Sharp bounds for indecomposable summands of graphical zonotopes

Caius Folkerts

Abstract. In dimension $n-1 \geq 3$, every indecomposable generalized permutahedron has at least $3n-6$ edge directions, and its direction graph has maximum degree at least $\min\{n-1, 6\}$. We prove these sharp bounds, attain both with one lattice construction on the third power of a path, and classify the direction graphs attaining the edge bound as the 3-trees in which each triangle belongs to at most three maximal cliques. A gluing theorem expresses indecomposability through positive relations among boundary responses, giving sharp separator capacities that can grow exponentially even at fixed clique number. For lattice summands without dilation, Boolean elemental differences show that the K_5-free graphs supporting full-dimensional indecomposable summands of dimension at least three are precisely the complete four-partite graphs with parts of size one or two. Allowing larger cliques, we construct graphical zonotopes as sums of two indecomposable lattice polytopes in every dimension at least three.

Contents

1. Introduction 3
1.1. Sharp direction bounds 3
1.2. Separators and exact capacities 4
1.3. Lattice summands and dilation 5
2. Deformation cones 6
3. An extremal path family 8
3.1. Elemental differences 8
3.2. Extremality 9
3.3. Dimension and the graph 10
4. Sharp bounds on the direction graph 11
4.1. A positive decomposition indexed by triangle classes 11
4.2. The minimum number of directions 14
4.3. The minimum maximum degree 14
5. Gluing across a boundary 15
5.1. Separator capacity 16
5.2. The space of boundary relations 17
5.3. Restriction gluing 19

6. A complete classification for 3-trees 22
6.1. Rooted branches and their boundary equations 22
6.2. Closing the boundary and proving extremality 27
6.3. All possible direction supports and the maximum dimension 30
6.4. Books, tree joins, and Hamiltonian graphs 32
7. Clique books 34
7.1. Clique separators and Boolean response modules 34
7.2. Constructive intervals for clique books 37
7.3. Sharp direction counts at saturated clique capacity 46
8. Four-partite constructions 51
8.1. An exact bound for a nonclique separator 51
8.2. An edge-minimal supporting family beyond 3-trees 55
8.3. Sharp exponential separator capacity in K_5-free graphs 58
8.4. Comparison with spanning 3-trees 65
9. Lattice summands 67
9.1. Sharp dilation for lattice summands on 3-trees 67
9.2. The sharp limit for lattice summands without dilation 69
9.3. Unit lattice summands from Boolean quadratic phases 73
9.4. The least dilation in the K_5-free case 80
Appendix A. A graphic model of the path family 81
Appendix B. Optimal lattice dilation on path joins 83
B.1. The elemental differences 84
B.2. One scale along the path 85
References 86

1. **Introduction**

Indecomposability imposes a sharp minimum on the directions of a generalized permutahedron. In dimension $n-1 \geq 3$, at least $3n-6$ directions are needed, and the equality cases are governed by a local condition on the triangles of a 3-tree. We prove these bounds and classify their equality cases, using a gluing principle that measures how many independently deformable pieces a separator can hold together.

The extremal examples also resolve a question about clique size. Poullot [1] proved that K_4-free graphs allow only segments and triangles as indecomposable summands, and conjectured that K_5-free graphs allow full-dimensional examples in every dimension at least three. The path construction below realizes these dimensions on planar graphs with the fewest possible directions. Its lattice structure then makes it possible to determine exactly when dilation is necessary.

For a graph G on $[n]$, put

$$Z_G = \sum_{ij \in E(G)} [e_i, e_j].$$

A polytope Q is a *weak Minkowski summand* of Z_G if $Q + R = \lambda Z_G$ for some polytope R and $\lambda > 0$. It is *indecomposable* if every Minkowski decomposition has summands homothetic to Q, up to translation. Modulo translations, the indecomposable weak summands are the extreme rays of the graphical deformation cone. We write $\operatorname{mdim}(Z_G)$ for their largest dimension, with value zero when G is edgeless. The *direction graph* of a generalized permutahedron has an edge ij for each actual edge direction $e_i - e_j$.

1.1. **Sharp direction bounds**

Every indecomposable generalized permutahedron of dimension $n-1$, $n \geq 4$, has a direction graph H satisfying

$$|E(H)| \geq 3n-6, \qquad \max_v \deg_H(v) \geq \min\{n-1, 6\}.$$

Both bounds are sharp, with no restriction on clique number. Equality in the edge bound forces H to be a 3-tree. The common reason is a positive decomposition indexed by interacting cliques, which forces the support of an indecomposable summand of dimension at least three into a triangle-connected family of tetrahedra.

The third power P_n^3 of a path attains the two bounds simultaneously. For $s_i = \mathbf{1}_{i \in S}$, define

$$z_n(S) = \sum_{i=1}^{n-2} s_i s_{i+1} s_{i+2} + \sum_{i=1}^{n-3} s_i s_{i+2} s_{i+3} (1 - s_{i+1}).$$

Its lower base polytope Q_n is a lattice indecomposable summand of $2Z_{P_n^3}$, has dimension $n-1$, and uses every direction of the planar, K_5-free graph P_n^3. Its elemental differences count components in edge-link trees, while their zero equations propagate a single scale along the path.

The equality cases are governed by one local condition. A 3-tree supports a full-dimensional indecomposable weak summand if and only if every triangle lies in at most three maximal K_4's (Theorem 6.1). The necessity comes from the number of independent branches that a triangle can support; the converse constructs all six nonconstant triangle responses on rooted branches. This gives a complete classification of actual direction supports in partial 3-trees and a linear-time algorithm for their maximum dimension. Among planar graphs on $n \geq 4$ vertices, a full-dimensional indecomposable weak summand exists exactly for stacked triangulations.

1.2. **Separators and exact capacities**

A separator limits how independently deformable pieces can combine into one extreme ray. Suppose the graph splits into branches meeting along an induced boundary graph H, with no edges between their private vertices. The deformation coefficients outside H then fall into independent blocks before the boundary equations are imposed. Put

$$d_H = \#\{C : C \text{ is a clique of } H,\ |C| \geq 2\}, \qquad m_H = \sum_{ij \in E(H)} 2^{|N_H(i) \cap N_H(j)|}.$$

An extreme ray has at most $m_H - d_H + 1$ nonzero private coefficient blocks across this boundary. Each boundary edge and conditioning pattern imposes at most one equation on the independent branch scales; the boundary coefficients account for the subtraction of d_H.

The construction that attains this bound records, for each branch, the conditional minima of its boundary slacks and the equalities between minimizing rows. When the private equations leave one scale per branch, the linear span of the minimal face of the glued deformation is isomorphic to the relation space among its boundary responses (Theorem 5.4). Extreme rays therefore arise from positive relations whose space is one-dimensional. This turns the separator bound into a construction principle.

For a clique boundary of size $s \geq 2$, the capacity is

$$b_s = \binom{s}{2} 2^{s-2} - (2^s - s - 1) + 1.$$

Boolean branch functions attain every count up to b_s, and hence, for $k \geq 1$,

$$\operatorname{mdim}(Z_{K_s * I_k}) = \min\{s + k - 1, s + b_s - 1\},$$

where I_k is an independent set. We also determine the least number of actual directions at saturation.

The same principle gives exponential capacities while keeping the clique number equal to four; the boundary $K_2 * I_p$ has capacity $2^p + p$. For $p, b \geq 1$, writing $a = \min\{p, b\}$, we obtain

$$\operatorname{mdim}(Z_{K_{1,1,p,b}}) = \min\{p + b + 1, 2^a + 2a + 1\}.$$

These dimensions can exceed those attainable on spanning admissible 3-trees. A related construction gives, for every $n \geq 10$, an edge-minimal supporting graph of treewidth four with $3n - 5$ edges. Thus the equality classification is the first case of a broader separator theory.

1.3. **Lattice summands and dilation**

Within the K_5-free class, dilation becomes necessary in large dimension. If G is K_5-free and a lattice indecomposable summand of the unscaled Z_G has dimension at least three, its actual directions form the edges of

$$K_{a_1,a_2,a_3,a_4}, \qquad a_i \in \{1, 2\},$$

and every such graph occurs, so there are at most eight active coordinates. Here integrality turns the elemental differences into Boolean functions; their quadratic interaction graphs then force the supporting tetrahedra to complete to a join of four parts.

Together with the path construction, this determines the least dilation; among K_5-free graphs on n vertices, the smallest factor admitting a full-dimensional lattice indecomposable summand is one for $4 \leq n \leq 8$ and two for $n \geq 9$.

For general graphical zonotopes, both members of a complementary pair of lattice summands can be indecomposable in every dimension at least three. Boolean quadratic phases produce such pairs on $K_{2,\dots,2}$ with $r \geq 4$ parts. Restricting one part to a singleton preserves both indecomposabilities and supplies the even dimensions at least six; the smaller four-partite cases complete the range (Corollary 9.11).

The path formula also has a graphic-matroid interpretation, in which a spanning-tree determinant explains its elemental differences (Appendix A). The alternative path joins $K_2 * P_{n-2}$ proposed by Poullot admit an optimal-dilation construction as well, developed in the final appendix; the path powers have the additional property of minimizing the maximum direction degree.

2. **Deformation cones**

For a function $z : 2^V \to \mathbb{R}$ with $z(\varnothing) = 0$, define its *elemental differences* by

$$\delta_{uv} z(S) = z(S \cup \{u, v\}) - z(S \cup \{u\}) - z(S \cup \{v\}) + z(S), \qquad u, v \notin S.$$

It is supermodular if and only if all these quantities are nonnegative. For such a function put

$$B(z) = \{x \in \mathbb{R}^V : x(V) = z(V),\ x(S) \geq z(S) \text{ for all } S \subseteq V\}, \qquad x(S) = \sum_{v \in S} x_v.$$

Lemma 2.1. *Let G be a graph on V. Modulo translations, its graphical deformation cone consists exactly of supermodular set functions with Möbius expansion*

$$z(S) = \sum_{C \subseteq S} y_C$$

supported on cliques of G of size at least two. Its complete inequalities are

$$L_{uv,S}(y) = \sum_{\{u,v\} \subseteq C \subseteq S \cup \{u,v\}} y_C \geq 0, \quad uv \in E(G), \quad S \subseteq N_G(u) \cap N_G(v). \tag{2.1}$$

The association $y \mapsto B(z)$ preserves addition and nonnegative scalar multiplication. Integral coefficients give an integral polytope. If all edge slacks are at most M, then

$$B(z) + B(Mz_G - z) = MZ_G, \qquad z_G(S) = |E(G[S])|.$$

Proof. This is the clique-coordinate description of the graphical deformation cone [2]. We recall its relation to support functions.

For an ordering $\pi = (v_1, \ldots, v_N)$, write $S_k = \{v_1, \ldots, v_k\}$ and define $q^\pi_{v_k} = z(S_k) - z(S_{k-1})$. Increasing marginal differences give $q^\pi(T) \geq z(T)$ for every T, with equality for the prefixes S_k, so $B(z)$ is nonempty; the coordinate bounds $z(\{v\}) \leq x_v \leq z(V) - z(V \setminus \{v\})$ make it bounded. If a linear functional increases in the order π, summation by parts shows q^π maximizes it. Consequently $B(z)$ is the convex hull of the greedy points q^π and is integral when z is integer valued. This support-function calculation also proves $B(z) + B(w) = B(z + w)$.

Swapping adjacent u, v after a prefix S changes the greedy point, up to sign, by $\delta_{uv} z(S)(e_u - e_v)$. Hence the support function is linear across all braid walls for a nonedge uv exactly when $\delta_{uv} z(S) = 0$ for every S. This is exactly the condition that its normal fan coarsen the graphical fan. A positive difference is an actual root edge direction; at a generic point of that braid wall, the exposed face is the nondegenerate segment between the two greedy points.

In the Möbius expansion, the difference is the sum in (2.1). Vanishing of this sum for all S for a fixed nonedge implies by Möbius inversion that every coefficient containing the nonedge is zero. Thus only clique coefficients remain, and conversely clique support makes every nonedge difference vanish. Removing singleton coefficients amounts to a translation. The remaining support also shows that an edge difference depends only on the common neighbors of its endpoints.

Any weak summand of Z_G has a normal fan coarsening the graphical fan, and therefore the braid fan; on a chamber where the support function of MZ_G is linear, the convex support functions of its two summands must both be linear. To see directly that such a summand Q has the required set function, put $z(S) = \min_{x \in Q} x(S)$. A vertex maximizing a functional increasing in the order π simultaneously minimizes every prefix sum, since $-\mathbf{1}_{S_k}$ is in the closure of that braid chamber. Its coordinates are consequently the greedy differences above. For the order beginning with S, u, v, its inequality at $S \cup \{v\}$ gives $z(S) + z(S \cup \{u, v\}) - z(S \cup \{u\}) \geq z(S \cup \{v\})$. Thus z is supermodular, and equality of the greedy support functions gives $Q = B(z)$. Adjacent braid chambers separated by a nonedge wall lie in one graphical chamber, so their greedy vertices coincide and the nonedge differences vanish. Hence this is the whole closed deformation cone.

$z_G(S) = |E(G[S])|$ has edge differences all equal to 1 and nonedge differences all 0. Thus $Mz_G - z$ is supermodular whenever the slacks of z lie in $[0, M]$, and the stated Minkowski identity follows. □

In particular, if $Q = Q' + Q''$ and $Q + R = MZ_G$, then $Q' + (Q'' + R) = MZ_G$ and $Q'' + (Q' + R) = MZ_G$. Every decomposition of Q therefore stays in this same positive cone. We write $\mathcal{C}(G)$ for this cone with all singleton coefficients normalized to zero.

Lemma 2.2 (Tight equations)**.** *The common kernel of the inequalities tight at $y \in \mathcal{C}(G)$ is the linear span of the minimal face containing y. A nonzero y spans an extreme ray if and only if this kernel is a line. It suffices to find any collection of tight inequalities whose common kernel is $\mathbb{R}y$.*

Proof. If $y = y' + y''$ with $y', y'' \in \mathcal{C}(G)$, every inequality vanishing at y vanishes at both summands, by nonnegativity. Conversely, a vector in the common kernel permits a sufficiently small perturbation of y in either direction; the tight inequalities remain zero, and the finitely many other inequalities remain positive. This proves the assertion about the minimal face. The cone is pointed, since vanishing of every elemental difference leaves only a modular function, whose normalized coefficients are zero. Its one-dimensional faces are therefore exactly its extreme rays. □

In the constructions that follow, nonnegative elemental differences produce a summand, and their zero equations must then leave only its scalar multiples. Positive differences record which directions actually occur, and an upper bound on the differences specifies a dilation of Z_G containing the summand.

We use y_C and $c_C(f)$ interchangeably for Möbius coefficients, and $\delta_e f$ for elemental differences. A root direction $e_i - e_j$ occurs in $B(f)$ precisely when $\delta_{ij} f$ is positive somewhere. The graph of these directions is denoted by $H(B(f))$. Its connected components determine the variation space, so a connected direction graph on n vertices gives dimension $n-1$.

3. **An extremal path family**

Let $n \geq 4$. On the vertex set $[n] = \{1, \dots, n\}$ take

$$G_n = P_n^3, \qquad E(G_n) = \{\{i,j\} : 1 \leq |i-j| \leq 3\}.$$

Its maximal cliques are the consecutive four-sets. Give each of them coefficient -1, give coefficient 1 to every consecutive triple and to every triple $\{i, i+2, i+3\}$, and set all other coefficients to zero. The consecutive triples are shared by neighboring four-sets; their common coefficients will force the local scales to agree. Thus, for $S \subseteq [n]$ and $s_i = \mathbf{1}_{i \in S}$, set

$$z_n(S) = \sum_{i=1}^{n-2} s_i s_{i+1} s_{i+2} + \sum_{i=1}^{n-3} s_i s_{i+2} s_{i+3} - \sum_{i=1}^{n-3} s_i s_{i+1} s_{i+2} s_{i+3}.$$

Define

$$Q_n = \left\{ x \in \mathbb{R}^n : \sum_{i=1}^{n} x_i = n-2, \quad \sum_{i \in S} x_i \geq z_n(S) \quad (S \subseteq [n]) \right\}. \tag{3.1}$$

Theorem 3.1. *The polytope Q_n is integral and Minkowski indecomposable, has dimension $n-1$, and is a summand of $2Z_{G_n}$ with complement $B(2z_{G_n} - z_n)$. Its actual direction graph is G_n, a planar K_5-free graph with $3n-6$ edges and maximum degree $\min(n-1, 6)$.*

3.1. **Elemental differences**

The $3n-8$ nonzero coefficients are supported on cliques of G_n, and the links of its edges explain why the elemental differences are nonnegative. For an edge e, let L_e be

the graph on its common neighbors, with uv an edge when $e \cup \{u, v\}$ is a four-clique. Successive four-cliques attach leaves to L_e, so it is a tree. The triangles with coefficient zero are precisely $\{i, i+1, i+3\}$; they are pairwise edge-disjoint, since for each edge length the endpoints determine i uniquely. Thus at most one vertex of L_e corresponds to such a triangle. For $U = S \cap V(L_e)$, let $\epsilon_e(U)$ indicate whether that vertex exists and belongs to U. Then

$$\delta_e z_n(S) = |U| - |E(L_e[U])| - \epsilon_e(U) \geq 0.$$

The first two terms count components of an induced forest, and the subtraction occurs only when that forest is nonempty.

For the quantitative bound, put $s_j = \mathbf{1}_{j \in S}$, with $s_j = 0$ when j lies outside $[n]$. The exact edge slacks are

$$\begin{aligned}
\delta_{i,i+1} z_n(S) &= s_{i-2}(1 - s_{i-1}) + s_{i-1}(1 - s_{i+2}) + s_{i+2}(1 - s_{i+3}), \\
\delta_{i,i+2} z_n(S) &= s_{i+1}(1 - s_{i-1}) + s_{i+3}(1 - s_{i+1}), \\
\delta_{i,i+3} z_n(S) &= s_{i+2}(1 - s_{i+1}).
\end{aligned} \tag{3.2}$$

The first expression counts transitions from one to zero along a four-bit word, hence is at most two; the two terms in the second cannot both be one, so the second and third expressions are at most one. Together with the vanishing nonedge differences, this proves supermodularity and gives

$$R_n = B(2z_{G_n} - z_n), \qquad Q_n + R_n = 2Z_{G_n}.$$

Since the coefficients are integers, both polytopes are integral. Also $z_n([n]) = n - 2$, as used in (3.1).

3.2. **Extremality**

The zero equations first remove all coefficients that were set to zero in the construction. On each consecutive four-set, the remaining equations identify its nonzero coefficients up to one scalar. The shared consecutive triples then identify the scalars of neighboring four-sets. We give the equations explicitly so that this propagation also covers the two ends of the path.

For a general point of the graphical cone, abbreviate all its clique coordinates as follows

$$\begin{aligned} &e_i = y_{i,i+1}, && f_i = y_{i,i+2}, && g_i = y_{i,i+3}, \\ &t_i = y_{i,i+1,i+2}, && u_i = y_{i,i+1,i+3}, && v_i = y_{i,i+2,i+3}, \\ &h_i = y_{i,i+1,i+2,i+3}. \end{aligned}$$

The indices run over the ranges for which all labels are in $[n]$. These are all the cliques, since every clique has largest and smallest labels at distance at most three. There are $7n - 17$ coordinates.

At the constructed vector, the following cone inequalities vanish

$$\begin{aligned} e_i = 0, \quad f_i = 0, \quad g_i = 0 \qquad & \text{for all their indices,} \\ g_i + u_i = 0 \qquad & (1 \le i \le n-3), \\ e_i + t_i + u_i + h_i = 0, & \\ f_{i+1} + u_i + t_{i+1} + h_i = 0, \qquad & (1 \le i \le n-3). \\ g_i + u_i + v_i + h_i = 0 & \end{aligned}$$

The first family uses the empty conditioning set on every edge. The second uses edge $i(i+3)$ with subset $\{i+1\}$. The last three equalities use, respectively, edges $i(i+1)$, $(i+1)(i+3)$, and $i(i+3)$, with conditioning sets $\{i+2, i+3\}$, $\{i, i+2\}$, and $\{i+1, i+2\}$.

The first two lines force all pair coefficients and all u_i to vanish. The remaining equations then reduce to

$$t_i = t_{i+1} = v_i = -h_i \qquad (1 \le i \le n-3).$$

The overlapping consecutive triples propagate one scalar λ; every t_i and v_i equals λ, and every h_i equals $-\lambda$. These account for every coordinate, so the common kernel of the selected equations is exactly one-dimensional. The edge 14 inequality conditioned on $\{3\}$ gives $v_1 = \lambda \ge 0$, identifying its intersection with the cone as the ray of Q_n. Lemma 2.2 therefore proves Minkowski indecomposability.

3.3. **Dimension and the graph**

For $1 \le i \le n-2$, choosing $S = \{i+2\}$ gives $\delta_{i,i+1} z_n(S) = 1$. For the last adjacent pair, use $S = \{n-2\}$. Therefore $e_i - e_{i+1}$ belongs to the difference space of Q_n. These $n-1$ vectors are independent, while (3.1) puts Q_n in one affine hyperplane. Hence $\dim Q_n = n - 1$.

Every distance-two direction has slack 1 at $S = \{i+1\}$. For a distance-three pair $i(i+3)$, take $S = \{i+2\}$ in (3.2), giving slack 1. Thus the actual direction graph is exactly G_n.

Five distinct integer labels have range at least four, so cannot form a clique in G_n. Its number of edges is $(n-1)+(n-2)+(n-3) = 3n-6$, and its maximum degree is $\min(n-1, 6)$. It is planar; start with the tetrahedral graph on $1, 2, 3, 4$ and insert each successive vertex j inside the triangular face $\{j-3, j-2, j-1\}$, joining it to that face's vertices. The triangle needed for the next insertion is one of the three new faces. This completes Theorem 3.1.

4. **Sharp bounds on the direction graph**

The path-cube construction uses the minimum possible number of edge directions and has the smallest possible maximum vertex degree among all indecomposable generalized permutahedra of the same dimension. Both bounds follow from the way its supporting four-cliques must be joined, without any restriction on the clique number of the ambient graph.

For a generalized permutahedron $Q \subseteq \mathbb{R}^n$, its *direction graph* $H(Q)$ has vertex set $[n]$ and an edge ij whenever an edge of Q is parallel to $e_i - e_j$. Throughout this section Q is full-dimensional in its affine hyperplane, so $\dim Q = n-1$. In particular, its direction graph uses all n coordinates.

Theorem 4.1. *Let $n \geq 4$, and let $Q \subseteq \mathbb{R}^n$ be a Minkowski-indecomposable generalized permutahedron of dimension $n-1$. Its direction graph H satisfies*

$$|E(H)| \geq 3n-6, \qquad \max_{v \in [n]} \deg_H(v) \geq \min\{n-1, 6\}.$$

If $|E(H)| = 3n-6$, then H is a 3-tree. The family Q_n on P_n^3 attains both bounds for every $n \geq 4$.

The 3-trees attaining equality are classified in Theorem 6.1.

4.1. **A positive decomposition indexed by triangle classes**

By Lemma 2.1, the normalized coefficients of $Q = B(z)$ are supported on cliques of its actual direction graph. Partition its triangles by declaring all triangles in each clique C with $y_C \neq 0$ and $|C| \geq 4$ equivalent. For every nonzero higher coefficient, let $\alpha(C)$ denote its class. One may also use the coarser partition generated by all four-cliques of H. The same argument applies to either partition; the refined one will be useful for lattice summands.

Simply separating these coefficient classes need not preserve supermodularity; a negative higher coefficient may need positive edge coefficients to compensate for it. The variables belonging to different classes are disjoint in each elemental difference. Their minima therefore tell us exactly how much of each edge coefficient to assign to each class, leaving only nonnegative segment summands.

Lemma 4.2. *Every generalized permutahedron whose direction graph is contained in H has, modulo translation, a Minkowski decomposition*

$$Q = \sum_{\alpha} Q_\alpha + \sum_{e \in E(H)} r_e \Delta_e, \qquad r_e \geq 0,$$

where the higher Möbius coefficients of Q_α are exactly those of Q belonging to the class α. Every Q_α is a generalized permutahedron with nonnegative elementary differences. If α is a singleton triangle class, Q_α has dimension at most two.

Proof. Fix an edge $e = ij$. Its common-neighbor variables x_k are partitioned into the disjoint sets

$$I_{e,\alpha} = \{k : e \cup \{k\} \text{ is a triangle in class } \alpha\}.$$

Every nonconstant monomial in the elemental difference at e uses variables from one $I_{e,\alpha}$. Indeed, a coefficient $y_C \neq 0$ with $C \supseteq e$ puts every triangle $e \cup \{k\}$, $k \in C \setminus e$, in the same class. Thus

$$\delta_e z(x) = y_e + \sum_{\alpha} F_{e,\alpha}(x_\alpha), \qquad F_{e,\alpha}(x_\alpha) = \sum_{\substack{C \supseteq e,\ |C| \geq 3,\ y_C \neq 0 \\ \alpha(C) = \alpha}} y_C \prod_{k \in C \setminus e} x_k.$$

The summands have disjoint variable sets and satisfy $F_{e,\alpha}(0) = 0$. Set

$$a_{e,\alpha} = -\min_{x_\alpha \in \{0,1\}^{I_{e,\alpha}}} F_{e,\alpha}(x_\alpha) \geq 0.$$

Since the variables are disjoint, their minima add. The complete positivity condition at e is therefore equivalent to

$$y_e \geq \sum_{\alpha} a_{e,\alpha}.$$

For each class α, retain the original coefficient y_C on every clique C of size at least three with $\alpha(C) = \alpha$, set the other higher coefficients to zero, and give edge e the coefficient $a_{e,\alpha}$. The resulting set function z_α has elementary differences

$$a_{e,\alpha} + F_{e,\alpha}(x_\alpha) \geq 0.$$

It is consequently supermodular and defines a generalized permutahedron Q_α. The residual coefficient

$$r_e = y_e - \sum_\alpha a_{e,\alpha}$$

is nonnegative. The decomposition follows coefficientwise from Lemma 2.1.

For an isolated triangle with coefficient t, the three edge coefficients in its piece are all $\max\{0, -t\}$. This piece is, modulo translation, $t\Delta_T$ when $t \geq 0$ and $(-t)(-\Delta_T)$ when $t \leq 0$. It has dimension at most two. □

Lemma 4.3. *If Q is Minkowski-indecomposable and $\dim Q = n - 1 \geq 3$, then the graph $H(Q)$ has a family of K_4 subgraphs, connected through shared triangles, whose union contains all n vertices.*

Proof. Apply Lemma 4.2 to the actual direction graph. Each nonzero summand must be a positive homothetic copy of Q. Since segments and isolated-triangle pieces have dimension at most two, a nonzero piece comes from a class containing a K_4.

Using the coarser partition, such a class corresponds to a connected component of the graph of four-cliques joined across common triangles. Every higher coefficient of the chosen piece belongs to a clique in this component. Its allocated edge coefficient vanishes on any edge contained in no such clique, because the corresponding $F_{e,\alpha}$ is identically zero. Thus the piece can vary only in coordinates belonging to the union of the component's K_4s. Since it is a positive homothetic copy of the full-dimensional Q, that union contains all n vertices. □

Corollary 4.4. *Let Q be indecomposable of dimension at least three, and suppose its actual direction graph H is K_5-free. The family*

$$\mathcal{C} = \{C : |C| = 4,\ y_C \neq 0\}$$

is nonempty, is connected through shared triangles, and satisfies $E(H) = \bigcup_{C \in \mathcal{C}} \binom{C}{2}$.

Proof. Use the refined partition in Lemma 4.2. A nonzero piece of dimension at least three must contain a nonzero four-clique coefficient and be homothetic to Q. Comparing higher coefficients shows that all nonzero four-clique coefficients belong to that piece. Its triangles and allocated edge coefficients are supported on their union. Conversely, each edge in a four-clique with nonzero coefficient has a nonzero elemental polynomial and hence a positive value somewhere. □

4.2. **The minimum number of directions**

Choose a connected family as in Lemma 4.3, and order its K_4s so that every member after the first shares a triangle with an earlier member. The first contributes six edges on four vertices. Each later member introduces at most one new vertex; whenever it introduces one, its three edges to the old shared triangle are new. A member introducing no vertex can only add further edges. Hence the union, and therefore $H(Q)$, has at least

$$6 + 3(n - 4) = 3n - 6$$

edges.

Suppose equality holds. Since the component's union covers all n vertices and has at least $3n - 6$ edges, it equals $H(Q)$ and every step of the count is an equality. In particular, a K_4 with no new vertex adds no edge. Such steps may safely be ignored; their edges were already present, so omitting them loses no triangle needed for a later attachment. Every remaining step adds one vertex and exactly its three edges to an old triangle. Starting from the initial K_4, these are precisely the defining operations of a 3-tree. This proves both the direction-count bound and its necessary equality condition.

4.3. **The minimum maximum degree**

Lemma 4.5. *In a graph of maximum degree at most d, where $d \in \{3, 4, 5\}$, every triangle-connected family of K_4 subgraphs has at most $d + 1$ vertices.*

Proof. For $d = 3$, all vertices of the initial K_4 are already saturated, so no new vertex can attach along a triangle.

For $d = 4$, upon introducing a fifth vertex the graph induced on the five vertices is K_5 minus a possible edge. The three common vertices have degree four. Only two vertices remain available to receive another neighbor, so there is no triangle along which a sixth vertex can attach. If the missing edge is present, all five vertices are saturated.

Consider $d = 5$, and suppose a component has reached five vertices. The first two K_4s which introduce them have union K_5 minus one possible edge. Thus the induced graph on these five vertices is either K_5 or $K_5 - de$. Intervening K_4s using only these vertices introduce no third possibility. A sixth vertex is adjacent to a triangle on the first five. We show that afterward no triangle on these six vertices can have all its vertices of degree at most four. This excludes a seventh vertex, because its attachment would require precisely such a triangle. All degrees used below are

already lower bounds from the displayed six-vertex subgraph; neighbors outside it can only strengthen the obstruction.

If the first five vertices induce K_5, the three attachment vertices already have degree five. The only remaining possible available triangle uses the two other old vertices and the new vertex. Both old vertices had degree four in K_5, so making this triangle gives each of them degree five, and gives the new vertex degree five as well. Thus this triangle cannot receive another neighbor.

Otherwise write the first five vertices as a central triangle abc and two tips d, e, with de missing. Up to symmetry, the attachment triangle of the sixth vertex f is either abc or abd.

If it is abc, those three vertices have degree five. The only possible available triangle is def. Each of its vertices already has three neighbors in abc, so if def is a triangle all three have degree five.

If the attachment triangle is abd, then a, b have degree five. Among the remaining vertices, the required edges form the path $e - c - d - f$. Its internal vertices c, d already have degree four, and its ends e, f have degree three. Every missing edge except ef touches an internal vertex, thereby giving that vertex degree five. Adding only ef makes a four-cycle, which has no triangle. Consequently no additional edges can create a triangle all of whose vertices remain of degree at most four.

These cases include every added edge on the six vertices. Steps through K_4s with no new vertex therefore cannot allow a seventh vertex to enter. □

By Lemma 4.3, the component must cover all n active vertices. The three cases of Lemma 4.5 therefore give the lower bounds three, four, five, and six for $n = 4$, $n = 5$, $n = 6$, and $n \geq 7$, respectively. Finally, P_n^3 has exactly $3n - 6$ edges and maximum degree $\min\{n - 1, 6\}$. The full-dimensional indecomposable family Q_n attains both bounds, completing the proof of Theorem 4.1.

5. **Gluing across a boundary**

The direction bounds show that four-cliques must be linked through triangles, but do not yet say which such assemblies support an indecomposable summand. A separator imposes a further restriction; each branch can be rescaled independently until a tight equation on the boundary relates it to the others. We first count the available boundary equations, then describe exactly what information a branch contributes to them. This will turn the classification and the later constructions into finite-dimensional relation problems.

5.1. **Separator capacity**

Lemma 5.1 (An arbitrary induced separator)**.** *Let $B \subseteq V(G)$, let $H = G[B]$, and put*

$$d_H = \#\{C \subseteq B : C \text{ is a clique of } H,\ |C| \geq 2\}, \qquad m_H = \sum_{ij \in E(H)} 2^{|N_H(i) \cap N_H(j)|}.$$

For an extreme ray of the graphical deformation cone, let r be the number of components of $G - B$ on which some private Möbius coefficient is nonzero. Here a coefficient is private to a component if its clique meets that component. Then

$$r \leq m_H - d_H + 1.$$

Proof. Normalize singleton coefficients to zero. Since every clique not contained in B meets only one component of $G - B$, we can scale the r nonzero private coefficient blocks independently while allowing all d_H boundary coefficients to vary. This gives a space of variations of dimension $r + d_H$. A zero slack on an edge not contained in B remains zero under these variations, because all its coefficients belong to one private block.

Fix a boundary edge $e = ij$ and the membership pattern S of its common neighbors inside H. Conditional on these bits, the full slack is a boundary constant plus a sum of functions on disjoint sets of component-private variables. At a zero slack each private function must attain its conditional minimum, since minimizing that function alone would otherwise make the slack negative. Thus all zero rows of this fixed boundary type give the same equation on the proposed variations, namely the varied boundary constant plus the sum of the scaled conditional minima. There is at most one such equation for each of the m_H boundary types. The common zero space therefore has dimension at least $r + d_H - m_H$.

An extreme ray has a one-dimensional common zero space after translations are removed. Hence $r + d_H - m_H \leq 1$, as required. □

Corollary 5.2 (Every separating set contains a triangle)**.** *If a graph G on $n \geq 4$ vertices supports an indecomposable weak summand of dimension $n - 1$, every vertex separator of G contains a triangle. In particular, G is 3-vertex-connected.*

Proof. Full dimension implies that every component outside the proposed separator has a nonzero private coefficient, for otherwise all coordinates in that component are fixed, after the singleton translations have been removed. Thus r in Lemma 5.1 is the number of components. If the induced boundary graph H is triangle-free, it has only edge cliques and every edge has no common neighbor. Hence $m_H = d_H = |E(H)|$,

and the lemma gives $r \leq 1$. This excludes every triangle-free separating set. Sets of at most two vertices are triangle-free, proving the connectivity assertion. The same conclusion applies to the actual direction graph of any indecomposable generalized permutahedron of dimension $n-1$. □

Corollary 5.3 (Triangle separators)**.** *Across a triangle separator, an extreme ray has at most three nonzero private coefficient blocks. In particular, a full-dimensional ray requires at most three components after deleting the triangle.*

Proof. For $H = K_3$, the boundary dimensions are $d_H = 4$ and $m_H = 6$. Apply Lemma 5.1. □

5.2. **The space of boundary relations**

The separator bound counts equations without specifying the branches. To attain it, we construct branches whose private equations leave one scalar each, then use the boundary to determine their ratios. A branch contributes its minimum slack for each fixed boundary context, since these are the values at which a global boundary inequality can become tight. Differences between minimizing rows carry additional private equations; when these and the internal tight rows leave one scalar, extremality becomes uniqueness of a positive relation among the vectors of conditional minima.

Let a graph be the union of induced branches $G[B \cup V_i]$, where $V_1, \ldots, V_r$ are distinct components of $G - B$, and put $H = G[B]$. There are no edges between different V_i. Normalize singleton coefficients to zero. Let d be the number of cliques of H of size at least two, and index the m boundary types by

$$(e, S), \qquad e = ab \in E(H), \quad S \subseteq N_H(a) \cap N_H(b).$$

The boundary slack map

$$A : \mathbb{R}^d \longrightarrow \mathbb{R}^m$$

is injective. Indeed, the empty-context differences first determine the pair coefficients. For a larger clique C, a difference on any pair $e \subset C$, conditioned on $C \setminus e$, determines y_C once the coefficients of its proper subcliques are known. Induction therefore recovers all normalized boundary coefficients from A.

For a branch function f_i, require all its slacks on edges not contained in B to be nonnegative. For each boundary type, let

$$m_i(e, S) = \min_{X \subseteq V_i} \delta_e f_i(S \cup X).$$

Its *private equations* consist of its tight internal slacks and the differences between any two boundary rows attaining the same conditional minimum. They involve only

coefficients of cliques not contained in B. The ambient private space includes every such graphical clique, including those whose chosen coefficient is zero. The minimizing rows are fixed at the chosen function f_i; these are ordinary linear equations, with no assumption that a minimizer remains optimal under a signed variation.

Adding an arbitrary boundary function h_i to f_i changes neither its private coefficients nor its private equations, but replaces m_i by $m_i + Ah_i$. Thus the response class $[m_i]$ in $\mathbb{R}^m/\operatorname{im} A$ is independent of the chosen boundary lift; the corresponding change in the closing function is $z \mapsto z - \sum_i \lambda_i h_i$.

Theorem 5.4 (Positive response gluing)**.** *Suppose each branch has a nonzero private coefficient vector and its private kernel is the line spanned by that vector. Let $\lambda_i > 0$ and suppose a boundary function z satisfies*

$$Az + \sum_{i=1}^{r} \lambda_i m_i = 0.$$

Then

$$f = z + \sum_{i=1}^{r} \lambda_i f_i$$

belongs to the graphical deformation cone. Modulo translations, the dimension of its complete common zero space is exactly

$$\dim \ker \left(\mathbb{R}^r \longrightarrow \mathbb{R}^m/\operatorname{im} A, \quad c \longmapsto \sum_i c_i [m_i] \right).$$

In particular, f spans an extreme ray if and only if this relation space is one-dimensional, in which case it is spanned by the strictly positive vector λ. The number of branches is then at most $m - d + 1$.

Proof. Internal-edge inequalities belong to their own branch and retain its nonnegativity under positive scaling. Conditional on a boundary type, the private-variable sets are disjoint. Hence

$$\delta_e f = \sum_i \lambda_i \big(\delta_e f_i - m_i(e, S) \big) \geq 0.$$

This proves all inequalities. Because every λ_i is positive, a boundary row is tight exactly when every branch independently attains its conditional minimum. Thus the boundary minimizer set is the Cartesian product of the individual minimizer sets.

Let g lie in the complete common zero space of f. Every internal zero equation of a branch is a global zero equation. To obtain a branch boundary-minimizer difference,

fix minimizing choices in all the other branches and subtract the two resulting global zero rows. Thus all private equations of every branch survive. By hypothesis the private part of g in branch i is c_i times the private part of f_i. Subtracting $\sum_i c_i f_i$ leaves a function z' supported on B. The remaining global zero equations are exactly

$$Az' + \sum_i c_i m_i = 0.$$

Conversely, every solution of this equation produces a vector in the complete common zero space; all internal tight rows vanish, and all boundary tight rows evaluate to the displayed aggregate expression. The nonzero, disjoint private blocks determine the scalars c_i uniquely, and injectivity of A makes z' unique when it exists. The common zero space is therefore isomorphic to the stated quotient-relation space, proving the dimension formula. Lemma 2.2 identifies this dimension with that of the minimal face and gives the extremality assertion. Finally, the quotient has dimension $m-d$, so a one-dimensional relation space among r vectors requires $r \le m-d+1$. □

The relevant boundary dimensions in the present constructions are

H	d	m	$m-d+1$
K_3	4	6	3
K_4	11	24	14
$K_4 - e$	7	12	6
$K_2 * I_p$	$1+3p$	2^p+4p	2^p+p

Each displayed capacity is attained by a positive construction below. The 3-tree induction supplies all needed triangle-response states, the four-clique books use Boolean modules, and the complete four-partite classification uses cube-cell indicators. In each case the retained equations come from actual conditional minima of nonnegative elemental slacks.

5.3. **Restriction gluing**

When the restrictions themselves are extreme, a nonzero overlap identifies their scales directly, without computing boundary responses. The following criterion makes this observation precise; the four-element example afterwards also supplies one of the complementary unit lattice summands used in Section 9.3.

Proposition 5.5 (Restriction gluing)**.** *Let $f \in \mathcal{C}(G)$, and suppose $U_1, \dots, U_m \subseteq V(G)$ cover every clique of size at least two. Assume that each restriction $f|_{2^{U_i}}$ is nonzero and spans an extreme ray of $\mathcal{C}(G[U_i])$. If the graph joining i and j whenever $f|_{2^{U_i \cap U_j}}$ is nonzero is connected, then f spans an extreme ray of $\mathcal{C}(G)$.*

Proof. Suppose $f = g + h$ with $g, h \in \mathcal{C}(G)$. Restriction preserves supermodularity and normalization, so local extremality gives

$$g|_{2^{U_i}} = \lambda_i f|_{2^{U_i}}, \qquad \lambda_i \geq 0.$$

If f is nonzero on $2^{U_i \cap U_j}$, evaluating the two identities at a subset on which it is nonzero gives $\lambda_i = \lambda_j$. Connectivity therefore gives one common scalar λ. Möbius inversion on each U_i shows that $c_A(g) = \lambda c_A(f)$ for every clique A covered by U_i. The covering hypothesis accounts for every possible nonzero coefficient; nonclique coefficients and coefficients of size at most one already vanish. Thus $g = \lambda f$. Applying the same argument to h proves the assertion. □

Without fixing normalization, the overlap requirement says that the common restriction is nonmodular, so it still identifies the local scaling parameters. This criterion addresses extremality once global supermodularity is known; agreement of extreme local restrictions alone does not imply the latter.

5.3.1. **A four-element block.** On $\{1, 2, 3, 4\}$ define p by the following nonzero Möbius coefficients

$$\begin{aligned} c_{12}(p) = c_{13}(p) = c_{23}(p) = c_{24}(p) = c_{34}(p) &= 1, \\ c_{123}(p) = c_{234}(p) &= -1, \\ c_{1234}(p) &= 1. \end{aligned}$$

In particular $c_{14} = c_{124} = c_{134} = 0$. If x_i denotes membership of i in the conditioning set, the six pair differences are

$\{i, j\}$	$\delta_{ij} p$
14	$x_2 x_3$
13	$1 - x_2 + x_2 x_4$
24	$1 - x_3 + x_1 x_3$
12	$1 - x_3 + x_3 x_4$
34	$1 - x_2 + x_1 x_2$
23	$(1 - x_1)(1 - x_4)$.

Every entry is nonnegative for zero-one arguments, so p is supermodular. In fact every entry is either zero or one.

To prove extremality, let g be any normalized supermodular function that satisfies all pair-difference equalities satisfied by p. Write its pair, triple, and quadruple coefficients as a_{ij}, b_{ijk}, and q. The zero differences for pair 14 at $\varnothing$, $\{2\}$, and $\{3\}$ give

$$a_{14} = b_{124} = b_{134} = 0.$$

The zero differences

$$\delta_{12}g(\{3\}) = \delta_{13}g(\{2\}) = \delta_{23}g(\{1\}) = 0$$

give $a_{12} = a_{13} = a_{23} = -b_{123}$, while

$$\delta_{24}g(\{3\}) = \delta_{34}g(\{2\}) = \delta_{23}g(\{4\}) = 0$$

give $a_{24} = a_{34} = a_{23} = -b_{234}$. Finally,

$$\delta_{23}g(\{1,4\}) = a_{23} + b_{123} + b_{234} + q = 0$$

gives $q = a_{23}$. Consequently every coefficient of g is determined by the single parameter $t = a_{23}$, and $g = tp$. Supermodularity gives $t = \delta_{23}g(\varnothing) \geq 0$. Thus the common zero set cuts out exactly the ray $\mathbb{R}_{\geq 0}p$ in the nonnegative supermodular cone.

The restrictions to $\{1,2,3\}$ and $\{2,3,4\}$ have all three pair coefficients equal to 1 and their triple coefficient equal to -1. Their base polytopes are

$$(1,1,1) - \Delta_{\{1,2,3\}}, \qquad (1,1,1) - \Delta_{\{2,3,4\}},$$

respectively, in the corresponding coordinate subspaces. In particular they are nonmodular, so either triangle can synchronize the scale with the next four-element block.

Geometric identification. The base polytope $B(p)$ is the thirteen-vertex polytope called the *persimmon* in [3, Section 3.3]. Indeed, let $H = K_{\{1,4\},\{2,3\}}$, so

$$Z_H = [e_1, e_2] + [e_1, e_3] + [e_4, e_2] + [e_4, e_3].$$

The set functions p and z_H agree on every subset except $\{2,3\}$, where $p(\{2,3\}) = 1$ and $z_H(\{2,3\}) = 0$. Indeed their difference is the Boolean function

$$\mathbf{1}_{\{2,3\}\subseteq S}(1 - \mathbf{1}_{1\in S})(1 - \mathbf{1}_{4\in S}).$$

Therefore

$$B(p) = Z_H \cap \{x_2 + x_3 \geq 1\}.$$

The four-cycle zonotope Z_H has fourteen vertices; its acyclic orientations are the 2^4 orientations of the cycle except the two directed cycles. The cut removes the unique vertex $(2, 0, 0, 2)$ at which $x_2 + x_3 = 0$. Every edge from that vertex is a translate of one of the four unit generating segments and meets the cutting hyperplane at its other endpoint. Hence the cut creates no new vertices and leaves thirteen vertices.

6. **A complete classification for 3-trees**

A 3-tree is obtained from K_4 by repeatedly adding a vertex whose earlier neighbors form a triangle, so it is K_5-free and has $3n - 6$ edges. The triangle-separator bound gives a necessary condition for such a graph to support a full-dimensional indecomposable summand. It is also sufficient.

Theorem 6.1. *Let G be a 3-tree on $n \geq 4$ vertices. The following are equivalent*

(i) *Z_G has a Minkowski-indecomposable weak summand of dimension $n - 1$;*

(ii) *every triangle of G belongs to at most three maximal K_4's;*

(iii) *deleting any triangle of G leaves at most three connected components.*

When these conditions hold, there is an integral coefficient vector on such an extreme ray with every maximal-clique coefficient equal to 1. Its polytope uses every direction of G and is an actual summand of $(n - 3)Z_G$. Every triangle coefficient is 0, -1, or -2.

The proof uses the entire conditional boundary response. A branch which is decomposable when its boundary coordinates are considered in isolation can still be rigid after its boundary minimizer equations are retained.

6.1. **Rooted branches and their boundary equations**

The bipartite incidence graph whose nodes are the maximal K_4's and the triangles of a 3-tree, with containment as incidence, is a tree. Indeed, the initial K_4 gives a four-leaf star. Each added vertex attaches a new clique node to one existing triangle node and creates three new triangle leaves. A clique node always has degree four; the degree of a triangle node is the number of maximal K_4's containing it.

Deleting a triangle node T from this incidence tree gives one branch for each incident maximal clique. The corresponding graph branches meet exactly in T, and their vertices outside T form the connected components of $G - T$. This also follows directly by induction; every new vertex in one branch attaches to a triangle containing an earlier vertex outside T, and there are no edges between different branches. This

proves the equivalence of (ii) and (iii). Corollary 5.3 proves the necessity of these conditions.

We prove sufficiency by constructing a branch with a prescribed response. Use the private coefficients and equations of Section 5.2. Normalize the coefficient of the boundary triangle T to zero, leaving its pair coefficients free. For a boundary edge $e \subset T$ and its third vertex k, write

$$m_e^\epsilon(f) = \min_{x \in \{0,1\}^{V \setminus T}} \delta_e f(x)\big|_{x_k = \epsilon}, \qquad \epsilon \in \{0,1\}.$$

The induction retains the minimizer differences for both values of ϵ; this is what allows the branch scales to be recovered after gluing.

There are six boundary types, two for each edge of T. A boundary pair coefficient shifts both values on its edge equally, so after removing pair coefficients the response is the vector of three differences $m_e^1 - m_e^0$. A boundary triangle coefficient shifts all three differences equally, giving the response quotient $\mathbb{R}^3/\langle \mathbf{1} \rangle$. Complementary nonconstant bit vectors for two branches, or the three unit vectors for three branches, have the unique positive relation required by the gluing theorem. The induction must produce all six nonconstant bit vectors, since the state needed by a child depends on the other children at its parent.

Lemma 6.2 (A module with any prescribed nonconstant response)**.** *Suppose a rooted branch has at most two child branches at every triangle. For every nonconstant vector*

$$v = (v_e : e \in \tbinom{T}{2}) \in \{0,1\}^3 \setminus \{000, 111\},$$

there is an integral branch function f with the following properties

(i) *every maximal-clique coefficient is* 1 *and the boundary triangle coefficient is zero;*

(ii) *every internal elemental slack is nonnegative;*

(iii) $m_e^{\epsilon}(f) = v_e(\epsilon - 1)$ *for all six boundary choices;*

(iv) *the common kernel of the private equations is precisely the line spanned by the private coefficient vector of f;*

(v) *each internal edge has a positive slack, and each boundary edge has a value strictly above its conditional minimum for some conditioning set.*

Every nonboundary triangle coefficient is $-d_F$, where $0 \le d_F \le \deg(F)$ and $d_F \in \{0,1,2\}$; here $\deg(F)$ is its maximal-clique multiplicity in the branch.

Proof. We give one induction which also covers a branch consisting of a single maximal clique. Its first clique is $Q = T \cup \{p\}$. Label the three vertices of T by $1, 2, 3$, put

$$F_i = Q \setminus \{i\}, \qquad E_i = T \setminus \{i\}, \qquad v_i = v_{E_i},$$

and let $k_i \in \{0,1,2\}$ be the number of child branches on F_i. Call F_i occupied when $k_i > 0$.

The bits chosen next have two purposes. Each occupied face must have a nonconstant row, so its children can be forced to the core's scale. On edges shared by occupied faces, the rows must agree, so that the reduced slack has enough zeros to expose both boundary states of every child. These requirements explain the following choices.

First choose a bit $b_i(e)$ for each of the three edges of F_i. Its coordinate on E_i is always v_i. If F_i is unoccupied, take its entire row b_i to be the constant vector $v_i\mathbf{1}$. At an intersection of an occupied and an unoccupied face, copy the bit of the unoccupied face to the occupied row. If exactly two faces are occupied and F_k is unoccupied, assign their common edge the bit $1 - v_k$ in both rows. If all three faces are occupied, assign the intersection $F_i \cap F_j$ the bit v_k, where $\{i,j,k\} = \{1,2,3\}$. Every occupied row is nonconstant. With one occupied face it is a permutation of v; with two, it contains both v_k and $1 - v_k$; with three, each row is again a permutation of v. Moreover,

$$b_i(F_i \cap F_j) = b_j(F_i \cap F_j) \quad \text{whenever at least one of } F_i, F_j \text{ is occupied.} \tag{6.1}$$

We now complete each occupied row to a constant vector by adding one or two child responses. Prescribe a nonconstant response r to each child branch as follows. For an unoccupied face set $d_i = v_i$. For one child set

$$r = \mathbf{1} - b_i, \qquad d_i = 1.$$

For two children, a nonconstant three-bit row is either a unit vector or the complement of a unit vector. If b_i is one unit vector, give the children the other two unit vectors and set $d_i = 1$. If b_i is one complementary unit vector, give them the other two complementary unit vectors and set $d_i = 2$. Write R_i for the sum of the prescribed child responses, taking $R_i = 0$ when there is no child. In every case,

$$d_i\mathbf{1} - R_i = b_i. \tag{6.2}$$

Apply the induction hypothesis separately to the child branches.

Sum their complete normalized functions and add the following core terms; coefficient 1 on Q, coefficient $-d_i$ on F_i, and pair coefficients

$$\begin{aligned} c_{E_i} &= R_i(E_i), \\ c_e &= R_i(e) + R_j(e) + \max\{b_i(e), b_j(e)\}, \quad e = F_i \cap F_j. \end{aligned} \tag{6.3}$$

The coefficient on T remains zero. These are integer coefficients; children contribute no coefficient on their boundary triangle, so each triangle coefficient is assigned exactly once.

For positivity, minimize a core-edge slack over all child-private variables. Contributions from different branches use disjoint variables, so their minima add. On $e = F_i \cap F_j$, put $a = b_i(e)$ and $b = b_j(e)$. After inserting (6.2)–(6.3), the result is

$$\max(a, b) - ax_j - bx_i + x_ix_j. \tag{6.4}$$

For $(a, b) = (0, 0)$ this is x_ix_j, for $(1, 1)$ it is $(1 - x_i)(1 - x_j)$, and for unequal bits it is $1 - x + xz$ after relabeling. Every case is nonnegative, has a zero, and has a positive value. Internal edges lying strictly inside a child branch retain their child's slacks. This proves internal nonnegativity and the internal part of (v).

On a boundary edge E_i, after minimizing child variables the slack is

$$-v_ix_p + x_ix_p. \tag{6.5}$$

Its minimum over x_p is $v_i(x_i - 1)$, as required. It also has a value strictly above that minimum; use $x_i = 1, x_p = 1$ when $v_i = 0$, and $x_i = 0, x_p = 0$ when $v_i = 1$. This proves (iii) and the remaining part of (v).

We next prove the one-dimensional kernel assertion. The first point is that all six conditional minimizer equations of every child are retained. Consider an occupied face F_i. On an edge $F_i \cap F_j$, equality (6.1) makes (6.4) either $x_i x_j$ or $(1-x_i)(1-x_j)$. Its zero set contains both values of the third vertex of F_i. Fixing either such value, choose a reduced zero and minimize all child contributions independently. Any two minimizers of the chosen child then give two zero rows of the parent; their difference is precisely the child's private equation.

For the remaining edge E_i of F_i, set $x_i = v_i$ in (6.5). Both values of x_p attain its conditional minimum. Thus both child boundary types are again exposed, now as differences of parent boundary minimizers. Child internal zero rows are retained directly. Consequently, any vector satisfying the parent's private equations restricts on each child-private block to a scalar λ times that child's vector.

Subtract these scaled normalized child functions from the variation. The remaining coefficients are supported on the core Q. Denote its quadruple coefficient by h, its negative triangle coefficient on F_i by d_i', and the child scales on this face by $\lambda_{i,j}$. Define the effective row

$$A_i = d_i' \mathbf{1} - \sum_j \lambda_{i,j} r_{i,j}.$$

On a core edge $e = F_i \cap F_j$ adjacent to an occupied face, the effective slack of the variation has the form

$$c - A_i(e) x_j - A_j(e) x_i + h x_i x_j.$$

If the chosen bits are zero, the values at $(0,0), (1,0), (0,1)$ vanish. If they are one, the zero rows are $(1,0), (0,1), (1,1)$. Substitution gives, respectively,

$$c = A_i(e) = A_j(e) = 0, \qquad c = A_i(e) = A_j(e) = h.$$

Thus both effective coefficients are h times their chosen bits. On E_i, set $x_i = v_i$ as above and subtract the rows for $x_p = 0$ and $x_p = 1$. Their difference is $-A_i(E_i) + h v_i$, so the root coordinate satisfies the same rule. Therefore, for every occupied face,

$$d_i' \mathbf{1} - \sum_j \lambda_{i,j} r_{i,j} = h b_i. \tag{6.6}$$

For an unoccupied face the boundary equation alone gives $d_i' = h v_i$.

If there is one child, substitute $r = \mathbf{1} - b_i$ into (6.6). Because b_i has both a zero and a one, it gives $d_i' = h$ and $\lambda = h$. With two unit responses, the coordinate of the third unit gives $d_i' = h$, and the other coordinates give both child scales equal to h. With two complementary unit responses, the three coordinates give $d_i' = 2h$

and again both scales equal to h. Thus every child scale is h, every core triangle coefficient is h times its chosen coefficient, and each internal core pair coefficient is fixed by one of its zero rows. Only the boundary pairs remain free, and these are not private coordinates. Hence the entire private variation equals h times the constructed vector. Conversely that vector satisfies all its private equations, proving (iv).

When there are no children, the same argument uses the boundary-minimizer ties to fix the three triangle coefficients and the internal zero rows to fix the three internal pairs; thus it supplies the base case as well. Finally, a nonboundary face with k_i children has multiplicity $k_i + 1$. The displayed choices of d_i give the stated coefficient bounds. □

6.2. **Closing the boundary and proving extremality**

Sufficiency in Theorem 6.1. For $G = K_4$, use the simplex coefficient $y_{1234} = 1$, with all other coefficients zero. Suppose G has at least two maximal cliques. Choose a shared triangle T; its multiplicity is two or three. Removing its node from the incidence tree gives two or three rooted branches of the type in Lemma 6.2.

For two branches, prescribe a nonconstant response v to one and $\mathbf{1} - v$ to the other. For three branches, prescribe the three unit vectors. Sum the resulting normalized branch functions, add coefficient 1 on each of the three edges of T, and add coefficient -1 on T. Every internal-edge slack is nonnegative by the module lemma. On a boundary edge and for either third-vertex bit ϵ, the minimum slack is exactly

$$1 - \epsilon + (\epsilon - 1) \sum_j (v_j)_e = 0,$$

since the prescribed responses sum to $\mathbf{1}$. This proves all elemental inequalities.

For a triangular boundary, quotienting by boundary pair coefficients leaves the three differences $m_e^1 - m_e^0$; the boundary triangle coefficient changes all three by the same amount. Thus the response quotient is $\mathbb{R}^3/\langle \mathbf{1} \rangle$, and the branch classes are represented by their prescribed vectors. Two complementary nonconstant vectors have a one-dimensional relation space in this quotient, as do the three unit vectors. Theorem 5.4 therefore proves extremality.

Every internal graph edge has a positive slack by the module lemma. For an edge of T, keep all but one branch at a conditional minimum and use the strict excess supplied by part (v) in the remaining branch. This gives a positive global slack on that edge as well. Thus every edge of G occurs as a polytope direction. Since G is connected, these directions span the $(n-1)$-dimensional zero-sum subspace. The polytope has dimension $n - 1$, and extremality gives Minkowski indecomposability.

All coefficients are integral, all quadruple coefficients are 1, and every triangle coefficient lies in $\{0, -1, -2\}$ by construction. Integral supermodular data have integral greedy vertices, so the resulting polytope is a lattice polytope up to translation.

For the dilation bound, fix a graph edge e. Its link L_e has one vertex for each triangle containing e, and one edge for each maximal K_4 containing e. It is a tree because the 3-tree construction starts its link as one edge and adds only leaves. Write $k_e = |E(L_e)|$. The constructed slack is

$$y_e - \sum_{w \in V(L_e)} d_w x_w + \sum_{uv \in E(L_e)} x_u x_v, \qquad 0 \le d_w \le \deg_{L_e}(w), \quad d_w \in \mathbb{Z}.$$

Every edge has a zero slack, as is explicit in the construction.

At each link vertex w, mark any d_w of its incident edges. On a link edge uv, let $a, b \in \{0, 1\}$ indicate whether it was marked at u, v. The nonconstant part of the slack is the sum, over link edges, of

$$x_u x_v - a x_u - b x_v.$$

Each such function has range of length one on its four binary inputs. The whole slack therefore has range of length at most k_e. Its minimum is zero, so every elemental slack lies between 0 and k_e. There are $n - 3$ maximal cliques, hence $k_e \le n - 3$. The elemental slacks of Z_G equal one on every graphical edge and zero on nonedges. Consequently $(n-3)Z_G - Q$ has nonnegative elemental slacks and is represented by a polytope. Thus Q is an actual Minkowski summand of $(n-3)Z_G$. □

Remark 6.3. The construction traverses the clique–triangle incidence tree and uses integer arithmetic. The coefficient -2 at a triangle of multiplicity three is what permits all six response directions to propagate; restricting every triangle coefficient to 0 or -1 would unnecessarily exclude some of the required boundary states.

Corollary 6.4 (Complete criterion for planar graphs)**.** *Let G be a simple planar graph on $n \ge 4$ vertices. Its graphical zonotope has a Minkowski-indecomposable weak summand of dimension $n-1$ if and only if G is a planar 3-tree, equivalently a stacked triangulation.*

Proof. Suppose such a summand exists and let $H \subseteq G$ be its actual direction graph. Theorem 4.1 and Euler's planar bound give

$$3n - 6 \le |E(H)| \le |E(G)| \le 3n - 6.$$

Equality in the first bound forces H to be a 3-tree, and equality throughout forces $G = H$.

Conversely, in a planar 3-tree no triangle can belong to three distinct maximal K_4s; its three vertices and the three distinct opposite vertices would contain a $K_{3,3}$ subgraph. Every triangle therefore belongs to at most two maximal K_4s, so Theorem 6.1 supplies the required full-dimensional indecomposable summand. The usual stacked triangulation description is also immediate; start with an embedded tetrahedron and attach each new vertex inside its attachment face; the multiplicity bound two ensures that no face is reused. □

Corollary 6.5 (Maximum dimension and the full spectrum for planar graphs)**.** *Let G be any simple planar graph. Form the graph whose vertices are the K_4 subgraphs of G, joining two when they share a triangle. For each connected component $\mathcal{C}$, let $V(\mathcal{C})$ be the union of the vertices in its tetrahedra. Then*

$$\operatorname{mdim}(Z_G) = \begin{cases} \max\limits_{\mathcal{C}} \big(|V(\mathcal{C})| - 1\big), & \text{if } G \text{ contains a } K_4, \\ 2, & \text{if } G \text{ has a triangle but no } K_4, \\ 1, & \text{if } G \text{ has an edge but no triangle}, \\ 0, & \text{if } G \text{ is edgeless.} \end{cases}$$

If this maximum is $d > 0$, the positive dimensions of Minkowski-indecomposable weak summands are exactly $1, 2, \ldots, d$. Every one of these dimensions is attained by a lattice polytope.

Proof. Consider the union H of a triangle-connected family of K_4s on v vertices. Start from one tetrahedron and traverse a spanning tree of the triangle-adjacency graph. Each newly encountered vertex brings at least three new edges, so $|E(H)| \geq 3v - 6$. Since H is planar, Euler's bound forces equality. The equality argument in Theorem 4.1 shows that H is a 3-tree; steps adding no vertex add no edge, and every new vertex attaches to exactly its shared triangle. It is a planar 3-tree and hence admissible by the preceding corollary.

The positive triangle-class decomposition (Lemma 4.2) says that an indecomposable summand is supported either on a segment direction, on an isolated triangle class, or on one triangle-connected K_4 component. Its dimension is therefore bounded by the corresponding value in the formula. Conversely, each such component is a planar 3-tree and has a full-dimensional lattice indecomposable ray, which embeds in the deformation cone of G. A triangle simplex and an edge segment give the lower-dimensional cases.

Finally, a planar 3-tree on v vertices has construction prefixes on every number of vertices from four through v. These prefixes are again planar 3-trees, so their full-dimensional rays attain every dimension from three through $v - 1$. Their triangles

and edges supply dimensions two and one. This proves the no-gap assertion, including the cases with no K_4. □

6.3. **All possible direction supports and the maximum dimension**

A *partial* 3*-tree* is a graph contained in a 3-tree on the same vertex set; equivalently, it has treewidth at most three. The usual small-vertex cases are understood separately. This class includes every chordal K_5-free graph.

Corollary 6.6. *Let G be a partial 3-tree. After isolated vertices are deleted, the actual direction graphs of its Minkowski-indecomposable weak summands of dimension at least three are exactly the induced subgraphs $G[U]$ which are 3-trees and in which every triangle belongs to at most three maximal K_4's. Such a summand has dimension $|U| - 1$.*

In particular, for $n \geq 4$ a partial 3-tree has a full-dimensional indecomposable weak summand if and only if it is itself a 3-tree satisfying the triangle condition of Theorem 6.1.

Proof. An indecomposable generalized permutahedron has only one nontrivial connected component in its actual direction graph; distinct components separate the coefficient function and give a Minkowski decomposition. Let U be that component's vertex set. The sharp direction bound, Theorem 4.1, gives at least $3|U| - 6$ actual directions when the dimension is at least three. On the other hand, an induced subgraph of a partial 3-tree on $|U| \geq 4$ vertices has at most $3|U| - 6$ edges. Equality is therefore forced. The actual graph equals $G[U]$, and the equality case in the direction bound makes it a 3-tree. Theorem 6.1 supplies its triangle condition.

Conversely, apply that theorem on any such $G[U]$ and fix all coordinates outside U. The resulting polytope remains indecomposable, since any Minkowski summands of a polytope in an affine subspace can be translated into the parallel subspace. It is a weak summand of Z_G, since unused graphical segments can be added to its complement. Its dimension is $|U| - 1$. □

For a 3-tree G, call a nonempty collection of its maximal cliques *admissible* if it is connected through shared triangles and uses at most three cliques at any triangle. A connected clique collection is itself a 3-tree; order its clique nodes outwards from any one node in the incidence tree. Each subsequent clique meets the earlier union in its attachment triangle and introduces one new vertex. The running intersection property, or the original 3-tree construction, excludes any repeated new vertex. If it contains m cliques, its union has $m+3$ vertices and $3m+3$ edges. There cannot be an

additional edge of G between these vertices, since that would exceed the partial-3-tree bound $3(m+3)-6$. Thus this union is an induced 3-tree.

Corollary 6.7 (Dimension spectrum and a tree algorithm)**.** *For a 3-tree G, let $M(G)$ be the largest number of maximal cliques in an admissible collection. Then the maximum dimension of an indecomposable weak summand of Z_G is*

$$\operatorname{mdim}(Z_G) = M(G) + 2.$$

Every integer dimension from 1 through $M(G)+2$ occurs. Both $M(G)$ and a maximizing support can be found in linear time on the clique–triangle incidence tree.

Proof. Corollary 4.4 says that the actual directions are the union of a triangle-connected family of nonzero four-clique coefficients. In a 3-tree this is a connected clique collection of the kind just described. By Corollary 6.6, its triangle multiplicities are at most three. Conversely every admissible collection supports the full-dimensional ray constructed above. A collection of m cliques has dimension $m+2$, proving the formula.

Starting with a maximizing collection and repeatedly removing a leaf clique yields admissible collections of every size from $M(G)$ down to one. They give all dimensions from three through $M(G)+2$. A graphical segment and a triangle simplex give dimensions one and two.

The dynamic program has the following form. For an incidence edge $C - T$, let $A_{C\to T}$ be the largest number of selected cliques in the branch on the C side, requiring C to be selected. Let $B_{T\to C}$ be the largest contribution on the T side when C is already selected. Then

$$A_{C\to T} = 1 + \sum_{F\subset C,\ F\neq T} B_{F\to C},$$

$$B_{T\to C} = \text{the sum of the two largest numbers among } \{A_{D\to T} : D \supset T,\ D \neq C\}.$$

Use all available numbers if there are fewer than two. Choosing a root clique C gives the value

$$1 + \sum_{T\subset C} B_{T\to C},$$

and the maximum of these values is $M(G)$. A bottom-up pass followed by a top-down pass computes all directed messages. At a triangle node it is enough to retain the three largest incoming values, so excluding any one neighbor and taking the best two is constant-time work per incidence. The incidence tree has linear size. Retaining the maximizing children also returns the clique collection and hence a concrete coefficient construction on that support. □

Corollary 6.8 (An intrinsic algorithm for partial 3-trees)**.** *Let G be a partial 3-tree. Its bipartite incidence graph between its K_4's and its triangles is a forest. If G contains a K_4, let M be the maximum number of clique nodes in a connected selection in this forest which uses at most three clique nodes at each triangle. Then*

$$\operatorname{mdim}(Z_G) = M + 2.$$

If there is no K_4, the maximum is two when a triangle exists, one when there is an edge but no triangle, and zero when there are no edges. These values and a maximizing support can be computed in linear time directly from G, without constructing a 3-tree completion.

Proof. For $|V(G)| < 4$ all assertions are immediate. Otherwise choose any 3-tree completion only for the proof. Every K_4 of G is a maximal clique of the completion, and every triangle of G is one of its triangle nodes. The asserted incidence graph is therefore a subgraph of the completion's incidence tree, hence a forest. A connected clique selection gives a 3-tree union, and the partial-3-tree edge bound forces that union to be induced in G. Corollary 6.6 gives both directions of the dimension formula. Without a K_4, that corollary excludes dimension at least three, and the stated simplices or segments attain the remaining bounds.

Run the top-two-message dynamic program of Corollary 6.7 separately on every component of the incidence forest and take the best answer. Its size is linear in $|V(G)|$. Indeed, a partial 3-tree is 3-degenerate; from an elimination order in which each vertex has at most three later neighbors, enumerate every clique by its earliest vertex and a subset of those later neighbors. There are at most eight candidates per vertex, and all edge checks take constant time using the bounded lists of later neighbors. The order itself is obtained by the usual degree-bucket elimination in $O(|V|+|E|) = O(|V|)$ time. Thus clique enumeration, forest construction, the dynamic program, and recovery of a maximizing support all take linear time. □

6.4. **Books, tree joins, and Hamiltonian graphs**

Several graph families now follow from the same triangle criterion. For a tree T, let $s_3(T)$ be the largest number of vertices in a connected subtree of maximum degree at most three.

Corollary 6.9. *For $k \geq 1$ and every nonempty tree T,*

$$\operatorname{mdim}(Z_{K_3 * I_k}) = \min\{k+2, 5\}, \qquad \operatorname{mdim}(Z_{K_2 * T}) = s_3(T) + 1.$$

Thus $K_2 * T$ supports a full-dimensional indecomposable weak summand exactly when T has maximum degree at most three. Every Hamiltonian 3-tree also supports such a summand.

Proof. The maximal cliques of $K_3 * I_k$ share one triangle, so an admissible collection has at most three members. In $K_2 * T$, with hubs a, b, the maximal cliques are $\{a, b, u, v\}$ for $uv \in E(T)$. A connected clique collection corresponds to a subtree U, and the multiplicity of $\{a, b, v\}$ is $\deg_U(v)$. The dimension formula of Corollary 6.7 gives both assertions. If T has one vertex, its join is a triangle and the formula is immediate. Finally, deleting three vertices from a Hamiltonian cycle leaves at most three paths. Every triangle of a Hamiltonian 3-tree therefore satisfies the component condition of Theorem 6.1. □

In particular, the common-triangle book on six vertices has an indecomposable summand of dimension five. This corrects the computational claim in [1, Example 4.3, version 5]; the first book without a full-dimensional ray has seven vertices.

The support classification also explains the linear ray counts in [1, Example 4.2]. Write $r_{n,d}$ for the number of d-dimensional rays on $K_2 * P_{n-2}$, and put $a_d = r_{d+1,d}$. For $3 \le d \le n-1$,

$$r_{n,d} = (n-d)a_d.$$

Indeed, a d-dimensional direction support consists of the two hubs and an interval of $d-1$ path vertices. There are $n-d$ such intervals. For each one, setting every omitted edge slack to zero identifies its deformation cone with a face of the full cone. Its full-dimensional rays remain extreme in that face and use the entire interval, so this counts each ray exactly once. Poullot's enumeration gives

$$a_3 = 23, \quad a_4 = 96, \quad a_5 = 378, \quad a_6 = 1542, \quad a_7 = 6450.$$

The segment and triangle counts are respectively $3(n-2)$ and $6n-16$.

7. Clique books

A clique book is a join $K_s * I_k$, with the independent outer vertices forming branches over the central clique. We now construct their responses so that the separator bound is attained for every admissible number of branches.

Theorem 7.1 (Exact clique-book classification)**.** *For every $s \geq 2$ and $k \geq 1$, let*

$$b_s = \binom{s}{2} 2^{s-2} - (2^s - s - 1) + 1.$$

Then

$$\operatorname{mdim}(Z_{K_s * I_k}) = \min\{s + k - 1,\ s + b_s - 1\}.$$

In particular, $K_s * I_k$ has a full-dimensional indecomposable weak summand if and only if $k \leq b_s$. A maximum-dimensional summand can always be chosen to be a lattice polytope and an actual summand of

$$(s-1)\min\{k, b_s\}\, Z_{K_s * I_k}.$$

The upper bound follows immediately from separator capacity. For the reverse inequality, each outer vertex will carry a monotone Boolean function on the hubs. Its zero spoke equations leave one scale, and its boundary response is an explicit vector. We construct positive circuits of these vectors of every size up to b_s.

7.1. Clique separators and Boolean response modules

Theorem 7.2 (A clique-separator bound)**.** *Let S be a clique of size $s \geq 2$ in a graph G. If an extreme deformation ray has nonzero private coefficient blocks in r components of $G - S$, then*

$$r \leq b_s.$$

Here b_s is the capacity in Theorem 7.1; in particular, $b_2 = 1$, $b_3 = 3$, and $b_4 = 14$.

Proof. There are $d_s = 2^s - s - 1$ boundary clique coefficients and $m_s = \binom{s}{2} 2^{s-2}$ boundary edge/context types. Apply Lemma 5.1. In a book, each outer vertex is a component. An outer vertex whose private block is zero has fixed coordinate after translation, giving the upper dimension bound in Theorem 7.1. □

7.1.1. **Boolean modules and positive response circuits.** Let $F : 2^S \to \{0, 1\}$ be monotone, with $F(\varnothing) = 0$ and $F(S) = 1$. For an outer vertex v, the function

$$f_v(A) = \mathbf{1}_{v \in A} F(A \cap S)$$

has clique support in $S \cup \{v\}$. Its spoke slacks are $F(T \cup \{a\}) - F(T) \in \{0, 1\}$.

The zero spoke equations force every private variation to be a scalar multiple of F; the zero-valued sets form a downward-closed family connected to $\varnothing$ in the Boolean cube, while the one-valued sets form an upward-closed family connected to S. The zero cover differences make the variation constant on each family; normalization sets the first constant to zero, leaving one scalar on the second family. This argument includes every unused private coefficient.

Index central elemental slacks by $e \in \binom{S}{2}$ and $T \subseteq S \setminus e$. Put

$$D_F(e, T) = \delta_e F(T), \qquad N_F(e, T) = \max\{0, -D_F(e, T)\}.$$

Because F is monotone Boolean, $D_F \in \{-1, 0, 1\}$. Let $A_s : \mathbb{R}^{d_s} \to \mathbb{R}^{m_s}$ be the injective map from central Möbius coefficients to their elemental slacks.

Lemma 7.3 (Positive circuits of Boolean modules)**.** *Suppose the classes of $N_{F_1}, \dots, N_{F_k}$ in $\mathbb{R}^{m_s}/\operatorname{im} A_s$ have a one-dimensional relation space spanned by a strictly positive vector w, with*

$$\sum_i w_i N_{F_i} = A_s z.$$

Then

$$f(A) = z(A \cap S) + \sum_i w_i \mathbf{1}_{v_i \in A} F_i(A \cap S)$$

spans an extreme ray of the deformation cone of $Z_{K_s * I_k}$. If the essential spoke directions form a connected spanning graph, its dimension is $s + k - 1$. In particular, this holds when every F_i depends on every hub.

Proof. Every spoke slack is nonnegative. A central edge slack, at fixed central context, is

$$\sum_i w_i \big(N_{F_i}(e, T) + x_{v_i} D_{F_i}(e, T)\big). \tag{7.1}$$

Each summand is 0, $w_i x_{v_i}$, or $w_i(1 - x_{v_i})$. This proves all inequalities and makes every central conditional minimum zero.

The spoke zeros already leave one private scale per module, so each private kernel is a line. The conditional response of f_{v_i} is $-N_{F_i}$. Theorem 5.4 therefore identifies the complete zero space with the stated positive relation, proving extremality. Each essential spoke has a positive slack, and connected essential spokes span the full zero-sum space. □

7.1.2. **The response quotient and pure signed coordinates.** For the constructions, we need coordinates in which boundary compatibility can be checked one subset of hubs at a time. Expanding each response in monomials does this; the same clique coefficient appears in several edge differences, and compatibility requires those copies to agree. After quotienting out their common value, different hub subsets become independent blocks.

Write each pair response as a multilinear polynomial in its remaining hub bits. Its monomial on $T \subseteq S \setminus e$ is indexed by $C = e \cup T$. A central coefficient y_C contributes that same coefficient to every pair $e \subset C$. These are exactly the integrability conditions. Consequently

$$\mathbb{R}^{m_s} / \operatorname{im} A_s \simeq \bigoplus_{\substack{C \subseteq S \\ |C| \geq 3}} \mathbb{R}^{\binom{C}{2}} / \langle \mathbf{1} \rangle.$$

Write $p_{C,e}$ for the class of a coordinate vector in the C block. Then

$$\sum_{e \in \binom{C}{2}} p_{C,e} = 0,$$

and omitting any one coordinate gives a basis of the block.

Lemma 7.4 (Both signs of a pure coordinate)**.** *For $e = \{i, j\} \subset C$, $|C| \geq 3$, define*

$$P_{C,e} = (x_i \vee x_j) \prod_{k \in C \setminus e} x_k,$$

$$R_{C,e} = \left(\prod_{i \in C \setminus e} x_i \right) \vee \bigvee_{\ell \in C \setminus e} \left(\prod_{i \in C \setminus \{\ell\}} x_i \right).$$

Here $\vee$ is Boolean disjunction. Their response classes are $+p_{C,e}$ and $-p_{C,e}$, respectively. Both functions depend essentially on every variable of C and ignore the other hubs.

Proof. The only negative second difference of $P_{C,e}$ is on its own pair e, when all other variables of C are one. Thus its negative response is the monomial for that one coordinate.

The minimal winning sets of $R_{C,e}$ have complements given by the block e and the singleton blocks of $C \setminus e$. Their pairwise unions are C. A negative second difference can therefore occur only in a full C context, and there precisely on pairs joining distinct blocks. Those are all pairs except e. The resulting class is $\sum_{f \neq e} p_{C,f} = -p_{C,e}$. The displayed formulas show essentiality. □

When a positive relation among these integer response vectors is specified, its central lift is constructive and integral; expand the weighted response sum in each pair coordinate and take the common coefficient in every C block. Formula (7.1) then bounds every slack by the sum of the positive weights. This gives an actual lattice Minkowski complement in that integer multiple of the graphical zonotope.

7.2. **Constructive intervals for clique books**

The number of outer vertices is the number of columns in a positive response circuit. We must therefore vary that number without enlarging the relation space. For small counts, one full-context Boolean response cancels a chosen set of pure coordinates. Threshold functions then reach many blocks at once, and splitting the cancelling columns fills larger intervals. Near capacity, several threshold functions are needed to reach every block; connector blocks first fix their relative scales. We construct these families in that order, then check that their ranges cover every $k \le b_s$. All modules are monotone Boolean functions, so their private kernels and positivity are controlled uniformly by Lemma 7.3.

7.2.1. **Every count up to the number of hub pairs.**

Theorem 7.5. *For every $s \ge 2$ and $1 \le k \le \binom{s}{2}$, the graph $K_s * I_k$ has a full-dimensional lattice Minkowski-indecomposable weak summand. The construction for $k \ge 2$ is an actual summand of $kZ_{K_s * I_k}$.*

Proof. We first note that every integer $0 \le m \le \binom{s}{2}$ is the number of edges of a complete multipartite graph on some subset of an s-element set S; vertices outside that subset are isolated. This follows by induction on s. Counts at most $\binom{s-1}{2}$ use the induction hypothesis. For a larger count, put $j = \binom{s}{2} - m$, so $0 \le j \le s-2$. Write $j = 3q + r$, with $0 \le r \le 2$. Partition S into q blocks of size three, r blocks of size two, and singleton blocks for all remaining vertices. This uses $3q + 2r = j + r \le s$ vertices before the singletons are added, and the sum of the numbers of within-block pairs is exactly $3q + r = j$. Its complete multipartite graph therefore has the required m edges.

The case $k = 1$ is the simplex on the $s+1$ vertices of the complete graph $K_s * I_1$. Suppose $k \ge 2$ and put $m = k - 1$, so $1 \le m < \binom{s}{2}$. Choose a complete multipartite graph J on a subset $U \subseteq S$ with exactly m edges, and let $B_1, \dots, B_r$ be its nonempty independent parts. Since $m > 0$, there are at least two parts. Define the monotone

Boolean function

$$F_J = \bigvee_{j=1}^{r} \left(\prod_{i \in S \setminus B_j} x_i \right).$$

Its minimal winning sets have pairwise union S, with pairwise disjoint complements B_j. Hence its negative second differences occur precisely on the edges of J, and only when every other hub is present. In response coordinates,

$$N_{F_J}(e, T) = \mathbf{1}_{e \in E(J)} \mathbf{1}_{T = S \setminus e}.$$

This also follows directly by considering the two missing hubs; removing one vertex from each of two different B_j destroys every winning set, while either single removal leaves one winning set.

Use one outer module F_J and one outer module $R_{S,e}$ for every $e \in E(J)$, all with weight one. Their total negative response is m on the full context of every hub pair and zero in every other context; a pair in $E(J)$ receives one contribution from F_J and $m - 1$ from the R modules, while every other pair receives m from the R modules. Thus the central function is explicitly

$$g = m \prod_{i \in S} x_i.$$

All conditional minima vanish after adding g, and the complete slacks are sums of k nonnegative zero–one literals.

The response relation is unique. Let c_0 be the scale of F_J and c_e that of $R_{S,e}$. Choose a hub pair outside $E(J)$, which exists because $m < \binom{s}{2}$. Comparing its full-context response with that of $e \in E(J)$ forces $c_e = c_0$. Thus the relation space is exactly the positive line with all scales equal. The Boolean private-kernel lemma now proves extremality in the nonnegative deformation cone.

Finally, every variable is essential in F_J; variables outside U occur in every winning set, and a variable in B_i can be made pivotal by omitting one variable from a different block. Every $R_{S,e}$ is also essential in all s hubs. Their positive spoke directions connect all $s + k$ vertices, so the dimension is $s + k - 1$. All coefficients are integral and every slack is at most k, proving the asserted actual lattice summand. □

7.2.2. **Signed-star interpolation.** A single threshold module reaches many response blocks at once. Once its scale is fixed, cancellation in a block of independent columns determines every other scale there. We can therefore add branches by splitting a column into independent columns with the same sum, preserving the unique positive relation. The following interpolation performs these splittings within each block and fills the required intervals of branch counts.

Lemma 7.6 (The threshold-star response)**.** *On a hub set S of size $s \geq 3$, let*

$$H_a(x) = x_a \ \vee \ \mathbf{1}_{\sum_{i \neq a} x_i \geq 2}.$$

This function depends essentially on every hub. In the response block indexed by $C \subseteq S$, $|C| \geq 3$, its quotient class is

$$[N_{H_a}]_C = \begin{cases} (-1)^{|C|-3} \displaystyle\sum_{e \ni a} p_{C,e}, & a \in C, \\ 0, & a \notin C. \end{cases} \tag{7.2}$$

Proof. On a pair containing a, the negative second difference is the indicator that exactly one other hub is present. On a pair not containing a, it is

$$(1 - x_a)\, \mathbf{1}_{\sum_{i \notin e \cup \{a\}} x_i = 1}.$$

The coefficient of a monomial of positive degree r in an exactly-one indicator is $r(-1)^{r-1}$; its constant coefficient is zero. Consequently, if $a \in C$, the C-block coefficients are

$$(|C| - 2)(-1)^{|C|-3} \quad (e \ni a), \qquad (|C| - 3)(-1)^{|C|-3} \quad (e \not\ni a).$$

Subtracting the common coefficient leaves the signed star in (7.2). If $a \notin C$, every pair coefficient is $(|C| - 2)(-1)^{|C|-3}$, so the quotient class is zero. Finally, a is pivotal when all other bits are zero, and any other hub is pivotal when $a = 0$ and exactly one further hub is present. □

Lemma 7.7 (Every positive basis size for a signed star)**.** *Let C have $r \geq 3$ vertices, fix $a \in C$, and write*

$$S_a = \sum_{i \in C \setminus \{a\}} p_{C,\{a,i\}}, \qquad d_C = \binom{r}{2} - 1.$$

For either sign, and every $1 \leq t \leq d_C$, the vector $\pm S_a$ is a sum of t linearly independent pure Boolean response columns, all with coefficient one. Every chosen module depends essentially on every vertex of C.

Proof. We use two basic pure response graphs. A partial star with center i and nonempty neighbor set $B \subseteq C \setminus \{i\}$ is realized by

$$\left(\prod_{j \in C \setminus \{i\}} x_j \right) \vee \left(\prod_{j \in C \setminus B} x_j \right).$$

A clique on a subset $U \subseteq C$, $|U| \geq 2$, is realized by

$$\bigvee_{i \in U} \left(\prod_{j \in C \setminus \{i\}} x_j \right).$$

In both cases the minterm complements are disjoint, so the negative response is exactly the indicated full-context graph. Variables outside that graph occur in every minterm and remain essential. We also have the pure negative-edge module $R_{C,e}$, with response $-p_{C,e}$.

For $+S_a$, begin with its single star module. Split its neighbor set into nonempty parts, one at a time, until all $r-1$ individual edges ai are separate columns. At every stage these partial stars have disjoint supports and are independent, since their union omits every pair inside $C \setminus \{a\}$, so no nonzero combination can be a constant vector on all pairs. This realizes every count from one to $r-1$.

Fix one reference pair e_* inside $C \setminus \{a\}$. Regard the individual edge ai as a partial star centered at i. For every other pair ij inside $C \setminus \{a\}$, assign it to one endpoint, say i, and perform the replacement

$$V_i \quad \longmapsto \quad (V_i + p_{C,\{i,j\}}), \quad -p_{C,\{i,j\}}.$$

The first new column is still a partial star centered at i; the second is $R_{C,\{i,j\}}$. The sum is unchanged and all coefficients remain one. The current span is the span of the original star edges and the pairs already activated. Since the reference pair is always omitted, the individual pair vectors in this span are independent. Each new pair therefore adds exactly one rank and one column. At the end every pair except e_* is in the span, so the number of columns is d_C. This fills all remaining counts.

For $-S_a$, use

$$-S_a = \sum_{\{i,j\} \subset C \setminus \{a\}} p_{C,\{i,j\}}.$$

The right side is one pure clique module. Split a clique into the star of one vertex and the remaining clique, and split partial stars into smaller partial stars. Continuing to the individual edges realizes every count from one to $\binom{r-1}{2}$. The supports form a partition of the clique edge set, so all these columns are independent.

Fix $q \in C \setminus \{a\}$. For each $i \in C \setminus \{a, q\}$, use the distinct seed edge iq, centered at i, and replace it by

$$p_{C,\{i,q\}} + p_{C,\{a,i\}}, \quad -p_{C,\{a,i\}}.$$

These are again a partial star and a negative-edge module. Every step increases rank by one, because the reference pair aq remains omitted. After $r-2$ steps the count is $\binom{r-1}{2}+r-2=d_C$. This completes both signs and all counts. □

Proposition 7.8 (One distinguished threshold module)**.** *For $s\geq 3$, put*

$$a_s=2^{s-1}-s+1, \qquad U_s=(s^2+s-10)2^{s-4}+2.$$

Every book K_s*I_k with $a_s\leq k\leq U_s$ has a full-dimensional lattice indecomposable weak summand. It can be chosen as an actual summand of $kZ_{K_s*I_k}$.

Proof. Use just one distinguished function $H_a=x_a\vee \mathbf{1}_{\sum_{i\neq a}x_i\geq 2}$, with weight one. Its response is zero in every block not containing a, and by (7.2) is $(-1)^{|C|-3}S_a$ in every block containing a. In each of those blocks, independently choose between one and d_C pure columns from Lemma 7.7 to cancel this signed star.

Each block's columns are independent. Therefore, once the scale of H_a is fixed, every other scale is uniquely determined; they are all that same positive scale. The complete response relation is a positive line. The Boolean circuit lemma proves positivity and extremality, and its literal formula bounds every slack by the number k of weight-one modules. Essentiality of H_a connects all hubs, and the other modules connect their outer vertices, proving full dimension.

The minimum count is

$$1+\#\{C\subseteq S: a\in C,\ |C|\geq 3\}=1+(2^{s-1}-s)=a_s.$$

The maximum is

$$1+\sum_{r=2}^{s-1}\binom{s-1}{r}\left(\binom{r+1}{2}-1\right)=(s^2+s-10)2^{s-4}+2=U_s.$$

The interpolation lemma changes one block count at a time, so every integer between these endpoints occurs. □

Proposition 7.9 (A two-threshold bridge)**.** *The books K_4*I_k have full-dimensional lattice indecomposable weak summands for $12\leq k\leq 14$, and the books K_5*I_k have such summands for $39\leq k\leq 53$. These constructions are actual summands of $2kZ_G$.*

Proof. The construction works on any hub set S of size at least three. Choose distinct hubs a,b,c and use H_a,H_b, both with weight one. In a block containing neither a nor b, their combined quotient response is zero. In a block containing exactly one, it is a signed star, so the preceding interpolation lemma supplies any count from one to d_C.

In a block containing both, before the common parity sign the pair values are two on ab, one on pairs meeting exactly one of a,b, and zero on pairs meeting neither. The

value on ab is unique. Omitting that reference pair therefore expresses the opposite response in the d_C independent columns $\pm p_{C,e}$, $e \neq ab$, with strictly positive weights one or two.

The sole exception occurs in the connector triple $C_0 = \{a, b, c\}$, where we use just $R_{C_0,\{a,b\}}$. If the two distinguished scales in an arbitrary relation are u, v, the connector's three pair values force $u = v$, and then fix its own scale. Every other selected block has independent columns, so all remaining scales are uniquely determined. Thus the complete relation space is exactly the displayed positive line. As before, the Boolean circuit lemma gives positivity and extremality, all coefficients are integral, and all weights are at most two. Both distinguished modules depend on all hubs, giving full dimension.

For $s = 4$, the blocks containing both a, b contribute one column from the connector triple, two from the other such triple, and five from the whole four-element block. Together with the two distinguished modules this is ten columns. The two exactly-one triples each contribute one or two, giving every total from 12 to 14.

For $s = 5$, the both-containing blocks contribute five columns from the three triples, fifteen from the three quadrilaterals, and nine from the whole five-element block. Including H_a, H_b gives 31 columns. There are six exactly-one triples and two exactly-one quadrilaterals; they contribute any total from 8 to 22. Thus every count from 39 to 53 occurs. □

7.2.3. **An interval ending at the separator capacity.** To reach the full capacity, one threshold module is no longer enough; we must reach every response block. We use $s-1$ threshold modules and force their scales to be in the ratios $1 : 2 : \cdots : (s-1)$. A small set of quadrilateral blocks imposes those ratios. All other blocks then serve only to cancel the fixed response, and can be filled to their full dimensions by the same splitting operation.

Theorem 7.10. *For $s \geq 5$, put*

$$L_s = (s-2)2^{s-1} + 2s - 2 - \binom{s}{2}.$$

For every integer $L_s \leq k \leq b_s$, the book $K_s * I_k$ has a full-dimensional lattice Minkowski-indecomposable weak summand. The construction is an actual summand of $(s-1)kZ_{K_s * I_k}$.

Proof. Label the hubs $0, 1, \ldots, s-1$. Introduce the distinguished modules H_a,

$1 \le a \le s-1$, with intended weights $c_a = a$, and set $c_0 = 0$. By the threshold-star lemma, their combined response in block C has pair values

$$(-1)^{|C|-3}(c_i + c_j), \qquad \{i,j\} \subset C,$$

modulo a common constant.

We first connect all their scales using $s-2$ quadrilateral blocks. Use the following three blocks and equal-sum pair choices

C	reference pair	other omitted pair
$\{0,1,2,3\}$	$\{0,3\}$	$\{1,2\}$
$\{0,1,3,4\}$	$\{0,4\}$	$\{1,3\}$
$\{1,2,3,4\}$	$\{1,4\}$	$\{2,3\}$.

For $5 \le j \le s-1$, also use $C = \{0,1,j-1,j\}$ with reference pair $\{0,j\}$ and other omitted pair $\{1,j-1\}$. These $s-2$ connector blocks are distinct. In each, introduce a pure column for each of the four retained pairs e. If

$$\beta_{C,e} = \sum_{i \in e} c_i - \sum_{i \in r_C} c_i,$$

use $P_{C,e}$ when $\beta_{C,e} > 0$ and $R_{C,e}$ when $\beta_{C,e} < 0$, with positive weight $|\beta_{C,e}|$. These columns cancel the distinguished response in that block, whose parity sign is minus. Exactly the two omitted pairs have the reference sum, so all four weights are nonzero.

For an arbitrary relation, write d_a for the distinguished scales and set $d_0 = 0$. The two omitted pairs in each connector must have the same value. The first three equations are

$$d_3 = d_1 + d_2, \qquad d_4 = d_1 + d_3, \qquad d_1 + d_4 = d_2 + d_3.$$

They give $d_2 = 2d_1, d_3 = 3d_1, d_4 = 4d_1$. The remaining equations are $d_j = d_1 + d_{j-1}$, so every scale is $d_j = jd_1$. Once these scales are fixed, the four retained pure pair columns in each connector are independent and uniquely determined.

For an ordinary block C, write

$$S_i = \sum_{j \in C \setminus \{i\}} p_{C,\{i,j\}}.$$

These star vectors satisfy $\sum_{i \in C} S_i = 0$. Any $|C|-1$ of them are independent; if an endpoint-sum vector $a_i + a_j$ is constant over all pairs, then all a_i are equal, and setting one $a_i = 0$ forces all of them to vanish. If $|C|$ is odd, let h have the largest

c_i in C; if $|C|$ is even, let ℓ have the smallest. The response needed to cancel the distinguished modules is, respectively,

$$-\sum_{i\in C} c_i S_i = \sum_{i\neq h}(c_h - c_i)S_i, \qquad \sum_{i\in C} c_i S_i = \sum_{i\neq \ell}(c_i - c_\ell)S_i.$$

Thus $|C| - 1$ independent star columns suffice, with all weights strictly positive integers at most $s - 1$.

A star, or any nonempty partial star, is a pure Boolean response module. For center $i \in C$ and nonempty $B \subseteq C \setminus \{i\}$, use

$$U_{C,i,B} = \left(\prod_{j\in C\setminus\{i\}} x_j\right) \vee \left(\prod_{j\in C\setminus B} x_j\right).$$

The minterm complements are $\{i\}$ and B, so the negative response is exactly the full-context partial star $\sum_{j\in B} p_{C,\{i,j\}}$. Every variable of C is essential.

We can increase the number of independent columns in this block one at a time, until it reaches the full block dimension $\binom{|C|}{2} - 1$. Initially the selected $|C| - 1$ stars cover every pair of C. Suppose their current partial-star columns span a proper subspace V of the block. Some pair vector $p_{C,e}$ lies outside V. It belongs to a current partial star with at least two edges, since a singleton would already put it in V. Split that column into $p_{C,e}$ and the remaining nonempty partial star. Both are available Boolean modules and inherit the old positive weight. Their sum is the old column, and $p_{C,e} \notin V$ makes the new columns independent, increasing rank and count by exactly one. Their supports still cover every pair, so this continues to full block dimension and realizes every intermediate count.

At every stage the ordinary columns remain independent and their span contains the original star subspace. Hence they impose no further condition on the distinguished scales, and their own scales are uniquely fixed once $d_i = id_1$. The entire response relation remains one-dimensional and strictly positive.

Initially the number of modules is

$$(s-1) + 4(s-2) + \sum_{\substack{|C|\geq 3\\ C \text{ ordinary}}} (|C| - 1) = \sum_{|C|\geq 3} (|C| - 1) + 2s - 3 = L_s.$$

Each connector is a quadrilateral, so its three initial stars are replaced by four fixed columns. At the end every ordinary block has full dimension and every connector has one fewer. Write $D_s = b_s - 1$ for the dimension of the response quotient. The final count is

$$(s-1) + D_s - (s-2) = D_s + 1 = b_s.$$

Performing the splits one at a time realizes every integer between these endpoints.

The Boolean circuit lemma now proves positivity and extremality in the actual nonnegative cone. The central coefficients, recovered from the weighted negative responses, are integral. Every weight is at most $s-1$; this holds for the stars and distinguished modules, while the connector differences are at most $s-2$. Splitting preserves a weight. Thus the literal formula bounds all slacks by $(s-1)k$. The distinguished modules connect all hubs, and every other module connects its outer vertex to the essential variables in C. The dimension is therefore $s+k-1$. □

7.2.4. **Completion of the classification.**

Proof of Theorem 7.1. The clique-separator bound gives the upper bound. We prove existence in full dimension for every $1 \le k \le b_s$ by induction on s. The case $s=2$, with $b_2 = 1$, is the triangle simplex. For $s=3$, the small-count theorem already covers $1 \le k \le 3 = b_3$. For $s=4$, the small-count theorem covers 1 through 6, the single-threshold interval covers 5 through 12, and the two-threshold bridge covers 12 through $14 = b_4$.

For $s=5$, first observe that $K_{s-1} * I_{k+1}$ is a spanning subgraph of $K_s * I_k$; move one of the s hubs into the independent part and delete its edges to the old outer vertices. An indecomposable deformation supported on that subgraph remains an extreme ray in the larger graphical deformation cone, since every extra edge has identically zero slack and these zero equations define a face. The already-proved $s=4$ case therefore covers every $k \le 13$ for $s=5$. The single-threshold interval is $[12, 42]$, the two-threshold bridge is $[39, 53]$, and Theorem 7.10 gives $[46, 55]$. These intervals cover every count through $b_5 = 55$.

For $s \ge 6$, the same embedding and induction cover every $k \le b_{s-1} - 1$. The single-threshold construction covers $[a_s, U_s]$, and the arithmetic upper interval covers $[L_s, b_s]$. These three ranges overlap. Indeed,

$$b_{s-1} - 1 - (a_s - 1) = 2^{s-3}\left(\binom{s-1}{2} - 8\right) + 2s > 0,$$

and

$$U_s - L_s = 2^{s-4}(s-1)(s-6) + \frac{s^2 - 5s + 8}{2} > 0.$$

This proves full-dimensional existence for every $k \le b_s$.

We also track the asserted dilation. The small-count and single-threshold constructions use weights one, the two-threshold bridge uses weights at most two, and the arithmetic upper interval uses weights at most $s-1$. Their bounds are therefore at most $(s-1)k$. In an inductive embedding, the bound is $(s-2)(k+1) \le (s-1)k$

whenever $k \geq s-2$; all smaller counts are covered directly by the small-count theorem. Thus the same bound holds in every case, with integral coefficients throughout.

Finally, if $k > b_s$, embed the saturated ray on any b_s outer vertices. The other outer coordinates are fixed up to translation. Its dimension is $s + b_s - 1$, and its dilation is $(s-1)b_s$. Together with the upper bound this proves both formulas. □

Corollary 7.11 (Dimension spectrum and prescribed coordinate sets)**.** *The positive dimensions of indecomposable weak summands of $Z_{K_s * I_k}$ are exactly*

$$\{1, 2, \ldots, \min\{s + k - 1, s + b_s - 1\}\}.$$

More precisely, let U be a prescribed set of at least two coordinates, containing h hubs and t outer vertices. There exists an indecomposable weak summand of dimension $|U| - 1$ fixing all coordinates outside U exactly in the following cases

$$t = 0; \qquad h = 1,\ t = 1; \qquad h \geq 2,\ 1 \leq t \leq b_h.$$

This is a statement about the varying coordinate set; it does not require every ambient graph edge on that set to be an active direction.

Proof. Simplices on hubs and, if needed, one outer vertex give every dimension from one to s. For $1 \leq t \leq \min\{k, b_s\}$, embed the full-dimensional ray on the induced subbook $K_s * I_t$; it gives dimension $s + t - 1$. The maximum-dimension theorem excludes every larger dimension.

For the refined statement, the induced graph on U is $K_h * I_t$. If $t = 0$, use its simplex. If $h = 1$, it is a star, whose graphical zonotope has only segment rays, so full dimension occurs precisely when $t = 1$. If $h \geq 2$, the exact clique-book theorem applies. With no hubs and at least two outer vertices there is no nonzero direction. Finally, any summand fixing the outside coordinates has all its active directions in this induced graph, so the same cases are necessary as well as sufficient. □

7.3. **Sharp direction counts at saturated clique capacity**

At the maximal branch count, the separator bound also controls the number of actual directions by limiting how many outer vertices can be nonneighbors of a given hub. The ambient book may still contain additional edges with identically zero slack.

Theorem 7.12 (Minimum directions in a saturated book)**.** *Let $s \geq 2$ and let b_s be as in Theorem 7.1. Put $U_2 = 1$ and, for $s \geq 3$,*

$$U_s = b_s - b_{s-1} + 1 = (s^2 + s - 10)2^{s-4} + 2.$$

Among all full-dimensional Minkowski-indecomposable weak summands of $Z_{K_s * I_{b_s}}$, the minimum number of actual edge directions is

$$\binom{s}{2} + sU_s.$$

Equality is attained by a lattice polytope. In every such full-dimensional summand the actual hub graph is complete when $s \geq 3$, and each hub has at least U_s active outer neighbors. The equality construction has exactly U_s active outer neighbors at every hub.

7.3.1. **A lower bound from the actual hub graph.** For a graph J on a fixed hub set, write

$$m(J) = \sum_{e \in E(J)} 2^{|N_J(e)|}, \qquad D(J) = m(J) - d(J).$$
$$d(J) = \#\{C : C \text{ is a clique of } J,\ |C| \geq 2\},$$

Here $N_J(e)$ is the common-neighbor set of the ends of e. These are precisely the response and boundary-coordinate dimensions in Lemma 5.1.

Lemma 7.13. *If $s \geq 3$ and J is a proper subgraph of K_s on the same vertices, then $D(J) < D(K_s)$.*

Proof. Add a missing edge uv. Let $W = N_J(u) \cap N_J(v)$ and $t = |W|$. The new cliques are exactly $\{u, v\} \cup C$, where C is a clique of $J[W]$, including the empty clique. Their number is at most 2^t. The new edge contributes 2^t to m. For each $w \in W$, the existing edges uw and vw each gain one common neighbor, increasing their contributions by $2^{|N_J(uw)|}$ and $2^{|N_J(vw)|}$, respectively. No other edge's common-neighbor set changes. Therefore D does not decrease, and it strictly increases whenever $t > 0$. Complete J to K_s in any order. The last added edge has $s - 2 > 0$ common neighbors, so at least that increase is strict. □

Proof of the lower bound. Let H be the actual direction graph of a full-dimensional indecomposable summand, and let S be its s hub vertices. Every outer vertex has a nonzero private coefficient block, since a zero block would fix that coordinate. The

components of $H - S$ are its b_s outer vertices. Apply the general separator bound to the boundary graph $J = H[S]$

$$b_s \le D(J) + 1.$$

For $s \ge 3$, the preceding lemma forces $J = K_s$.

Fix a hub i. In $H - (S \setminus \{i\})$, every outer nonneighbor of i is an isolated component, and the hub i together with its outer neighbors is one further component. All these components have nonzero private blocks, again by full dimension. The separator $S \setminus \{i\}$ is a clique of size $s - 1$, so

$$\#\{\text{outer nonneighbors of } i\} + 1 \le b_{s-1}.$$

Thus i has at least $b_s - b_{s-1} + 1 = U_s$ outer neighbors. Summing over hubs counts every spoke once; the complete hub graph contributes another $\binom{s}{2}$ directions.

For $s = 2$, the saturated book is K_3. If an actual direction were missing, its clique-supported function would be a sum of at most two edge segments. A full-dimensional such sum is decomposable. Hence an indecomposable two-dimensional summand has all three directions, which is the claimed value $\binom{2}{2} + 2U_2 = 3$. □

7.3.2. **A full-block matching circuit.** The direction bound imposes a different requirement from the branch count alone. To attain it, each proper block C will contribute exactly $d_C = \binom{|C|}{2} - 1$ modules, essential only on C. The distinguished module scales must therefore be coupled entirely inside the full hub block. The matching construction below supplies that coupling, leaving every proper block free to use a basis.

Besides H_a from Lemma 7.6, consider

$$J_a = x_a \left(\bigvee_{i \ne a} x_i \right).$$

Its negative response on a pair not containing a is $x_a \prod_{i \notin e \cup \{a\}} (1 - x_i)$, and it is zero on pairs containing a. The same monomial calculation gives

$$[N_{J_a}]_C = -[N_{H_a}]_C$$

in every response block. Both functions are monotone Boolean and essential in every hub variable.

For a matching M in a hub set C, partition C into the matched pairs and the remaining singleton parts. The Boolean function whose minimal winning sets are the

complements of these parts has negative response graph $K_C - M$, only in its full C context. Therefore its quotient response is

$$-\sum_{e\in M} p_{C,e}.$$

We call this a negative-matching module. A nonempty submatching has the same type of realization. These functions are essential on C whenever the partition has at least two parts, as it always does in our construction below.

Equality for $s \geq 5$. Label the hubs $0, 1, \ldots, s-1$. For $1 \leq a \leq s-1$, choose

$$A_a = \begin{cases} H_a, & s \text{ odd}, \\ J_a, & s \text{ even}, \end{cases} \qquad c_a = a, \quad c_0 = 0.$$

In a block C containing a, its response is $(-1)^{|C|-s} S_a$, where $S_a = \sum_{e\ni a} p_{C,e}$. In particular, the response in the full block $C = S$ is $+S_a$.

For $1 \leq t \leq 2s-3$, let

$$M_t = \{\{i, j\} \subseteq S : i + j = t\}.$$

These are nonempty matchings which partition all hub pairs, and $M_1 = \{\{0, 1\}\}$. Omit M_1 and take a negative-matching module for each M_t, $2 \leq t \leq 2s-3$, with positive weight $t-1$. Together with the distinguished modules of weights $c_a = a$, these have a zero response class in the full block; on a pair ij with $i + j = t \geq 2$, the response is

$$i + j - (t-1) = 1,$$

and it is also one on the omitted pair 01. A constant pair vector is zero in the quotient.

This full-block relation is unique. Write d_a for the distinguished scales, w_t for the matching scales, and $d_0 = 0$. The common pair value is d_1, read from 01. Pairs in the same matching must have equal endpoint sums. In particular,

$$d_3 = d_1 + d_2, \qquad d_4 = d_1 + d_3, \qquad d_1 + d_4 = d_2 + d_3,$$

using the pairs $03, 12$, then $04, 13$, then $14, 23$. These force $d_a = a d_1$ for $a \leq 4$. For $j \geq 5$, the pairs $0j$ and $1(j-1)$ give $d_j = d_1 + d_{j-1}$, proving the same formula for all a. Finally $w_t = (t-1)d_1$ for every matching. The relation space is exactly the specified positive line.

Initially this circuit has $3s-5$ columns and rank $3s-6$. We enlarge it inside the full block until its rank is $d_S = \binom{s}{2} - 1$. The negative matching columns partition

every pair except 01. If the current total span is proper, some pair vector p_e, $e \neq 01$, lies outside it, since these pair vectors form a basis. Its current matching piece has at least two edges; otherwise $-p_e$ would already be a column. Split that piece into $\{e\}$ and the remaining nonempty submatching, giving both new modules the old positive weight. Their response sum is unchanged, while p_e outside the span increases rank and column count by exactly one. This preserves the unique positive relation and can continue until full block rank. At the end there are $d_S + 1 = \binom{s}{2}$ full-hub modules, including the distinguished ones.

For every proper block $C \subsetneq S$, $|C| \geq 3$, use exactly $d_C = \binom{|C|}{2} - 1$ pure signed pair modules. Choose as reference pair the two largest hub labels in C, whose endpoint sum is unique. The combined distinguished response has pair values $(-1)^{|C|-s}(i+j)$. Subtract its reference value and cancel every remaining coordinate with $P_{C,e}$ or $R_{C,e}$ of the appropriate sign, with weight equal to the nonzero absolute difference. All these weights are positive integers, and the selected columns are a basis of that block.

The full-block circuit fixes every distinguished scale. Each proper block then fixes all of its own scales uniquely. Thus the global response relation is a positive line. The Boolean circuit lemma gives an integral indecomposable deformation with nonnegative slacks. Its module count is

$$\binom{s}{2} + \sum_{\substack{C \subsetneq S \\ |C| \geq 3}} d_C = (d_S + 1) + (D(K_s) - d_S) = b_s.$$

The distinguished modules connect every hub and all remaining modules have essential spokes, so its dimension is $s + b_s - 1$. Every weight is at most $2s - 4$, so, for example, $(2s-4)b_s$ is an explicit integral dilation bound.

A proper-block module is essential exactly on C, while every full-block module is essential on all hubs. Therefore the number of outer nonneighbors of any hub i is

$$\sum_{\substack{C \subseteq S \setminus \{i\} \\ |C| \geq 3}} d_C = D(K_{s-1}) = b_{s-1} - 1.$$

The actual hub graph is complete by the already-proved lower-bound argument. Each hub consequently has exactly U_s outer neighbors, so the construction attains the asserted number of directions. □

7.3.3. **The smaller clique sizes.** For $s = 3$, the saturated book is $K_3 * I_3$. The complete clique-book classification supplies a full-dimensional lattice ray, and the lower bound already forces every ambient direction to be active. Its twelve edges equal $\binom{3}{2} + 3U_3$. For $s = 2$, use the triangle simplex.

For $s = 4$, the following structural circuit attains 54 directions. Use distinguished functions H_2, H_3, both with weight one. In the full four-hub block, use the four additional positive pure columns

$$p_{02} + p_{23}, \qquad p_{12} + p_{23}, \qquad p_{03}, \qquad p_{13}.$$

The first two are partial stars and the last two are individual pairs. Their sum is $S_2 + S_3$, cancelling the distinguished full-block response $-S_2 - S_3$.

This six-column full-block relation is unique. The pair 01 has coefficient zero in every column, so the common quotient constant is zero. The pairs $02, 12$ fix the first two pure scales to the scale of H_2; the pairs $03, 13$ fix the last two to the scale of H_3. The pair 23 then equates the two distinguished scales. All six scales are therefore the same positive scalar.

In each of the four triple blocks the distinguished response is a nonzero signed star; if exactly one of $2, 3$ belongs to the triple, it is that star; if both belong, their sum is the negative star of the third vertex. Use two independent positive pure columns from Lemma 7.7 to cancel it. The entire response relation is again a positive line, with all weights one. There are eight modules essential on triples and six essential on all four hubs, hence fourteen outer vertices and

$$8 \cdot 3 + 6 \cdot 4 = 48$$

active spokes. The actual hub graph is complete, giving $48 + 6 = 54$ directions. The Boolean circuit lemma proves positivity, lattice integrality, and indecomposability; the essential spokes give full dimension. This completes all cases of Theorem 7.12.

8. **Four-partite constructions**

A bound on clique size does not bound separator capacity. For a boundary $K_2 * I_p$, the hub edge has 2^p conditioning patterns, and the capacity is $2^p + p$. We first treat $p = 2$, where a sparse construction also gives edge-minimal supporting graphs beyond 3-trees, and then realize the general capacity.

8.1. **An exact bound for a nonclique separator**

The boundary need not be a clique for the response construction to attain the separator bound. The diamond provides the first example in which the resulting full-dimensional ray cannot be supported on any spanning 3-tree; its modules will also give an edge-minimal family just above the sharp direction bound.

For a complete boundary on h vertices the bound is $\binom{h}{2}2^{h-2} - (2^h - h - 1) + 1$. For the diamond $H = K_4 - uv$, with other vertices a, b, it is

$$m_H - d_H + 1 = (4 + 4 \cdot 2) - (5 + 2) + 1 = 6.$$

Theorem 8.1. *For $k \geq 1$, let $G_k = K_{1,1,2,k}$, with singleton parts $\{a\}, \{b\}$ and other parts $\{u, v\}$ and $\{t_1, \ldots, t_k\}$. Then*

$$\operatorname{mdim}(Z_{G_k}) = \min\{k + 3, 9\}.$$

In particular G_6 has a nine-dimensional lattice indecomposable summand of $4Z_{G_6}$. Its actual direction graph has $25 = 3 \cdot 10 - 5$ edges. However, G_6 has no spanning 3-tree which supports a full-dimensional indecomposable deformation.

8.1.1. **Six private modules and their positive closure.** The diamond response quotient has dimension five, so attaining its capacity requires a positive relation among six branch responses. We use four modules depending on only three boundary variables and two depending on all four, thereby also controlling the number of actual directions. The last two modules are integer-valued rather than Boolean, but their connected level sets still leave only one private scale.

Use a, b, u, v, t_i also for zero–one indicator variables. Define

$$\begin{aligned} q_1 &= av + bv - abv, & q_2 &= ab + bv - abv, \\ q_3 &= ab + au - abu, & q_4 &= b + au - abu, \\ q_5 &= b + u - bu + abu + av - abv, \\ q_6 &= a + b - ab + u - bu + abu + v - av + abv. \end{aligned}$$

Put

$$w = (1, 2, 1, 1, 1, 1), \qquad z_0 = ab + av + 2bu + abv, \qquad z = z_0 + \sum_{i=1}^{6} w_i t_i q_i.$$

Every monomial of z is a clique of G_6.

Each q_i is nondecreasing in its four Boolean variables. For $i \leq 4$ it is a monotone Boolean function; respectively it is $v(a \vee b)$, $b(a \vee v)$, $a(b \vee u)$, and $b \vee (au)$. For the two remaining functions the four first differences are

	δ_a	δ_b	δ_u	δ_v
q_5	$bu + (1-b)v$	$(1-a)(1-u) + a(1-v)$	$1 - b + ab$	$a(1-b)$
q_6	$(1-b)(1-v) + bu$	$(1-a)(1-u) + av$	$1 - b + ab$	$1 - a + ab$

and lie in $\{0,1\}$. Thus all slacks on edges incident with a private vertex are nonnegative, and are at most two after weighting.

For a boundary edge e, put $N_e(q) = \max(0, -\delta_e q)$. The negative responses are

e	$N_e(q_1)$	$N_e(q_2)$	$N_e(q_3)$	$N_e(q_4)$	$N_e(q_5)$	$N_e(q_6)$
ab	v	0	0	u	$v(1-u)$	$(1-u)(1-v)$
au	0	0	0	0	0	0
av	0	b	0	0	0	$1-b$
bu	0	0	a	a	$1-a$	$1-a$
bv	0	0	0	0	a	0

Their weighted sum is exactly the boundary slack of z_0

$$\sum_i w_i N_{ab}(q_i) = 1+v, \quad \sum_i w_i N_{av}(q_i) = 1+b,$$
$$\sum_i w_i N_{bu}(q_i) = 2, \qquad \sum_i w_i N_{bv}(q_i) = a,$$

and both sides vanish on au. Consequently

$$\delta_e z = \sum_i w_i \big(N_e(q_i) + t_i \delta_e q_i\big) \geq 0,$$

and every boundary conditional minimum is zero. The maximum of this expression over the private bits is $\sum_i w_i |\delta_e q_i|$. On ab it equals

$$3 - v + |u - v| + |u + v - 1| = 4 - v;$$

on au, bu it is two, and on av, bv it is respectively $3-b, 3-a$. Thus all boundary slacks are at most four. Together with the spoke bounds, this makes $B(z)$ a lattice summand of $4Z_{G_6}$.

8.1.2. **Private level sets and the boundary relation.** The zero first differences of $q_1, \ldots, q_4$ leave one scale per module, by the Boolean-level-set argument in Lemma 7.3. The same principle applies to the two integer-valued modules. Their values at fixed (a, b) are

(a,b)	q_5	q_6
$(0,0)$	u	$u+v$
$(1,0)$	$u+v$	$1+u$
$(0,1)$	1	$1+v$
$(1,1)$	$1+u$	$1+u+v$.

Each level set is connected along cube edges on which the function is constant. For q_5, the middle level joins the face $a = 0, b = 1$; its bottom and top levels are a down-set and an up-set. For q_6, the two middle levels join respectively through $(1, 1, 0, 0)$ and $(0, 0, 1, 1)$, as the table shows; the extreme levels are singletons.

A private variation with the same zero spoke rows is therefore a function $\varphi(q_i)$, with $\varphi(0) = 0$. The absence of the edge uv forces its mixed u, v difference to vanish. For q_5, the slice $(a, b) = (1, 0)$ gives $\varphi(2) = 2\varphi(1)$. For q_6, the slices $(0, 0)$ and $(1, 1)$ give

$$\varphi(2) = 2\varphi(1), \qquad \varphi(3) - 2\varphi(2) + \varphi(1) = 0.$$

Thus the variation is a scalar multiple of q_i in both cases, including all unused private coordinates.

It remains to identify the relations among the six negative responses. Give them scales ρ_i and ask when their sum is a boundary slack vector. Its au row is zero, so the central coefficients on au and abu vanish. Comparing the u and uv coefficients in its ab row then gives

$$\rho_4 = \rho_6 = \rho_5.$$

The bu row gives $\rho_3 = \rho_5$. Agreement of the abv coefficient in the ab, bv, av rows gives

$$\rho_1 = \rho_5, \qquad \rho_2 = 2\rho_5.$$

Hence the response relation space is exactly $\mathbb{R}(1, 2, 1, 1, 1, 1)$. Its central lift is z_0, so Theorem 5.4 proves extremality.

All five boundary edges have positive slacks somewhere. Modules q_1, q_2 use a, b, v, modules q_3, q_4 use a, b, u, and modules q_5, q_6 use all four boundary variables. There are therefore $5 + 4 \cdot 3 + 2 \cdot 4 = 25$ actual directions, forming a connected graph on all ten vertices. The polytope has dimension nine.

8.1.3. **Maximum dimension and the obstruction to a spanning 3-tree.** Apply Lemma 5.1 with $B = \{a, b, u, v\}$. Every indecomposable ray has at most six active private vertices, so its dimension is at most $4+6-1 = 9$, as well as at most $k+3$. For $k \geq 6$, the construction above, using six private vertices and setting the remaining coefficients to zero, attains nine.

For $1 \leq k \leq 5$, take a spanning tree T of $K_{2,k}$ as follows. Use one private vertex to form the path $u - t_1 - v$, and attach the remaining at most four private vertices to u, v, with at most two additional leaves at each. This is a subcubic tree. The graph $K_2 * T$ is a spanning 3-tree contained in G_k, and every triangle has multiplicity at most three. Theorem 6.1 supplies a full-dimensional indecomposable weak summand, proving the formula for all k.

Finally, every K_4 of G_6 consists of a, b, one of u, v, and one private vertex. Any spanning 3-tree contained in G_6 is consequently $K_2 * T$ for a spanning tree T of $K_{2,6}$. Such a tree has seven edges. If it supported a full-dimensional indecomposable deformation, the triangle-capacity criterion would force $\deg_T(u), \deg_T(v) \leq 3$. But every edge of T has one endpoint in $\{u, v\}$, giving $7 = \deg_T(u) + \deg_T(v) \leq 6$, a contradiction. Thus no spanning admissible 3-tree occurs. All graphs used in this theorem are K_5-free, since G_k is complete four-partite.

8.2. **An edge-minimal supporting family beyond 3-trees**

The ten-vertex ray has a triangular boundary through which the rooted modules of the 3-tree construction can be attached indefinitely. The extended graphs have one more edge than the absolute minimum, with every edge still indispensable.

Theorem 8.2. *For every $n \geq 10$, there is a K_5-free graph L_n with exactly $3n - 5$ edges and treewidth four, and a full-dimensional lattice Minkowski-indecomposable summand of $(n-6)Z_{L_n}$ whose actual direction graph is L_n. No proper edge-subgraph of L_n supports a Minkowski-indecomposable weak summand of dimension $n - 1$. In particular L_n has no spanning 3-tree supporting such a ray.*

Define the base graph L_{10} to be the actual direction graph in Theorem 8.1. More explicitly, it is $K_2 * B$, with hubs a, b, where B has vertices $u, v, t_1, \ldots, t_6$ and edges

$$vt_1,\ vt_2,\ ut_3,\ ut_4,\ ut_5,\ vt_5,\ ut_6,\ vt_6.$$

Thus B is a four-cycle with two leaves at each of two opposite vertices. Let $T = \{a, b, t_3\}$. To obtain L_{10+m}, add vertices $w_1, \ldots, w_m$ successively, giving w_i the three earlier neighbors a, b, w_{i-1}, where $w_0 = t_3$.

Lemma 8.3 (A rigid boundary in the base ray)**.** *Let z be the ten-vertex polynomial constructed above, and put $f = z - at_3$. Relative to boundary $T = \{a, b, t_3\}$, its private coefficients have a one-dimensional private zero space. Its conditional minima, in the edge order (ab, at_3, bt_3), are*

$$(m_e^{\epsilon}(f) : e \subset T) = (0, 1, 0)(\epsilon - 1).$$

Every internal edge of L_{10} retains a positive slack.

Proof. Subtracting a boundary pair changes neither an internal slack nor a private equation. For the response, the base polynomial gives

$$\delta_{at_3} z = b + (1 - b)u, \qquad \delta_{bt_3} z = a(1 - u).$$

Their conditional minima are respectively ϵ and zero. For the edge ab, all four choices of (u, v) can attain a zero when $t_3 = 0$; the third private module has $\delta_{ab} q_3 = 1 - u \geq 0$, so its zero bit always minimizes it, and the other private bits can be minimized independently. When $t_3 = 1$, choosing $u = 1$ still attains zero. Hence both conditional minima at ab are zero, proving the stated response after subtracting at_3.

For the private-kernel assertion, each of the five modules other than q_3 retains all its spoke zero equations, since none of its spoke edges lies in T. The proof above therefore fixes these five blocks up to scales ρ_i. For the third module, its private part has the form

$$t_3(Uu + Qau + Sbu + Habu).$$

The internal edge ut_3 has zero slack at $(a, b) = (0, 0), (0, 1), (1, 1)$. These equations give $U = S = 0$ and $H = -Q$, so this private block also has one scale ρ_3. We may complete it by $\rho_3 q_3 t_3$; the discrepancy is supported entirely on the four boundary cliques of T.

The old central edges au, av, bu, bv are internal relative to T. Their eight conditional minimum equations are retained. For the boundary edge ab with $t_3 = 0$, the four choices of (u, v) are all conditional minimizers, as just observed. Differences between these four minimizing rows determine the two central triangle slopes and impose the four-context compatibility at ab. Together with the eight internal minimum equations, they give the same coefficient comparisons as the preceding boundary relation

$$\rho_4 = \rho_6 = \rho_5, \qquad \rho_3 = \rho_5, \qquad \rho_1 = \rho_5, \qquad \rho_2 = 2\rho_5.$$

Thus

$$\rho = \lambda(1, 2, 1, 1, 1, 1).$$

The four old central edge coefficients other than ab, and the two old central triangle coefficients, are then the same scalar times their original values. These are all the remaining private coordinates. Hence the private zero space is precisely the line spanned by f's private vector. Boundary pair and triangle coefficients are excluded from this assertion, as required. □

Proof of Theorem 8.2. The base case is Theorem 8.1. For $m \geq 1$, let g_m be the rooted 3-tree module supplied by Lemma 6.2 on the branch with successive cliques $\{a, b, w_{i-1}, w_i\}$ and boundary T. Prescribe its response to be $(1, 0, 1)$ in the edge order (ab, at_3, bt_3). Define

$$z_{10+m} = f + g_m + ab + at_3 + bt_3 - abt_3 = z + g_m + ab + bt_3 - abt_3.$$

This is an integral, uniform recursive construction.

All internal slacks are nonnegative. On each boundary edge and each third-vertex bit, the two branch responses add to $\epsilon - 1$, and the closing boundary terms contribute $1 - \epsilon$. Therefore all six boundary minima are zero. The two complementary nonconstant responses have a unique relation in $\mathbb{R}^3/\langle(1,1,1)\rangle$, so Theorem 5.4 gives extremality.

Every old internal direction is retained, and every new internal direction occurs by the module lemma. Its strict boundary excess also makes all three edges of T actual directions. No other directions occur, since all coefficients are supported on cliques of L_{10+m}. The graph is connected, so the polytope has dimension $9 + m = n - 1$.

For the asserted dilation, every base slack has range at most four. Subtracting the boundary edge at_3 does not change any of these ranges. Every slack of g_m has range at most m; apply the edge-marking oscillation argument from the quantitative bound for rooted 3-tree modules, using its m quadruple coefficients equal to one and triangle coefficients $-d_F$ with $0 \le d_F \le \deg(F)$. This argument also applies to boundary edges, whose boundary-triangle coefficient is zero. Conditional on the third boundary bit, the new slack has minimum zero and range at most $4 + m$. Internal slacks have the same or smaller bound. Hence

$$0 \le \delta_e z_{10+m} \le m + 4 = n - 6$$

on every graph edge, proving that the polytope is an actual summand of $(n-6)Z_{L_n}$. Integral greedy vertices give the lattice claim.

Each added vertex is simplicial with three earlier neighbors, so K_5-freeness is preserved and $|E(L_n)| = 25 + 3(n - 10) = 3n - 5$. The base graph has a tree decomposition with bags $\{a, b, u, v, t_i\}$, all sharing $\{a, b, u, v\}$. Appending the bag $\{a, b, w_{i-1}, w_i\}$ at each step keeps treewidth at most four. A graph of treewidth at most three has at most $3n - 6$ edges, so the displayed edge count makes the treewidth exactly four.

There is no spanning admissible 3-tree in L_{10}, since it is a subgraph of G_6. If one existed in L_n, its last added vertex w_m would have degree exactly three; every vertex of a 3-tree has degree at least three, and w_m has only three available neighbors. Deleting this simplicial vertex leaves a 3-tree and cannot increase any triangle multiplicity. Repeating the deletion would produce a spanning admissible 3-tree in L_{10}, a contradiction.

Finally, a full-dimensional indecomposable ray on a proper edge-subgraph would have at most $3n - 6$ actual directions. The sharp direction bound forces equality and makes its direction graph a spanning 3-tree; the classification makes it admissible. The preceding obstruction excludes this possibility. Thus deleting any nonempty set of edges destroys full-dimensional supporting capacity, completing the proof. □

8.3. **Sharp exponential separator capacity in K_5-free graphs**

The diamond construction uses the two independent boundary vertices to separate its responses. For a boundary $K_2 * I_p$, there are 2^p such contexts, and a branch can be assigned to each one. Another p branches balance their side responses, attaining the full separator capacity without increasing the clique number.

For $p \geq 1$, put

$$B_p = 2^p + p, \qquad D_p = 2^p + 2p + 1.$$

These boundary capacities are attained by K_5-free graphs.

Theorem 8.4. *For every $p \geq 1$, the graph $G_p = K_{1,1,p,B_p}$ has a D_p-dimensional lattice indecomposable summand of $M_p Z_{G_p}$, where $M_p = (p-1)2^p + 2$. Moreover, for every $b \geq B_p$,*

$$\operatorname{mdim}(Z_{K_{1,1,p,b}}) = D_p.$$

8.3.1. **The explicit functions.** Let the two singleton parts be a, b, and let the third part be $U = \{u_1, \ldots, u_p\}$. Split the last part into vertices v_ϵ, indexed by $\epsilon \in \{0,1\}^p$, and additional vertices $w_1, \ldots, w_p$. In polynomial formulas, each vertex symbol denotes its zero–one membership indicator. Put $\kappa = 2^{p-1}$.

The module at v_ϵ must have negative hub response only at $u = \epsilon$, which is exactly the negative part of the affine function $\operatorname{dist}(u, \epsilon) - 1$. We realize this function as the mixed ab difference of a polynomial whose first differences are nonnegative; the extra modules then balance the remaining boundary responses.

For $\epsilon \in \{0,1\}^p$, write $h = |\epsilon|$ and $\eta_\epsilon = \mathbf{1}_{\epsilon=0}$. Define

$$\begin{aligned} q_\epsilon(a,b,u) &= \eta_\epsilon a + (p-h)b + (h-1)ab \\ &\quad + \sum_{i:\epsilon_i=0} u_i(1 - b + ab) + \sum_{i:\epsilon_i=1} u_i(a + b - ab), \qquad (8.1) \\ r_i(a,b,u) &= a(b \vee u_i) = ab + au_i - abu_i. \end{aligned}$$

Here $b \vee u_i = b + u_i - bu_i$. Take the central function

$$g(a,b,u) = ab + \kappa \sum_{i=1}^{p} bu_i$$

and set

$$f_p = g + \sum_{\epsilon \in \{0,1\}^p} v_\epsilon q_\epsilon + \kappa \sum_{i=1}^{p} w_i r_i.$$

Every monomial in this polynomial is a clique of G_p of size at most four. All coefficients are integers, and there are no constant or singleton coefficients. In particular,

$$f_p(V) = (3p-1)2^p + 2.$$

The polytope is the generalized permutahedron with lower set-function f_p.

8.3.2. **All elemental inequalities.** First consider a single module q_ϵ. Its four values as a function of a, b are

(a,b)	$q_\epsilon(a,b,u)$
$(0,0)$	$\sum_{\epsilon_i=0} u_i$
$(1,0)$	$\eta_\epsilon + \sum_i u_i$
$(0,1)$	$p-h+\sum_{\epsilon_i=1} u_i$
$(1,1)$	$\eta_\epsilon + p - 1 + \sum_i u_i$.

It is nondecreasing in every variable. For example, its increments in a at $b = 0, 1$ are

$$\eta_\epsilon + \sum_{\epsilon_i=1} u_i, \qquad \eta_\epsilon + h - 1 + \sum_{\epsilon_i=0} u_i.$$

The second is nonnegative because either $h \geq 1$ or $h = 0, \eta_\epsilon = 1$. Its increments in b at $a = 0, 1$ are

$$p - h - \sum_{\epsilon_i=0} u_i + \sum_{\epsilon_i=1} u_i, \qquad p-1,$$

and each coefficient of u_i is a zero–one function of a, b. The functions r_i are monotone Boolean functions. Therefore every spoke slack on an edge incident to a vertex v_ϵ or w_i is nonnegative.

The boundary graph on $\{a, b\} \cup U$ is $K_2 * I_p$. Its negative conditional responses are especially simple. Write $N_e(q) = \max\{0, -\delta_e q\}$. Direct differentiation gives

$$\delta_{ab} q_\epsilon(u) = \operatorname{dist}(u, \epsilon) - 1, \qquad N_{ab}(q_\epsilon) = \mathbf{1}_{u=\epsilon}, \tag{8.2}$$

$$N_{au_i}(q_\epsilon) = 0, \qquad N_{bu_i}(q_\epsilon) = (1-\epsilon_i)(1-a),$$

$$N_{bu_i}(r_i) = a, \qquad N_e(r_i) = 0 \quad (e \neq bu_i). \tag{8.3}$$

The nonzero derivatives on au_i are b when $\epsilon_i = 0$ and $1 - b$ when $\epsilon_i = 1$, so they have no negative part.

For each fixed u, exactly one module has negative ab response. For each i, exactly κ binary vectors have $\epsilon_i = 0$. It follows from (8.2)–(8.3) that the central coefficients of g cancel the sum of all negative responses exactly

$$\delta_{ab} g = 1, \qquad \delta_{au_i} g = 0, \qquad \delta_{bu_i} g = \kappa(1-a) + \kappa a = \kappa.$$

Every boundary conditional minimum is therefore zero. More explicitly, its full slack is the sum, over outer modules, of

$$N_e(q) + x_v \delta_e q$$

with the positive module weights included. Each term is nonnegative; if the derivative is negative it is $N_e(q)(1 - x_v)$, and otherwise it is $x_v \delta_e q$. This proves all elemental inequalities for all choices of conditioning variables.

8.3.3. **The zero equations force one positive scale.** We first establish the required private rigidity of each q_ϵ. Write a general variation of its private coefficient function as

$$q' = Aa + Bb + Cab + \sum_{i=1}^{p} u_i P_i(a, b),$$

where each P_i is an arbitrary bilinear function of a, b. The original P_i has three values equal to one and one value equal to zero. If $\epsilon_i = 0$, its zero is at $(a, b) = (0, 1)$; its two zero boundary derivatives occur on au_i at $b = 0$ and on bu_i at $a = 1$. If $\epsilon_i = 1$, its zero is at $(0, 0)$; its zero derivatives occur on au_i at $b = 1$ and on bu_i at $a = 1$.

The zero spoke row fixes the zero value of P_i. The two zero boundary derivatives equate its three other values. All these equations are present in the global zero face; when a module derivative vanishes, its outer membership bit can be switched between zero and one while every other module is kept at its conditional minimum. Their difference is the required equation. Consequently

$$P_i = \begin{cases} \gamma_i(1 - b + ab), & \epsilon_i = 0, \\ \gamma_i(a + b - ab), & \epsilon_i = 1. \end{cases}$$

The ab derivative vanishes at each of the p Hamming neighbors of ϵ. Switching the corresponding outer bit in these zero contexts gives

$$C - \sum_{i:\epsilon_i = 1} \gamma_i + \gamma_j = 0 \qquad (j = 1, \dots, p).$$

Thus all γ_i equal a common γ, and $C = (h - 1)\gamma$. For $p = 1$ the same single equation gives the latter identity, with only one γ_i to consider.

The zero b-spoke context at $a = 0, u = \mathbf{1} - \epsilon$ gives $B = (p - h)\gamma$. If $h > 0$, the zero a-spoke context at $b = 0$ with all $u_i = 0$ gives $A = 0$. If $h = 0$, use $b = 1, u = 0$ instead, giving $A + C = 0$ and hence $A = \gamma$. Therefore $q' = \gamma q_\epsilon$. This proves the private kernel is a line for every $p \geq 1$.

Each extra function r_i is monotone Boolean. Its zero spoke equations already force one private scale, by the connected-zero-region and connected-one-region argument of Lemma 7.3. Its ignored variables simply have identically zero spoke derivatives and introduce no additional freedom.

Now let a vector satisfy every global tight equation of f_p. The preceding argument leaves a private scale c_ϵ for each q_ϵ and a scale d_i multiplying each unweighted r_i. Write its central coefficients as

$$y_{ab} = C_0, \qquad y_{au_i} = A_i, \qquad y_{bu_i} = B_i, \qquad y_{abu_i} = t_i.$$

The two conditional minimum equations on au_i are simply

$$A_i = 0, \qquad A_i + t_i = 0,$$

because no module has a negative au_i response. Hence $A_i = t_i = 0$. The ab minimum equation at context u is then

$$C_0 - c_u = 0.$$

All 2^p main module scales equal C_0. Finally the bu_i equations at $a = 0, 1$ give

$$B_i = \sum_{\epsilon_i = 0} c_\epsilon = \kappa C_0, \qquad d_i = B_i = \kappa C_0.$$

Every coefficient has therefore been forced to the same common multiple of the constructed vector. The latter is nonzero and belongs to the nonnegative cone, so it spans an extreme ray and is Minkowski-indecomposable.

8.3.4. **Dimension, dilation, and sharpness.** All $p+2$ spokes of each main module v_ϵ are active. The u_i-spoke has a zero–one coefficient taking value one. The hub spokes are positive by the displayed monotonicity formulas, including the case $p = 1$ by choosing a suitable u-context. Each extra vertex w_i has positive directions to a, b, u_i. These directions form a connected spanning graph, so the dimension is $|V| - 1 = D_p$.

The actual direction graph omits precisely the edges $w_i u_j$ with $i \neq j$ from G_p. The edge ab and every bu_i have positive central coefficients; each au_i has a positive module derivative. Thus the actual number of directions is

$$1 + 5p + (p + 2)2^p.$$

Every clique of G_p uses at most one vertex from each of its four independent parts, verifying the K_5-free condition directly.

To obtain the quantitative dilation, fix u in an ab slack. Among the main modules, the sum of the Hamming distances is $p2^{p-1}$, and the only negative derivative is -1, at $\epsilon = u$. Hence

$$\sum_{\epsilon} \big(\operatorname{dist}(u,\epsilon) - 1\big)_+ = (p-2)2^{p-1} + 1.$$

The central coefficient contributes one, and the extra modules contribute at most $\kappa(p - |u|)$. The largest ab slack at this fixed context is therefore

$$(p-1)2^p + 2 - \kappa|u|.$$

It attains M_p at $u = 0$. The au_i and bu_i slacks are at most 2^p. Main-module spoke slacks are at most p, and extra-module spoke slacks are at most κ. For $p \geq 1$ all are at most M_p. Thus $0 \leq \delta_e f_p \leq M_p$ on graph edges and nonedge slacks vanish. This gives an actual Minkowski complement in $M_p Z_{G_p}$. All data are integral, so both generalized permutahedra are lattice polytopes up to translation.

Finally, take the separator $S = \{a, b\} \cup U$. Its induced graph has $1 + 3p$ clique coefficients of size at least two. There are 2^p boundary contexts on ab and two on each of its $2p$ other edges, for a total of $2^p + 4p$. Lemma 5.1 therefore bounds the number of active outer blocks by

$$(2^p + 4p) - (1 + 3p) + 1 = 2^p + p = B_p.$$

Any indecomposable weak summand of $Z_{K_{1,1,p,b}}$ consequently has dimension at most $p + 1 + B_p = D_p$. For $b \geq B_p$, embed the constructed polytope on any B_p of its outer vertices. Its dimension and indecomposability are unchanged, and unused graphical segments can be added to the complement. This proves the claimed exact stabilization and completes Theorem 8.4.

Corollary 8.5. *For every $p \geq 2$, the graph G_p has a full-dimensional indecomposable weak summand but contains no spanning 3-tree satisfying the triangle condition of Theorem 6.1.*

Proof. Every K_4 of G_p contains both universal vertices a, b, one vertex of U, and one vertex of the last part. A spanning 3-tree must therefore be $K_2 * T$, where T is a spanning tree of the bipartite graph between the two nonsingleton parts. Its triangle multiplicities are the vertex degrees of T, so the triangle condition requires T to be subcubic. But T has $p + B_p - 1$ edges, all incident to the p vertices of U. Subcubicity would give

$$p + B_p - 1 \leq 3p, \qquad \text{or} \qquad 2^p \leq p + 1,$$

which is false for $p \geq 2$. □

8.3.5. **Decision-tree partitions and every intermediate size.** The point indicators in (8.2) can be replaced by the cells of a decision-tree partition. This provides every remaining full-dimensional case.

Lemma 8.6 (A cube partition construction)**.** *Let a full binary decision tree partition $\{0,1\}^p$ into k cells. An internal node tests a coordinate, no root-to-leaf path tests a coordinate twice, and every coordinate is tested somewhere in the tree. For a leaf C, write its cell as*

$$C = \{u : u|_{J_C} = \epsilon_C\}, \qquad \ell_C = |J_C|,$$

and put $L = \sum_C \ell_C$. Then $K_{1,1,p,k+p}$ has a full-dimensional lattice indecomposable weak summand which is an actual summand of

$$(L - k + 2)Z_{K_{1,1,p,k+p}}.$$

Proof. For each cell, use the module (8.1) on its fixed coordinates J_C, replacing p by ℓ_C and ignoring all other u_i. Its negative ab response is exactly the indicator of C. Its au_i response has no negative part, and its negative bu_i response is $(1 - a)$ when $i \in J_C, \epsilon_{C,i} = 0$, and zero otherwise.

Let m_i and n_i be the numbers of leaves fixing coordinate i to zero and one, respectively. Because coordinate i is tested at an internal node, both its child subtrees have leaves, so $m_i, n_i > 0$. Add the p extra modules r_i with weights m_i, and use central function

$$g = ab + \sum_i m_i b u_i.$$

The cell indicators sum to one. The bu_i responses sum to $m_i(1 - a) + m_i a = m_i$, and the au_i negative responses vanish. Thus every boundary conditional minimum is zero, and the same positive factorization proves all inequalities.

Each cell module has the one-dimensional private kernel already proved, with p replaced by $\ell_C \geq 1$. Ignored coordinates have identically zero spoke derivatives and give no additional freedom. The au_i minimum equations kill all central au_i and abu_i coefficients in a variation. On a given cube cell, the ab minimum equation then equates its unique module scale with the central ab coefficient. All cell scales are equal, and the bu_i equations force the extra scales to be m_i times this common value. Hence the complete zero space is one-dimensional. Every vertex belongs to a nonzero module and the positive directions connect all vertices through the hubs, so the polytope is full-dimensional.

We record the exact dilation. Fix u. Exactly one cell module has ab derivative -1, and all the others have nonnegative derivative. Consequently

$$\sum_C \big(\operatorname{dist}(u|_{J_C}, \epsilon_C) - 1\big)_+ = \sum_C \operatorname{dist}(u|_{J_C}, \epsilon_C) - k + 1.$$

After adding the central constant and the positive contributions of the extra modules, the largest ab slack at this u is

$$2 - k + L - \sum_i n_i u_i.$$

Its maximum is $L - k + 2$, attained at $u = 0$. The other boundary slacks are at most $m_i + n_i \le k$, the extra spoke slacks are at most $m_i \le k$, and a cell-module spoke is at most its depth $\ell_C \le k - 1$. A full binary tree with k leaves has $k - 1$ internal nodes, each with at least two descendant leaves. Counting leaf depths by ancestors gives $L \ge 2k - 2$, hence $L - k + 2 \ge k$. This proves the claimed exact bound on every slack and the actual Minkowski-summand assertion. □

Theorem 8.7 (Complete four-partite classification with two singleton parts)**.** *Let $p, b \ge 1$ and put $s = \min(p, b)$. Then*

$$\operatorname{mdim}(Z_{K_{1,1,p,b}}) = \min\{p + b + 1,\ 2^s + 2s + 1\}.$$

Equivalently, this graph has a full-dimensional Minkowski-indecomposable weak summand if and only if

$$\max(p, b) \le 2^{\min(p,b)} + \min(p, b).$$

All dimensions in the formula are attained by lattice polytopes.

Proof. The graph is symmetric in its two nonsingleton parts, so assume $p \le b$. The separator bound gives the upper bound in the formula. Theorem 8.4 already gives equality for $b \ge 2^p + p$. It remains to construct full-dimensional rays when $p \le b \le 2^p + p$.

For $p \le b \le 2p + 1$, take a subcubic spanning tree T of $K_{p,b}$. An explicit choice starts with the alternating path through all p vertices of the smaller part and $p - 1$ vertices of the larger part. Attach the remaining $b - p + 1$ vertices as leaves. For $p \ge 2$, the path endpoints have room for two leaves each and its other smaller-part vertices have room for one each, giving total capacity $p + 2$. The inequality $b \le 2p + 1$ is exactly what is needed. For $p = 1$, use the star with at most three leaves. Then $K_2 * T$ is a spanning 3-tree with triangle multiplicities at most three. Its full-dimensional ray from Theorem 6.1 is also a weak summand of the larger graphical zonotope.

For $2p+1 \le b \le 2^p + p$, put $k = b - p$. There is a binary decision-tree partition using every coordinate with any prescribed number

$$p + 1 \le k \le 2^p$$

of leaves. Begin with a full binary tree with p internal nodes and label these nodes by the p different coordinates. Then repeatedly split a leaf using a coordinate not yet tested along its path. A split adds one leaf and preserves the partition. Until all 2^p point cells are reached there is a leaf of depth less than p, so a split is always possible. Lemma 8.6 supplies a full-dimensional ray on precisely $p + k = b$ outer vertices. This proves every case of the formula. □

Remark 8.8 (A balanced quantitative choice)**.** The initial p-internal-node tree can be chosen balanced, and thereafter one can always split a shallowest leaf. Its depths are then $\lfloor \log_2 k \rfloor$ and $\lceil \log_2 k \rceil$. Initially label the internal nodes by distinct coordinates; subsequent splits use any coordinate not already on the chosen path. The path depth remains at most p, and every coordinate remains in use. With $h = \lceil \log_2 k \rceil$, this gives

$$L = k(h+1) - 2^h, \qquad M = L - k + 2 = kh - 2^h + 2.$$

Thus the intermediate-size construction is an actual integral summand of MZ_G, and its largest elemental slack is exactly M.

8.4. **Comparison with spanning 3-trees**

The exponential capacity cannot be reached by using only the sparse supports of Section 6. The following criterion gives the exact limitation of that approach on every complete four-partite graph.

Theorem 8.9 (Spanning admissible 3-trees)**.** *Let $n_1, n_2, n_3, n_4 \ge 1$ and $N = \sum_i n_i$. The graph K_{n_1,n_2,n_3,n_4} contains a spanning 3-tree in which every triangle belongs to at most three maximal cliques if and only if*

$$3 \max_i n_i \le 2N - 3.$$

If L is the sum of the three smallest part sizes, the largest dimension of an indecomposable weak summand whose actual direction graph is a 3-tree is

$$\min\{N - 1,\, 3L - 4\}.$$

This last formula restricts the direction graph to a 3-tree; it is not a formula for all indecomposable weak summands of the complete four-partite graph.

Proof. Every tetrahedron uses all four colors. In a properly four-colored 3-tree, the number of triangles missing color i is $N-n_i-2$; the initial tetrahedron contributes one, and each subsequent vertex of another color contributes one new such triangle. The $N-3$ tetrahedra each contain exactly one triangle missing color i. The multiplicity bound therefore implies

$$N-3 \le 3(N-n_i-2),$$

which is the stated necessary inequality.

Conversely start with one vertex of each color, forming a tetrahedron. Every face initially has two available attachment slots. After T_0 additional vertices have been inserted, with t_i of color i, the number of available slots of missing color i is

$$2+2T_0-3t_i.$$

Attach any still-required color with a positive number of slots. Two colors cannot both have zero slots, since that would imply $3(t_i+t_j) = 4+4T_0 > 3T_0$. If the procedure stopped before reaching the target, a single required color would remain and its final availability would be negative. The assumed inequalities say precisely that every final availability $2N-3n_i-3$ is nonnegative. Thus the procedure reaches the prescribed counts without using any triangle more than three times.

For the dimension statement, retain all vertices in the three smaller parts and

$$\min\{n_{\max},\, 2L-3\}$$

vertices in the largest part. This count is at least each of the other three part sizes. The preceding criterion gives an admissible 3-tree on exactly $\min\{N, 3L-3\}$ vertices. Its full-dimensional indecomposable ray exists by the 3-tree classification. Conversely, any admissible 3-tree support uses some number $L' \le L$ of vertices outside the original largest part. The necessary inequality for that part bounds its total number of vertices by $3L'-3 \le 3L-3$. This proves the sharp restricted dimension formula. □

For example, when $b \ge p$, the largest dimension obtainable from an admissible 3-tree support in $K_{1,1,p,b}$ is

$$\min\{p+b+1,\, 3p+2\}.$$

For $b \ge 2^p+p$, the unrestricted answer from Theorem 8.7 is 2^p+2p+1. The difference from the best 3-tree support is therefore 2^p-p-1.

9. **Lattice summands**

9.1. **Sharp dilation for lattice summands on 3-trees**

For weak summands, a dilation factor has no intrinsic meaning unless a normalization is imposed, since one can always rescale the summand. Here a lattice polytope means one with vertices in the fixed standard lattice $\mathbb{Z}^n$, and the same conclusions hold for translates of such polytopes. With this normalization, the uniform factor two in Theorem 3.1 is optimal on graphs with the minimum number of directions.

Theorem 9.1. *Let G be a 3-tree on $n \geq 6$ vertices. If Q is a lattice indecomposable polytope of dimension $n-1$ and $Q + R = \lambda Z_G$ for a polytope R and $\lambda > 0$, then $\lambda \geq 2$. Equality is attained by Q_n when $G = P_n^3$. More precisely, the four-clique Möbius coefficients y_C of the lower-support function of Q are nonzero integers, and every edge e of G satisfies*

$$\lambda \geq \left\lceil \frac{1}{2} \sum_{\substack{C \supseteq e \\ |C|=4}} |y_C| \right\rceil .$$

Lemma 9.2. *If a full-dimensional indecomposable generalized permutahedron has a 3-tree as its direction graph, every four-clique Möbius coefficient is nonzero.*

Proof. Corollary 4.4 gives a triangle-connected family of nonzero four-clique coefficients whose union is the direction graph. A connected family of k four-cliques covers at most $k+3$ vertices; start with four and introduce at most one vertex at each subsequent shared-triangle step. A 3-tree has exactly $n-3$ four-cliques. Thus the family must contain all of them. □

The edge and triangle coefficients supply the constant and linear terms of an edge slack, so a lower bound depending only on tetrahedra must eliminate their contribution. Averaging at opposite sign vectors removes the linear terms; on a tree, the remaining interaction signs can then be chosen independently to give the following oscillation bound.

Lemma 9.3. *Let T be a finite tree, and let*

$$F(x) = a + \sum_{v \in V(T)} b_v x_v + \sum_{uv \in E(T)} c_{uv} x_u x_v, \qquad x \in \{0,1\}^{V(T)}.$$

Then

$$\max F - \min F \;\geq\; \frac{1}{2} \sum_{uv \in E(T)} |c_{uv}|.$$

If F is integer-valued, the right-hand side may be rounded up to the next integer.

Proof. Put $x_v = (1 + s_v)/2$, with $s_v \in \{-1, 1\}$, and average the values at s and $-s$. The linear terms cancel, leaving

$$\frac{F((1+s)/2) + F((1-s)/2)}{2} = a' + \frac{1}{4} \sum_{uv \in E(T)} c_{uv} s_u s_v$$

for a constant a'. Every such average lies between $\min F$ and $\max F$. On a tree, arbitrary prescribed edge products $s_u s_v \in \{-1, 1\}$ are realizable; fix one root sign and propagate the signs along the unique paths. We may therefore make all summands simultaneously $|c_{uv}|$, and also simultaneously $-|c_{uv}|$. The difference of these two averages is $\frac{1}{2} \sum |c_{uv}|$, proving the bound. Integer-valuedness makes the oscillation an integer. □

Proof of Theorem 9.1. The actual direction graph of Q is contained in G. By Theorem 4.1, it has at least $3n - 6$ edges, while a 3-tree has exactly $3n - 6$. Hence the actual direction graph is G itself. Lemma 9.2 makes every four-clique coefficient nonzero. These coefficients are integers; the minimum of $x(S)$ over a lattice polytope is an integer, and Möbius inversion preserves integrality. This remains true if Q is only a translate of a lattice polytope, because translations affect only singleton coefficients.

For an edge e of a 3-tree, its common-neighbor graph is a tree. This can be seen directly from a 3-tree construction. When e first appears, it has two common neighbors, joined by an edge. A later vertex becomes a common neighbor of both ends of e precisely when it is attached to a triangle containing e; in the common-neighbor graph, it is then attached as a leaf to the third vertex of that triangle. These are the only changes. The edges of this tree are in bijection with the four-cliques containing e.

The elementary difference at e, as a function of its common-neighbor indicators, is therefore a Boolean quadratic polynomial whose interaction graph is this tree

and whose quadratic coefficients are exactly the y_C for four-cliques $C \supseteq e$. The decomposition $Q + R = \lambda Z_G$ gives

$$0 \leq \sigma_e(S) \leq \lambda$$

for every conditioning set S, because the corresponding elementary difference of Z_G is one. Lemma 9.3 now gives the displayed edgewise lower bound.

Finally, a 3-tree on at least six vertices has at least three four-cliques. Their shared-triangle adjacency graph is connected, so choose three distinct cliques C_1, C_2, C_3 with C_1 adjacent to C_2 and C_2 adjacent to C_3. The triangles $C_1 \cap C_2$ and $C_2 \cap C_3$ are three-element subsets of the four-element set C_2; they intersect in an edge. That edge belongs to all three cliques. Its three nonzero integer coefficients have absolute-value sum at least three, so

$$\lambda \geq \lceil 3/2 \rceil = 2.$$

Theorem 3.1 supplies equality for P_n^3. □

The same argument gives the exact integral dilation for the original hub-and-path family.

Corollary 9.4. *For $G = K_2 * P_m$, $m \geq 2$, the smallest real $\lambda > 0$ for which λZ_G has a full-dimensional lattice indecomposable summand is*

$$\lambda_{\min} = \lfloor m/2 \rfloor.$$

In particular the smallest factor is one on four and five vertices. For the path-cube graphs P_n^3, the smallest factor is one for $n = 4, 5$ and two for every $n \geq 6$.

Proof. The graph $K_2 * P_m$ is a 3-tree. Its hub edge belongs to all $m - 1$ four-cliques, and its common-neighbor graph is P_m. The proof of the edgewise inequality above applies also for $m = 2, 3$ and gives

$$\lambda \geq \left\lceil \frac{m-1}{2} \right\rceil = \lfloor m/2 \rfloor.$$

Theorem B.1 attains this bound. On four vertices both graph families are K_4; on five vertices both are isomorphic to K_5 with one edge deleted. The remaining path-cube cases follow from Theorem 9.1. □

9.2. **The sharp limit for lattice summands without dilation**

For the unscaled zonotope, integrality requires every elemental difference to be zero or one. In a K_5-free graph these differences are quadratic, and that combination forces a

complete classification of the possible direction supports. As above, integrality refers to the fixed standard lattice $\mathbb{Z}^n$.

Theorem 9.5. *Let G be a K_5-free graph on $n \geq 4$ vertices, and suppose that a lattice polytope Q is a Minkowski-indecomposable summand of Z_G with $\dim Q = n-1$. Then G is a complete four-partite graph in which every part has size one or two, and every edge of G is an actual direction of Q. In particular, $n \leq 8$.*

Equality is attained by a seven-dimensional indecomposable lattice summand of $Z_{K_{2,2,2,2}}$.

The upper bound also holds for a translate of a lattice polytope. The intermediate part-size patterns are attained as well, as proved below.

Let $\mathcal{C}$ be the family of nonzero four-clique coefficients. Corollary 4.4 says that this family is triangle-connected and its edge union is the actual direction graph.

9.2.1. **Boolean quadratic links.** For an actual lattice summand Q of Z_G, every elementary difference satisfies

$$\sigma_e(S) \in \{0, 1\}.$$

Indeed, it is a nonnegative integer and is at most the corresponding elementary difference of Z_G, which is one on a graph edge. On the common-neighbor variables of an edge e, write

$$F_e(x) = y_e + \sum_k y_{e\cup\{k\}} x_k + \sum_{k\ell} y_{e\cup\{k,\ell\}} x_k x_\ell.$$

The *interaction graph* L_e has an edge $k\ell$ precisely when the last coefficient is nonzero. Thus its edges correspond to the members of $\mathcal{C}$ containing e. We use only nonzero interactions, not every four-clique of the ambient graph.

Lemma 9.6. *Under the hypotheses of Theorem 9.5, if $n > 4$ then every nonzero four-clique coefficient is 1 or -1. For every edge in the actual direction graph, L_e is a single edge, a path with two edges, or a four-cycle whose coefficient signs have product -1. Every triangle is contained in at most two members of $\mathcal{C}$.*

Proof. The difference of a Boolean-valued quadratic polynomial between $x_k = 1$ and $x_k = 0$ is an affine function taking values in $\{-1, 0, 1\}$. Its oscillation is the sum of the absolute values of its linear coefficients. Consequently, for every triangle T,

$$\sum_{C \supseteq T} |y_C| \leq 2.$$

In particular, L_e has maximum degree at most two. If $|y_C| = 2$, each triangle of C has no other nonzero four-clique extension. Such a C is an isolated member of the shared-triangle graph of $\mathcal{C}$. Corollary 4.4 would then force $n = 4$. Thus for $n > 4$ every nonzero four-clique coefficient is ± 1, and each triangle has at most two extensions.

The graph L_e is connected on its nonisolated vertices. Otherwise F_e would split into polynomials on two or more disjoint nontrivial variable sets, together with any isolated linear terms. Each polynomial with a nonzero quadratic coefficient is a nonconstant integer-valued function and has oscillation at least one. Oscillations of sums on disjoint variable sets add, contradicting the total oscillation at most one. By Corollary 4.4, every actual direction is contained in a nonzero four-clique, so L_e is nonempty.

A connected graph of maximum degree two is a path or a cycle. A path with at least three edges is excluded by Lemma 9.3; its oscillation is at least $3/2$. A three-cycle in L_e would make a K_5 on its three vertices together with the ends of e.

For a cycle of length k, pass to signs $s_v \in \{-1, 1\}$ and average the values at s and $-s$, as in Lemma 9.3. The nonconstant part of this average is

$$\frac{1}{4} \sum_{uv \in E(L_e)} y_{euv} s_u s_v.$$

On a cycle, all desired edge products can be prescribed except, at worst, one. Hence the maximum of the sum is at least $k - 2$, and its minimum is at most $-k + 2$. Its contribution to the oscillation is at least $(k - 2)/2$, which exceeds one when $k \geq 5$. For a four-cycle with sign product $+1$, both the all-positive and all-negative signed contributions are simultaneously realizable, giving oscillation two. The only remaining cycle is a four-cycle with sign product -1. □

9.2.2. **Completion to a join of four small parts.** The local link restrictions force a global product structure because two face changes from one four-clique give three consecutive edges in a pair link. Such a path must close to a square, so the two changes commute. Repeated square completion will turn the four original vertices into four independent choices, each with at most two possibilities.

Lemma 9.7. *Let $\mathcal{C}$ be a finite family of four-element sets, connected through shared three-element sets. Suppose each three-element set belongs to at most two members. For each pair e contained in a member, form the link with edges uv whenever $e \cup \{u, v\} \in \mathcal{C}$; assume that this link is a single edge, a two-edge path, or a four-cycle. Then there are four pairwise disjoint sets A_1, A_2, A_3, A_4, each of size one or two, such that*

$$\mathcal{C} = \{\{a_1, a_2, a_3, a_4\} : a_i \in A_i\}.$$

Proof. First consider a member C with neighbors across two distinct faces. Write its four vertices as v_1, v_2, v_3, v_4, and suppose the neighbors replace v_i by w_i and v_j by w_j, where $i \neq j$. For the pair $e = C \setminus \{v_i, v_j\}$, the interaction graph contains the three edges

$$w_i v_j, \qquad v_j v_i, \qquad v_i w_j.$$

If $w_i = w_j$, these would form a three-cycle, which is not one of the allowed links. Thus they form a three-edge path. The allowed-link condition forces its completion to the four-cycle, so

$$e \cup \{w_i, w_j\} \in \mathcal{C}.$$

This fourth member completes a unique commuting square of face changes.

Start with one member, regarded as the complete join of four singleton parts. Inductively suppose that $\mathcal{C}$ contains the complete join of four disjoint parts of sizes one or two. If this is not the whole shared-face component, choose an outside member adjacent to a member of the join. It cannot replace a vertex in a doubled part; the shared face already has its two extensions within the join. It therefore replaces a singleton vertex, say v_i, by a vertex w.

The vertex w is new to the entire join. It cannot equal a chosen vertex of another part, because the outside member has four distinct vertices. If it were the unused choice of a doubled part, the change to that choice and the change from v_i to w would give a triangular interaction link, which is forbidden. Thus w is distinct from all vertices already used.

Apply the commuting-square completion successively with every doubled part. For every old choice of one vertex from each part, it supplies the member obtained by replacing v_i by the same w. Hence the enlarged family is the complete join with the singleton part $\{v_i\}$ replaced by $\{v_i, w\}$. The vertex labels of each completed member are explicit, so the construction is consistent regardless of the order of the squares.

Each enlargement doubles one of only four parts. After at most four enlargements, either the whole connected family has been obtained, or all parts are doubled. In the latter case every face already has two extensions and no outside member can be

adjacent, so connectedness again implies that the family is complete. This proves the assertion. □

Proof of Theorem 9.5. Consider the four-cliques with nonzero coefficient. By Corollary 4.4, if $n = 4$ their edge union is K_4, giving the assertion. Otherwise Lemma 9.6 gives all hypotheses of Lemma 9.7. The edge union is precisely the actual direction graph, so this graph is complete four-partite with part sizes at most two. Its vertex count is at most eight. Any extra edge of G would join the two vertices of a part; with one vertex from each other part, it would form a K_5. Thus there are no extra edges. □

9.3. **Unit lattice summands from Boolean quadratic phases**

The four-partite restriction came from combining Boolean slacks with clique number four. Keeping the Boolean slacks while allowing more parts gives lattice indecomposable summands in arbitrarily high dimension, with indecomposable complements as well. We construct these pairs from quadratic sign functions; the four-part case also supplies the sharp examples in the K_5-free classification.

Theorem 9.8. *Let $G_r = K_{2,\dots,2}$ have $r \geq 4$ parts. The zonotope Z_{G_r} is the sum of two Minkowski-indecomposable lattice polytopes of dimension $2r - 1$. Both summands have Boolean elemental slacks and actual direction graph G_r.*

For even r, both are centrally symmetric up to translation. For odd r, the complement of a summand Q_r is a translate of $-Q_r$; thus Z_{G_r} is a translate of the difference body $Q_r - Q_r$.

We seek Boolean edge slacks whose zero sets determine them up to scale. On a complete multipartite graph, each slack belongs to a tensor product of three-dimensional polynomial spaces, where a nonseparable sign function will give precisely the required one-dimensional face.

Lemma 9.9 (A Boolean face in a multipartite polynomial space)**.** *For $d \geq 1$, let*

$$\mathcal{V}_d = \bigotimes_{i=1}^{d} \operatorname{span}_{\mathbb{R}}\{1, s_{i,0}, s_{i,1}\}, \qquad s_{i,\alpha} \in \{-1, 1\}.$$

Write $u_i = s_{i,0}$ and $v_i = s_{i,0}s_{i,1}$, so that all u_i, v_i are independent signs. Let $G(v_1, \dots, v_d)$ take values in $\{-1, 1\}$ and suppose it does not factor across any nonempty proper partition of its variables. Put

$$g(u, v) = \Big(\prod_{i=1}^{d} u_i\Big) G(v), \qquad f = \frac{1 + g}{2}.$$

Then the space of polynomials in $\mathcal{V}_d$ vanishing wherever $f = 0$ is exactly the line $\mathbb{R}f$. For $d = 1$ the nonfactorization hypothesis is empty, and the same conclusion holds.

Proof. In one part, the basis becomes $1, u_i, u_i v_i$. Hence choosing the parts on which a monomial contains u_i leaves an arbitrary function of exactly those v_i. Every $p \in \mathcal{V}_d$ therefore has a unique expansion

$$p(u, v) = \sum_{A \subseteq [d]} u_A P_A(v_A), \qquad u_A = \prod_{i \in A} u_i,$$

where P_A is an arbitrary function of the signs indexed by A. Vanishing on $g = -1$ is equivalent to the pointwise identity $gp = p$. Comparison of the u-parity components gives

$$P_A(v_A) = G(v) P_{A^c}(v_{A^c}) \qquad (A \subseteq [d]). \tag{9.1}$$

Suppose A is nonempty and proper. If either side is nonzero at any point, fixing that point on one side of the partition shows that the function on the other side is nowhere zero and has constant absolute value. Repeating in the other direction gives the same property for both functions. Equation (9.1) would then factor G across $A \mid A^c$, contrary to the hypothesis. Hence both components vanish.

Only $A = \varnothing, [d]$ remain. If $P_\varnothing = c$, then $P_{[d]} = cG$, so $p = c(1 + g) = 2cf$. Conversely every multiple of f vanishes on its zero set. □

We also record the elementary quadratic-sum calculation used for positivity and integrality. Over $\mathbf{F}_2$, put

$$q_m(\epsilon) = \sum_{1 \le i < j \le m} \epsilon_i \epsilon_j, \qquad W_m(\ell) = \sum_{\epsilon \in \{0,1\}^m} (-1)^{q_m(\epsilon) + \ell \cdot \epsilon}.$$

The associated bilinear matrix is $A_m = I + J$, where J is the all-ones matrix. For even m it is invertible and $A_m^{-1} = A_m$. For odd m its kernel consists of zero and the all-ones vector. Squaring the sum and setting the difference of its two summation variables to z gives

$$W_m(\ell)^2 = 2^m \sum_{z \in \ker A_m} (-1)^{q_m(z) + \ell \cdot z}.$$

Consequently

$$\begin{cases} |W_m(\ell)| = 2^{m/2}, & m \text{ even}, \\ W_m(\ell) \in \{0, \pm 2^{(m+1)/2}\}, & m \text{ odd}. \end{cases} \tag{9.2}$$

For even m, the normalized transform is a quadratic sign function. Completing the square shows that the exponent of $2^{-m/2} W_m(\ell + w)$, as a function of w, has quadratic

part $q_m(w)$; its associated bilinear matrix is $A_m^{-1} = A_m$. Its constant and linear terms may differ. Such a sign function is nonseparable across every nontrivial partition; a variable on each side has mixed Boolean second difference one in the exponent. Equivalently, the product of its values at the four corners obtained by flipping those two variables is minus one, whereas a separable sign function would give plus one.

9.3.1. **The even case.** Let $r \geq 4$ be even, and label the two vertices in part i by $(i,0),(i,1)$. For Boolean variables $x_{i,\alpha}$ put $s_{i,\alpha} = 1 - 2x_{i,\alpha}$. Define

$$h_\epsilon = (-1)^{q_r(\epsilon)}, \qquad \epsilon \in \{0,1\}^r.$$

The quadratic graph term below gives every edge slack $1/2$. The signed products will change each slack by exactly $1/2$ in absolute value; the quadratic-sum formula determines their common normalization. Consider

$$H(x) = \frac{1}{2} \sum_{uv \in E(G_r)} x_u x_v + 2^{-r/2-2} \sum_{\epsilon \in \{0,1\}^r} h_\epsilon \prod_{i=1}^{r} s_{i,\epsilon_i}. \tag{9.3}$$

Discard its constant and singleton terms, and let $h(S)$ be the resulting polynomial evaluated at the indicator of S. This normalizes $h(\varnothing) = h(\{v\}) = 0$. Every remaining monomial is supported on a clique of G_r.

Fix an edge between vertices (a,α) and (b,β). Taking its two Boolean differences in (9.3) gives

$$\delta_e = \frac{1}{2} + 2^{-r/2} \sum_{\eta \in \{0,1\}^{r-2}} h_{\alpha,\beta,\eta} \prod_{i \notin \{a,b\}} s_{i,\eta_i}.$$

On the remaining $d = r-2$ parts, write $s_{i,0} = u_i$, $s_{i,1} = u_i v_i$. The sum is a quadratic Walsh sum multiplied by $\prod_i u_i$. Since d is even, (9.2) yields

$$\delta_e = \frac{1}{2} + \frac{1}{2}\Big(\prod_{i=1}^{d} u_i\Big) G_e(v), \tag{9.4}$$

where G_e is a sign function with complete quadratic part q_d. Thus every edge slack is zero or one. Every nonedge slack is zero because no monomial contains two vertices of one part.

The function h is therefore supermodular, as is $g_{G_r} - h$, where $g_{G_r}(S)$ counts the graph edges induced by S. Their base polytopes satisfy

$$B(h) + B(g_{G_r} - h) = Z_{G_r}.$$

This is an actual summand identity for the unweighted zonotope.

To check integrality, let C be a clique of size $k \geq 2$, with its vertex choices fixed. Its Möbius coefficient is

$$y_C = \frac{1}{2}\mathbf{1}_{k=2} + 2^{-r/2-2}(-2)^k \sum_{\epsilon \text{ extending } C} h_\epsilon.$$

The extension sum has $m = r - k$ free bits. If k is even, its magnitude is $2^{m/2}$, so the second term has magnitude $2^{k/2-2}$. At $k = 2$ this is one half, and the displayed coefficient is zero or one; at even $k \geq 4$ it is integral. If k is odd, the extension sum is zero or has magnitude $2^{(m+1)/2}$, so the coefficient is zero or has magnitude $2^{(k-3)/2}$, again integral for $k \geq 3$. The normalized singleton and constant coefficients are zero. Both summands therefore have integral set functions and integral greedy vertices.

For extremality, let z satisfy all elemental equalities tight at h. For a fixed edge, its slack polynomial belongs to $\mathcal{V}_{r-2}$. The function G_e in (9.4) is nonseparable, so Lemma 9.9 forces the slack of z to equal $\lambda_e \delta_e$. Every maximal clique has a nonzero coefficient

$$y_{C_\epsilon} = 2^{r/2-2} h_\epsilon.$$

For every $e \subset C_\epsilon$, the coefficient of the remaining variables in its slack is the original clique coefficient, so

$$c_{C_\epsilon}(z) = \lambda_e c_{C_\epsilon}(h).$$

Since this coefficient is nonzero, all edge scales in the clique agree. Maximal cliques whose choice vectors differ in one bit share $r - 1 \geq 3$ vertices, so their scales also agree. The Boolean cube of choice vectors is connected, and every graph edge lies in one of these cliques. Hence all λ_e are one common scalar. Every nontranslation coefficient contains an edge, so $z = \lambda y$. Nonnegative slacks imply $\lambda \geq 0$.

The same argument applies to the complementary slacks $1 - \delta_e$, using $-G_e$, and proves that both summands are extreme rays. Every edge slack is nonconstant and hence positive somewhere. Both actual direction graphs are G_r, so both dimensions are $2r - 1$. Lemma 2.1 gives Minkowski indecomposability.

9.3.2. **The odd case.** Let $r \geq 5$ be odd. Start with the construction on $r + 1$ parts and restrict its set function to the vertices of any r parts, deleting one entire part. In the spin polynomial this sets both spins of the deleted part to one. Integrality and the zero-or-one slack property are inherited.

For any surviving edge, formula (9.4) becomes

$$\delta_e = \frac{1}{2} + \frac{1}{2}\Big(\prod_{i=1}^{r-2} u_i\Big)\widetilde{G}_e(v),$$

where $\widetilde{G}_e$ is obtained by fixing one variable of a complete quadratic sign function. Its quadratic part remains complete on $r-2 \geq 3$ variables. It is still nonseparable across every nontrivial partition, so the local one-dimensional-face lemma remains applicable.

The surviving maximal-clique coefficients are nonzero exactly on the even-parity choice vectors. Indeed, for a choice vector ϵ on the r remaining parts,

$$h_{(\epsilon,0)} + h_{(\epsilon,1)} = h_\epsilon\big(1 + (-1)^{|\epsilon|}\big).$$

The even-parity vectors are connected by moves flipping two bits. Two corresponding maximal cliques share $r-2 \geq 3$ vertices, so their common edge coefficients again synchronize all local scales. Every edge is contained in an even-parity maximal clique. The same extremality and dimension proof therefore works for the restricted function and its complement.

9.3.3. **Symmetry and difference bodies.** Reflecting a generalized permutahedron complements all conditioning bits in its elemental slacks. In the representation above, complementing both bits of a remaining part sends u_i to $-u_i$ and leaves v_i unchanged. It consequently multiplies g_e by $(-1)^{r-2}$.

For even r, the reflected polytope has the same elemental slacks as the original. Their set functions differ only by a modular function, so they are translates; this is central symmetry. The same applies to the complementary summand. For odd r, reflection instead exchanges δ_e and $1-\delta_e$. Thus the complement is a translate of the reflection, and Z_{G_r} is a translate of $Q_r - Q_r$. The center of Z_{G_r} is $(r-1)\mathbf{1}$, so in the odd case the identity can be written explicitly as

$$Z_{G_r} = Q_r - Q_r + (r-1)\mathbf{1}.$$

Finally, every clique of an r-partite graph uses at most one vertex from each part, so G_r is K_{r+1}-free. This completes the proof of Theorem 9.8.

Corollary 9.10 (One singleton part)**.** *For every $r \geq 4$, the unscaled graphical zonotope of $K_{1,2,\dots,2}$, with r parts, has a lattice summand of dimension $2r-2$ such that both it and its complement are indecomposable and have the entire graph as their actual direction graph.*

Proof. Deleting a vertex preserves integrality and the Boolean slack property, but may remove equations needed for extremality. We first fix a common scale on edges at the singleton. For every other edge, the two values of the singleton bit then give two distinct one-dimensional spaces for the same residual slack, forcing it to vanish.

Restrict the function of Theorem 9.8 by retaining only $a = (1, 0)$ in the first part. Integrality and Boolean slacks are inherited. For every edge incident to a, the slack still has the form (9.4) on $r - 2$ doubled parts. Hence its tight equations force one scale. The nonzero maximal-clique coefficients synchronize these scales, and their choices form a cube for even r and the even-parity cube, connected by two-bit flips, for odd r. Adjacent choices share an edge at a, and every edge at a lies in such a clique.

Subtract this common multiple from a vector in the tight kernel. All its coefficients containing a now vanish. For an edge avoiding a, its residual slack p is therefore independent of x_a and belongs to $\mathcal{V}_d$, where $d = r - 3 \geq 1$. Put $z = 1 - 2x_a$; the missing vertex has spin one. The two original slack slices are

$$f_\varepsilon = \frac{1 + \varepsilon u_1 \cdots u_d G(\varepsilon, v)}{2}, \qquad \varepsilon \in \{-1, 1\}.$$

Both phases have complete quadratic part on the d remaining variables. Moreover,

$$\frac{G(-1, v)}{G(1, v)} = \sigma \prod_{i=1}^{d} v_i, \qquad \sigma \in \{-1, 1\},$$

so f_1 and f_{-1} are distinct. The residual p vanishes on the zero set of each slice. Lemma 9.9 puts it in $\mathbb{R}f_1 \cap \mathbb{R}f_{-1} = \{0\}$. Thus the complete tight kernel is one-dimensional. Replacing every phase by its negative proves the same assertion for the complement. Every slice is nonconstant, so every edge occurs in both summands. The cone criterion and connectedness finish the proof. □

Corollary 9.11 (The four-part cases)**.** *For every $a_1, a_2, a_3, a_4 \in \{1, 2\}$, the unweighted graphical zonotope of K_{a_1,a_2,a_3,a_4} is the sum of two full-dimensional indecomposable lattice polytopes with Boolean elemental slacks. Consequently such a decomposition exists in every dimension at least three.*

Proof. The preceding results cover three or four doubled parts. With two, let a, b be the singleton vertices. Restrict Q_4 by retaining their zero-labelled vertices. Its ab slack belongs to $\mathcal{V}_2$ and fixes one scale by Lemma 9.9. Subtract that scale from a vector in the tight kernel. Its coefficients containing ab vanish. For an edge from either singleton to a doubled part, the residual slack is independent of the other singleton. The two slices belong to distinct Boolean-face lines in $\mathcal{V}_1$, as in the preceding proof, so the residual is zero. The remaining edge slacks are constant; their original slacks have a nonconstant quadratic phase on a, b, so a tight zero kills each constant. The argument applies equally to the complement.

For the two smaller graphs, write vertex names also for their Boolean variables and use

$$h_4 = ab + acd + bcd - abcd$$

on K_4. Its complement is the persimmon block of Section 5.3.1, after relabelling. The edge slacks of h_4, up to interchanging a, b or c, d, are

$$\delta_{ab}h_4 = 1 - cd, \qquad \delta_{ac}h_4 = d(1-b), \qquad \delta_{cd}h_4 = a + b - ab.$$

Their empty-context zeros kill every pair coefficient except ab. The zeros on ac and ad with b present kill abc and abd. Finally the ab zero with c, d present, and the ac, bc zeros with both remaining vertices present, leave the line

$$y_{acd} = y_{bcd} = y_{ab} = -y_{abcd}.$$

On $K_{1,1,1,2}$, with singleton vertices a, b, c and doubled part $\{d, e\}$, take

$$h_5 = ab + ad + bd - abd + ace + bce + abcd - abce.$$

Its restriction to $\{a, b, c, e\}$ is h_4. After subtracting the scale fixed there, a vector in the tight kernel has the form

$$d(Aa + Bb + Cc + Dab + Eac + Fbc + Habc).$$

The relevant original slacks are

$$\begin{aligned} &\delta_{ab}h_5 = 1 - d + c(d-e), \quad \delta_{cd}h_5 = ab, \\ &\delta_{ad}h_5 = 1 - b + bc, \qquad\quad \delta_{bd}h_5 = 1 - a + ac, \\ &\delta_{ac}h_5 = e(1-b) + bd. \end{aligned}$$

The cd zeros at $(a, b) = (0, 0), (1, 0), (0, 1)$ give $C = E = F = 0$. The ab zero with $c = 0, d = 1$ gives $D = 0$, and the zero with $c = d = e = 1$ then gives $H = 0$. The ad and bd zeros give $A = B = 0$. For the complement, the contexts where $\delta_{ad}h_5 = \delta_{bd}h_5 = 1$ give $A = B = E = F = 0$ and $D + H = 0$. Taking $b = d = 1$ in $\delta_{ac}h_5$ gives $H = 0$, and taking $a = b = 1$ in $\delta_{cd}h_5$ gives $C = 0$. Both tight kernels are therefore lines.

The larger cases inherit integrality and Boolean slacks from Q_4. For h_4 these properties follow directly from its coefficients. For h_5, besides the displayed slacks and their images under $a \leftrightarrow b$, the remaining ones are

$$\delta_{ae}h_5 = c(1-b), \qquad \delta_{be}h_5 = c(1-a), \qquad \delta_{ce}h_5 = a + b - ab.$$

All edge slacks in these cases are nonconstant and Boolean, so both summands have the entire graph as their direction graph. These give dimensions three through seven; larger dimensions follow from Theorem 9.8 and Corollary 9.10. □

9.4. **The least dilation in the K_5-free case**

Corollary 9.11 completes Theorem 9.5; every permitted part-size pattern occurs, with both complementary summands indecomposable. The same classification separates dilation one from dilation two in all dimensions.

Corollary 9.12. *For every $n \geq 4$, the smallest real dilation factor for which some K_5-free graphical zonotope on n vertices has a full-dimensional lattice indecomposable summand is one when $4 \leq n \leq 8$, and two when $n \geq 9$.*

Proof. If such a summand satisfied $Q + R = \lambda Z_G$ with $\lambda < 2$, all its elementary differences would be integers in $[0, \lambda]$, hence zero or one. Then $z_G - z_Q$ would be supermodular, making Q an actual summand of Z_G. Theorem 9.5 rules this out for $n \geq 9$. The path-cube family supplies equality with dilation two. For $4 \leq n \leq 8$, Corollary 9.11 supplies dilation one. A smaller factor is impossible, since a non-point lattice generalized permutahedron has a positive integer elementary difference. □

Appendix A. A graphic model of the path family

The two types of monomials in z_n count free intervals in a path of linear constraints, realizing Q_n as a projection of a graphic matroid base polytope over every field. The same graph gives a tridiagonal determinant with Newton polytope Q_n, whose indecomposability then controls factorization.

Proposition A.1 (A graphic 2-polymatroid model)**.** *Put $m = n - 2$. Over any field K, let $f_1, \dots, f_m$ be the standard basis of K^m, with $f_i = 0$ outside $[m]$, and set*

$$V_j = \operatorname{span}_K\{f_{j-2},\ f_{j-1} - f_j\} \qquad (1 \le j \le n).$$

Then, for every $T \subseteq [n]$,

$$r(T) := \dim_K \sum_{j \in T} V_j = m - z_n([n] \setminus T).$$

In particular, the same integer rank function is represented over every field, $r(\{j\}) \le 2$, and

$$Q_n = \{x \in \mathbb{R}^n_{\ge 0} : x([n]) = m,\ x(T) \le r(T) \text{ for all } T \subseteq [n]\}.$$

More precisely, the representation is a grouping of the columns of a reduced incidence matrix of a graph. Thus Q_n is the image of a spanning-tree polytope under summing edge coordinates within the groups.

Proof. Write $S = [n] \setminus T$, and identify the dual of K^m with coordinates $a_1, \dots, a_m$, extended by zero outside $[m]$. The annihilator of $\sum_{j \in T} V_j$ is defined by

$$a_{j-2} = 0, \qquad a_{j-1} = a_j \qquad (j \in T).$$

The equality constraints join coordinates along a path with fixed zero endpoints $a_0 = a_{m+1} = 0$. Each component disjoint from these endpoints and from all marked zero coordinates contributes one free parameter.

Such a free component is an interval $[p, q] \subseteq [m]$. Its boundary constraints require $p, q+1 \in S$, its internal edges require $p+1, \dots, q \notin S$, and the absence of marked zeros requires $p+2, \dots, q+2 \in S$. If $q \ge p+2$, the last two requirements conflict, so a free component has one or two coordinates. A singleton i contributes precisely when $\{i, i+1, i+2\} \subseteq S$, and a pair $i, i+1$ contributes precisely when $\{i, i+2, i+3\} \subseteq S$ and $i+1 \notin S$. Writing $s_i = \mathbf{1}_{i \in S}$, the annihilator dimension is therefore

$$\sum_{i=1}^{m} s_i s_{i+1} s_{i+2} + \sum_{i=1}^{m-1} s_i s_{i+2} s_{i+3}(1 - s_{i+1}) = z_n(S).$$

This proves the rank formula in every characteristic. The upper rank inequalities and the defining lower inequalities of Q_n are equivalent after imposing $x([n]) = m$. Nonnegativity follows also from $z_n(\{j\}) = 0$.

For the graphic description, take vertices $0, 1, \dots, m$, all spokes $0i$, all path edges $i(i+1)$, and one additional parallel copy of each endpoint spoke 01 and $0m$. Delete the row for vertex 0 from an oriented incidence matrix. Its nonzero columns, grouped as $(f_{j-2}, f_{j-1} - f_j)$, are exactly the vectors above. There are $2m+1$ edges, in n groups of at most two. This is a totally unimodular incidence representation; its matroid is graphic and regular. Grouping edge coordinates in the matroid base polytope gives the rank function r. Greedy basis extension along any ordering of the groups realizes every greedy vertex of the polymatroid base, proving the asserted equality of polytopes. □

Proposition A.2 (Stability and absolute irreducibility)**.** *Let $D_m(t_1, \dots, t_{m+2})$ be the determinant of the m by m tridiagonal matrix whose ith diagonal entry is $t_i + t_{i+1} + t_{i+2}$ and whose $(i, i+1)$ and $(i+1, i)$ entries are $-t_{i+1}$. Over $\mathbb{R}$, its Newton polytope is Q_n, and it is a homogeneous real stable polynomial of degree m. Moreover, for every field K, the polynomial D_m is absolutely irreducible over K.*

Proof. Let M be the incidence matrix just constructed. Give every edge in group j weight t_j. The displayed tridiagonal matrix is

$$L(t) = \sum_{j=1}^{n} t_j \big(f_{j-2} f_{j-2}^{\mathsf{T}} + (f_{j-1} - f_j)(f_{j-1} - f_j)^{\mathsf{T}} \big).$$

Cauchy–Binet expresses its determinant as

$$D_m(t) = \sum_{F \text{ a spanning tree}} \prod_{e \in F} t_{\mathrm{group}(e)},$$

since the nonzero maximal incidence minors are ± 1 and correspond exactly to spanning trees. Over $\mathbb{R}$, the coefficients count trees and are positive on their support. Proposition A.1 therefore gives $\mathrm{Newt}(D_m) = Q_n$.

If every t_j has positive imaginary part, then for every nonzero complex vector u,

$$u^* \operatorname{Im} L(t) u = \sum_j \operatorname{Im}(t_j) \big(|f_{j-2}^{\mathsf{T}} u|^2 + |(f_{j-1} - f_j)^{\mathsf{T}} u|^2 \big) > 0,$$

because the columns span $\mathbb{R}^m$. Hence $L(t)$ is nonsingular, proving real stability. Homogeneity and degree are immediate from the determinant or the spanning-tree formula.

Absolute irreducibility has a short characteristic-independent proof. Set $D_0 = 1$ and $D_1 = t_1 + t_2 + t_3$. Expansion along the first row gives, for $m \geq 2$,

$$D_m = t_1 A + B, \qquad A = D_{m-1}(t_2, \ldots, t_{m+2}), \qquad B = (t_2 + t_3)A - t_2^2 C,$$

where $C = D_{m-2}(t_3, \ldots, t_{m+2})$. Over any field, induction makes A irreducible. It is not associate to t_2; for $m = 2$ this follows from its displayed linear form, and for $m > 2$ it follows from its degree. Also $\deg C < \deg A$. Thus $\gcd(A, B) = \gcd(A, t_2^2 C) = 1$. The polynomial $t_1 A + B$ is primitive and linear in t_1, so Gauss's lemma proves it irreducible. Applying the same argument over an algebraic closure proves absolute irreducibility. □

The Newton-polytope assertion is made in characteristic zero. Coefficients can disappear after reduction. For example,

$$D_2 = (t_1 + t_2 + t_3)(t_2 + t_3 + t_4) - t_2^2$$

has coefficient two on $t_2 t_3$. The subspace representation, its rank function, and the absolute irreducibility assertion remain valid over every field.

Corollary A.3 (Every polynomial with this Newton polytope)**.** *Over any field K, every polynomial whose Newton polytope is exactly Q_n is absolutely irreducible.*

Proof. Work over an algebraic closure and suppose $F = AB$ with both factors nonzero. The Newton polytope of a product is the Minkowski sum of the Newton polytopes; maximizing any linear weight reduces this assertion to the nonvanishing product of the two initial forms. Indecomposability of Q_n therefore makes the two factor polytopes translates of λQ_n and $(1 - \lambda)Q_n$, with $0 \leq \lambda \leq 1$.

The displayed slack-one contexts in the main construction give a primitive lattice edge of Q_n. If $\lambda > 0$, the corresponding edge of the lattice polytope $\lambda Q_n + t$ has vector $\lambda(e_i - e_j)$, forcing λ to be an integer. Hence one factor has point Newton polytope and is a monomial. The minimum of each coordinate on Q_n is $z_n(\{i\}) = 0$, so no variable divides F. That monomial factor is constant, proving the assertion. □

Appendix B. Optimal lattice dilation on path joins

The path-cube family attains a uniform dilation factor two, whereas for the path joins $K_2 * P_m$ the least factor grows as $\lfloor m/2 \rfloor$. An endpoint correction to the complementary summand attains the lower bound of Corollary 9.4.

Let $G = K_2 * P_m$, with hubs a, b and path vertices $1, \dots, m$, where $m \geq 2$. For the path edge $i(i+1)$, put $h_i = a$ when i is odd and $h_i = b$ when i is even. Define

$$g(S) = \sum_{i=1}^{m} s_a s_b s_i + \sum_{i=1}^{m-1} s_{h_i} s_i s_{i+1} - \sum_{i=1}^{m-1} s_a s_b s_i s_{i+1}, \qquad s_v = \mathbf{1}_{v \in S}.$$

Put $c = \lfloor m/2 \rfloor$ and define

$$f(S) = z_G(S) - g(S) + (c-1)s_a s_b - s_b s_1 + s_a s_b s_1.$$

The three correction terms change only coefficients on $\{a, b, 1\}$. The hub difference of the uncorrected complement counts selected path blocks with a negative sign. Removing the contribution of the first block makes its largest possible count $\lfloor m/2 \rfloor$, exactly the dilation forced by the lower bound. The endpoint terms make this change while keeping all nonhub differences Boolean. We have $f(V(G)) = c + 2m - 1$. Write $\widehat{Q}_m = B(f)$.

Theorem B.1. *The polytope $\widehat{Q}_m$ is integral, indecomposable, and has dimension $m+1$. Its actual direction graph is $K_2 * P_m$, and*

$$\widehat{Q}_m + B(cz_G - f) = cZ_G.$$

Thus it attains the optimal lattice dilation in Corollary 9.4.

B.1. **The elemental differences**

The hub difference of g counts the nonempty selected path blocks. The endpoint correction subtracts the block containing 1, if present. Thus, with $x_i = \mathbf{1}_{i \in S}$,

$$\delta_{ab} f(S) = c - \sum_{i=2}^{m} x_i + \sum_{i=1}^{m-1} x_i x_{i+1} = c - \sum_{i=2}^{m} (1 - x_{i-1}) x_i.$$

The sum counts 01 transitions in a binary word. Their pairs of positions are disjoint, so there are at most c; both zero and c occur. Hence the difference lies in $[0, c]$.

The nonhub differences of g are Boolean. On a path edge they are $s_{h_i}(1 - s_{\bar{h}_i})$, where $\bar{h}_i$ is the other hub. For a path vertex v, let $N_a(v)$ and $N_b(v)$ be its path neighbors on edges labelled respectively a and b. Alternation makes both sets have size at most one, and

$$\delta_{av} g = \begin{cases} \sum_{u \in N_a(v)} x_u, & b \notin S, \\ 1 - \sum_{u \in N_b(v)} x_u, & b \in S. \end{cases}$$

The formula for bv interchanges the hubs. Thus every nonhub edge other than $a1, b1$ has $\delta_e f = 1 - \delta_e g \in \{0, 1\}$. The remaining two differences are

$$\delta_{a1} f(S) = \begin{cases} 1 - x_2, & b \notin S, \\ 1, & b \in S, \end{cases} \qquad \delta_{b1} f(S) = \mathbf{1}_{a \in S} x_2.$$

Nonedge differences vanish by clique support. All coefficients are integral, and Lemma 2.1 gives the asserted integral complement.

B.2. **One scale along the path**

For a vector satisfying the tight equations of f, write its coefficients as

$$A = y_{ab}, \quad B_i = y_{ai}, \quad C_i = y_{bi}, \quad E_i = y_{abi}, \quad D_i = y_{i(i+1)}, \quad H_i = y_{abi(i+1)}.$$

On path edge i, denote by N_i the triangle coordinate that is negative in the construction, and by Z_i the other one. Thus $(N_i, Z_i) = (y_{ai(i+1)}, y_{bi(i+1)})$ for odd i, with the order reversed for even i.

The tight left-end rows give

$$C_1 = E_1 = Z_1 = 0, \qquad B_1 = D_1 = -N_1.$$

Explicitly, use the $b1$ rows with contexts $\varnothing, \{a\}, \{2\}$, the $a1$ row with context $\{2\}$, and the 12 row with context $\{a\}$. These start the propagation without any assumption about an edge to the left of the path. For $2 \le i \le m$, let U_i be the spoke coordinate at the hub of N_{i-1} and V_i the other spoke coordinate. Tight spoke rows give

$$B_i + E_i = C_i + E_i = 0, \qquad U_i + N_{i-1} = 0, \qquad U_i + E_i + N_{i-1} + H_{i-1} = 0.$$

The first two use just the other hub as context. The next two are the rows on the U_i spoke conditioned on $i - 1$, first without and then with the other hub. For $2 \le i < m$, the remaining needed rows are

$$V_i + N_i = 0, \qquad U_i + N_{i-1} + Z_i = 0, \qquad D_i + N_i = 0.$$

These are respectively the V_i spoke row conditioned on $i + 1$, the U_i spoke row conditioned on both adjacent path vertices, and the path-edge row conditioned on h_i. The complement and endpoint formulas show that all these rows are tight at f.

Put $B_i = C_i = -E_i = t_i$ for $i \ge 2$. The left and right triangle equations give $N_{i-1} = -t_i = N_i$ where both occur. Thus all t_i equal one scalar t, and

$$B_1 = D_i = H_i = t, \qquad N_i = -t, \qquad Z_i = 0.$$

The hub-edge equality conditioned on $\{2, 4, \ldots, 2c\}$ then gives $A = ct$. All coefficients are consequently t times their prescribed values. The elimination has unit pivots and one free coordinate, also for $m = 2$, when no internal rows occur. Since $D_1 = t \geq 0$, the tight equations cut out precisely the ray of f. Lemma 2.1 proves indecomposability.

Every edge except $b1$ has positive empty-context difference; for $b1$ use $S = \{a, 2\}$. Hence the direction graph is G and the dimension is $m + 1$. The graph is a planar 3-tree; starting from $\{a, b, 1, 2\}$, successively attach $i + 1$ to $\{a, b, i\}$.

References

[1] G. Poullot, *Deformation cones of graphical zonotopes for K_4-free graphs*, arXiv 2404.02669v5, 23 October 2025.

[2] A. Padrol, V. Pilaud, and G. Poullot, *Deformed graphical zonotopes*, Discrete & Computational Geometry **73** (2025), 447–465; arXiv 2111.12422. DOI 10.1007/s00454-023-00586-x.

[3] A. Padrol and G. Poullot, *The graph of implicit edge dependencies for indecomposability and beyond*, arXiv 2512.05307v5, 26 May 2026.