hexagon:2611.01023v1

A constant gap for the maximal cp-rank

Yair Lavi

25 November 2026

Abstract

We prove that the maximal cp-rank p_N of completely positive matrices of order N satisfies $p_N = \frac{1}{2}N^2 - \frac{1}{2}N + O(1)$. More precisely, we prove $p_N \le \binom{N}{2}$ for every $N \ge 3$ using an analytic argument based on coordinate scaling. A computer-assisted construction gives $p_N \ge \binom{N}{2} - 2$ for $14 \le N \le 208$ and for every larger $N \not\equiv 1 \pmod 4$. In the remaining residue class, it gives $p_N \ge \binom{N}{2} - 3$.

Keywords. Completely positive matrix; cp-rank; copositive matrix; exposed face.
2020 Mathematics Subject Classification. 15B48, 15A23.

1 Introduction and main theorem

We determine the maximal cp-rank up to an additive constant independent of the matrix order.

A real symmetric matrix A is **completely positive** if $A = BB^{\mathsf{T}}$ for an entrywise nonnegative real matrix B with finitely many columns. Its **cp-rank**, $\operatorname{cpr}(A)$, is the least possible number of columns; $\operatorname{cpr}(0) = 0$. Write CP_N for the cone of such matrices and

$$p_N = \max_{A \in \mathrm{CP}_N} \operatorname{cpr}(A).$$

The maximum exists because cp-rank takes nonnegative integer values and conic Carathéodory gives a finite uniform upper bound. Our main result is

$$p_N = \tfrac{1}{2}N^2 - \tfrac{1}{2}N + O(1) \qquad (N \to \infty).$$

We call $\binom{N}{2} - p_N$ the **cp-rank gap**. Thus this gap stays bounded by a constant independent of N. The following bounds give explicit constants.

Theorem 1.1 (upper bound)**.** *For every integer $N \ge 3$,*

$$p_N \le \binom{N}{2}.$$

Theorem 1.2 (constant gap)**.** *For every integer $N \ge 14$,*

$$0 \le \binom{N}{2} - p_N \le \begin{cases} 3, & N \ge 209 \text{ and } N \equiv 1 \pmod 4, \\ 2, & \text{otherwise.} \end{cases} \tag{1.1}$$

The theorem bounds the unknown maximum; it does not assert equality with either bound in (1.1).

The starting order $N = 14$ comes from the consecutive finite coverage established in Appendix D; we do not claim that this cutoff is sharp. At $N = 15, 17, 19$, the finite-order constructions give the stronger bounds $p_{15} \geq 104$, $p_{17} \geq 135$, $p_{19} \geq 170$, each equal to $\binom{N}{2} - 1$.

Write $\mathbb{R}^N_+ = \{x \in \mathbb{R}^N : x_i \geq 0\}$, and write $x \geq 0$ for membership in this set. A real symmetric matrix S is **copositive** if $x^\mathsf{T} S x \geq 0$ for every $x \geq 0$. A nonzero $u \geq 0$ with $u^\mathsf{T} S u = 0$ is a **zero** of S; its **zero ray** is $\mathbb{R}_+ u$. The lower bound is obtained from copositive matrices with finitely many such rays.

The Drew–Johnson–Loewy conjecture $p_N = \lfloor N^2/4 \rfloor$ for $N \geq 4$ [1] was disproved by Bomze, Schachinger and Ullrich [2, 3]. Their constructions showed that p_N has leading term $N^2/2$. Schachinger [4, Theorem 7.4] reduced the difference between the then standard upper and lower bounds to $O(N \log\log N)$. The construction in [5] gives

$$p_N \geq \frac{N(N-3)}{2} - \mathbf{1}_{2|N} = \binom{N}{2} - N - \mathbf{1}_{2|N} \qquad (N \geq 5),$$

where $\mathbf{1}_{2|N}$ is 1 for even N and 0 for odd N. Together with the previously known upper bound, this yields $p_N = \frac{1}{2}N^2 + O(N)$. None of those estimates fixed the coefficient of N in p_N. For $N \geq 7$, Theorem 1.1 sharpens the previously used general bound $p_N \leq \binom{N+1}{2} - 4$, valid for $N \geq 5$ [6, Corollary 5.2]; the stronger individual bounds at orders 5 and 6 are recalled below. Combined with the new lower construction, it fixes the linear coefficient and reduces the remaining uncertainty to $O(1)$.

For $1 \leq N \leq 4$, one has $p_N = N$; see [6] and the references therein. The equality $p_5 = 6$ was proved by Shaked-Monderer, Bomze, Jarre and Schachinger [7], building on Loewy and Tam's [8] bound for matrices with a zero entry. For order six, Shaked-Monderer's zero-entry reduction [9, Corollary 3.2], with Dickinson's correction [10], combined with [6, Theorem 5.3], yields $p_6 \leq 2p_5 = 12$. Holden [11] has posted a proposed proof that $p_6 = 9$. The inequality $p_N \leq \binom{N}{2}$ at $N = 3, 4$ is also recorded in [6, Remark 5.1].

The **cp-plus-rank** of A is the least number of columns in a factorization $A = BB^\mathsf{T}$ with B entrywise strictly positive [12]. Let p^+_N denote its maximum finite value among matrices of order N. Bomze, Dickinson and Still [12] established $p_N \leq p^+_N \leq p_N + 1$, so the lower bounds of Theorem 1.2 also hold for p^+_N. A recent preprint of Ferudun [13, Theorem 1] proves $p^+_N = p_N$, which also transfers the upper bound of Theorem 1.1.

Section 2 proves the upper bound by coordinate scaling. For the lower bound, Section 3 explains how finitely many copositive zero rays force large cp-rank through the dimension spanned by their outer products, and gives the perturbation tool used to produce these rays.

Sections 4 and 5 start from the rank-three Gram matrix of points on two parallel regular polygons. A perturbation yields a copositive matrix with prescribed four-coordinate zero rays. We call this matrix the parent and write n for its order, reserving N for the final order. The zero outer products span $\binom{n}{2} - 2$ dimensions when $n \equiv 2 \pmod 4$ and $n \geq 202$.

Section 6 adjoins coordinates whose base Gram vectors lie on the axis perpendicular to the polygon planes; these are the axial columns. Adding one and then two opposite axial columns preserves the two-unit gap, while duplicating one gives a three-unit gap at order $n + 3$. Appendix D supplies exact finite certificates for orders 14 through 201 and a separate extension at order 205. The latter uses a new base Gram vector parallel to the polygon planes, which we call a horizontal extension.

At order 205, the general duplication step guarantees only $p_{205} \geq \binom{205}{2} - 3$; the separate horizontal extension improves this to $\binom{205}{2} - 2$.

The paper contains all mathematical proofs and reductions, with the detailed potential, character-rank, axial and finite-order construction arguments in Appendices A–D. Appendix E identifies the exact arithmetic statements checked by each program and their roles in the proofs. The accompanying computational supplement contains code, certificate data, computational results and reproduction instructions.

2 An upper bound by coordinate scaling

For a completely positive matrix A orthogonal to a copositive matrix S with $k \geq 2$ positive diagonal entries, [6, Theorem 5.1] gives $\operatorname{cpr}(A) \leq \binom{N+1}{2} - k$. Thus a positive-diagonal supporting matrix already gives the bound $\binom{N}{2}$. The coordinate-scaling argument below handles zero diagonal entries and gives the graph bounds needed for Theorem 1.1.

Ito and Lourenço [14] bound the number of extreme-ray summands needed in a pointed closed convex cone by one less than the length of a longest chain of nonempty faces. Nishijima [15] proves that the longest chain of nonempty faces of CP_N has length $\binom{N+1}{2} + 1$. Consequently, applying this general estimate directly to CP_N gives $\operatorname{cpr}(A) \leq \binom{N+1}{2}$. Our argument also uses independent coordinate scaling, which supplies the additional saving counted below. The following proof is self-contained.

Independent scaling can save one summand per scalable coordinate. We work with a restricted set of vectors because the induction passes to an exposed face.

For a set $V \subseteq \mathbb{R}^{r+t}_+$, call its elements **permitted vectors** and their outer products **atoms**, and put

$$K(V) = \operatorname{cone}\{vv^{\mathsf{T}} : v \in V\}, \qquad D(V) = \dim \operatorname{span} K(V).$$

Write $v = (z, y)$, with $z \in \mathbb{R}^r$ and $y \in \mathbb{R}^t$. Regard all z coordinates as a single **core** vertex and each y_i as a separate vertex. The **grouped graph** joins two vertices if some permitted vector is nonzero in both groups. A y coordinate is **scalable** if multiplying it by any nonnegative number preserves membership in V.

We use $\langle S, A\rangle = \operatorname{tr}(SA)$ for the matrix inner product. The relative interior of a cone is taken in its own linear span. A **supporting functional** at a boundary point B of K is a linear functional nonzero on $\operatorname{span} K$, nonnegative on K, and zero at B. Its zero set in K is an **exposed face**. A symmetric matrix represents such a functional by the trace inner product.

The proof subtracts a permitted atom supported entirely in the core until it reaches the boundary. Coordinates with positive supporting diagonal then join the core and supply independent equations. The other coordinates remain scalable, and one additional equation pays for restoring the subtracted atom.

Lemma 2.1 (coordinate scaling)**.** *Let $V \subseteq \mathbb{R}^{r+t}_+$ be closed and stable under multiplication by nonnegative scalars. Suppose that its t designated y coordinates are independently scalable, that some permitted vector has nonzero core part, and that its grouped graph is connected. Then every matrix in $K(V)$ is a sum of at most $D(V) - t$ permitted outer products.*

Proof. The permitted unit vectors form a compact set. Since each of their outer products has trace one, $K(V)$ is closed. For each fixed m, the set of sums of at most m permitted outer products is closed as well: in a convergent sequence of such sums, trace bounds every column norm; pad by zero columns and take a subsequence. Thus it is enough to prove the assertion for matrices in the relative interior of $K(V)$.

We induct on $D = D(V)$, allowing the core and the number of scalable coordinates to change. At $t = 0$ the claim is conic Carathéodory. For $t > 0$, connectedness says that every y_i occurs in a permitted vector. Scaling that coordinate to infinity and normalizing gives the axis $e_i \in V$ by closedness. Turning all y coordinates off in a vector with nonzero core gives a permitted vector $w = (z_*, 0) \neq 0$.

Fix $A \in \operatorname{relint} K(V)$. Subtract the largest possible $\lambda w w^{\mathsf{T}}$ while staying in $K(V)$, and call the endpoint B. Then $\lambda > 0$ and B is on the relative boundary. A nonzero relative supporting functional S at B, represented by a symmetric matrix, obeys $v^{\mathsf{T}} S v \geq 0$ for $v \in V$, $\langle S, B \rangle = 0$, and

$$\langle S, A \rangle = \lambda w^{\mathsf{T}} S w > 0. \tag{2.1}$$

Each scalable diagonal S_{ii} is nonnegative because $e_i \in V$. We may erase the row and column at each scalable index where $S_{ii} = 0$. To see this, let Q_i contain just the off-diagonal part of that row and column. If v_i is replaced by $a v_i$, the quadratic form $v \mapsto v^{\mathsf{T}} S v$ becomes $v^{\mathsf{T}}(S - Q_i)v + a\, v^{\mathsf{T}} Q_i v \geq 0$ for all $a \geq 0$. Setting $a = 0$ and letting $a \to \infty$ shows that both $S - Q_i$ and Q_i are nonnegative on $K(V)$. Both vanish at B: their sum does, and both are nonnegative. Since Q_i vanishes at $w w^{\mathsf{T}}$, it also vanishes at A. Relative interior then forces Q_i to vanish on all of $\operatorname{span} K(V)$. Repeating the operation gives a symmetric matrix R representing the same functional as S on $\operatorname{span} K(V)$, with the entire row and column zero at every scalable index where $R_{ii} = 0$.

Let I be the remaining scalable indices, with $R_{ii} > 0$, and put $k = |I|$. Put $V' = \{v \in V : v^{\mathsf{T}} R v = 0\}$. This set is closed and homogeneous, as the zero set of a continuous homogeneous quadratic form within V. Then

$$K(V') = \{X \in K(V) : \langle R, X \rangle = 0\},$$

because each nonnegative summand $v^{\mathsf{T}} R v$ in a factorization must vanish when their sum does. In particular $B \in K(V')$. For $v \in V'$, varying a positive coordinate v_i in both directions shows, for $i \in I$,

$$v_i (R v)_i = 0. \tag{2.2}$$

This is the restricted-cone analogue of the column-orthogonality relation in [6]. Here R need only be nonnegative on permitted vectors, so the direct argument is needed. The k linear functionals $X \mapsto (RX)_{ii}$, $i \in I$, vanish on $K(V')$ because $(R v v^{\mathsf{T}})_{ii} = v_i (Rv)_i$. They are independent on $\operatorname{span} K(V)$: on the axis outer products $e_i e_i^{\mathsf{T}}$ their evaluation matrix is diagonal with positive diagonal. The supporting functional $X \mapsto \langle R, X \rangle$ gives one more independent equation, because it is positive at $w w^{\mathsf{T}}$ by (2.1), while each of these k functionals vanishes there. Hence $\dim \operatorname{span} K(V') \leq D - k - 1$.

Move the k coordinates of I into the core. The remaining $t - k$ coordinates still scale independently in V' because their rows and columns in R are zero. To check the graph hypothesis, an old grouped edge has a permitted vector positive in both endpoint groups and hence some positive matrix entry between its groups. That entry functional is nonnegative on $K(V)$ and nonzero on its span, so it is strictly positive at $A \in \operatorname{relint} K(V)$. Every such edge involves an old scalable coordinate; subtracting the pure-core atom leaves the entry unchanged in B. In a factorization of B into atoms from V', at least one vector therefore witnesses the edge. Contracting the moved vertices into the core preserves connectedness. Since $t > 0$, the old core has a neighbor, and its preserved edge also witnesses a nonzero new core. All the induction hypotheses are satisfied. Induction factors B in at most $(D - k - 1) - (t - k)$ columns. Restoring $\sqrt{\lambda} w$ uses one more, for a total of at most $D - t$. Closedness extends the result from relative interior to all of $K(V)$. □

For a simple graph G on N vertices, define

$$\mathrm{CP}(G) = \{A \in \mathrm{CP}_N : A_{ij} = 0 \text{ for every nonedge } ij,\ i \neq j\}.$$

The **support graph** of A has edge ij exactly when $A_{ij} > 0$, $i \neq j$. By a permitted vector for $\mathrm{CP}(G)$, we mean a nonnegative vector whose positive support is contained in a clique of G, a set of pairwise adjacent vertices.

Corollary 2.2 (graph bound)**.** *If $A \in \mathrm{CP}(G)$, then*

$$\mathrm{cpr}(A) \leq |E(G)| + c(G), \tag{2.3}$$

where $c(G)$ is the number of connected components, including isolated vertices.

Proof. In each connected component G_0 choose one vertex as the core; all other coordinates are scalable. The axes e_i and edge vectors $e_i + e_j$ show that the outer-product span has dimension $|V(G_0)|+|E(G_0)|$. Lemma 2.1 gives $|E(G_0)|+1$ summands in that component. Sum over components. □

Hannah and Laffey [16] give a bound in terms of the rank and the number of zeros in a nonsingular principal submatrix of full rank. For nonsingular A with support graph G, this becomes $N + |E(G)|$. Shaked-Monderer [17] gives a new proof and an improvement for matrices with a zero entry. Corollary 2.2 gives $|E(G)| + c(G)$ without a rank assumption, saving $N - c(G)$ over the displayed nonsingular bound.

For a connected triangle-free graph G that is the actual support graph of A, sharper classical results are available: $\mathrm{cpr}(A) = \mathrm{rank}(A)$ when G is a tree, by Berman and Hershkowitz [18], and $\mathrm{cpr}(A) = |E(G)|$ otherwise, by Drew, Johnson and Loewy [1]; see also [19, Proposition 2.4]. Bound (2.3) is sharp for nonsingular matrices supported on trees and applies without restrictions on the triangles or cycles of G.

Lemma 2.3 (boundary graph bound)**.** *Let G be connected. If A lies on the relative boundary of $\mathrm{CP}(G)$ and has positive diagonal and positive entries on every edge of G, then $\mathrm{cpr}(A) \leq |E(G)|$.*

Proof. Let V consist of the nonnegative vectors whose positive support is contained in a clique of G. Then $K(V) = \mathrm{CP}(G)$ and its span has dimension $N + |E(G)|$. Represent a nonzero relative supporting functional at A by a symmetric matrix R whose nonedge entries are zero. It is nonnegative on $K(V)$ and satisfies $\langle R, A\rangle = 0$. Thus $A \in K(V')$, where $V' = \{v \in V : v^{\mathsf{T}} R v = 0\}$.

The axes show $R_{ii} \geq 0$. If $R_{ii} = 0$, testing nonnegative vectors on an incident edge gives $R_{ij} \geq 0$. As in Lemma 2.1, let Q_i be the off-diagonal part of row and column i. Scaling that coordinate shows that both Q_i and $R - Q_i$ are nonnegative functionals on $K(V)$. Their values at A sum to zero and hence both vanish. Since every edge entry A_{ij} is positive and every corresponding R_{ij} is nonnegative, $\langle Q_i, A\rangle = 0$ forces all those coefficients to be zero. The full row and column of R at i are therefore zero. At least one diagonal is positive, since otherwise R would vanish on the graph span.

Let $I = \{i : R_{ii} > 0\}$ and $k = |I| \geq 1$. Coordinate variation at $v \in V'$ gives $v_i(Rv)_i = 0$ for $i \in I$. The corresponding linear functionals $X \mapsto (RX)_{ii}$ are independent on the original span, by evaluation on the axis outer products. Consequently

$$D(V') \leq N + |E(G)| - k.$$

Here $\langle R, X\rangle = \sum_{i\in I}(RX)_{ii}$, so the supporting functional is already the sum of these equations and is not an additional constraint.

Take the coordinates in I as the core. The other $N - k$ coordinates remain independently scalable in the closed homogeneous set V', because the corresponding rows of R are zero. Positivity of the diagonal of A witnesses a nonzero core, and positivity of its edge entries witnesses every edge

in a factorization using V'. Contracting the core thus preserves connectedness. Lemma 2.1 gives at most

$$D(V') - (N - k) \le |E(G)|$$

summands. □

Maximal cp-rank is known to be attained at a positive-definite boundary matrix [7, Theorem 3.4]. The next proof includes a direct shear argument, also used by Holden [11, Lemma 3.4], that transports a boundary factorization back to the original matrix.

Proof of Theorem 1.1. If A has a zero entry, use its support graph G. A connected graph that is not complete has $|E(G)| \le \binom{N}{2} - 1$, so Corollary 2.2 suffices. If G has $c \ge 2$ components, at least c edges are missing when $N \ge 3$ (for two components at least $N - 1 \ge 2$ are missing; for three or more use the cross-component pairs). Thus $|E(G)| + c \le \binom{N}{2}$.

Suppose next that A is entrywise positive and on the boundary of CP_N. It has the complete support graph, so Lemma 2.3 gives $\mathrm{cpr}(A) \le |E(K_N)| = \binom{N}{2}$.

Finally, let A be interior, choose $i \ne j$, and set $L_s = I - se_ie_j^{\mathsf{T}}$. The feasible set

$$\mathcal{F} = \{s \ge 0 : L_sAL_s^{\mathsf{T}} \in \mathrm{CP}_N\}$$

contains zero, is closed, and is bounded by $s \le A_{ij}/A_{jj}$, since $(L_sAL_s^{\mathsf{T}})_{ij} = A_{ij} - sA_{jj} \ge 0$ for $s \in \mathcal{F}$. Thus it has a maximum s_*. Put $B = L_{s_*}AL_{s_*}^{\mathsf{T}}$. If B were interior, continuity would allow a larger feasible parameter, contradicting maximality. Hence B is boundary. If it has a zero entry, the first case of the proof applies; otherwise the entrywise-positive boundary case applies. In either case $B = CC^{\mathsf{T}}$ for some $C \ge 0$ with at most $\binom{N}{2}$ columns. Since $L_{s_*}^{-1} = I + s_*e_ie_j^{\mathsf{T}} \ge 0$, the factor $L_{s_*}^{-1}C$ has the same number of columns and factors A. □

3 From copositive zeros to a lower bound

If a copositive matrix S has finitely many zero rays, choose one generator u_i from each and set

$$h(S) = \dim_{\mathbb{R}} \operatorname{span}\{u_iu_i^{\mathsf{T}}\}.$$

We call $h(S)$ the **outer-product span dimension**. The following finite-zero face lemma is due to Bomze, Schachinger and Ullrich [3, Lemma 2.2 in the author preprint]. We include its short proof for completeness.

Lemma 3.1 (finite zero rays)**.** *If S is copositive of order N and has finitely many zero rays, then $p_N \ge h(S)$.*

Proof. The exposed face $F = \{A \in \mathrm{CP}_N : \langle S, A\rangle = 0\}$ is generated by those finitely many outer products. In any nonnegative factorization of a matrix in F, every nonzero column u satisfies $u^{\mathsf{T}}Su = 0$ and therefore lies on a listed zero ray. Sums using fewer than $h(S)$ rays lie in a finite union of proper linear subspaces of $\operatorname{span} F$. That union cannot contain the relative interior of F. A matrix outside it needs at least $h(S)$ columns. □

Lemma 3.2 (duplication of a coordinate)**.** *Let S be a copositive matrix of order N with finitely many zero rays, and fix a coordinate i. Put*

$$\rho_i = \dim \operatorname{span}\{u \in \mathbb{R}^N_+ : u^{\mathsf{T}}Su = 0,\ u_i > 0\}.$$

Replace coordinate i by two copies i_1, i_2. For a vector x in the enlarged space let Cx collapse the copies by $(Cx)_i = x_{i_1} + x_{i_2}$ and leave the other coordinates unchanged. For any $\lambda_{\mathrm{dup}} > 0$, define S' by

$$x^{\mathsf{T}} S' x = (Cx)^{\mathsf{T}} S (Cx) + 2\lambda_{\mathrm{dup}} x_{i_1} x_{i_2}.$$

Then S' is copositive, has finitely many zero rays, and

$$h(S') = h(S) + \rho_i.$$

Proof. For $x \geq 0$ both terms defining the quadratic are nonnegative. Thus a nonzero x is a zero precisely when Cx is a parent zero and at most one copy is positive. Each parent ray through i gives two rays, obtained by putting its ith coordinate entirely in either copy; a ray avoiding i gives one.

Let J embed the old space by always choosing the first copy, so CJ is the identity. The fixed-choice outer products span an $h(S)$-dimensional subspace. All remaining outer products are obtained by adding differences

$$D(u) = u^{(2)}(u^{(2)})^{\mathsf{T}} - u^{(1)}(u^{(1)})^{\mathsf{T}}.$$

Each difference is a fixed injective linear image of $u_i u$: the second-copy row records $u_i u_j$ at every other old coordinate and u_i^2 on its diagonal, and the first-copy entries have the opposite signs. Consequently these differences span ρ_i dimensions. They lie in the kernel of $X \mapsto CXC^{\mathsf{T}}$, whereas this map is injective on the fixed-choice subspace. The two subspaces therefore form a direct sum, proving the formula. □

We will start with a rank-three positive semidefinite matrix P whose nonnegative kernel cone is $Z = \ker P \cap \mathbb{R}^n_+$. A **circuit** is a nonzero vector $z \in Z$ whose support is minimal under inclusion among supports of nonzero vectors in Z. Its ray $\mathbb{R}_+ z$ is a **circuit ray**. Circuit rays are exactly the extreme rays of Z; the bases used below have many four-point circuit rays. A symmetric perturbation Q is chosen so that $q(z) = z^{\mathsf{T}} Q z$ is nonnegative on Z and vanishes on precisely the wanted rays. For an index set I, write $P[I]$ for the principal submatrix indexed by I. We write $\mathbf{1}$ for the all-ones vector in the ambient space, so $\mathbf{1}^{\perp} = \{x : \sum_i x_i = 0\}$. We call q the **first-order form** and a ray $\mathbb{R}_+ z$ with nonzero $z \in Z$ and $q(z) = 0$ a **first-order contact**. Such a contact does not by itself give a zero of $P + \varepsilon Q$: the zero must be continued and the quadratic form of the perturbed matrix must remain nonnegative on the entire nonnegative orthant. Dickinson and Hildebrand [20] characterize feasible linear perturbations of a copositive matrix in terms of its zeros. Manainen, Seliugin, Tarasov and Hildebrand [21] construct related local families of extreme copositive matrices near matrices associated with COP-irreducible graphs of stability number three. The lemma below uses a related simplex-minimizer argument, together with analytic corrections, to continue selected zero rays and transfer a lower bound on the dimension of their outer-product span.

Lemma 3.3 (perturbation lemma)**.** *Let $P \succeq 0$, $K = \ker P$, $Z = K \cap \mathbb{R}^n_+$, and let Q be symmetric. Let Γ be a group of permutation matrices satisfying $\gamma P \gamma^{\mathsf{T}} = P$ and $\gamma Q \gamma^{\mathsf{T}} = Q$ for every $\gamma \in \Gamma$. Let $\mathcal{E}$ be a linear space of symmetric matrices each fixed by every $\gamma \in \Gamma$: $\gamma E \gamma^{\mathsf{T}} = E$. Assume:*

(1) $q \geq 0$ on Z, and its zero rays are precisely finitely many distinct rays $\mathbb{R}_+ z_i$, with nonzero representatives z_i;

(2) for each $I_i = \operatorname{supp} z_i$, the positive semidefinite block $P[I_i]$ has a one-dimensional kernel;

(3) for each i, there exists w_i supported on I_i such that $P[I_i](w_i|_{I_i}) = -(Qz_i)|_{I_i}$, and $g_i := Qz_i + Pw_i$ satisfies $g_i(j) > 0$ for every $j \notin I_i$;

(4) if $i_1, \dots, i_r$ represent the Γ-orbits of the rays, the map $\mathcal{E} \to \mathbb{R}^r$, $E \mapsto (z_{i_s}^\mathsf{T} E z_{i_s})_{s=1}^r$, is surjective.

Then, for all sufficiently small $\varepsilon > 0$, there is $E_\varepsilon \in \mathcal{E}$, analytic with $E_0 = 0$, for which $S_\varepsilon = P + \varepsilon(Q + E_\varepsilon)$ is copositive and has exactly the zero rays $\mathbb{R}_+ u_i(\varepsilon)$, where u_i is analytic, $u_i(0) = z_i$, $\operatorname{supp} u_i(\varepsilon) = I_i$, and $\langle u_i(\varepsilon), z_i \rangle = \|z_i\|^2$. Moreover, with $t_i = P w_i$,

$$h(S_\varepsilon) \geq \dim \operatorname{span}\{(z_i + t_i)(z_i + t_i)^\mathsf{T}\}. \tag{3.1}$$

We call g_i the **first-order entering gradient**. The conditions $g_i(j) > 0$ for $j \notin I_i$ are the **entering inequalities**; entering refers to adding mass outside the zero support. The analytic deformation $u_i(\varepsilon)$ is the **continuation** of the first-order contact.

Proof. The choice in (3) is well-defined: by (2), $\ker P[I_i] = \operatorname{span}\{z_i|_{I_i}\}$, and the right side is orthogonal to this kernel because $z_i^\mathsf{T} Q z_i = 0$. Thus the equation is solvable. Two supported solutions differ by a multiple of z_i; since $P z_i = 0$, they give the same g_i.

The least eigenvalue of $P[I_i] + X[I_i]$ is analytic near $X = 0$, because its zero eigenvalue is simple. Its derivative at zero is $z_i^\mathsf{T} X z_i / \|z_i\|^2$. Write this eigenvalue as $F_i(X)$. For Γ-invariant X, the eigenvalue equation is identical within each orbit. Since $z_i^\mathsf{T} Q z_i = 0$, the quotient $F_i(\varepsilon(Q+E))/\varepsilon$ extends analytically to $\varepsilon = 0$, with value $z_i^\mathsf{T} E z_i / \|z_i\|^2$. Its derivative in E is onto on the orbit representatives by (4). The implicit function theorem on a complement of its kernel gives an analytic $E_\varepsilon = O(\varepsilon)$ making every selected eigenvalue zero. The resulting kernel vectors $u_i(\varepsilon)$ are positive on their original supports and converge to z_i. Differentiating $S_\varepsilon[I_i](u_i(\varepsilon)|_{I_i}) = 0$ at $\varepsilon = 0$ and using $E_0 = 0$ gives

$$P[I_i](u_i'(0)|_{I_i}) = -(Q z_i)|_{I_i}.$$

Comparison with condition (3), together with $\ker P[I_i] = \operatorname{span}\{z_i|_{I_i}\}$ and the fact that both vectors are supported on I_i, shows that $u_i'(0) - w_i \in \operatorname{span}\{z_i\}$. Since $P z_i = 0$, it follows that $P u_i'(0) = P w_i$. Consequently $S_\varepsilon u_i(\varepsilon) = \varepsilon g_i + O(\varepsilon^2)$, which is strictly positive off the support for small positive ε.

To exclude other zeros and negative values, minimize $x^\mathsf{T} S_\varepsilon x$ on the nonnegative simplex $\sum x_i = 1$. Suppose a sequence of minimizers has value at most zero as $\varepsilon \downarrow 0$. Compactness gives a limit x_*; the leading P-term forces $x_* \in Z$, and the first-order inequality forces $q(x_*) \leq 0$. Thus x_* is one of the normalized z_i. For a sufficiently late minimizer x, its support contains I_i. The simplex optimality equation is $S_\varepsilon x = v\mathbf{1} + \nu$, with $\nu \geq 0$ and $\nu = 0$ on the support of x. Pair it with $u_i(\varepsilon)$: since $S_\varepsilon u_i(\varepsilon) \geq 0$, the value v must be nonnegative. It was at most zero, so $v = 0$; strict positivity of $S_\varepsilon u_i(\varepsilon)$ off I_i then confines x to I_i. The simple kernel of $S_\varepsilon[I_i]$ makes x that listed ray. This proves both copositivity and completeness.

For the outer-product span dimension, let $\operatorname{pr}_K, \operatorname{pr}_{K^\perp}$ denote the orthogonal projections and put

$$T_\varepsilon = \operatorname{pr}_K + \varepsilon^{-1} \operatorname{pr}_{K^\perp}.$$

For $\varepsilon > 0$ this map is invertible, so congruence $X \mapsto T_\varepsilon X T_\varepsilon^\mathsf{T}$ preserves the dimension of the span of the zero outer products. The derivative identity proved above gives

$$T_\varepsilon u_i(\varepsilon) \longrightarrow z_i + \operatorname{pr}_{K^\perp} w_i.$$

Choose fixed coordinates on the space of symmetric matrices. A maximal nonzero minor of the limiting outer-product coordinate matrix remains nonzero for all sufficiently small $\varepsilon > 0$, by continuity. Hence the actual zero outer-product span has dimension at least that of these limiting outer products. Finally, $\operatorname{pr}_K + P \operatorname{pr}_{K^\perp}$ is invertible, since P is positive definite on $K^\perp$. It sends the limiting vectors to $z_i + P w_i$; its congruence also preserves span dimension. This proves (3.1). □

We will compute the dimension of the span of the limiting outer products in (3.1). The lemma transfers this dimension to the outer-product span of the actual zeros.

4 Four-point zeros in orders $4m+2$

Fix $n = 2d \equiv 2 \pmod 4$ with odd $d \geq 101$ and write $\theta = \pi/n$. The rank-three base can be represented by points on two parallel copies of a regular d-gon. In cyclic indexing its Gram matrix is

$$P_{ij} = (-1)^{i-j}\big(\kappa + (1-\kappa)\cos(2\theta(i-j))\big), \qquad 0 < \kappa < 1. \tag{4.1}$$

For $f \in \mathbb{Z}/n\mathbb{Z}$, define $a_f \in \mathbb{C}^n$ by $a_f(r) = e^{2\pi \mathrm{i} fr/n}$, $r = 0, \dots, n-1$. We call f the frequency of a_f. Over $\mathbb{C}$, the range of P is spanned by a_d, a_{d-1}, a_{d+1}. The Gram columns are

$$\mathbf{b}_i = (-1)^i\big(\sqrt{\kappa}, \sqrt{1-\kappa}\cos(2\theta i), \sqrt{1-\kappa}\sin(2\theta i)\big).$$

Lemma 4.1 (four-point circuits)**.** *For the Gram columns in* (4.1)*, at any even order $n \geq 4$ and with $0 < \kappa < 1$, any three distinct columns are linearly independent. Every circuit of $Z = \ker P \cap \mathbb{R}^n_+$ has exactly four positive coordinates. A four-point support $I = \{r_0 < r_1 < r_2 < r_3\} \subseteq \{0, \dots, n-1\}$ admits a nonnegative dependence if and only if its four cyclic gaps*

$$r_1 - r_0, \quad r_2 - r_1, \quad r_3 - r_2, \quad n + r_0 - r_3$$

are odd. In that case its nonnegative kernel is the ray generated by

$$z_I(r) = \prod_{\ell \in I \setminus \{r\}} \sin(\theta|r-\ell|)^{-1} \quad (r \in I), \qquad z_I(r) = 0 \quad (r \notin I), \qquad \theta = \pi/n. \tag{4.2}$$

Proof. Set $\omega_r = e^{2\pi \mathrm{i} r/n}$. An invertible change of rows over $\mathbb{C}$ takes the Gram columns to $(-1)^r(\omega_r^{-1}, 1, \omega_r)^{\mathsf{T}}$. Multiplying column r by $(-1)^r\omega_r$ gives the Vandermonde column $(1, \omega_r, \omega_r^2)^{\mathsf{T}}$. Its nodes are distinct, so every three columns are independent. Supports of size at most three are therefore independent. On a support of size greater than four, the kernel has dimension at least two, so a nonproportional kernel perturbation can be continued within the nonnegative orthant until a coordinate vanishes, leaving a nonzero vector with smaller support. Thus every circuit has four positive coordinates. On four columns the kernel is generated by

$$\alpha_r = \frac{1}{\prod_{\ell \in I \setminus \{r\}} (\omega_r - \omega_\ell)}.$$

Indeed, Lagrange interpolation of $1, x, x^2$ at the four nodes gives $\sum_{r \in I} \omega_r^k \alpha_r = 0$ for $k = 0, 1, 2$ by comparison of the cubic coefficients. Returning to the original columns gives coefficients $(-1)^r\omega_r\alpha_r$. Since

$$\omega_r - \omega_\ell = 2\mathrm{i}\, e^{\mathrm{i}\theta(r+\ell)} \sin(\theta(r-\ell)),$$

these coefficients have a common nonzero factor; after removing it, a real kernel generator is

$$\lambda_r = \frac{(-1)^r}{\prod_{\ell \in I \setminus \{r\}} \sin(\theta(r-\ell))}.$$

For $r = r_j$ its sign is $(-1)^{r_j+3-j}$. Thus consecutive coefficients have the same sign exactly when each of the three consecutive differences is odd. The closing gap is then odd because n is even. Conversely, a nonnegative dependence must have all four coefficients strictly positive, since any three columns are independent, and hence all four gaps must be odd. Taking absolute values of the signed generator gives (4.2). □

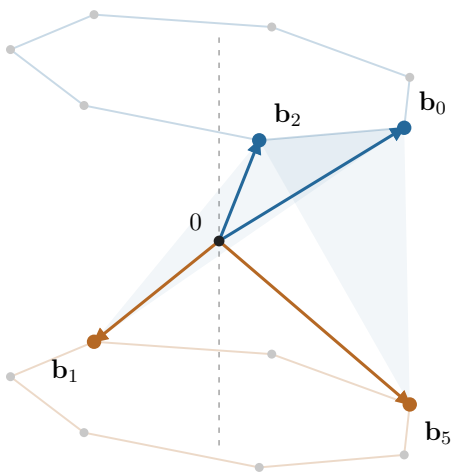

Figure 1: A positive dependence among four Gram columns at order $n = 14$: $I = \{0, 1, 2, 5\}$ has odd cyclic gaps $(1, 1, 3, 9)$. This illustrates the base kernel geometry; the corresponding zero rays are continued in the perturbed matrices.

Figure 1 illustrates (4.2) at the smaller order $n = 14$, and Example 4.2 gives the corresponding finite-order construction; the uniform construction here has $n \geq 202$. We call the four odd cyclic gaps, up to rotation and reversal, the **type** of I.

Given a real symmetric circulant matrix Q, we seek support types for which $z_I^{\mathsf{T}} Q z_I = 0$. Write its real n-periodic symbol as $Q(x)$, so $Q_{ij} = Q(i - j)$ and $Q(n - x) = Q(x)$. Put $s(x) = \sin(\theta x)$ and define the **potential associated with** Q, for $1 \leq x < n$, by

$$R(x) = \frac{(-1)^x Q(x) - Q(0)}{s(x)^2}, \qquad R(n - x) = R(x).$$

For a separation $x \in \{1, \ldots, n - 1\}$, its **folded distance** is $\min\{x, n - x\} \in \{1, \ldots, d\}$, since $n = 2d$.

For gaps (g_1, g_2, g_3, g_4) with sum n, the diagonal distances are $g_1 + g_2$ and $g_2 + g_3$. Put $\Pi_I = \prod_{i<j,\ i,j\in I} \sin(\theta|i - j|)$, the product of the six sine chords. To derive the functional, write $Q(x) = (-1)^x[Q(0) + s(x)^2 R(x)]$ for $x \neq 0$. Including the diagonal terms, the contribution of $Q(0)$ is $Q(0)(\sum_{r\in I}(-1)^r z_I(r))^2 = 0$, since the first Gram coordinate gives $\sum_{r\in I}(-1)^r z_I(r) = 0$. For distinct $a, b \in I$, write $I \setminus \{a, b\} = \{u, v\}$ with $u < v$. Formula (4.2) gives the complementary-chord identity

$$\Pi_I z_I(a) z_I(b) = \frac{s(v - u)}{s(|a - b|)}.$$

Consequently

$$\frac{\Pi_I}{2} z_I^{\mathsf{T}} Q z_I = \sum_{\substack{a,b\in I \\ a<b}} (-1)^{b-a} s(b - a) s(v - u) R(b - a),$$

where u, v denote the complementary pair in each summand. The four side pairs have odd index difference and the two diagonal pairs have even index difference. Grouping complementary pairs and using $s(n-x) = s(x)$ and $R(n-x) = R(x)$ gives

$$\begin{aligned}\tfrac{1}{2}\Pi_I z_I^{\mathsf{T}} Q z_I = {} & s(g_1+g_2)s(g_2+g_3)[R(g_1+g_2)+R(g_2+g_3)] \\ & - s(g_1)s(g_3)[R(g_1)+R(g_3)] - s(g_2)s(g_4)[R(g_2)+R(g_4)].\end{aligned} \tag{4.3}$$

Define the **type slack** $\mathcal{D}_g(R)$ as the right-hand side of (4.3); thus $\mathcal{D}_g(R) = (\Pi_I/2) z_I^{\mathsf{T}} Q z_I$. The first-order zero problem is a finite family of Ptolemy-type inequalities at each order.

Example 4.2 (finite-order construction at order 14)**.** To illustrate the circuit and type formulas, take the finite construction of Subsection D.1 at $n = 14$. Its five tight gap types and numbers of distinct supports are listed in Table 1. This finite construction is separate from the uniform four-band construction, whose proved range begins at $n = 202$.

Table 1: The order-14 construction: tight types and support counts.

Type	Number of supports
$(1,1,3,9)$	28
$(1,1,5,7)$	28
$(1,3,5,5)$	28
$(1,3,7,3)$	14
$(3,3,3,5)$	14

For instance, $I = \{0,1,2,5\}$ shown in Figure 1 has the first type. Writing $s_j = \sin(\pi j/14)$, formula (4.2) gives its nonzero coordinates as

$$\left(\frac{1}{s_1 s_2 s_5}, \frac{1}{s_1^2 s_4}, \frac{1}{s_1 s_2 s_3}, \frac{1}{s_3 s_4 s_5}\right).$$

These support counts give 112 continued zero rays. We return to their outer-product span in Section 5. ◇

We now return to the uniform construction for $n = 2d \geq 202$, with d odd. Take the following eight **anchor types**, the fixed patterns used to start the band equations:

$$\begin{gathered}(1,1,d-4,d+2),\ (1,1,d-2,d),\ (1,3,d-2,d-2),\ (1,3,d,d-4), \\ (1,3,d+2,d-6),\ (1,5,d-2,d-4),\ (1,d-6,11,d-6),\ (5,5,d-6,d-4),\end{gathered} \tag{4.4}$$

and four bands $(g_1, g_2, g_3, n-g_1-g_2-g_3)$, with g_2 running over all odd integers in the closed intervals of Table 2.

Table 2: Ranges for the four band families.

(g_1, g_3)	Lower bound for g_2	Upper bound for g_2
$(3, d-4)$	3	$(d+1)/2$
$(3, d-2)$	5	$(d-1)/2$
$(5, d-4)$	7	$(d-1)/2$
$(5, d-2)$	7	$(d-3)/2$

Omit $(3, 9, d-2, d-10)$ from the second band: the potential constructed below has strictly positive slack at this type for every odd $d \geq 101$, so it is not a first-order contact (Subsection A.2). Denote the resulting set of dihedral types by $\mathcal{T}_d$. The $d-2$ contact equations of $\mathcal{T}_d$ have a two-dimensional potential kernel, spanned by constants and the constructed nonconstant potential (Lemma 4.4).

For $d = 4q+1$, the four band counts before the deletion are $q, q-2, q-3, q-3$; for $d = 4q+3$ they are $q, q-1, q-2, q-3$. Adding eight anchors and deleting the one type gives $|\mathcal{T}_d| = d-2$ in both cases. The proof below is uniform for odd $d \geq 101$, the common range of its sign and rank certificates; smaller orders are covered separately in Appendix D.

The list $\mathcal{T}_d$ was suggested by exploratory numerical continuation of tight-type lists, starting from a certified order-14 seed and increasing the order by four. The eight anchors and four bands were extracted from a surviving order-82 candidate and then written as the formulas above for general odd d. Uniform feasibility and the required ranks are proved by the exact identities and certificates in Appendices A and B.

We call the property $x^{\mathsf{T}}Qx < 0$ for every nonzero $x \in \ker P \cap \mathbf{1}^{\perp}$ **strict negative type**.

Lemma 4.3 (uniform feasibility)**.** *For every odd $d \geq 101$ there is a real symmetric circulant Q such that $z_I^{\mathsf{T}} Q z_I \geq 0$ for every four-point support I with odd cyclic gaps, with equality exactly at the types in $\mathcal{T}_d$. Its quadratic form is strictly negative on* $\ker P \cap \mathbf{1}^{\perp}$ *outside zero.*

Proof outline; see Subsections A.1–A.2. One explicit rational-function potential works for every odd $d \geq 101$. Fixed identities verify its contact equations; two bulk inequalities and 552 endpoint cases cover all other types, and two spectral polynomial sign checks prove strict negative type.

(1) *Construct the real potential.* Adding a constant to R leaves (4.3) unchanged by the sine-chord Ptolemy identity. We fix this freedom by omitting an additive constant from the following ansatz; $u = u(d)$ and $v = v(d)$ are real coefficients determined by the contact equations. On ordinary odd distances $9 \leq x \leq d-10$ and ordinary even distances $14 \leq x \leq d-5$, respectively, use

$$R(x) = (du - x)\cot(\theta x), \qquad R(x) = (dv - x)\cot(\theta x), \tag{4.5}$$

with separate values at the seventeen exceptional distances listed in Appendix A. In this normalization, as d tends to infinity through odd integers,

$$u(d) \longrightarrow \frac{1645}{3664}, \qquad v(d) \longrightarrow \frac{1645}{1832}.$$

For example, $u(101) \approx 0.449114428465$ and $v(101) \approx 0.897350628683$.

Put $Q_0(0) = 0$ and $Q_0(x) = (-1)^x \sin^2(\theta x) R(x)$ for $1 \leq x < n$. Subsection A.1 directly solves the fixed 72 contact equations by rational functions. The band recurrence extends these identities to every selected type. Coefficient-reversal identities prove reality, and the denominator is nonzero throughout the stated range. This construction does not require contact independence; that is proved afterwards.

(2) *Prove the type inequalities and determine equality.* Types with only ordinary distances satisfy the two bulk inequalities (A.3). For the other types, fixing suitable distances leaves one parameter. Sine interpolation, concavity or monotonicity reduces each regular interval to its endpoints. The uniform affine and parity checks leave 552 cases: 20 zero identities in $\mathcal{T}_d$ and 532 strict signs. Lemma A.5 excludes further zeros between endpoints. Subsection A.2 gives the reductions and sign certificates.

(3) *Prove strict negative type.* Fourier diagonalization gives the eigenvalue formulas in Proposition A.1. The positive denominators and two positive spectral numerator polynomials make every nonconstant kernel-mode eigenvalue negative. The signs follow from rational Bernstein certificates on a fixed rectangle containing all corresponding pairs (c, χ) for odd $d \geq 101$ and the relevant frequencies.

□

Lemma 4.4 (independent contact equations)**.** *For every odd $d \geq 101$, the $d-2$ functionals $E \mapsto z_I^{\mathsf{T}} E z_I$, one per type in $\mathcal{T}_d$, are linearly independent on symmetric circulant matrices.*

Proof outline; see Subsection A.3. (1) *Determine the initial freedom.* The eight anchors and initial band rows give 22 independent equations in 27 distance coordinates, leaving five parameters. Subsection A.3 and Proposition A.6 specify the local minor and its nonvanishing certificate.

(2) *Propagate the potential values.* Each later band equation determines one new distance with a nonzero sine-product coefficient. The four families successively determine distances $g_2, d+2-g_2, g_2+5, d-3-g_2$. Propagation from the two ends of the distance range can give two expressions for the same coordinate; they must agree wherever they overlap.

(3) *Match and count the kernel.* Three matching equations impose three independent conditions on the five parameters, in both congruence classes of odd d. The surviving solution space is spanned by the constant potential $R \equiv 1$ and the nonconstant potential already constructed in Lemma 4.3. Thus its dimension is two, proving rank $d-2$. Only now do we conclude that removing the additive constant and fixing the coefficient of $x\cot(\theta x)$ to -1 determines the normalized potential uniquely. Proposition A.6 certifies the closing minors.

□

Proposition 4.5 (finite-zero parent)**.** *For every $n = 2d \geq 202$ with d odd, there is a positive-diagonal copositive matrix S of order n with finitely many zero rays, exactly the continued zero rays whose support types belong to $\mathcal{T}_d$.*

Proof. Strict negative type matters twice. A strictly concave quadratic on $Z \cap \{x : \sum x_i = 1\}$ can have a minimum only at an extreme ray, so Lemma 4.3 identifies all first-order contacts. For a tight four-point support I and $j \notin I$, inserting j splits an odd gap into one odd and one even gap. On $I \cup \{j\}$ there is exactly one even gap. Its two endpoints can be deleted to obtain the two circuit rays of this two-dimensional nonnegative kernel cone: the original ray and a second ray that includes j. Normalize their generators z, z' to have sum one. Since $z - z' \in \ker P \cap \mathbf{1}^{\perp}$ is nonzero, strict negative type gives

$$z^{\mathsf{T}} Q z' > \tfrac{1}{2}(z^{\mathsf{T}} Q z + z'^{\mathsf{T}} Q z') \geq 0.$$

For a supported correction w, the vector $g = Qz + Pw$ vanishes on I. Because $Pz' = 0$ and z' is supported in $I \cup \{j\}$,

$$z^{\mathsf{T}} Q z' = g^{\mathsf{T}} z' = g(j) z'(j) > 0.$$

As $z'(j) > 0$, this proves the strict entering inequality $g(j) > 0$. Also $P[I]$ has a simple kernel by independence of any three base columns. These verify conditions (1)–(3) of Lemma 3.3.

Lemma 4.4 supplies condition (4). Apply Lemma 3.3 with the dihedral group and the space of symmetric circulant perturbations. The diagonal remains positive by continuity from (4.1). □

5 The dimension of the zero outer products

To use Lemma 3.1 we must count outer products, not just zero rays. We do this before choosing a small perturbation parameter. For each selected first-order contact $\mathbb{R}_+ z_I$, let $t_I \in \operatorname{range} P$ agree with $-Qz_I$ on its four-point support, as in Lemma 3.3, and put $v_I = z_I + t_I$. Define $W = \operatorname{span}_{\mathbb{C}}\{v_I v_I^{\mathsf{T}}\}$, using both orientations of every selected dihedral type and all their cyclic translates. For $\rho_j = e^{2\pi \mathrm{i} j/n}$ put

$$W_j = \{X \in W : X_{a+1,b+1} = \rho_j X_{ab} \text{ for all } a, b\}, \qquad r_j = \dim_{\mathbb{C}} W_j.$$

We call r_j the **character rank** of W at j.

All indices are modulo n. The cyclic shift is diagonalizable, so $W = \bigoplus_{j \in \mathbb{Z}_n} W_j$ and $\sum_j r_j$ is the real dimension of the original span. Reality gives $r_j = r_{-j}$. Extend the trace pairing complex-bilinearly. If H and X have characters j and k, invariance under cyclic shift gives $\langle H, X\rangle = \rho_j \rho_k \langle H, X\rangle$. The pairing therefore vanishes unless $k = -j$. The opposite ambient character spaces pair nondegenerately, so restriction of character-j functionals to W has rank $r_{-j} = r_j$.

An ambient symmetric character matrix has $H_{a,a+x} = \rho_j^a \mu_x$, with $\mu_{n-x} = \rho_j^{-x} \mu_x$. In particular $\mu_d = 0$ for odd j. The ambient space of complex symmetric character-j matrices has dimension $D_j = d + 1$ for even j and $D_j = d$ for odd j. Let $\mathcal{R}$ contain one representative of each selected support modulo cyclic rotation, including reflected orientations whenever they give different cyclic orbits. Using the independent coordinates $\mu_0, \dots, \mu_d$ (omitting μ_d for odd j), define the evaluation matrices by

$$(\mathcal{E}_j^{\text{base}}\mu)_I = z_I^{\mathsf{T}} H_j(\mu) z_I,$$
$$(\mathcal{E}_j^{\text{lift}}\mu)_I = v_I^{\mathsf{T}} H_j(\mu) v_I, \qquad I \in \mathcal{R}.$$

The rows are supports, the columns are those independent coordinates, and each entry is the coefficient of its coordinate in the displayed quadratic evaluation. Rotating a support multiplies the row by a nonzero phase, so these representatives suffice and $\operatorname{rank} \mathcal{E}_j^{\text{lift}} = r_j$. A **generic character** below means $j \notin \{0, \pm 1, \pm 2\}$.

To compare the base and lifted evaluations, use the Fourier vectors a_f defined in Section 4, and write $\widehat{v}(f) = a_f^{\mathsf{T}} v$, with frequency indices modulo n. For one representative of each unordered pair $\{f, j - f\}$ put

$$B_f = a_f a_{j-f}^{\mathsf{T}} + a_{j-f} a_f^{\mathsf{T}}, \qquad v^{\mathsf{T}} B_f v = 2\widehat{v}(f)\widehat{v}(j - f).$$

The Fourier vectors form a basis of $\mathbb{C}^n$, so these symmetric products form a basis of the ambient character-j matrix space. Passing from the μ-coordinates to this basis is an invertible change of columns in both evaluation matrices and preserves their ranks. The positive-exponent convention agrees with the character convention above.

Since $z_I \in \ker P$ and $t_I \in \operatorname{range} P$, we have $\widehat{z}_I(f) = 0$ for $f \in \{d - 1, d, d + 1\}$ and $\widehat{t}_I(f) = 0$ for all other frequencies. Thus columns for which both frequencies lie outside this set have the same base and lifted evaluations. At a generic character, exactly three unordered pairs have one frequency in this set; choosing that frequency as f, their base evaluations vanish and their lifted evaluations are

$$L_f(I) = 2\widehat{t}_I(f)\widehat{z}_I(j - f), \qquad f \in \{d - 1, d, d + 1\}.$$

These are the three additional evaluation columns. A pair with both frequencies in $\{d - 1, d, d + 1\}$ can occur only when their sum j belongs to $\{0, \pm 1, \pm 2\}$ modulo n. Such columns have quadratic lift terms $2\widehat{t}_I(f)\widehat{t}_I(j - f)$ and are treated separately. The kernel-by-kernel columns remain unchanged at these characters as well.

Lemma 5.1 (base rows and propagation)**.** *At every nonzero character j, the evaluations of the base terms $z_I z_I^{\mathsf{T}}$ have rank at least $D_j - 3$, where $D_j = d + 1$ for even j and $D_j = d$ for odd j. For $j \neq \pm 1$ this is their exact rank; at $j = \pm 1$ the exact rank is $d - 2$.*

Proof outline; see Subsection B.1. (1) *Describe the evaluation kernel.* Fix a character j and abbreviate $\rho = \rho_j = e^{2\pi \mathrm{i} j/n}$. A character-$j$ symmetric matrix has entries $H_{a,a+x} = \rho^a \mu_x$, with $\mu_{n-x} = \rho^{-x}\mu_x$; its antipodal coordinate vanishes for odd j. A **stress** is such a matrix whose quadratic evaluation vanishes on every base ray. Products of base Fourier modes with complementary modes give three stresses:

$$\mu_x = (-1)^x\big(w^{2sx} + \rho^x w^{-2sx}\big), \qquad s = -1, 0, 1, \quad w = e^{\pi \mathrm{i}/n}. \tag{5.1}$$

(2) *Propagate all possible stresses.* Eight anchors give nonzero 13×13 and 12×12 minors in even and odd parity, after removing two distance columns. Band rows then determine each new folded-distance coordinate with a nonzero sine-product coefficient. Propagating from the two ends can give two expressions for the same coordinate; matching requires agreement wherever they overlap. Subsection B.1 gives the rows and meeting positions for both classes of odd d; Proposition B.1 supplies the initial minor certificates.

(3) *Count the global kernel.* For $j \neq \pm 1$, the three stresses in (5.1) are independent and exhaust the kernel, giving rank $D_j - 3$. At $j = \pm 1$ two stresses coincide. Their divided difference gives another local mode, but its two propagated expressions at an overlapping distance differ by a nonzero multiple of d. It fails global matching, leaving rank $d - 2$.

□

The base span for this uniform family has dimension

$$(d-2)(2d-1) = \binom{n}{2} - 2n + 2.$$

Indeed its rank is $d - 2$ at the trivial character by Lemma 4.4, at each of the $d - 1$ nonzero even characters, and at $j = \pm 1$; it is $d - 3$ at the remaining $d - 2$ odd characters by Lemma 5.1.

The following lemma refines the dimension-counting argument in the proof of [6, Theorem 5.1], with $k = n$, by decomposing the constraints into cyclic characters.

Lemma 5.2 (character ceilings)**.** *For every character j,*

$$r_j \leq D_j - 1 = \begin{cases} d, & j \text{ even}, \\ d-1, & j \text{ odd}. \end{cases}$$

Proof. If a copositive S has positive diagonal and u is a zero, put $E_k = e_k e_k^{\mathsf{T}}$. By [6, Theorem 2.1(a)], applied to uu^{T} and S,

$$\langle E_k S + S E_k, uu^{\mathsf{T}} \rangle = 2u_k(Su)_k = 0.$$

The n matrices $E_k S + S E_k$ are independent, as their diagonals show. For circulant S, rotation permutes them cyclically, giving one independent constraint in each character and a ceiling $D_j - 1$ for the actual zero outer products. The rescalings in Lemma 3.3 commute with rotation; hence the rank in each limit character cannot exceed its actual rank, and the same ceiling bounds r_j. □

Lemma 5.3 (lifted character ranks)**.** *For every odd $d \geq 101$,*

$$r_0 = d - 2, \qquad r_j = d \ (j \neq 0 \ \textit{even}), \qquad r_j = d - 1 \ (j \ \textit{odd}). \tag{5.2}$$

Proof outline; see Appendix B. (1) *Settle the trivial character.* Let pr_K be the orthogonal projection onto $K = \ker P$. For a symmetric circulant H, the matrix $H_K = \mathrm{pr}_K H \, \mathrm{pr}_K$ is symmetric circulant and

$$v_I^{\mathsf{T}} H_K v_I = z_I^{\mathsf{T}} H z_I,$$

because $z_I \in K$ and $t_I \in K^{\perp}$. Lemma 4.4 gives $r_0 \geq d - 2$. Each cyclic orbit contributes at most one averaged outer product, and reflected supports have the same average: the average is symmetric circulant and fixed by reflection. The $d - 2$ dihedral types therefore give $r_0 = d - 2$.

(2) *Test the new directions at generic characters.* Let M be the base anchor evaluation matrix. Choose a matrix Z whose rows are its relations, so $ZM = 0$, and let L contain the three additional evaluation columns defined above. The rows of Z cancel the base evaluations, so nonzero 2×2 minors of the residual matrix ZL certify two new directions. The base rank $D_j - 3$ and ceiling $D_j - 1$ permit exactly this gain. The band pivots then carry it to the full character matrix; see Subsection B.2.

(3) *Verify the generic even characters.* For nonzero even $j \neq \pm 2$, the selected residual minor factors into nonzero elementary factors and $G_{\mathrm{even}}(c, \chi)$, where $c = \cos(2\theta)$ and $\chi = \cos(2\pi j/n)$. Proposition B.2 proves $G_{\mathrm{even}} > 341$ by a Bernstein certificate on $1999/2000 \leq c \leq 1$, $-1 \leq \chi \leq 1$.

(4) *Verify the generic odd characters.* For odd $j \neq \pm 1$, simultaneous vanishing of two residual minors would make w a root of an integer resultant of degree 2,620. Since w is a primitive root of unity of order $4d$ and $\varphi(4d) = 2\varphi(d) \geq 2\sqrt{d}$, only $d \leq 1{,}716{,}100$ can occur. The degree restriction leaves 803 candidate odd $d \geq 101$. Finite-field nonvanishing certificates exclude them and the relevant scalar factors and denominators. Propositions B.3–B.4 give the identities and exhaustive exclusions.

(5) *Verify the special characters.* At $j = \pm 1$, the base rank $d - 2$ gains one from a quadratic lift term: the relevant quotient is nonreal and cannot cancel d. At $j = \pm 2$, a cross term and a quadratic term give a minor proportional to $d^2(C_2 + dC_3)$, with $C_3 \neq 0$ and $|C_2/C_3| < 100 < d$. Proposition B.5 proves these exclusions and interval bounds. These lower bounds attain Lemma 5.2's ceilings at every nonzero character and complete (5.2).

□

Proposition 5.4 (parent lower bound)**.** *For every $n = 2d \geq 202$ with d odd, $p_n \geq \binom{n}{2} - 2$.*

Proof. There are $d - 1$ nonzero even characters and d odd ones. Lemma 5.3 gives

$$\sum_j r_j = (d-2) + (d-1)d + d(d-1) = 2d^2 - d - 2 = \binom{n}{2} - 2.$$

Lemma 3.3 transfers this limit rank to the actual finite-zero matrix of Proposition 4.5, and Lemma 3.1 gives the cp-rank lower bound. □

Lemma 5.3 attains the ceiling of Lemma 5.2 at every nonzero character. At the trivial character $r_0 = d - 2$ is two below $D_0 - 1 = d$. The lifted span $\binom{n}{2} - 2$ therefore has $2n - 4$ more dimensions

than the base span. Both missing dimensions of this constructed face occur in the trivial character; this does not determine the exact unrestricted cp-rank gap.

Returning to the finite-order construction of Example 4.2, its five contact equations determine the five spectral coefficients in the finite construction (D.1). The 112 continued zero rays have certified character ranks $r_0 = 5$, $r_j = 7$ at the six nonzero even characters, and $r_j = 6$ at the seven odd characters. Hence their outer products span

$$5 + 6 \cdot 7 + 7 \cdot 6 = 89 = \binom{14}{2} - 2$$

dimensions, proving $p_{14} \geq 89$.

6 Axial extensions and proof of Theorem 1.2

For the axial construction set $\kappa = 1/2$ in (4.1) and multiply P by two, giving Gram columns $(\pm 1, \boldsymbol{u}_j)$, where $\boldsymbol{u}_j$ runs over a regular odd d-gon in $\mathbb{R}^2$; the two signs form the two **rails**. It has the same kernel and four-point rays, and the lifted vector t_I is unchanged: it is the unique range vector agreeing with $-Qz_I$ on the support I. Hence Proposition 5.4 still applies. An **axial column** is a new Gram column on the axis perpendicular to the polygon planes; an extension by a new row and column is also called a **border**. The **mass** on a set of coordinates is the sum of their coefficients.

Fix odd $d \geq 101$ and $n = 2d$. Let $\sigma_i = (-1)^i$ and, using the raw symbol Q_0 of Section 4, put

$$\gamma = \frac{\sigma^{\mathsf{T}} Q_0 \sigma}{n^2}, \qquad Q = Q_0 - \gamma \sigma \sigma^{\mathsf{T}}.$$

Since Q_0 is circulant, $Q\sigma = 0$. For $z \in \ker P$, $\sigma^{\mathsf{T}} z = 0$, so $Qz = Q_0 z$. Consequently the kernel form, supported corrections, lifted vectors and preceding rank conclusions are unchanged. Throughout this section use the dihedrally invariant family

$$S_\varepsilon = P + \varepsilon(Q + E_\varepsilon), \qquad E_\varepsilon = O(\varepsilon),$$

obtained from Lemma 3.3 with this normalized P and this Q.

Adding the Gram column $(1, 0, 0)$ gives the base matrix

$$\widehat{P} = \begin{pmatrix} P & \sigma \\ \sigma^{\mathsf{T}} & 1 \end{pmatrix}.$$

In a first-order extension, a zero through an axial column is a nonnegative dependence consisting of that axial column and three old columns whose spatial directions surround the origin. The three old columns are either all on the negative rail or two negative and one positive. For three polygon vertices, their **spatial gaps** are the positive numbers of polygon edges between successive vertices in cyclic order; they sum to d. They are distinct from the matrix-index gaps in Section 4, which sum to $2d$. Figure 2 illustrates an axial circuit at $d = 7$. This finite list of surrounding triangles makes the new border problem concrete.

Lemma 6.1 (one axial column)**.** *Let $S = S_\varepsilon$ be the normalized dihedrally invariant parent family just specified, with odd $d \geq 101$, $n = 2d$, and sufficiently small $\varepsilon > 0$, and set $h = h(S)$. There is a copositive extension by one axial column with finitely many zero rays, retaining the old zero rays and having outer-product span at least $h + n$.*

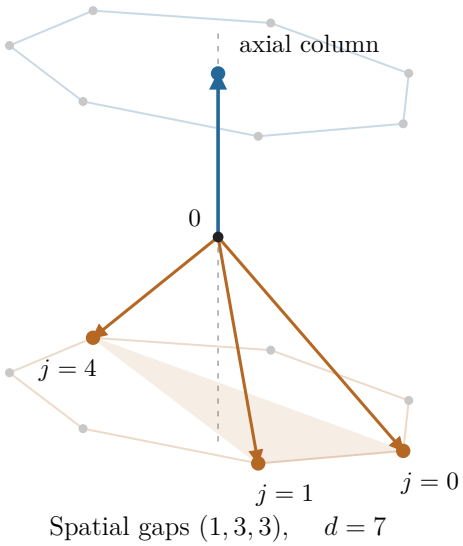

Figure 2: An axial circuit at $d = 7$: the column $(1, 0, 0)$ and three negative-rail columns with spatial indices $0, 1, 4$. Their spatial gaps are $(1, 3, 3)$, and their triangle contains the origin. The drawing illustrates the base Gram geometry.

Proof outline; see Subsections C.1–C.2. (1) *Choose the first-order border.* Normalize every circuit containing the new axial column to axial mass one. Its old part x has $\sigma^{\mathsf{T}} x = -1$; therefore $x^{\mathsf{T}} Q x = x^{\mathsf{T}} Q_0 x - \gamma$, the same shift for every circuit. Define M as the minimum of the **raw** value $x^{\mathsf{T}} Q_0 x$ on all-negative circuits. For mixed circuits let $\ell = \sum_i x_i > 1$ and set

$$b = \max_{\text{mixed}} \frac{M - x^{\mathsf{T}} Q_0 x}{2(\ell - 1)}, \qquad c_* = -M + \gamma - 2b. \tag{6.1}$$

The common shift cancels from the ratio defining b. The first-order extension is

$$\widehat{Q} = \begin{pmatrix} Q & b\mathbf{1} \\ b\mathbf{1}^{\mathsf{T}} & c_* \end{pmatrix}.$$

It is nonnegative on all base circuits and is tight on the minimizing all-negative orbit and at least one mixed orbit. Subsection C.1 proves that the minimizing spatial gaps are uniquely $(1, (d-1)/2, (d-1)/2)$, up to the polygon symmetries, and proves $b > 0$ and $c_* < 0$ for every odd $d \geq 101$.

(2) *Realize actual finite-zero extensions.* Subsection C.2 realizes these first-order contacts at a positive parameter. It first minimizes the actual parent quadratic on the negative-rail simplex and chooses the axial diagonal to make those minimizers exact zeros. It then varies the positive-rail coupling $1 + \varepsilon(b + \lambda)$ to its first copositive value λ_ε. Strict negative curvature on a suitable weighted simplex restricts minimizers to at most four coordinates; a mixed zero appears at that threshold. The same support argument proves that the zero list is finite and complete.

(3) *Count the new outer-product directions.* The minimizing negative triangle has a weight with nonzero derivative, so all Fourier coefficients of its translated old parts are nonzero at sufficiently small positive parameter. Those translates span the d negative-rail coordinates. A mixed zero has one positive-rail coordinate, and its translates span the remaining d coordinates after projection away from the negative rail. Projection onto the old-to-axial entries therefore adds n dimensions to the old outer-product span. □

Lemma 6.2 (opposite axial columns and duplication)**.** *With this same normalized family $S = S_\varepsilon$, there is a copositive extension by two opposite axial columns with finitely many zero rays and outer-product span at least $h + 2n + 1$. Duplicating either axial column in this extension gives a matrix of order $n + 3$ with span at least $h + 3n + 2$.*

Proof. Let $a_\varepsilon > 0$ denote the diagonal entry of the new axial column in the extension of Lemma 6.1. Obtain the opposite axial column by rotating the old indices once; this exchanges the rails. Give the two axial columns the same diagonal a_ε and mutual entry $-a_\varepsilon$. Their couplings to any old coordinate add to a strictly positive number. For a nonnegative vector with axial masses t_1, t_2, subtract $\min(t_1, t_2)$ from both. This leaves the quadratic on the two axial coordinates $a_\varepsilon(t_1 - t_2)^2$ unchanged and decreases the old-to-axial terms. Copositivity follows from the two extensions by one axial column. With nonzero old mass, a zero has at most one axial column; with zero old mass, the only new ray is the pure axial pair. The two families of zeros using one axial column contribute n independent cross entries each, in disjoint blocks; the pure pair contributes the entry between the axial columns. This gives $h + 2n + 1$.

The zero rays through either axial coordinate project onto all n old coordinates, and the pure axial pair supplies another independent direction. Their ordinary span therefore has dimension at least $n + 1$. Complementarity at the positive axial diagonal puts them in one hyperplane of the order-$(n + 2)$ space, giving the reverse inequality. Hence $\rho_i = n + 1$. Lemma 3.2 now gives an order-$(n + 3)$ matrix whose outer-product span is at least

$$(h + 2n + 1) + (n + 1) = h + 3n + 2.$$ □

These lemmas yield the four rows of Table 3, needed for the all-order statement. For each parent $n = 202, 206, 210, \ldots$, Proposition 5.4 gives $h \geq \binom{n}{2} - 2$.

Table 3: Dimension and cp-rank gap bounds from axial extensions and duplication.

Order	Construction	Span at least	Cp-rank gap at most
n	Four-band parent	h	2
$n + 1$	One axial column	$h + n$	2
$n + 2$	Opposite axial columns	$h + 2n + 1$	2
$n + 3$	Duplicate an axial column	$h + 3n + 2$	3

Proposition 6.3 (finite bridge)**.** *For every $14 \leq N \leq 201$, and also at $N = 205$, there is a completely positive matrix of order N with cp-rank at least $\binom{N}{2} - 2$.*

Proof. Proposition D.1 gives finite-zero constructions at the 94 even orders $14, 16, \ldots, 200$ and the 90 odd orders $23, 25, \ldots, 201$. Proposition D.2 supplies the four remaining orders $15, 17, 19, 21$. These disjoint lists cover all 188 orders from 14 through 201, and their outer-product span dimensions are at least $\binom{N}{2} - 2$. The construction given in Subsection D.2, using Proposition D.3, extends the order-204 witness and adds 204 dimensions, giving $20{,}704 + 204 = 20{,}908 = \binom{205}{2} - 2$. Apply the finite-zero Lemma 3.1 in each case. □

Proof of Theorem 1.2. Theorem 1.1 gives $\binom{N}{2} - p_N \geq 0$. Proposition 6.3 supplies the two-unit lower bound through 201. The parent at $n = 202$ and Lemmas 6.1–6.2 supply it at 202, 203 and 204; the separate certificate supplies it at 205. The parent at $n = 206$ supplies it at 206, 207 and 208. Every later order lies in one of the four-row blocks beginning at $n = 206, 210, 214, \ldots$. The first three rows have gap at most two, while the fourth has gap at most three and occurs exactly at $N \equiv 1 \pmod 4$. The first such order beyond the bridge is 209. This proves (1.1). □

7 Remarks and open questions

A natural next target is $p_N \geq \binom{N}{2} - 2$ for $N \geq 209$ with $N \equiv 1 \pmod 4$. The extra unit in the present guarantee arises at the duplication step. The horizontal extension at order 205 suggests investigating replacement extensions valid throughout this residue class. More generally, what are the exact values of $\binom{N}{2} - p_N$ for large N, and which values occur infinitely often? In particular, is the upper bound $\binom{N}{2}$ attained for arbitrarily large orders?

Computational supplement

The archive `certificates-v6.tar.gz` supplies the code, exact inputs, computational results and reproduction instructions. Its file sizes and SHA-256 hashes are recorded in `certificate-manifest-v6.json`. Both files are archived together with this paper on Zenodo at https://doi.org/10.5281/zenodo.24517737. Every archive path in the paper is relative to its extraction root. All mathematical proofs and all reductions justifying the certificates are in this paper, including Appendices A–E.

Conventions for the appendices

The following proofs distinguish symbolic identities, continuous sign certificates and finite-order witnesses. Throughout, $c = \cos(\pi/d) = \cos(2\theta)$. For odd $d \geq 101$, this value belongs to each of $[99/100, 1]$, $[1999/2000, 1]$ and $[2047/2048, 1]$; separate certificates use different containing intervals. For instance, $1 - c < \pi^2/(2 \cdot 101^2) < 1/2048$.

All paths below are relative to the extracted computational supplement. In Appendix A, bare certificate data filenames refer to `code/runs/uniform-potential/exact/` and verifier filenames to `code/tests/uniform-potential/`. In Appendix B, the corresponding directories are `code/runs/character-ranks/` and `code/tests/character-ranks/`.

The finite arithmetic certificates have the following elementary interpretations. A polynomial expressed in the Bernstein basis on a rectangle is at least its smallest Bernstein coefficient, since those basis functions are nonnegative and sum to one. Subdividing the rectangle and repeating this test is a finite proof when the subrectangles cover it. For a univariate expansion about an endpoint, bounding each negative Taylor term by its largest possible absolute value gives a rational lower bound. Exact powers of a vanishing endpoint factor are removed before claiming strict positivity in the interior. To prove that a polynomial $P(w)$ of degree m has no zero on the stated arc, put

$$H(s_{\mathrm{arc}}) = (1 - \mathrm{i}s_{\mathrm{arc}})^m P\left(\frac{1 + \mathrm{i}s_{\mathrm{arc}}}{1 - \mathrm{i}s_{\mathrm{arc}}}\right).$$

This is a polynomial in the real parameter s_{arc}, and the clearing factor is nonzero. It therefore suffices to prove that either $\operatorname{Re} H$ or $\operatorname{Im} H$ has a strict constant sign throughout the arc. If that

component has an exact factor s_{arc}^k, remove it and certify the remaining real polynomial by the rational Taylor or Bernstein bounds above; the resulting nonvanishing conclusion holds for $s_{\mathrm{arc}} > 0$, with the endpoint 0 excluded when $k > 0$. All intervals, coefficients and subdivision trees in the certificates are rational. These checks establish statements on continuous intervals; checking individual polygon orders would not suffice.

A Four-band feasibility

A.1 Potential and contact equations

Let $n = 2d$, where $d \geq 101$ is odd, and set $\theta = \pi/n$, $w = e^{\mathrm{i}\theta}$, $c = \cos(2\theta)$. For $1 \leq x \leq d$ put $s(x) = \sin(\theta x)$. We use the raw zero-diagonal symmetric circulant symbol $Q_0(x) = (-1)^x s(x)^2 R(x)$ of Section 4; adding a constant to R leaves the associated quadratic form unchanged on the nonnegative kernel. The later normalization $Q = Q_0 - \gamma\sigma\sigma^{\mathsf{T}}$ in Section 6 has the same kernel form and lifted vectors. The four-point ray and its Ptolemy functional are (4.2)–(4.3). In particular, the sign of $z_I^{\mathsf{T}} Q z_I$ is the sign of the right-hand side of (4.3), since every factor of Π_I is positive.

For ordinary distances the potential is

$$\begin{aligned} R(x) &= (du - x)\cot(\theta x), \quad x \text{ odd}, \quad 9 \leq x \leq d - 10, \\ R(x) &= (dv - x)\cot(\theta x), \quad x \text{ even}, \quad 14 \leq x \leq d - 5. \end{aligned} \tag{A.1}$$

The seventeen exceptional positive distances are

$$1, 2, 3, 4, 5, 6, 7, 8, 10, 12; \quad d-8, d-6, d-4, d-3, d-2, d-1, d. \tag{A.2}$$

Define $k_x = u$ for odd x and $k_x = v$ for even x, and write the exceptional corrections as

$$Q_0(x) = \tfrac{1}{2}(-1)^x (dk_x - x)\sin(\pi x/d) + d\sin(\pi/d)\delta_x,$$

with $\delta_x = 0$ at ordinary distances. At $d = 101$, the corrections at the seventeen distances in (A.2) are approximately

$$\begin{aligned} (\delta_1, \delta_2, \delta_3) &= (0.00051230845, -0.0017023469, 0.00046361538), \\ (\delta_4, \delta_5, \delta_6) &= (-0.0018138574, -0.00017660988, -0.00080833262), \\ (\delta_7, \delta_8, \delta_{10}) &= (-0.00042576798, 0.00052198887, 0.0010725457), \\ (\delta_{12}, \delta_{93}, \delta_{95}) &= (0.00072049529, -0.027533973, -0.050765812), \\ (\delta_{97}, \delta_{98}, \delta_{99}) &= (-0.017688278, 0.016282255, 0.18697790), \\ (\delta_{100}, \delta_{101}) &= (-0.0045113946, 0.92143235). \end{aligned}$$

These rounded values illustrate the normalized potential. A separate unknown value is assigned to each exceptional distance. We use $R(x)/\mathrm{i}$ as the coordinate at each exceptional distance so that it matches the scaling of the ordinary potential. Indeed, with $w = e^{\mathrm{i}\theta}$,

$$\frac{\cot(\theta x)}{\mathrm{i}} = \frac{w^x + w^{-x}}{w^x - w^{-x}}.$$

For the affine distances $x = rd + k$ occurring here, substitute $w^d = \mathrm{i}$. The ordinary formulas $R(x)/\mathrm{i} = (dk_x - x)\cot(\theta x)/\mathrm{i}$ then have rational-function coefficients in w. Thus the nineteen coordinates are du, dv and the seventeen exceptional values $R(x)/\mathrm{i}$. Removing a common phase

from each equation and clearing Laurent powers gives polynomial coefficients over $\mathbb{Q}$, allowing exact elimination over $\mathbb{Q}(w)$. This rescaling is a choice of coordinates; reality of R is proved below.

A band row is **irregular** if at least one pair distance varying with g_2, after folding, belongs to the exceptional list (A.2); all other band rows are **regular**. Fixed band sides are exceptional in every row and do not determine this classification. We use the eight anchors (4.4) and the bands in Table 2, omitting the second-band row $g_2 = 9$. Following the order of the four bands in Table 2, the eight irregular rows have $g_2 = 3, 5, 7, 9$ in the first band, $g_2 = 5, 7$ in the second, and $g_2 = 7$ in each of the third and fourth bands. For the regular rows take

$$\begin{aligned} g_2 &= 11, 13, \dots, 37 \quad \text{in each of the first two bands,} \\ g_2 &= 9, 11, \dots, 35 \quad \text{in each of the last two bands.} \end{aligned}$$

Each list contains fourteen consecutive odd parameters and lies within its band range for every odd $d \geq 101$. The deleted second-band row $g_2 = 9$ is excluded. Together with the eight anchor equations (4.4), these give exactly $8 + 8 + 4 \cdot 14 = 72$ equations. The recurrence proved below propagates the fourteen regular equations through the regular part of each band; the irregular rows have been included separately.

Fix the coefficient of $x \cot(\theta x)$ to -1 as in (A.1). In the coordinate order du, dv followed by $R(x)/\mathrm{i}$ at the seventeen distances (A.2), the candidate solution is

$$y = \frac{N_0(w) + dN_1(w)}{D(w)}.$$

The file `potential-solution.json` gives its integer numerator vectors, degree-204 common denominator and canceled polynomial row factors. The following proposition verifies the candidate and its spectral properties.

Proposition A.1 (potential and spectral certificates)**.** *The 72 defining equations in Subsection A.1 have the stated real rational solution for every odd $d \geq 101$. Substitution makes each equation identically zero, separately in its constant and linear coefficients in d. The common denominator is nonzero. The eigenvalues of Q_0 on nonconstant Fourier modes in $\ker P$ are strictly negative.*

Proof using certificates. On a regular part of a band, let I be the corresponding four-point support and set

$$F(g_2) = 2\Pi_I \mathcal{D}_g(R) = \Pi_I^2 z_I^{\mathsf{T}} Q_0 z_I.$$

Since $\Pi_I > 0$, $F(g_2) = 0$ is equivalent to the contact equation. The slack has frequencies $e^{2\mathrm{i}r\theta g_2}$ with $|r| \leq 1$, and Π_I has $|r| \leq 2$. Consequently the cleared expression has the form

$$F(g_2) = \sum_{r=-3}^{3} (a_r + g_2 b_r) e^{2\mathrm{i}r\theta g_2}.$$

If T shifts g_2 by two, this sequence is annihilated by

$$\prod_{r=-3}^{3} (T - e^{4\mathrm{i}r\theta})^2.$$

The seven roots are distinct for $d \geq 101$, and the recurrence has a nonzero constant coefficient. Fourteen consecutive zero values therefore force all values on the same regular branch to vanish, propagating in either direction. The irregular rows are included separately among the 72 defining equations.

Substitution of the candidate into the 72 selected equations makes every numerator vanish separately in its constant and linear coefficients in d. The recurrence therefore proves all selected band contact equations. The denominator's noncyclotomic reciprocal factor is the degree-32 polynomial $H_c(c)$, greater than 1509 on $[99/100, 1]$ by `spectral-H.json`. Together with the canceled sine and polynomial factors recorded in the solution certificate, this proves that the denominator is nonzero at every required w.

Reality follows directly from the coefficient identities. Order the nineteen coordinates as du, dv, followed by the seventeen values $R(x)/\mathrm{i}$ in the order of (A.2). In the representation $(N_0(w) + dN_1(w))/D(w)$, the integer coefficient arrays satisfy

$$w^{204}D(w^{-1}) = -D(w), \qquad w^{204}N_k(w^{-1}) = \operatorname{diag}(-1,-1,1,\dots,1)N_k(w), \quad k = 0, 1.$$

These are exact coefficient identities, checked by reversal within degree 204. At the required unit-circle values of w, conjugation sends w to w^{-1}. The denominator is nonzero by the preceding certificate. Hence du, dv are real and each ratio representing $R(x)/\mathrm{i}$ is purely imaginary. Thus all exceptional values $R(x)$ are real; the ordinary values are real by the cotangent formulas.

The construction and these reality identities precede the contact-independence argument. Proposition A.6 uses the constructed nonconstant solution together with the constant solution and proves that they span the full contact-equation kernel. Uniqueness after fixing addition of a constant and scaling follows from that later rank statement.

The rational functions for u, v in `spectral-rational.json` have a common denominator equal to 3664 at $c = 1$. Evaluation gives the limits $1645/3664$ and $1645/1832$ stated after (4.5).

Let λ_k be the eigenvalue of the raw circulant Q_0 at the Fourier vector a_k, and put

$$c = \cos(\pi/d), \qquad \chi = \cos(2\pi k/n) = \cos(\pi k/d).$$

Put $\alpha = \pi/d$ and, with $\delta_x = 0$ at ordinary distances, define

$$\mathcal{C}_k = 2\sum_{x=1}^{d-1} \delta_x \cos(k\alpha x) + (-1)^k \delta_d.$$

For $1 \le k \le d-2$, the explicit Fourier sums are

$$\frac{\lambda_k}{d\sin\alpha} = \begin{cases} -\dfrac{1}{2(c+\chi)} - \dfrac{v-u}{2(c-\chi)} + \mathcal{C}_k, & k \text{ even}, \\ -\dfrac{u+v-1}{2(c+\chi)} + \mathcal{C}_k, & k \text{ odd}. \end{cases}$$

Indeed, symmetry of the first row gives

$$\lambda_k = 2\sum_{x=1}^{d-1} Q_0(x)\cos(k\alpha x) + (-1)^k Q_0(d).$$

For any integer m not divisible by $2d$, finite geometric sums give

$$\sum_{x=1}^{d-1} \sin(m\alpha x) = \frac{1-(-1)^m}{2}\cot\frac{m\alpha}{2},$$
$$\sum_{x=1}^{d-1} x\sin(m\alpha x) = -\frac{d(-1)^m}{2}\cot\frac{m\alpha}{2}.$$

Both sums vanish when $2d \mid m$. Substitute $Q_0(x) = \frac{1}{2}(-1)^x(dk_x - x)\sin(\alpha x) + d\sin\alpha\,\delta_x$, split $k_x = (u+v)/2 + (v-u)(-1)^x/2$, and use the sine–cosine product identity. The two parity cases give the displayed formulas.

Writing $T_r(\cos t) = \cos(rt)$ for the Chebyshev polynomials, one has $T_{d-r}(\chi) = (-1)^k T_r(\chi)$. Thus the exceptional corrections fold to a polynomial in χ: low distances contribute $2\delta_x T_x(\chi)$, distances $d-r$ contribute $2(-1)^k\delta_{d-r}T_r(\chi)$, and distance d contributes $(-1)^k\delta_d$. Substituting the certified rational functions for u, v, δ_x and clearing their common denominator gives

$$\lambda_k = -\frac{d\sin(\pi/d)}{H_c(c)} \begin{cases} \dfrac{F_{\mathrm{even}}(c,\chi)}{c^2-\chi^2}, & k \text{ even}, \\ \dfrac{F_{\mathrm{odd}}(c,\chi)}{c+\chi}, & k \text{ odd}. \end{cases}$$

Here $H_c, F_{\mathrm{even}}, F_{\mathrm{odd}}$ are the certificate-normalized polynomials in `spectral-H.json`, `spectral-even.json` and `spectral-odd.json`; their degrees are respectively 32, $(41, 14)$ and $(40, 13)$, and $H_c(1) = 3664$. The rational-function identities are specified in `spectral-rational.json` and reconstructed by the independent verifier. Its certificate variable t is the displayed χ.

The Bernstein subdivisions prove $H_c > 1509$, $F_{\mathrm{even}} > 13$ and $F_{\mathrm{odd}} > 234$ on $[99/100, 1] \times [-1, 1]$, with H_c independent of the second variable. At the required even frequencies, $-c < \chi < c$, so $c^2 - \chi^2 > 0$. At the required odd frequencies, $\chi > -c$, so $c + \chi > 0$. Therefore every displayed eigenvalue is strictly negative. Symmetry gives $\lambda_{n-k} = \lambda_k$, and these pairs exhaust the nonconstant modes in $\ker P$; the range frequencies $d-1, d, d+1$ are excluded. Hence Q_0 is strictly negative on $\ker P \cap \mathbf{1}^\perp$. The rational Bernstein certificates in `spectral-H.json`, `spectral-even.json` and `spectral-odd.json` prove the three bounds, respectively. The independent verifier reconstructs the identities and checks the subdivisions and coefficient reversals.

The constant eigenvalue is positive: a first-order contact has positive sum, no component on the base frequencies and a nonzero low-frequency component, so its zero quadratic requires a positive constant contribution. □

A.2 Exhausting the other types

For a type $g = (g_1, g_2, g_3, g_4)$, recall the slack $\mathcal{D}_g(R)$ defined after equation (4.3). All gaps are positive odd integers with sum $2d$. Put $\alpha = 1 - 2u$, $\beta = 1 - v$ and $\Delta = v - 2u$. The parameter certificates prove $u, \alpha, \beta, v - u > 0$ and $\Delta < 0$ on $1999/2000 \le c < 1$.

Lemma A.2 (bulk reduction)**.** *If none of the six folded pair distances is exceptional, then $\mathcal{D}_g(R) > 0$ for every odd $d \ge 101$.*

Proof. First suppose all gaps are below d. Substitute (A.1) into $\mathcal{D}_g$, expand the sines of sums and use $g_1 + g_2 + g_3 + g_4 = 2d$. The coefficients of the distance-dependent linear terms combine to

$$\mathcal{D}_g(R)/d = \alpha S_g - \beta F, \qquad S_g = \sin(\theta(g_1 + g_3)),$$

where

$$\begin{aligned} F &= |\cos(\theta(g_1+g_2))|\sin(\theta(g_2+g_3)) \\ &\quad + |\cos(\theta(g_2+g_3))|\sin(\theta(g_1+g_2)) \\ &= \max\{|\sin(\theta(g_1-g_3))|, |\sin(\theta(g_2-g_4))|\}. \end{aligned}$$

The last equality follows by the four possible signs of the two diagonal cosines. For $0 < A, C < \pi/2$,

$$\frac{|\sin(A-C)|}{\sin(A+C)} = \frac{|\tan A - \tan C|}{\tan A + \tan C}.$$

This is largest when the two arguments are as far apart as allowed. Each bulk side is between 9 and $d-10$, so $F/S_g \le \cos(19\theta)/\cos\theta$.

If one gap $g_{\max}$ exceeds d, let g_{opp} denote its opposite gap in this paragraph. The same expansion, now folding $g_{\max}$ to $2d - g_{\max}$, gives

$$\mathcal{D}_g(R)/d = v[-\cos(\theta g_{\max})]\sin(\theta g_{\mathrm{opp}}) + \Delta\sin(\theta g_{\max})\cos(\theta g_{\mathrm{opp}}).$$

The folded large gap is at most $d-10$, and $g_{\mathrm{opp}} \ge 9$. After dividing by positive cosine factors, the minimum is bounded by the case with folded large gap $d-10$ and opposite gap 9, because $\Delta < 0$ and tangent is increasing. Positivity therefore follows from

$$v\sin(10\theta)\sin(9\theta) + \Delta\cos(10\theta)\cos(9\theta) > 0.$$

The two cases reduce respectively to

$$\alpha - \beta\frac{\cos(19\theta)}{\cos\theta} > 0, \qquad (v-u) - u\frac{\cos(19\theta)}{\cos\theta} > 0. \tag{A.3}$$

As $d \to \infty$ through odd integers, the two left-hand sides of (A.3) are respectively

$$\frac{27057\pi^2}{7328d^2} + O(d^{-4}), \qquad \frac{141447\pi^2}{7328d^2} + O(d^{-4}),$$

with leading coefficients approximately 36.4413 and 190.5057. These expansions follow by Taylor expansion of the rational functions at $c = 1$ and $1 - \cos(\pi/d) = \pi^2/(2d^2) + O(d^{-4})$. The certificates below prove positivity throughout the required range.

Here $\cos(19\theta)/\cos\theta$ is a degree-nine polynomial in $c = \cos(2\theta)$. Clearing the positive solution denominator and removing one factor $1 - c$ gives the rational Taylor certificates in `bulk-positivity.json`. They prove (A.3) on the whole required interval. □

A **distance slot** is one of $g_1, g_2, g_3, g_4, g_1+g_2, g_2+g_3$, before folding. Fixing two slots leaves a one-parameter family after imposing the sum of the gaps. A **bulk branch** specifies the parity and the side of the folding point d of each varying distance, with its folded value in an ordinary interval of (A.1). A **regular interval** of the parameter contains no change of bulk branch and no additional exceptional distance; retain its first and last admissible odd parameters. Branch changes at d are separated by exceptions: d for a side, and $d-1, d+1$ for an even diagonal.

Lemma A.3 (endpoint reduction)**.** *In a regular family with two fixed distance slots, nonnegative slack at both endpoints implies nonnegative slack throughout.*

Proof. Up to rotation and reversal there are three cases.

Adjacent fixed gaps. With two adjacent gaps fixed, including the equivalent case of a side and an adjacent diagonal fixed, substitution shows that the slack has the form $f(t) = A\sin(\theta t) + B\cos(\theta t)$. If the endpoints are t_0, t_1, with $0 < \theta(t_1 - t_0) < \pi$, then

$$f(t) = \frac{\sin(\theta(t_1 - t))}{\sin(\theta(t_1 - t_0))}f(t_0) + \frac{\sin(\theta(t - t_0))}{\sin(\theta(t_1 - t_0))}f(t_1).$$

Both weights are positive in the interior.

Opposite fixed gaps. With two opposite gaps fixed and all gaps below d, the bulk quantity S_g and every exception correction are constant. The variable part of F is an absolute sine of the difference of the other two gaps; this difference lies in $(-d, d)$, so its maximum on an interval occurs

at an endpoint. If a gap greater than d is fixed the slack is constant. If that gap varies, its variable contribution is $(v-u)\sin(\theta(g_1+g_3+2g_2))$ after relabeling; its angle is in $(0,\pi)$, so it is concave.

Fixed diagonals. With both diagonals fixed, F and every exceptional correction are constant. When all gaps are below d the variable term is $\alpha\sin(\theta(g_1+g_3))$, again concave. When one gap exceeds d it is $-u\sin(\theta(g_1+g_3))$; the side sum remains wholly in $(0,d)$ or wholly in $(d,2d)$, where the sine is monotone.

Each case puts the minimum at an endpoint. These computations remain valid when a fixed slot has an ordinary bulk value. □

To turn this reduction into a fixed finite enumeration, use the **critical distances**: the exceptional distances in (A.2) together with their immediate neighbors of the same parity within $\{1,\dots,d\}$. They are

$$\begin{aligned}\text{odd: } & 1,3,5,7,9;\ d-10,d-8,d-6,d-4,d-2,d,\\ \text{even: } & 2,4,6,8,10,12,14;\ d-5,d-3,d-1.\end{aligned}$$

Here is why two critical slots suffice. Suppose first that an exceptional side occurs. Hold that side and either adjacent side fixed, and vary the other two gaps with their sum fixed. If the exception is instead a diagonal, hold one of its adjacent sides fixed as well; the other side of that diagonal is then fixed too. In either case Lemma A.3 applies to an adjacent-fixed-gap family and preserves the original exception. Divide its feasible odd parameters at every additional exception. At a feasibility endpoint a varying side is 1; at an exceptional-distance event there is another exceptional slot; and immediately next to such an event a varying slot is a parity neighbor of an exception. Thus every retained endpoint has a second critical slot. No folding change is missed: a side crosses d through the exception d, and an even diagonal crosses it between the exceptions $d-1$ and $d+1$. Applying the sine interpolation on each regular interval reduces the original value to values with two critical slots, at least one still an exception. If both chosen slots were already critical, this preliminary step is unnecessary.

We now give the finite recipe, including the tests that make it uniform. Unfold each listed critical distance x to both x and $2d-x$. For every pair among the six slots and every assignment of unfolded critical values to it, retain the assignment only if at least one value is an exception. Impose these two equations and $g_1+g_2+g_3+g_4=2d$. Choose a gap as a free parameter t so that the resulting 4×4 system is nonsingular. Exact elimination gives

$$g_i = a_i d + b_i + s_i t, \qquad a_i,b_i,s_i\in\mathbb{Z},\quad |s_i|\le 1.$$

The inverse matrices are integral, and the defining equations are checked by substitution. Since d,t are odd, parity is the fixed test $a_i+b_i+s_i\equiv 1 \pmod 2$. A constant gap must be positive. The remaining inequalities $g_i\ge 1$ give affine odd bounds $L(d)\le t\le U(d)$; empty intervals are discarded.

For each varying slot, solve its equality to every unfolded exceptional value. The parameter coefficient is 1 or -1, so each event is integral affine in d and odd. Retain those in $[L,U]$, together with L,U, and order the distinct events. Between consecutive events $a<b$, retain $a+2,b-2$ when $b-a\ge 4$; if $b-a=2$ there is no interior odd parameter. Thus the retained list contains every exception and the first and last odd parameters in each regular interval. A singleton interval is retained as well.

Every comparison is checked as an affine inequality for the entire ray $d\ge 101$: $ad+b\ge q$ follows exactly from $a\ge 0$ and $101a+b\ge q$. This includes discarded assignments, out-of-range event roots and empty intervals. Endpoint gap positivity, parity and the gap sum are rechecked. Each tuple is identified under its four rotations and four reversed rotations. The chosen affine representative is checked to be lexicographically least throughout the ray, including any tie at 101.

Finally, solving the four affine equalities between distinct representatives excludes a collision at an isolated admissible odd order. These finite rational and integer tests justify the enumeration for every odd $d \geq 101$.

For example, fixing $g_1 = g_2 = 1$ gives $(1, 1, t, 2d-2-t)$, with $1 \leq t \leq 2d-3$, t odd. Its event parameters are

$$1, 3, 5, 7, 9, 11; \qquad d-10, d-8, \ldots, d+8;$$
$$2d-13, 2d-11, \ldots, 2d-3.$$

The two long intervening intervals also contribute $13, d-12, d+10, 2d-15$. The relevant ordering difference is $d - 21 \geq 80$, while the other consecutive event differences are two. Thus this family has 22 events and 26 retained parameters, uniformly in d.

Proposition A.4 (finite endpoint certificate)**.** *This construction gives 3,224 feasible affine families, 25,568 boundary/exception events, 29,192 retained parameter occurrences after adding the regular-interval endpoints, and 552 distinct dihedral endpoint types. Exactly 20 endpoint slacks vanish identically, and each is a selected type. Every other endpoint slack is strictly positive for odd $d \geq 101$.*

Proof using certificates. The complete family equations, ordered events, retained parameters and 172 distinct affine inequalities are in `endpoint-coverage.json`. The independent standard-library checker `verify_endpoint_coverage.py` reconstructs them by the preceding recipe without importing the producer. Of 6,406 slot/value assignments, 240 have no exception, 1,616 force a nonpositive gap, and 1,326 have an empty feasible interval; none fails parity. The remaining 3,224 families give the stated counts. Their dihedral list is checked against both `endpoints.json` and the 552 entries in `endpoint-sign-polynomials.json`. Thus every needed endpoint is among the sign-certified expressions.

We spell out the conversion from the quadratic expression to the sign polynomial. Let $D(w)$ be the common degree-204 solution denominator. The Laurent identity is

$$D(w) = w^{102}(w^2 - w^{-2})E(c), \qquad E(1) = 196976640.$$

Here

$$\begin{aligned} E(c) &= C_{\mathrm{den}}(c)H_c(c), \\ C_{\mathrm{den}}(c) &= 512c(4c^2-1)(2c^2-1)(4c^2-3) \\ &\quad \cdot (16c^4 - 12c^2 + 1)(8c^4 - 8c^2 + 1)(8c^3 + 4c^2 - 4c - 1). \end{aligned}$$

Each displayed factor is positive on $[99/100, 1]$. Thus H_c is the degree-32 noncyclotomic factor, while the degree-50 polynomial E includes these additional degree-18 factors.

For the scaled ray $\widetilde{z}_I = \Pi_I z_I$, use $\widehat{s}(x) = 2\mathrm{i}\sin(\theta x)$ in the exact calculation, so $\widehat{\Pi}_I = (2\mathrm{i})^6\Pi_I = -64\Pi_I$ and $\widehat{\widetilde{z}}_I = (2\mathrm{i})^3\widetilde{z}_I$. The stored symbol numerators are $4\mathrm{i}D(w)Q_0(x)$. Therefore their quadratic expression numerator $\mathcal{R}_g$ satisfies

$$\mathcal{R}_g = -256\mathrm{i}D(w)\widetilde{z}_I^{\mathsf{T}}Q_0\widetilde{z}_I = 8\mathrm{i}D(w)\widehat{\Pi}_I\mathcal{D}_g(R).$$

Each endpoint identity has zero constant coefficient in d. After dividing its linear coefficient by $\widehat{\Pi}_I$, the checked Laurent reflection and parity identities give one of the two forms

$$\frac{\mathcal{R}_g}{d\widehat{\Pi}_I} = 2w^{102}J(c) \quad \text{or} \quad 2\mathrm{i}w^{102}(w - w^{-1})K(c).$$

Since $w^2 - w^{-2} = 2\mathrm{i}\sin(2\theta)$ and $w - w^{-1} = 2\mathrm{i}\sin\theta$, cancellation gives, respectively,

$$\frac{\mathcal{D}_g(R)}{d} = -\frac{J(c)}{8\sin(2\theta)E(c)} \quad \text{or} \quad \frac{\mathcal{D}_g(R)}{d} = \frac{K(c)}{8\cos\theta E(c)}.$$

All canceled sine factors and Π_I are positive for the actual gaps. The denominator certificate proves $E(c) > 0$ on $[1999/2000, 1]$; also $\sin(2\theta) > 0$ and $\cos\theta > 0$.

For a nonzero endpoint the stored polynomial $p(c)$ is the sign-oriented, endpoint-reduced numerator: $-J(c) = (1-c)^m p(c)$ in the first case, or $K(c) = (1-c)^m p(c)$ in the second, with the integer multiplicity m recorded. The verifier reconstructs these identities before checking positivity. Expanding $p(1-y) = \sum_k a_k y^k$, the rational Taylor test proves

$$a_0 + \sum_{k\geq 1} \min(a_k, 0)(1/2000)^k > 0,$$

or checks the recorded rational Bernstein subdivision. This bounds the whole interval $0 \leq y \leq 1/2000$. Because $c < 1$ at each finite order, every removed factor $(1-c)^m$ is positive there.

As a worked endpoint, $g = (1, 1, 1, 2d-3)$ has

$$\mathcal{D}_g(R) = 2s(2)^2 R(2) - 2s(1)^2 R(1) - s(1)s(3)[R(1) + R(3)].$$

Substituting the three exceptional solution values gives

$$\frac{\mathcal{D}_g(R)}{d} = \frac{(1-c)p(c)}{8\sin(2\theta)E(c)},$$

where the recorded p has degree 52 and $p(1) = 1414963200$. This is the first case with $J = -(1-c)p$ and multiplicity one. The same explicit substitution and the stored rational lower bound prove its strict sign on the whole interval, rather than only near $c = 1$.

The files `endpoint-sign-polynomials.json` and `endpoint-positivity.json` provide the exact coefficients and bounds for every endpoint. The verifier reconstructs all 552 expressions; 20 vanish identically and are checked to belong to $\mathcal{T}_d$, while all other 532 have positive reduced numerators. Finally $\cos(2\pi/n) > 1999/2000$ for $n \geq 202$, so the certified interval contains every required order. □

Lemma A.5 (equality cases)**.** *Let $d \geq 101$ be odd and let R be the potential constructed in Subsection A.1. For every positive odd gap tuple $g = (g_1, g_2, g_3, g_4)$ with sum $2d$, one has $\mathcal{D}_g(R) = 0$ if and only if its dihedral type belongs to $\mathcal{T}_d$.*

Proof. The selected types have zero slack by the identities in Subsection A.1. Lemmas A.2–A.3, the exceptional-to-critical reduction and Proposition A.4 give nonnegativity for every type. We identify its equality cases.

First consider the enumerated families with two fixed critical slots. Proposition A.4 classifies the retained event parameters and interval endpoints: their slack is zero precisely at selected types. Singleton intervals are settled by the same checks. On every remaining regular interval, all varying folded distances are ordinary. Every anchor has at least three exceptional sides, and every band type has an exceptional short side 3 or 5 opposite an exceptional long side $d-4$ or $d-2$.

Adjacent fixed gaps. This includes a fixed side and a fixed diagonal, which determine two adjacent gaps. At least one of the required exceptional sides of any selected type would vary, so a selected type cannot be an endpoint of such a regular interval. Both endpoint slacks are therefore strictly positive. The positive sine interpolation in Lemma A.3 gives strict positivity throughout.

Opposite fixed gaps. If neither endpoint is selected, both endpoint slacks are positive and Lemma A.3 gives strict positivity throughout. If an endpoint is selected, it cannot be an anchor, since at least one of its three exceptional sides would vary. It must be a band type whose short and long exceptional sides are the two fixed gaps. The two remaining gaps are ordinary and below d.

After exchanging them if necessary, they lie in the corresponding band range of Table 2; thus the entire regular interval consists of types in that band, whose slacks vanish by the band identities. The deleted second-band type has a varying side 9 and adjacent diagonal $3 + 9 = 12$, so it is an excluded exceptional event rather than a point of a regular interval. It is not restored by this argument.

Fixed diagonals. All four sides vary, so no selected type can be an endpoint of a regular interval. Both endpoint slacks are strictly positive. The concavity or monotonicity established in Lemma A.3 then gives strict positivity throughout.

This proves the classification on all two-critical-slot families. Types with no exceptional distance have positive slack by Lemma A.2. For any remaining type, apply the preliminary adjacent-gap reduction that produces a second critical slot. Its retained points belong to the families already classified. If the original type lies strictly inside one of its regular intervals, the adjacent-gap argument above excludes selected types at both endpoints; their slacks are positive, and positive sine interpolation gives positive slack at the original type. All remaining points are retained and already classified. This proves the asserted equality set. □

A.3 Independence of the contact evaluations

Write $A_{g_2}, B_{g_2}, C_{g_2}, D_{g_2}$ for the functionals $\mathcal{D}_g$ with gap tuples

$$A_{g_2} : (3, g_2, d-4, d+1-g_2), \quad B_{g_2} : (3, g_2, d-2, d-1-g_2),$$
$$C_{g_2} : (5, g_2, d-4, d-1-g_2), \quad D_{g_2} : (5, g_2, d-2, d-3-g_2).$$

The index g_2 runs over the odd integers specified in Table 2, with $g_2 = 9$ omitted in the second band. These equations are explicit six-term linear equations in the values $R(x)$, using (4.3). For example, the A_{g_2} equation is

$$\begin{aligned} 0 = {} & s(g_2+3)s(g_2+d-4)[R(g_2+3) + R(g_2+d-4)] \\ & - s(3)s(d-4)[R(3) + R(d-4)] \\ & - s(g_2)s(d+1-g_2)[R(g_2) + R(d+1-g_2)]. \end{aligned}$$

For a regular A_{g_2} row the displayed six slots fold to

$$3, \quad g_2, \quad d-4, \quad d+1-g_2, \quad g_2+3, \quad d+4-g_2.$$

For example, at $g_2 = 13$ the two varying side distances are 13 and $d - 12$, and the two folded diagonal distances are 16 and $d - 9$. The folded coefficient of $R(g_2)$ is $-s(g_2)s(d + 1 - g_2) \neq 0$, so once the other distances in this row are known it determines that new coordinate. The next rows give the other pivots listed below; each is checked in its indicated ordinary or exceptional branch. All distances are folded to $1, \ldots, d$. Apply the rows in increasing band index, in the order A, B, C, D when present. After the preceding rows, the new distances are $g_2, d + 2 - g_2, g_2 + 5, d - 3 - g_2$ respectively. Their coefficients are nonzero sine products. The deliberately omitted B_9 leaves a free boundary coordinate.

The eight anchors and band rows through $g_2 = 11$ give 22 equations on 27 distance coordinates. A **left label** writes a distance as $x = k$ and a **right label** as $x = d + k$. These are two propagated descriptions of a folded coordinate, not two different coordinates in the matrix. Before the recursions meet, solve for their coefficients separately; at an overlap require equality of the two values for the actual distance x. A **formal row** is the explicit six-term equation after substituting these affine labels, the chosen ordinary or exceptional branches, and $w^d = \mathrm{i}$, before imposing the matching equations. Thus each branch and every coefficient in the following model is specified. Use the form

$$R(x)/\mathrm{i} = c_{\ell,p} + (a_{\ell,p} + b_{\ell,p}k)\cot(\theta x)/\mathrm{i},$$

where ℓ is the left or right label, p its parity, and $k = x$ on the left and $k = x - d$ on the right. Include 17 independent exceptional values. Thus the model has 29 coefficients. The row equations force its four slopes to one common b and give

$$c_{L,\mathrm{even}} + c_{R,\mathrm{even}} = c_{L,\mathrm{odd}} + c_{R,\mathrm{odd}}.$$

Proposition A.6 (local solution and closing certificates)**.** *The local 22-row matrix has rank 22. The 29-coefficient model has five independent rational-function solutions satisfying every formal anchor and band equation. Their restriction to the local 27 coordinates is injective. The three closing constraints below have rank three on this five-dimensional space for either residue class of odd $d \geq 101$.*

Proof using certificates. In `uniform-type-rank.json`, `local_variables` lists the ordered affine coordinates (r, k) representing $rd+k$, and `local_columns` selects the 22-column minor of the matrix defined above. `local_determinant` and `local_arc` give its nonvanishing proof. `basis_numerators` lists five vectors of 29 polynomial numerators. The first twelve coordinates are (c, a, b) (stored as (C, A, B) in the certificate) for left-even, left-odd, right-even, right-odd in that order; the final seventeen are the exceptional values in (A.2). The five parameter coordinates, numbered from zero, are $(0, 1, 2, 3, 6)$, and the five basis vectors restrict to the identity matrix there. `basis_denominator` is their common polynomial denominator. Coefficient arrays are in descending powers of w. The verifier reconstructs the model and its parameter coordinates, checks all five normalized vectors, all 360 row identities obtained from the explicit formulas, the denominator, and the two closing minors. The certificate variable z denotes the parameter s_{arc}. All arc checks use $w = (1 + \mathrm{i}s_{\mathrm{arc}})/(1 - \mathrm{i}s_{\mathrm{arc}})$, $0 < s_{\mathrm{arc}} \leq 1/100$.

Derivation of exhaustiveness and closing. If a model solution vanishes on the local coordinates, triangular propagation through $g_2 = 13$ makes three left bulk values of each parity vanish; three right values are already local. The functions $1, \cot(\theta x), k\cot(\theta x)$ are independent at three distinct bulk distances: multiplying by $\tan(\theta x)$ reduces the assertion to strict convexity of tangent on $(0, \pi/2)$. Hence all branch coefficients vanish, as do the exceptional values. Restriction is injective. Since the local nullity is $27 - 22 = 5$, the five model solutions exhaust it. Each subsequent nonzero pivot forces the same continuation, so they exhaust all labeled band solutions.

For $d = 4q + 1$, the final rows identify the odd labels at $2q - 1, 2q + 1$ and the even labels at $2q + 4$. For $d = 4q + 3$, they identify the even labels at $2q + 4, 2q + 6$ and the odd labels at $2q + 1$. Put

$$M_{p,0} = c_{L,p} - c_{R,p}, \qquad M_{p,1} = a_{L,p} - a_{R,p} + bd.$$

Equality of the two labels at distance x is $M_{p,0} + M_{p,1}\cot(\theta x)/\mathrm{i} = 0$. The two matches of one parity, with strictly different cotangents, force both its M coefficients to zero. The remaining equation is

$$\begin{cases} M_{\mathrm{even},0} + M_{\mathrm{even},1}\cot(\theta(d+7)/2)/\mathrm{i} = 0, & d \equiv 1 \pmod 4, \\ M_{\mathrm{odd},0} + M_{\mathrm{odd},1}\cot(\theta(d-1)/2)/\mathrm{i} = 0, & d \equiv 3 \pmod 4. \end{cases}$$

On the four-dimensional subspace $b = 0$, the certificate proves a selected 3×3 closing minor is nonzero in either case. Thus the full closing map has rank at least three. Its nullity is at most two; the constant potential and the constructed potential of nonzero slope give two independent surviving solutions. The full kernel therefore has dimension two. Since there are d potential coordinates and $d - 2$ rows, those rows are independent. Finally $\tan(\pi/(2n)) < 1/100$ for $n \geq 202$, so every required order lies on the certified arc. □

B Character ranks

Throughout Appendix B, $n = 2d$, $d \geq 101$ is odd, $\theta = \pi/n$, $w = e^{\mathrm{i}\theta}$ and $\rho = e^{2\pi \mathrm{i} j/n}$ for a character j. Thus $w^d = \mathrm{i}$ and $\rho^d = (-1)^j$. Here a formal character row is the quadratic evaluation defined in Section 5, expressed in affine support indices and folded distance coordinates, with the left/right labels of Subsection A.3 and the two phase relations just stated. Every displayed rational-function identity is first checked with these substitutions, retaining all scalar factors. The nonvanishing arguments then justify its specialization to the actual roots of unity.

B.1 Explicit base rows, propagation and the resonant character

For a four-point support $I = \{p_0 < p_1 < p_2 < p_3\}$ with odd cyclic gaps put $\widetilde{z}_I = \Pi_I z_I$. Its coordinate at p_a is the product of the three sine chords of the complementary triangle. A character matrix has $H_{a,a+x} = \rho^a \mu_x$. The three matrices

$$\mu_x^{(s)} = (-1)^x (w^{2sx} + \rho^x w^{-2sx}), \qquad s = -1, 0, 1,$$

annihilate every base ray because they are symmetrized products of Fourier vectors, one in range P and the other at the complementary frequency.

Subtract half the diagonal coefficient times the $s = 0$ matrix, leaving $\mu_0 = 0$. For $x > 0$ put $\mu_x = (-1)^x \sin^2(\theta x) R_x$. In $\widetilde{z}_I^{\mathsf{T}} H \widetilde{z}_I / (2\Pi_I)$ the pair $p_a < p_b$, with complementary pair $p_u < p_v$, has coefficient

$$(-1)^{p_b - p_a} \sin(\theta(p_b - p_a)) \sin(\theta(p_v - p_u)) \rho^{p_a} \tag{B.1}$$

on $R_{p_b - p_a}$. If $p_b - p_a > d$, fold the distance and replace the phase by ρ^{p_b}. The complementary-chord identity preceding (4.3) gives this coefficient, since $\widetilde{z}_I(p_a)\widetilde{z}_I(p_b)/\Pi_I = \sin(\theta(p_v - p_u))/\sin(\theta(p_b - p_a))$. There is no antipodal coordinate in odd parity.

Apply (B.1) to the eight anchor tuples (4.4), each followed by its reversed tuple. This fixes the order of 16 rows. Their ordered potential coordinates are

$$1, 2, 3, 4, 5, 6, 10, 11, d-6, d-5, d-4, d-3, d-2, d-1, d,$$

with the last coordinate omitted for odd j.

Proposition B.1 (anchor certificates)**.** *Deleting columns R_1, R_2 gives a nonzero 13×13 minor in even parity, using rows $0, 1, 2, 3, 4, 5, 6, 8, 9, 10, 11, 12, 14$, and a nonzero 12×12 minor in odd parity, using rows $0, 1, 2, 3, 4, 5, 8, 9, 10, 11, 12, 14$, for every required nonzero character.*

Proof using certificates. The matrices obtained directly from the formulas, chosen indices and row scalings are in `code/runs/character-ranks/core-even.json` and `core-odd.json`. A row scaling recorded as `wshift=a`, `rshift=b`, `removed=f` means multiplication by $4w^a\rho^b/f(w)$. Every removed factor is retained in the factorization and is nonzero at the required roots. Up to nonzero constants, monomials and cyclotomic w factors of orders at most 24, the determinant factors are

$$\begin{aligned} \text{even: } & (\rho - 1)^5 (\rho - w^2)(\rho w^2 - 1)(2\cos(2\theta) - 1 - \rho), \\ \text{odd: } & (\rho - 1)^5. \end{aligned}$$

The order of w is $4d \geq 404$, so none of the listed cyclotomic factors vanishes. The real number $2\cos(2\theta) - 1$ lies in $(0, 1)$ and cannot equal ρ. The remaining even phase zeros correspond to $j = 0, \pm 1$, with ± 1 odd. Thus both chosen minors are nonzero. `verify_core.py` reconstructs the rows by (B.1), their scalings and determinants, and checks the factorizations exactly. □

To propagate from the anchors, use the following initial rows and the new coordinates they determine, where the band subscript is g_2:

$$A_5 : 8, \quad A_7 : 7, \quad B_7 : d-8, \quad C_7 : 12, \quad D_7 : d-10,$$
$$A_9 : 9, \quad C_9 : 14, \quad A_{11} : d-7, \quad D_9 : d-12,$$
$$B_{11} : d-9, \quad C_{11} : 16, \quad D_{11} : d-14.$$

The rows A_3 and B_5 are absent from this pivot list because they introduce no further distance coordinate: A_3 uses only anchor coordinates, and the sole additional coordinate of B_5, R_8, is already determined by A_5. Both remain among the selected contact equations.

Every indicated coefficient in (B.1) is a nonzero sine product times a unit phase. After these seeds, process each odd $g_2 \geq 13$ in the order A, B, C, D when present. Their new coordinates are $g_2, d+2-g_2, g_2+5, d-3-g_2$; every other coordinate in that row is already known. Once ranges meet, an already known coordinate needs no new determination.

For $d = 4q+1$, the final band indices are $(2q+1, 2q-1, 2q-1, 2q-1)$ in this order. The odd ranges overlap at $2q-1, 2q+1$, and the even ranges meet at $2q+4$. For $d = 4q+3$, the final indices are $(2q+1, 2q+1, 2q+1, 2q-1)$; the odd ranges meet at $2q+1$ and the even ranges overlap at $2q+4, 2q+6$. This reaches every distance, excluding the odd antipodal coordinate. Hence the base rank is at least $D_j - 3$. Except at $j = \pm 1$, the three range stresses have distinct unordered Fourier pairs and are independent, giving the matching upper bound.

At $j = 1$, so $\rho = w^2$, the $s = 1$ and $s = 0$ stresses coincide. Besides the actual difference $\mu^{\mathrm{diff}} = \mu^{(-1)} - \mu^{(0)}$, divide the formal difference $\mu^{(1)} - \mu^{(0)}$ by $\rho - w^2$ before specialization. At an affine distance $x = rd + k$ this gives a local solution

$$V_\mu(r,k) = (-1)^x k w^{-2}(1-w^{2x}), \qquad V_R(r,k) = \frac{k w^{-2}(1-w^{2x})}{\sin^2(\theta x)}.$$

Differentiating the identically zero row evaluations proves that this solves every formal row. The μ^{diff}, V determinant at distances 1 and 2 is $16(w-1)(w+1) \neq 0$, so they exhaust the two-dimensional local kernel after diagonal removal. At an overlapping distance $0 < x < d$ the two labels give

$$V_R(0,x) - V_R(1,x-d) = \frac{d w^{-2}(1-w^{2x})}{\sin^2(\theta x)} \neq 0.$$

The actual stress μ^{diff} agrees across labels, so the coefficient of V must vanish. Restoring the diagonal stress leaves a two-dimensional global base kernel and rank $d-2$. Conjugation treats $j = -1$. The identities and small determinant are checked by `verify_closing.py`; the explicit range argument above supplies the all-order conclusion.

B.2 Lift residuals and generic characters

Recall from Section 5 that $v_I = z_I + t_I$, where $t_I \in \operatorname{range} P$ agrees with $-Qz_I$ on I. We compute the character evaluations of the outer products $v_I v_I^{\mathsf{T}}$. Put $\widetilde{z}_I = \Pi_I z_I$ and $\widetilde{t}_I = \Pi_I t_I$, and write

$$\widetilde{t}_I(p) = (-1)^p (A + B\cos(2\theta p) + C\sin(2\theta p)), \quad \widehat{c}_0 = A, \quad \widehat{c}_1 = (B - \mathrm{i}C)/2, \quad \widehat{c}_{-1} = (B + \mathrm{i}C)/2.$$

Three independent support columns determine A, B, C; the fourth interpolation equation is checked as an identity. These coefficients are affine in d. The certificate names γ_I, t_{γ_I} correspond to $\widetilde{z}_I, \widetilde{t}_I$, respectively.

For a generic character, as defined in Section 5, $j \notin \{0, \pm 1, \pm 2\}$ and there is no range-by-range block. The three evaluation columns for the kernel–range block, after removing common nonzero column factors n and the row factor $2\Pi_I$, are

$$L_s = \frac{\widehat{c}_{-s}}{\Pi_I} \sum_{p \in I} \widetilde{z}_I(p)(-1)^p \rho^p w^{-2sp}, \qquad s = -1, 0, 1. \tag{B.2}$$

Here $s = f - d$ is the frequency offset. Thus the column indexed by s corresponds to the frequency $f = d + s$ in Section 5, with the row and column scalings described above.

The base columns may be represented by their projections onto the kernel-by-kernel block, which preserves their evaluations on z_I; thus they do not acquire unwanted lift terms.

Let $M_\pm$ be the scaled 16-row anchor matrix from Subsection B.1. The row vectors in `relations-even.json` and `relations-odd.json` form matrices Z_+ of size 3×16 and Z_- of size 4×16, with $Z_\pm M_\pm = 0$. They are primitive polynomial representatives of the left kernels; all rational units are retained in applying them. Scale each row of (B.2) by the same factor as its base row and each lift column by the common nonzero factor ρ^{12}. This defines a 16×3 matrix $L_\pm$ unambiguously. The residual matrix $Z_\pm L_\pm$ has zero constant coefficient in d and equals $d\,T_\pm(w, \rho)$.

For the even matrix select residual rows 0 and 2, and lift columns $s = -1, 0$ (column indices 0,1). For the odd matrix select row pairs 0,1 and 0,3, with those same columns. The corresponding `lift-minor-...json` files encode each determinant of T as

$$\det T[\text{rows}, \text{columns}] = \frac{a(w)P(w, \rho)}{b(w)},$$

where the fields `content`, `primitive`, `denominator` are exactly a, P, b. The determinant for the actual residuals has the further factor d^2.

Proposition B.2 (even lift certificate)**.** *The selected even residual minor is nonzero at every nonzero even character other than $j = \pm 2$, for odd $d \geq 101$.*

Proof using certificates. Its primitive polynomial factors as

$$\begin{aligned} P(w, \rho) = {} & 4\rho^{25}(w^4 - \rho)(\rho - 1)^3(\rho + 1)(\rho w^4 - 1) \\ & \cdot (w^4 - w^2 + 1 - \rho w^2)\, w^{20} \rho^4 G_{\text{even}}(c, \chi), \end{aligned}$$

where $c = (w^2 + w^{-2})/2$ and $\chi = (\rho + \rho^{-1})/2$. The variable χ is stored as t in the polynomial certificates. The polynomial G_{even} has degrees $(10, 4)$. The Bernstein certificate on two subrectangles in `even-angle-polynomial.json` proves

$$G_{\text{even}}(c, \chi) > 341$$

on $[1999/2000, 1] \times [-1, 1]$. All displayed phase factors are nonzero in this case: $\rho = -1$ is the odd character $j = d$, and $w^4 - w^2 + 1 - \rho w^2 = w^2(2c - 1 - \rho)$ with $0 < 2c - 1 < 1$. The scalar $a(w)$ and denominator $b(w)$ are nonzero by Proposition B.4 below. □

Proposition B.3 (odd lift certificate)**.** *At every odd character other than $j = \pm 1$, the two selected odd residual minors cannot vanish simultaneously.*

Proof using certificates. Remove the explicit common phase factor $\rho^{25}(\rho - w^2)^2(\rho w^2 - 1)^2$ from their primitive polynomials. The remaining polynomials F, G, defined in `odd-angle-01.json` and `odd-angle-03.json`, have bidegrees $(152, 4)$ and $(204, 12)$ in (w, ρ). Their resultant

$$\mathcal{R}(w) = \text{Res}_\rho(F, G)$$

is a nonzero polynomial of degree 2,620. Proposition B.4 verifies its exact coefficients and proves that it and both scalar numerator/denominator pairs are nonzero at every required w. A common root in ρ would make this resultant zero. The displayed phase factor vanishes only at $j = \pm 1$, which are excluded here. □

Two independent residual columns increase the even anchor rank from 13 to 15 and the odd anchor rank from 12 to 14. Each propagation pivot then adds one independent row and distance column: that new distance column vanishes on the anchors and earlier pivot rows. Therefore the gain of two dimensions persists through all distances and reaches $D_j - 1$. Section 5 gives the matching ceiling from positive-diagonal scaling equations.

B.3 Resultants and all-order root exclusions

Proposition B.4 (exact elimination and nonvanishing)**.** *The stored resultant is the exact resultant of the two polynomials in Proposition B.3. It, and every scalar content and denominator used in Subsections B.2 and B.4, is nonzero at $w = e^{\pi \mathrm{i}/(2d)}$ for every odd $d \geq 101$.*

Proof using certificates. **Resultant identity.** The degree of the resultant is at most $12 \cdot 152 + 4 \cdot 204 = 2{,}640$. Clear scalar rational denominators. If L_F, L_G are the sums of the absolute integer coefficients of the two input polynomials, the Sylvester determinant has every coefficient bounded by $L_F^{12} L_G^4$: expand the determinant and use the row sums of polynomial coefficient norms. Add the largest absolute coefficient of the candidate resultant to bound every coefficient of a possible error polynomial.

At each of 14 primes whose primality is verified by trial division, the independent verifier computes resultants by a separate finite-field Euclidean algorithm and compares 2,641 distinct values of w, avoiding zeros of the input leading coefficients. The degree bound proves equality of the two polynomials modulo that prime. The product of the primes has 419 bits and exceeds twice the exact coefficient error bound, whose threshold has 408 bits. Every possible error coefficient is divisible by that product and smaller than half of it in absolute value, hence is zero. This proves the integer polynomial identity deterministically.

Nonvanishing. If a rational polynomial of degree D vanishes at this primitive root of unity of order $4d$, it is divisible by Φ_{4d}, whose degree is $2\varphi(d)$. For odd prime powers,

$$\frac{\varphi(p^a)^2}{p^a} = p^{a-2}(p-1)^2 \geq 1.$$

Multiplication gives $\varphi(d) \geq \sqrt{d}$ for odd d. Thus only $d \leq (D/2)^2$ with $2\varphi(d) \leq D$ can be exceptions. For a polynomial over $\mathbb{Q}(\mathrm{i})$, first multiply it by its coefficient conjugate to obtain a rational polynomial.

For each candidate d, the certificate supplies a prime $p \equiv 1 \pmod{4d}$ and an element $\omega \in \mathbb{F}_p$ of exact order $4d$ where the primitive integer polynomial has nonzero value. Exact-order checks use the prime divisors of $4d$. Such a value excludes divisibility by the monic integer polynomial Φ_{4d}, and therefore excludes the complex zero. For the degree-2,620 resultant, the conditions $101 \leq d \leq 1{,}716{,}100$, d odd, and $2\varphi(d) \leq 2{,}620$ leave exactly 803 candidates. Every candidate has a certified nonzero finite-field evaluation. The same construction covers all eleven polynomial exclusions used for contents, denominators and the special characters, with 1,660 finite-field nonvanishing certificates in total. The independent checker uses its own totient sieve and trial primality tests. □

The inputs, coefficient bounds, primes and nonvanishing certificates are in the character directory's `odd-resultant.json` and `*-cyclotomic.json` files. The checker `verify_resultant.py` reconstructs the coefficient bounds and checks the resultant identity; `verify_cyclotomic.py` checks the cyclotomic root exclusions.

B.4 The special characters and the total dimension

For $j = 1$ and $j = 2$, the missing quadratic terms on the range block are, respectively,

$$\widehat{c}_0\widehat{c}_{-1}/\Pi_I, \qquad \widehat{c}_{-1}^2/\Pi_I,$$

up to the nonzero column factors n^2 or $n^2/2$ and the already specified row scalings. They are quadratic terms of the actual scaled lift $\widetilde{t}_I$.

Proposition B.5 (special-character certificates)**.** *At $j = 1$, odd relation 0 applied to the quadratic term has residual $d\,a_1(w) + d^2 b_1(w)$ with $b_1(w) \neq 0$ and $a_1(w)/b_1(w) \notin \mathbb{R}$. At $j = 2$, even relations 0 and 2 applied to the $s = -1$ cross term and quadratic term give a determinant $d^2(C_2(w) + dC_3(w))$, with $C_3(w) \neq 0$ and $|C_2(w)/C_3(w)| < 100$.*

Proof using certificates. `special-character-1.json` and `special-character-2.json` define the explicit interpolation and quadratic coefficients. The identities in `j1-nonreal-test.json` preserve the full rational ratio and compute its difference from its conjugate, where conjugation sends w to $1/w$. Proposition B.4 excludes zeros of b_1 and of that difference. Thus $d + a_1/b_1$ cannot vanish for real d.

For $j = 2$, `j2-rank-test.json` gives C_2, C_3 and the real quotient $N(c)/D(c)$. Its rational Taylor bounds prove $D > 0$, $100D + N > 0$ and $100D - N > 0$ on $1999/2000 \le c \le 1$. Since $d \ge 101$, the determinant is nonzero. These residuals annihilate all base columns, including distance columns absent on the anchors, so the gains persist globally. The character-1 base rank $d - 2$ gains one; the character-2 base rank $d - 2$ gains two. Conjugation treats $-1, -2$. □

At the trivial character, projection by $H \mapsto \mathrm{pr}_K\, H\, \mathrm{pr}_K$ transfers the $d - 2$ independent base evaluations to the lifted evaluations, as in the proof of Lemma 5.3. There is at most one independent averaged outer product per dihedral type, so the rank is exactly $d - 2$. Combining Propositions B.1–B.5 with the ceilings in Lemma 5.2 gives

$$r_0 = d - 2, \qquad r_j = d \quad (j \neq 0 \text{ even}), \qquad r_j = d - 1 \quad (j \text{ odd}).$$

Their sum is $(d-2) + (d-1)d + d(d-1) = \binom{2d}{2} - 2$. The equivariant rescaling in Lemma 3.3 transfers this rank to actual positive-parameter zeros.

C One axial column and an opposite pair

C.1 Uniform estimates

Throughout this subsection put $\psi = \pi/d = 2\theta$, so $c = \cos\psi$.

Use the Gram columns $(\pm 1, \boldsymbol{u}_j)$ of Section 6, where

$$\boldsymbol{u}_j = (\cos(2\pi j/d), \sin(2\pi j/d)).$$

A circuit involving the new column $(1, 0, 0)$ has three old spatial directions surrounding the origin. The corresponding spatial barycentric masses are each less than $1/2$: equality or a larger mass contradicts the strict triangle inequality for the three distinct unit vectors. The odd polygon has no antipodal pair. Thus the old Gram columns are either all on the negative rail or two are negative and one is positive. Conversely these assignments give positive dependencies after normalization of the new coordinate to one. The three old columns are independent, since their unique spatial relation has nonzero vertical sum.

Keep the raw zero-diagonal symbol Q_0 distinct from $Q = Q_0 - \gamma\sigma\sigma^{\mathsf{T}}$, where $\sigma_i = (-1)^i$ and $\gamma = \sigma^{\mathsf{T}}Q_0\sigma/n^2$. Then $Q\sigma = 0$. For a normalized new circuit $(x, 1)$, write $\ell = \sum_i x_i$; its vertical equation is $\sigma^{\mathsf{T}}x = -1$. Hence $x^{\mathsf{T}}Qx = x^{\mathsf{T}}Q_0x - \gamma$, a common shift for every such circuit. Recall the parameters of equation (6.1):

$$M = \min_{\text{all-negative circuits}} x^{\mathsf{T}}Q_0x, \qquad b = \max_{\text{mixed circuits}} \frac{M - x^{\mathsf{T}}Q_0x}{2(\ell - 1)}, \qquad c_* = -M + \gamma - 2b.$$

The first-order extension has quadratic form

$$F(x, t) = x^{\mathsf{T}}Qx + 2bt\sum_i x_i + c_*t^2.$$

It is nonnegative on every base circuit, with equality on the minimizing negative orbit and on at least one mixed orbit. We next establish $b > 0$, $c_* < 0$ and uniqueness of the negative orbit.

The symbol, with the corrections defined in Subsection A.1, is

$$Q_0(x) = \tfrac{1}{2}(-1)^x(dk_x - x)\sin(\pi x/d) + d\sin(\pi/d)\delta_x, \quad k_x = v \ (x \text{ even}),\ k_x = u \ (x \text{ odd}). \tag{C.1}$$

Proposition C.1 (uniform axial estimates)**.** *For* $2047/2048 \le c = \cos(\pi/d) \le 1$ *the rational functions satisfy* $0.448 < u < 0.451$, $0.896 < v < 0.900$, $|\delta_x| < 1$, $|\delta_{\text{even}}| < 0.018$, $|\delta_2| < 0.004$, *and* $-0.006 < \delta_{d-1} < -0.003$. *For every odd* $d \ge 101$, *the unique minimizing negative orbit has spatial gaps* $(1, (d-1)/2, (d-1)/2)$, *and* $M > 0$, $b/d > 0.035$, $\gamma/d < 0.063$.

Proof using certificates. All decimals denote rationals. The 57 polynomial sign checks in `code/runs/axial-extensions/axial-signs.json`, including denominator signs, give the displayed function bounds by Taylor expansion at one and Bernstein coefficients. The checker is `code/tests/axial-extensions/verify_axial_signs.py`. Since $\psi < 1/32$, the cosine interval contains every required order.

For spatial gaps $g_1, g_2, g_3 < d/2$ summing to d, put $\phi_i = \psi g_i$. The barycentric masses are proportional to $\sin(2\phi_i)$. Using $\sum_i \sin(2\phi_i) = 4\prod_i \sin\phi_i$ in (C.1) gives the raw circuit value $x^{\mathsf{T}}Q_0x$, with axial mass normalized to one:

$$d\prod_{i=1}^{3}\cot\phi_i\left(\tfrac{3v}{2} - 1 + \sum_{i=1}^{3}\frac{\sin\psi\,\delta_{2g_i}}{\sin(2\psi g_i)}\right). \tag{C.2}$$

The parenthesis lies between $K_- = 3(0.896)/2 - 1 - 3(0.018) > 0$ and $K_+ = 3(0.900)/2 - 1 + 3(0.018)$. Define the complementary odd integers $r_i = d - 2g_i > 0$. They sum to d, and $\cot\phi_i = \tan(\psi r_i/2)$. At fixed sum less than $\pi/2$, a product of two tangents increases as their arguments move toward equality, by the cosine sum-and-difference formula. Successively transferring two units from a smaller odd part to a larger one therefore reduces the product. Its minimum is $(r_1, r_2, r_3) = (1, 1, d-2)$ up to permutation; among the remaining triples it is $(1, 3, d-4)$. Their ratio is

$$\frac{\tan(3\psi/2)}{\tan(\psi/2)}\frac{\tan\psi}{\tan(2\psi)} \ge \tfrac{3}{2}(1 - \tan^2\psi) > K_+/K_-.$$

The last inequality follows from $\tan\psi \le \psi/(1 - \psi^2/2)$ and $\psi < 1/32$. Thus even the correction in (C.2) cannot change the minimizing orbit. Its masses are proportional to $(1, 1, 2c)$. We use $\operatorname{sinc}(x) = \sin(x)/x$ for real $x \ne 0$, with $\operatorname{sinc}(0) = 1$. Put $A_0 = Q_0(2)/\pi$, $B_0 = Q_0(d-1)/\pi$, $M_0 = M/\pi$. Direct substitution gives

$$A_0 = (v - 2/d)\operatorname{sinc}(2\psi) + \delta_2\operatorname{sinc}(\psi), \quad B_0 = ((v - 1 + 1/d)/2 + \delta_{d-1})\operatorname{sinc}(\psi), \quad M_0 = \frac{A_0 + 4cB_0}{2(1+c)^2}.$$

In particular $A_0 > 0.87$, $-0.059 < B_0 < -0.047$ and $M > 0$. For the mixed competitor set

$$g_1 = \lfloor (d+1)/4 \rfloor, \qquad g_2 = \lceil (d+1)/4 \rceil, \qquad g_3 = (d-1)/2,$$

and use spatial vertices $0, g_1, g_1 + g_2$. Choose the positive-rail column at g_1, opposite the last gap, and the negative-rail columns at the other two vertices. In the original cyclic indexing, the columns $(1, \boldsymbol{u}_j)$ and $(-1, \boldsymbol{u}_j)$ have indices $2j$ and $2j + d$ modulo $2d$, respectively. Define

$$\kappa_r = 2g_r - d/2 \quad (r = 1, 2), \qquad (\kappa_1, \kappa_2) = \begin{cases} (-1/2, 3/2), & d \equiv 1 \pmod 4, \\ (1/2, 1/2), & d \equiv 3 \pmod 4. \end{cases}$$

Thus $2\psi g_r = \pi/2 + \kappa_r \psi$ and $\kappa_1 + \kappa_2 = 1$.

Normalize the new axial coefficient to one. Write m_+ for the positive-rail coefficient, and y_0, y_2 for the negative-rail coefficients at vertices $0, g_1 + g_2$. The denominator D_0 normalizes their vertical equation $y_0 + y_2 - m_+ = 1$; spatial cancellation gives

$$D_0 = \cos(\kappa_1 \psi) + \cos(\kappa_2 \psi) - \sin \psi, \quad m_+ = \frac{\sin \psi}{D_0}, \quad y_0 = \frac{\cos(\kappa_2 \psi)}{D_0}, \quad y_2 = \frac{\cos(\kappa_1 \psi)}{D_0}.$$

For $r = 1, 2$, the cross-rail matrix distance is $d - 2g_r = d/2 - \kappa_r$, an ordinary odd distance. The quantity W_{κ_r} is m_+ times the corresponding raw matrix entry, divided by π. If x denotes the old coefficient vector of this circuit, set $q_0 = x^{\mathsf{T}} Q_0 x/\pi$. The negative pair has matrix distance $d - 1$, so

$$W_\kappa = \frac{\operatorname{sinc}(\psi)(1/2 - u - \kappa/d)\cos(\kappa\psi)}{2D_0}, \qquad q_0 = 2B_0 y_0 y_2 + 2y_0 W_{\kappa_1} + 2y_2 W_{\kappa_2}.$$

The old total mass is $\ell = y_0 + m_+ + y_2 = 1 + 2m_+$. Substituting this one competitor in the maximum defining b yields

$$\frac{b}{d} \geq \frac{\pi(M_0 - q_0)}{4dm_+} = \frac{(M_0 - q_0)D_0}{4\operatorname{sinc}(\psi)} > 0.035.$$

This last bound uses $2(1 - (3\psi/2)^2/2) - \psi \leq D_0 \leq 2$, the corresponding bounds on y_0, y_2, and the rational inequalities above; the checker evaluates these bounds at $\psi \leq 1/32$.

Finally pair folded distances $x, d - x$, using $\sum_{x=1}^{d-1} \sin(\psi x) = \cot(\psi/2)$. If $w_x = 2$ except $w_d = 1$, then

$$\frac{\gamma}{d} = \frac{(u + v - 1)\cot(\psi/2)}{4d} + \frac{\sin \psi}{2d} \sum_{x \text{ exceptional}} w_x (-1)^x \delta_x < 0.063.$$

Use $\cot(\psi/2) \leq 2/\psi$, $\sum w_x \leq 34$ and $d \geq 101$. Consequently $M + 2b - \gamma > 0$, as required. The separate `axial-geometry.json` and `verify_axial_geometry.py` in the same run/test directories reconstruct the barycentric identities; the uniform inequalities follow from the continuous sign certificates. □

C.2 Positive parameter and dimension gains

The minimization argument uses the following curvature estimate. Every vector in the enlarged kernel is uniquely

$$(x, t) = (y, 0) + t(-\sigma/n, 1), \qquad y \in \ker P.$$

Since $Q\sigma = 0$, on $\sum_i x_i = 0$ its first-order value is $y^{\mathsf{T}} Q y + c_* t^2 < 0$ away from zero: $\sum_i y_i = 0$, the old negative-type inequality applies to y, and $c_* < 0$. Compactness on the unit sphere of this subspace and continuity give some $\tau > 0$ such that the same strict negativity holds on the enlarged

kernel intersected with $\ker L$, where $L(x,t) = \sum_i x_i + \tau t$. It persists when a sufficiently small term $2\lambda t \sum_{i:\,\sigma_i=1} x_i$ is added. Fix $0 < \eta < b$ so that this holds throughout $|\lambda| \le \eta$.

Let $S_\varepsilon = P + \varepsilon Q + O(\varepsilon^2)$ be the old dihedrally invariant family. On its negative-rail simplex set

$$\mu_\varepsilon = \min_{x \ge 0,\ \sum x_i = 1} x^{\mathsf{T}} S_\varepsilon x.$$

A minimizer cannot have four positive entries: those entries admit a nonzero direction preserving the spatial sum and total mass, with strictly negative curvature for small positive ε. At zero parameter the minimum is one and the spatial sum must vanish. Thus, for sufficiently small parameter, only the finitely many surrounding triples can minimize. On each triple the base block is positive definite; its simplex minimizer is analytic near its positive barycentric vector. Proposition C.1 gives a strict separation between its first-order value and those of all other negative-rail orbits. It follows that

$$\mu_\varepsilon = 1 + \varepsilon(M - \gamma) + O(\varepsilon^2),$$

and precisely the chosen orbit minimizes. Define the extension $\widehat{S}_\varepsilon(\lambda)$ to have old block S_ε, axial diagonal a_ε, and constant couplings B_- and $B_+(\lambda)$ on the negative and positive rails, respectively, where

$$B_- = -1 + \varepsilon b, \quad a_\varepsilon = B_-^2/\mu_\varepsilon = 1 + \varepsilon c_* + O(\varepsilon^2), \quad B_+(\lambda) = 1 + \varepsilon(b + \lambda).$$

The block on the negative rail and the axial coordinate is copositive: writing $r = \sum x_i$, its value is at least

$$\mu_\varepsilon r^2 + 2B_- rt + a_\varepsilon t^2 = (\sqrt{\mu_\varepsilon}\, r + B_- t/\sqrt{\mu_\varepsilon})^2.$$

Equality occurs at the minimizing triples with the appropriate positive mass t.

A mixed zero at the threshold. Minimize the full quadratic on the compact simplex $L = 1$. A support larger than four admits a nonzero two-sided direction annihilated by both the three-row enlarged Gram matrix and L. Its second derivative is strictly negative, uniformly for $|\lambda| \le \eta$ and small positive ε. Every minimizer, and hence every zero of a copositive member, has at most four positive coordinates.

At $\lambda = \eta$ the full matrix is copositive. The old block and negative-rail block have already been handled. On any remaining support of size at most four, a base zero must be a mixed circuit containing the new column. Its simple kernel has strictly positive first-order value at η, so the corresponding block is positive definite for small ε. A support without a nonnegative base zero has positive minimum on its compact simplex, which persists. At $\lambda = -\eta$ a maximizing mixed circuit has negative first-order value and hence negative value for small positive ε.

Increasing λ increases the quadratic on the nonnegative orthant. Copositivity is closed, so the **copositivity threshold**

$$\lambda_\varepsilon = \min\{\lambda \in [-\eta, \eta] : \widehat{S}_\varepsilon(\lambda) \text{ is copositive}\}$$

belongs to $(-\eta, \eta]$. Moreover, $\lambda_\varepsilon \to 0$: for every fixed $0 < \delta < \eta$, the preceding endpoint arguments apply at $\lambda = \pm\delta$ for all sufficiently small $\varepsilon > 0$. Monotonicity gives $-\delta < \lambda_\varepsilon \le \delta$, proving the claim. Take negative normalized minimizers as $\lambda \uparrow \lambda_\varepsilon$. The support reduction confines them, after a subsequence, to one mixed four-coordinate support. No proper subsupport has a nonnegative base zero. Therefore, after choosing ε sufficiently small, these minimizers remain uniformly away from that support's boundary. Their limit is a zero with both positive-rail mass and positive new-coordinate mass. The threshold argument supplies a mixed zero even if several mixed circuit types attain the maximum in (6.1).

Each new zero lies on a rank-three base block whose three nonzero eigenvalues remain positive. Its perturbed singular block has a one-dimensional kernel. There are finitely many supports, and the old zero list is finite, so the full extension has finitely many zero rays. The construction preserves rotation of the d-gon.

The new cross entries span n dimensions. Let $i_1, i_2, i_3 \in \mathbb{Z}/d\mathbb{Z}$ be the spatial indices of the minimizing triple, with i_1, i_2 adjacent. Write its normalized weights, in this order, as $m_\varepsilon, m_\varepsilon, 1-2m_\varepsilon$. At zero, $m_0 = 1/(2(1+c))$. The difference between an adjacent and the opposite component of Qx is

$$m_0\pi(A_0 - 2(1-c)B_0) > 0.$$

The old base quadratic has positive Hessian on the simplex tangent. Differentiating stationarity therefore gives $m_0' \neq 0$. Put $\zeta = e^{2\pi \mathrm{i}/d}$. For $k = 0, \dots, d-1$, the Fourier coefficient of this triple is

$$m_\varepsilon(\zeta^{ki_1} + \zeta^{ki_2} - 2\zeta^{ki_3}) + \zeta^{ki_3}.$$

It is affine in m_ε, with nonzero constant term. If the coefficient of m_ε is zero, the Fourier coefficient equals the nonzero constant ζ^{ki_3}; otherwise $m_0' \neq 0$ shows that it has no zeros for sufficiently small positive ε, after possibly excluding the endpoint zero. In particular the spatial characters $k = \pm 1$, which vanish at zero, become nonzero. There are only d characters, so one common positive interval works for all. The cyclic translates span the entire negative rail. A mixed zero has one positive-rail coordinate; its translates span the other rail modulo the first. Normalizing the axial mass to one, projection of their outer products onto the old-to-new entries thus has rank n and kills all old outer products. The gain is at least n.

Opposite pair and duplication. Give both axial columns the same diagonal a_ε, exchange their rail couplings by one old-coordinate rotation, and set their mutual entry to $-a_\varepsilon$. Their coupling sum at every old coordinate is

$$\varepsilon(2b + \lambda_\varepsilon) > 0.$$

The reduction and dimension count in Lemma 6.2 prove copositivity, finiteness and the gain $2n+1$ for the opposite pair. Its application of Lemma 3.2 adds $n+1$ dimensions on duplication. Together with the one-column gain above, the three gains are n, $2n+1$ and $3n+2$.

D Finite coverage and order 205

D.1 Orders 14 through 201

The interval $14 \le n \le 201$ is covered by 184 negative-type constructions and four algebraic symmetric circulants. We first define the negative-type constructions: 94 at even orders $14, 16, \dots, 200$ and 90 at odd orders $23, 25, \dots, 201$. The remaining orders $15, 17, 19, 21$ are treated by the algebraic circulants below.

For n even or odd put $d = \lfloor n/2 \rfloor$ and $C_k(i,j) = \cos(2\pi k(i-j)/n)$. Let $J = \mathbf{1}\mathbf{1}^{\mathsf{T}}$ be the $n \times n$ all-ones matrix. Fix $0 < \kappa < 1$ and set

$$P = \kappa C_d + (1-\kappa)C_{d-1}, \qquad Q = J - 2\sum_{k=1}^{d-2} \xi_k C_k. \tag{D.1}$$

The base has rank three when n is even and rank four when n is odd. Its nonnegative kernel circuits are, respectively, four-point and five-point supports with all cyclic gaps odd. On such a support I their positive generators are

$$z_I(i) = \prod_{\ell \in I \setminus \{i\}} \sin(\pi|i-\ell|/n)^{-1} \quad (i \in I), \qquad z_I(i) = 0 \quad (i \notin I).$$

For even n, Lemma 4.1 gives the circuit supports and their generators. For odd $n = 2d + 1$, set $\omega_r = e^{2\pi \mathrm{i} r/n}$. The complex Gram row space of (D.1) consists of the four consecutive modes with exponents $d-1, d, d+1, d+2$. Multiplication of column r by $\omega_r^{-(d-1)}$ therefore gives four Vandermonde rows. Every four columns are independent. On five columns Lagrange interpolation, followed by the same sine-difference identity, gives a real dependence proportional to

$$\frac{(-1)^r}{\prod_{\ell \in I \setminus \{r\}} \sin(\pi(r-\ell)/n)}.$$

Here the phase depending on r is $e^{-\pi \mathrm{i} r} = (-1)^r$, since $n = 2d + 1$. Along the sorted support the denominator signs alternate, so the coefficients have one common sign exactly when all consecutive index differences are odd. The closing gap is odd as well, because n is odd and there are four consecutive differences. Taking absolute values gives the displayed positive-generator formula. Every circuit has five coordinates: smaller supports are independent, while a nonnegative kernel vector on a larger support admits a nonproportional kernel perturbation until a coordinate vanishes. Each circuit block consequently has a one-dimensional kernel.

Normalize each generator to total mass one and write $a_k(I) = 2|\sum_{r=0}^{n-1} z_I(r) e^{2\pi \mathrm{i} k r/n}|^2$. For each of these orders, the $d-2$ tight dihedral types are specified by the `tight_types` field of the corresponding input: `code/runs/finite-even-witnesses/certified/negative-type-n.json` for even n, and `code/runs/finite-odd-witnesses/certified/odd-negative-type-n.json` for odd n. The integer order replaces `n` in each filename. These types define the spectral vector ξ by the following square linear system.

$$\sum_{k=1}^{d-2} a_k(I)\xi_k = 1 \qquad (I \text{ tight}).$$

The enclosures prove unique solvability, $\xi_k > 0$ for every k, and $1 - a(I) \cdot \xi > 0$ for every omitted type. With the mass-one normalization above, the interval certificates give $1 - a(I) \cdot \xi > 3.5 \times 10^{-9}$ for every omitted type and $\xi_k > 6.4 \times 10^{-8}$ for every k, throughout these 184 constructions. The smallest recorded certified lower bounds for these two quantities occur at $n = 199$ and $n = 193$, respectively. Consequently Q is strictly negative on $\ker P \cap \mathbf{1}^\perp$, and is nonnegative on the nonnegative base kernel with precisely the specified first-order contacts as zeros. The nonsingular system also proves independence of the tight evaluation functionals on the invariant space of symmetric circulants.

For even orders the entering-gradient proof is the one given in Section 4. The same proof applies to the odd rank-four base: inserting a point into an odd-gap five-point circuit gives a six-point support with a sole even gap. Every four base columns are independent, so its kernel is two-dimensional. Deleting either endpoint of that even gap gives its two positive circuit generators, one of which contains the inserted point. The strict negative-type inequality and the identity $z^{\mathsf{T}} Q z' = g(j) z'(j)$ from Section 4 therefore give $g(j) > 0$ here as well. Together with the nonsingular tightness system, these verify all hypotheses of Lemma 3.3 in both parities.

Proposition D.1 (finite negative-type certificates)**.** *For each even* $14 \le n \le 200$ *and odd* $23 \le n \le 201$*, the certificate-specified tight types define a family as above whose zero outer products span at least* $\binom{n}{2} - 2$ *dimensions.*

Proof using certificates. The `arithmetic` fields specify interval precision and a prime/root specialization. The checkers `verify_negative_type_independent.py` and `verify_odd_negative_type_independent.py`, in the correspondingly named `code/tests/` directories, rebuild the trigonometric data, solve and enclose the systems, enumerate every other type, and prove the strict inequalities.

They also reconstruct the limiting character matrices of the zero outer products and find the recorded nonzero modular minors. The primitive-root and denominator checks make each nonzero minor a certificate in characteristic zero. The same equivariant rescaling and continuity argument as in Section 5 transfers these minors to positive parameter. This proves the proposition. □

The four remaining orders use algebraic symmetric circulants rather than (D.1). Their diagonal is one. Write their first-row entries as $c_j = c_{n-j}$, $1 \le j \le \lfloor n/2 \rfloor$. The fixed entries and certified ranks are listed in Table 4.

Table 4: Certified algebraic circulant witnesses at orders 15, 17, 19 and 21.

n	Fixed entries	Certified outer-product span dimension r
15	$c_5 = 1/2$	104
17	$c_5 = 1/4$	135
19	$c_6 = -3/5$	170
21	$c_4 = -9/50$, $c_7 = 24/25$	208

For example, at order 15 the first row is

$$(1, c_1, c_2, c_3, c_4, c_5, c_6, c_7, c_7, c_6, c_5, c_4, c_3, c_2, c_1),$$

where

$$(c_1, \dots, c_7) \approx (-0.9140916,\ 0.6779640,\ 0.0240408,\ -0.3248302, \tfrac{1}{2},\ -0.3970986,\ 0.1563606).$$

For this order-15 matrix, the zero supports, up to rotations and reflection, have cyclic gaps

$$(1,1,3,5,5), \quad (1,1,5,1,7), \quad (1,1,5,3,5),$$
$$(1,3,1,3,7), \quad (1,3,3,5,3), \quad (1,3,5,1,5).$$

The third type has 15 distinct symmetry images and each other type has 30, giving 165 zero rays in total. The displayed decimals illustrate the matrix; its exact algebraic definition is the isolating-box system below.

All remaining c_j are the unique solution, in the rational box of radius 10^{-60} about the supplied center, of $\det S[I] = 0$ for the listed support types I. Centers, free coordinates and supports are in `code/runs/finite-circulant-witnesses/certified/circulant-n-hr-isolating-box-seed.json`, with (n, r) from Table 4. The companion file `circulant-n-hr-rank-minor.json` in the same directory specifies a nonzero minor.

Proposition D.2 (four algebraic circulants)**.** *These four systems define copositive matrices with finitely many zero rays and the ranks in Table 4.*

Proof using certificates. `code/tests/finite-circulant-witnesses/verify_isolating_box_parallel.py` proves a contraction on each rational box. It then enumerates every nonempty support up to cyclic rotation. On each support it solves the bordered stationarity equations $S[I]x = \lambda \mathbf{1}$, $\mathbf{1}^{\mathsf{T}} x = 1$, with rational Neumann bounds valid throughout the box. The bordered matrix is invertible. Every support is certified as one of: a listed zero with strictly positive coordinates; a stationary point with a strictly negative coordinate; or a stationary point with strictly positive value. The determinant equations make the listed values exactly zero. A negative simplex minimum would have

a nonnegative stationary vector on its positive support and negative value, which the exhaustive classification excludes. The same classification proves completeness and finiteness of the zero rays. Finally the checker encloses those rays and proves the specified outer-product minor nonsingular by an inverse perturbation bound smaller than one. These are algebraic certificates; the decimal centers merely specify the rational boxes. □

The two families and four circulants cover every integer $14 \le n \le 201$ exactly once. The starting order 14 is the endpoint of this coverage, without a sharpness claim. In particular the lower bounds 104, 135 and 170 at orders 15, 17 and 19 are stronger than $\binom{n}{2} - 2$.

D.2 Horizontal extension at order 205

The order-204 matrix of Subsection C.2 has outer-product span at least 20,704. Let U contain its 202 rail columns and two axial columns A, B. Add the horizontal column $\mathbf{c} = (0, 1, 0)$. Write the old family as $S_\varepsilon = U^\mathsf{T} U + \varepsilon Q_\varepsilon$, with $Q_\varepsilon \to Q$ by the expansions and the convergence $\lambda_\varepsilon \to 0$ proved in Subsection C.2. Here Q is the order-204 first-order matrix at $\lambda = 0$. Write the new border as $U^\mathsf{T}\mathbf{c} + \varepsilon f$ and its diagonal as $1 + \varepsilon s_0$. The **border coefficients**, or heights, are the entries f_i of this first-order border; they are distinct from the Gram coordinates in U.

The selected contacts. The 52 support triples in `code/runs/horizontal-extension-205/order205-contact-basis.json` each consist of three independent old columns. Solve $U_I x_I = -\mathbf{c}$, and then the 52 invariant equations $2f^\mathsf{T} x_I = -x_I^\mathsf{T} Q x_I$ with $s_0 = 0$. Reflection of spatial indices and exchange of rails leave 51 rail heights and one common axial height. Restoring both symmetries gives 208 distinct supports. Denote the initial height solution by f^0.

For each selected support, solve

$$Q_{\varepsilon,I} x + f_I = U_I^\mathsf{T} y, \qquad U_I x + \varepsilon y = -\mathbf{c}. \tag{D.2}$$

For a selected support I, fix the new coordinate to 1 and minimize over the old coefficients $x \in \mathbb{R}^I$. To explain the dependence on $\mathbf{c}$, temporarily write the new diagonal as $\|\mathbf{c}\|^2 + \varepsilon s_0$; this agrees with $1 + \varepsilon s_0$ at the prescribed unit vector. The quadratic to minimize is

$$\Phi_{\varepsilon,I}(x) = \|U_I x + \mathbf{c}\|^2 + \varepsilon(x^\mathsf{T} Q_{\varepsilon,I} x + 2f_I^\mathsf{T} x + s_0).$$

Its old-coordinate Hessian is twice the positive definite old support block. Its unique minimizer remains strictly positive for small positive ε, since its limit solves $U_I x = -\mathbf{c}$ and is strictly positive. Introducing y by $U_I x + \mathbf{c} = -\varepsilon y$ gives precisely the stationarity system (D.2).

At $\varepsilon = 0$ the homogeneous system in (D.2) is invertible because U_I is invertible. Its solution (x, y) consequently depends smoothly on the independent variables $(\varepsilon, \mathcal{Q}, f, s_0, \mathbf{c})$ near $(0, Q, f^0, 0, (0, 1, 0))$, with $\mathcal{Q}$ subsequently set to Q_ε. The normalized minimum

$$M_I = \frac{\min_x \Phi_{\varepsilon,I}(x)}{\varepsilon} = \varepsilon\|y\|^2 + x^\mathsf{T} Q_{\varepsilon,I} x + 2f_I^\mathsf{T} x + s_0$$

extends smoothly to $\varepsilon = 0$ through the expression on the right. For fixed $\varepsilon, \mathcal{Q}$, differentiating the minimum and using $U_I x + \mathbf{c} = -\varepsilon y$ gives

$$dM_I = 2x^\mathsf{T}\, df + ds_0 - 2y^\mathsf{T}\, d\mathbf{c},$$

where x is extended by zero outside I. The formula at zero follows by continuity. The last term describes an auxiliary variation of the Gram column; the actual symmetric construction keeps $\mathbf{c} = (0, 1, 0)$ and $s_0 = 0$ fixed.

For the simultaneous continuation, the unknowns are the 51 invariant rail heights and the common axial height. Impose one equation $M_I = 0$ for each of the 52 selected support representatives. On these height variables the derivative has rows $2x_I^{\mathsf{T}}$, with coordinates grouped by the prescribed symmetries, and has rank 52. The implicit function theorem therefore supplies these heights as smooth functions of the independent parameters $(\varepsilon, \mathcal{Q})$. Substituting the prescribed $Q_\varepsilon \to Q$ continues the selected contacts as exact support zeros while leaving the old matrix fixed. This substitution requires no differentiability of Q_ε or of the parent threshold.

Reduction to two axial cases. Normalize the old axial diagonals and their mutual entry to $1, 1, -1$ by positive diagonal congruence. Let p_A, p_B be their rail couplings, and α, β their couplings to the new column. For rail vector x, old axial masses a_1, a_2, and new mass t, subtract $a_{\min} = \min(a_1, a_2)$ from both old axial masses. The quadratic decreases by

$$2a_{\min}((p_A + p_B)^{\mathsf{T}} x + (\alpha + \beta)t). \tag{D.3}$$

Here $p_A + p_B > 0$ by Subsection C.2. The certified common new height at A, B exceeds $1/25$, so $\alpha + \beta > 0$ for sufficiently small positive parameter, also after normalization. Thus it suffices to prove copositivity on the two branches obtained by omitting one old axial column. A zero using both old axial coordinates must be the old pure equal pair.

On each branch let x collect its 203 old coordinates. Use the positive functional L from Subsection C.2 and minimize on $L(x) + t = 1$. If at least five old coordinates are positive, there is a nonzero two-sided variation supported on them satisfying

$$U_{\text{branch}}\delta x = 0, \qquad L(\delta x) = 0, \qquad \delta t = 0.$$

The first two conditions impose at most four linear constraints on the old-coordinate variation. With t fixed, the new border terms have zero second derivative along this direction, whereas the old block has strictly negative curvature by Subsection C.2. This contradicts minimality. Hence every minimizer uses at most four old coordinates plus the new coordinate.

In the next proposition, let (x_I, y_I) be the solutions of (D.2) at $\varepsilon = 0$ for all 208 restored contacts, with each x_I extended by zero to the 204 old coordinates. Write $\mathrm{pr}_{\ker U}$ and $\mathrm{pr}_{\mathbf{c}^\perp}$ for the orthogonal projections. Use fixed bases in these subspaces to write their projected row coordinates.

Proposition D.3 (horizontal contact and enumeration certificates)**.** *The 52 defining contacts have positive masses. The coefficient matrix of their 52 contact equations in the 52 invariant border heights has rank 52. After restoring all 208 contacts, the matrix with rows x_I has rank 202, and the matrix with rows $(\mathrm{pr}_{\ker U}\, x_I, 1, \mathrm{pr}_{\mathbf{c}^\perp}\, y_I)$ has rank 204. On each branch the first-order form is nonnegative on every nonnegative base kernel with at most four old coordinates. Its new zeros are precisely the 170 selected contacts lying on that branch. All cross terms between distinct extreme rays in these blocks are strictly positive, and all selected contacts have strictly positive entering gradients.*

Proof using certificates. The directory `code/tests/horizontal-extension-205/` supplies `verify_contact_basis.py`, `replay_projected_rank.py`, `verify_binary_blocks.cpp` and `replay_contact_neighbors.cpp`. Inputs and reports are in the correspondingly named `code/runs/` directory. The basis reconstruction proves the positive masses, common height $> 1/25$, and entering gradients, normalized by $1 + \sum x_i$, greater than 0.0014. It separates the parent's maximizing mixed triangle, with gaps $25, 26, 50$, from the other 649 triangle orbits with a distinguished positive-rail vertex. The three ranks are certified modulo 536871157 and independently modulo 1010809; checks of exact root order 404 and denominator nonvanishing justify passage to characteristic zero.

There are $\binom{203}{4} = 68{,}685{,}050$ blocks on four old coordinates on a branch. Every such old block has rank three, so adjoining the new column gives a two-dimensional kernel. Its nonnegative cone

has at most two extreme rays. They are obtained by signed three-column cofactors after omitting a column, allowing zero cofactors. This includes smaller supports: in particular there are 50 circuits with just two old columns (25 spatial pairs and two rail assignments). It is only the selected continuation supports that require three independent old columns. A singular old triple containing an axial column has both rail columns at the same spatial position and cannot supply $-\mathbf{c}$; three rail columns are independent by the circle argument in Section 4.

The exhaustive block enumeration finds 169,275 distinct new base circuits, exactly 170 prescribed zero contacts, and 21,596,275 distinct strictly positive cross terms, including 8,842,550 old/new pairs. The order-204 family S_ε constructed in Lemma 6.2 and Subsection C.2 is copositive. Hence, for $z \geq 0$ with $Uz = 0$,

$$z^{\mathsf{T}} Q z = \lim_{\varepsilon \downarrow 0} \frac{z^{\mathsf{T}} S_\varepsilon z}{\varepsilon} \geq 0.$$

Consequently the first-order quadratic is nonnegative on every cone generated by at most two extreme rays and vanishes only on its listed zero rays. An independent neighboring-circuit enumeration on the other branch gives the same pair counts.

These inequalities use explicit integer error bounds. Sines have error at most 2^{-57}; each cofactor uses at most three signed sines. An unresolved sign is accepted as zero only after formal cancellation. For absolute cofactor centers of sum S in these units, normalizing and truncating on a grid $W = 2^{40}$ has ℓ^1 error at most $4/W + 24/(S - 12)$. First-order entries have error at most 2^{-39}. If two normalized mass errors are at most $e_1/W, e_2/W$ and B bounds the entry magnitudes in units of 2^{-39}, the integer bilinear evaluation has a certified upper bound for the absolute error of $2 + \lceil B(e_1 + e_2)/W \rceil$. Both truncated masses sum to at most one; $B < 2^{45}$ ensures the accumulators fit in signed 128-bit integers. Only the selected exact contact equations are treated as identically zero. The reports `order205-basis-proof.json`, `order205-block-census.json` and `order205-neighbor-replay.json` record the bounds and exhaustive counts. □

Why the finite checks imply global copositivity. Suppose there are nonpositive simplex minimizers at parameters tending to zero. The branch and support reductions give a subsequence on one fixed block. Its limit lies in the base kernel, with first-order value at most zero. If it uses the new column, Proposition D.3 makes it a selected contact. The continued exact zero has positive entering gradient off its support. At a simplex minimum of value $v \leq 0$, the half-gradient equals $v\nabla(L + t)$ on its positive coordinates. Pair this equation with the continued zero. Symmetry of the matrix makes the pairing nonnegative and strictly positive if extra coordinates occur. Thus $v = 0$, no extra coordinates occur, and the minimizer is the unique normalized zero of that block.

If the limit is an old tight circuit w, its four old coordinates exhaust the block. Choose the positive unit least eigenvector w_ε of the old support block tending to w. Its eigenvalue is nonnegative by old copositivity and is $o(\varepsilon)$. Writing $U_I^{\mathsf{T}} y_w = Q_I w$, the eigenvector equation gives $U_I w_\varepsilon/\varepsilon \to -y_w$. The new-coordinate gradient divided by ε tends to $f^{\mathsf{T}} w - \mathbf{c}^{\mathsf{T}} y_w > 0$. To see the strict sign, solve $U_I x = -\mathbf{c}$ and add a large multiple of w to make $x \geq 0$; the resulting cone generated by at most two extreme rays contains a new extreme ray. Its checked cross term against w is a positive multiple of this expression. Pairing the minimum equations with $(w_\varepsilon, 0)$ now excludes positive new mass. Old-only zeros are already known. Together with equation (D.3), this proves copositivity and completeness of the finite zero list.

The additional 204 dimensions. For every selected contact form the limiting row

$$(\mathrm{pr}_{\ker U}\, x,\ 1,\ \mathrm{pr}_{\mathbf{c}^\perp}\, y).$$

These rows have rank 204 by Proposition D.3. Indeed $(\mathrm{pr}_{\ker U}\, x, Ux)$ are invertible coordinates on the 204-dimensional old space. Equation (D.2) says that the longitudinal component of Ux tends

to -1 and the two transverse components are $-\varepsilon \operatorname{pr}_{\mathbf{c}^\perp} y$. Dividing those two columns by ε gives the displayed limit, up to signs. A nonzero limiting minor therefore proves full old-projection rank at small positive parameter. Projection onto the new cross entries kills all old outer products and adds 204 dimensions. The finite-zero lemma yields

$$p_{205} \geq 20{,}704 + 204 = 20{,}908 = \binom{205}{2} - 2.$$

E Certificate map and reproduction

All paths are relative to the extracted computational supplement. The archive contains computational material and instructions; all mathematical reductions and implications are proved in this paper.

The following map distinguishes the arithmetic assertion from its role in the proof.

- **Potential and all-order feasibility (Subsections A.1–A.2).** `verify_endpoint_coverage.py` reconstructs the family list in `endpoint-coverage.json` and verifies every affine inequality for all odd $d \geq 101$. `verify_uniform_potential.py` checks that coverage, the 72 defining identities obtained from the formulas, all 552 endpoint identities, the rational Taylor and Bernstein sign bounds, and the spectral formulas. Identities are formal rational-function equalities, while the sign checks hold on the stated continuous intervals and rectangles. Lemmas A.2–A.3 and the exceptional-to-critical reduction prove that these finite statements exhaust all four-point supports with odd cyclic gaps; the equality analysis then identifies exactly the selected zero types.

- **Contact independence (Proposition A.6).** `uniform-type-rank.json`, checked by `verify_uniform_potential.py --rank`, gives the local rank 22, five normalized model solutions, 360 row identities obtained from the explicit formulas and two rank-three closing minors. The nonvanishing tests hold on the complete root-of-unity arc. The propagation, injectivity and matching proof in Subsection A.3 turns these checks into rank $d - 2$ for every odd $d \geq 101$.

- **All character ranks (Propositions B.1–B.5).** The seven programs in `code/tests/character-ranks/` reconstruct the base matrices, lift interpolations, residual minors, special-character terms and their identities. Continuous polynomial bounds and complete cyclotomic exclusions prove their nonvanishing for every required order and character. The resultant checker proves the entire integer polynomial identity by the degree and coefficient bounds of Subsection B.3, then the exclusion checker tests every possible exceptional order allowed by that degree bound. The row propagation and rank ceilings in Section 5 and Appendix B prove the asserted character ranks; finite-order rank controls are additional tests only.

- **Axial extensions (Subsections C.1–C.2).** `axial-signs.json` and `axial-geometry.json` are replayed by `verify_axial_signs.py` and `verify_axial_geometry.py` in `code/tests/axial-extensions/`. The first checks 57 rational polynomial signs on $2047/2048 \leq \cos(\pi/d) \leq 1$ and the resulting scalar bounds; the second checks the specified geometric identities. The minimization, continuation, finite-zero and dimension arguments are the proofs in Appendix C, which apply to the normalized family of Section 6 for all odd $d \geq 101$.

- **Finite negative-type witnesses (Proposition D.1).** The even and odd witness checkers verify the 184 listed orders separately, using rational trigonometric enclosures for tightness,

unique solvability, spectral signs and every omitted type, and exact finite-field minors for contact and character ranks. Their root-order and denominator checks justify each characteristic-zero rank implication. Proposition D.1 and Lemma 3.3 supply the continuation and copositivity arguments. This is finite coverage, not a uniform extrapolation.

- **Four algebraic circulants (Proposition D.2).** `verify_isolating_box_parallel.py` in `code/tests/finite-circulant-witnesses/` verifies the four rational contraction boxes, every cyclic support orbit and the four outer-product minors. Proposition D.2 proves why the exhaustive stationary-point classification yields copositivity and a complete finite zero list at orders 15, 17, 19 and 21, with respective ranks 104, 135, 170 and 208.
- **Horizontal extension at 205 (Proposition D.3).** The four checkers in `code/tests/horizontal-extension-205/` reconstruct the 52-contact system, two modular rank witnesses, and the two exhaustive circuit/cross-term enumerations, including their integer error bounds. These are calculations at this one order. The branch reduction, support bound, continuation, old-limit case and projection proof in Subsection D.2 establish copositivity, completeness and the additional 204 dimensions from those calculations.

The computational supplement's `README.md` gives dependency setup and commands. Its `reproduce.py` runs checks in dependency order, saves fresh reports and logs, and resumes completed stages using input and output hashes. The manifest records the size and SHA-256 digest of every released file. Report input identifiers are relative to the archive root, with POSIX separators; timings are diagnostic.

Acknowledgments

The proofs in this paper were developed by OpenAI's GPT models (Sol and Astra) and Anthropic's Claude Fable models under the author's guidance. Responsibility for the final text rests with the author.

References

[1] J. H. Drew, C. R. Johnson and R. Loewy, "Completely positive matrices associated with M-matrices," *Linear Multilinear Algebra* **37** (1994), 303–310. doi:10.1080/03081089408818334.

[2] I. M. Bomze, W. Schachinger and R. Ullrich, "From seven to eleven: completely positive matrices with high cp-rank," *Linear Algebra Appl.* **459** (2014), 208–221. doi:10.1016/j.laa.2014.06.025.

[3] I. M. Bomze, W. Schachinger and R. Ullrich, "New lower bounds and asymptotics for the cp-rank," *SIAM J. Matrix Anal. Appl.* **36** (2015), 20–37. doi:10.1137/140973207; author preprint: `https://optimization-online.org/wp-content/uploads/2014/06/4392.pdf`.

[4] W. Schachinger, "Lower bounds for maximal cp-ranks of completely positive matrices and tensors," *Electron. J. Linear Algebra* **36** (2020), 519–541. doi:10.13001/ela.2020.5323.

[5] Y. Lavi, "A linear gap for the maximal cp-rank," (2026), arXiv:2610.02289.

[6] N. Shaked-Monderer, A. Berman, I. M. Bomze, F. Jarre and W. Schachinger, "New results on the cp-rank and related properties of co(mpletely)positive matrices," *Linear Multilinear Algebra* **63** (2015), 384–396. doi:10.1080/03081087.2013.869591; author preprint, arXiv:1305.0737v2. Numbering of cited results follows this preprint.

[7] N. Shaked-Monderer, I. M. Bomze, F. Jarre and W. Schachinger, "On the cp-rank and minimal cp factorizations of a completely positive matrix," *SIAM J. Matrix Anal. Appl.* **34** (2013), 355–368. doi:10.1137/120885759.

[8] R. Loewy and B.-S. Tam, "CP rank of completely positive matrices of order 5," *Linear Algebra Appl.* **363** (2003), 161–176. doi:10.1016/S0024-3795(01)00588-2.

[9] N. Shaked-Monderer, "On the DJL conjecture for order 6," *Oper. Matrices* **11** (2017), 71–88. doi:10.7153/oam-11-05; corrected author preprint, arXiv:1501.02426v3.

[10] P. J. C. Dickinson, "Erratum to: On the DJL conjecture for order 6," *Oper. Matrices* **11** (2017), 1197–1200. doi:10.7153/oam-2017-11-81.

[11] S. Holden, "PF-04: A proposed proof that the maximum cp-rank in order six is nine," manuscript, 12 September 2026. https://github.com/ajt60gaibb/OpenProblemsInNLA/blob/3fececca7c8dbdc594270e07626f6af7ff9e7f14/references/holden-pf04-2026-09-13/submission/PF04_proposed_proof.pdf.

[12] I. M. Bomze, P. J. C. Dickinson and G. Still, "The structure of completely positive matrices according to their CP-rank and CP-plus-rank," *Linear Algebra Appl.* **482** (2015), 191–206. doi:10.1016/j.laa.2015.05.021.

[13] A. Ferudun, "The Maximal Finite CP-Plus-Rank Equals the Maximal CP-Rank," EulerSolve Research Papers, preprint, version 1.0, 13 October 2026. doi:10.5281/zenodo.23350364.

[14] M. Ito and B. F. Lourenço, "A bound on the Carathéodory number," *Linear Algebra Appl.* **532** (2017), 347–363. doi:10.1016/j.laa.2017.06.043.

[15] M. Nishijima, "On the longest chain of faces of the completely positive and copositive cones," *Linear Algebra Appl.* **698** (2024), 479–491. doi:10.1016/j.laa.2024.06.012; arXiv:2305.13640.

[16] J. Hannah and T. J. Laffey, "Nonnegative factorization of completely positive matrices," *Linear Algebra Appl.* **55** (1983), 1–9. doi:10.1016/0024-3795(83)90162-3.

[17] N. Shaked-Monderer, "A note on upper bounds on the cp-rank," *Linear Algebra Appl.* **431** (2009), 2407–2413. doi:10.1016/j.laa.2009.04.015.

[18] A. Berman and D. Hershkowitz, "Combinatorial results on completely positive matrices," *Linear Algebra Appl.* **95** (1987), 111–125. doi:10.1016/0024-3795(87)90029-2.

[19] N. Shaked-Monderer, "Bounding the CP-rank by graph parameters," *Electron. J. Linear Algebra* **28** (2015), 99–116. doi:10.13001/1081-3810.3009.

[20] P. J. C. Dickinson and R. Hildebrand, "Considering copositivity locally," *J. Math. Anal. Appl.* **437** (2016), 1184–1195. doi:10.1016/j.jmaa.2016.01.063; author preprint: https://optimization-online.org/wp-content/uploads/2014/04/4315.pdf.

[21] M. Manainen, M. Seliugin, R. Tarasov and R. Hildebrand, "Generating extreme copositive matrices near matrices obtained from COP-irreducible graphs," *Linear Algebra Appl.* **693** (2024), 297–323. doi:10.1016/j.laa.2023.09.026.