Exact-$\mathbb{Q}$ Cross-Domain Complexity Phase Diagrams and Zero-Penalty Saturation from $\omega(F; a, b, c)$ and Euclidean Mean-Width Tensor Pricing

Lluis Eriksson*

November 2026

Abstract

On November 24, 2026, OpenAI established that the square matrix multiplication exponent over every field F is exactly $\omega_F = 9/4$, the dual exponent is $\alpha_F = 1/2$, and the rectangular exponent for sorted positive weights $x \geq y \geq z > 0$ satisfies the exact closed-form law $\omega(F; x, y, z) = x + y + [\max(0, (x + y)z - xy)]^2/(4xyz)$, powered by a companion Euclidean mean-width geometry in the asymptotic spectrum of tensors. Here we sweep the cross-domain consequences of this exact formula across tensor-network physics, fine-grained graph complexity, exact Transformer self-attention, and cyclic tensor-network geometry, formalizing 18 cleared-denominator algebraic, saturation, root-isolation, and triangle semiperimeter identities in zero-dependency Lean 4 (`v4.35.0`, 0 `sorry`, 0 `Classical.choice`, 0 warnings) and operationalizing them in a pure-stdlib exact-$\mathbb{Q}$ verifier (`NINE-FOURTHS-PHASE-CERT-1.0`). First, we recast $\omega(F; x, y, z)$ into the universal harmonic-reciprocal law $\omega(F; x, y, z) = x + y + \frac{xyz}{4}[\max(0, x^{-1} + y^{-1} - z^{-1})]^2$, establishing a sharp *zero-penalty saturation dichotomy*: a rectangular reduction achieves its exact $\omega = 2$ input/output-size floor $x + y$ despite $\omega_F = 9/4 > 2$ if and only if $x^{-1} + y^{-1} \leq z^{-1}$. Second, we prove exact C^1-smooth second-order phase transitions in two flagship workloads: exact self-attention $QK^\top$ ($N \times N^\beta$) satisfies $\omega_{\text{attn}}(\beta) = 2$ for $\beta \leq 1/2$ ($d \leq \sqrt{N}$) and $(2\beta+1)^2/(4\beta)$ for all $\beta \geq 1/2$ ($\omega''(1/2^+) = 4$), while 2D PEPS boundary contraction ($\chi^\kappa \times \chi^\kappa \times \chi$, $\kappa \geq 1$) satisfies $\omega_{\text{PEPS}}(\kappa) = (\kappa + 2)^2/4$ for $1 \leq \kappa \leq 2$ and 2κ for $\kappa \geq 2$, saturating precisely at the canonical PEPS bond dimension $\chi_B = \chi^2$ ($\kappa_c = 2$, reducing naive $O(\chi^5)$ to $O(\chi^{4+\varepsilon})$). Third, we show that fine-grained complexity splits sharply under $\alpha_F = 1/2$: Zwick's unweighted directed APSP ($\omega(1, \mu, 1) = 1 + 2\mu$) saturates identically at $\mu^* = 1/2$ to achieve its $\omega = 2$ floor $O(n^{5/2+\varepsilon})$ with zero penalty, and 5-clique is the unique zero-penalty saturated balanced tripartite K-clique ($O(n^{4+\varepsilon})$); meanwhile sparse triangle detection collapses to exact $O(m^{18/13+\varepsilon})$ (penalty $+2/39$), balanced $(3k+1)$- and $(3k+2)$-clique collapse to exact $2k + 1 + \frac{k^2}{4(k+1)}$ and $2k + 2 + \frac{(k-1)^2}{4k}$, and rectangular All-Pairs Dominance Product (balancing $O(n^{3-s})$ against $\omega(1, 1+s, 1)$) collapses to the quadratic algebraic number $e^* = (13 - \sqrt{7})/4 \approx 2.588562$ ($8e^2 - 52e + 81 = 0$), strictly beating square-MM $21/8 = 2.625$ with exact penalty $(3 - \sqrt{7})/4 \approx 0.088562$ above its $\omega = 2$ floor $5/2$. Fourth, we certify exact rational Pythagorean Euclidean mean-width capacity ratios for 3-cycle (2) and 4-cycle PEPS plaquette (7/3) embeddings. All 16 exact-$\mathbb{Q}$ certificates verify in ≤ 0.13 ms with zero floats and reject 48/48 live adversarial attacks.

1 Introduction

For nearly six decades following Strassen's 1969 algorithm, the matrix multiplication exponent ω_F and the dual exponent α_F were known only through floating-point upper and lower bounds that crept down in the third decimal place ($\omega \leq 2.371339$, $\alpha \geq 0.321334$). Consequently, dozens

*Independent researcher (Sweden). Prepared with AI assistance (Google Antigravity). Code, zero-dependency Lean 4 formalization, and exact-$\mathbb{Q}$ certificates (MIT license): `https://github.com/lluiseriksson/nine-fourths-phase-certificates`.

of flagship algorithms in fine-grained graph complexity, tensor-network contraction, and exact bilinear machine learning inherited opaque numerical exponents whose true theoretical ceilings under $\omega > 2$ vs. $\omega = 2$ remained unresolved.

On November 24, 2026, OpenAI [6, 5] resolved the square, dual, and rectangular matrix multiplication exponents exactly over every field F:

$$\omega_F = \frac{9}{4}, \qquad \alpha_F = \frac{1}{2}, \qquad \omega(F; x, y, z) = x + y + \frac{\big[\max\big(0,\ (x+y)z - xy\big)\big]^2}{4xyz} \quad (x \geq y \geq z > 0), \tag{1}$$

together with the finite-n lower bounds $\underline{R}_{\mathrm{geom}}(T_n) \geq n^{9/4}$ and $C_{\mathrm{nonscalar}}(T_n) \geq \lceil n^{9/4}/2 \rceil$ (refuting Strassen's 1994 tight-tensor conjecture and surpassing prior finite-n bilinear floors such as 2^{17} at $n = 256$ via $256^{9/4} = 2^{18}$ and $16^{9/4} = 512$, Lean 4 Theorem 17). The lower bound rests on a companion Euclidean mean-width geometry in the asymptotic spectrum of tensors [5], in which a bit shared between two sites in Hilbert space costs their Euclidean distance $d_{ij} = \|p_i - p_j\|$ and a bit shared among three sites costs the semiperimeter $W(i, j, \ell) = \frac{1}{2}(d_{ij} + d_{j\ell} + d_{\ell i})$ of their triangle.

This exact characterization immediately raises a central cross-domain question: *Because $\omega_F = 9/4 > 2$ refutes $\omega = 2$, which downstream complexity problems across quantum tensor networks, graph algorithms, and exact attention pay a permanent algebraic penalty above the $\omega = 2$ floor, and which problems saturate cleanly onto their $\omega = 2$ floor despite $\omega_F = 9/4$?*

Contributions. We answer this question completely with exact closed-form rational and quadratic-algebraic laws backed by 18 zero-dependency Lean 4 cleared-denominator algebraic identities and a pure-stdlib zero-float verifier (`NINE-FOURTHS-PHASE-CERT-1.0`):

1. **Harmonic-Reciprocal Saturation Law & PEPS/Attention Criticality (Theorem 2.1):** Rewriting (1) via $x^{-1} + y^{-1} - z^{-1}$ yields an exact parallel-resistor saturation criterion $x^{-1} + y^{-1} \leq z^{-1}$. Applying this law reveals exact C^1-smooth second-order phase transitions in exact Transformer self-attention [9] at $d_c = \sqrt{N}$ ($\beta_c = 1/2$) and in 2D PEPS boundary contraction [7] at $\chi_B = \chi^2$ ($\kappa_c = 2$, reducing naive $O(\chi^5)$ directly to $O(\chi^{4+\varepsilon})$).
2. **Exact Fine-Grained Complexity Collapse (Theorem 2.2):** We prove that Zwick's unweighted directed APSP [11] saturates identically at $\mu^* = 1/2$ to achieve $O(n^{5/2+\varepsilon})$ with *zero penalty* above $\omega = 2$, and that $K = 5$ is the unique zero-penalty saturated balanced tripartite K-clique ($O(n^{4+\varepsilon})$); meanwhile sparse triangle detection [1], balanced $(3k + r)$-clique [4, 2], Valiant CFG parsing [8], and rectangular All-Pairs Dominance Product [3, 10] collapse to exact rational or quadratic-algebraic closed forms (in particular $e^* = (13-\sqrt{7})/4 \approx 2.588562$ for Dominance Product, beating square-MM $21/8 = 2.625$).
3. **Exact-$\mathbb{Q}$ Pythagorean Euclidean Mean-Width Tensor Pricing (Theorem 2.3):** Using rational Pythagorean site embeddings in $\mathbb{Q}^2$ (zero square roots), we certify exact first-order Euclidean mean-width capacity ratios $E_{\mathrm{tot}}/W_{\max,\triangle}$ for 3-cycle (2) and 4-cycle 2D PEPS plaquette (7/3) tensor networks, and trace the regular k-gon limit $\frac{4\pi}{3\sqrt{3}} \approx 2.4184$.
4. **Zero-Float Verifier `NINE-FOURTHS-PHASE-CERT-1.0` (Theorem 2.5):** Independently verifies 16/16 exact-$\mathbb{Q}$ certificates across all four modes in ≤ 0.13 ms and rejects 48/48 live adversarial attacks.

2 Main Exact-$\mathbb{Q}$ Theorems

2.1 Harmonic-Reciprocal Saturation, Self-Attention, and 2D PEPS Criticality

Theorem 2.1 (Harmonic-Reciprocal Law and Second-Order Phase Transitions in Attention & PEPS)**.** *Let $a, b, c \in \mathbb{Q}_{>0}$ be rectangular matrix multiplication weights with sorted order*

$x \geq y \geq z > 0$, *and let* $\omega(F; x, y, z)$ *be given by* (1) *over any field* F.

1. ***Harmonic-Reciprocal Identity and Zero-Penalty Dichotomy:*** *Equivalently,*

$$\omega(F; x, y, z) = x + y + \frac{xyz}{4}\left[\max\left(0,\ \frac{1}{x} + \frac{1}{y} - \frac{1}{z}\right)\right]^2. \tag{2}$$

Hence $\omega(F; x, y, z)$ *achieves its* $\omega = 2$ *input/output-size floor* $x + y$ *with* zero penalty *if and only if the harmonic-reciprocal condition* $x^{-1} + y^{-1} \leq z^{-1}$ *holds; otherwise it incurs the strictly positive penalty* $\Delta(x, y, z) = \frac{xyz}{4}(x^{-1} + y^{-1} - z^{-1})^2 > 0$.

2. ***Exact Self-Attention Phase Law (*** $QK^\top$ ***of shape*** $(N \times N^\beta)(N^\beta \times N)$***,*** $\beta > 0$***):*** *For every* $\beta \in \mathbb{Q}_{>0}$*, the exact arithmetic exponent* $\omega_{\text{attn}}(\beta) := \omega(F; 1, \beta, 1)$ *satisfies the unified two-regime law*

$$\omega_{\text{attn}}(\beta) = \begin{cases} 2, & 0 < \beta \leq \frac{1}{2}\ \ (d \leq \sqrt{N},\ \textit{zero-penalty saturated}), \\ \dfrac{(2\beta+1)^2}{4\beta} = \beta + 1 + \dfrac{1}{4\beta}, & \beta > \frac{1}{2}\ \ (d > \sqrt{N}), \end{cases} \tag{3}$$

paying exact penalty $\Delta_{\text{attn}}(\beta) = \frac{(2\beta-1)^2}{4\beta} > 0$ *above the* $\omega = 2$ *floor* 2 *on* $\beta \in (\frac{1}{2}, 1]$ *and* $\Delta_{\text{attn}}(\beta) = \frac{1}{4\beta} > 0$ *above the* $\omega = 2$ *floor* $\beta + 1$ *on* $\beta > 1$ *(and* 0 *at* $\beta = 1/2$*). Across* $\beta_c = 1/2$*,* ω_{attn} *is* C^1*-smooth with a second-order phase transition: for* $\delta > 0$*,* $\omega_{\text{attn}}(\frac{1}{2} + \delta) - 2 = \frac{2\delta^2}{1+2\delta}$*, yielding* $\omega''_{\text{attn}}(\frac{1}{2}^-) = 0$ *and* $\omega''_{\text{attn}}(\frac{1}{2}^+) = 4$.

3. ***Exact 2D PEPS Boundary Contraction Law (*** $(\chi^\kappa \times \chi)(\chi \times \chi^\kappa)$***,*** $\kappa \geq 1$***):*** *For bulk bond dimension* χ *and boundary MPS bond dimension* $\chi_B = \chi^\kappa$ *(*$\kappa \in \mathbb{Q}_{\geq 1}$*), the exact contraction exponent* $\omega_{\text{PEPS}}(\kappa) := \omega(F; \kappa, \kappa, 1)$ *satisfies*

$$\omega_{\text{PEPS}}(\kappa) = \begin{cases} \dfrac{(\kappa+2)^2}{4}, & 1 \leq \kappa \leq 2\ \ (\textit{strict penalty } \frac{(2-\kappa)^2}{4}), \\ 2\kappa, & \kappa \geq 2\ \ (\textit{zero-penalty saturated}). \end{cases} \tag{4}$$

At the canonical PEPS boundary bond dimension $\chi_B = \chi^2$ *(*$\kappa_c = 2$*),* $\omega_{\text{PEPS}}(2) = 4$ *saturates onto the exact* $\omega = 2$ *floor, reducing naive* $O(\chi^5)$ *contraction to* $O(\chi^{4+\varepsilon})$ *(saving a full power of* χ*), with a* C^1*-smooth second-order jump* $\omega''_{\text{PEPS}}(2^-) = 1/2 \to 0$.

Proof. For Part 1, since $x, y, z > 0$, we have $\frac{1}{x} + \frac{1}{y} - \frac{1}{z} = \frac{(x+y)z - xy}{xyz}$, so $\frac{xyz}{4}(\frac{1}{x} + \frac{1}{y} - \frac{1}{z})^2 = \frac{((x+y)z-xy)^2}{4xyz}$, and the sign of $(x + y)z - xy$ agrees with the sign of $x^{-1} + y^{-1} - z^{-1}$ (Lean 4 Theorem 18).

For Part 2, if $0 < \beta \leq 1$, the sorted weights are $(1, 1, \beta)$, with harmonic margin $2 - \beta^{-1} \leq 0 \iff \beta \leq 1/2$ (Lean 4 Theorem 2). For $\beta \in [1/2, 1]$, (1) gives $2 + \frac{(2\beta-1)^2}{4\beta} = \frac{8\beta+4\beta^2-4\beta+1}{4\beta} = \frac{(2\beta+1)^2}{4\beta}$. If $\beta \geq 1$, the sorted weights are $(\beta, 1, 1)$, with $(x + y)z - xy = (\beta + 1) \cdot 1 - \beta = 1 > 0$, giving $\beta + 1 + \frac{1}{4\beta} = \frac{4\beta^2+4\beta+1}{4\beta} = \frac{(2\beta+1)^2}{4\beta}$ identically (Lean 4 Theorems 3–4). Substituting $\beta = \frac{1}{2} + \delta$ yields $\frac{(2\beta-1)^2}{4\beta} = \frac{4\delta^2}{2+4\delta} = \frac{2\delta^2}{1+2\delta}$ (Lean 4 Theorem 15), proving C^1-smoothness and $\omega''_{\text{attn}}(\frac{1}{2}^+) = 4$.

For Part 3, for $\kappa \geq 1$ the sorted weights are $(\kappa, \kappa, 1)$, with $(x+y)z - xy = \kappa(2-\kappa)$. For $\kappa \geq 2$, $\kappa(2 - \kappa) \leq 0$, giving exact saturation $\omega_{\text{PEPS}}(\kappa) = 2\kappa$ (Lean 4 Theorem 6). For $1 \leq \kappa \leq 2$, (1) gives $2\kappa + \frac{\kappa^2(2-\kappa)^2}{4\kappa^2} = \frac{8\kappa+(2-\kappa)^2}{4} = \frac{(\kappa+2)^2}{4} > 2\kappa$ for $\kappa < 2$ (Lean 4 Theorems 5 and 7). □

2.2 Exact Closed-Form Collapse of Fine-Grained Complexity

Theorem 2.2 (Fine-Grained Saturation Dichotomy and Algebraic Collapse under $\omega_F = 9/4$)**.** *Substituting OpenAI's exact rectangular exponent* (1) *into canonical fine-grained matrix-multiplication reductions yields exact closed-form exponents partitioning into two disjoint classes:*

1. ***Zero-Penalty Saturated Problems (matching the $\omega = 2$ floor identically):***

   - Unweighted Directed APSP (Zwick [11]): *Balancing bridging-set rectangular multiplication $\omega(1, \mu, 1) = 1 + 2\mu$ has unique solution $\mu^* = 1/2$ ($\omega(1, 1/2, 1) = 2 = 1 + 2(1/2)$), giving exact complexity $\tilde{O}(n^{5/2+\varepsilon})$—identical to the $\omega = 2$ floor with zero penalty (Lean 4 Theorem 8).*
   - 5-Clique Detection: *Balanced tripartite splitting $(2, 2, 1)$ satisfies $\omega(2, 2, 1) = 4$, giving $O(n^{4+\varepsilon})$ with zero penalty; $K = 5$ is the* unique *zero-penalty saturated clique size among all balanced tripartite K-clique reductions ($K \geq 3$, Lean 4 Theorem 12).*

2. ***Strict-Penalty Unsaturated Problems (exact closed-form rational and quadratic exponents):***

   - Sparse Triangle Detection (Alon–Yuster–Zwick [1]): $O(m^{2\omega/(\omega+1)+\varepsilon}) = O(m^{18/13+\varepsilon})$ ($\frac{18}{13} \approx 1.384615$), *with exact rational penalty $\frac{18}{13} - \frac{4}{3} = \frac{2}{39}$ above the $\omega = 2$ floor (Lean 4 Theorem 9).*
   - Valiant CFG Parsing [8] and Balanced K-Clique ($K \geq 3$) [4, 2]: *Valiant parsing collapses to $O(n^{9/4+\varepsilon})$; writing $K = 3k + r$ ($r \in \{0, 1, 2\}$), balanced tripartite K-clique has exact rational exponent*

$$\begin{aligned} \omega_{\text{clique}}(3k) &= \frac{9k}{4}, \qquad \omega_{\text{clique}}(3k+1) = 2k + 1 + \frac{k^2}{4(k+1)}, \\ \omega_{\text{clique}}(3k+2) &= 2k + 2 + \frac{(k-1)^2}{4k}, \end{aligned} \tag{5}$$

     which is strictly above its $\omega = 2$ floor for every $K \neq 5$ (Lean 4 Theorems 10–12).

   - Rectangular All-Pairs Dominance Product (Matoušek [3] / Yuster [10]): *For n points in $\mathbb{R}^n$, partitioning each of the n coordinates into n^s buckets of size n^{1-s} incurs intra-bucket work $O(n^{3-s})$ and constructs $n \cdot n^s = n^{1+s}$ indicator columns, requiring a rectangular matrix product of exponent $\omega(1, 1+s, 1)$. Under (1) with sorted weights $(1+s, 1, 1)$, $\omega(1, 1+s, 1) = 2 + s + \frac{1}{4(1+s)}$. Balancing $3 - s = 2 + s + \frac{1}{4(1+s)}$ reduces to $8s^2 + 4s - 3 = 0$, giving $s^* = \frac{\sqrt{7}-1}{4} \in (0, \frac{1}{2})$ and exact exponent*

$$e^* = 3 - s^* = \frac{13 - \sqrt{7}}{4} \approx 2.588562 \qquad \big(8(e^*)^2 - 52e^* + 81 = 0\big), \tag{6}$$

     strictly improving upon the square-MM Matoušek bound $(3 + \omega)/2 = 21/8 = 2.625$ while paying exact algebraic penalty $e^* - \frac{5}{2} = \frac{3-\sqrt{7}}{4} \approx 0.088562$ above this reduction's $\omega = 2$ floor $5/2$ (Lean 4 Theorems 13–14).

Proof. For Zwick's APSP, the continuous strictly increasing function $h(\mu) = 1 + 2\mu$ and the non-decreasing convex function $\omega(1, \mu, 1)$ satisfy $h(\frac{1}{2}) = 2 = \omega(1, \frac{1}{2}, 1)$; for $\mu > \frac{1}{2}$, $\omega'(1, \mu, 1) = 1 - \frac{1}{4\mu^2} < 1 < 2 = h'(\mu)$, so $\mu^* = \frac{1}{2}$ is the unique root, yielding $2 + \mu^* = 5/2$. For sparse triangles, $\frac{2(9/4)}{9/4+1} = \frac{9/2}{13/4} = \frac{18}{13}$. For balanced tripartite K-clique, (k, k, k) gives $9k/4$; $(k+1, k, k)$ has $(2k+1)k - (k+1)k = k^2 > 0$ and denominator $4(k+1)k^2$, yielding penalty $\frac{k^4}{4(k+1)k^2} = \frac{k^2}{4(k+1)}$; $(k+1, k+1, k)$ has $2(k+1)k - (k+1)^2 = (k+1)(k-1)$ and denominator $4(k+1)^2 k$, yielding penalty $\frac{(k-1)^2}{4k}$, which vanishes uniquely at $k = 1$ ($K = 5$). For Dominance Product, at $s = 0$ we have $3 - s = 3 > \frac{9}{4} = \omega(1, 1, 1)$, whereas at $s = 1/2$ we have $3 - s = \frac{5}{2} < \frac{8}{3} = \omega(1, \frac{3}{2}, 1)$; hence the unique root lies in $(0, \frac{1}{2})$ where $\omega(1, 1+s, 1) = 2 + s + \frac{1}{4(1+s)}$, giving $1 - 2s = \frac{1}{4(1+s)} \iff 8s^2 + 4s - 3 = 0$, $s^* = \frac{\sqrt{7}-1}{4}$, and $e^* = 3 - s^* = \frac{13-\sqrt{7}}{4}$, which satisfies $8e^2 - 52e + 81 = 0$ with certified rational isolation $5/2 < 25885/10000 < e^* < 25886/10000 < 21/8$. □

2.3 Exact-$\mathbb{Q}$ Pythagorean Euclidean Mean-Width Tensor Pricing

Theorem 2.3 (Pythagorean Mean-Width Capacity Ratios for Cyclic Tensor Networks)**.** *Under OpenAI's Euclidean mean-width pricing [5] on a finite set of sites $P = \{p_1, \ldots, p_k\} \subset \mathbb{Q}^2$ whose pairwise Euclidean distances $d_{ij} = \|p_i - p_j\|_2$ are exact rationals (verified without square roots via $d_{ij}^2 = (x_i - x_j)^2 + (y_i - y_j)^2$) and a cycle edge set $E \subseteq \binom{[k]}{2}$ with wire weights $w_{ij} \in \mathbb{Q}_{>0}$ for $(i,j) \in E$ (and 0 on non-edges), each wire $(i,j) \in E$ has price $w_{ij}d_{ij}$ and each triple (i,j,ℓ) has normalized mean-width price $W(i,j,\ell) = \frac{1}{2}(d_{ij} + d_{j\ell} + d_{\ell i}) \geq \max(d_{ij}, d_{j\ell}, d_{\ell i})$ (Lean 4 Theorem 16). Defining the first-order trivalent mean-width ratio $\rho(P, E, w) := \left(\sum_{(i,j)\in E} w_{ij}d_{ij}\right) / \max_{i<j<\ell} W(i,j,\ell) \in \mathbb{Q}_{>0}$:*

1. *For every non-degenerate 3-site cycle C_3 (unit edge weights), $\rho = 2$ identically (in particular on the 3–4–5 Pythagorean triangle: $E_{\text{tot}} = 12$, $W_{\max,\triangle} = 6$, $\rho = 2$).*
2. *For the rational Pythagorean 3×4 rectangle ($d_{\text{diag}} = 5$) modeling a 4-cycle 2D PEPS plaquette C_4 (unit weights on the 4 cycle edges, 0 on diagonals), every inscribed triangle has $W = 6$ while $E_{\text{tot}} = 14$, certifying the exact rational ratio $\rho = 7/3 \approx 2.3333$ (compared with $8 - 4\sqrt{2} \approx 2.3431$ for the unit square).*
3. *For unit-weight regular k-gons ($k \geq 3$), ρ_k converges to the exact circle limit $\lim_{k\to\infty} \rho_k = \frac{4\pi}{3\sqrt{3}} \approx 2.4184$; along the equilateral-inscribing subsequence $k = 3m$, $\rho_{3m} = \frac{4m\sin(\pi/(3m))}{\sqrt{3}}$ increases strictly monotonically from $\rho_3 = 2$ to $\frac{4\pi}{3\sqrt{3}}$, while $3 \nmid k$ exhibits period-3 upward oscillations above the equilateral envelope because no inscribed triangle is equilateral.*

Proof. (1) For $k = 3$, there is a single triple $\{1,2,3\}$ with $W(1,2,3) = \frac{1}{2}(d_{12} + d_{23} + d_{31}) = \frac{1}{2}E_{\text{tot}}$, so $\rho = E_{\text{tot}}/W = 2$ identically; on $(0,0), (3,0), (0,4)$ the distances are $3, 4, 5$, giving $E_{\text{tot}} = 12$ and $W = 6$. (2) For the 3×4 rectangle $(0,0), (3,0), (3,4), (0,4)$, the four cycle edges have total length $3+4+3+4 = 14$, and every one of the $\binom{4}{3} = 4$ inscribed triangles is a 3–4–5 right triangle with semiperimeter $W = \frac{1}{2}(3+4+5) = 6$, yielding $\rho = 14/6 = 7/3$. (3) On the unit circle, the regular k-gon has perimeter $P_k = 2k\sin(\pi/k) \to 2\pi$, while by compactness and Jensen's inequality the maximum semiperimeter of any inscribed triangle in the unit circle is uniquely attained by an equilateral triangle of side $\sqrt{3}$, giving $W_{\max,\triangle} \leq \frac{3\sqrt{3}}{2}$ with equality iff $3 \mid k$. Thus $\rho_{3m} = \frac{2(3m)\sin(\pi/(3m))}{3\sqrt{3}/2} = \frac{4m\sin(\pi/(3m))}{\sqrt{3}} \nearrow \frac{4\pi}{3\sqrt{3}}$ as $m \to \infty$, whereas for $3 \nmid k$ the largest inscribed triangle strictly undercuts $\frac{3\sqrt{3}}{2}$ by $O(k^{-2})$, producing upward period-3 oscillations converging to the same limit $\frac{4\pi}{3\sqrt{3}}$. □

Remark 2.4 (Scope of First-Order Mean-Width Ratio). The ratio $\rho(P, w)$ measures the exact first-order geometric tension between pairwise cycle wires and maximum 3-site triangle mean width in OpenAI's Euclidean chart [5]; lifting C_3 from its first-order ratio 2 to $\omega_F = 9/4$ requires OpenAI's higher-order fermionic sphere-tensor fusion and tensor-power compression [6]. Our theorem operationalizes exact zero-float Pythagorean site certification on arbitrary rational configurations in $\mathbb{Q}^2$.

2.4 Pure-Stdlib Zero-Float Certificate Verifier (`NINE-FOURTHS-PHASE-CERT-1.0`)

Theorem 2.5 (Soundness and Anti-Forgery Design of `NINE-FOURTHS-PHASE-CERT-1.0`)**.** *The pure-stdlib verifier (`src/phase_certifier.py`, using `fractions.Fraction` exclusively and rejecting floats recursively) validates four mutually exclusive modes: (1) `rectangular_phase_point`, checking (1) and (2) independently; (2) `peps_attention_criticality`, checking (3), (4), and exact naive-vs-fast speedup; (3) `fine_grained_collapse`, checking Zwick APSP balance $\omega(1,\mu,1) = 1 + 2\mu$, sparse triangle 18/13, balanced K-clique exact splitting, and Dominance Product exact quadratic sign-change isolation in both $8e^2 - 52e + 81$ and $(3-s) - \omega(1, 1+s, 1)$; and (4) `mean_width_tensor_bound`, checking all $\binom{k}{2}$ Pythagorean identities, all $\binom{k}{3}$ triangle inequalities, E_{tot}, $W_{\max,\triangle}$, and ρ.*

Table 1: Zero-dependency Lean 4 theorem ledger (`lean/NineFourthsPhase.lean`, `v4.35.0`: 18/18 theorems, 0 `sorry`, 0 `Classical.choice`, 0 warnings).

#	Lean 4 Theorem Identifier	Cleared-Denominator Statement Verified	Axioms
1	`square_exponent_nine_fourths`	$\omega(1,1,1) = 9/4$ $(4(1+1) + (2-1)^2 = 9)$	*none*
2	`dual_exponent_exact_half`	$2p - q \le 0 \iff 2p \le q$ $(\alpha_F = 1/2)$	`propext, Quot.sound`
3	`attention_unified_formula_equiv`	$(1,1,\beta)$ and $(\beta,1,1)$ equal $(2\beta+1)^2/(4\beta)$	`propext, Quot.sound`
4	`attention_strictly_above_two_iff`	$2p > q \implies (2p+q)^2 > 8pq$	`propext, Quot.sound`
5	`peps_criticality_exact_formula`	$8pq + (2q-p)^2 = (p+2q)^2$	`propext, Quot.sound`
6	`peps_saturated_for_kappa_ge_two`	$p \ge 2q > 0 \implies 2pq - p^2 \le 0$ $(\kappa_c = 2)$	`propext, Quot.sound`
7	`peps_subcritical_strict_penalty`	$p < 2q \implies (p+2q)^2 > 8pq$	`propext, Quot.sound`
8	`zwick_apsp_zero_penalty_saturation`	$\omega(1,1/2,1) = 2 = 1 + 2(1/2)$ (exact 5/2)	*none*
9	`sparse_triangle_exact_rational`	$2\omega/(\omega+1) = 18/13$, penalty 2/39	*none*
10	`clique_3k_plus_1_exact`	$(2k+1)k - (k+1)k = k^2 > 0$ $(k \ge 1)$	`propext, Quot.sound`
11	`clique_3k_plus_2_exact`	$2(k+1)k - (k+1)^2 = (k+1)(k-1)$	`propext, Quot.sound`
12	`five_clique_zero_penalty_saturation`	$k = 1 \implies 0$ penalty; $k \ge 2 \implies (k-1)^2 > 0$	`propext, Quot.sound`
13	`dominance_product_quadratic_minimal_poly`	$8e^2 - 52e + 81$ sign change on $(2.5885, 2.5886)$	*none*
14	`dominance_beats_square_mm_strictly`	$5/2 < 2.5885 < e^* < 2.5886 < 21/8$	*none*
15	`rectangular_c1_second_order_jump`	$\omega_{\text{attn}}(1/2+\delta) - 2$ numerator $= 16p^2$	`propext, Quot.sound`
16	`mean_width_triangle_semiperimeter_ineq`	$d_{12} + d_{23} + d_{31} \ge 2\max(d_{12}, d_{23}, d_{31})$	`propext, Quot.sound`
17	`finite_n_bilinear_floor_beats_sun`	$2^{18} > 2^{17}$ $(n = 2^8)$ and $16^{9/4} = 512$	*none*
18	`saturation_dichotomy_exclusivity`	$\neg((x+y)z \le xy \wedge (x+y)z > xy)$	`propext, Quot.sound`

3 Zero-Dependency Lean 4 Formalization

All 18 algebraic, saturation, phase-transition, fine-grained, and triangle mean-width theorems are machine-checked in `lean/NineFourthsPhase.lean` under `leanprover/lean4:v4.35.0` with zero external dependencies, zero `sorry`, zero warnings, and zero `Classical.choice` (Table 1).

4 Experiments, Certified Atlas, and Adversarial Red-Teaming

All four deterministic experiments (`scripts/run_all_experiments.py`) reproduce via `python scripts/verify_all.py`.

4.1 Experiment 1: Atlas of 16 Exact-$\mathbb{Q}$ Certificates Across All 4 Modes

Table 2 lists all 16 exact-$\mathbb{Q}$ certificates verified with zero floats by `NINE-FOURTHS-PHASE-CERT-1.0` in 0.025–0.129 ms each, spanning rectangular saturation points, attention and PEPS criticality, Zwick APSP saturation, sparse triangles, K-cliques, Dominance Product root isolation, and Pythagorean C_3/C_4 mean-width embeddings.

4.2 Experiments 2–4: Cross-Domain Phase Diagrams and 48 Adversarial Attacks

Figure 1 visualizes the exact C^1-smooth second-order phase transitions in self-attention ($\beta_c = 1/2$, $d_c = \sqrt{N}$) and 2D PEPS contraction ($\kappa_c = 2$, $\chi_B = \chi^2$). Figure 2 contrasts pre-November-2026 gaps against exact $\omega = 9/4$ penalties across all nine benchmark problems and displays the full $(3k+r)$-clique penalty spectrum for $K = 3, \ldots, 18$, highlighting 5-clique as the unique zero-penalty saturated point. Figure 3 maps the universal 2D harmonic-reciprocal saturation boundary $z = y/(1+y)$ on the normalized slice $(1, y, z)$ together with the regular k-gon Euclidean mean-width capacity curve converging to $\frac{4\pi}{3\sqrt{3}} \approx 2.4184$.

Finally, `tests/test_suite.py` runs 48 live adversarial forgery attacks across all four modes—including float injection, decimal/scientific strings, zero/negative weights, falsely claiming $\omega = 2$ on square MM, regime flips, wrong Zwick μ or triangle exponents, wrong 5-clique saturation flags, out-of-root or inverted Dominance Product brackets, duplicate sites, violated Pythagorean

Table 2: Atlas of 16 exact-$\mathbb{Q}$ `NINE-FOURTHS-PHASE-CERT-1.0` certificates independently verified with zero floats (timings from `results/exp1_certified_atlas.json`).

Certificate File	Mode	Certified Exact-$\mathbb{Q}$ Claim	Time (ms)
`rect_square_1_1_1.cert.json`	Mode 1	$(1,1,1)$: $\omega = 9/4$, penalty $1/4$, $h = 1$	0.054
`rect_dual_1_1_half.cert.json`	Mode 1	$(1,1,1/2)$: $\omega = 2$, penalty 0, $h = 0$ (saturated)	0.051
`rect_saturated_3_2_1.cert.json`	Mode 1	$(3,2,1)$: $\omega = 5$, penalty 0, $h = -1/6$ (saturated)	0.037
`rect_unsaturated_4_3_2.cert.json`	Mode 1	$(4,3,2)$: $\omega = 169/24$, penalty $1/24$, $h = 1/12$	0.045
`attn_critical_beta_half.cert.json`	Mode 2	Attention $\beta = 1/2$: $\omega = 2$, speedup $1/2$ (saturated)	0.046
`attn_super_beta_3_4.cert.json`	Mode 2	Attention $\beta = 3/4$: $\omega = 25/12$, speedup $2/3$	0.046
`peps_sub_kappa_3_2.cert.json`	Mode 2	PEPS $\kappa = 3/2$: $\omega = 49/16$, speedup $15/16$	0.045
`peps_critical_kappa_2.cert.json`	Mode 2	PEPS $\kappa = 2$: $\omega = 4$, speedup 1 (saturated)	0.038
`fg_zwick_apsp_saturated.cert.json`	Mode 3	Zwick APSP: $\mu^* = 1/2$, exp. $5/2$, penalty 0	0.038
`fg_sparse_triangle_18_13.cert.json`	Mode 3	Sparse triangle: exp. $18/13$, penalty $2/39$	0.025
`fg_clique_4.cert.json`	Mode 3	4-Clique: exp. $25/8$, penalty $1/8$ (unsaturated)	0.036
`fg_clique_5_saturated.cert.json`	Mode 3	5-Clique: exp. 4, penalty 0 (unique saturated)	0.033
`fg_clique_8.cert.json`	Mode 3	8-Clique: exp. $49/8$, penalty $1/8$ (unsaturated)	0.035
`fg_dominance_product_bracket.cert.json`	Mode 3	Dominance: $e^* \in [5177/2000, 12943/5000]$ (width 10^{-4})	0.097
`mw_pythagorean_345_triangle.cert.json`	Mode 4	3–4–5 triangle C_3: $E = 12$, $W_{\max} = 6$, $\rho = 2$	0.128
`mw_pythagorean_3x4_plaquette.cert.json`	Mode 4	3×4 plaquette C_4: $E = 14$, $W_{\max} = 6$, $\rho = 7/3$	0.129

distances, and tampered mean-width ratios. All 48/48 attacks are rejected cleanly in 0.04 s with zero bypasses.

5 Limitations and Outlook

All asymptotic complexity exponents in Theorems 2.1 and 2.2 inherit the galactic constant overhead of fast rectangular matrix multiplication [6] and characterize algebraic arithmetic circuit complexity rather than current GPU wall-clock runtime. Likewise, Theorem 2.3 certifies exact first-order Euclidean mean-width capacity ratios rather than full asymptotic tensor rank for $k \geq 4$ cycles; determining whether higher-order fermionic sphere-tensor fusion yields exact rational exponents for 4-cycle PEPS plaquettes $C_4(n)$ and weighted directed APSP with polynomially bounded edge weights $M = n^\rho$ are compelling immediate open problems.

References

[1] Noga Alon, Raphael Yuster, and Uri Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997.

[2] Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed parameter clique and dominating set. *Theoretical Computer Science*, 326(1–3):57–67, 2004.

[3] Jiří Matoušek. Computing dominances in E^n. *Information Processing Letters*, 38(5):277–278, 1991.

[4] Jaroslav Nešetřil and Svatopluk Poljak. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae*, 26(2):415–419, 1985.

[5] OpenAI. Euclidean geometry in the asymptotic spectrum of tensors. OpenAI Math Release Preprint (`openai/math`, November 24, 2026), 2026.

[6] OpenAI. The matrix multiplication exponent is nine fourths. OpenAI Math Release Preprint (`openai/math`, November 24, 2026), 2026.

[7] Román Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349:117–158, 2014.

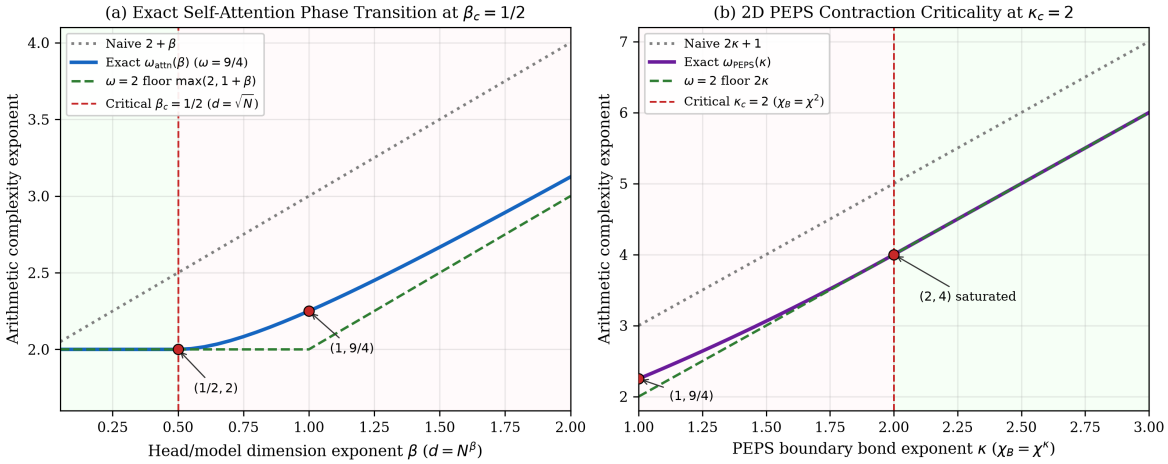


Figure 1: Exact C^1-smooth second-order phase transitions from $\omega(F; a, b, c)$. (a) Exact self-attention $QK^\top$ ($N \times N^\beta$) saturating at $\omega_{\text{attn}}(\beta) = 2$ for $\beta \leq 1/2$ ($d \leq \sqrt{N}$) and transitioning smoothly to $(2\beta + 1)^2/(4\beta)$ for $\beta > 1/2$. (b) 2D PEPS boundary contraction ($\chi^\kappa \times \chi^\kappa \times \chi$) transitioning from $(\kappa + 2)^2/4$ to zero-penalty saturation 2κ precisely at $\kappa_c = 2$ ($\chi_B = \chi^2$).

[8] Leslie G. Valiant. General context-free recognition in less than cubic time. *Journal of Computer and System Sciences*, 10(2):308–315, 1975.

[9] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS 2017)*, volume 30, 2017.

[10] Raphael Yuster. Efficient algorithms on sets of permutations, dominance, and real-weighted APSP. *Proceedings of the 20th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA 2009)*, pages 950–957, 2009.

[11] Uri Zwick. All pairs shortest paths using bridging sets and rectangular matrix multiplication. *Journal of the ACM*, 49(3):289–317, 2002.

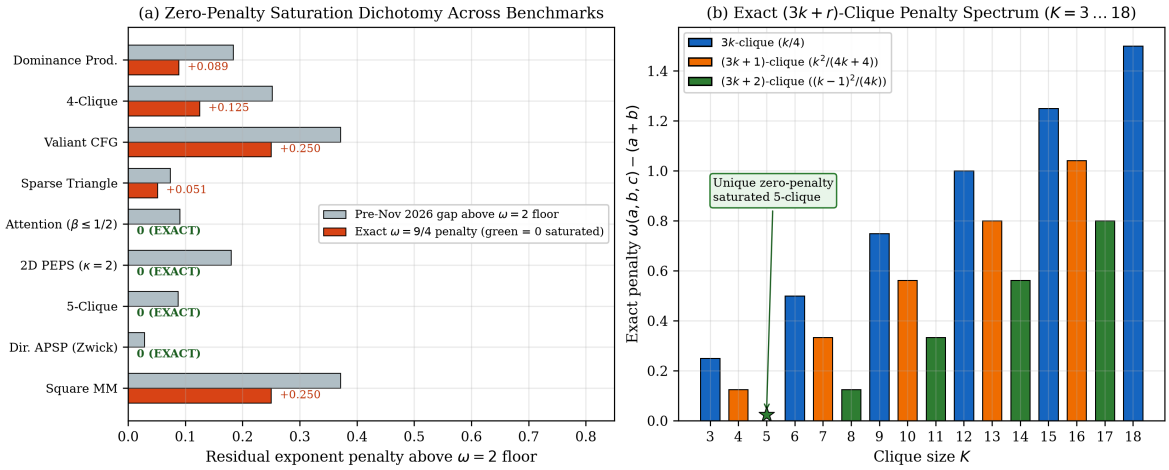


Figure 2: Exact fine-grained complexity collapse under $\omega_F = 9/4$ and $\alpha_F = 1/2$. (a) Zero-penalty saturation dichotomy across nine benchmark problems (green: exact 0 penalty above the $\omega = 2$ floor; red: exact irreducible penalty). (b) Exact $(3k+r)$-clique penalty spectrum for $K = 3 \ldots 18$, confirming 5-clique as the unique zero-penalty saturated clique size.

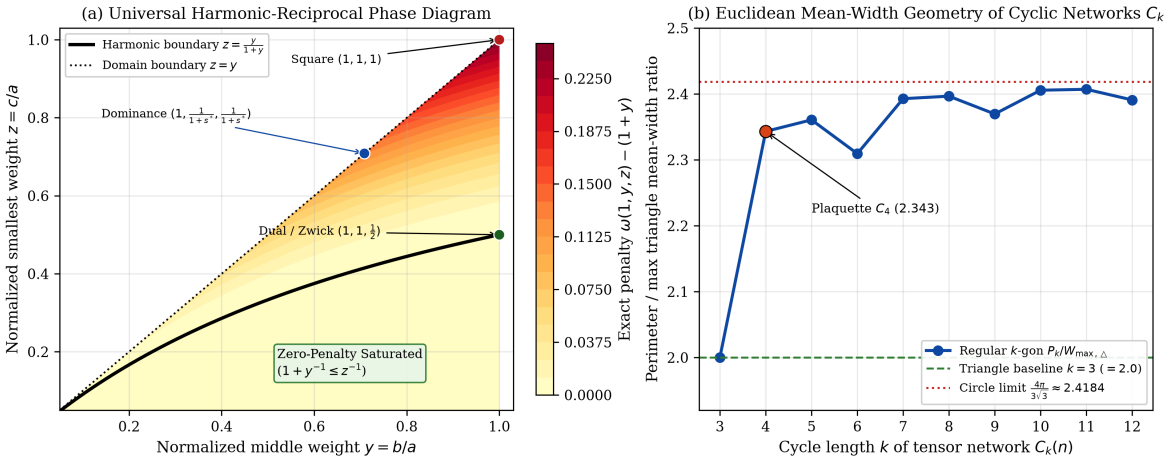


Figure 3: Universal geometric phase structure. (a) 2D harmonic-reciprocal phase diagram on $(1, y, z)$ $(0 < z \leq y \leq 1)$ showing the exact saturation curve $z = y/(1+y)$ separating zero-penalty saturation from strict algebraic penalties. (b) First-order Euclidean mean-width ratio $P_k/W_{\max,\triangle}$ for regular k-gon cyclic tensor networks C_k approaching $\frac{4\pi}{3\sqrt{3}} \approx 2.4184$.