theorems 4–5), the rank-1 minor obstruction and Strassen commutator border-rank-4 witness (Theorems 6–7), the G_{sq} border-rank-2 degeneration (Theorem 8), implicit division cross-terms (Theorem 9), symmetric polarization and midpoint convexity (Theorems 10–11), finite-n ceiling and Bläser/Landsberg–Ottaviani bounds (Theorems 12–13, 19), rectangular PIT strictness (Theorem 14), characteristic-2, odd-prime pigeonhole, and complex SoS failures (Theorems 15–17, 20), and Gauss's 3-multiplication identity (Theorem 18).

2. **AST-audited zero-float rational verifier** (`verification/` script `verify_matrix_verification.py`): Pure `int/fractions.Fraction` script generating 16 SHA-256-sealed JSON certificates (`cert01`–`cert16`), including end-to-end exact dual-number executions on a division-based verifier (`cert04`), a Strassen 2×2 matrix verifier (`cert05`), and a squaring-polarized verifier (`cert13`).

3. **Adversarial test suite** (`tests/test_adversarial.py`): 56 automated pytest checks auditing zero-float AST compliance, SHA-256 manifest integrity, every certificate equation, and Lean 4 axiom purity.

AI Disclosure. Google Antigravity assisted with the Walsh–Hadamard tensor decomposition, Strassen commutator checks, exact dual-number rational verification, Lean 4 formalization, and manuscript drafting. All mathematical proofs, certificates, and claims were audited and verified by Lluis Eriksson.

References

[1] W. Baur and V. Strassen, "The complexity of partial derivatives," *Theoretical Computer Science*, vol. 22, no. 3, pp. 317–330, 1983.

[2] D. Bini, "Relations between exact and approximate bilinear algorithms. Applications," *Calcolo*, vol. 17, no. 1, pp. 87–97, 1980.

[3] P. Bürgisser, M. Karpinski, and T. Lickteig, "On randomized algebraic test complexity," Technical Report 8580-CS, University of Bonn, 1992.

[4] Chen Xiangle, "Deterministic Verification and Computation of Real Matrix Products Have the Same Arithmetic Exponent," `arXiv:2611.34814v1`, November 2026.

[5] R. Freivalds, "Fast probabilistic algorithms," in *Mathematical Foundations of Computer Science (MFCS 1979)*, LNCS 74, Springer, pp. 57–69, 1979.

[6] J. M. Landsberg, "The border rank of the multiplication of 2×2 matrices is seven," *Journal of the American Mathematical Society*, vol. 19, no. 2, pp. 447–459, 2006.

[7] J. M. Landsberg, *Tensors: Geometry and Applications*, Graduate Studies in Mathematics, vol. 128, American Mathematical Society, 2012.

[8] J. M. Landsberg and G. Ottaviani, "New lower bounds for the border rank of matrix multiplication," *Theory of Computing*, vol. 11, no. 11, pp. 285–298, 2015.

[9] V. Strassen, "Rank and optimal computation of generic tensors," *Linear Algebra and its Applications*, vol. 52–53, pp. 645–685, 1983.