A super-cubic lower bound for the formula size of the permanent

Abstract

We prove that over any field of characteristic zero the arithmetic formula size of the $n \times n$ permanent is $\Omega(n^4/\log n)$, and in particular $\omega(n^3)$. The proof combines the Nechiporuk–Kalorkoti transfer principle (the number of leaves of a formula carrying variables from a set Y bounds the transcendence degree of the coefficient family obtained by expanding the computed polynomial in Y) with a new Jacobian certificate: for a block of $\Theta(\log n)$ cells lying in distinct rows and distinct columns, the associated family of padded permanental minors has transcendence degree $\Omega(n^2)$. The certificate is obtained by evaluating the Jacobian at a single explicit point built from roots of unity and identifying its rows with distinct characters of $\mathbb{Z}_M^2$.

Theorem. *Let $\mathbb{F}$ be any field of characteristic zero and let $L(\mathrm{perm}_n)$ denote the minimum number of leaves of an arithmetic formula over $\mathbb{F}$ (a rooted binary tree whose leaves are labelled by variables or constants and whose internal nodes are labelled $+$ or $\times$) computing*

$$\mathrm{perm}_n \;=\; \sum_{\sigma \in S_n} \prod_{i=1}^{n} x_{i\sigma(i)}.$$

Then

$$L(\mathrm{perm}_n) \;=\; \Omega\!\left(\frac{n^4}{\log n}\right).$$

In particular $L(\mathrm{perm}_n) = \omega(n^3)$.

Together with Ryser's formula, which gives $L(\mathrm{perm}_n) = 2^{O(n)}$ (see the end of Section 5), the theorem locates the growth rate of $L(\mathrm{perm}_n)$ between $n^4/\log n$ and $2^{O(n)}$.

The proof is organized as follows. Section 1 fixes notation and records the elementary direction of the Jacobian criterion. Section 2 proves the Nechiporuk–Kalorkoti transfer lemma. Section 3 sets up a partition of the matrix of variables into $\Theta(n^2/\log n)$ "diagonal arcs" of size $\Theta(\log n)$ and identifies the per-arc coefficient family (padded permanental minors). Section 4 contains the Key Lemma: each arc family has transcendence degree $\Omega(n^2)$. Section 5 assembles the theorem, and Section 6 contains concluding remarks.

Throughout, $[n] = \{1, \dots, n\}$; for a matrix A and index sets R, C, the matrix "A with rows R and columns C deleted" is the submatrix of A on the remaining rows and columns; the permanent of the 0×0 matrix is 1. We abbreviate transcendence degree by td.

1 Preliminaries

Formulas. A formula Φ over $\mathbb{F}$ in variables $x_1, \dots, x_N$ is a rooted binary tree; each leaf is labelled by a variable or by a constant of $\mathbb{F}$, and each internal node is labelled $+$ or $\times$ and

has exactly two children. We write $\widehat{\Phi}$ for the polynomial computed at the root in the obvious bottom-up fashion, and $|\Phi|$ for the number of leaves of Φ. The formula size of f is $L(f) := \min\{|\Phi| : \widehat{\Phi} = f\}$.

Transcendence degree. For a family $\mathcal{C} \subseteq \mathbb{F}[Z]$ of polynomials in a set of variables Z, we write $\mathrm{td}_{\mathbb{F}}\mathcal{C}$ for the transcendence degree over $\mathbb{F}$ of the subfield $\mathbb{F}(\mathcal{C}) \subseteq \mathbb{F}(Z)$. We use two standard facts: a field extension generated by N elements has transcendence degree at most N; and the following Jacobian criterion.

Lemma 1.1 (Jacobian certificate, characteristic 0)**.** *Let $\mathbb{F}$ be a field with $\mathrm{char}\,\mathbb{F} = 0$ and let $c_1, \dots, c_r \in \mathbb{F}[Z]$. If some $r \times r$ minor of the Jacobian matrix $\big(\partial c_i/\partial z\big)_{i\in[r],\, z\in Z}$ is a nonzero element of $\mathbb{F}[Z]$, then $c_1, \dots, c_r$ are algebraically independent over $\mathbb{F}$; that is, $\mathrm{td}_{\mathbb{F}}\{c_1, \dots, c_r\} = r$.*

Proof. Suppose not, and let $P \in \mathbb{F}[y_1, \dots, y_r]$ be a nonzero polynomial of minimal total degree with $P(c_1, \dots, c_r) = 0$. Applying $\partial/\partial z$ for each $z \in Z$ and using the chain rule,

$$\sum_{i=1}^{r} (\partial_{y_i} P)(c_1, \dots, c_r) \cdot \frac{\partial c_i}{\partial z} = 0 \qquad \text{for every } z \in Z.$$

Thus the row vector $\big((\partial_{y_1} P)(\mathbf{c}), \dots, (\partial_{y_r} P)(\mathbf{c})\big)$, with $\mathbf{c} = (c_1, \dots, c_r)$, annihilates the Jacobian matrix over the field $\mathbb{F}(Z)$. This vector is nonzero: since $\deg P \geq 1$ and $\mathrm{char}\,\mathbb{F} = 0$, some $\partial_{y_i} P$ is a nonzero polynomial, and it has strictly smaller degree than P, so $(\partial_{y_i} P)(\mathbf{c}) \neq 0$ by minimality of $\deg P$. A nonzero vector over $\mathbb{F}(Z)$ annihilating the matrix forces every $r \times r$ minor to vanish identically, contradicting the hypothesis. □

Remark 1.2 (Base field). In Section 4 we exhibit a minor with integer coefficients and prove that it is nonzero by evaluating it at a point whose coordinates lie in a cyclotomic field $\mathbb{Q}(\omega)$. If a polynomial with rational coefficients takes a nonzero value at some point of $\mathbb{Q}(\omega)^N$, then it is a nonzero polynomial over $\mathbb{Q}$, hence also nonzero as an element of $\mathbb{F}[Z]$ for every field $\mathbb{F} \supseteq \mathbb{Q}$, i.e. for every field of characteristic zero. Consequently all transcendence-degree lower bounds proved below hold simultaneously over every characteristic-zero field.

2 The transfer lemma

Lemma 2.1 (Nechiporuk–Kalorkoti transfer)**.** *Let Φ be a formula computing f, let Y be a set of variables, and let Z be the set of all other variables occurring in Φ. Expand*

$$f \;=\; \sum_{\alpha} c_\alpha(Z)\, Y^\alpha, \qquad c_\alpha \in \mathbb{F}[Z],$$

the sum being over monomials Y^α in the variables of Y (only finitely many c_α are nonzero). Let $s_Y(\Phi)$ denote the number of leaves of Φ labelled by a variable of Y. Then

$$\mathrm{td}_{\mathbb{F}}\{c_\alpha\}_\alpha \;\leq\; 4\, s_Y(\Phi) + 2.$$

Proof. Write $s = s_Y(\Phi)$. If $s = 0$ then $f = c_0 \in \mathbb{F}[Z]$ and the coefficient family is a single polynomial, of transcendence degree at most $1 \leq 2$. Assume from now on $s \geq 1$.

Let W be the set of Y-labelled leaves, $|W| = s$, and let $P \subseteq \Phi$ be the union of the root-to-leaf paths ending in the leaves of W; thus P is a subtree of Φ whose leaf set is exactly W. Call

$v \in P$ a *branch vertex* if both children of v lie in P; a tree with s leaves has at most $s-1$ branch vertices. Put

$$V' \;=\; W \;\cup\; \{\text{branch vertices of } P\} \;\cup\; \{\text{root}\}, \qquad |V'| \le 2s.$$

Contracting P onto V' yields a tree T' with vertex set V' and at most $2s-1$ edges: each edge e of T' joins a vertex $v \in V'$ to the first vertex $v^+ \in V'$ strictly above it in P. Define the *gate set* of e to consist of all vertices of P lying strictly between v and v^+, together with v^+ itself in the single case where v^+ is the root and the root is not a branch vertex. In this way every vertex of P that is neither a W-leaf nor a branch vertex belongs to the gate set of exactly one edge.

Fix an edge e of T' and list its gates $w_1, \dots, w_t$ in upward order (possibly $t = 0$). Each w_j is an internal node of Φ lying on P which is not a branch vertex, so exactly one of its two children lies in P; the subtree hanging at the other child contains no Y-leaf and therefore computes some polynomial $u_j \in \mathbb{F}[Z]$. If w_j is a $+$-node it transforms the value g arriving from its P-child into $g + u_j$; if it is a $\times$-node, into $u_j\, g$. Composing these maps along the path yields an affine map

$$g \;\longmapsto\; A_e\, g + B_e, \qquad A_e, B_e \in \mathbb{F}[Z]$$

(with $A_e = 1$, $B_e = 0$ when $t = 0$).

Let $R \subseteq \mathbb{F}[Z]$ be the $\mathbb{F}$-subalgebra generated by the at most $2(2s-1) = 4s-2$ elements $\{A_e, B_e : e \in E(T')\}$. We claim, by induction upward along T', that the polynomial computed at each non-root vertex of V' lies in $R[Y]$, and that $f \in R[Y]$.

Base case. At a leaf $w \in W$ the computed value is a variable of Y, which lies in $R[Y]$.

Induction step at a branch vertex. Let $v \in V'$ be a branch vertex, with operation $\circ \in \{+, \times\}$, and let e_1, e_2 be the two T'-edges entering v from below, coming from vertices $v_1, v_2 \in V'$ whose values G_1, G_2 lie in $R[Y]$ by induction. Then the value computed at v equals

$$(A_{e_1} G_1 + B_{e_1}) \circ (A_{e_2} G_2 + B_{e_2}) \;\in\; R[Y].$$

The root. If the root is a branch vertex, the previous step applied at the root exhibits $f \in R[Y]$. If the root is neither a branch vertex nor a leaf, let e_0 be the unique top edge of T', with lower endpoint v_0 and value $G \in R[Y]$; the gate set of e_0 contains the root, so the affine map $A_{e_0}(\cdot) + B_{e_0}$ accounts for the root's own operation together with all gates below it, and $f = A_{e_0} G + B_{e_0} \in R[Y]$. Finally, if the root is itself a W-leaf then Φ is a single leaf and f is a variable of Y.

Since $Y \cap Z = \emptyset$, the expansion of an element of $\mathbb{F}[Z][Y]$ in monomials of Y is unique, and the coefficients of an element of $R[Y]$ lie in R. Hence every c_α lies in R, and

$$\mathrm{td}_{\mathbb{F}}\{c_\alpha\} \;\le\; \mathrm{td}_{\mathbb{F}}\, \mathbb{F}\big(\{A_e, B_e\}_{e \in E(T')}\big) \;\le\; 2(2s-1) \;=\; 4s-2 \;\le\; 4s+2. \qquad \square$$

Corollary 2.2 (Partition bound)**.** *Let $Y_1, \dots, Y_p$ be pairwise disjoint sets of variables and let f be a polynomial. For each i, let τ_i be the transcendence degree of the family of coefficients obtained by expanding f in the variables of Y_i (the coefficients being polynomials in all remaining variables). Then for every subset $G \subseteq [p]$,*

$$L(f) \;\ge\; \sum_{i \in G} \frac{\tau_i - 2}{4}.$$

Proof. Fix an optimal formula Φ for f, so that $L(f) = |\Phi|$ is its number of leaves. Each leaf carries at most one variable label, and by disjointness of the Y_i that variable belongs to at most one Y_i; hence the leaf sets counted by $s_{Y_i}(\Phi)$, $i \in G$, are pairwise disjoint and $L(f) \geq \sum_{i \in G} s_{Y_i}(\Phi)$. Applying Lemma 2.1 to the same formula Φ for each i (with $Y := Y_i$ and $Z :=$ all other variables occurring in Φ) gives $s_{Y_i}(\Phi) \geq (\tau_i - 2)/4$. Summing over $i \in G$ completes the proof. □

3 Arc blocks and padded permanental minors

Fix n and a block size $k \leq n$. For $0 \leq r \leq n-1$ let

$$D_r \;=\; \big\{\, (i,\; 1 + ((i-1+r) \bmod n)) \;:\; i \in [n] \,\big\}$$

be the r-th cyclic diagonal of the grid $[n]\times[n]$. The sets $D_0, \ldots, D_{n-1}$ partition the grid, and each of them is the support of a permutation matrix. Cut each D_r into $\lfloor n/k \rfloor$ *arcs* of k consecutive cells (namely the cells with row indices $jk+1, \ldots, jk+k$ for $0 \leq j < \lfloor n/k \rfloor$), and place the at most $k-1$ leftover cells of each diagonal into a residual set. This produces

$$p_0 \;=\; n\,\lfloor n/k \rfloor$$

pairwise disjoint *arc blocks*, each consisting of k cells with pairwise distinct rows and pairwise distinct columns, together with residual cells; all together they partition the n^2 variables. Only the arc blocks will be used, via the subset G in Corollary 2.2.

Lemma 3.1 (Symmetry)**.** *Let $Y = \{(a_1, b_1), \ldots, (a_k, b_k)\}$ be any set of k cells with pairwise distinct rows and pairwise distinct columns. The transcendence degree of the coefficient family obtained by expanding* perm_n *in the variables of Y equals the corresponding quantity for the diagonal prefix $Y_0 = \{(1,1), \ldots, (k,k)\}$.*

Proof. Choose $\pi, \rho \in S_n$ with $\pi(a_j) = j$ and $\rho(b_j) = j$ for $j \in [k]$; this is possible since the a_j are distinct and the b_j are distinct. The substitution $x_{ab} \mapsto x_{\pi(a)\rho(b)}$ is a bijective relabelling of the variables, hence induces an $\mathbb{F}$-algebra automorphism Ψ of $\mathbb{F}[x_{ab}]$; it maps perm_n to perm_n, because permuting rows and columns leaves the permanent unchanged, and it maps the variables of Y bijectively onto the variables of Y_0 and the remaining variables bijectively onto the remaining variables. Applying Ψ to the expansion $\mathrm{perm}_n = \sum_\alpha c_\alpha Y^\alpha$ gives $\mathrm{perm}_n = \sum_\alpha \Psi(c_\alpha) Y_0^\alpha$, so the coefficient family for Y_0 is the image under Ψ of the coefficient family for Y. Automorphisms preserve transcendence degree. □

The chunk family. Let $Y_0 = \{(1,1), \ldots, (k,k)\}$ (the *chunk*) and let Z denote the remaining $n^2 - k$ variables. Since perm_n is multilinear and the chunk cells lie in distinct rows and distinct columns, the monomials of perm_n in the chunk variables are exactly $\prod_{i \in T} x_{ii}$ for $T \subseteq [k]$, and

$$\mathrm{perm}_n \;=\; \sum_{T \subseteq [k]} \Big(\prod_{i \in T} x_{ii} \Big)\, c_T,$$

where

$$c_T \;:=\; \mathrm{perm}\big(X \text{ with rows and columns } T \text{ deleted, and } x_{ii} \mapsto 0 \text{ for } i \in [k] \setminus T\big) \;\in\; \mathbb{F}[Z].$$

Indeed, the permutations contributing to the coefficient of $\prod_{i\in T} x_{ii}$ are exactly those using all chunk cells of T and avoiding all other chunk cells; deleting the rows and columns indexed by T accounts for the former, and the substitution $x_{ii} \mapsto 0$ ($i \in [k] \setminus T$) accounts for the latter. We write

$$\tau_n(k) \;:=\; \mathrm{td}_{\mathbb{F}}\{c_T :\; T \subseteq [k]\},$$

so that, by Lemma 3.1, every arc block has coefficient family of transcendence degree exactly $\tau_n(k)$.

For the Jacobian computation we record the following: for a non-chunk cell (a, b) with $a, b \notin T$,

$$\frac{\partial c_T}{\partial x_{ab}} \;=\; \mathrm{perm}\big(X \text{ with rows } T\cup\{a\} \text{ and columns } T\cup\{b\} \text{ deleted, and } x_{ii}\mapsto 0 \ (i\in[k]\setminus T)\big), \tag{3.1}$$

since the permanent is multilinear with $\partial\,\mathrm{perm}\,/\partial x_{ab}$ equal to the permanental minor at (a, b), and the padding substitution commutes with this differentiation (the padded variables are distinct from x_{ab}). (If $a \in T$ or $b \in T$, the variable x_{ab} does not occur in c_T and the derivative is 0.)

4 The Key Lemma

Lemma 4.1 (Key Lemma)**.** *There is an absolute constant n_0 such that for all $n \geq n_0$ and every integer k with $2\log_2 n + 4 \leq k \leq n/2$,*

$$\tau_n(k) \;\geq\; \frac{3}{64}\,(n-k+1)^2 \;\geq\; \frac{3}{256}\,n^2.$$

Remark. Numerical computation of generic Jacobian ranks (not needed below) suggests that the truth is exactly $\min\big(2^k,\ (n-1)^2+1\big)$; the constant $3/256$ is what the construction below yields and is ample for our purposes.

4.1 The evaluation point

Fix n and k as in the statement and put $m := n - k$ ($\geq n/2$). Let

$$a := \lfloor \log_2(m+1) \rfloor \quad (\text{so } 2^a \leq m+1 < 2^{a+1}), \qquad M := 2^a - 1 \quad (\text{odd},\ \tfrac{m+1}{2} \leq 2^a,\ M \leq m),$$

and let $\omega \in \mathbb{C}$ be a primitive M-th root of unity. For $n \geq n_0$ (any $n_0 \geq 2^{10}$ works) we have $m \geq n/2 \geq 2^9$ and hence $a \geq 9$; moreover $m + 1 \leq n$, so

$$2(a-1) \;\leq\; 2\log_2(m+1) - 2 \;\leq\; 2\log_2 n \;\leq\; k.$$

We may therefore fix two disjoint subsets $I_1 = \{\alpha_1 < \cdots < \alpha_{a-1}\}$ and $I_2 = \{\beta_1 < \cdots < \beta_{a-1}\}$ of $[k]$. Define weight functions $g, h : I_1 \cup I_2 \to \mathbb{Z}_{\geq 1}$ by

$$g(\alpha_j) = 2^{j-1}, \quad g(\beta_j) = 2^{j-1}, \qquad h(\alpha_j) = 2^{j-1}, \quad h(\beta_j) = 2^j \qquad (1 \leq j \leq a-1),$$

and for $B \subseteq I_1 \cup I_2$ write $\sigma_g(B) = \sum_{i\in B} g(i)$ and $\sigma_h(B) = \sum_{i\in B} h(i)$.

Define a point $P^{\#}$, an $n \times n$ matrix with entries in $\mathbb{Q}(\omega)$, whose rows and columns are indexed by the *chunk indices* $1, \ldots, k$ followed by the *bottom indices* $k + u$ for $u \in [m]$:

- top-left $k \times k$ block: identically 0 (in particular all pad entries x_{ii}, $i \in [k]$, are 0);
- top-right cross: $P^{\#}[\,i,\, k+u\,] = \omega^{\,g(i)\,u}$ for $i \in I_1 \cup I_2$ and $u \le M$; all other top-right entries are 0 (i.e. for $u > M$, and for $i \notin I_1 \cup I_2$ the whole row is set to 0; its value is irrelevant, since such rows i will always be deleted);
- bottom-left cross: $P^{\#}[\,k+u,\, i\,] = \omega^{\,h(i)\,u}$ for $i \in I_1 \cup I_2$ and $u \le M$; all other bottom-left entries are 0;
- bottom-right $m \times m$ block: all entries equal to 1.

We shall use only the following part of the Jacobian matrix of the family $\{c_T\}$:

- **rows** indexed by $T_K := [k] \setminus K$ for $K \subseteq I_1 \cup I_2$;
- **columns** indexed by the bottom-block cells $(k+s,\, k+t)$ with $s, t \in [M]$ (these are genuine variables of Z, since only the k chunk-diagonal cells are excluded from Z).

By (3.1), and because all pad entries of $P^{\#}$ are already 0, the corresponding Jacobian entry evaluated at $P^{\#}$ equals

$$J\big[T_K,\ (k+s, k+t)\big]\Big|_{P^{\#}} \;=\; \operatorname{perm}\big(N_{K,s,t}\big), \tag{4.1}$$

where $N_{K,s,t}$ is the matrix $P^{\#}$ with rows $T_K \cup \{k+s\}$ and columns $T_K \cup \{k+t\}$ deleted.

4.2 Exact evaluation of the entries

Throughout we set $\kappa := |K| \le 2(a-1)$, and we freely identify $[M] = \{1, \dots, M\}$ with $\mathbb{Z}_M$ via $u \mapsto u \bmod M$ (legitimate because every quantity $\omega^{\gamma u}$ depends only on $u \bmod M$).

Lemma 4.2 (Factorization)**.** *Let $K \subseteq I_1 \cup I_2$ with $\kappa = |K| \le m-1$, and let $s, t \in [M]$. Then*

$$\operatorname{perm}(N_{K,s,t}) \;=\; (m-1-\kappa)! \left(\sum_{\pi\colon K \hookrightarrow [M]\setminus\{t\}} \prod_{i \in K} \omega^{g(i)\pi(i)} \right) \left(\sum_{\rho\colon K \hookrightarrow [M]\setminus\{s\}} \prod_{i \in K} \omega^{h(i)\rho(i)} \right),$$

where both sums range over injections (an empty set of injections, which occurs when $\kappa > M-1$, contributes the value 0; for $K = \emptyset$ each sum equals 1).

Proof. The rows of $N_{K,s,t}$ are indexed by $K \sqcup \{k+u : u \in [m] \setminus \{s\}\}$ and its columns by $K \sqcup \{k+v : v \in [m] \setminus \{t\}\}$. Consider a bijection φ from rows to columns whose associated product of entries is nonzero.

A top row $i \in K$ has nonzero entries only in columns $k+v$ with $v \le M$ and $v \ne t$ (the top-left block of $P^{\#}$ is zero, and the top-right cross entries vanish for $v > M$). Hence the restriction of φ to the rows in K determines an injection $\pi : K \hookrightarrow [M] \setminus \{t\}$, contributing the entry product $\prod_{i \in K} \omega^{g(i)\pi(i)}$.

Dually, a top column $i \in K$ can only be hit by a bottom row $k+u$ with $u \le M$ and $u \ne s$. Hence the restriction of φ^{-1} to the columns in K determines an injection $\rho : K \hookrightarrow [M] \setminus \{s\}$, contributing $\prod_{i \in K} \omega^{h(i)\rho(i)}$.

The bottom columns used by π and the bottom rows used by ρ are independent resources; after they are fixed, the remaining $m-1-\kappa$ bottom rows must be matched bijectively with the

remaining $m-1-\kappa$ bottom columns, and every such matching lies inside the all-ones block, contributing a factor 1. There are $(m-1-\kappa)!$ matchings. Conversely, every triple (injection π, injection ρ, bijective completion) arises from exactly one bijection φ with nonzero entry product. Summing the products over all φ gives the stated identity. □

The two injection sums are completely decoupled, and each is evaluated in closed form by the following character computation, which is the heart of the proof.

Lemma 4.3 (Character collapse)**.** *Let $M \geq 2$, let $\omega \in \mathbb{C}$ be a primitive M-th root of unity, let K be a finite set with integer weights $(\gamma_i)_{i\in K}$, and put $\kappa = |K|$ and $\sigma(B) := \sum_{i\in B} \gamma_i$. Assume*

$$\sigma(B) \not\equiv 0 \pmod{M} \qquad \text{for every nonempty } B \subseteq K. \tag{4.2}$$

Then $\kappa \leq M-1$, and for every $t \in \mathbb{Z}_M$,

$$\sum_{\pi\colon K\hookrightarrow \mathbb{Z}_M\setminus\{t\}} \omega^{\sum_{i\in K}\gamma_i\pi(i)} \;=\; (-1)^\kappa\,\kappa!\;\omega^{\,t\,\sigma(K)}.$$

Proof. First, (4.2) forces $\kappa \leq M-1$: order K as $i_1,\dots,i_\kappa$ and consider the $\kappa+1$ prefix sums $0, \gamma_{i_1}, \gamma_{i_1}+\gamma_{i_2}, \dots$; if $\kappa \geq M$, two of them would be congruent modulo M, and the corresponding nonempty block $B = \{i_{p+1},\dots,i_q\}$ would satisfy $\sigma(B) \equiv 0 \pmod{M}$.

Work in the $\mathbb{C}$-algebra

$$S := \mathbb{C}[x_i : i \in K]/(x_i^2 : i \in K),$$

which has $\mathbb{C}$-basis $\{x_B := \prod_{i\in B} x_i \;:\; B \subseteq K\}$; write $[x_B]$ for the operator extracting the coefficient of x_B. For $u \in \mathbb{Z}_M$ set $L_u := \sum_{i\in K} x_i\,\omega^{\gamma_i u} \in S$.

(a) *Injection sums are coefficients.* For any subset $E \subseteq \mathbb{Z}_M$,

$$[x_K]\prod_{u\in E}(1+L_u) \;=\; \sum_{\pi:K\hookrightarrow E}\omega^{\sum_{i\in K}\gamma_i\pi(i)}.$$

Indeed, expanding the product, each factor indexed by u contributes either 1 or a single term $x_i\omega^{\gamma_i u}$; the coefficient of x_K collects exactly those choices in which each $i \in K$ is selected precisely once, necessarily at pairwise distinct values of u, i.e. the injections $\pi : K \hookrightarrow E$ (with $\pi(i)$ the index of the factor contributing x_i).

(b) *The full product equals* 1. Let $P := \prod_{u\in\mathbb{Z}_M}(1+L_u)$. The assignment $x_i \mapsto \omega^{\gamma_i}x_i$ extends to an algebra automorphism of S which maps L_u to L_{u+1}; it therefore permutes the factors of P cyclically and fixes P. On the other hand it multiplies the coefficient of x_B by $\omega^{\sigma(B)}$. Hence $[x_B]P = \omega^{\sigma(B)}\,[x_B]P$ for every B, so $[x_B]P = 0$ whenever $\sigma(B) \not\equiv 0 \pmod{M}$, which by (4.2) is the case for every nonempty B. Since $[x_\emptyset]P = 1$, we conclude $P = 1$ in S.

(c) *Dividing out one factor.* The element L_t is nilpotent with $L_t^{\kappa+1} = 0$ (every monomial of degree $\kappa+1$ in the x_i repeats a variable), so $1+L_t$ is invertible with $(1+L_t)^{-1} = \sum_{r=0}^{\kappa}(-1)^r L_t^r$. Therefore

$$\prod_{u\in\mathbb{Z}_M\setminus\{t\}}(1+L_u) \;=\; (1+L_t)^{-1}P \;=\; \sum_{r=0}^{\kappa}(-1)^r L_t^r.$$

Extracting $[x_K]$, only the term $r=\kappa$ contributes, and by the multinomial expansion (all $\kappa!$ orderings of the κ distinct variables)

$$[x_K]\,L_t^\kappa \;=\; \kappa!\prod_{i\in K}\omega^{\gamma_i t} \;=\; \kappa!\,\omega^{t\sigma(K)}.$$

Combining with (a) applied to $E = \mathbb{Z}_M \setminus \{t\}$ yields the claim. □

Corollary 4.4 (Entries are scaled characters)**.** *Let $K \subseteq I_1 \cup I_2$ be such that condition* (4.2) *holds modulo M for both weight systems $g|_K$ and $h|_K$. Then for all $s, t \in [M]$,*

$$J\big[T_K,\ (k+s, k+t)\big]\Big|_{P^\#} \;=\; \gamma_K\ \omega^{\,t\,\sigma_g(K)+s\,\sigma_h(K)}, \qquad \gamma_K := (m-1-\kappa)!\,(\kappa!)^2 \;>\; 0.$$

Proof. By Lemma 4.3 (applied to the weights $g|_K$) we have $\kappa \le M-1 \le m-1$, so the hypothesis of Lemma 4.2 is satisfied. Combining (4.1), Lemma 4.2, and Lemma 4.3 — once with weights $g|_K$ and excluded residue t, once with weights $h|_K$ and excluded residue s — gives

$$J\big[T_K, (k+s, k+t)\big]\Big|_{P^\#} = (m-1-\kappa)! \cdot \big((-1)^\kappa \kappa!\, \omega^{t\sigma_g(K)}\big) \cdot \big((-1)^\kappa \kappa!\, \omega^{s\sigma_h(K)}\big),$$

and the two signs $(-1)^\kappa$ multiply to $+1$. □

4.3 The arithmetic design

Encode subsets $B \subseteq I_1 \cup I_2$ by the pair of integers

$$x_B := \sum_{j:\,\alpha_j \in B} 2^{j-1}, \qquad y_B := \sum_{j:\,\beta_j \in B} 2^{j-1}, \qquad x_B, y_B \in [0, 2^{a-1}).$$

The map $B \mapsto (x_B, y_B)$ is a bijection onto $[0, 2^{a-1})^2$, and by construction

$$\sigma_g(B) = x_B + y_B, \qquad \sigma_h(B) = x_B + 2y_B.$$

(F1) The weight system g never vanishes. For nonempty B we have $1 \le \sigma_g(B) \le 2(2^{a-1}-1) = M-1$, so $\sigma_g(B) \not\equiv 0 \pmod M$. Thus (4.2) holds for g on every $K \subseteq I_1 \cup I_2$.

(F2) Bad sets for h are large and few. Call a nonempty $B \subseteq I_1 \cup I_2$ *bad* if $\sigma_h(B) \equiv 0 \pmod M$. Since $1 \le \sigma_h(B) \le 3(2^{a-1}-1) < 2M$, badness is equivalent to $x_B + 2y_B = M = 2^a - 1$. Writing $w(\cdot)$ for the binary digit sum and using subadditivity $w(p+q) \le w(p) + w(q)$ together with $w(2y) = w(y)$, a bad B satisfies

$$|B| \;=\; w(x_B) + w(y_B) \;\ge\; w(x_B + 2y_B) \;=\; w(2^a - 1) \;=\; a.$$

Moreover the solutions of $x + 2y = 2^a - 1$ with $x, y \in [0, 2^{a-1})$ are exactly $y \in [2^{a-2}, 2^{a-1})$ with $x = M - 2y$; hence there are exactly 2^{a-2} bad sets.

(F3) The good family is large. Let

$$\mathcal{G} := \{K \subseteq I_1 \cup I_2 :\ \text{no subset } B \subseteq K \text{ is bad}\}.$$

Every $K \in \mathcal{G}$ satisfies (4.2) for g (by (F1)) and for h (by the definition of $\mathcal{G}$). Each bad B is contained in exactly $2^{(2a-2)-|B|} \le 2^{a-2}$ subsets of $I_1 \cup I_2$, so by a union bound over the 2^{a-2} bad sets,

$$|\mathcal{G}| \;\ge\; 2^{2a-2} - 2^{a-2} \cdot 2^{a-2} \;=\; \tfrac{3}{4}\, 2^{2a-2}.$$

(F4) Distinct character pairs. The map $K \mapsto \big(\sigma_g(K) \bmod M,\ \sigma_h(K) \bmod M\big)$ is injective on subsets of $I_1 \cup I_2$. Indeed, $u := \sigma_g(K) = x_K + y_K$ lies in $[0, M-1]$, so u is determined by its residue; next, $\sigma_h(K) - \sigma_g(K) = y_K$ lies in $[0, 2^{a-1}) \subseteq [0, M)$, so y_K is the unique representative in $[0, M)$ of the residue $(\sigma_h(K) - \sigma_g(K)) \bmod M$; finally $x_K = u - y_K$, and (x_K, y_K) determines K.

4.4 Proof of the Key Lemma

Proof of Lemma 4.1. Consider the submatrix of the Jacobian matrix of the family $\{c_T\}_{T\subseteq[k]}$, evaluated at $P^{\#}$, with rows indexed by $\{T_K : K \in \mathcal{G}\}$ and columns indexed by $\{(k+s, k+t) : (s,t) \in [M]^2\}$. By (F1), (F3) and Corollary 4.4, the row corresponding to $K \in \mathcal{G}$ is

$$\gamma_K \cdot \chi_K, \qquad \chi_K(s,t) := \omega^{\,tA_K + sB_K}, \qquad (A_K, B_K) := \big(\sigma_g(K), \sigma_h(K)\big) \bmod M,$$

with $\gamma_K > 0$. As (s,t) ranges over $[M]^2 \cong \mathbb{Z}_M \times \mathbb{Z}_M$, the function χ_K is exactly the character $(s,t) \mapsto \omega^{tA_K+sB_K}$ of the group $\mathbb{Z}_M^2$, evaluated at all group elements. By (F4) the pairs (A_K, B_K), $K \in \mathcal{G}$, are pairwise distinct, so the χ_K are pairwise distinct characters of $\mathbb{Z}_M^2$; distinct characters of a finite abelian group are linearly independent over $\mathbb{C}$ (orthogonality: $\sum_{(s,t)} \chi(s,t)\overline{\chi'(s,t)} = M^2\delta_{\chi,\chi'}$). Hence the displayed submatrix has rank $|\mathcal{G}|$ over $\mathbb{C}$, and therefore some $|\mathcal{G}| \times |\mathcal{G}|$ minor of it is nonzero.

That minor is the evaluation at the point $P^{\#}$ of the corresponding $|\mathcal{G}| \times |\mathcal{G}|$ minor of the Jacobian matrix of the family $\{c_T\}$, which is a polynomial in the variables Z with integer coefficients. Being nonzero at a point of $\mathbb{Q}(\omega)^{|Z|}$, it is a nonzero polynomial, and hence nonzero over every characteristic-zero field $\mathbb{F}$ (Remark 1.2). By Lemma 1.1 the $|\mathcal{G}|$ polynomials c_{T_K}, $K \in \mathcal{G}$, are algebraically independent over $\mathbb{F}$. Using (F3) and $2^a > \frac{m+1}{2}$ we conclude

$$\tau_n(k) \;\ge\; |\mathcal{G}| \;\ge\; \tfrac{3}{4}\cdot 2^{2a-2} \;=\; \tfrac{3}{16}\,(2^a)^2 \;\ge\; \tfrac{3}{64}(m+1)^2 \;=\; \tfrac{3}{64}(n-k+1)^2 \;\ge\; \tfrac{3}{256}\,n^2,$$

the last inequality because $n-k+1 > n-k \ge n/2$ for $k \le n/2$. □

5 Assembly: proof of the Theorem

Proof of the Theorem. Let $n \ge \max(n_0, 2^{12})$ and set $k := \lceil 2\log_2 n \rceil + 4$; then $2\log_2 n + 4 \le k \le n/2$. Apply Corollary 2.2 to $f = \mathrm{perm}_n$ with the partition of the n^2 variables into the $p_0 = n\lfloor n/k \rfloor$ arc blocks of Section 3 together with the residual cells (taken, say, as singleton blocks), and with G the set of arc blocks. Each arc block consists of k cells in pairwise distinct rows and pairwise distinct columns, so by Lemma 3.1 and Lemma 4.1 its associated transcendence degree is $\tau_i = \tau_n(k) \ge \frac{3}{256}n^2$. Hence, using $\lfloor n/k \rfloor \ge n/(2k)$ (valid since $k \le n$) and $\frac{3}{256}n^2 - 2 \ge \frac{1}{100}n^2$ for all large n,

$$L(\mathrm{perm}_n) \;\ge\; \frac{p_0\big(\frac{3}{256}n^2 - 2\big)}{4} \;\ge\; \frac{n\cdot\frac{n}{2k}\cdot\frac{1}{100}n^2}{4} \;=\; \frac{n^4}{800\,k} \;=\; \Omega\!\left(\frac{n^4}{\log n}\right),$$

because $k = \Theta(\log n)$. Since $n^4/\log n = \omega(n^3)$, this proves the Theorem. □

For completeness of the picture, Ryser's formula

$$\mathrm{perm}_n = (-1)^n \sum_{\emptyset \ne S \subseteq [n]} (-1)^{|S|} \prod_{i=1}^{n} \sum_{j \in S} x_{ij}$$

is a depth-3 formula of size $O(n^2 2^n)$, so $L(\mathrm{perm}_n) = 2^{O(n)}$. Thus

$$\Omega\!\left(\frac{n^4}{\log n}\right) \;\le\; L(\mathrm{perm}_n) \;\le\; 2^{O(n)}.$$

6 Remarks

Remark 6.1 (Where the hypotheses are used). Characteristic zero enters three times: in Lemma 1.1 (the minimal-degree differentiation argument), in Remark 1.2 (so that evaluations in $\mathbb{Q}(\omega)$ certify nonvanishing over $\mathbb{F}$), and through the positivity of the factorial factors $\gamma_K = (m-1-\kappa)!\,(\kappa!)^2$ in Corollary 4.4: over a field of small positive characteristic these factorials could vanish and the certificate would be destroyed.

Remark 6.2 (The argument is genuinely permanental). For the determinant, the analogue of the factor $(m-1-\kappa)!$ in Lemma 4.2 is the determinant of an all-ones block, which vanishes as soon as $m-1-\kappa \geq 2$; the certificate would therefore vanish identically. This is not an artifact of the construction: by Jacobi's complementary-minor identity, determinantal minors satisfy polynomial relations which collapse the determinantal analogue of the family $\{c_T\}$ to transcendence degree $O(k^2)$. This is of course consistent with the existence of quasipolynomial-size formulas for $\det_n$, against which no bound of the above strength could hold. The positive sum over the $(m-1-\kappa)!$ completions inside the all-ones block — precisely where the determinant would cancel — is the permanental leverage of the argument.

Remark 6.3 (Sharpness within the method). The measure $\sum_{\text{blocks}} \mathrm{td}(\text{coefficient family})$ used in Corollary 2.2 is capped by $\sum_i \min\left(2^{|Y_i|},\, n^2\right)$, which over all partitions of the n^2 cells is $O(n^4/\log n)$, the optimum being attained by blocks of size $\Theta(\log n)$ consisting of cells in distinct rows and columns. The Theorem thus extracts the full strength of the Nechiporuk–Kalorkoti framework for the permanent: the classical $\Omega(n^3)$ bound corresponds to blocks at the break-even line ($\mathrm{td} = \Theta(n \cdot |Y_i|)$, e.g. whole diagonals or pairs of columns), whereas the blocks of Section 4 beat break-even by a factor $\Theta(n/\log n)$, which is the maximum possible.

Remark 6.4 (Tightness of the Key Lemma). The bound $\tau_n(k) \geq \frac{3}{64}(n-k+1)^2$ is within a constant factor of optimal, since $\tau_n(k) \leq n^2 - k - (2n-k-2)$. Indeed the coefficient family $\{c_T\}$ is preserved by the torus action

$$\left\{\, x_{ab} \mapsto \lambda_a \mu_b\, x_{ab} \;:\; \lambda_i \mu_i = 1 \ (i \leq k),\ \ \textstyle\prod_{a>k} \lambda_a \mu_a = 1 \,\right\},$$

which acts on the space of Z-variables with generic orbits of dimension $2n-k-2$, forcing the generic fibers of the coefficient map to have dimension at least $2n-k-2$. This upper bound is not needed for the Theorem.