Stretched-exponential circuit lower bounds for the noncommutative permanent and determinant

Yilin Liu

24 November 2026

Abstract

Over every field, the row-ordered permanent and determinant of an $n \times n$ matrix have border circuit size greater than 2^{n^δ} for some $\delta > 0$ and all sufficiently large n. The bound persists when only $O(n)$ entries are noncentral and each fails to commute with at most three others. Two increasing runs of a permutation encode a repeated binary word. If no planar crossing admits a strict shortening, a circuit decomposition along chords of a square gives the lower bound $2^{n/2-O(\sqrt{n \log n})}$. If a shortening exists, a matrix evaluation transfers OpenAI's stretch theorem through directed Hamiltonian cycles. We also construct hard polynomials on a path whose commutative coefficient tables have linear-size formulas. On a matching, the row-ordered determinant has quasipolynomial-size circuits over every field.

hexagon:2611.00994v1

1 Introduction

For an $n \times n$ matrix $X = (x_{ij})$ of free generators over a field K, write

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

The row order in each word fixes the Cayley convention for the determinant. A circuit has binary addition, subtraction and ordered multiplication, arbitrary field constants and unrestricted reuse of gates. We count arithmetic gates. For homogeneous f of degree d, its *border circuit size* $\underline{C}(f)$ is the least s for which f belongs to the Zariski closure of the degree-d output components of circuits with at most s gates over $\overline{K}$. Intermediate degrees are unrestricted. An exact circuit over K bounds $\underline{C}(f)$ from above.

Theorem 1.1. *In each characteristic there are constants $\delta > 0$ and n_0 such that, over every field of that characteristic and for $n \geq n_0$,*

$$\min\{\underline{C}(\mathrm{Perm}_n), \underline{C}(\mathrm{Det}_n)\} > 2^{n^\delta}.$$

If no planar crossing over $\overline{K}$ admits a strict shortening, then each of these sizes s satisfies $\log_2(s+1) \geq n/2 - O(\sqrt{n \log(n+1)})$.

Nisan proved exponential lower bounds for noncommutative formulas and algebraic branching programs [6]. For circuits with shared gates, Shastri and Willumsgaard proved a tight quadratic bound for the palindrome polynomial when degree and alphabet size are comparable [8]. Hrubeš, Wigderson and Yehudayoff related the permanent to the sum-of-squares problem by encoding a repeated binary word [4]. We use four consecutive blocks of the word as the terminals of a planar crossing. A cut at a circuit gate determines a chord of a square. Long chords give low-cost tensor networks. To transfer the resulting estimate to matrix order n, we read an n-bit subset from the columns of a permutation and read it again from its complement.

Our second result locates the same lower bound in partially commutative algebras. For a graph G on a set of generators V, write

$$A_G(K) = K\langle V \rangle / \langle xy - yx : \ x \neq y, \ \{x, y\} \notin E(G) \rangle.$$

Its word basis consists of equivalence classes under interchanges of adjacent commuting letters, called *traces*. Edges record the pairs that need not commute. We use the same definition of border size in the degree-d trace basis.

Theorem 1.2. *There are explicit graphs G_n on the n^2 entries of X, with $O(n)$ nonisolated vertices and maximum degree three, for which both* Perm_n *and* Det_n *in $A_{G_n}(K)$ have border size greater than 2^{n^η} over every field, for some $\eta > 0$ depending only on the characteristic and all sufficiently large n. In the absence of a planar strict shortening, each size s satisfies*

$$\log_2(s+1) \geq n/12 - O(\sqrt{n \log(n+1)}).$$

On a path, Theorem 4.1 gives a hard polynomial of individual degree two whose commutative coefficient table has a linear-size formula. On a matching, Corollary 5.2 gives quasipolynomial-size circuits for the row determinant over every field.

2 Circuits and planar crossings

All logarithms are to base two. Work over $\overline{K}$. At the four points $e_1, e_2, -e_1, -e_2$ of the Euclidean unit disk place binary strings a, b, c, d of length r. The crossing tensor C_r is the indicator of $a = c,\ b = d$. A bit tensor network has arbitrary local tables. Joining two bit ports contracts their indices, and a cable of h bits costs h times its length. Local operations at a terminal have zero cost. A *strict shortening* of C_r expresses it as a sum of A networks in the disk, of cost at most w each, with $w + 4 \log A < 4r$ [7, Definition 2.1].

We use the planar instance of the stretch theorem [7, Theorem 3.1]. If a strict shortening exists for some r in the given characteristic, an explicit homogeneous family f_j in commutative VNP has degree-j border size greater than 2^{j^σ} for some $\sigma > 0$ and all sufficiently large j. By the Nullstellensatz, existence of a shortening over an algebraically closed field depends only on its characteristic.

The cost test also applies to the closure of sums on fixed layouts. If C_r lies in that closure, then $C_r^{\otimes q}$ lies in the linear span of the q-th powers of actual sums. This follows by polynomial continuity, since a linear subspace is closed. In a fixed tensor space the span of pure q-th powers has dimension polynomial in q, by counting coordinate monomials. Expanding gives C_{qr} as a sum of at most $\mathrm{poly}(q)A^q$ networks of cost at most qw. Thus $w + 4 \log A < 4r$ gives an exact strict shortening for large q.

Lemma 2.1. *Let $P \in K\langle z_0, z_1\rangle$ be homogeneous of degree $4r$. Group its word coefficients into four consecutive blocks of r bits. Suppose the resulting tensor is C_r up to permutations of bit positions within the terminals. If no planar crossing admits a strict shortening, then*

$$\log(1 + \underline{C}(P)) \ \geq\ r - O(\sqrt{r \log(r+1)}).$$

Proof. Give the $4r$ positions the successive directions $e_1, e_2, -e_1, -e_2$, each repeated r times. Their prefix sums $p_0, \ldots, p_{4r}$ traverse the boundary of a square of side r. For $i \leq j$ write $Q_{ij} = \mathrm{conv}\{p_i, \ldots, p_j\}$. By convexity and the boundary order, Q_{ij} is the union of Q_{ik}, the triangle $p_i p_k p_j$, and Q_{kj} whenever $i \leq k \leq j$.

A triangle of boundary points that contains the center has an edge of length at least r. Indeed, if all pairs have distance less than r, the three points meet no pair of opposite sides. They therefore lie on two adjacent sides. In coordinates based at their common corner, their hull lies in $x + y \leq r$. Its intersection with $x + y = r$ can contain the center only if both $(r, 0)$ and $(0, r)$ occur. These points are more than r apart.

Put $d = 4r$. From a size-s circuit, form its homogeneous components through degree d by convolution at each product, using $H = O((d+1)^2(s+1))$ gates. Expand its degree-d output in terms UgV, with a homogeneous gate g in positions $I = (i, j]$, where $i = \deg U$ and $j = i + \deg g$. Start with the full interval. Stop a term as soon as $\|p_j - p_i\| \geq r$. At a sum expand both children with their signs. At a product follow one child, absorbing the other in

U or V. Choose a child with chord length at least r whenever possible. Otherwise all three chords of the splitting triangle are shorter than r. The triangle misses the center, so a child hull contains it whenever the parent hull does. Follow that child. The root hull contains the center. A leaf hull is a point or a unit boundary segment, so every descent in the acyclic circuit stops at a long chord.

For fixed I and g, all stopped occurrences share the coefficient vector of g on I. Summing their signed contexts supplies one vector on I^c. Thus for each interval the collected tensor has flattening rank at most H across $I \mid I^c$, regardless of how often a gate is reused. There are at most $A = (d+1)^2$ intervals.

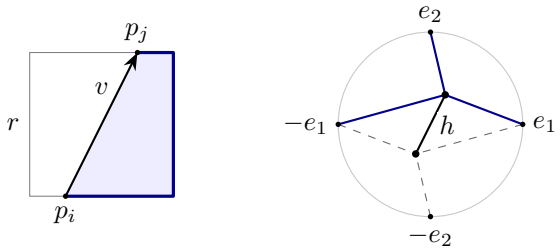

Figure 1: An interval on the square and its rank factorization in the disk. Solid terminal cables carry the bits in the interval, dashed cables its complement. The two hubs lie on the line spanned by the chord v.

Set $v = p_j - p_i$, $a = v/\|v\|$, and $h = \lceil \log H \rceil$. A rank factorization places the two factors at $ta, -ta$ for $0 < t \le 1$, joined by an h-bit cable (Figure 1). Send the bits in I to ta and the remaining bits to $-ta$. Their direction sums are v and $-v$. Since $\|u - ta\| \le 1 - t\langle u, a\rangle + t^2/2$ for unit u, a, each layout costs at most

$$4r - 2t\|v\| + 2rt^2 + 2th \;\le\; 4r + 2t(h - r) + 2rt^2.$$

The layouts depend only on the intervals and s. They cover every degree-d output tensor above, and their sum closure contains the tensor of P when $s = \underline{C}(P)$. Reordering bit ports at a terminal gives C_r at the same cost. The closure cost test now gives $h \ge r - rt - 2\log(A)/t$. Taking $t = \sqrt{2\log(A)/r}$ for large r proves the bound. □

Two crossing cuts. The same cost estimate gives a sharp exponent for a classical problem connected with noncommutative circuits [4]. Suppose no planar crossing over $\mathbb{C}$ admits a strict shortening. If C_r is a sum of two tensors of rank at most R, one across $ab \mid cd$ and one across $bc \mid da$, then

$$\lceil \log R \rceil \ge \sqrt{2}\, r - 2\sqrt{2r} \qquad (r \ge 2).$$

Indeed each interval consists of two adjacent terminal blocks, so $\|v\| = \sqrt{2}\, r$ and $A = 2$ in the preceding cost estimate. Take $t = \sqrt{2/r}$. Now let $R(q)$ be the least number of squares of complex bilinear forms expressing $(\sum_{i=1}^{q} a_i^2)(\sum_{j=1}^{q} b_j^2)$. Polarizing in the two arguments gives

$$2\langle a, c\rangle\langle b, d\rangle = \sum_{\ell} \big(B_\ell(a, b)B_\ell(c, d) + B_\ell(a, d)B_\ell(c, b)\big),$$

where the brackets are the standard bilinear pairing. At $q = 2^r$ the two sums give precisely the two cuts above. Restricting to the preceding power of two gives $R(q) \ge q^{\sqrt{2}-o(1)}$. With Khoruzhii's upper bound [5], this gives $R(q) = q^{\sqrt{2}+o(1)}$ under the same planar hypothesis.

3 Two increasing runs and Hamiltonian cycles

We use matrix substitutions throughout. Evaluating a circuit of size s on $a \times a$ matrices costs $O(a^3(s+1))$ arithmetic gates, in addition to the cost of their entries. In a trace algebra the images of nonadjacent generators must commute. For a homogeneous source of degree d, first extract that component at cost $O((d+1)^2(s+1))$. Substitution, selection of a matrix entry and projection to a target degree are linear on the resulting coefficient space. They preserve Zariski closure, also when some variables map to constants. A sum of polynomially many such evaluations has polynomial overhead.

For a bit string $w = w_1 \cdots w_n$, put $z_w = z_{w_1} \cdots z_{w_n}$ and $D_n = \sum_{w \in \{0,1\}^n} z_w z_w$. For even n, Hrubeš, Wigderson and Yehudayoff obtain D_n from the row permanent of order $2n$ by a scalar and variable substitution [4, Lemma 2.3]. The following reduction uses order n for both signs.

Lemma 3.1. *For $P_n = \mathrm{Perm}_n$ or Det_n, one has $\underline{C}(D_n) + 1 \le n^{O(1)}(\underline{C}(P_n) + 1)$. The same simulation holds for exact circuits.*

Proof. Fix $0 < m < n$. Keep a permutation if its first m columns and its last $n - m$ columns are separately increasing. Such a permutation is the increasing list of a subset $S \subseteq [n]$ followed by the increasing list of S^c, and occurs once for each S of size m.

Use matrices on states $0, \ldots, n$ to remember the previous column p, initially zero in each run. At column $c > p$ emit $z_0^{c-p-1} z_1$ in the first run and $z_1^{c-p-1} z_0$ in the second. The next state is c. On the last row of each run append z_0^{n-c} or z_1^{n-c}, respectively, and return to zero. All other transitions have weight zero. The $(0,0)$ entry emits the incidence word of S twice. For example, the lists $2,4$ and $1,3,5$ both emit 01010.

For the determinant multiply the transition in row $i \le m$, column c, by $(-1)^{c-i}$. Exactly $c - i$ smaller columns lie in the complement, so these signs cancel every inversion of the shuffle. Each surviving word has coefficient one. Sum over m and add z_0^{2n}, z_1^{2n} from the two constant strings. The matrices have dimension $n+1$, with signed words of length $O(n)$ as entries. The matrix simulation and its border version give the claim. □

Proof of Theorem 1.1. If no planar crossing admits a strict shortening, apply Lemma 2.1 to D_{2r}, whose four blocks have coefficient tensor C_r. Lemma 3.1 gives the stated estimate at $n = 2r$. For odd n, fixing the last row and column to an identity block projects either polynomial to the preceding order.

If a planar strict shortening exists, the stretch theorem gives the integral commutative family f_j of Section 2. Let HC_q sum the edge monomials of the directed Hamiltonian cycles on $[q]$, counted up to cyclic rotation. By Valiant's projection completeness theorem over the prime field [9] (see also [3, Section 2.8]), f_j is a projection of $\mathrm{HC}_{q(j)}$ with $q(j) \le Cj^a$ for fixed $C, a > 0$. For $n \ge q(j)$, split one vertex into a mandatory directed chain to project HC_n to $\mathrm{HC}_{q(j)}$. Give the chain arcs weight one, attach incoming and outgoing arcs at the two ends, and set unused arcs to zero. Homogenizing before these substitutions preserves the border reduction. Taking $j = \lfloor (n/C)^{1/a} \rfloor$ gives $\underline{C}(\mathrm{HC}_n) > 2^{n^\gamma}$ for every fixed $0 < \gamma < \sigma/a$ and all large n.

Let E_{ij} denote matrix units and $M_j = \sum_{a=1}^{n} y_{aj} E_{aj}$ over the commutative polynomial ring $K[y]$. Substitute M_j in column j of Perm_n, with all entries of row one except its first set to zero. The trace of the term indexed by $(1, j_2, \ldots, j_n)$ is

$$y_{1j_2} y_{j_2 j_3} \cdots y_{j_{n-1} j_n} y_{j_n 1}.$$

Every directed Hamiltonian cycle appears once, including in characteristic two. The trace is HC_n, so the polynomial simulation transfers its lower bound to Perm_n.

For Det_{2n} use the interleaving reduction of Arvind and Srinivasan [2]. In row $2i-1$, column $j \le n$, put $x_{ij} E_{0j}$. In row $2i$, column $n+j$, put E_{j0}, with all other entries zero. In the $(0,0)$

entry the column pairs are $(\pi(i), n+\pi(i))$. Their sign is a fixed interleaving sign times $\operatorname{sgn}(\pi)^2$, giving $\pm\operatorname{Perm}_n$. Identity padding covers odd orders. The polynomial losses are absorbed by decreasing the exponent.

The constants can be chosen once per characteristic. For fixed n, s, the finitely many circuit shapes give polynomial maps from their scalar parameters to output coefficients, all defined over the prime field. Each image-closure ideal is the kernel of the corresponding map on coordinate rings and commutes with field extension by flatness. Thus the border size of a prime-field polynomial is the same in every algebraically closed extension. □

4 Sparse commutation relations

Two matrix evaluations will serve for both the row polynomials and the path construction. Take two disjoint $k \times k$ grids U, V, put $r = k^2$, and list $N = 4r$ visits to their cells in the order

$$U \text{ by rows}, \qquad V \text{ by rows}, \qquad U \text{ by columns}, \qquad V \text{ by columns}.$$

Each cell is visited twice. Write $\lambda(t)$ for the cell at visit t. Given two matrices $M_{t,0}, M_{t,1}$ of a common dimension at every visit, form the ordered matrix product sum

$$W(M) = \sum_{b \in \{0,1\}^{U \cup V}} M_{1,b_{\lambda(1)}} \cdots M_{N,b_{\lambda(N)}}.$$

The following two choices will be used.

- With scalar entries $M_{t,b} = z_b$, the four blocks of the free polynomial $W(M)$ have tensor C_r up to the transposition of bit positions in each of the last two blocks.
- With two states, one entry of $W(M)$ is the free polynomial Perm_k. Force all bits of V to zero by using $M_{t,0} = I_2$, $M_{t,1} = 0$ there. During the first visit to U require one chosen bit in every row, emitting x_{ij} at the chosen cell (i, j). During the second visit require one chosen bit in every column, with all weights one. The two states record whether the current row or column has seen a chosen bit. A second choice has weight zero. At the end of each row or column accept a count of one and reset to zero. The entry from zero to zero sums each permutation word once.

Proof of Theorem 1.2. For $n < 10$ take the edgeless graph. Otherwise choose k maximal with $n_0 = 3N - 2 = 12k^2 - 2 \le n$. Insert two single rows between successive visits. In these n_0 rows, let τ exchange the two visit rows of each cell. A visit row i has two active entries, in columns $i, \tau(i)$, for bits zero and one. Each inserted row has just its diagonal entry active. Join every active entry to every active entry of the next row, and include no further edges. The active counts are $2, 1, 1, 2, 1, 1, \ldots, 2$, so the maximum degree is three and at most $4N - 2$ entries are noncentral. Each interior visit gives a four-cycle on its two entries and their neighboring diagonal entries. Consecutive four-cycles are joined by single edges.

Every permutation on the active board is a product of disjoint transpositions. The two visits to each cell either both use their diagonals or exchange columns. A matrix substitution makes these choices read out $W(M)$. In core row i, at bit b of visit t, use $E_{i,i+1} \otimes M_{t,b}$. On an inserted diagonal use $E_{i,i+1} \otimes I$. Use identity matrices on padding diagonals and zero elsewhere. Distinct active entries in one row, or in rows at distance at least two, have product zero in either order. The required commutations hold. The block from address 1 to $n_0 + 1$ is $W(M)$. For the determinant multiply the bit-one matrix on the first visit to each cell by -1. This cancels exactly the sign of the corresponding transposition.

All matrices have order $O(n)$. The second choice of M and Theorem 1.1 give the stretched-exponential bound since $k = \Theta(\sqrt{n})$. Under the planar hypothesis, the first choice and

Lemma 2.1 give $\log(1+s) \geq k^2 - O(\sqrt{k^2 \log(k+1)}) - O(\log n)$. The choice of k gives $k^2 = n/12 - O(\sqrt{n})$, as required. □

The multilinear restriction to the active board, with identity padding, has the same two evaluations. Its abelianization is the product of the inserted diagonal variables and the binomials $x_{ii}x_{jj} + \varepsilon x_{ij}x_{ji}$ over paired visit rows $\{i, j\}$. Taking $\varepsilon = 1$ for the permanent and $\varepsilon = -1$ for the determinant gives linear-size formulas for both.

Permutation words on a path

Keep the $N = 4k^2$ visit rows and omit the inserted rows. In the $N \times N$ board let D consist of the four cells $\{i, j\} \times \{i, j\}$ for each pair of visits to a common bit. These squares use disjoint row sets and disjoint column sets. List the $m = 2N$ cells of D in row-major order. Give them variables $z_1, \ldots, z_m$ with noncommutation graph the path in that order. If cell h is (i, j), put $\chi_h(\pi) = [\pi(i) = j]$. Define

$$F_D = \sum_{\substack{\pi \in S_N \\ (i,\pi(i)) \in D \ \forall i}} z_1^{1+\chi_1(\pi)} \cdots z_m^{1+\chi_m(\pi)}.$$

The commutative *coefficient table* assigns a marker Y_h to exponent two and records the supported patterns as ordinary squarefree monomials.

Theorem 4.1. *The homogeneous path polynomial F_D has individual degree two. Over each field, its border circuit size s exceeds 2^{m^η} for some $\eta > 0$ and all large k. In the absence of a planar strict shortening,*

$$\log(1+s) \geq m/8 - O(\sqrt{m \log(m+1)}).$$

The coefficient table of F_D and its abelianization F_D^{ab} both have linear-size formulas.

Proof. In a visit row t, let $B_h = M_{t,0}$ on its diagonal cell and $B_h = M_{t,1}$ on its other cell. Use addresses $1, \ldots, m+1$ and one distinct extra address u_h for each h. Substitute

$$Z_h = E_{h,h+1} \otimes I + E_{h,u_h} \otimes B_h + E_{u_h,h+1} \otimes I \qquad \text{for } z_h.$$

Matrices at nonconsecutive positions have disjoint address supports. The square of Z_h is $E_{h,h+1} \otimes B_h$. In a first power only its direct h-to-$h+1$ step can reach the final address. Thus each row contributes $M_{t,b}$ at its chosen cell and the identity at its other cell. The block from 1 to $m+1$ is exactly $W(M)$. The two evaluations above prove the bounds, using $r = m/8$.

Each permitted permutation chooses one of two matchings in every square independently. Its coefficient table is therefore the product of $Y_{ii}Y_{jj} + Y_{ij}Y_{ji}$ over the squares, with markers named by their cells. Multiplication by $\prod_h z_h$, after replacing each marker by its variable, gives the abelianization. The degree of every monomial of F_D is $m + N$. □

Listing all N^2 cells in row-major order gives another path. Set $B_h = 0$ off D. The same evaluations, with $r = N/4$ in the crossing bound, apply both to the sum supported on D and to the sum over all permutations, since a doubled off-board letter gives a zero matrix. For the restricted sum the abelianization includes one extra factor z_h per off-board cell. For the full sum of any order N, putting $B_h = x_{ij}$ at every cell (i, j) recovers Perm_N directly.

5 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 5.1. *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 5.2. *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 [10]. Apply Proposition 5.1. 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 5.3. *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 $\Delta(x_i) = a_i b_i$, 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 1.2, orient the edges by row order. The image has linear-size circuits, while the original restriction has stretched-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] V. Arvind and S. Srinivasan. On the hardness of the noncommutative determinant. *Comput. Complexity* 27 (2018), 1–29.

[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 (2026), revised 4 November.

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

[7] OpenAI. VP and VNP are different. Preprint (23 November 2026).

[8] P. Shastri and S. V. Willumsgaard. A quadratic lower bound for noncommutative circuits. arXiv:2604.20575v4 (2026).

[9] L. G. Valiant. Completeness classes in algebra. *STOC 1979*, 249–261.

[10] 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.