Sharp exponential circuit lower bounds for the noncommutative permanent and determinant

Yilin Liu

26 November 2026

Abstract

The row-ordered permanent and determinant of an $n \times n$ matrix have exact and border circuit size $2^{n+o(n)}$ over every field. We prove a transfer theorem from planar tensor prices to ordered circuits with unrestricted gate reuse. A two-site uncrossing identity gives the full mean-width price of unsigned partition tensors; exterior multiplication gives the signed bound. Rectangular wire networks convert these prices into local encodings while keeping their auxiliary cost constant at every circuit cut. Regular polygon placements yield the exponent one. We also determine the rank exponent of unsigned balanced partition tensors. Universal matching boards and a central-coordinate determinant construction locate the exact noncentral-variable budget threshold: the full exponent is attainable for both row polynomials under a budget $A(n)$ if and only if $A(n)/n \to \infty$.

hexagon:2611.00994v2

1 Introduction

For free generators x_{ij} over a field K, put

$$w_\pi = x_{1,\pi(1)} \cdots x_{n,\pi(n)}, \qquad \mathrm{Perm}_n = \sum_{\pi \in S_n} w_\pi, \qquad \mathrm{Det}_n = \sum_{\pi \in S_n} \mathrm{sgn}(\pi) w_\pi.$$

A circuit uses binary sums, differences and ordered products, arbitrary constants, and unrestricted gate reuse. Its size $C_K(f)$ is its number of arithmetic gates. For homogeneous f of degree d, the border size $\underline{C}_K(f)$ is the least s for which f lies in the Zariski closure, over $\overline{K}$, of the degree-d output components of size-s circuits. Intermediate degrees and depth are unrestricted. We omit the subscript when the field is fixed. All logarithms are in base two.

Theorem 1.1. *For every $\epsilon > 0$ there is N_ϵ such that, over every field and for $n \geq N_\epsilon$,*

$$\min\{\underline{C}(\mathrm{Perm}_n), \underline{C}(\mathrm{Det}_n)\} \geq 2^{(1-\epsilon)n}.$$

Both polynomials have exact circuit size $O(n2^n)$. Thus each exact and border size is $2^{n+o(n)}$.

The upper bound expands one polynomial for every column subset, adding a last-row entry on the right. Cofactor signs give the determinant bound. Our lower bound measures a complete partition tensor at its planar placement. For an r-run tile on b columns, the transfer coefficient is its Euclidean price divided by πb. The normalized prices of balanced tiles are at least the half-perimeter of a regular unit r-gon up to a vanishing error. Consequently the coefficients

$$\frac{r \sin(\pi/r)}{\pi}$$

are all attainable, and tend to one.

Three constructions establish this transfer. First, a two-site uncrossing inequality makes the one-excitation tensor attain its full mean-width price on every planar polygon. A signed version follows by regularizing exterior multiplication. Second, balanced count sectors retain these prices up to a logarithmic loss and admit radial supporting functionals. Third, a rectangle of auxiliary wires keeps each tensor input in a prescribed flat square. Its cost is identical at every point of the square. We apply the resulting local maps after decomposing the circuit into three intervals, so each gate label pays only the width of its three hubs.

The partition calculation also determines a tensor-rank exponent. For fixed $r \ge 3$, let $B_{r,a}$ have one coordinate for each a-subset of $[ra]$ at each party, with coefficient one exactly on ordered partitions. Writing R and $\widetilde{\mathrm{R}}$ for tensor rank and asymptotic rank, Theorem 3.2 gives

$$\frac{2^{ra}}{\binom{a+r-1}{r-1}^{r}} \;\le\; \widetilde{\mathrm{R}}(B_{r,a}) \;\le\; \mathrm{R}(B_{r,a}) \;\le\; 2^{ra}. \tag{1}$$

For $r = 3$, the exponent relative to the local dimension is therefore $1/h_2(1/3)$, where h_2 is binary entropy. Pratt related this tensor family to the Set Cover Conjecture [10]. The displayed value determines the balanced tripartitioning exponent studied by Björklund, Kaski, Koana and Nederlof [2].

Nisan's rank method gives lower bounds for formulas and branching programs [7]. Hrubeš, Wigderson and Yehudayoff connected shared-gate circuits to repeated words and bilinear sums of squares [4]. The earlier version of this work treated three-interval cuts and sparse commutation graphs [6]. Here we use the isometry-invariant Euclidean price system of OpenAI [8] and its price-extension theorem [9]. Our uncrossing and rectangle constructions yield the full exponential rate. Sra combined the Euclidean price theorem with the earlier cut of Liu to obtain coefficient $1/2$ over every field, and established multiquadratic composition bounds using Khoruzhii's crossing construction [11, 5]. Appendix D gives the bilinear specialization.

Section 6 constructs universal matching boards whose degree is independent of the tensor block size. At maximum noncommutation degree D, the exponent is at least $1 - O((\log D/D)^{2/3})$, with only $O_D(n)$ nonisolated entries. A central change of coordinates gives an explicit determinant upper bound at fixed noncentral density. Together they characterize the budgets attaining exponent one for both row polynomials: $A(n)/n \to \infty$ is necessary and sufficient, and the maximum degree may tend to infinity arbitrarily slowly (Corollary B.2). We retain the degree-three and path constructions, and the matching-algebra simulations.

2 Tensor prices and interval cuts

A spectral point is monotone under partywise linear maps, multiplicative on tensor products and additive on direct sums [12]. We use a compatible planar price H over $\overline{K}$ from [8, Theorems 1.1 and 6.1]. On a finite set of sites S of positive diameter it is $W(S)$ times the logarithm of a normalized spectral point, where $W(S)$ is half the perimeter of its convex hull, or its length for a segment. In particular, a dimension-q wire costs its endpoint distance times $\log q$, and a common q-valued label on S costs $W(S)\log q$.

Prices add on tensor products and decrease under local maps and contractions at a common site. They are nonnegative on nonzero tensors, vanish when all legs coincide, and are jointly convex and positively homogeneous in their placements. Congruent labeled placements have equal prices. Moving a leg of dimension d through distance ℓ changes the price by at most $\ell \log d$. For tensors in fixed output spaces on S,

$$H\left(\sum_{j=1}^{A} T_j\right) \le \max_j H(T_j) + W(S)\log A. \tag{2}$$

The constructions of the summands may use different auxiliary sites. In each fixed format these bounds pass to Zariski closure [8, Lemma 2.1].

Lemma 2.1 (Three-interval cut)**.** *Let a homogeneous word polynomial of degree $N \ge 3$ have a circuit of size s, and let $q_0, \dots, q_N = q_0$ traverse a circle in order, with consecutive angles less than π. Its coefficient tensor is a sum of at most $(N+1)^3$ tensors. Each summand has rank at most $H_0 = O((N+1)^2(s+1))$ across*

$$(i,k], \quad (k,j], \quad (i,j]^c,$$

for a triple whose triangle $q_i q_k q_j$ contains the circle centre.

Proof. Homogenize through degree N. Follow an occurrence of a homogeneous gate in positions $(i, j]$, keeping its left and right context. At a sum expand both children with their signs. At a product split at k. The hull of the consecutive vertices $q_i, \dots, q_j$ is the union of the two child hulls and $q_i q_k q_j$. Stop when that triangle contains the centre. Otherwise follow a child hull containing it and absorb the other factor into the context. A leaf spans a point or one short edge, so every descent stops.

For a fixed triple and a fixed stopped product gate, its two child coefficient vectors are fixed. Sum the signed contexts of all its occurrences into one vector on the complementary positions, whose two pieces retain their order. Thus there are at most H_0 rank-one terms for that triple, including when a gate has arbitrarily many occurrences. There are at most $(N+1)^3$ triples. □

We record the geometric calculation used with this cut. Suppose $p_j = \sum_{h \le j} v_h$ agrees, after translation, with the circle vertices q_j of radius R. For a stopping triple put $Q_1 = q_k, Q_2 = q_j, Q_3 = q_i$. If o is the centre, choose $\lambda_r \ge 0$ with $\sum \lambda_r = 2$ and $\sum \lambda_r (Q_r - o) = 0$, and set cyclically

$$z_r - z_{r+1} = \lambda_r (Q_r - o)/R, \qquad z_1 = 0.$$

Then $W(\{z_1, z_2, z_3\}) = 1$, $\|z_r\| \le 2$, and summation by parts gives

$$\sum_{h=1}^{N} \langle v_h, z_{\mathrm{part}(h)} \rangle = 2R. \tag{3}$$

Degenerate triangles satisfy the same identities. If the prefix vertices have error at most E, the right side decreases by at most $2E$.

Corollary 2.2. *For the repeated word $D_m = \sum_{w \in \{0,1\}^m} z_w z_w$,*

$$\log(1 + \underline{C}(D_m)) \ge \frac{2m}{\pi} - O(\sqrt{m \log(m+1)}).$$

Proof. The $2m$ coefficient positions form m independent binary wires. Place them successively at $u_j = \exp(\pi \mathrm{i} j/m)$, using complex notation only for the real plane. Paired positions are opposite, so the target price is $2m$. Their prefix sums form a regular polygon of radius $R = 1/(2\sin(\pi/(2m)))$. For each cut place its three factors at tz_1, tz_2, tz_3. Their common label costs $t \log H_0$. Terminal transport costs at most $2m - 2tR + 4mt^2$, by $\|u - tz\| \le 1 - t\langle u, z \rangle + t^2 \|z\|^2/2$. Equation (2) adds $O(\log(m+1))$. Compare with the target, apply closure, and take $t = \sqrt{\log(m+1)/m}$. □

3 Partition tensors on a plane

The tensor $T^{\pm}_{\alpha}$, for $\alpha = (a_1, \dots, a_r)$ with $\sum a_i = b$, is supported on ordered partitions $(A_1, \dots, A_r)$ of $[b]$ with $|A_i| = a_i$. Its entries there are 1 for the plus sign and the sign of the concatenated increasing lists for the minus sign. Write $U^{\pm}_b = \sum_{\alpha} T^{\pm}_{\alpha}$ in the full subset spaces. Use the homogeneous isometry-invariant Euclidean price H of [8, Theorem 6.1]. Write $W(S)$ for the mean width of the convex hull of the sites S, equal in the plane to half its perimeter, or to its length for a segment.

Lemma 3.1. *At r planar sites S,*

$$H(U^{+}_b; S) = bW(S), \qquad H(U^{-}_b; S) \ge (b - r)W(S).$$

Proof. Write V_S for the sum of the binary basis tensors with one excited site. Then $U^{+}_b = V_S^{\otimes b}$. Put $h(S) = H(V_S)$. The difference $\bigotimes_{s \in S}(e_0 + te_1) - \bigotimes_{s \in S} e_0$ has linear coefficient V_S, so a binary diagonal degenerates to V_S and $h(S) \le W(S)$. Adding a site at an existing location is locally reversible, since the two new local states are 00 and $10 + 01$. Projecting an added site to its vacuum recovers the old tensor. Convexity in the new position therefore makes h constant

throughout the old hull. We write $h(K)$ for the resulting function of a convex polygon K. Two sites give a wire.

Suppose $I = A \cap B = \{i, j\}$ and $U = A \cup B$. In $V_A \otimes V_B$ denote the excitations by x, y. At $s \in I$ send the local basis $1, x, y, xy$ to $1, tu + \epsilon_s t^{-1} v, t^{-1} v, uv$, respectively, where $\epsilon_i = 1,\ \epsilon_j = -1$. At each unshared site send the excitation to tu and preserve the vacuum. The degree -2 coefficient is $(\epsilon_i + \epsilon_j) v_i v_j = 0$. At degree zero, u occurs once on U and v once on I. Given the location s of v, the coefficient of u_a is 1 for $a \in A$ and ϵ_s for $a \in B \setminus A$. Thus the specialization has two blocks locally equivalent to V_U. Their bit is visible at both i, j, since their v-flags are complementary. Copying this bit to the other sites costs at most $W(U) - |i - j|$ by telescoping the incremental copying bound [8, Proposition 5.2]. The resulting full direct sum costs $h(U) + W(U)$. Hence

$$h(A) + h(B) \geq h(U) + |i - j|. \tag{4}$$

Every rectangle has full width price. For the endpoints S of a segment and a perpendicular vector v, put $S_j = S + jv$ and $R = \operatorname{conv}(S_0 \cup S_1)$. Successively append $S_{j-1} \cup S_j$ to the previous union, with intersection S_{j-1}. By (4), $m|v| \leq h(\bigcup_{j=0}^{m} S_j) \leq h(S) + m(h(R) - h(S))$. Division by m gives $h(R) \geq h(S) + |v| = W(R)$. If a line cuts a full-price polygon R into P, Q along a chord $[i, j]$, their vertex sets augmented by i, j meet at these two endpoints. Equation (4) gives $h(P) + h(Q) \geq W(R) + |i - j| = W(P) + W(Q)$. Both pieces have full price. Cutting an enclosing rectangle by the sides of any polygon proves the unsigned formula.

For the signed tensor take local spaces V_s with basis $e_{s,1}, \ldots, e_{s,b}$. In site order, U_b^- is the expansion of $\bigwedge_{j=1}^{b} (\sum_s e_{s,j})$ in $\bigwedge(\bigoplus_s V_s) \simeq \bigotimes_s \bigwedge V_s$. Combining two column families by the exterior product introduces the sign $(-1)^{\sum_{s>t} p_s q_t}$, for local degree parities p_s, q_t in the two families. Fixing the vector p makes the sign local. Each of the ordinary and exterior products is consequently a sum of at most 2^r local images of the other. Their prices differ by at most $c = rW(S)$. The same parity argument bounds the cost of reordering sites independently of b. Thus $H_b = H(U_b^-; S)$ obeys $|H_{b+d} - H_b - H_d| \leq c$. Fekete's lemma applied to $H_b \pm c$ gives $f(S) = \lim_b H_b / b$, and taking multiples of b yields $|H_b - b f(S)| \leq c$.

For nonempty $I = A \cap B$, put $U = A \cup B$. In one column let $X = \sum_{s \in A} x_s$, $Y = \sum_{s \in B} y_s$, $u_D = \sum_{s \in D} u_s$, and $v_I = \sum_{s \in I} v_s$. Send x_s to tu_s on $A \setminus B$ and $tu_s + t^{-1} v_s$ on I. Send y_s to $-tu_s$ on $B \setminus A$ and $t^{-1} v_s$ on I. Taking exterior powers locally sends

$$X \wedge Y \ \longmapsto\ u_U \wedge v_I - t^2 u_A \wedge u_{B \setminus A},$$

since $v_I \wedge v_I = 0$ over every field. Apply this map in all b columns and specialize at zero. The comparison of ordinary and exterior products costs $O_{|U|}(W(U))$, so division by b gives $f(A) + f(B) \geq f(U) + f(I)$.

The function f is monotone by vacuum projection. Colocated sites merge reversibly using the exterior power of $e_j \mapsto e_j' + e_j''$, after reordering them to be adjacent. Joint convexity and homogeneity pass to the limit. Thus f depends only on the hull as above. For Minkowski addition, index sites by pairs of vertices (x_i, y_j) of K, L, placing each pair first at x_i, then at y_j. Midpoint convexity gives $f(K + L) \leq f(K) + f(L)$. On a segment U_b^- is a dimension-2^b wire up to local signs. The preceding translate argument now applies to the vertex set of any polygon K and any v, since the overlap contains the preceding translate of that set. It gives $f(K + [0, v]) \geq f(K) + |v|$, with equality by Minkowski subadditivity. The difference body $K - K$ is a centrally symmetric polygon and hence a sum of segments. Therefore $2f(K) \geq f(K - K) = W(K - K) = 2W(K)$ by isometry invariance. The finite error bound for H_b proves the lemma. □

Theorem 3.2 (Balanced partition tensors)**.** *Let $r \geq 3$, let $e = (e_1, \ldots, e_r)$ be a regular unit r-gon, and put $w = r \sin(\pi/r)$, $N_b = \binom{b+r-1}{r-1}$, $B^{\pm} = T^{\pm}_{b,\ldots,b}$. Set $c_+ = 0$, $c_- = r$. Then*

$$H(B^{\pm}; e) \geq rw(b - c_{\pm} - \log N_b), \qquad H(B^{\pm}; x) \geq \frac{H(B^{\pm}; e)}{r} \sum_i \langle e_i, x_i \rangle$$

for all planar placements x. In particular, over every field,

$$2^{rb} N_b^{-r} \le \widetilde{\mathrm{R}}(B^+) \le \mathrm{R}(B^+) \le 2^{rb}.$$

Proof. There are N_b sectors in $U_b^{\pm}$. The sum bound and Lemma 3.1 give one with price at least $w(b - c_{\pm} - \log N_b)$. Its r cyclic permutations have equal price by isometry invariance, since at fixed counts a site permutation changes the shuffle sign by a constant. Their product is a restriction of $B^{\pm}$ on r consecutive b-column blocks, prescribing the cyclic counts block by block. Each site has total count b and all interblock sign contributions are constant. This proves the first bound. Average a subgradient of $x \mapsto H(B^{\pm}; x)$ at e over the dihedral symmetries. Each reflection fixing e_i forces its component to be radial, and rotations give the same radial coefficient at every site. Homogeneity identifies this coefficient as $H(B^{\pm}; e)/r$, proving the second bound.

On these sites the price is w times the logarithm of a normalized spectral point, so the unsigned lower bound follows. For the ordinary rank upper bound, sum over $S \subseteq [rb]$ the products of the r local sums $\sum_{A \subseteq S,\, |A|=b} x_{i,A}$, weighted by $(-1)^{rb-|S|}$. Inclusion–exclusion selects exactly the partitions. □

4 Holding tensor sites in a square

Let $R = [-3, 3]^2$ and let l_-, l_+ and u_-, u_+ be its lower and upper corners. A wire has price equal to its endpoint distance. We work over the algebraic closure. Six sets of wires constrain a tensor site to R and can be placed at the same cost at every point of it (Figure 1).

Lemma 4.1. *In a complete CAT(0) space containing R isometrically, put*

$$h = d(l_-, a) + d(a, l_+) + d(u_-, b) + d(b, u_+) - 12, \quad v = d(a, y) + d(y, b) - 6, \quad P_Q = Q^4 h + Q v.$$

For each $y \in R$ there are $a, b \in R$ with $h = v = 0$. For arbitrary a, b, y and integers $Q \ge 1$ one has $P_Q \ge -3/Q^2$. If $P_Q \le C$ for a fixed $C \ge 0$, then $d(y, R) = O_C(Q^{-1/2})$.

Proof. For $y \in R$ take a, b directly below and above it on the horizontal sides. For the converse, suppose a point's distances to the ends of a segment of length D sum to $D + e$. Divide the segment in the ratio of these two distances. The CAT(0) squared-distance inequality at the dividing point bounds its distance from the given point by $\frac{1}{2}\sqrt{2De + e^2}$. Thus there are a_0, b_0 on the lower and upper sides with $d(a, a_0) + d(b, b_0) \le \sqrt{6h} + h/2$. Since $d(a_0, b_0) \ge 6$,

$$P_Q \ge (Q^4 - Q/2)h - Q\sqrt{6h} \ge -3/Q^2.$$

The same quadratic in $\sqrt{h}$ gives $h = O_C(Q^{-4})$ when $P_Q \le C$. It follows that $d(a, a_0) + d(b, b_0) = O_C(Q^{-2})$ and $v \le C/Q$. The distance sum from y to a_0, b_0 exceeds their distance by $O_C(Q^{-1})$. The segment estimate on $[a_0, b_0] \subset R$ proves the claim. □

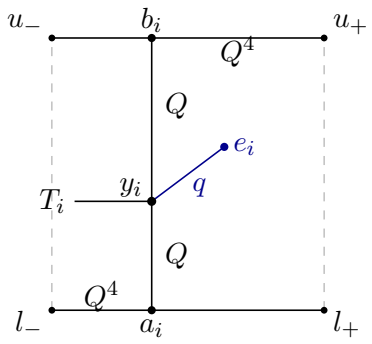


Figure 1: Each input has four horizontal bundles of weight Q^4, two vertical bundles of weight Q, and a signal bundle of weight q. The auxiliary cost is $12Q^4 + 6Q$ throughout the square.

Proposition 4.2. *Let T be a nonzero r-tensor, e_i unit vectors in the plane, and q a positive integer. Suppose the flat price H satisfies*

$$H(T; x_1, \dots, x_r) \geq q \sum_i \langle e_i, x_i \rangle \qquad (x_i \in R).$$

For each $\epsilon > 0$ there exist integers $Q, k \geq 1$ and a local linear encoding of $T^{\otimes k}$ with output J at the e_i and the four corners, such that

$$H(J) \geq k(B + rq - \epsilon), \qquad B = r(12Q^4 + 6Q).$$

The map on party i is given by one functional at each of three sites y_i, a_i, b_i. It uses kq binary wires from y_i to e_i, kQ from y_i to each of a_i, b_i, and kQ^4 from a_i to each lower corner and from b_i to each upper corner. Its input is the ith leg at y_i.

Proof. Fix the price on the finite terminal set of corners and e_i. For a tensor U with extra sites, the price-extension theorem identifies $\inf_{H'} H'(U)$ over extensions of this price with $\sup_{k,\ell} H(\ell(U^{\otimes k}))/k$, where ℓ consists of one functional at each extra site. For the trivial valuation, the natural-log prices in [9, Theorem 5.1] are $(\ln 2)H$ with the same wire distances. Every extension embeds in a complete CAT(0) spectral space. The closed convex hull of the base has its original flat distances and tensor prices, by geodesic determination [8, Theorem 4.3], [9, Lemma 5.2]. Thus each extension contains the same priced square R.

Form U by taking the tensor product of T at the y_i with the wires specified at $k = 1$. Its extended price less B is

$$H'(T; y) + q \sum_i d(y_i, e_i) + \sum_i P_Q(a_i, b_i, y_i).$$

The placement $y_i = 0$ attains rq. Whenever the displayed expression is at most $rq + 1$, each individual penalty is bounded above independently of Q, since the first two terms are nonnegative and the remaining penalties are at least $-3/Q^2$. By Lemma 4.1, each y_i is within $O(Q^{-1/2})$ of some $x_i \in R$. Transporting a leg of dimension d_i costs at most $d(y_i, x_i) \log_2 d_i$. The signal distances also change by $o(1)$. The displayed price is therefore at least $q \sum_i (\langle e_i, x_i \rangle + \|e_i - x_i\|) - o(1) \geq rq - o(1)$. Choose Q and then a power and functionals in the supremum to obtain the asserted bound. Each functional may couple all k copies at its site. The three sites belonging to a party meet only that party's input and its own wires, so their contractions define the stated local map. □

5 From a planar price to a circuit lower bound

Let $T = T^{\pm}_{a,\dots,a}$ be the ordered-partition tensor with r parties, $a \geq 1$ elements in each part and $b = ra$ columns. Write P_n for the corresponding row permanent or determinant. Place the parties at the ordered vertices $e_1, \dots, e_r$ of a regular r-gon on the unit circle. We use the isometry-invariant planar price H over the algebraic closure of the ground field.

Theorem 5.1 (Planar price transfer)**.** *For every $r \geq 3$ and $a \geq 1$,*

$$\liminf_{n \to \infty} \frac{\log_2(1 + \underline{C}(P_n))}{n} \;\geq\; \frac{H(T; e_1, \dots, e_r)}{\pi b}.$$

Proof. By Theorem 3.2,

$$H(T; x) \geq c \sum_i \langle e_i, x_i \rangle, \qquad c = H(T; e)/r.$$

We may assume $c > 0$. For an integer k large enough that $q = \lfloor kc \rfloor \geq 1$, the tensor $T^{\otimes k}$ satisfies the same inequality with coefficient q: use nonnegativity of the price when the sum on the right is negative.

Apply Proposition 4.2 to $T^{\otimes k}$, the square $[-3,3]^2$ and the integer q at every party. For any $\epsilon > 0$ it gives a power m, fixed local functionals at the openings, and an output tensor J with

$$H(J) \geq m(B + rq - \epsilon). \tag{5}$$

Here B is the total nominal cost of the r rectangle constructions before taking the mth power. The input to J is $T^{\otimes km}$. Keep the factorization of its local maps from that proposition. If an input is supplied at any point p of the square, the other two openings of that party can be placed vertically below and above p on its horizontal sides. All rectangle wires then have exactly their nominal cost, regardless of p.

Take $L = \lfloor n/(bkm) \rfloor$ disjoint groups of km copies of T, padding the remaining rows and columns by an identity. Rotate group ℓ and its rectangle by angle $2\pi\ell/(rL)$, for $0 \leq \ell < L$. Read all $N = rL$ party visits in angular order. At each visit assign a rows to each of the km corresponding column blocks. Only the b columns of the designated block are allowed at those rows, and each local subset must occur in increasing order. A finite-state matrix evaluation stores the km subsets of that visit, emits one letter for their tuple at its last row, and resets for the next visit. For fixed r, a, k, m the storage per visit is fixed. Row-dependent transition matrices can enforce its block and row assignments. The input permutation uses every column once, so its surviving words are exactly the independent ordered partitions in all the blocks. For the determinant, rearranging the fixed-length lists contributes only a constant shuffle sign. After that sign is removed, the coefficient tensor of the degree-N word on the visit letters is $(T^{\otimes km})^{\otimes L}$ in the indicated order. The alphabet for a visit may be indexed by that visit; projecting a coefficient position to its prescribed alphabet is a local map. Extract the degree-n source component first, evaluate its matrices, take the specified entry, and then extract degree N. These operations have polynomial overhead in n and are linear on the prescribed source coefficient space. Thus they also give the required border simulation in fixed formats.

We apply the three-interval circuit cut to this word before applying the local maps producing J. Each visit is now one coefficient position. Its unit signal direction u_v runs through the N regular directions in order. The prefix sums of the vectors $mq\,u_v$ are the vertices of a regular polygon of circumradius

$$R = \frac{mq}{2\sin(\pi/N)}.$$

For a size-s approximating circuit, the homogenized simulation has $H_0 \leq \mathrm{poly}(n)(s+1)$ gates. Stop at a triangle containing the polygon centre, as in the three-interval cut. There are at most $(N+1)^3$ stopping triples. For each triple, summing the contexts of all occurrences of a stopped gate gives a tensor of rank at most H_0 on the three cyclic intervals. The circle calculation supplies hub vectors z_1, z_2, z_3 with norm at most two and common-label width one, such that

$$mq \sum_v \langle u_v, z_{\mathrm{part}(v)} \rangle = 2R.$$

It applies to these exact prefix sums regardless of the size of their common step.

Place the three factors of a rank decomposition at tz_1, tz_2, tz_3, where $0 < t < 1$. The gate label costs at most $t \log_2 H_0$. Every visit lies wholly in one interval. Its input functional can therefore consume its entire leg at that interval's hub, inside its rotated square. Place its two auxiliary openings as above. The total rectangle cost is LmB in every summand. The signal wires together cost at most

$$Lmrq - 2tR + 2Lmrq\,t^2$$

by $\|u - tz\| \leq 1 - t\langle u, z\rangle + t^2\|z\|^2/2$ for unit u. All physical output sites of all groups lie in a fixed disk. Taking the ordinary sum over the triples therefore adds only $O(\log(N+1))$ to the largest price bound. This upper bound passes to border limits after the same fixed local maps. The intended output is a product of L rotated copies of J and has price $LH(J)$. Comparing with (5) gives

$$\log_2 H_0 \;\geq\; 2R - Lm(\epsilon/t + 2rqt) - O(\log(N+1)/t).$$

Fix k and t, then fix the hinge precision ϵ and the resulting network and power m, before letting n tend to infinity. The normalized lower bound is $[rq/\pi - \epsilon/t - 2rqt]/(kb)$. Send ϵ to zero, then t to zero, and finally k to infinity. □

Proof of Theorem 1.1. Theorem 3.2 gives, for either sign and fixed r,

$$H(T^{\pm}_{a,\dots,a}; e) \geq ra\, r \sin(\pi/r)\,[1 - O_r(\log(a+1)/a)].$$

Apply Theorem 5.1 and let a tend to infinity. The normalized logarithmic border size is at least $r\sin(\pi/r)/\pi$. Letting r tend to infinity gives one. The column-subset expansion gives $C(P_n) = O(n2^n)$. Appendix A proves that the lower threshold is uniform in the field. □

6 Sparse commutation relations

In the partially commutative algebra $A_G(K)$, distinct generators commute exactly when they are nonadjacent in G. Its trace basis is the set of words modulo adjacent commuting swaps. We take border closure in this basis. A row board is a set of allowed matrix cells; its restriction sets all other entries to zero.

Theorem 6.1. *There is an absolute C such that, for each sufficiently large integer D, a single finite board template determines graphs $G_{D,n}$ on the n^2 matrix entries, of maximum degree at most D and with at most Dn nonisolated vertices. For every field and both row polynomials, the full polynomial and its board restriction in $A_{G_{D,n}}(K)$ satisfy*

$$\liminf_{n\to\infty} \frac{\log(1 + \underline{C}(P_n))}{n} \geq 1 - C\left(\frac{\log D}{D}\right)^{2/3}.$$

The template depends only on D, and the bound is uniform in the field. The restricted polynomials have $O_D(n)$-size formulas after abelianization.

Lemma 6.2. *Let $b = ra$, $r \geq 3$ and $0 < \gamma \leq 1/2$. Put $e = 2\lceil \gamma a/2 \rceil$ and $m = a + e$, with $a \geq 2/\gamma$. There is a bipartite board on m rows, b ordinary columns and e private columns, with at most $O(r\log(2r)/\gamma)$ cells in each row, such that the rows admit a perfect matching to every choice of a ordinary columns together with all e private columns.*

Proof. We first construct a bipartite graph on $M \geq N$ left and right vertices in which every left set of size $t \leq (1-\eta)N$ has at least t neighbors, for $0 < \eta < 1$. Pad the left to $u = \lceil M/N \rceil$ groups of N vertices. From each group choose d independent uniform perfect matchings to the right, where

$$d = 2 + \left\lceil \frac{\ln(4\mathrm{e}^2 u)}{-\ln(1-\eta)} \right\rceil.$$

For fixed sets of t left and $t-1$ right vertices, the probability that all neighbors of the former lie in the latter is at most $(t/N)^{dt}$ by sampling without replacement. There are at most $(\mathrm{e}^2 uN^2/t^2)^t$ such pairs, since $\binom{N}{t-1} \leq (\mathrm{e}N/t)^t$ by monotonicity of $x\ln(\mathrm{e}N/x)$ on $[0, N]$. The failure probability is at most $\sum_{t\geq 1} 4^{-t} < 1$. Remove the padding and repeated edges. The left and right degrees are at most d and ud, respectively.

Take the union of two such graphs. The first goes from the m rows to the e private columns with $\eta = 1/2$, at a cost of $O(\log(2/\gamma))$ cells per row. The second goes from all $b+e$ columns to the m rows with $\eta = e/(2m)$, at a cost of $O(r\log(2r)/\gamma)$ cells per row. Fix a chosen ordinary columns. Any collection of at most $e/2$ rows has at least as many private neighbors. If $y > e/2$ rows had fewer than y neighbors among the m chosen and private columns, at least $m - y + 1$ of those columns would have all neighbors in the complementary $m - y$ rows. Since $m - y + 1 \leq m - e/2$, the second expansion property rules this out. Hall's theorem gives the matching. □

Proof of Theorem 6.1. Choose r of order $(D/\log D)^{1/3}$ with a sufficiently small absolute multiplier and put $\gamma = r^{-2}$. Next fix one $a = a(D) \geq 2/\gamma$ large enough that

$$r + \log \binom{a+r-1}{r-1} \leq \gamma a.$$

Use the integers b, e, m of Lemma 6.2. A base block has $B = rm$ columns, consisting of b ordinary columns and e private columns for each party. Give party i a copy of the lemma's m rows, connected to the ordinary columns and its own private columns. All these choices are made before choosing a field or a sign.

For matrix order n put $J = \lfloor n/B \rfloor$. Use J disjoint column blocks. In party order, visit these blocks in increasing order with m rows per visit. Pad the final $n - JB$ rows with their diagonal cells. This is the row board. Join every allowed cell of row v to every allowed cell of row $v+1$, with no further edges. If s is the maximum number of allowed cells in one row (including padding), then the maximum graph degree is at most $2s \leq D$ by our choice of r. There are at most $sn \leq Dn$ nonisolated generators.

Inside any base block a supported permutation must use all e private columns of each party there, hence exactly a ordinary columns at that party. For every a-subset choose the lexicographically first perfect matching supplied by the lemma. The transducer stores its m observed column indices and accepts only this representative. Thus every ordered partition of the ordinary columns occurs once. For the determinant, order the ordinary columns first and multiply the local transition by the sign of sorting the observed columns. The private columns of this party pass a fixed number of ordinary columns at later parties. Private sets are fixed, and the base-block visit sequence fixes all inversions between blocks. The remaining sign is the ordered partition sign up to one fixed factor. The permanent uses unit coefficients.

Write $T = T^{\pm}_{a,\dots,a}$. For the finite power h required by the encoder in the proof of Theorem 5.1, group h consecutive base blocks. At a party visit to a group, store its h ordinary subsets and emit a single letter at the end. The number of local states depends only on B and h, for example it is bounded by $(B+1)^{mh+1}$ for storing a padded list. The row clock uses polynomially many states in n. Freeze the last $J \bmod h$ blocks to one fixed supported permutation per block across all parties; accept the padding diagonals and emit no letters for either. Their sign is fixed. The $\lfloor J/h \rfloor$ groups now have exactly the macro-word tensor $(T^{\otimes h})^{\otimes \lfloor J/h \rfloor}$ in the required angular visit order. Changing h changes only the transition matrices and the frozen suffix within each party.

If M_{vj} is the transition matrix at an allowed cell, map its variable to $E_{v,v+1} \otimes M_{vj}$ and the forbidden ones to zero. Every required commuting pair has zero product in both orders. The $(1, n+1)$ block recovers the ordered evaluation. Extracting the source degree, evaluating fixed matrices, taking this block, and extracting the desired output degree preserve border closure with polynomial circuit overhead, as in Theorem 5.1. The same argument applies after restricting to the board.

Each base block consumes rm rows. The transfer proof therefore gives normalized logarithmic border size at least $H(T; e_1, \dots, e_r)/(\pi rm)$; the rows outside the complete groups number less than $hB + B$ for fixed h. By Theorem 3.2 this ratio is at least

$$\frac{r \sin(\pi/r)}{\pi} \frac{a}{a+e} (1-\gamma) \geq 1 - O(r^{-2} + \gamma).$$

Here $e \leq 2\gamma a$. This proves the bound. The template stays fixed in the ultraproduct argument of Appendix A, which gives a threshold uniform in the field. Finally a commutative row permutation makes the restricted matrix block diagonal, with $B \times B$ blocks and singleton padding. Its permanent or determinant is their product up to a fixed determinant sign. The fixed block size gives an $O_D(n)$ formula. □

A Uniformity in the field

One can take the remainder $o(n)$ uniformly over fields. Otherwise choose algebraically closed fields K_j and orders $n_j \to \infty$ violating the bound at some fixed coefficient $c < 1$. Pass to a subsequence on which the same sign violates it. Select the isometry-invariant homogeneous Euclidean prices H_j over these fields. Fix a nonprincipal ultrafilter and let $K = \prod_{\mathcal{U}} K_j$. This is an algebraically closed field. For a nonzero tensor $T = [T_j]$ of fixed finite format over K and a fixed real tuple x of plane sites, set

$$H(T; x) = \lim_{\mathcal{U}} H_j(T_j; x).$$

The tensor T_j is nonzero almost everywhere. Transport to one site gives a common finite bound. A different choice of coordinate representatives agrees almost everywhere. On every finite site set of positive width D the function $2^{H/D}$ is a spectral point: the restriction, direct-sum, product, zero and unit laws follow coordinatewise, since each linear map over K likewise has coordinate representatives. The configurations have the same wire distances and shared-bit width on every field. Compatibility, Euclidean invariance, homogeneity and convexity also pass to this limit. Thus H is an admissible flat price for the argument above.

Choose the fixed parameters r, b, k, t, Q, m and the finitely many coordinates of a closing network over K with strict margin over c. Representatives of these coordinates specify networks of the same dimensions over K_j. The output is a polynomial in those coordinates with integer coefficients, so its H_j-price tends to its H-price along the ultrafilter. The same strict margin therefore holds almost everywhere after a small fixed weakening. All remaining bounds on the row simulation, degree homogenization, number of circle cuts and output-site widths are uniform once these finite formats are fixed. Hence the circle inequality at n_j gives a coefficient strictly above c almost everywhere, a contradiction.

B Determinant circuits from central coordinates

A central change of coordinates converts a linear budget of noncentral entries into disjoint column constraints.

Theorem B.1. *Let $D \geq 0$ be an integer. If at most Dn entries of an $n \times n$ matrix of generators are noncentral, its row determinant has an exact circuit of size*

$$n^{O(1)}(2^{D+1} - 1)^{n/(D+1)}$$

over every field. The same bound holds for the row permanent in characteristic two.

Proof. Write X_i for the ith row, and let A be the algebra generated by the entries. The subring Z generated by K and the central entries lies in the center of A. Put $q = D + 1$ and $s = \lfloor n/q \rfloor$. There are at most $\lfloor Dn/q \rfloor$ rows containing at least q noncentral entries, so at least $n - \lfloor Dn/q \rfloor \geq s$ rows contain at most $q - 1$ such entries. Select s of them and assign to each a disjoint set $I_i \subseteq [n]$ of q coordinate positions.

For each selected row write $X_i = \sum_{u \in I_i} x_{iu} v_{iu}$, with $x_{iu} \in A$ and $v_{iu} \in Z^n$. Use one coefficient 1 for the row consisting of its central entries and zero elsewhere. For each noncentral entry use that entry as coefficient and its standard coordinate vector as v_{iu}. Pad with zero terms. With a central indeterminate t, replace v_{iu} by

$$v_{iu}(t) = (1 - t)e_u + t v_{iu}.$$

This gives the deformed matrix $X(t)$; the other rows are unchanged. Let row u of a central matrix $B(t)$ be $v_{iu}(t)$ when $u \in I_i$, and e_u at each unassigned coordinate. Thus $B(0) = I$. The entries of $X(t)$ are affine in t, so $\mathrm{Det}(X(t))$ is a polynomial in t of degree at most n and $\mathrm{Det}(X(1)) = \mathrm{Det}(X)$.

Compute over $A[t]/(t^{n+1})$. The matrix B over the central subring $Z[t]/(t^{n+1})$ is invertible, with the explicit formula

$$B^{-1} = \sum_{j=0}^{n} (I - B)^j.$$

To compute $\det B$, ordinary Gaussian elimination in this central ring needs only pivots u with constant term one, since $B(0) = I$. Replace each reciprocal by $\sum_{j=0}^{n}(1-u)^j$ and multiply the pivots. Put $Y = X(t)B^{-1}$. The selected row $Y_i = \sum_{u \in I_i} x_{iu} e_u$ is supported in I_i. For every matrix U over an associative algebra and every central matrix C, expansion in the given row order gives $\mathrm{Det}(UC) = \mathrm{Det}(U)\det(C)$. Indeed, the central coefficient of $U_{1j_1} \cdots U_{nj_n}$ in the expansion is the determinant of C with ordered rows $j_1, \ldots, j_n$. It vanishes for a repetition, and otherwise is the permutation sign times $\det(C)$. Applying the identity to $X(t) = YB$ gives

$$\mathrm{Det}(X(t)) = \mathrm{Det}(Y)\det(B).$$

Use the ordinary column-subset recurrence for $\mathrm{Det}(Y)$, appending the next entry on the right with its inversion sign. At each layer retain only subsets S for which $S \cap I_i \neq \varnothing$ for every selected row already processed and $I_i \setminus S \neq \varnothing$ for each selected row still to come. Every nonzero permutation term meets these requirements at every layer, since row i must use a column of I_i. Hence this pruning preserves the answer. In each disjoint q-set at least one of its 2^q membership patterns is forbidden, regardless of the layer. Each layer therefore has at most

$$(2^q - 1)^s 2^{n-qs} = 2^n (1 - 2^{-q})^s \leq (1 - 2^{-q})^{-1} (2^q - 1)^{n/q}$$

states. There are $n + 1$ layers and only polynomially many operations per state.

Finally, arithmetic modulo t^{n+1} is simulated by coefficient arrays of length $n + 1$, using addition and truncated convolution. Multiplication by $\det B$ recovers the coefficients of the original $\mathrm{Det}(X(t))$ through degree n, which are all its coefficients. Adding them evaluates this polynomial at $t = 1$ exactly over any field. All matrix and coefficient-array operations have polynomial overhead and require no division. □

For unequal row budgets, let d_i be the number of noncentral entries in row i. For any set I with $\sum_{i \in I}(d_i + 1) \leq n$, assigning $d_i + 1$ positions to row i gives

$$C(\mathrm{Det}(X)) \leq n^{O(1)} 2^n \prod_{i \in I} (1 - 2^{-d_i - 1}).$$

In particular, at most $k \leq n$ noncentral entries give an $n^{O(1)}2^k$-size circuit by selecting the fully central rows.

Corollary B.2. *Let $A(n)$ be integer budgets with $0 \leq A(n) \leq n^2$. There is a single graph family on the matrix entries, with at most $A(n)$ noncentral generators, for which both row polynomials have exact and border size $2^{n+o(n)}$ uniformly over all fields, if and only if*

$$A(n)/n \longrightarrow \infty.$$

In this case its maximum degree may also be bounded by any prescribed positive integer sequence $D(n) \to \infty$.

Proof. If $A(n)/n$ does not tend to infinity, some infinite subsequence has $A(n) \leq Dn$ for a fixed integer D. Theorem B.1 gives a determinant exponent strictly below one on that subsequence. For the converse, choose fixed degrees D_j in Theorem 6.1 whose limiting lower exponents exceed $1 - 1/(2j)$. Choose strictly increasing thresholds N_j so that, for every $n \geq N_j$, both lower bounds exceed $2^{(1-1/j)n}$ over every field, $D_j n \leq A(n)$, and $D_j \leq D(n)$. At order n use the template for the largest j with $N_j \leq n$, and the edgeless graph below N_1. This index tends to infinity. The usual $O(n2^n)$ exact upper bound completes the claim. □

C Degree-three graphs and paths

An $a \times a$ matrix evaluation replaces each arithmetic gate by $O(a^3)$ scalar gates. Degree extraction before substitution and taking a matrix entry afterwards preserve border closure. The following constructions use matrices over the free algebra on z_0, z_1.

Theorem C.1. *For every n there is a graph of maximum degree three on the n^2 matrix entries, with $O(n)$ nonisolated vertices, such that over every field*

$$\min\{\log(1+\underline{C}(\mathrm{Perm}_n)), \log(1+\underline{C}(\mathrm{Det}_n))\} \geq \frac{n}{2\pi} - O(\sqrt{n\log(n+1)}).$$

The same estimates hold for multilinear restrictions of the two polynomials whose abelianizations have linear-size formulas.

Proof. Use the edgeless graph for $n = 1$. For $n \geq 2$ set $k = \lfloor (n+2)/4 \rfloor$. In order, visit the labels $1, \ldots, k, 1, \ldots, k$. Group consecutive visits into pairs, and place two singleton rows between successive pairs. The core has $n_0 = 4k - 2$ rows. Let σ exchange the two rows visiting each label. A visit row i has two active entries, its diagonal for bit zero and column $\sigma(i)$ for bit one. A singleton has only its diagonal. Join every active entry of one row to every active entry of the next, with no other edges. The successive row sizes repeat $2, 2, 1, 1$, so the two rows neighboring any row have total size at most three.

Every permutation supported on this board chooses a disjoint transposition or two diagonals for each label. In visit row i substitute $E_{i,i+1}z_b$ at bit b. At a singleton use $E_{i,i+1}$. Put identities on the padding diagonals and zero elsewhere. In the determinant negate the bit-one entry at the first visit to each label. This cancels exactly the sign of the chosen transpositions. Among core rows, images in the same row or in rows at distance at least two have zero products in both orders. The padding identities commute with every image. The $(1, n_0 + 1)$ entry is D_k for both signs. Corollary 2.2 proves the estimate. On the active board the abelianization is the product of the inserted diagonals and $x_{ii}x_{jj} \pm x_{ij}x_{ji}$ over the pairs of visits. □

Theorem C.2. *For every even m the path on m generators supports a homogeneous polynomial F_m of individual degree two with*

$$\log(1+\underline{C}(F_m)) \geq m/\pi - O(\sqrt{m\log(m+1)}).$$

Both its abelianization and its commutative coefficient table have linear-size formulas.

Proof. Let $m = 2k$, give consecutive $z_1, \ldots, z_m$ the edges of the path, and take

$$F_m = \sum_{b\in\{0,1\}^k} \left(\prod_{j=1}^{k} z_j^{1+b_j}\right)\left(\prod_{j=1}^{k} z_{k+j}^{2-b_j}\right),$$

with products in increasing order. Each word has degree $3k$. Mark an exponent of two at position j by Y_j. The coefficient table is $\prod_{j=1}^{k}(Y_j + Y_{k+j})$. Its abelianization is $(\prod_{h=1}^{m} z_h)\prod_{j=1}^{k}(z_j + z_{k+j})$.

Use distinct auxiliary addresses u_h outside $[m+1]$, and substitute

$$Z_h = E_{h,h+1}M_{h,0} + E_{h,u_h}M_{h,1} + E_{u_h,h+1}.$$

The direct step of Z_h emits $M_{h,0}$ and its square emits $M_{h,1}$. A single step ending at u_h cannot reach the final address. Nonconsecutive matrices have disjoint address supports. Take $(M_{h,0}, M_{h,1}) = (z_0, z_1)$ for $h \leq k$, and (z_1, z_0) afterwards. The $(1, m+1)$ entry is D_k. Apply Corollary 2.2. □

The same construction gives a path family indexed by permutations. Take two visit rows for each of k bits in repeated-word order and keep their two diagonal and two exchange entries. Give the resulting $4k$ cells path generators in row-major order. For each supported permutation write

the ordered monomial with exponent two on its chosen cells and one elsewhere. In the preceding evaluation use $M_{h,0} = 1$ at every cell, and $M_{h,1} = z_0$ on diagonals or z_1 on exchanges. The sum extracts D_k and hence has log border size at least $2k/\pi - O(\sqrt{k \log(k+1)})$. Its coefficient table is the product of one quadratic binomial per bit. If instead all cells of an $N \times N$ matrix form the path and all permutations are used, set $M_{h,1} = x_{ij}$ at cell (i, j). The extracted entry is Perm_N, so Theorem 1.1 gives the corresponding bound in matrix order, up to a polynomial simulation factor.

D Bilinear sums of squares

Sra treats the corresponding problem for any fixed number of factors [11]. Here let $R(q)$ be the fewest complex bilinear forms whose squares sum to $(\sum_{i=1}^{q} a_i^2)(\sum_{j=1}^{q} b_j^2)$.

Theorem D.1. *One has* $R(q) = q^{\sqrt{2}+o(1)}$.

Proof. Polarizing a representation with R forms in both arguments gives

$$2\langle a, c\rangle\langle b, d\rangle = \sum_{j=1}^{R} \bigl(B_j(a, b)B_j(c, d) + B_j(a, d)B_j(c, b)\bigr).$$

For $q = 2^r$, place the four vector indices at $e_1, e_2, -e_1, -e_2$ in that order. The left side has price $4r$. Each of the two sums on the right has rank at most R across its cut of the square into two adjacent pairs. For that cut let v be the normalized sum of the unit vectors in the first pair and place the two factors at $tv, -tv$. The four terminal legs cost at most $4r - 2\sqrt{2}\, rt + 2rt^2$ in total. The label costs $2t \log R$, and summing the two tensors costs at most $2\sqrt{2}$. Thus $\log R \geq \sqrt{2}\, r - O(\sqrt{r})$ on taking $t = r^{-1/2}$. Restriction to the preceding power of two gives $R(q) \geq q^{\sqrt{2}-o(1)}$ for all q. Khoruzhii's construction [5, Corollary 3] supplies the upper bound. □

E Disjoint noncommuting pairs

Write $L_a(u) = au$ and $R_b(u) = ub$ for left and right multiplication in an algebra. Suppose we assign commuting linear operators to ordinary variables. An arithmetic formula in these variables can be evaluated on an argument u by passing that argument to both branches of a sum. At a product apply the right subtree first, then the left. Each node is visited once. The resulting arithmetic circuit may reuse its argument at any sum.

Proposition E.1. *Partition the noncentral generators into two internally commuting sets L, R. Suppose every word in F can be ordered, using the defining commutations, with all letters of L before those of R. An ordinary formula of size s for F^{ab} gives a circuit of size $O(s + 1)$ for F in A_G.*

Proof. Assign L_a to $a \in L$, R_b to $b \in R$ and multiplication to central letters. The operators commute since left and right multiplication commute and each class is internally commuting. Applied to the unit, every monomial has its prescribed trace. Evaluating the formula on the unit proves the bound. □

Corollary E.2. *For any sequence of matchings, the row permanent has quasipolynomial-size circuits exactly when the ordinary permanent does. The equivalence holds for border size as well. The row determinant has quasipolynomial-size circuits over every field, and so does the row permanent in characteristic two.*

Proof. In each pair put the earlier-row endpoint in L and the later one in R, breaking ties arbitrarily. Two entries from the same row never occur together. Every swap needed to move the L letters first therefore exchanges commuting letters. A size-s ordinary circuit of polynomial

degree d gives a formula of size $(s+d+1)^{O(\log(d+1))}$ by arithmetic depth reduction [13]. Apply Proposition E.1. For border size, extract the degree-n component of each approximant first. Placing L before R defines a linear map on all ordinary degree-n coefficients, so the lift preserves closure. Abelianization gives the converse. The ordinary determinant has polynomial-size circuits [3] and equals the permanent in characteristic two. □

We give two further simulations for a matching $\{x_i, y_i\}_{i=1}^{q}$. They allow the order within a pair to vary across the support. Coefficients may also involve central variables.

Theorem E.3. *Suppose each word in F uses each noncentral letter at most once. For pair i let $\mathcal{W}_i = \{1, x_i, y_i, x_iy_i, y_ix_i\}$, and write*

$$F = \sum_{\mathbf{w}} c_{\mathbf{w}} \prod_i w_i, \qquad G = \sum_{\mathbf{w}} c_{\mathbf{w}} \prod_i Y_{i,w_i}.$$

An ordinary formula of size s for G gives a matching-algebra circuit of size $O((s+1)^2)$ for F.

If each pair has just two prescribed local words $w_{i,0}, w_{i,1}$ of arbitrary lengths, a size-s ordinary circuit *for the binary table $\sum_b c_b \prod_i Z_i^{b_i}$ gives a circuit for $\sum_b c_b \prod_i w_{i,b_i}$ of size polynomial in $s+q+D$, where $D = \sum_i \max_j |w_{i,j}|$.*

Proof. For the first assertion take a central parameter t and set

$$A_i = L_{x_i+ty_i}, \qquad B_i = R_{x_i}, \qquad C_i = (A_i - B_i)/t.$$

These operators commute over $K(t)$, and central variables act by multiplication. For the five local words in the order displayed, substitute $\mathrm{Id}, B_i, C_i, A_iC_i, B_iC_i$. On the unit these give $1, x_i, y_i, x_iy_i + ty_i^2, y_ix_i$, respectively. Every table monomial uses one choice per pair and distinct pairs commute. Its constant coefficient in t, after evaluation on the unit, is the desired trace. Compile the formula by the recursion above. Each leaf shifts the Laurent degree of its argument by $O(1)$, and at a product the two shift bounds add. All intermediate exponents therefore lie in an interval $[-O(s+1), O(s+1)]$. Storing its coefficients at each of the $O(s+1)$ steps implements the computation over K.

For the second assertion replace the noncentral letters x by $1+tx$ modulo t^{D+1}, obtaining word polynomials $W_{i,j}$. Each $W_{i,0}$ has constant term one. The elements $T_i = W_{i,0}^{-1}W_{i,1}$ commute across distinct pairs, so the ordinary circuit may be evaluated at the T_i. Multiplying on the left by $\prod_i W_{i,0}$ gives $F(1+tX)$. Truncated convolution and the recursion for an inverse with constant term one use polynomially many gates. This output has degree at most D in t, so its coefficients are complete. Substitute $x-1$ for every noncentral x and sum the coefficients to recover F. □

The same operator argument quantifies the effect of the standard embedding of a trace algebra in a matching algebra [1, Lemma 3.1]. For each edge $e = \{i, j\}$, give x_i, x_j private copies x_i^e, x_j^e forming one noncommuting pair. Map x_i to the product $\Delta(x_i)$ of its incident copies and keep isolated letters as central variables. The projections onto edge pairs, together with the isolated-letter counts, determine a trace. Indeed a first letter of one representative can be moved to the front of any other with the same projections, and induction removes the remaining letters. This proves injectivity.

Suppose the edges admit an orientation such that, on each edge, all occurrences of its tail precede all occurrences of its head in every trace of F. Write $\Delta(x_i) = a_ib_i$, with outgoing copies in a_i and incoming copies in b_i. At an isolated vertex take $a_i = x_i,\ b_i = 1$. The operators $L_{a_i}R_{b_i}$ commute and give $\Delta(F)$ when the ordinary polynomial F^{ab} is evaluated on the unit. Hence a formula of size s for F^{ab} gives a circuit of size $O(s+|V|+|E|)$ for $\Delta(F)$. Conversely, sending one copy of each generator to that generator in the ordinary polynomial ring and its other copies to one recovers F^{ab}. For the multilinear board restriction following Theorem C.1, orient the edges by row order. The image has linear-size circuits, while the original restriction has exponential border size.

Acknowledgment. The author used OpenAI's GPT-6.2 Astra to formulate results, develop and review proofs, and prepare the manuscript.

References

[1] V. Arvind, A. Chatterjee, R. Datta and P. Mukhopadhyay. Equivalence testing of weighted automata over partially commutative monoids. *MFCS 2021*, LIPIcs 202, article 10.

[2] A. Björklund, P. Kaski, T. Koana and J. Nederlof. Kronecker scaling of tensors with applications to arithmetic circuits and algorithms. *ICALP 2026*, LIPIcs 374, article 36.

[3] P. Bürgisser. Completeness classes in algebraic complexity theory. arXiv:2406.06217 (2024).

[4] P. Hrubeš, A. Wigderson and A. Yehudayoff. Non-commutative circuits and the sum-of-squares problem. *J. Amer. Math. Soc.* 24 (2011), 871–898.

[5] K. Khoruzhii. Crossing decomposition for subquadratic noncommutative circuits. Lemmata, revised 4 November 2026.

[6] Y. Liu. Stretched-exponential circuit lower bounds for the noncommutative permanent and determinant. Hexagon 2611.00994v1, 24 November 2026.

[7] N. Nisan. Lower bounds for non-commutative computation. *STOC 1991*, 410–418.

[8] OpenAI. Euclidean geometry in the asymptotic spectrum of tensors. Preprint, 24 November 2026.

[9] OpenAI. The matrix multiplication exponent is nine fourths. Preprint, 24 November 2026.

[10] K. Pratt. A stronger connection between the asymptotic rank conjecture and the Set Cover Conjecture. *STOC 2024*, 871–874.

[11] S. Sra. Euclidean tensor prices and sharp multiquadratic composition bounds. Hexagon 2611.01066v1, 25 November 2026.

[12] V. Strassen. The asymptotic spectrum of tensors. *J. Reine Angew. Math.* 384 (1988), 102–152.

[13] L. G. Valiant, S. Skyum, S. Berkowitz and C. Rackoff. Fast parallel computation of polynomials using few processors. *SIAM J. Comput.* 12 (1983), 641–644.