FLAG γ-VECTORS THAT ARE NOT f-VECTORS, TOTAL NONNEGATIVITY FOR BOSE–BURTON GEOMETRIES, AND CHOW POLYNOMIALS OF PROJECTIVE DELETIONS

KOREA SUPERINTELLIGENCE LABS

Abstract. For a matroid M and a building set $\mathcal{G}$ of its lattice of flats, the Chow polynomial $\mathcal{H}(M,\mathcal{G})(t)$ is palindromic, and Coron, Ferroni and Li proved that its γ-vector is nonnegative when $(M,\mathcal{G})$ is flag, and is the f-vector of a simplicial complex when $(M,\mathcal{G})$ is complete. They ask whether the γ-vector of a flag built matroid is the f-vector of a flag simplicial complex, and name "satisfies the Kruskal–Katona inequalities" as an intermediate step. We answer both negatively: a flag, non-complete building set of 19 flats on the uniform matroid $U(7,8)$ has $\mathcal{H} = 1{+}11t{+}37t^2{+}55t^3{+}37t^4{+}11t^5{+}t^6$ and $\gamma = (1,5,2,1)$, and no simplicial complex has two edges and one triangle. Coloops give such examples in ranks 7 through 11, and lengthening four lines gives twelve of rank 7 with unbounded ground set, eleven non-uniform. We then locate the failure: in the k-block family on $U(2k{-}1,2k)$, of which the counterexample is the $k=4$ member, the top Kruskal–Katona inequality fails at exactly one pair (k,m), and the flag member with fewest edges is tight exactly for $5 \le k \le 8$, because Mantel's quadratic bound crosses the two linear Kruskal–Katona thresholds between $k=4$ and $k=9$.

In the second half we take the maximal building set and turn to real-rootedness. We prove that the $(r+1)\times(r+1)$ matrix of Gaussian binomials whose bottom row is replaced by the Whitney numbers of the Bose–Burton geometry $B_{r,k}(q)$ — the projective space $PG(r-1,q)$ with the points of a k-dimensional subspace deleted — is totally nonnegative, for every r, every $0 \le k \le r-1$ and every real $q > 1$. The proof exhibits the forced resolving array of Brändén and Saud-Maia-Leite in closed form, with manifestly positive coefficients. Equivalently, the dual of the lattice of flats of every Bose–Burton geometry is a TN-poset, confirming for this family their conjecture that the dual of any upper combinatorially uniform geometric lattice is a TN-poset; and for $2 \le k \le r-2$ these lattices are neither perfect matroid designs, nor supersolvable, nor paving, so they lie outside the classes named alongside that conjecture. With a theorem of Brändén and Vecchi it follows that the Chow polynomial and the augmented Chow polynomial of $B_{r,k}(q)$ are real-rooted — instances of a conjecture of Ferroni–Schröter and Huh–Stevens which, for $2 \le k \le r-2$, no published sufficient condition we could find reaches, apart from the Chow polynomials in rank at most six.

The third part perturbs the geometry. Cheng and Liu have recently refuted unimodality of matroid Kazhdan–Lusztig polynomials by deleting from a Bose–Burton geometry a further point set S that is line-separated from W. We show that the chain sum defining the Chow polynomial never reads a flat of rank 1, so that the Chow polynomial of a projective deletion depends on the deleted set only through the flat spans of dimension at least 2; that the Cheng–Liu perturbation is therefore invisible to the Chow polynomial as soon as every projective line outside W retains two surviving points; and that over a field with at least three elements a cap has that property, while over $\mathbb{F}_2$ it fails as soon as $W \ne 0$ and one further point is deleted, or two points of a chart are. Consequently, for every prime power $q \ge 3$ there is an explicit matroid whose Kazhdan–Lusztig polynomial is not unimodal and whose Chow polynomial is that of a Bose–Burton geometry, hence real-rooted: the Cheng–Liu refutation does not transfer. On the other side of that dichotomy we compute a census of 26 rank-7 deletions of $PG(6,2)$ from their lattices of up to 29 211 flats; the 21 whose lattice of flats is not upper rank uniform lie outside every published sufficient condition we could find, and all 21 Chow polynomials are real-rooted.

The coefficient data is computed from the Feichtner–Yuzvinsky formula, from the second part on inside the proof assistant itself; every arithmetic assertion is machine-checked in Lean 4, and the transfer theorems of the third part are formalised there over an arbitrary field.

1. Introduction

1.1. **Chow polynomials of built matroids.** Let M be a matroid on a finite ground set E, with rank function rk and lattice of flats $\mathcal{L}(M)$, with bottom and top $\hat{0} = \mathrm{cl}(\emptyset)$ and $\hat{1} = E$. A *building set* $\mathcal{G} \subseteq \mathcal{L}(M) \setminus \{\hat{0}\}$ is a family of flats for which the join map factors the interval structure of $\mathcal{L}(M)$ into products; we use throughout the criterion of [12, Prop. 2.2], after Backman and Danner: $\mathcal{G}$ is a building set if and only if $\mathcal{G}_{\min} \subseteq \mathcal{G}$ and $G \wedge G' \neq \hat{0}$ implies $G \vee G' \in \mathcal{G}$, where $\mathcal{G}_{\min}$ is the minimal building set. The pair $(M, \mathcal{G})$ is a *built matroid*; the extremes are $\mathcal{G}_{\max} = \mathcal{L}(M) \setminus \{\hat{0}\}$ and $\mathcal{G}_{\min}$.

Feichtner and Yuzvinsky [14] attach to $(M, \mathcal{G})$ a graded ring $CH(M, \mathcal{G})$ — for $\mathcal{G} = \mathcal{G}_{\max}$ the Chow ring of [1] — whose Hilbert–Poincaré series $\mathcal{H}(M, \mathcal{G})(t) = \sum_i \dim CH^i(M, \mathcal{G})\, t^i$ is called the *Chow polynomial* of $(M, \mathcal{G})$ [15, 12]. It is palindromic, by Poincaré duality for $CH(M, \mathcal{G})$ [12, §1.1], and when $E \in \mathcal{G}$ it has degree $\mathrm{rk}\, M - 1$, so it has a γ-vector: with $d = \mathrm{rk}\, M - 1$, the polynomials $t^i(1+t)^{d-2i}$, $0 \leq i \leq \lfloor d/2 \rfloor$, are triangular in the monomial basis, so there is a unique $\gamma = (\gamma_0, \ldots, \gamma_{\lfloor d/2 \rfloor})$ with

$$\mathcal{H}(M, \mathcal{G})(t) \;=\; \sum_{i=0}^{\lfloor d/2 \rfloor} \gamma_i\, t^i (1+t)^{d-2i}, \qquad \gamma_0 = 1. \tag{1}$$

1.2. **The question.** Gal's conjecture [18] predicts that the γ-vector of a flag simplicial sphere is nonnegative, and Nevo and Petersen [25] conjectured the sharper statement that it is the f-vector of a simplicial complex — equivalently, by the Kruskal–Katona theorem, that it satisfies the numerical inequalities recalled in §2. Coron, Ferroni and Li [12] transport both to built matroids. Two of their results and one of their questions are what this paper is about. All numbered references to [12] below are to its version 3 of 10 July 2026; the numbering of the earlier versions differs.

- (i) If $(M, \mathcal{G})$ is *complete*, then $\gamma \geq 0$ [12, Thm. 1.6], and moreover γ is the f-vector of a simplicial complex [12, Thm. 1.7].
- (ii) If $(M, \mathcal{G})$ is *flag*, then $\gamma \geq 0$ [12, Thm. 1.8].
- (iii) [12, Question 9.8]: *Let $(M, \mathcal{G})$ be a flag built lattice. Is the γ-vector of $\mathcal{H}(M, \mathcal{G})(t)$ the f-vector of a flag simplicial complex?* They add that when M is Boolean — the setting of flag nestohedra — the answer is affirmative by a construction of Aisbett [2], that in their manuscript they prove "only the nonnegativity of the γ-vector", and that "a natural intermediate step would therefore be to show that these γ-vectors satisfy the Kruskal–Katona inequalities".

The completeness hypothesis in (i) is not vacuous and not implied by flagness: $(M, \mathcal{G}_{\max})$ is complete for every M, but flag built matroids need not be complete, and it is exactly on those that (iii) asks a new question.

1.3. **Results.** We answer (iii) negatively, in the strong form: the γ-vector of a flag built matroid need not be the f-vector of *any* simplicial complex, so the Kruskal–Katona intermediate step fails as well.

Theorem A (Theorem 3.2 and Corollary 3.4). Split the ground set of $M = U(7, 8)$ by a perfect matching into four pairs A, B, C, D and let $\mathcal{G}$ consist of the eight atoms, the four pairs, the two opposite unions $A \cup C$ and $B \cup D$, the four unions of three pairs, and E — nineteen flats. Then $(M, \mathcal{G})$ is a flag built matroid, it is not complete, $\mathcal{H}(M, \mathcal{G})(t) = 1 + 11t + 37t^2 + 55t^3 + 37t^4 + 11t^5 + t^6$ and $\gamma = (1, 5, 2, 1)$, which is the f-vector of no simplicial complex, a complex with one triangle having at least three edges. Adjoining j coloops with a chain of flats above E gives flag built matroids of rank $7 + j$ with the same γ-vector, $j \leq 4$.

The example is the smallest member of a rank-7 family whose ground set is unbounded (Proposition 3.7): replacing the four pairs by four lines $U(2, s_i)$ and truncating to rank 7 leaves $\mathcal{H}$ and γ unchanged in the twelve cases computed, eleven of them non-uniform. The failure is thus neither an artefact of uniform matroids nor of a small ground set.

The second result says where, in the natural two-parameter family containing the counterexample, such a failure can occur — and that the answer is: only there. For $k \geq 3$ let $M_k = U(2k-1, 2k)$ with its ground set partitioned into k blocks $B_1, \ldots, B_k$ of size 2, let Γ be a graph on $[k]$ with $m = |\Gamma|$ edges, and let

$$\begin{aligned} \mathcal{G}(k, \Gamma) = \{\text{the } 2k \text{ atoms}\} \cup \{B_1, \ldots, B_k\} \cup \{B_i \vee B_j : ij \in \Gamma\} \\ \cup \{\text{all joins of} \geq 3 \text{ blocks}\} \cup \{E\}. \end{aligned} \tag{2}$$

The counterexample of Theorem A is the member $k = 4$, Γ a perfect matching, $m = 2$.

Theorem B (Theorem 5.5). Let $k \geq 4$ and let Γ be a graph on $[k]$ such that $(M_k, \mathcal{G}(k, \Gamma))$ is flag, and assume the γ-vector of $\mathcal{H}(M_k, \mathcal{G}(k, \Gamma))$ has the shape $(1, *, \ldots, *, m, 1)$ of length k. Then the top Kruskal–Katona inequality $\gamma_{k-1} \leq \gamma_{k-2}^{(k-2)}$ holds if and only if $(k, m) \neq (4, 2)$; and for the member with the fewest edges it is tight, with slack exactly 0, if and only if $5 \leq k \leq 8$.

The mechanism is a collision of two bounds of different degrees. Flagness of $(M_k, \mathcal{G}(k, \Gamma))$ is equivalent to the complement of Γ being triangle-free (Lemma 5.1), so Mantel's theorem forces $m \geq \lfloor (k-1)^2/4 \rfloor$, a *quadratic* lower bound; while the arithmetic of the Kruskal–Katona pseudopower makes the top inequality hold exactly when $m \geq k-1$ and be tight exactly when moreover $m \leq 2k-4$, two *linear* thresholds (Theorem 4.1). The quadratic bound is below the first threshold only at $k = 4$, and below the second only up to $k = 8$.

1.4. **Real-rootedness, and a matrix positivity conjecture.** The second half of the paper is about a different property of the same polynomials, for the maximal building set, where $\mathcal{H}(M, \mathcal{G}_{\max})$ is the Chow polynomial $\underline{\mathcal{H}}_M$ of [15]. Ferroni and Schröter conjectured that $\underline{\mathcal{H}}_M$ has only real zeros for every matroid M [17, Conj. 8.18], and Huh and Stevens conjectured the same, together with the counterpart for the augmented Chow ring [27, Conjs. 4.1.3 and 4.3.3]; the first of these is [12, Conj. 1.1], and both are open. Two recent frameworks prove them for large classes. Brändén and Vecchi [8] call a poset P whose matrix of lower Whitney numbers $R(P)$ is totally nonnegative a *TN-poset*, and prove that the Chow polynomials of a TN-poset — or of the *dual* of one — are real-rooted [8, Thm. 6.1]; their list of examples of TN-posets includes perfect matroid designs and duals of Dowling lattices, and a separate argument reaches the lattices of flats of paving matroids [8, Cor. 7.2]. Coron, Ferroni and Li prove it for *UMEL-shellable* posets — rank-uniform posets carrying a monotonic edge-lexicographic labelling — [13, Thm. 1.2], and specialise that to uniform matroids, projective and affine geometries, braid matroids of types A and B, Dowling geometries, and supersolvable rank-uniform geometric lattices, together with the rank-selected subposets of these [13, Thm. 1.8]. Brändén and Saud-Maia-Leite [6, Conj. 4.3] conjecture that *the dual of any upper combinatorially uniform geometric lattice is a TN-poset*; the two classes they name in that neighbourhood are the perfect matroid designs and the Dowling lattices, and for the latter they prove that the dual is a TN-poset [6, Thm. 4.1].

We prove that conjecture for one infinite family, and the family is outside those named classes. For a prime power q, $r \geq 1$ and $0 \leq k \leq r-1$, the *Bose–Burton geometry* $B_{r,k}(q)$ is the simple rank-r matroid on the points of $PG(r-1,q)$ that do not lie in a fixed k-dimensional subspace W [4]; $k=0$ gives the projective geometry and $k=r-1$ the affine geometry $AG(r-1,q)$. Write $W_d = \begin{bmatrix} r \\ d \end{bmatrix}_q - \begin{bmatrix} k \\ d \end{bmatrix}_q$ for its Whitney numbers of the second kind, $W_0 = 1$, in terms of the Gaussian binomial coefficients.

Theorem C (Theorem 9.2 and Corollary 9.3). Let $q > 1$ be real, $r \geq 1$, $0 \leq k \leq r-1$, and put $s = r-k$. The $(r+1)\times(r+1)$ lower unitriangular matrix

$$M_{n,j} = \begin{bmatrix} n \\ j \end{bmatrix}_q \quad (0 \leq n \leq r-1), \qquad M_{r,j} = W_{r-j},$$

is resolvable in the sense of [5], with the positive resolving coefficients

$$\lambda_j = q^j \quad (j < s), \qquad \lambda_j = q^{\binom{s}{2} - \binom{j}{2} - j(r-1-j)} \Big(\prod_{i=s}^{j} (q^i - 1)\Big)\, \Sigma(r,s,j) \quad (j \geq s),$$

where $\Sigma(r,s,j) = \sum_{m+n=r-j-1} q^{jn} \begin{bmatrix} j-s+m \\ m \end{bmatrix}_q > 0$; and the resolving array is the only one. In particular M is totally nonnegative, so for every prime power q the dual of the lattice of flats of $B_{r,k}(q)$ is a TN-poset: [6, Conj. 4.3] holds for the Bose–Burton geometries.

For $2 \leq k \leq r-2$ — a range that is nonempty from $r=4$ on — these lattices are neither perfect matroid designs nor supersolvable, and $B_{r,k}(q)$ is not paving (Lemma 8.2 and Proposition 8.3), so they escape the hypotheses named above — with the caveats recorded in Remark 9.4. Feeding Theorem C into [8, Thm. 6.1] gives:

Theorem D (Corollary 10.1). For every prime power q, every $r \geq 1$ and every $0 \leq k \leq r-1$, the Chow polynomial and the augmented Chow polynomial of $B_{r,k}(q)$ are real-rooted. For $2 \leq k \leq r-2$ these are instances of [12, Conj. 1.1] and of its augmented companion for which we could find no published sufficient condition, apart from the Chow polynomials in rank at most 6, which [12, Cor. 8.5] covers; see Remark 9.4.

The formal development records 41 of those instances — $q \in \{2,3,4,5,7,9\}$, $r \leq 8$, 29 of them with $2 \leq k \leq r-2$ — with independent sign-alternation certificates, and with the γ-vectors and the Kruskal–Katona verdicts of §2 alongside (Proposition 10.2).

1.5. **Projective deletions, and a refutation that does not transfer.** The third part of the paper perturbs the Bose–Burton geometry. Cheng and Liu [10] have recently shown that the Kazhdan–Lusztig polynomial P_M of a matroid need not be unimodal, and their counterexamples are exactly Bose–Burton geometries with an extra point set removed: they delete from $PG(r-1,q)$ the set $D = P(W) \sqcup$

S, where $\dim W = k$ and S is *line-separated* from W, meaning that every projective line contained in $P(W) \cup S$ already lies in $P(W)$ [10, Def. 2.5]. All of the failure of unimodality comes from S: their Proposition 2.6 gives $P_M(t) = \sum_j \begin{bmatrix} k \\ j \end{bmatrix}_q t^j + s\,t$ with $s = |S|$, which at $s = 0$ is the row of Gaussian binomials and is unimodal. Since the Chow polynomial and the Kazhdan–Lusztig polynomial of a matroid are both conjectured to be real-rooted, and one of the two conjectures has now been refuted in its unimodality form, it is natural to ask what the same perturbation does to $\underline{\mathcal{H}}_M$. The answer is a dichotomy, and both of its sides are theorems.

Call a subspace $U \le V$ a *flat span* of the deletion $M = P(V) \setminus D$ if the surviving points lying in U span U; the flats of M are exactly the sets of surviving points of flat spans, of rank $\dim U$ [10, §2]. Theorem E holds over an arbitrary field K, finite or not; there $B_{r,k}(q)$ is to be read as the deletion $D = P(W)$, which for $K = \mathbb{F}_q$ is the Bose–Burton geometry of §8. Theorem F counts the points of a line and is a statement about a finite field of order q.

Theorem E (Theorems 11.1 and 11.2). Let K be a field, V a finite-dimensional K-vector space and $W \le V$.

(i) The Chow polynomial of $P(V) \setminus D$ depends on D only through the flat spans of dimension ≥ 2: if two deleted sets have the same flat spans in every dimension ≥ 2, their Chow polynomials are equal. No hypothesis is placed on either deletion.

(ii) Suppose $P(W) \subseteq D$, that every projective line not contained in W retains a surviving point, and that every projective line through a surviving point is spanned by its surviving points. Then a subspace of dimension ≥ 2 is a flat span of $P(V) \setminus D$ if and only if it is not contained in W — the same condition as for the Bose–Burton geometry — and consequently $\underline{\mathcal{H}}_{P(V)\setminus D} = \underline{\mathcal{H}}_{B_{r,k}(q)}$, identically.

Part (i) is the whole engine, and it is one line of the chain formula (6): a chain step of rank gap 1 carries the empty sum, so no flat of rank 1 occurs in any chain of nonzero weight. Part (ii) says that the deleted points can be moved around freely, so long as no projective line is left with fewer than two survivors. The next theorem says when that happens.

Theorem F (Theorem 11.4 and Corollary 11.5). Let K be finite of order q, let $V = W \oplus X$ and let $S \subseteq P(X)$, $D = P(W) \sqcup S$.

(i) If $q \ge 3$ and S is a *cap* — no projective line carries three of its points — then the hypotheses of Theorem E(ii) hold, so $\underline{\mathcal{H}}_M = \underline{\mathcal{H}}_{B_{r,k}(q)}$; and conversely, if some line not contained in W fails to retain two survivors, then S carries three collinear points.

(ii) If $q \ge 3$, $H < X$ is a hyperplane, S lies in the affine chart $P(X) \setminus P(H)$ and S contains no full affine line of that chart, the same conclusion holds, with the matching converse. For $q \ge 4$ this hypothesis is strictly weaker than (i).

(iii) If $q = 2$ the conclusion fails as soon as $k \ge 1$ and $s \ge 1$, or $k = 0$ and $s \ge 2$ with S inside a chart: the lattices of flat spans of M and of $B_{r,k}(2)$ already differ in rank 2.

Part (iii) is why the next theorem asks for $q \geq 3$, and why the binary counterexample that [10] display falls on the other side of the dichotomy. Their construction runs over every finite field; what part (iii) settles is its binary half.

Theorem G (Theorem 12.4). Let $q \geq 3$ be a prime power. Put $k = 6$ and $s = \begin{bmatrix} 6 \\ 2 \end{bmatrix}_q - \begin{bmatrix} 6 \\ 1 \end{bmatrix}_q + 1$, let $r = 25$, let $V = W \oplus X$ with $\dim W = 6$, and let S be a cap of size s inside an affine chart $AG(18, q)$ of $P(X) \cong PG(18, q)$. Then $M = PG(24, q) \setminus (P(W) \sqcup S)$ is a matroid whose Kazhdan–Lusztig polynomial is not unimodal and whose Chow polynomial is $\underline{\mathcal{H}}_M = \underline{\mathcal{H}}_{B_{25,6}(q)}$, hence is real-rooted by Theorem D. For $q = 3$ the same holds at rank 21.

So the refutation of [10] does not transfer: on their own construction, over every field with more than two elements, one can keep the Kazhdan–Lusztig failure and still land on a Chow polynomial that this paper proves real-rooted. What it costs is rank: Cheng and Liu need only s points in a chart, while a cap of size s needs a chart of roughly twice the dimension, so at $q = 3$ their minimal rank 16 becomes 21. We also record that $k = 6$ is forced: for $k \leq 5$ the polynomial of their Proposition 2.6 is unimodal for every s (Remark 12.5).

The other side of the dichotomy is where the interesting instances are, and it is not accessible in closed form. Once a line is left with a single survivor the lattice of flats can stop being upper rank uniform, and with it the hypotheses of every framework named in §1.4. We computed the first rank at which that situation is not already settled by other means.

Theorem H (Theorem 13.2). Of the 26 deletions $PG(6, 2) \setminus (P(W) \sqcup S)$ tabulated in §13, exactly 21 have a lattice of flats that is not upper rank uniform. All 21 have real-rooted Chow polynomials, each certified by a sign alternation on seven rational points.

Rank 7 is the first rank not covered by [12, Cor. 8.5], and for these 21 matroids neither the total-nonnegativity route of [8, 6] nor the UMEL route of [13] applies, since both require the lattice of flats to be upper rank uniform; nor are they uniform, perfect matroid designs or paving. Theorem E(i) is what makes this a claim about instances rather than about classes: it produces deletions whose Chow polynomial *does* coincide with a Bose–Burton value although their lattice is not upper rank uniform (Corollary 11.3), so failure of upper rank uniformity by itself certifies nothing, and each of the 21 polynomials had to be compared against the Bose–Burton values one by one.

1.6. **What is computed and what is proved.** Every Chow polynomial in the first part of the paper is a computed value, obtained from the Feichtner–Yuzvinsky nested-set sum (3) by programs written from the definitions; the load-bearing values were recomputed by independent implementations and calibrated against published ones (§3.6). So are the matroid-theoretic predicates: that a stated set of flats is a building set, is flag, and is not complete are exhaustive finite computations, described where they are used. What is machine-checked in Lean 4 (§14) is everything downstream of those values — the γ-expansions, the Kruskal–Katona and Macaulay verdicts, the arithmetic of the pseudopower, and the two threshold laws of Theorem B, which contain no data at all — together with the fact that a complex with one triangle has three edges. There is no formal model of built matroids behind these statements, and we claim none; the matroid content enters as data. Two further limitations are stated where they arise: the shape hypothesis of Theorem B is data for $3 \le k \le 12$ (Remark 5.4), and the coloop extension of Theorem A is verified for five ranks rather than by induction.

The second part is the other way round. Theorem C and its corollaries are ordinary mathematics, proved by hand for all r, k and q; what the formal development contributes there is verification, not proof — the closed-form resolving array is checked against the matrix for 381 matrices, total nonnegativity is checked from the definition, minor by minor, at the sizes where that is affordable, the Chow polynomials are *computed* in the kernel from the Whitney numbers rather than supplied as data, and the real-rootedness of 41 of them is certified a second time, elementarily, by sign alternations and by Kurtz's criterion. The deduction of Theorem D from Theorem C is by citation, of [8, Thm. 6.1], and is not formalised.

The third part is the one place where the proofs themselves are formal. Theorems E and F are proved inside the proof assistant over an arbitrary field — an arbitrary finite field for Theorem F, which counts the points of a line — and an arbitrary finite-dimensional vector space, with the chain sum written over an abstract finite ranked poset and an arbitrary step weight killing a rank gap of 1; nothing about them is data. What is data there is the calibration (§11.5), where the abstract chain sum is evaluated in the kernel on the computed lattices of $PG(2,2)$, $PG(2,3)$ and seven of their deletions and compared with the recursion (7). Theorem G is a proof by hand from Theorem E, Theorem F, Theorem D and one cited input — [10, Prop. 2.6, Cor. 1.3] — whose arithmetic (the parameters, both hypotheses of that proposition, the strict local valley, and the resolvability certificate of Theorem C at exactly the ranks used) is decided in the kernel for 17 prime powers. Theorem H and the census of §13 are exhaustive computation throughout: the lattice of flats of each of the 26 deletions is enumerated in full, and only the resulting coefficient vectors and their certificates enter the kernel. Three limitations are stated where they arise: Cheng and Liu's Proposition 2.6 and Brändén and Vecchi's Theorem 6.1 are cited and not reproved (Remark 12.8); the identification of the bitmask model used for the calibration with the subspace lattice of $\mathbb{F}_q^3$ is mathematics and not a formal theorem (§11.5); and the exclusion of the published frameworks in Theorem H leaves one loophole, recorded in Remark 13.3.

1.7. **Organisation.** §2 fixes notation; §3 proves Theorem A and reports the search behind it; §4 proves the pseudopower thresholds and §5 Theorem B; §6 discusses the one Kruskal–Katona inequality that survives in low rank; §7 records the calibration data. The second part is independent of the first: §8 sets up the Bose–Burton geometries and shows where they sit, §9 proves Theorem C, and §10 draws Theorem D and the finite checks from it. The third part rests on the second: §11 proves Theorems E and F and calibrates them, §12 proves Theorem G, and §13 reports the rank-7 census and proves Theorem H. §14 records the formal development.

2. Preliminaries

2.1. **Nested sets and the Chow polynomial.** Let $(M, \mathcal{G})$ be a built matroid. A subset $S \subseteq \mathcal{G}$ is *nested* if for every antichain $\{G_1, \dots, G_r\} \subseteq S$ with $r \ge 2$ the join $G_1 \vee \cdots \vee G_r$ does not lie in $\mathcal{G}$. We write $c\mathcal{N}(\mathcal{L}, \mathcal{G})$ for the collection of all nested subsets of $\mathcal{G}$, maximal elements of $\mathcal{G}$ included, and $\mathcal{N}(\mathcal{L}, \mathcal{G})$ for the subcomplex of those avoiding them. For $S \in c\mathcal{N}(\mathcal{L}, \mathcal{G})$ and $F \in S$ put $J_F^S = \bigvee\{G \in S : G < F\}$, with $J_F^S = \hat{0}$ when the set is empty. The Feichtner–Yuzvinsky formula reads

$$\mathcal{H}(M,\mathcal{G})(t) \;=\; \sum_{S\in c\mathcal{N}(\mathcal{L},\mathcal{G})} \prod_{F\in S} \bigl(t + t^2 + \cdots + t^{\operatorname{rk} F - \operatorname{rk} J_F^S - 1}\bigr), \tag{3}$$

the empty product being 1 and a factor being empty — so that the whole term vanishes — as soon as $\operatorname{rk} F - \operatorname{rk} J_F^S \le 1$.

Remark 2.1. The sum in (3) runs over $c\mathcal{N}$, not over $\mathcal{N}$. Restricting it to $\mathcal{N}$ — that is, dropping the maximal elements of $\mathcal{G}$ from the index set — returns $1+3t$ on $(B_3, \mathcal{G}_{\max})$, where the Chow polynomial is the Eulerian polynomial $A_3 = 1+4t+t^2$, and $1 + 10t + 4t^2$ on $(B_4, \mathcal{G}_{\max})$, where it is $A_4 = 1 + 11t + 11t^2 + t^3$. The error is self-detecting — the answer is no longer palindromic of degree $\operatorname{rk} M - 1$ — and the $c\mathcal{N}$ reading is the only one reproducing the published values of §3.6.

Two conditions on $(M, \mathcal{G})$ appear throughout. It is *flag* when $\mathcal{N}(\mathcal{L}, \mathcal{G})$ is a flag simplicial complex [12, Def. 3.13]; by [12, Lem. 3.15] this holds if and only if for every antichain $\{G_1, \dots, G_r\} \subseteq \mathcal{G}$ with $r \ge 2$ and $G_1 \vee \cdots \vee G_r \in \mathcal{G}$ there are $i \ne j$ with $G_i \vee G_j \in \mathcal{G}$. Completeness is [12, Def. 3.1], which we restate because the counterexample below turns on it.

Definition 2.2 ([12, Def. 3.1])**.** Let $(M, \mathcal{G})$ be a loopless built matroid on E and let $\lhd$ be a total order on E. For $F \in \mathcal{L}(M)$ and $G \in \mathcal{G}$ with $F \le G$, write $G \setminus F = \{e_1 \lhd e_2 \lhd \cdots \lhd e_k\}$ and $F_i = \operatorname{cl}_M(F \cup \{e_1, \dots, e_i\})$; the $\lhd$-*chain from F to G* is the saturated chain $\operatorname{ch}_\lhd(F, G) = \{F \le F_1 \le \cdots \le F_k = G\}$. Then $(M, \mathcal{G})$ is $\lhd$-*complete* if $\operatorname{ch}_\lhd(F, G) \subseteq \mathcal{G}_F^G \cup \{F\}$ for every $F \in \mathcal{L}(M)$ and every $G \in \mathcal{G}$ with $F \le G$, and *complete* if it is $\lhd$-complete for some total order $\lhd$ on E. Here $\mathcal{G}_F^G$ is the building set that $\mathcal{G}$ induces on the interval $[F, G]$; at $F = \hat{0}$ it is $\mathcal{G} \cap [\hat{0}, G]$.

Only the case $F = \hat{0}$ is used below, in both directions. Failure of the displayed inclusion at $F = \hat{0}$ for every order certifies non-completeness, being a special case of the definition; conversely [12, Rem. 3.5] shows that $\operatorname{ch}_\lhd(\hat{0}, G) \subseteq \mathcal{G} \cup \{\hat{0}\}$ for all $G \in \mathcal{G}$ already implies $\lhd$-completeness. Every $(M, \mathcal{G}_{\max})$ is complete, and completeness neither implies nor is implied by flagness [12, Ex. 3.14].

For the built matroids below, all of which contain every atom of $\mathcal{L}(M)$ in $\mathcal{G}$, the linear coefficient of $\mathcal{H}$ satisfies $h_1 = |\mathcal{G}| - |E|$, the dimension of CH^1 read off the Feichtner–Yuzvinsky presentation; we use this only as a consistency check on (3).

2.2. **The Kruskal–Katona and Macaulay conditions.** For integers $m \ge 0$ and $i \ge 1$ there is a unique *i-cascade representation*

$$m = \binom{a_i}{i} + \binom{a_{i-1}}{i-1} + \cdots + \binom{a_j}{j}, \qquad a_i > a_{i-1} > \cdots > a_j \ge j \ge 1,$$

and one sets $m^{(i)} = \binom{a_i}{i+1} + \binom{a_{i-1}}{i} + \cdots + \binom{a_j}{j+1}$ and $m^{\langle i\rangle} = \binom{a_i+1}{i+1} + \cdots + \binom{a_j+1}{j+1}$. The Kruskal–Katona theorem [22, 20] says that a vector $(f_0, \dots, f_D)$ of nonnegative integers with $f_0 = 1$ is the f-vector of a simplicial complex — f_i counting the faces with exactly i vertices, the empty face included — if and only if

$$f_{i+1} \le f_i^{(i)} \qquad (1 \le i \le D - 1); \tag{4}$$

Macaulay's theorem [26] says it is the Hilbert function of a standard graded algebra — an M-vector, or O-sequence — if and only if the same inequalities hold with $f_i^{\langle i\rangle}$ on the right. We write $\mathrm{KK}(f)$, $\mathrm{Mac}(f)$ for the two conditions, call (4) with

$i = D-1$ the *top* inequality and $f_i^{(i)} - f_{i+1}$ its *slack*. As Mac is the weaker condition, a vector with Mac true and KK false is an M-vector that is not an f-vector.

Remark 2.3. A γ-vector of a palindromic polynomial of degree d has $\lfloor d/2 \rfloor + 1$ entries, so (4) is an empty condition unless $\lfloor d/2 \rfloor \geq 2$. With $d = \operatorname{rk} M - 1$ this means: no inequality at all below rank 5; one inequality, $\gamma_2 \leq \binom{\gamma_1}{2}$, in ranks 5 and 6; and the second inequality $\gamma_3 \leq \gamma_2^{(2)}$ only from rank 7 on. The second is the sharp one, because $m^{(2)} = 0$ for $m \leq 2$ and $m^{(2)} = 1$ for $m \in \{3, 4\}$.

Proposition 2.4. *On a two-entry vector the Kruskal–Katona condition is exactly nonnegativity:* $\mathrm{KK}(1, a)$ *holds if and only if* $a \geq 0$. *Consequently, for a built matroid of rank at most* 4 *with* $E \in \mathcal{G}$*, the conclusion of* [12, Thm. 1.7] *follows from* γ*-positivity alone, and no search over such built matroids can bear on* [12, Question 9.8].

Proof. The conditions (4) range over $1 \leq i \leq D-1$ and $D = 1$ here, so there are none, and what remains of KK is $f_0 = 1$ and nonnegativity. A built matroid of rank $r \leq 4$ with $E \in \mathcal{G}$ has $d = r - 1 \leq 3$, hence at most two γ-entries, and γ-positivity is [12, Thms. 1.6, 1.8]. $\square$

3. A flag built matroid whose γ-vector is not an f-vector

3.1. **Where to look.** Proposition 2.4 says the target has rank at least 5, and Remark 2.3 that the sharp inequality appears only from rank 7. Two further constraints narrow the search. First, M must not be Boolean, by Aisbett's affirmative answer for flag nestohedra [2], so the smallest available ground set in rank 7 has $|E| = 8$: the only rank-7 matroid on 7 elements is B_7. Second, $\gamma_1 = h_1 - d = |\mathcal{G}| - |E| - d$, so the ceilings $\binom{\gamma_1}{2}$ and $\gamma_2^{(2)}$ grow with $|\mathcal{G}|$ and the tight regime consists of the *smallest* flag building sets. The shape of a refutation was fixed before any example was found: $\mathrm{KK}(1, \gamma_1, 2, 1)$ is false for every γ_1, because $2^{(2)} = 0$.

Remark 3.1. An exhaustive enumeration covers all building sets of $U(3,4)$ and $U(3,5)$: every subset of $\mathcal{L}(M) \setminus \{\hat{0}\}$ containing $\mathcal{G}_{\min}$ passes the criterion of [12, Prop. 2.2] here, so there are $64+1024 = 1088$ built matroids, of which $41+388 = 429$ are flag; exactly two fail γ-positivity, namely the two minimal building sets, with $\gamma = (1, -1)$, and both are non-flag as [12, Thm. 1.8] requires. The enumeration lies outside the formal development and, by Proposition 2.4, could not have borne on [12, Question 9.8]: everything in that range has rank 3.

3.2. **The counterexample.**

Theorem 3.2. *Let* $M = U(7,8)$ *on* $E = \{0, \ldots, 7\}$*, let*

$$A = \{0,5\}, \quad B = \{1,3\}, \quad C = \{2,6\}, \quad D = \{4,7\}$$

be a perfect matching of E *into rank-2 flats, and let* $\mathcal{G}$ *consist of the eight atoms, the four pairs* A, B, C, D*, the two "opposite" unions* $A \cup C$ *and* $B \cup D$*, the four unions of three pairs, and* E*; thus* $|\mathcal{G}| = 19$*. Then*

(i) $\mathcal{G}$ *is a building set and* $(M, \mathcal{G})$ *is flag;*
(ii) $(M, \mathcal{G})$ *is not complete;*
(iii) $\mathcal{H}(M, \mathcal{G})(t) = 1 + 11t + 37t^2 + 55t^3 + 37t^4 + 11t^5 + t^6$*, which is palindromic of degree* 6*, with* γ*-vector* $(1, 5, 2, 1) \geq 0$*;*

(iv) $(1,5,2,1)$ *is not the f-vector of any simplicial complex.*

In particular the answer to [12, Question 9.8] *is negative, and so is the answer to the Kruskal–Katona intermediate step named there.*

Part (iv) needs no Kruskal–Katona machinery; it is the following statement, which is the one place where the refutation is a proof rather than a computation.

Lemma 3.3. *Let $\mathcal{S}$ be a simplicial complex, i.e. a downward-closed family of finite sets, and let f_i be the number of members of $\mathcal{S}$ of cardinality i. If $f_3 \geq 1$ then $f_2 \geq 3$.*

Proof. A member $s \in \mathcal{S}$ with $|s| = 3$ has three subsets of size 2, and all of them lie in $\mathcal{S}$ by downward closure. □

Proof of Theorem 3.2. Claims (i) and (ii) are exhaustive finite computations. That $\mathcal{G}$ is a building set is the criterion of [12, Prop. 2.2] checked on all ordered pairs of members. Flagness is [12, Lem. 3.15] applied literally to all 2^{19} subsets of $\mathcal{G}$: for each antichain whose total join lies in $\mathcal{G}$ some pairwise join lies in $\mathcal{G}$. Non-completeness is certified by the $F = \hat{0}$, $G = E$ case of Definition 2.2, quantified over all $8! = 40320$ total orders of E and all 248 flats: for no order does $\mathrm{ch}_{\lhd}(\hat{0}, E)$ stay inside $\mathcal{G}$. This last computation is what distinguishes a result from an error, since [12, Thm. 1.7] *proves* the f-vector conclusion for complete built matroids.

For (iii), the value of $\mathcal{H}$ is the sum (3), evaluated in three independent ways that agree: a pruned depth-first search over nested sets; a brute-force sum over every subset of $\mathcal{G}$ of size at most 7 with nestedness decided by quantifying over all antichains of the subset; and an earlier, separately written implementation. Four consistency checks that the computation does not impose are satisfied: $\mathcal{H}$ is palindromic, it is unimodal, its linear coefficient is $h_1 = |\mathcal{G}| - |E| = 19 - 8 = 11$, and its γ-vector is nonnegative as [12, Thm. 1.8] requires. Given the coefficient vector, the identity $(1+t)^6 + 5t(1+t)^4 + 2t^2(1+t)^2 + t^3 = \mathcal{H}$, the palindromicity and the nonnegativity are decided by kernel computation (§14).

For (iv), $\gamma_3 = 1$ and $\gamma_2 = 2$, so Lemma 3.3 applies. Numerically the same verdict is $\gamma_2^{(2)} = 2^{(2)} = 0 < 1 = \gamma_3$, the second Kruskal–Katona inequality, the one Remark 2.3 identifies as sharp. Since the f-vector of a flag complex is in particular an f-vector, both parts of [12, Question 9.8] are answered. □

Corollary 3.4. *Adjoin to $(M, \mathcal{G})$ of Theorem 3.2 j coloops $e_1, \dots, e_j$ together with the flats $E \cup \{e_1\} \subset \cdots \subset E \cup \{e_1, \dots, e_j\}$. For $j = 0, 1, 2, 3, 4$ the result is a flag built matroid of rank $7 + j$ on $8 + j$ elements with $|\mathcal{G}| = 19 + 2j$, whose Chow polynomial is $(1+t)^j \mathcal{H}(M, \mathcal{G})(t)$ and whose γ-vector is again $(1,5,2,1)$, padded with zeros. Hence a flag built matroid whose γ-vector is not an f-vector exists in each of the ranks* $7, 8, 9, 10, 11$.

Proof. Each of the five Chow polynomials was computed directly from (3) on the extended built matroid, not from its predecessor, and agrees with $(1+t)^j\mathcal{H}$; the tabulated values are

j	$\lvert\mathcal{G}\rvert$	$\mathcal{H}_j$
0	19	$1, 11, 37, 55, 37, 11, 1$
1	21	$1, 12, 48, 92, 92, 48, 12, 1$
2	23	$1, 13, 60, 140, 184, 140, 60, 13, 1$
3	25	$1, 14, 73, 200, 324, 324, 200, 73, 14, 1$
4	27	$1, 15, 87, 273, 524, 648, 524, 273, 87, 15, 1$

Multiplication by $(1+t)$ raises d by one and leaves (1) unchanged, so the γ-vector is preserved; this identity, and the failure of (4) for each padded vector, are checked by kernel computation. The pattern evidently continues, but no induction on j is formalised here, so the claim made is the finite one. $\square$

Remark 3.5. A second building set on the same $U(7,8)$, found by a separate descent, has 20 flats, is flag and not complete, and has $\gamma = (1,6,2,1)$, failing in the same way; it is recorded as a companion computation and is not needed above.

3.3. **What the counterexample does not kill.**

Proposition 3.6. $(1,5,2,1)$ *is an M-vector:* $\mathrm{Mac}(1,5,2,1)$ *holds, since* $2^{\langle 2\rangle} = 2 \geq 1$ *while* $2^{(2)} = 0$*. The same is true of every γ-vector computed in §3 and §5. Conversely* $\mathrm{Mac}(1,6,10,21)$ *fails, so the two conditions are separated in both directions.*

Thus [12, Question 9.8] fails exactly at the f-vector step and nowhere earlier; the O-sequence form of the question is untouched, consistently with the unconditional M-vector statement that CFL quote for the g-vector of a nested set complex.

3.4. **The rank stays at 7 and the ground set grows.** The coloop construction of Corollary 3.4 raises the rank while keeping $|E| = \mathrm{rk}\, M + 1$. The following moves in the independent direction.

Proposition 3.7. *Fix $s_1, \dots, s_4 \geq 2$ and let $M(s) = T_7\big(U(2,s_1) \oplus U(2,s_2) \oplus U(2,s_3) \oplus U(2,s_4)\big)$ be the rank-7 truncation of a direct sum of four lines, on $n = s_1 + \cdots + s_4$ elements, with $B_1, \dots, B_4$ the four lines as rank-2 flats. Let*

$$\begin{aligned}\mathcal{G}(s) = \{\text{the } n \text{ atoms}\} &\cup \{B_1, \dots, B_4\} \cup \{B_1 \vee B_2,\ B_3 \vee B_4\}\\ &\cup \{\text{the four triple joins}\} \cup \{E\},\end{aligned}$$

so $|\mathcal{G}(s)| = n + 11$. For each of the twelve size vectors s with $8 \leq n \leq 12$ given here, $(M(s), \mathcal{G}(s))$ is a flag, non-complete built matroid with $\mathcal{H} = 1 + 11t + 37t^2 + 55t^3 + 37t^4 + 11t^5 + t^6$ and $\gamma = (1,5,2,1)$, which is not an f-vector; eleven of the twelve are non-uniform. The size vectors are $(2,2,2,2)$, $(3,2,2,2)$, $(3,3,2,2)$, $(4,2,2,2)$, $(3,3,3,2)$, $(4,3,2,2)$, $(5,2,2,2)$, $(3,3,3,3)$, $(4,4,2,2)$, $(6,2,2,2)$, $(4,3,3,2)$ *and* $(5,3,2,2)$.

Proof. The member $s = (2,2,2,2)$ is Theorem 3.2, since $T_7(U(2,2)^{\oplus 4}) = U(7,8)$. For each of the twelve, $\mathcal{H}$ was computed from (3) by a pruned depth-first search; for the three smallest, also by the brute-force sum over all subsets of $\mathcal{G}(s)$ of size at most 7 with nestedness decided by quantification over antichains. Flagness was decided by the literal criterion [12, Lem. 3.15] over all $2^{|\mathcal{G}(s)|}$ subsets, for $|\mathcal{G}(s)| = 19, 20, 21$, and non-completeness by the exhaustive order search described in the proof of Theorem 3.2. The consistency checks hold on every member, including $h_1 = |\mathcal{G}(s)| - n = 11$. The tabulated coefficient vectors, their γ-expansions, γ-positivity, the failure of (4) and the truth of Mac are checked by kernel computation. $\square$

Remark 3.8. That $\mathcal{H}$ does not depend on s has a short argument: an atom a has $\mathrm{rk}\, a - \mathrm{rk}\, J_a^S - 1 = 0$, so its factor in (3) is empty, and an atom below some B_i empties B_i's factor as well; hence no atom occurs in a nonvanishing term and the sum runs over nested subsets of the fixed 11-element subposet $\{B_1, \dots, B_4,\ B_1 \vee B_2,\ B_3 \vee B_4, \text{four triple joins}, E\}$, whose joins and membership in $\mathcal{G}(s)$ do not see the s_i. This argument is not formalised, and it would need a formal model of built matroids; the claim recorded in Proposition 3.7 is the twelve instances.

3.5. **How the example was found, and how common it is.** Exhaustive enumeration by binary augmentation from $\mathcal{G}_{\min}$ is complete — by [12, Prop. 3.21] every flag building set is reached from $\mathcal{G}_{\min}$ one element at a time — and enumerates tightest first, but it does not scale: on $U(6,7)$ the level $|\mathcal{G}| = 15$ contains 318 920 building sets and not one is flag. The strategy that worked deletes instead. Starting from the always-flag $\mathcal{G}_{\max}$, repeatedly remove an element that leaves both a building set and a flag one, in random order, until nothing can be removed; the result is a *minimal* flag building set, which is the tight regime. On $U(5,6)$, where exhaustive enumeration is feasible, the descent reproduces its answers.

M	rk	method	flag	worst slack	violations
$U(5,6)$	5	exhaustive, $\lvert\mathcal{G}\rvert \le 15$	12 445	2	0
$U(5,6)$	5	descent	1 479	2	0
$U(6,7)$	6	descent	979	5	0
$U(7,8)$	7	descent	135	-1	1

The pattern is the one Remark 2.3 predicts: ranks 5 and 6 carry only the loose inequality $\gamma_2 \le \binom{\gamma_1}{2}$, rank 7 adds the sharp one, and the minimal flag building sets of $U(7,8)$ sit at slack 0 routinely. Counterexamples there are not rare: two runs sampling 221 of them met six violations, all with $|\mathcal{G}| = 19$ and the same $\mathcal{H}$, under different matchings. Whether $|\mathcal{G}| = 19$ is minimal in rank 7 is not settled here.

3.6. **Anchors.** Because everything rests on computed values of (3), the implementations were calibrated against published data, all reproduced exactly: the Eulerian polynomials $A_1, \ldots, A_6$ for $(B_n, \mathcal{G}_{\max})$; the series $\mathcal{H}(U(n-1,n), \mathcal{G}_{\min}) = 1 + t + \cdots + t^{n-2}$ for $2 \le n \le 7$, printed in [12, §1.1], whose γ-vectors — among them $(1,-3,1)$ and $(1,-4,3)$ — are CFL's own non-γ-positive examples; twelve entries of Table 1 of [15] for $U(k,k+1)$ and $U(k,k+2)$, for instance $U(7,9) \mapsto (1, 457, 8269, 20155, 8269, 457, 1)$; the Chow polynomial $(1, 13, 48, 73, 48, 13, 1)$ of the 20-element building set on B_7 in [12, Ex. 9.5], with $\gamma = (1,7,5,1)$; and the two built matroids of [12, Ex. 3.14], one flag and not complete, the other complete and not flag, on which our classifiers agree in the right directions.

4. The Kruskal–Katona pseudopower, in general

Theorem B is a statement about all k, so the arithmetic behind (4) has to be available as theorems rather than as evaluations on cases. The facts below are classical properties of the cascade representation; what is new here is only that they are proved about the computable pseudopower used in the verification, which computes the leading term of the cascade representation by binary search.

Theorem 4.1. *Let $i \ge 1$ and $m < 2^{128}$. Then*

(i) the binary search returns the largest a with $\binom{a}{i} \le m$; that is, $\binom{a}{i} \le m < \binom{a+1}{i}$ and $i \le a$ determine it;
(ii) $m^{(i)} = 0$ whenever $m \le i$, and $m^{(1)} = \binom{m}{2}$;
(iii) $m^{(i)} \ge 1$ if and only if $m \ge i + 1$;
(iv) $m^{(i)} = 1$ if and only if $i + 1 \le m \le 2i$;
(v) $m^{(i)} \ge 2$ whenever $m \ge 2i + 1$.

Proof sketch. (i) The search runs on $[i,\, m+i+1)$, which is safe because $\binom{m+i+1}{i} \geq m+2$ for $i \geq 1$; three invariants — that the lower endpoint always satisfies $\binom{\mathrm{lo}}{i} \leq m$, that the returned value satisfies it, and that no admissible a in the interval exceeds it — give the characterisation. The bound $m < 2^{128}$ is a genuine hypothesis, not an artefact: the search is run with a fixed budget of 128 halvings, so it is correct exactly while the interval fits. No number in this paper is anywhere near it. (ii) For $m \leq i$ the greedy representation is $\binom{i}{i} + \binom{i-1}{i-1} + \dots$, all of whose shifted terms $\binom{j}{j+1}$ vanish; for $i = 1$ the representation is $m = \binom{m}{1}$ and the shift is $\binom{m}{2}$. (iii) and (v) follow from monotonicity of $a \mapsto \binom{a}{i}$ together with $\binom{i+1}{i} = i+1$ and $\binom{i+2}{i+1} = i+2$; (iv) combines (iii) with (v) and the computation of the leading term for $i+1 \leq m \leq 2i$, where $a = i+1$ and the remainder $m-(i+1) \leq i-1$ contributes nothing by (ii). □

The two thresholds of Theorem 4.1(iii),(iv) are what Theorem B applies. Part (ii) identifies the level-1 inequality of (4) on a γ-vector as literally $\gamma_2 \leq \binom{\gamma_1}{2}$; see §6.

5. The k-block family: one law for every k

Throughout this section $k \geq 3$, $M_k = U(2k-1, 2k)$ with ground set partitioned into blocks $B_1, \dots, B_k$ of size 2, Γ is a graph on $[k]$ with $m = |\Gamma|$ edges, and $\mathcal{G}(k,\Gamma)$ is as in (2). The flats of M_k are the subsets of size at most $2k-2$ together with E, so a union of $r \leq k-1$ blocks is a flat of rank $2r$, and $\mathcal{G}(k,\Gamma)$ is a building set for every Γ. Its size is $|\mathcal{G}(k,\Gamma)| = 2^k - 1 + 2k - \binom{k}{2} + m$; for $k = 4$ and $m = 2$ this is 19, and indeed the built matroid of Theorem 3.2 is $(M_4, \mathcal{G}(4,\Gamma))$ with Γ a perfect matching of $[4]$.

5.1. Flagness is Mantel's condition.

Lemma 5.1. *For $k \geq 4$, $(M_k, \mathcal{G}(k,\Gamma))$ is flag if and only if the complement of Γ is triangle-free. Consequently*

$$\textit{flag} \implies m \geq \mathrm{mantel}(k) := \binom{k}{2} - \left\lfloor \frac{k^2}{4} \right\rfloor = \left\lfloor \frac{(k-1)^2}{4} \right\rfloor,$$

and the bound is attained, by $\Gamma = K_a \sqcup K_b$ with $a + b = k$ and $a = \lceil k/2 \rceil$, whose complement is the triangle-free $K_{a,b}$.

Proof. Necessity. If i, j, l are pairwise non-adjacent in Γ, then B_i, B_j, B_l are pairwise incomparable and no pairwise join lies in $\mathcal{G}(k,\Gamma)$, while their triple join does; by [12, Lem. 3.15] the built matroid is not flag. Hence the complement of Γ has no triangle.

Sufficiency. Write P for the non-atom part of $\mathcal{G}(k,\Gamma)$, whose members are unions of blocks. Two members of P that are unions of at least two blocks are either comparable or join to a union of at least three blocks, which lies in $\mathcal{G}(k,\Gamma)$; the same holds for a single block and a member of P of size at least two. So a pairwise-nested antichain is of one of three kinds: (I) atoms only, one per block; (II) atoms together with single blocks pairwise non-adjacent in Γ; (III) atoms together with one union S of $2 \leq r \leq k-2$ blocks, all atoms outside S — an atom and a union of $k-1$ blocks join to $E \in \mathcal{G}(k,\Gamma)$ and are not nested. In case (I) the join has at most k elements, one per block, so it is not a union of blocks and not in $\mathcal{G}(k,\Gamma)$. In case (III) the antichain has at least three members, hence $p \geq 2$ atoms from blocks outside S, and the join has $p + 2r \leq k + r \leq 2k - 2$ elements, again not a union of blocks. In case (II) the join lies in $\mathcal{G}(k, \Gamma)$ only if there are at least three blocks and no atoms — a triangle in the complement of Γ — or if it has at least $2k - 1$ elements, which forces $k - 1 \geq 3$ pairwise non-adjacent blocks and hence again a triangle.

The numerical consequence is Mantel's theorem: a triangle-free graph on k vertices has at most $\lfloor k^2/4 \rfloor$ edges, and $\binom{k}{2} - \lfloor k^2/4 \rfloor = \lfloor (k-1)^2/4 \rfloor$. □

Remark 5.2. For $k = 3$ the criterion degenerates: only $\Gamma = K_3$ gives a flag built matroid, whereas $\mathrm{mantel}(3) = 1$. This is why Theorem 5.5 is stated for $k \geq 4$. Lemma 5.1 is a hand proof and, unlike everything else recorded in §5, is not part of the formal development, which has no model of built matroids; before it was proved it was known as a computation over $k \leq 6$ (Proposition 5.6).

5.2. **The shape of the γ-vector.**

Proposition 5.3. *For every k with $3 \leq k \leq 12$ and every m in the flag range $\mathrm{mantel}(k) \leq m \leq \binom{k}{2}$ — 168 pairs in all — the Chow polynomial $\mathcal{H}(M_k, \mathcal{G}(k, \Gamma))$ depends only on (k, m) and not otherwise on Γ; it is palindromic of degree $2k - 2$; its γ-vector is nonnegative, is an M-vector, and has length k with*

$$\gamma_{k-2} = m, \qquad \gamma_{k-1} = 1.$$

Kruskal–Katona fails for exactly one of the 168 vectors, namely $(k, m) = (4, 2)$, where $\gamma = (1, 5, 2, 1)$. For $k = 4, 5, 6$ these values reproduce the ones of Proposition 5.6; the laws are, for instance, $\gamma = (1, 3 + m, m, 1)$ at $k = 4$, $(1, 13 + m, 5 + 3m, m, 1)$ at $k = 5$ and $(1, 38 + m, 49 + 9m, 5 + 6m, m, 1)$ at $k = 6$.

Proof. The values were computed twice. First by a literal sum over nested subsets of $\mathcal{G}(k, \Gamma)$, nestedness decided by quantification over all antichains; this is exponential in $|\mathcal{G}(k, \Gamma)|$ and was run for $k \leq 5$ and for the smallest $k = 6$ instance ($|\mathcal{G}| = 66$, 433 seconds). Second by a reduced sum obtained by classifying the nested sets: atoms never contribute, by the argument of Remark 3.8, so the sum runs over subsets of the non-atom part P, whose nested members consist of a chain of unions of at least two blocks together with at most two blocks inside its minimum, two only if they are non-adjacent in Γ. The reduced sum reproduces every value computed the first way and all anchors of §3.6; the nine remaining $k = 6$ instances were priced at about an hour of processor time each for a third confirmation and deliberately not run. Given the values, every assertion of the statement is decided by kernel computation over the whole table. □

Remark 5.4. The shape $\gamma_{k-2} = m$, $\gamma_{k-1} = 1$ is the one input of Theorem 5.5 that is data rather than theorem. The reduced description above makes $\mathcal{H}$ an *affine* function of m for fixed k, verified directly (vanishing second difference in m) for $3 \leq k \leq 40$; as the γ-transform is linear, two values of m then determine the whole flag range at those k, where the shape was confirmed as well. Extracting γ_{k-2} and γ_{k-1} from that description in closed form is a finite computation in binomials with k symbolic; it was not attempted, and doing it would make Theorem 5.5 unconditional.

5.3. **The law.**

Theorem 5.5. *Let $k \geq 4$ and let γ be a vector of length k with $\gamma_{k-2} = m < 2^{128}$ and $\gamma_{k-1} = 1$, and suppose m lies in the flag range, $m \geq \mathrm{mantel}(k)$. Then:*

(i) *the top Kruskal–Katona inequality* $\gamma_{k-1} \le \gamma_{k-2}^{(k-2)}$ *holds if and only if* $m \ge k-1$, *and is tight, with slack* 0, *if and only if* $k-1 \le m \le 2k-4$;

(ii) $\mathrm{mantel}(k) \ge k-1$ *if and only if* $k \ge 5$, *and* $\mathrm{mantel}(k) \le 2k-4$ *if and only if* $k \le 8$;

(iii) *consequently the top inequality holds if and only if* $(k,m) \ne (4,2)$; *and the member with the fewest edges,* $m = \mathrm{mantel}(k)$, *has slack exactly* 0 *if and only if* $5 \le k \le 8$.

In particular, in the family (2) *the counterexample of Theorem 3.2 is the only violation of the top inequality, for every* k *and not merely for the computed range.*

Proof. (i) is Theorem 4.1(iii),(iv) with $i = k-2 \ge 2$ and the two hypotheses $\gamma_{k-1} = 1$, $\gamma_{k-2} = m$. (ii) is elementary: with $k = j+2$, $\mathrm{mantel}(k) = \lfloor (j+1)^2/4 \rfloor \ge j+1$ reduces to $j^2 \ge 3j$, i.e. $j \ge 3$; with $k = j+4$, $\mathrm{mantel}(k) \le 2k-4$ reduces to $j^2 \le 4j$, i.e. $j \le 4$. For (iii), if $k = 4$ then $\mathrm{mantel}(4) = 2 < 3 = k-1$, so the inequality fails exactly at $m = 2$, the smallest flag value; if $k \ge 5$ then $m \ge \mathrm{mantel}(k) \ge k-1$ by (ii), so it always holds. Tightness at $m = \mathrm{mantel}(k)$ needs both $\mathrm{mantel}(k) \ge k-1$ and $\mathrm{mantel}(k) \le 2k-4$, which by (ii) is $5 \le k \le 8$. $\square$

The shape hypothesis of Theorem 5.5 is supplied by Proposition 5.3 for $3 \le k \le 12$ and by Remark 5.4 beyond it; the conclusion itself contains no data. Comparing the general law with the computed verdict on all 168 rows of Proposition 5.3 gives agreement in every case.

So the family fails once, hugs the boundary for four values of k, and then leaves it for good. The tabulated slacks of the members with the fewest edges are

k	4	5	6	7	8	9	10	11	12
$\mathrm{mantel}(k)$	2	4	6	9	12	16	20	25	30
top slack	-1	0	0	0	0	1	1	1	2

5.4. **The scan that preceded the law, and controls.**

Proposition 5.6. *Over* $k = 3, \dots, 6$ *the graphs* Γ *with triangle-free complement fall into* 62 *isomorphism classes,* 60 *giving flag built matroids, whose Chow polynomials collapse onto* 23 *distinct pairs* (m, γ), *each satisfying* $h_1 = |\mathcal{G}| - 2k$. *At* $k = 4$ *the* γ*-vector is* $(1, 3+m, m, 1)$ *and* (4) *holds if and only if* $m \ge 3$; *at* $k = 5, 6$ *the minimal flag member has top slack exactly* 0. *So* $k = 4$ *is the unique crossing point in the scanned range: growing* k *gives no second counterexample.*

Proposition 5.7. *The following controls hold. (i) The thresholds of Theorem 4.1 are not vacuous at the relevant levels:* $2^{(2)} = 0$, $3^{(2)} = 4^{(2)} = 1$, $5^{(2)} = 2$, $4^{(3)} = 6^{(3)} = 1$ *and* $7^{(3)} = 2$. *(ii) Tightness is sharp:* $\mathrm{KK}(1,17,17,4,1)$ *holds while* $\mathrm{KK}(1,17,17,4,2)$ *fails, so the minimal* $k=5$ *member really sits on the boundary; and* $\mathrm{KK}(1,5,3,1)$ *holds, so raising* γ_2 *by one repairs the counterexample — the failure is a single unit at a single place. (iii)* $\mathrm{mantel}(4) = 2$ *is exactly the failing number of edges, and* $k = 4$ *and* $k = 9$ *are the two crossing points. (iv) A correction to an earlier record: at* $k = 6$ *the top slack is* 0 *for* $m = 6, 7, 8$ *and* not *for* $m = 9$, *where* $\gamma = (1, 47, 130, 59, 9, 1)$ *and* $9^{(4)} = 2$, *so the slack is* 1.

Proposition 5.8. *Three further controls locate the hypotheses. (i) Replacing blocks of size* 2 *by blocks of size* 3 *or more — size vectors* $(3,2,2,2)$, $(3,3,2,2)$, $(3,3,3,3)$ *— gives* 17 *built matroids, none of which is flag, and not one of which is* γ*-positive; for instance* $\gamma = (1, 1, -8, 11, -5, 1)$ *on* $U(11, 12)$. *This is* [12, Thm. 1.8] *in contrapositive on real data. (ii) A failure of* (4) *off the flag class proves nothing: the* $k = 3$ *member with* $m = 1$ *has* $\gamma = (1, 1, 1)$*, which fails* (4)*, and it is not flag, while the flag* $k = 3$ *member (*$\Gamma = K_3$*) passes. (iii) The truncation to rank* 7 *in Proposition 3.7 is what bites: truncating the same four lines to rank* 8 *gives* $\gamma = (1, 4, 0, 0)$ *and to rank* 6 *gives* $\gamma = (1, 2, 1)$*, both of which pass* (4) *(the rank-*6 *member is not flag).*

6. The one inequality that survives in low rank

Proposition 6.1. *The level-*1 *inequality of* (4) *on a* γ*-vector is literally* $\gamma_2 \le \binom{\gamma_1}{2}$*, by Theorem 4.1(ii). It holds, with margin at least* 2*, on all* 234 *nonnegative* γ*-vectors computed in the course of this work and on all* 168 *of Proposition 5.3; the smallest margin is* 2*, attained at* $\gamma = (1, 3, 1)$*, and the next smallest is* 8*. It does not, however, follow from the numerical properties a Chow polynomial is known to have:* $h = (1, 6, 12, 6, 1)$ *is palindromic, is an* M*-vector, and has the nonnegative* γ*-vector* $(1, 2, 2)$*, which is itself an* M*-vector, yet* $\gamma_2 = 2 > 1 = \binom{\gamma_1}{2}$*; the same failure occurs in degree* 6 *with* $h = (1, 8, 25, 38, 25, 8, 1)$ *and* $\gamma = (1, 2, 2, 2)$*.*

By Remark 2.3 this inequality is the whole of the Kruskal–Katona condition in ranks 5 and 6, which is why those ranks stayed clean in §3. Proposition 6.1 says that proving $\gamma_2 \le \binom{\gamma_1}{2}$ for all flag built matroids — a genuine strengthening of [12, Thm. 1.8], untouched by Theorem 3.2, where $2 \le \binom{5}{2}$ — cannot be done numerically: it needs an interpretation of γ_1 as a vertex set and γ_2 as a set of pairs, a Macaulay estimate on h alone giving only $\gamma_2 \le \binom{\gamma_1}{2} + 3\gamma_1 + d$. We claim no progress on the inequality itself.

7. Verification of published statements on 121 matroids

The computational layer above was calibrated on instances of published theorems, using the maximal building set, for which $\mathcal{H}(M, \mathcal{G}_{\max})$ is the Chow polynomial $\underline{\mathcal{H}}_M$ of [15]. None of what follows is new mathematics. The 121 matroids are the nontrivial uniform matroids $U(r, n)$ with $2 \le r \le n \le 14$; the graphic matroids of K_4, K_5, K_6, K_7 and the one-edge deletions $K_4 - e$, $K_5 - e$, $K_6 - e$, of $K_{3,3}$, $K_{3,4}$, $K_{4,4}$, of the wheels W_4, W_5, W_6, and of the prism, the octahedron, the cube and the Petersen graph; and the Fano, non-Fano, $AG(3, 2)$ and Vámos matroids together with several sparse paving matroids and a handful of further small matroids of ranks 3 and 4.

Proposition 7.1. *Each of the* 121 *computed Chow polynomials is palindromic, its stated* γ*-vector satisfies* (1)*, and that vector is nonnegative.*

Palindromicity is Poincaré duality, not imposed by the computation, so those checks are real; the γ-positivity is an instance of [12, Thm. 1.6], as $(M, \mathcal{G}_{\max})$ is complete for every M.

Proposition 7.2. *Each of the* 121 γ*-vectors satisfies the Kruskal–Katona inequalities* (4)*, hence is the* f*-vector of a simplicial complex — an instance of* [12, Thm. 1.7].

Proposition 7.3. *Each of the* 121 *Chow polynomials is real-rooted, with* $\operatorname{rk} M - 1$ *distinct real roots.*

The certificates are sign alternations: a list of rationals $x_0 < \cdots < x_d$ at which P alternates in sign, carried as integer pairs and evaluated on the homogenisation $q^d P(p/q)$ so that all arithmetic stays in $\mathbb{Z}$; the intermediate value theorem gives a root in each interval, the intervals are disjoint, and $\deg P \leq d$ closes the count.

These 121 statements are instances of [12, Conj. 1.1]: that $\mathcal{H}(M, \mathcal{G}_{\max})(t)$ is real-rooted for every loopless matroid M. CFL attribute it to Huh–Stevens and to Ferroni–Schröter. In the primary sources it is [17, Conj. 8.18] — "the Hilbert–Poincaré series of the Chow ring of any matroid is a real-rooted polynomial" — and [27, Conj. 4.1.3], whose companion [27, Conj. 4.3.3] makes the same assertion for the augmented Chow ring; both Huh–Stevens statements add that deletion interlaces. The two are the first two items of [15, Conj. 1.6]. The numbering of the Ferroni–Schröter statement has moved: versions 1 and 2 of [15] cite it as Conjecture 10.19, the number it has in the first two arXiv postings of [17]; the current versions of [15], [8] and [13] all say Conjecture 8.18, which is the number it has in the last two. It is open in general. How much of the table above it leaves open is considerably less than the literature of a year ago would suggest, and we spell that out, because it is what makes this section a calibration rather than a contribution. First, the instances of rank at most 6 are a theorem, [12, Cor. 8.5], proved there from the balanced form of [12, Thm. 1.7] and the Frankl–Füredi–Kalai inequalities. Second, the uniform matroids — the bulk of the table — are a theorem as well: Brändén and Vecchi [7] prove that the Chow polynomial of every uniform matroid is real-rooted. What is *not* available for that purpose is [15, Thm. 1.10], which concerns the *augmented* Chow polynomial, a different series, and [15] record that they know of no analogue of it for the Chow ring itself. Third, every rank-3 entry of the table is covered by [8, Cor. 7.2]: its lattice of flats is the lattice of flats of a simple rank-3 matroid, and a simple matroid of rank 3 is paving. The same corollary covers the sparse paving entries $\mathrm{SP}(4,7)$, $\mathrm{SP}(4,8)$ and the Vámos matroid, and [8, Thm. 6.1] covers $AG(3,2)$, an affine geometry being a perfect matroid design; the graphic matroids of $K_4, \ldots, K_7$ are the braid matroids of type A, covered by [13, Thm. 1.8].

Running the whole table past these criteria — uniform, rank at most 6, paving, rank-uniform, braid — leaves exactly three entries uncovered, and all three are graphic: $M(Q_3)$ and $M(K_{4,4})$, of rank 7, and $M(\text{Petersen})$, of rank 9. None of the three is paving, each having a circuit shorter than its rank (4, 4 and 5 against 7, 7 and 9), and none has a rank-uniform lattice of flats: in each of the three graphs a chordless cycle of length g and an induced forest with $g-1$ edges are both flats of rank $g-1$, with g and $g-1$ atoms below them. We claim nothing further about those three beyond what Propositions 7.1–7.3 check, and we have not searched the literature exhaustively; the point of the audit is only that the other 118 entries are calibration and not evidence. Finally, real-rootedness does *not* extend to the hypotheses of (i) and (ii) of the introduction: [12, Ex. 9.5] — the built matroid of the B_7 anchor above — is flag and complete and its Chow polynomial is not real-rooted, and [12, Thm. 9.6] turns it into an infinite family. Conjecture 1.1, being about the maximal building set, is untouched by those examples, and so is everything in §3–§6. The Bose–Burton sections §8–§10 also use only the maximal building set.

Proposition 7.4. *Five negative controls hold. (i) The check is tight on real data: the γ-vector* $(1,7,5,1)$ *of* [12, Ex. 9.5] *sits one unit below the bound, and* $(1,7,5,3)$

is rejected. (ii) $(1,6,10,21)$ *is nonnegative and fails* (4)*, so the conclusion of* [12, Thm. 1.7] *is strictly stronger than γ-positivity. (iii) The same verdict for* $(1,6,10,21)$ *follows from a formalisation of the Kruskal–Katona theorem independent of our arithmetic:* 21 *triangles force at least* 15 *edges, and there are* 10*. (iv) The built matroids* $(U(n-1,n),\mathcal{G}_{\min})$*,* $n=4,\dots,8$*, have γ-vectors with negative entries and are rejected by the same machinery, so the flag or complete hypothesis is necessary. (v) Newton's inequalities* hold *for the non-real-rooted example of* [12, Ex. 9.5]*, so that cheap test cannot detect such examples; recorded as a limitation.*

8. Bose–Burton geometries

From here on the building set is always maximal, so that $\mathcal{H}(M,\mathcal{G}_{\max}) = \underline{\mathcal{H}}_M$, and the question is real-rootedness rather than the Kruskal–Katona inequalities.

8.1. **The geometry, its lattice and its Whitney numbers.** Fix a prime power q, an integer $r \geq 1$ and an integer k with $0 \leq k \leq r-1$. Let $V = \mathbb{F}_q^r$ and fix a subspace $W \leq V$ with $\dim W = k$. The *Bose–Burton geometry* $B_{r,k}(q)$ is the simple rank-r matroid whose ground set is the set of points of the projective space $PG(V) = PG(r-1,q)$ that do not lie in $PG(W)$ [4]. So $B_{r,0}(q) = PG(r-1,q)$ and $B_{r,r-1}(q) = AG(r-1,q)$; for $1 \leq k \leq r-2$ it is neither. These are the projective subspace deletions that appear in the work of Cheng and Liu on matroid Kazhdan–Lusztig polynomials [10], which is where they came to our attention; §8–§10 are independent of that work, and §11 onwards returns to it.

Lemma 8.1. *For $U \leq V$ write F_U for the set of points of $PG(U)$ outside $PG(W)$. The lattice of flats of $B_{r,k}(q)$ is*

$$\mathcal{L}(B_{r,k}(q)) \;=\; \{\hat{0}\} \cup \{F_U \,:\, U \leq V,\ U \not\leq W\},$$

with $F_U \leq F_{U'}$ if and only if $U \leq U'$, and $\operatorname{rk} F_U = \dim U$. For every $U \not\leq W$ the interval $[F_U, \hat{1}]$ is isomorphic to the lattice of subspaces of V/U. Consequently the number $s(b,d)$ of flats of rank d lying above a fixed flat of rank b depends only on b and d:

(5)
$$s(0,d) = W_d := \begin{bmatrix} r \\ d \end{bmatrix}_q - \begin{bmatrix} k \\ d \end{bmatrix}_q \quad (d \geq 1), \qquad W_0 := 1, \qquad s(b,d) = \begin{bmatrix} r-b \\ d-b \end{bmatrix}_q \quad (b \geq 1),$$

where $\begin{bmatrix} n \\ j \end{bmatrix}_q$ denotes the Gaussian binomial coefficient.

Proof. If $U \leq W$ then $F_U = \emptyset$. If $U \not\leq W$ then F_U spans U: choose $u \in U \setminus W$; a $v \in U$ either lies outside W, or lies in W, in which case $u+v \notin W$ and $v = (u+v)-u$ is a difference of two points of F_U. So the subspaces arising as spans of flats are exactly those not contained in W, $\operatorname{rk} F_U = \dim U$, and the order, meet and join are induced from the subspace lattice.

For the interval, it is enough to know that no point of $PG(V/U)$ is deleted, i.e. that for every L with $U < L \leq V$ and $\dim L = \dim U + 1$ some point of $L \setminus U$ lies outside W. Suppose not, and pick $u \in U \setminus W$ and $x \in L \setminus U$. Then $x \in W$, and $u + x$ again lies in $L \setminus U$, so $u + x \in W$; hence $u \in W$, a contradiction. Therefore the contraction of $B_{r,k}(q)$ by F_U is all of $PG(V/U)$, and $[F_U, \hat{1}]$ is the lattice of subspaces of V/U. The counts follow: W_d is the number of d-dimensional subspaces of V not contained in W, and above a nonzero flat of rank b the rank-d flats correspond to the $(d-b)$-dimensional subspaces of an $(r-b)$-dimensional space. □

A geometric lattice with the property just proved — the number of flats of rank d above a flat of rank b depends only on (b, d) — is *upper combinatorially uniform*, the term used by Brylawski [9]; it says exactly that the dual lattice is rank uniform, and [13] call such lattices *rank-uniform*. These are the lattices of [6, Conj. 4.3]. We write *upper rank uniform* for this property and *lower rank uniform* for its mirror, that all flats of a given rank have the same number of points — the perfect matroid design condition — and we keep the two words wherever the distinction matters, since the unqualified term is used for the first property in [13] and for the second in the total-nonnegativity literature.

8.2. **Modular flats, and why this family is in a gap.**

Lemma 8.2. *A flat $F_U \neq \hat{1}$ of $B_{r,k}(q)$ is modular if and only if $U \cap W = 0$. Hence the ranks of the modular flats are exactly $\{0, 1, \ldots, r-k\} \cup \{r\}$, and $B_{r,k}(q)$ is supersolvable if and only if $k \leq 1$.*

Proof. For flat spans U, U' one has $F_U \vee F_{U'} = F_{U+U'}$, while $F_U \wedge F_{U'}$ is $F_{U\cap U'}$ if $U \cap U' \not\leq W$ and $\hat{0}$ otherwise; and $\dim U + \dim U' = \dim(U + U') + \dim(U \cap U')$ always holds. So F_U is modular exactly when no flat span U' has $U \cap U'$ nonzero and contained in W.

If $U \cap W \neq 0$, pick $0 \neq x \in U \cap W$ and $y \notin U \cup W$ — possible, since a vector space is not the union of two proper subspaces — and set $U' = \langle x, y\rangle$. Then $U' \not\leq W$, so U' is a flat span, and $U \cap U' = \langle x\rangle$ because $y \notin U$; this is a nonzero subspace of W, so F_U is not modular. If $U \cap W = 0$, then $U \cap U' \leq W$ forces $U \cap U' = 0$, and F_U is modular.

A subspace meeting W trivially has dimension at most $r - k$, which gives the list of modular ranks, the top element being modular always. A maximal chain of modular flats requires a modular flat of rank $r-1$; for $k \geq 2$ every hyperplane meets W in dimension at least $k-1 \geq 1$, so none exists. For $k \leq 1$ there is a hyperplane H with $H \cap W = 0$, and any complete flag of subspaces inside H consists of modular flats. □

Proposition 8.3. *Let $r \geq 4$ and $2 \leq k \leq r-2$. The lattice $\mathcal{L}(B_{r,k}(q))$ is upper combinatorially uniform, by Lemma 8.1, but*

(i) it is not supersolvable, by Lemma 8.2; in particular it is not a Dowling lattice, Dowling lattices being supersolvable;

(ii) $B_{r,k}(q)$ is not a perfect matroid design — equivalently, $\mathcal{L}(B_{r,k}(q))$ is not lower rank uniform: a line of $PG(V)$ disjoint from W gives a rank-2 flat with $q+1$ points and a line meeting W in a point contributes one with q points, and both kinds exist precisely when $1 \leq k \leq r-2$;

(iii) $B_{r,k}(q)$ is not paving: a line disjoint from W is a dependent flat of rank $2 \leq r-2$.

The counts (5) *and the three verdicts were also obtained by direct enumeration of the lattice of flats inside $\mathbb{F}_q^r$ — every subspace enumerated, flats filtered, joins by span and meets by intersection — for $q \in \{2, 3\}$, $r \leq 5$ and every k, twenty-three lattices in all. There the lattice came out upper combinatorially uniform with the predicted counts in every case, lower rank uniform exactly for $k \in \{0, r-1\}$, and supersolvable exactly for $k \leq 1$, with modular ranks $\{0, \ldots, r-k\} \cup \{r\}$.*

Consequently the class hypotheses of the published sufficient conditions for real-rootedness of the Chow polynomial all fail for $B_{r,k}(q)$ when $2 \leq k \leq r-2$: the matroid production rules of [13, Thm. 1.8] end at supersolvable rank-uniform lattices, the relevant entry in the list of TN-posets in [8] is perfect matroid designs, and their Corollary 7.2 is about paving matroids. What this does *not* say is that the instances themselves are all open: [12, Cor. 8.5] settles every matroid of rank at most 6 regardless of class, and Remark 9.4 records a second caveat. The range is nonempty exactly for $r \geq 4$. The other values of k are classical or degenerate: $k = 0$ and $k = r-1$ give the projective and affine geometries, both perfect matroid designs, and $k = 1$ carries no information at all, by Proposition 8.4(iii) below.

8.3. **Computing the Chow polynomial.** For the maximal building set the nested sets of (3) are exactly the chains of nonzero flats, so

$$\underline{\mathcal{H}}_M(t) \;=\; \sum_{\hat{0}=F_0<F_1<\cdots<F_j} \prod_{i=1}^{j} \big(t + t^2 + \cdots + t^{\,\mathrm{rk}\,F_i - \mathrm{rk}\,F_{i-1} - 1}\big), \tag{6}$$

a chain in which some rank gap equals 1 contributing 0. This is the case $g = \zeta$ of the non-recursive chain formula of Ferroni, Matherne and Vecchi [16, Thm. 3.10]; we use it in the form [8, Cor. 2.7], which is that theorem specialised to a scalar g — there $g_{z,y} = 1$ and each factor $t^{\rho-1} - 1$ over $t-1$, multiplied by the t that [8, Cor. 2.7] carries outside the product, is $t + \cdots + t^{\rho-1}$. Corollary 2.9 of [8] is a different identity, expressing $H_{x,y}$ through the Chow polynomials of the truncations of $[x, y]$, and is not what is used here. On a lattice whose upper intervals are as in Lemma 8.1, (6) collapses to an $O(r^2)$ recursion in the numbers (5):

$$C_0 = 1, \qquad C_d = \sum_{b \leq d-2} C_b \cdot s(b,d) \cdot \big(t + \cdots + t^{\,d-b-1}\big), \qquad \underline{\mathcal{H}}_{B_{r,k}(q)} = \sum_{d=0}^{r} C_d. \tag{7}$$

Proposition 8.4. *Evaluated inside the Lean kernel from the Gaussian binomials, the recursion* (7)

- *(i) reproduces all* 41 *Chow polynomials used in §10;*
- *(ii) agrees with the published and previously recorded values it can be compared against: the Fano matroid* $B_{3,0}(2)$ *giving* $1+8t+t^2$*, the affine plane* $B_{3,2}(2) = AG(2,2) = U(3,4)$ *giving* $1+7t+t^2$*, and the projective geometries* $PG(2,3) \mapsto 1+14t+t^2$*,* $PG(2,4) \mapsto 1+22t+t^2$*,* $PG(2,5) \mapsto 1+32t+t^2$*,* $PG(2,7) \mapsto 1+58t+t^2$*,* $PG(3,2) \mapsto 1+51t+51t^2+t^3$*,* $PG(4,2) \mapsto 1+342t+1582t^2+342t^3+t^4$ *and* $PG(6,2) \mapsto 1+29084t+1638047t^2+5547488t^3+\ldots$*;*
- *(iii) gives* $\underline{\mathcal{H}}_{B_{r,1}(q)} = \underline{\mathcal{H}}_{PG(r-1,q)}$ *for* $q \in \{2,3,4,5\}$ *and* $2 \leq r \leq 8$*;*
- *(iv) and is sensitive to its input: feeding it the Whitney numbers of a different* k *changes the answer.*

Part (iii) has a one-line reason, and it is why the novel range starts at $k = 2$: a rank-1 flat sits at gap 1 above $\hat{0}$, so it contributes weight 0 to every chain in (6), and deleting one point from $PG(r-1,q)$ removes only an atom. Part (ii) is the calibration; $\underline{\mathcal{H}}_{PG(r-1,q)}$ is the Shareshian–Wachs maj-exc q-Eulerian polynomial, by Hameister, Rao and Simpson [19], and the first two values sit verbatim in the data of §7. Nothing in Proposition 8.4 is new mathematics; what it changes is that the coefficients used below are *derived* in the kernel from the Whitney numbers instead of being asserted. The smallest instances of the range $2 \le k \le r-2$ are

$$\begin{array}{r|l} B_{4,2}(2) & 1 + 50t + 50t^2 + t^3 \\ B_{4,2}(3) & 1 + 170t + 170t^2 + t^3 \\ B_{5,2}(2) & 1 + 341t + 1574t^2 + 341t^3 + t^4 \\ B_{5,3}(2) & 1 + 334t + 1524t^2 + 334t^3 + t^4 \\ B_{7,3}(2) & 1 + 29076t + 1635601t^2 + 5536312t^3 + \dots \end{array}$$

which may be compared with $PG(3,2)$, $PG(4,2)$ and $PG(6,2)$ above: in each of these instances deleting the subspace lowers every interior coefficient, and only slightly.

9. The dual of a Bose–Burton lattice is a TN-poset

9.1. **Totally nonnegative posets.** We recall the definitions of [5]. Let P be a locally finite poset with a least element $\hat{0}$, and let $\rho(x)$ be the largest m for which there is a chain $\hat{0} = x_0 < x_1 < \cdots < x_m = x$. Then P is *quasi-rank uniform* if $|\{z \le x : \rho(z) = j\}|$ depends only on $\rho(x)$ and j; writing $r_{n,j}$ for the common value, this defines a lower unitriangular matrix $R(P) = (r_{n,j})$. The poset P is a *TN-poset* when $R(P)$ is totally nonnegative, i.e. all its minors are nonnegative.

A lower unitriangular matrix $R = (r_{n,j})$, with row polynomials $R_n(t) = \sum_j r_{n,j} t^j$, is *resolvable* if there are an array $(\lambda_{n,j})$ of *nonnegative* numbers and monic polynomials $R_{n,j}$, $0 \le j \le n$, with

$$R_{n,0} = R_n, \qquad R_{n,n} = t^n, \qquad t^j \mid R_{n,j}, \qquad R_{n+1,j} = R_{n+1,j+1} + \lambda_{n,j} R_{n,j}.$$

The tool we use is [5, Thm. 2.6]: a lower unitriangular matrix is totally nonnegative if and only if it is resolvable. Conjecture 4.3 of [6] reads, verbatim, "The dual of any upper combinatorially uniform geometric lattice is a TN-poset". It is offered there immediately after [6, Thm. 4.1], *the dual of any Dowling lattice is a TN-poset*, and after the observation that perfect matroid designs are upper combinatorially uniform as well; of those, [6, Thm. 3.7] proves that a triangular semimodular lattice — hence a perfect matroid design, by [6, Prop. 3.3] — is itself a TN-poset. The three classes named in the abstract of [6] — perfect matroid designs, Dowling lattices, and a class of geometric lattices containing the lattices of flats of all paving matroids — are the classes for which the paper settles a *different* conjecture, that of Athanasiadis and Kalampogia-Evangelinou [3, Conj. 1.2] on chain polynomials of geometric lattices; see Remark 9.6.

Remark 9.1. The quotation is from version 3 of [6], of 11 December 2025. Versions 1 and 2, of August 2025, carry a misprinted form of the conjecture, asserting instead that upper combinatorially uniform geometric lattices are themselves TN-posets. Coron, Ferroni and Li flag that misprint and record, on Brändén's authority, the intended statement: "they conjecture that $\mathcal{L}(M)^*$ is a TN-poset whenever $\mathcal{L}(M)$ is rank-uniform" [13, §8.2, footnote] — which, since [13] use *rank-uniform* for Brylawski's upper combinatorially uniform, is the statement quoted above, the one now printed in version 3, and the one proved here for the Bose–Burton family. The conjecture is numbered 4.3 in all three versions.

By Lemma 8.1 the dual $\mathcal{L}^* = \mathcal{L}(B_{r,k}(q))^*$ is quasi-rank uniform: an element of $\mathcal{L}^*$ of quasi-rank n is a flat of $\mathcal{L}$ of rank $r - n$, and the elements below it of quasi-rank j are the flats of rank $r-j$ above it, of which there are $s(r-n, r-j)$. With (5) and the symmetry $\begin{bmatrix} n \\ n-j \end{bmatrix}_q = \begin{bmatrix} n \\ j \end{bmatrix}_q$ this gives

$$R(\mathcal{L}^*)_{n,j} = \begin{bmatrix} n \\ j \end{bmatrix}_q \quad (0 \le n \le r-1), \qquad R(\mathcal{L}^*)_{r,j} = W_{r-j}. \tag{8}$$

So the whole question is a statement about one explicit matrix: the q-Pascal matrix with its bottom row replaced by the Whitney numbers.

9.2. **The theorem.**

Theorem 9.2. *Let $q > 1$ be real, let $r \ge 1$ and $0 \le k \le r-1$, and put $s = r-k$. Let M be the $(r+1) \times (r+1)$ lower unitriangular matrix of* (8)*. Then M is resolvable, its resolving array is unique, and its coefficients are $\lambda_{n,j} = q^j$ for $n \le r-2$, while the last row of the array is*

$$\lambda_{r-1,j} \;=\; \lambda_j \;:=\; \begin{cases} q^j, & 0 \le j < s, \\ q^{\binom{s}{2} - \binom{j}{2} - j(r-1-j)} \Big(\displaystyle\prod_{i=s}^{j} (q^i - 1) \Big) \Sigma(r,s,j), & s \le j \le r-1, \end{cases}$$

where

$$\Sigma(r,s,j) \;=\; \sum_{m+n=r-j-1} q^{jn} \begin{bmatrix} j-s+m \\ m \end{bmatrix}_q.$$

Every λ_j is strictly positive. In particular M is totally nonnegative.

Proof sketch. Write $P_n(t) = \sum_j \begin{bmatrix} n \\ j \end{bmatrix}_q t^j$, let T be multiplication by t and let α be the diagonal operator $\alpha(t^m) = q^m t^m$ on $\mathbb{R}[t]$.

(1) *The q-Pascal part resolves with $\lambda_{n,j} = q^j$.* The q-Pascal recursion is exactly $(T+\alpha)P_{n-1} = P_n$, so $P_n = (T+\alpha)^n 1$, and [5, Thm. 2.6] gives $R_{n,j} = (T+\alpha)^{n-j} t^j$. As $\alpha T = qT\alpha$, the noncommutative q-binomial theorem yields the closed form $R_{n,j}(t) = \sum_{i \ge j} \begin{bmatrix} n-j \\ i-j \end{bmatrix}_q q^{j(n-i)} t^i$.

(2) *The array is forced.* For each n the polynomials $R_{n,j}$, $0 \le j \le n$, form a basis of the space of polynomials of degree at most n: the differences $R_{n,j} - R_{n,j+1}$ have lowest term t^j with positive coefficient, and $R_{n,n} = t^n$. The resolvability relations express $R_{n+1,j}$ as t^{n+1} plus a combination of the $R_{n,i}$ with $i \ge j$, so the single row $R_{n+1,0} = R_{n+1}$ determines the whole of $\lambda_{n,\cdot}$ and hence the next row of the array. Starting from $R_{0,0} = 1$, the array is unique. Rows $0, \ldots, r-1$ of M are q-Pascal, so the only unknowns are $\lambda_j := \lambda_{r-1,j}$, determined by $\sum_j \lambda_j R_{r-1,j} = f$, where $f(t) = \sum_i M_{r,i} t^i - t^r$. By [5, Thm. 2.6], M is totally nonnegative if and only if these forced λ_j are all nonnegative.

(3) *The target.* $M_{r,i} = W_{r-i} = \begin{bmatrix} r \\ i \end{bmatrix}_q - \begin{bmatrix} k \\ r-i \end{bmatrix}_q$ for $i < r$, and $\sum_i \begin{bmatrix} k \\ r-i \end{bmatrix}_q t^i = t^s P_k(t)$, so $f = P_r - t^s P_{r-s}$.

(4) *Inversion.* Put $N = r-1$. Since $R_{N,j} = (T+\alpha) R_{N-1,j}$ for $j \le N-1$ and $R_{N,N} = t^N$, writing $f = (T+\alpha)g + \lambda_N t^N$ and comparing coefficients gives $f_i = g_{i-1} + q^i g_i$ for $i \le N-1$, and then one recurses on g. Passing to the generating function $\hat{f}(z) = \sum_i q^{\binom{i}{2}} f_i z^i$ turns the recursion into multiplication by $(1+z)$ after a rescaling $z \mapsto z/q$, and unwinding gives

$$\lambda_j \;=\; q^{-\binom{j}{2} - j(N-j)} \, [z^j] \Big(\hat{f}(z) \Big/ \prod_{b=0}^{N-j} (1 + q^b z) \Big). \tag{9}$$

(5) $\hat{f}$ *factorises.* The q-binomial theorem gives $\sum_i q^{\binom{i}{2}} \left[\begin{smallmatrix} n \\ i \end{smallmatrix}\right]_q z^i = \prod_{i=0}^{n-1}(1+q^i z)$, and only the coefficients of z^j with $j \le N < r$ occur in (9); so modulo z^r

$$\hat{f}(z) \;\equiv\; E(z)\prod_{i=s}^{r-1}(1+q^i z), \qquad D(z) := \prod_{b=0}^{s-1}(1+q^b z),$$
$$E(z) := \sum_{m=0}^{s-1} q^{\binom{m}{2}} \left[\begin{smallmatrix} s \\ m \end{smallmatrix}\right]_q z^m \;=\; D(z) - q^{\binom{s}{2}} z^s.$$

Cancelling D against the first s factors of the denominator of (9) leaves

$$q^{\binom{j}{2}+j(N-j)}\lambda_j \;=\; [z^j]\Big((E/D)\prod_{i=r-j}^{r-1}(1+q^i z)\Big).$$

(6) *Partial fractions.* As $\deg E < \deg D$, the quotient E/D is a sum of simple fractions $A_b/(1+q^b z)$, $0 \le b \le s-1$, whose residues are computed from $E(-q^{-b}) = (-1)^{s+1} q^{\binom{s}{2}-bs}$. Extracting the coefficient of z^j — the contribution of the term b vanishing when $b \ge r-j$ — and writing $[m]! = \prod_{i=1}^m (q^i - 1) > 0$ gives

$$q^{\binom{j}{2}+j(N-j)}\lambda_j \;=\; q^{\binom{s}{2}}\,\frac{[j]!}{[s-1]!}\,\Sigma(r,s,j),$$
$$\Sigma(r,s,j) \;=\; \sum_{b=0}^{s-1}(-1)^b q^{\binom{b+1}{2}+b(j-s)} \left[\begin{smallmatrix} s-1 \\ b \end{smallmatrix}\right]_q \left[\begin{smallmatrix} r-1-b \\ j \end{smallmatrix}\right]_q,$$

which is the stated formula once $[j]!/[s-1]! = \prod_{i=s}^{j}(q^i - 1)$ is substituted.

(7) *The alternating sum is positive.* Summing over r and using $\sum_m \left[\begin{smallmatrix} m \\ j \end{smallmatrix}\right]_q z^m = z^j / \prod_{i=0}^{j}(1-q^i z)$ together with $\sum_b (-1)^b q^{\binom{b}{2}} \left[\begin{smallmatrix} s-1 \\ b \end{smallmatrix}\right]_q x^b = \prod_{i=0}^{s-2}(1-q^i x)$ at $x = q^{j-s+1}z$, the numerator cancels all but two factors of the denominator, and

$$\sum_{r\ge 0}\Sigma(r,s,j)z^r \;=\; \frac{z^{j+1}}{(1-q^j z)\prod_{i=0}^{j-s}(1-q^i z)},$$

legitimately, because $j \ge s \ge 1$ in this case. Both remaining factors expand with nonnegative coefficients, $1/(1-q^j z) = \sum_n q^{jn} z^n$ and $\prod_{i=0}^{j-s}(1-q^i z)^{-1} = \sum_m \left[\begin{smallmatrix} j-s+m \\ m \end{smallmatrix}\right]_q z^m$, and reading off the coefficient of z^r gives the displayed $\Sigma(r,s,j)$, all of whose terms are positive — the term with $n = r-j-1$, $m = 0$ alone is $q^{j(r-j-1)} > 0$. For $j < s$ step (5) degenerates and (9) returns $\lambda_j = q^j$ directly. $\square$

Corollary 9.3. *For every prime power q, every $r \ge 1$ and every $0 \le k \le r-1$, the dual of the lattice of flats of $B_{r,k}(q)$ is a TN-poset. Thus* [6, Conj. 4.3] *holds for the Bose–Burton geometries; and by Proposition 8.3, for $2 \le k \le r-2$ these lattices are neither perfect matroid designs nor Dowling lattices, and $B_{r,k}(q)$ is not paving.*

Proof. Combine (8), Theorem 9.2 and the definition of a TN-poset. $\square$

Remark 9.4. Two caveats on the word “outside”, both about classes rather than about $B_{r,k}(q)$ itself.

First, the largest class [6] treat is the *L-paving lattices*: given a finite geometric lattice L and a generalised d-partition $H \subseteq L$, the lattice $L(H, \hat{1}, d)$ obtained by keeping the elements of L of rank at most $d-1$, putting H at rank d and $\hat{1}$ at rank

$d+1$ [6, §5.1]. For L Boolean this is a paving lattice; [6, Thm. 5.9] treats every L-paving lattice with L rank uniform. That class does not contain ours either: a lattice of the form $L(H, \hat{1}, d)$ of rank r has $d = r-1$ and agrees with L in all ranks $\leq r-2$, so if $\mathcal{L}(B_{r,k}(q))$ were one, with $r \geq 4$ and $1 \leq k \leq r-2$, the rank-uniform lattice L would have to carry the rank-2 flats of $B_{r,k}(q)$, some with $q+1$ atoms and some with q (Proposition 8.3(ii)), which a rank-uniform lattice cannot. The conclusion of [6, Thm. 5.9] is in any case about chain polynomials, not about [6, Conj. 4.3].

Second, [13, Thm. 1.2] proves real-rootedness for every UMEL-shellable poset, and being UMEL-shellable means being rank uniform — which $\mathcal{L}(B_{r,k}(q))$ is, by Lemma 8.1 — and carrying an edge-lexicographic labelling that is uniform and monotonic. We have not decided whether $\mathcal{L}(B_{r,k}(q))$ admits such a labelling, and we do not claim it does not; what we claim is only that $B_{r,k}(q)$ with $2 \leq k \leq r-2$ is in none of the classes for which [13, Thm. 1.8] draws the conclusion. Coron, Ferroni and Li themselves conjecture that a matroid whose lattice of flats is a TN-poset has a UMEL-shellable lattice of flats [13, Conj. 8.7], and ask the converse question [13, Ques. 8.8], whether $\mathcal{L}(M)^*$ is a TN-poset when $\mathcal{L}(M)$ is UMEL-shellable — which they note is close to [6, Conj. 4.3]. The Bose–Burton geometries sit exactly where those two meet.

Remark 9.5. At $q=1$ the coefficients of Theorem 9.2 collapse to $(1, \ldots, 1, 0, \ldots, 0)$ with s ones. For $q \neq 1$ they are not integers — at $q=2$, $r=4$, $s=1$ they are $(1, \frac{7}{4}, \frac{21}{8}, \frac{21}{8})$ — so no counting or shelling argument produces them, and a closed form is what the proof needs. In the formal development the common denominator $q^{\binom{r-1}{2}}$ is cleared, so that an identity verified between natural numbers exhibits a *nonnegative* resolving array, which is the hypothesis [5, Thm. 2.6] asks for.

Remark 9.6. Theorem 9.2 feeds a second published machine as well. Let Q be the poset of nonempty flats of $B_{r,k}(q)$ ordered by reverse inclusion, so that Q is $\mathcal{L}^*$ with its top element removed and $\widehat{Q} = \mathcal{L}^*$. Every lower interval of Q of quasi-rank n is an interval $[F, \hat{1}]$ of $\mathcal{L}$ with $\operatorname{rk} F = r-n$, hence a subspace lattice of dimension n (Lemma 8.1); so Q is a P-poset in the sense of [5, Def. 6.1] for P the subspace lattice of $\mathbb{F}_q^{r-1}$, and $\widehat{Q}$ is TN by Corollary 9.3, so Q is P-positive by [5, Thm. 6.4]. Then [5, Thm. 6.6] puts every zero of the chain polynomial c_Q in $[-1, 0]$. Adjoining a bounding element multiplies a chain polynomial by $1+t$, and a poset and its dual have the same chain polynomial, so $c_{\mathcal{L}} = c_{\mathcal{L}^*} = (1+t)c_Q$ is $[-1,0]$-rooted too. These are instances of the conjecture of Athanasiadis and Kalampogia-Evangelinou [3, Conj. 1.2], that the chain polynomial of every geometric lattice is real-rooted. We record the consequence and do not develop it: it is not part of the formal development, and while Proposition 8.3 and Remark 9.4 put the family outside the classes for which [6] prove that conjecture, the UMEL question of Remark 9.4 is open here too, so we claim no novelty for it.

9.3. **What is machine-checked.**

Proposition 9.7. *The closed form of Theorem 9.2 is verified against the matrix* (8) *by kernel computation for* 381 *matrices: $q=2$ with $r \leq 12$; $q=3$ with $r \leq 10$; $q=4,5$ with $r \leq 8$; $q=7,8,9$ with $r \leq 6$; $q=11,13,16,25,27,32$ with $r \leq 5$; and $q=49,81,121,128$ with $r \leq 4$; in each case for every k with $0 \leq k \leq r-1$. For the fourteen matrices with $q=2$, $2 \leq r \leq 5$ and $0 \leq k \leq r-1$ — among them $B_{4,2}(2)$,* *the smallest member of the range* $2 \le k \le r-2$*, and the rank-*5 *instance* $B_{5,2}(2)$ *— total nonnegativity is in addition decided directly from the definition, over every one of the* $\binom{2n}{n}$ *minors of the* $n \times n$ *matrix, as it is for seventeen further* 5×5 *matrices arising in the controls of Proposition 10.4.*

The two routes are independent: the certificate of Theorem 9.2 is an identity between polynomials with natural-number coefficients, while the direct check forms every minor and computes its determinant. They agree everywhere both were run, which is what makes the cheap certificate usable at the sizes where the direct check is not: one rank-6 matrix costs about 8 gigabytes in the kernel, and the certificate does all of $q = 2$, $r \le 12$ in seconds.

10. Real-rootedness for the Bose–Burton geometries

Corollary 10.1. *Let* q *be a prime power,* $r \ge 1$ *and* $0 \le k \le r-1$*. Then all four polynomials that* [8, Thm. 6.1] *attaches to* $\mathcal{L}(B_{r,k}(q))$ *are real-rooted: the Chow polynomial* $\underline{\mathcal{H}}_{B_{r,k}(q)}$*, the augmented Chow polynomial, the Chow-derangement polynomial and the Chow-Eulerian polynomial.*

Proof. [8, Thm. 6.1] states: *Suppose* P *is a TN-poset, or the dual of a TN-poset, of rank* n*. Then the Chow polynomials* H_P*,* d_P*,* G_P *and* A_P *are real-rooted.* Apply it to $P = \mathcal{L}(B_{r,k}(q))$, which is the dual of the TN-poset of Corollary 9.3. Here g is the zeta function, so H_P is the characteristic Chow polynomial of P, which for $P = \mathcal{L}(M)$ is the Hilbert–Poincaré series $\underline{\mathcal{H}}_M$ of the Chow ring [16, Thm. 4.9], and G_P, the *augmented* Chow polynomial of [8, §2.1], is the Hilbert series of the augmented Chow ring [16, Thm. 4.11]. The remaining two, d_P and A_P, are the Chow-derangement and the Chow-Eulerian polynomial of P [8, §2]; they carry no separate matroid names. □

For $2 \le k \le r-2$ this is a family of instances of [12, Conj. 1.1] and of its augmented companion, for which we could find no published sufficient condition, subject to the caveats of Remark 9.4 and to one overlap: [12, Cor. 8.5] reaches every matroid of rank at most 6, and the range $2 \le k \le r-2$ is nonempty from $r = 4$ on, so the two overlap but neither contains the other. That corollary is about the Chow polynomial only, so in the whole range $2 \le k \le r-2$ the augmented statement is untouched by it, and the augmented Chow polynomial is a series untouched elsewhere in this paper. Two things this does *not* say, and we are explicit about both. It says nothing about [6, Conj. 4.3] for upper combinatorially uniform geometric lattices in general, which remains open. And the Bose–Burton geometries are not a family anybody has asked about in print: they are a natural instance of an open conjecture, not a posed question.

10.1. The finite verification, and a second proof of it.

Proposition 10.2. *For each of* 41 *Bose–Burton geometries —* $q \in \{2,3,4,5,7,9\}$ *and* $r \le 8$*, of which* 29 *have* $2 \le k \le r-2$ *— the Chow polynomial computed by Proposition 8.4 is palindromic; its stated* γ*-vector satisfies* (1)*, is nonnegative and satisfies the Kruskal–Katona inequalities* (4)*; Newton's inequalities hold; and the polynomial is real-rooted with* $r-1$ *distinct real roots, by a sign-alternation certificate. The certificates are the ones of §7: lists of rationals carried as integer pairs and evaluated on the homogenisation* $q^d P(p/q)$*.*

Palindromicity is Poincaré duality and is not imposed by (7), so those checks are real; γ-positivity and the Kruskal–Katona inequalities are instances of [12, Thms. 1.6 and 1.7], the maximal building set being complete. Only the real-rootedness half is what Corollary 10.1 is about, and there the 41 certificates are an independent confirmation of a theorem proved by other means — the point of running them is that they are checked by a different mechanism from the proof.

Proposition 10.3. *Kurtz's criterion* [21] *— a polynomial with positive coefficients a_i satisfying $a_i^2 > 4a_{i-1}a_{i+1}$ at every interior i has only real, simple roots — holds on all* 41 *γ-vectors above. A palindromic polynomial with nonnegative coefficients is real-rooted if and only if its γ-polynomial is* [8, Lem. 5.5]*, so each of those* 41 *instances has a second, entirely elementary proof of real-rootedness, independent of total nonnegativity. The criterion is strictly sufficient: it fails on the γ-vector* $(1, 7, 5, 1)$ *of* [12, Ex. 9.5]*, correctly, that Chow polynomial not being real-rooted; and it also fails on the γ-vector of the Eulerian polynomial A_{15}, which is real-rooted.*

Kurtz's criterion is elementary but does not scale. Measured over the Bose–Burton family with $q \leq 5$, the ratio $\gamma_i^2/4\gamma_{i-1}\gamma_{i+1}$, minimised over the interior i, falls from about 29.8 at rank 5 to about 1.96 at rank 13, the minimum always being attained at $k = r - 1$, $q = 2$. For the Boolean matroids the same ratio has already fallen below 1 at rank 15, where the criterion fails on a polynomial that is real-rooted. So the test covers the instances tabulated here and cannot be expected to reach much further, which is exactly why Corollary 10.1 is worth having.

Proposition 10.4. *The following controls hold. (i) Total nonnegativity is a genuine constraint on the bottom row of* (8)*: of sixty random nonnegative rows of the right shape,* 45 *destroy it at $r = 4$ and* 59 *of sixty at $r = 5$; four such failures are decided in the formal development, two of them with an explicit negative* 2×2 *minor. (ii) It is nevertheless not delicate: for $B_{4,2}(2)$, whose true bottom row is* $(1, 15, 34, 12, 1)$*, raising any single entry by one leaves the matrix totally nonnegative. (iii) The admissible window is finite in both directions and the Whitney values sit strictly inside it: with the other entries fixed, the second entry may run over* $[7, 25]$ *and the third over* $[19, 55]$*, and one step outside either range destroys total nonnegativity. (iv) The resolvability identity is not vacuous: the certificate built for one value of k fails against the Whitney row of another. (v) The alternation-certificate checker rejects a truncated, a mis-signed and a reversed certificate. (vi) Kruskal–Katona has enormous slack on this family, unlike on* [12, Ex. 9.5]*: for $\gamma(B_{7,3}(2)) = (1, 29070, 1519306, 2323260)$ the top pseudopower bound is $\gamma_3 \leq 881\,698\,919$, sharp at that value, a factor of about* 380 *above the actual γ_3, whereas* $(1, 7, 5, 1)$ *sits one unit under its own bound.*

Control (vi) is the honest reading of Proposition 10.2: the Kruskal–Katona verdicts there say nothing about how tight [12, Thm. 1.7] is. The example that does say something about that is the one in §3.

11. Projective deletions: what the Chow polynomial cannot see

11.1. **Flat spans.** Fix a field K, a finite-dimensional K-vector space V, and a set D of points of $P(V)$. Write $E = P(V) \setminus D$ for the surviving points and $M = M(D)$ for the simple matroid on E induced by linear dependence. Following Cheng and Liu [10, §2], call a subspace $U \leq V$ a *flat span* of D if

$$\operatorname{span}_K\big(E \cap P(U)\big) \;=\; U,$$

that is, if the points of U that survive still span it. The flats of M are exactly the sets $E \cap P(U)$ for flat spans U, and $\operatorname{rk}\big(E \cap P(U)\big) = \dim U$; so the lattice of flats of M is the poset of flat spans ordered by inclusion, graded by dimension. For $D = P(W)$ with $\dim W = k$ this is the Bose–Burton geometry $B_{r,k}(q)$ of §8, whose flat spans of dimension ≥ 1 are, by Lemma 8.1, exactly the subspaces not contained in W. Theorems 11.1 and 11.2 and Corollary 11.3 do not use finiteness of K, and there $B_{r,k}(q)$ is to be read as the deletion $D = P(W)$; Theorem 11.4 and Corollary 11.5 count points on a line and so are statements about $|K| = q$.

For a surviving point $p \in E$ put

$$\delta(p) \;=\; \#\{\text{lines } L \ni p \,:\, P(L) \setminus \{p\} \subseteq D\}, \tag{10}$$

the number of projective lines through p whose other q points are all deleted. The condition $\delta \equiv 0$ says that every line through a surviving point is spanned by its surviving points; the number of rank-2 flats of M above the atom $\{p\}$ is $\left[{r-1 \atop 1}\right]_q - \delta(p)$.

11.2. **The chain sum is blind to rank one.** The first theorem is not about geometry at all. Let F be a finite set carrying a rank function $\operatorname{rk} : F \to \mathbb{N}$ and a strict order $<$ with $x < y \Rightarrow \operatorname{rk} x < \operatorname{rk} y$, and let $w : \mathbb{N} \to R$ be any weight with values in a commutative semiring. Define

$$C(U) \;=\; \begin{cases} 1 & \operatorname{rk} U = 0, \\ \displaystyle\sum_{x \in F,\; x < U} C(x)\, w(\operatorname{rk} U - \operatorname{rk} x) & \text{otherwise,} \end{cases} \qquad \Sigma(F) \;=\; \sum_{U \in F} C(U). \tag{11}$$

We write C_F when the ambient set has to be named. For F the lattice of flats of a matroid, $w(m) = t + t^2 + \cdots + t^{m-1}$ and rk the matroid rank, $\Sigma(F)$ is the chain sum (6), so $\Sigma(F) = \underline{\mathcal{H}}_M$.

Theorem 11.1. *Suppose $w(1) = 0$. Then $C(U) = 0$ for every U of rank 1; and if F_1, F_2 are two such sets inside a common ranked poset which contain the same elements of every rank $\neq 1$, then $C_{F_1}(U) = C_{F_2}(U)$ for every U, and $\Sigma(F_1) = \Sigma(F_2)$. In particular, two sets D, D' of deleted points of $P(V)$ with the same flat spans in every dimension ≥ 2 have the same Chow polynomial,*

$$\underline{\mathcal{H}}_{M(D)} \;=\; \underline{\mathcal{H}}_{M(D')},$$

with no hypothesis whatever on D or D'.

Proof. If $\operatorname{rk} U = 1$ then every $x < U$ has $\operatorname{rk} x = 0$, so $C(x) = 1$ and the step weight is $w(\operatorname{rk} U - \operatorname{rk} x) = w(1) = 0$; hence $C(U) = 0$. For the second claim, induct on $\operatorname{rk} U$. The sum defining $C(U)$ may be restricted to the $x < U$ of rank $\neq 1$, since the rank-1 terms vanish by the first claim, and those x lie in F_1 exactly when they lie in F_2, with the same order and rank; by induction their C-values agree. Summing over U and again dropping the rank-1 terms gives $\Sigma(F_1) = \Sigma(F_2)$. For the last claim take F_i to be the sets of flat spans, ordered by inclusion and graded by dimension, and $w(m) = t + \cdots + t^{m-1}$, for which $w(1)$ is the empty sum 0; the flat spans of dimension 0 agree, both being $\{0\}$. $\square$

Concretely: a flat of rank 1 sits at rank gap 1 above $\hat{0}$, so it contributes weight 0 to every chain, and the chain formula never reads an atom. The degree-1 shadow of Theorem 11.1 — that the linear coefficient of the Chow polynomial of a poset is the number of its elements of rank at least two — is used in print, as an evident step and without proof, inside the proof of [11, Cor. 1.2]; we have not found the statement itself anywhere. Its usefulness is that it is exactly the hypothesis a transfer statement needs, and it is what the next theorem transfers along.

11.3. **Moving the deleted points.**

Theorem 11.2. *Let $W \le V$ with $\dim W = k$, and let $D \supseteq P(W)$ be a set of deleted points such that*

(i) every projective line not contained in W retains a surviving point, and

(ii) every projective line through a surviving point is spanned by its surviving points, i.e. $\delta \equiv 0$.

Then for every subspace U with $\dim U \ge 2$,

$$U \text{ is a flat span of } D \iff U \not\le W,$$

which is the same condition as for $B_{r,k}(q)$. Consequently $\underline{\mathcal{H}}_{M(D)} = \underline{\mathcal{H}}_{B_{r,k}(q)}$ identically, and for every U and every $d \ge 2$ the two lattices have literally the same flat spans of dimension d above U; in particular $\mathcal{L}(M(D))$ is upper combinatorially uniform, with the same upper intervals as $\mathcal{L}(B_{r,k}(q))$.

Proof. If $U \le W$ then $E \cap P(U) = \emptyset$, because $P(W) \subseteq D$, so the span of the survivors of U is $0 \ne U$ and U is not a flat span; the same argument applies to $B_{r,k}(q)$.

Let $U \not\le W$ with $\dim U \ge 2$. Choose $u \in U \setminus W$ and then $v \in U \setminus \langle u \rangle$, possible since $\dim U \ge 2$. The line $\langle u, v\rangle$ lies in U and is not contained in W, so by (i) it retains a surviving point $y \in E \cap P(U)$. Now let $x \in U$ be arbitrary. The line $\langle y, x\rangle$ lies in U and passes through the survivor y, so by (ii) it is spanned by its surviving points, all of which lie in $E \cap P(U)$; hence $x \in \operatorname{span}(E \cap P(U))$. Therefore U is a flat span. Note that (i) was used exactly once, to produce the single survivor y; no induction on $\dim U$ is needed.

The Bose–Burton geometry satisfies (i) and (ii): a line not contained in W meets W in at most one point, so retains at least $q \ge 2$ points, and two distinct points of a line span it; and a line through a point outside W is not contained in W. So the flat spans of dimension ≥ 2 of D and of $P(W)$ coincide, and Theorem 11.1 gives the identity of Chow polynomials. The statement about upper intervals is the displayed equivalence read above a fixed U, and $\mathcal{L}(B_{r,k}(q))$ is upper combinatorially uniform by Lemma 8.1; the flats of rank ≥ 1 of $M(D)$ are indexed by the same subspaces as those of $B_{r,k}(q)$ once $d \ge 2$, which is all that an upper interval of a nonzero flat sees. ☐

Theorem 11.1 contains the identity $\underline{\mathcal{H}}_{B_{r,1}(q)} = \underline{\mathcal{H}}_{PG(r-1,q)}$ of Proposition 8.4(iii), which §8 could only check on a table: deleting a single projective point leaves every subspace of dimension ≥ 2 a flat span, exactly as the empty deletion does, so the two agree outside rank 1. It also explains the identity — deleting one projective point changes only the atoms, and the chain sum never reads an atom — and it holds for every rank and every field.

Hypothesis (ii) is not necessary for the conclusion, and it is important for §13 that it is not.

Corollary 11.3. *Let $W \le V$ with $\dim W = 2$ — a projective line — let p be one of its points and let $D = P(W) \setminus \{p\}$. Then*

$$\underline{\mathcal{H}}_{M(D)} \;=\; \underline{\mathcal{H}}_{B_{r,2}(q)}$$

identically, although $M(D)$ has one point more than $B_{r,2}(q)$ and $\delta(p) \geq 1$, so that hypothesis (ii) of Theorem 11.2 fails.

Proof. Since $D \subseteq P(W)$, every $U \not\leq W$ is a flat span of D: pick $y \in U \setminus W$, which survives; a deleted $x \in U$ lies in W, so $x+y \notin W$ survives, and $x = (x+y)-y$ lies in the span of the survivors of U. Conversely the only subspace of W of dimension 2 is W itself, whose survivors are the points of $\langle p \rangle$ alone, spanning a line of dimension 1; so W is not a flat span. The flat spans of dimension ≥ 2 are therefore the $U \not\leq W$, exactly as for $B_{r,2}(q)$, and Theorem 11.1 applies. Finally the line W passes through the survivor p and has no other survivor, so $\delta(p) \geq 1$. □

For $q = 2$ a projective line has three points, so Corollary 11.3 deletes two points from $PG(r-1, 2)$; for $q = 3$ it deletes a full affine line. In both cases the lattice of flats of $M(D)$ is *not* upper combinatorially uniform — the atom p carries one rank-2 flat fewer than the others — and yet the Chow polynomial is a Bose–Burton value. This is why failure of upper rank uniformity, on its own, certifies nothing about novelty, and why the 21 instances of Theorem H had to be compared with the Bose–Burton values one by one.

11.4. **When the hypotheses hold, and when they cannot.** We now specialise to the shape of deletion that occurs in [10]: $V = W \oplus X$ and $D = P(W) \sqcup S$ with S a set of points of $P(X)$, $s = |S|$. Call S a *cap* if no projective line carries three of its points.

Theorem 11.4. *Let $V = W \oplus X$ and $D = P(W) \sqcup S$ with $S \subseteq P(X)$.*

- *(i) If $q \geq 3$ and S is a cap, then every projective line not contained in W retains at least two surviving points. Hence hypotheses (i) and (ii) of Theorem 11.2 hold and $\underline{\mathcal{H}}_{M(D)} = \underline{\mathcal{H}}_{B_{r,k}(q)}$ identically.*
- *(ii) Conversely, if $q \geq 3$ and some projective line not contained in W retains at most one surviving point, then S carries three collinear points.*
- *(iii) If $q = 2$, then as soon as $k \geq 1$ and $s \geq 1$, or $k = 0$ and $s \geq 2$ with S inside an affine chart of $P(X)$, the lattices of flat spans of $M(D)$ and of $B_{r,k}(2)$ differ already in dimension 2; in particular $\delta \not\equiv 0$ and Theorem 11.2 never applies.*

Proof. (i) Let L be a line with $L \not\leq W$; it has $q+1 \geq 4$ points. At most one of them lies in W, since $L \cap W$ is a subspace of dimension ≤ 1. At most two lie in S, since S is a cap. Moreover if some point of L lies in W then at most *one* lies in S: two points of S on L would force $L \leq X$, and $L \cap W \leq X \cap W = 0$. In every case at most two points of L are deleted, so at least $q-1 \geq 2$ survive. Hypothesis (i) of Theorem 11.2 is immediate; for (ii), a surviving point lies outside W, so every line through it is not contained in W, retains two survivors, and two distinct points of a line span it.

(ii) Suppose $L \not\leq W$ retains at most one survivor, so at least q of its $q+1$ points are deleted. If L met W, the count in (i) would leave at least $q-1 \geq 2$ survivors; so $L \cap W = 0$ and all the deleted points of L lie in S. There are at least $q \geq 3$ of them, and they are collinear.

(iii) Let $q = 2$, so that $|K| = 2$ and a projective line has three points. If $k \geq 1$ and $s \geq 1$, take $w \in W \setminus 0$ and $y \in S \setminus 0$; then $y \notin W$ because $S \subseteq X$ and $W \cap X = 0$, so the line $\langle w, y \rangle$ is not contained in W and is a flat span of $B_{r,k}(2)$. Its third point $w + y$ lies neither in W (else $y \in W$) nor in $X \supseteq S$ (else $w \in X$), so it is its only survivor, and the line is not a flat span of $M(D)$. If $k = 0$ and $s \geq 2$, with S inside the chart $P(X) \setminus P(H)$ for a hyperplane $H < X$, take two distinct $y, y' \in S$; then $y + y' \in H$ because H is the kernel of a linear form taking the value 1 at y and at y', so $y + y' \notin S$ and it survives, while y, y' do not: again the line $\langle y, y' \rangle$ has exactly one survivor. In both cases the survivor p of that line has $\delta(p) \geq 1$. ☐

The chart form of the hypothesis, which is the one stated in [10, §2] for the existence of S, also suffices, and for $q \geq 4$ it is strictly weaker: a set may carry three collinear points without carrying a whole affine line of q points.

Corollary 11.5. *Let $q \geq 3$, let $V = W \oplus X$, let $H < X$ be a hyperplane, and let S be a set of points of the affine chart $P(X) \setminus P(H)$ containing no full affine line of that chart. Then every projective line not contained in W retains at least two surviving points, so $\underline{\mathcal{H}}_{M(D)} = \underline{\mathcal{H}}_{B_{r,k}(q)}$ identically; and conversely, if some line not contained in W retains at most one survivor, then S contains a full affine line of the chart.*

Proof. Let $L \not\leq W$ retain at most one survivor. The second observation in the proof of Theorem 11.4(i) does not use the cap hypothesis: if L met W then at most one of its points would lie in S, and at least $q - 1 \geq 2$ would survive. So $L \cap W = 0$, and at least $q \geq 3$ of its points lie in $S \subseteq X$, whence $L \leq X$. Since H is a hyperplane of X and $L \not\leq H$ — the points of S on L lie outside H — the intersection $L \cap H$ is a single point, and the remaining q points of L are the points of an affine line of the chart, all of them deleted and hence in S. That contradicts the hypothesis, and the converse is the same computation read backwards. ☐

Part (iii) of Theorem 11.4 is a genuine obstruction and not an artefact of the proof: over $\mathbb{F}_2$ a line has only three points, and removing any two of them from a line outside W leaves a single survivor. It is the reason Theorem G asks for $q \geq 3$, and the reason the binary counterexamples of [10] themselves fall on the other side of the dichotomy (Remark 12.7).

11.5. **Calibration, and controls.** Theorems 11.1–11.4 are proved in the formal development over an arbitrary field, with the chain sum (11) written over an abstract finite ranked poset. What has to be checked separately is that this abstract chain sum really is the Chow polynomial this paper has been computing since §8.

Proposition 11.6. *Evaluated in the kernel over lattices of flat spans computed from the definition, the chain sum* (11) *with $w(m) = t + \cdots + t^{m-1}$ reproduces the recursion* (7) *of Proposition 8.4 in two characteristics:*

D	*matroid*	*flats*	*Σ at $t = 1, 2, 3$*	*target*
$\emptyset$	$PG(2,2)$	16	$10, 21, 34$	$1 + 8t + t^2$
a point	$B_{3,1}(2)$	15	$10, 21, 34$	*the same*
two points	—	13	$9, 19, 31$	$1 + 7t + t^2$
a line	$AG(2,2)$	12	$9, 19, 31$	*the same*
$\emptyset$	$PG(2,3)$	28	$16, 33, 52$	$1 + 14t + t^2$
a point	$B_{3,1}(3)$	27	$16, 33, 52$	*the same*
two points of a line	—	26	$16, 33, 52$	*the same*
three points of a line	—	24	$15, 31, 49$	$1 + 13t + t^2$
a line	$AG(2,3)$	23	$15, 31, 49$	*the same*

and over the Boolean lattice of rank 3 *it returns the Eulerian polynomial* $A_3 = 1 + 4t + t^2$*. The targets are the values of* (7) *itself, not transcribed numbers. The input data is verified rather than asserted: the* 16 *listed subspaces of* $\mathbb{F}_2^3$ *are proved to be exactly the* XOR*-closed subsets of its seven nonzero vectors; the flat spans are computed from the definition; the hypothesis* $x < y \Rightarrow \operatorname{rk} x < \operatorname{rk} y$ *is checked rather than assumed; and the* 13 *lines listed at* $q = 3$ *are proved to be* 13 *four-point lines on* 13 *points with any two points on exactly one line, i.e. a projective plane of order* 3*, which is unique.*

Several rows are not merely computed but derived. In each block the single-point deletion is Theorem 11.2 applied, and the last two rows of each block are related by Corollary 11.3. For $q = 2$ both the identity between the one-point deletion and $PG(2, 2)$ and the identity between the two-point deletion and the line deletion are obtained from Theorem 11.1, its hypothesis discharged by exhaustive comparison of the two lattices; so they are identities in $\mathbb{Z}[t]$, not merely agreements at three values of t. The third row of the $q = 3$ block is the boundary the theorem draws: leave two survivors on the deleted line and it is still a flat span, so nothing moves at all. At $q = 3$ the computation is affordable only because Theorem 11.1 allows the rank-1 flats to be dropped first, which cuts the poset from 28 elements to 15 — the theorem paying for its own verification. The identification of the bitmask model of $\mathbb{F}_q^3$ used here with the abstract subspace lattice over which Theorems 11.1–11.4 are proved is mathematics, not a formal theorem; both layers are separately machine-checked, and only their identification is by hand.

Proposition 11.7. *The following controls hold. (i) The hypothesis* $w(1) = 0$ *in Theorem 11.1 is load-bearing: with* $w(m) = m$ *the conclusion fails on the data of Proposition 11.6, the two sides being* 95 *and* 86*. (ii) The conclusion is not universally true:* $PG(2, 2)$ *and* $AG(2, 2)$ *have different Chow polynomials,* 10 *against* 9 *at* $t = 1$*. (iii) The hypothesis of Theorem 11.1 is strictly weaker than equality of lattices: the two-point deletion and the line deletion of* $PG(2, 2)$ *have* 13 *and* 12 *flat spans and equal Chow polynomials. (iv)* $\delta \not\equiv 0$ *is realised, both in rank* 3 *by computation and in every rank by Theorem 11.4(iii). (v) The hypotheses of Theorem 11.2 are non-vacuous on both sides:* $B_{r,k}(q)$ *satisfies them, a single deleted projective point over a field with at least three elements satisfies them, and the same one-point perturbation over* $\mathbb{F}_2$ *breaks them. (vi) A subspace contained in* W *is never a flat span.*

Proposition 11.8. *Theorems 11.2, 11.4 and Corollary 11.3 were also tested against lattices enumerated in full, for* $q \in \{2, 3, 5\}$ *and ranks* 4*,* 5 *and* 7*. On all ten instances with* $\delta \equiv 0$ *the identity* $\underline{\mathcal{H}}_{M(D)} = \underline{\mathcal{H}}_{B_{r,k}(q)}$ *holds while the Kazhdan–Lusztig polynomial moves by* $s\,t$*; on the three instances of Corollary 11.3, where* $\delta \not\equiv 0$*, it holds as well; and on all five further instances with* $\delta \not\equiv 0$ *it fails. The thresholds of Theorem 11.4 are matched exactly by that data: at* $q = 2$ *the identity breaks at the first perturbation, at* $q = 3$ *at a collinear triple but not at a generic triple, and at* $q = 5$ *at five collinear points.*

The three instances of Corollary 11.3 verified there are

$$\underline{\mathcal{H}}_{PG(4,2)\setminus\{x,y\}} = \underline{\mathcal{H}}_{B_{5,2}(2)}, \quad \underline{\mathcal{H}}_{PG(6,2)\setminus\{x,y\}} = \underline{\mathcal{H}}_{B_{7,2}(2)}, \quad \underline{\mathcal{H}}_{PG(4,3)\setminus\ell} = \underline{\mathcal{H}}_{B_{5,2}(3)}$$

for a full affine line ℓ; in each case the lattice of flats is not upper combinatorially uniform.

12. A refutation that does not transfer

12.1. **The Cheng–Liu construction.** Cheng and Liu [10] refute unimodality — hence log-concavity, hence real-rootedness — of matroid Kazhdan–Lusztig polynomials over every finite field. Their matroids are projective deletions in the sense of §11. We quote the two statements we use, in their numbering.

Definition 12.1 ([10, Def. 2.5])**.** $S \subseteq P(V) \setminus P(W)$ is *line-separated from* W if every projective line contained in $P(W) \cup S$ is contained in $P(W)$.

Proposition 12.2 ([10, Prop. 2.6])**.** *Let $0 \le k = \dim W < r/2$, let S be line-separated from W with $|S| = s$, put $D = P(W) \cup S$ and suppose $|D| = \left[\begin{smallmatrix} k \\ 1 \end{smallmatrix}\right]_q + s < q^{r-2}$. Then the Kazhdan–Lusztig polynomial of $M = P(V) \setminus D$ is*

$$P_M(t) \;=\; \sum_{j=0}^{k} \left[\begin{smallmatrix} k \\ j \end{smallmatrix}\right]_q t^j \;+\; s\,t.$$

At $s = 0$ the right-hand side is the row of Gaussian binomials, which is unimodal; the whole failure comes from S. Taking $k = 6$ and $s = \left[\begin{smallmatrix} 6 \\ 2 \end{smallmatrix}\right]_q - \left[\begin{smallmatrix} 6 \\ 1 \end{smallmatrix}\right]_q + 1$ makes the coefficients of t, t^2, t^3 equal to $\left[\begin{smallmatrix} 6 \\ 2 \end{smallmatrix}\right]_q + 1$, $\left[\begin{smallmatrix} 6 \\ 2 \end{smallmatrix}\right]_q$ and $\left[\begin{smallmatrix} 6 \\ 3 \end{smallmatrix}\right]_q$, and $\left[\begin{smallmatrix} 6 \\ 3 \end{smallmatrix}\right]_q > \left[\begin{smallmatrix} 6 \\ 2 \end{smallmatrix}\right]_q$, so there is a strict local valley at index 2; this is [10, Cor. 1.3], whose displayed instance is $q = 2$, $s = 589$, rank 17, with

$$P_M \;=\; 1 + 652t + 651t^2 + 1395t^3 + 651t^4 + 63t^5 + t^6.$$

Their existence argument places S inside an affine chart $P(X) \setminus P(H)$ of a complement X of W: any subset of such a chart is line-separated from W.

12.2. **Caps.** The construction we need differs from theirs in one point: S must be a cap.

Lemma 12.3. *For every prime power q the parabola $\{(a, a^2) : a \in \mathbb{F}_q\}$ is a cap of $AG(2, q)$, of size q; $AG(2, 3)$ contains the cap $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$ of size 4; and if $AG(2, q)$ contains a cap of size c, then $AG(2t, q)$ contains one of size c^t.*

Proof. Three distinct points $(a, a^2), (b, b^2), (c, c^2)$ are collinear if and only if the Vandermonde determinant $(b - a)(c - a)(c - b)$ vanishes, which it does not; the computation involves no division and is valid in every characteristic. The four listed points of $AG(2, 3)$ contain no three collinear, by inspection. For the last claim it is enough to know that a product of caps is a cap, and then to induct on t along $AG(2t, q) = AG(2, q) \times AG(2t - 2, q)$. So let $S_1 \subseteq A_1$ and $S_2 \subseteq A_2$ be caps in affine spaces over $\mathbb{F}_q$ and suppose three distinct points of $S_1 \times S_2$ lie on an affine line ℓ of $A_1 \times A_2$. The image $\pi_1(\ell)$ is either a line, in which case π_1 is injective on ℓ and the three points project to three distinct collinear points of S_1; or a single point, in which case ℓ lies in a fibre and the three points project to three distinct collinear points of S_2. Either way one of the factors fails to be a cap. $\square$

Lemma 12.3 was also verified by exhaustive search over $\mathbb{F}_q$ itself, in a polynomial basis, for all seventeen prime powers q in the range $3 \le q \le 32$. We record a trap that the search caught: an implementation working in $\mathbb{Z}/q\mathbb{Z}$ rather than in $\mathbb{F}_q$ reports the parabola as *not* a cap at $q = 8$. That is the arithmetic being wrong, not the geometry.

12.3. **The theorem.**

Theorem 12.4. *Let $q \geq 3$ be a prime power. Put*

$$k = 6, \qquad s = \begin{bmatrix} 6 \\ 2 \end{bmatrix}_q - \begin{bmatrix} 6 \\ 1 \end{bmatrix}_q + 1, \qquad r = 25,$$

let $V = \mathbb{F}_q^{25} = W \oplus X$ with $\dim W = 6$, let $H < X$ be a hyperplane, and let S be a cap of size s inside the affine chart $P(X) \setminus P(H) \cong AG(18, q)$. Put $D = P(W) \sqcup S$ and $M = PG(24, q) \setminus D$. Then

(i) $P_M(t) = \sum_{j=0}^{6} \begin{bmatrix} 6 \\ j \end{bmatrix}_q t^j + s\, t$ is not *unimodal, with a strict local valley at index 2;*

(ii) $\underline{\mathcal{H}}_M = \underline{\mathcal{H}}_{B_{25,6}(q)}$, identically; and

(iii) $\underline{\mathcal{H}}_M$ is real-rooted.

For $q = 3$ the same three conclusions hold at $r = 21$, with S a cap of size $s = 10\,648$ inside $AG(14, 3)$.

Proof. The cap exists: by Lemma 12.3, $AG(18, q)$ contains a cap of size q^9, and $s \leq \begin{bmatrix} 6 \\ 2 \end{bmatrix}_q < q^9$, because

$$\begin{bmatrix} 6 \\ 2 \end{bmatrix}_q = \frac{(q^6-1)(q^5-1)}{(q^2-1)(q-1)} < \frac{q^{11}}{(q^2-1)(q-1)} \leq \tfrac{27}{16}\, q^8 < q^9$$

for $q \geq 3$, using $(1 - q^{-2})(1 - q^{-1}) \geq \frac{16}{27}$; a subset of a cap is a cap. For $q = 3$ use instead the 4-point cap of $AG(2, 3)$: $4^7 = 16\,384 \geq 10\,648 = s$, so $AG(14, 3)$ carries a cap of size s.

(i) The hypotheses of Proposition 12.2 hold. Any subset of the chart is line-separated from W; $k = 6 < r/2$; and $|D| = \begin{bmatrix} 6 \\ 1 \end{bmatrix}_q + s < q^6 + q^9 < q^{\,r-2}$ since $r - 2 \geq 19$. So P_M is as displayed, and its coefficients at t, t^2, t^3 are $\begin{bmatrix} 6 \\ 2 \end{bmatrix}_q + 1 > \begin{bmatrix} 6 \\ 2 \end{bmatrix}_q < \begin{bmatrix} 6 \\ 3 \end{bmatrix}_q$: a strict local valley, hence non-unimodality ([10, Cor. 1.3]).

(ii) S is a cap contained in X, and $W \cap X = 0$, so Theorem 11.4(i) applies: $\underline{\mathcal{H}}_M = \underline{\mathcal{H}}_{B_{r,6}(q)}$.

(iii) Theorem D, that is Corollary 10.1, applied to $B_{25,6}(q)$ — for which $2 \leq k \leq r - 2$ — gives real-rootedness. We do not claim that the roots are simple; the sign-alternation certificates of Proposition 10.2 give simplicity only for the tabulated instances, which stop at rank 8. ☐

The single rank $r = 25$ serves every $q \geq 3$ because the estimate above is uniform in q: the product cap of Lemma 12.3 fills $AG(18, q)$ with q^9 points and $s \leq \begin{bmatrix} 6 \\ 2 \end{bmatrix}_q$ stays below that for every q, the two sides being closest at $q = 3$, where $s = 10\,648$ against $q^9 = 19\,683$. The columns base and t of the table record which cap is used: base $= q$ and $t = 9$ on every row but the first, where the four-point cap of $AG(2, 3)$ is denser than the parabola and $4^7 = 16\,384 \geq s$ lowers the rank to 21. We do not claim these ranks are least possible — any denser cap of $AG(2, q)$ would lower them — only that one rank works for every prime power $q \geq 3$.

The parameters are tabulated for seventeen prime powers, and for each of them the following are decided by kernel computation: that $s = \begin{bmatrix} 6 \\ 2 \end{bmatrix}_q - \begin{bmatrix} 6 \\ 1 \end{bmatrix}_q + 1$; that $m = 2t$, $r = 7 + m$, $2k < r$, $\begin{bmatrix} 6 \\ 1 \end{bmatrix}_q + s < q^{\,r-2}$, $s \leq \mathrm{base}^{\,t}$ and $\mathrm{base}^{\,t} \leq q^m$; that P_M is not unimodal, with the coefficients of t, t^2, t^3 equal to $\begin{bmatrix} 6 \\ 2 \end{bmatrix}_q + 1$, $\begin{bmatrix} 6 \\ 2 \end{bmatrix}_q$, $\begin{bmatrix} 6 \\ 3 \end{bmatrix}_q$; and that the resolvability certificate of Theorem 9.2 holds for $B_{r,6}(q)$ at exactly the rank r

used, that is at $r = 21$ for $q = 3$ and $r = 25$ for the other sixteen. Cheng and Liu's own headline value at $q = 2$, $s = 589$, is reproduced in the kernel as well.

q	$\begin{bmatrix}6\\1\end{bmatrix}_q$	s	base	t	m	r	$\lvert D\rvert = \begin{bmatrix}6\\1\end{bmatrix}_q + s$
3	364	10 648	4	7	14	21	$11\,012 < 3^{19}$
4	1 365	91 729	4	9	18	25	$93\,094 < 4^{23}$
5	3 906	504 526	5	9	18	25	$508\,432 < 5^{23}$
7	19 608	6 845 644	7	9	18	25	$6\,865\,252 < 7^{23}$
8	37 449	19 440 193	8	9	18	25	$19\,477\,642 < 8^{23}$
9	66 430	48 965 554	9	9	18	25	$49\,031\,984 < 9^{23}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	9	18	25	$\vdots$
32	34 636 833	1 136 054 535 169	32	9	18	25	$1\,136\,089\,172\,002 < 32^{23}$

Remark 12.5. The value $k = 6$ is forced, for every s and not only for the one chosen here. Write a_j for the coefficients of the polynomial of Proposition 12.2: so $a_0 = 1$, $a_1 = \begin{bmatrix}k\\1\end{bmatrix}_q + s$, and $a_j = \begin{bmatrix}k\\j\end{bmatrix}_q$ for $j \geq 2$. A failure of unimodality is a valley, indices $i < j < l$ with $a_i > a_j < a_l$. Here $a_0 = 1$ is minimal and only the entry at index 1 has been raised, so any valley has $j \geq 2$; and $a_j < a_l$ for some $l > j \geq 2$ forces the row of Gaussian binomials to increase somewhere from index 2 on, that is $\begin{bmatrix}k\\2\end{bmatrix}_q < \begin{bmatrix}k\\3\end{bmatrix}_q$, which holds exactly for $k \geq 6$. For $k \leq 5$ that row is nonincreasing from index 2 on — already at $k = 5$ one has $\begin{bmatrix}5\\2\end{bmatrix}_q = \begin{bmatrix}5\\3\end{bmatrix}_q$ — so the polynomial is unimodal for every s. A machine check confirms this for the recipe $s = \begin{bmatrix}k\\2\end{bmatrix}_q - \begin{bmatrix}k\\1\end{bmatrix}_q + 1$ at $k = 4$ and $k = 5$ over every q in the table above. Cheng and Liu take $k = 6$ without remarking that it is minimal; it is.

Remark 12.6. The price of arranging $\delta \equiv 0$ is rank. Proposition 12.2 needs only s points somewhere in a chart of q^{r-k-1} points, so at $q = 3$ its smallest instance has $r = 16$. A cap of size s in an affine space over $\mathbb{F}_q$ needs, by Lemma 12.3, about twice as many dimensions as an arbitrary set of the same size, and Theorem 12.4 lands at $r = 21$. That gap is the whole cost of making the perturbation invisible to the Chow polynomial.

Remark 12.7. The construction of [10] is at its sharpest over $\mathbb{F}_2$, and there Theorem 12.4 has nothing to say. By Theorem 11.4(iii) a binary instance of their construction has a line carrying a single survivor as soon as $k \geq 1$ and $s \geq 1$, or $k = 0$ and $s \geq 2$ in a chart; their counterexamples have $k = 6$, so hypothesis (ii) of Theorem 11.2 fails for every one of them and the Chow polynomial is not forced onto a Bose–Burton value. The one binary perturbation that escapes is a single deleted point with $W = 0$, for which Theorem 11.2 does apply and the Chow polynomial stays $\underline{\mathcal{H}}_{PG(r-1,2)}$; it is the row $(k, s) = (0, 1)$ of §13. In particular we do not know whether the Chow polynomial of their rank-17 counterexample is real-rooted. What fails there is the hypothesis of Theorem 11.2; whether the lattice of flats of that matroid is upper rank uniform we have not determined, the measurements of §13 being at rank 7. At rank 7 the same construction is computable, and §13 reports what happens there.

Remark 12.8. Two inputs to Theorem 12.4 are cited and not reproved: Proposition 12.2 together with [10, Cor. 1.3], and, through Theorem D, [8, Thm. 6.1]. Both were exercised against our own computations rather than taken on trust: Proposition 12.2 is reproduced from its own definitions on all 26 lattices of §13, and the polynomial displayed above at $s = 589$ is recomputed in the kernel.

12.4. **Prior work.** We looked for the combination in the literature and did not find it. In a corpus of arXiv mathematics, of the 43 papers containing the phrase "Chow polynomial", none contains "Bose", "projective deletion" or "line-separated", and none cites [10]; the word "Chow" does not occur anywhere in [10]; and the two papers we found citing [10], arXiv:2608.02303 and arXiv:2608.13836, do not contain the word at all. Nor does any of [12, 15, 8, 13] contain a statement to the effect that $\underline{\mathcal{H}}_M$ depends on the deleted points only through the flats of rank at least two; the invariance statements there are the standard one, that $\underline{\mathcal{H}}_M$ depends only on the lattice of flats. What we claim from this is not that the question was asked and left open — as far as we can tell it was not asked — but that Theorems 11.1–12.4 combine published statements that had not been combined.

13. The rank-7 census

Theorem 11.4(iii) puts every binary Cheng–Liu deletion on the side of the dichotomy where Theorem 11.2 is silent. Rank 7 is the first rank on that side which is not settled by [12, Cor. 8.5] and is still within reach of a direct enumeration: $\mathbb{F}_2^7$ has 29 212 subspaces, and rank 8 has 417 199.

Proposition 13.1. *Let $q = 2$, $r = 7$, $0 \leq k \leq 3$, let $W \leq \mathbb{F}_2^7$ have dimension k, let H be a hyperplane of a complement X, and let S consist of the first s points of the chart $P(X) \setminus P(H)$ in a fixed enumeration. For each of the 26 pairs (k, s) of the table below, the lattice of flat spans of $D = P(W) \sqcup S$ — up to 29 211 flats — was enumerated in full and the Chow polynomial computed from (6) over that lattice. Every one of the 26 resulting polynomials is palindromic, its γ-vector satisfies (1), is nonnegative, satisfies the Kruskal–Katona inequalities (4) and Newton's inequalities, and passes Kurtz's criterion; and on every row the subspace polynomial computed from the lattice agrees with Proposition 12.2.*

The enumeration was run twice, by two independently written programs, agreeing on all 130 compared quantities — the coefficient vector, the γ-vector, the number of flats, the size of the ground set, upper rank uniformity, real-rootedness and the subspace polynomial, on each of the 26 rows. It cost about 2 300 seconds of processor time in total. As a further calibration, three values it produces are values this paper computed earlier by the unrelated recursion (7): the coefficient vector and γ-vector of $B_{7,3}(2)$ and, twice, the coefficient vector of $PG(6, 2)$. The γ-positivity and Kruskal–Katona halves of the proposition are instances of [12, Thms. 1.6 and 1.7], the maximal building set being complete, and are re-derivations; the computation and the real-rootedness are not.

k	s	$\lvert E\rvert$	flats	γ_1	γ_2	γ_3	u.r.u.	Kurtz
3	0	120	29197	29070	1519306	2323260	yes	8.545
3	1	119	29189	29063	1516940	2316960	no	8.543
3	2	118	29180	29055	1514236	2309760	no	8.542
3	3	117	29170	29046	1511194	2301660	no	8.540
3	4	116	29158	29035	1507766	2292660	no	8.538
3	5	115	29146	29024	1504048	2282760	no	8.536
3	6	114	29131	29010	1499896	2271960	no	8.533
3	7	113	29113	28993	1495310	2260260	no	8.530
3	8	112	29090	28971	1490230	2247624	no	8.526
2	0	124	29208	29077	1521382	2328660	yes	8.546
2	1	123	29204	29074	1520368	2325960	no	8.545
2	2	122	29199	29070	1519016	2322360	no	8.545
2	4	120	29185	29058	1515250	2312460	no	8.542
2	8	116	29133	29010	1503122	2281824	no	8.533
2	12	112	29052	28933	1484998	2236752	no	8.519
2	16	108	28853	28738	1457470	2176380	no	8.491
1	0	126	29211	29078	1521720	2329560	yes	8.546
1	2	124	29206	29075	1520706	2326860	no	8.546
1	8	118	29152	29027	1508868	2297124	no	8.536
1	20	106	28757	28644	1441212	2139336	no	8.474
1	$32^{\dagger}$	94	26728	26627	1272982	1831140	no	8.309
0	1	126	29211	29078	1521720	2329560	yes	8.546
0	4	123	29201	29071	1519644	2324160	no	8.545
0	16	111	28905	28787	1474032	2220480	no	8.498
0	$48^{\dagger}$	79	22780	22694	951018	1247940	no	7.984
0	$64^{\dagger\ddagger}$	63	2825	2756	38844	—	yes	48.885

$\gamma_0 = 1$ throughout; the column "u.r.u." records whether the lattice of flats is upper rank uniform, and the last column is $\min_i \gamma_i^2/4\gamma_{i-1}\gamma_{i+1}$, the margin in Kurtz's criterion. Three rows are not certified instances of [10]. Those marked † have $|D| = \left[{k \atop 1}\right]_q + s \geq 2^{\,r-2} = 32$, so the size hypothesis of Proposition 12.2 fails and its formula for P_M is not established for them; the row marked ‡ deletes the whole chart, so that the survivors span only a hyperplane and the matroid is $PG(5,2)$ in rank 6. They are honest matroids with honest Chow polynomials and are kept as controls. The remaining 23 rows are certified Cheng–Liu deletions.

Two rows deserve separate mention because they are the dichotomy showing itself inside the table: $(k,s) = (1,0)$ and $(0,1)$ have the same Chow polynomial, namely that of $PG(6,2)$. The first is Proposition 8.4(iii) and the second is Theorem 11.2 for a one-point deletion. By $(0,4)$ the deleted points have created a line with a single survivor — already the second of them does, by Theorem 11.4(iii) — upper rank uniformity has failed, and the γ-vector has moved.

None of the 26 is a counterexample to unimodality, and this is a theorem rather than an omission: all 26 Kazhdan–Lusztig polynomials are unimodal. By Remark 12.5 non-unimodality in this construction requires $k \geq 6$, and then the hypothesis $k < r/2$ of Proposition 12.2 forces rank at least 13 — far past the reach of a direct enumeration. The census is evidence about the side of the dichotomy that no published framework reaches, not about non-unimodality; its rows must not be quoted as Cheng–Liu counterexamples.

Theorem 13.2. *Exactly* 21 *of the* 26 *deletions of Proposition 13.1 have a lattice of flats that is not upper rank uniform. Each of those* 21 *Chow polynomials is real-rooted, with six distinct real roots, certified by a sign alternation on seven rational points. None of the* 21 *equals* $\underline{\mathcal{H}}_{B_{7,k}(2)}$ *for any* $0 \le k \le 6$ *or* $\underline{\mathcal{H}}_{B_{6,k}(2)}$ *for any* $0 \le k \le 5$*, while each of the remaining five rows equals one of those values.*

Proof sketch. Rank uniformity is decided from the enumerated lattice, flat by flat, for each of the 26 rows; it holds on the five rows marked "yes" in the table and fails on the other 21. The γ-vectors are the ones tabulated, pinned by (1), and the real-rootedness certificates are the sign alternations of §7, one list of seven rationals per row, checked in the kernel on the homogenised polynomial. The comparison with the Bose–Burton values is exhaustive: each of the 21 coefficient vectors is compared with all thirteen values $\underline{\mathcal{H}}_{B_{7,k}(2)}$, $0 \le k \le 6$, and $\underline{\mathcal{H}}_{B_{6,k}(2)}$, $0 \le k \le 5$, computed by (7), and no pair is equal; the five upper rank uniform rows all match, as Theorem 11.2 predicts. □

Nineteen of the 21 are certified Cheng–Liu deletions; the other two are the rows $(k, s) = (1, 32)$ and $(0, 48)$, which break the size hypothesis of Proposition 12.2 and so carry no claim about P_M, but are projective deletions with Chow polynomials like any other. These 21 instances of [12, Conj. 1.1] are, as far as we could determine, not reached by any published sufficient condition. Rank 7 is past [12, Cor. 8.5]. The total-nonnegativity route is shut: a TN-poset is lower rank uniform, and by Brylawski's theorem [9, Thm. 4.5] a lower rank uniform geometric lattice is upper rank uniform, so $\mathcal{L}(M)$ is not a TN-poset; and $\mathcal{L}(M)^*$ being a TN-poset would require $\mathcal{L}(M)$ to be upper rank uniform directly. So neither [8, Thm. 6.1] nor [6] applies. The UMEL route is shut for the same reason: the letter U of UMEL is upper rank uniformity, so [13, Thm. 1.2] does not apply. Two more classes go with it: a perfect matroid design is lower rank uniform, hence upper rank uniform by [9, Thm. 4.5], and the lattice of flats of a uniform matroid is a perfect matroid design, so none of the 21 is either of those. The supersolvable production rules of [13, Thm. 1.8] end at supersolvable *rank-uniform* lattices, so they do not reach these matroids either, whatever their modular flats do; we did not test supersolvability of the perturbed matroids and do not use it. That leaves paving, the hypothesis of [8, Cor. 7.2], which upper rank uniformity does not settle; but the tabulated flat counts do. In a paving matroid of rank r every set of at most $r-2$ elements is a flat. Indeed a set of fewer than r elements is independent, since a dependent set contains a circuit and every circuit has at least r elements; so adjoining any element to a set of size at most $r-2$ leaves it independent and raises its rank, and the set is closed. A paving matroid of rank 7 on $|E|$ points therefore has $\binom{|E|}{5}$ flats of rank 5 alone, which exceeds 2×10^7 for every $|E| \ge 79$, against the at most 29 211 flats any row of the table has. Every one of the 21 has $|E| \ge 79$, the smallest being the row $(0, 48)$, so none of them is paving. Directly, and in agreement: on the row $(k, s) = (3, 1)$ the flat sizes are $\{2, 3\}$ in rank 2, $\{3, 4, 5, 6, 7\}$ in rank 3 and $\{7, 8, 11, 12, 13, 14, 15\}$ in rank 4, so a rank-2 flat there carries three points, a circuit of length 3.

Remark 13.3. One loophole remains open and we state it rather than closing it, as we did the two of Remark 9.4: [13, Thm. 1.6] proves real-rootedness for the *rank-selected subposets* of a UMEL-shellable poset, and a rank selection need not be rank uniform itself, so the failure of upper rank uniformity does not exclude one. We have not excluded $\mathcal{L}(M)$ being a rank selection of a larger UMEL-shellable poset, and we know of no published construction that would produce one.

Remark 13.4. The margin in Kurtz's criterion runs the wrong way for a counterexample hunt. Across the 21 rows of Theorem 13.2 it lies between 7.98 and 8.55, and it barely moves with s. For comparison, the Boolean matroids of ranks 7 and 8, whose Chow polynomials are the Eulerian polynomials A_7 and A_8, have $\gamma = (1, 114, 720, 272)$ and $\gamma = (1, 240, 3072, 3968)$, on which the same quantity is 4.18 and 2.48. These deletions therefore sit about twice as far inside the criterion as the Boolean matroid of the same rank, and the perturbation does not move them towards the boundary of real-rootedness.

Two margins should be kept apart here. The first is Kurtz's, the margin in his *sufficient* criterion: it is $\min_i \gamma_i^2/4\gamma_{i-1}\gamma_{i+1}$, the quantity just quoted and the one in the last column of the table. The margin in real-rootedness is a little larger, Kurtz's criterion not being necessary: with (γ_1, γ_2) held at the Boolean values, the cubic $1+\gamma_1 u+\gamma_2 u^2+\gamma_3 u^3$ still has three real roots for every γ_3 up to 1 170 at rank 7 and up to 10 109 at rank 8. Those are factors 4.30 and 2.55 above the Boolean γ_3, where Kurtz's bound allows only 4.18 and 2.48. The comparison above is the Kurtz reading throughout; the two readings agree on direction and nearly agree on size.

Proposition 13.5. *The following controls hold. (i) The certificate checker is not degenerate: for the row* $(k, s) = (3, 0)$ *it accepts the true certificate and rejects a truncated list, a reversed list, and the same list with one point moved into a region where no sign change occurs. (ii) The* γ*-vectors are pinned rather than fitted: changing either of the entries* γ_1, γ_3 *of that row by one breaks* (1)*. (iii) Kurtz's criterion is far from tight here: with* (γ_1, γ_2) *fixed at the values of that row it admits every* $\gamma_3 \leq 19\,851\,141$ *and rejects* $19\,851\,142$*, while the actual value is* $2\,323\,260$*. (iv) The resolvability certificate of Theorem 9.2 used in Theorem 12.4 is not vacuous at those sizes: pairing the resolving array of one deletion with the Whitney row of another fails, at rank* 21 *and at rank* 25*. (v) The unimodality predicate is non-vacuous on both sides. (vi) A too-large and a too-small claim are both refuted: the* 26 *rows do not all share one Chow polynomial, and two of them do.*

Control (iii) is the honest reading of the last column of the table: Kurtz's criterion holds with a factor of about 8.5 to spare on every row, so it is a proof of real-rootedness here and not a signal that anything is close to failing.

14. Formal verification

All arithmetic assertions in this paper are formalised in Lean 4 [23] against Mathlib [24] and checked by the Lean kernel. Polynomials are coefficient lists over $\mathbb{Z}$, low degree first; (1), palindromicity, nonnegativity, the Kruskal–Katona and Macaulay conditions and the pseudopower are computable, and every statement about a tabulated value — including the whole-table statements of Proposition 5.3, over 168 entries — is closed by kernel evaluation (`decide`). Compiled evaluation (`native_decide`) is not used anywhere, so no assertion rests on an unaudited evaluator; binomial coefficients are computed as descending factorials, the recursive definition being unable to evaluate $\binom{2627083776}{6}$ in the kernel, and certificate points are carried as integer pairs, rational arithmetic not reducing there. Theorems 4.1 and 5.5 are proved in the ordinary sense for all k and m below the stated bound, Lemma 3.3 from the definition of a downward-closed family, and control (iii) of Proposition 7.4 from Mathlib's own Kruskal–Katona theorem.

The Bose–Burton layer of §8–§10 adds four modules, and differs from the rest in that its input is computed rather than tabulated. The Gaussian binomials, the Whitney numbers (5) and the recursion (7) are definitions, so the Chow polynomials of Proposition 8.4 are derived inside the kernel; the closed form of Theorem 9.2 is a definition too, and the resolvability identity it satisfies is decided over the 381 matrices of Proposition 9.7, with the denominator cleared so that the coefficients are natural numbers and a positive verdict exhibits a nonnegative resolving array (Remark 9.5). Total nonnegativity from the definition — form the matrix, enumerate every one of the $\binom{2n}{n}$ minors, expand each determinant — is the one expensive check, and is run where it is affordable. The γ-vectors, the Kruskal–Katona and Newton verdicts, the 41 alternation certificates and Kurtz's criterion follow the same pattern as the earlier sections.

The layer of §11 is different again, and is the only one whose *proofs* are formal. Theorems 11.1, 11.2, 11.4 and Corollaries 11.3 and 11.5 are stated and proved over `[Field K] [AddCommGroup V] [Module K V]`, using Mathlib's own lattice of submodules; the first two and Corollary 11.3 carry no finiteness hypothesis beyond the use of dim, while Theorem 11.4 and Corollary 11.5 add `[Fintype K]` with $|K| \geq 3$, or $|K| = 2$ for the last part, since they count the points of a line. The chain sum (11) is written over an abstract finite ranked poset with values in an arbitrary commutative semiring, and the only property of the weight ever used is $w(1) = 0$. Two points of friction are worth recording, because they shape what the statements look like. First, Mathlib has no `Fintype` instance on the submodules of a finite-dimensional space over a finite field, so the finite sets of flat spans appear as hypotheses in the general theorems; over a finite V a separate module supplies them and the conclusions become identities between values of a function chowPoly of the deleted set. Second, that function is noncomputable, so there is no formal bridge from the general theorems to the calibration of Proposition 11.6, which runs on an independent bitmask model of $\mathbb{F}_q^3$ whose completeness is itself decided; the identification of the two is mathematics, as §11.5 says. The layer of §12–§13 returns to the earlier pattern: the parameter tables, the non-unimodality verdicts, the resolvability certificates at ranks 21 and 25, the 26 coefficient vectors and γ-vectors of Proposition 13.1 and their alternation certificates are all closed by kernel evaluation, on data produced by the lattice enumerations described there.

For scale: the projective-deletion layer is eight modules, about 2 100 lines and 91 theorems, and rechecks in about 217 seconds. The whole development is free of `sorry`; the axioms used are propositional extensionality, quotient soundness and, where Mathlib lemmas require it, choice. What is *not* formalised is listed in §1.6: the values of $\mathcal{H}$ for non-maximal building sets, the building-set, flagness and completeness verdicts, Lemma 5.1, the s-independence of Remark 3.8, the induction on coloops in Corollary 3.4; in the second part, Lemmas 8.1 and 8.2, the proof of Theorem 9.2 itself, and the deduction of Corollary 10.1 from [8, Thm. 6.1]; and in the third part, the lattice enumerations of Propositions 11.8 and 13.1, whose output enters the kernel as data, together with the two cited inputs of Remark 12.8. Repository reference to be added at submission.

References

[1] K. Adiprasito, J. Huh and E. Katz, *Hodge theory for combinatorial geometries*, Ann. of Math. (2) **188** (2018), no. 2, 381–452.

[2] N. Aisbett, *Frankl–Füredi–Kalai inequalities on the γ-vectors of flag nestohedra*, Discrete Comput. Geom. **51** (2014), no. 2, 323–336; arXiv:1203.4715.

[3] C. A. Athanasiadis and K. Kalampogia-Evangelinou, *Chain enumeration, partition lattices and polynomials with only real roots*, Comb. Theory **3** (2023), no. 1, Paper No. 12, 21 pp.; arXiv:2205.03796v2. Conjecture 1.2 is numbered from that version.

[4] R. C. Bose and R. C. Burton, *A characterization of flat spaces in a finite geometry and the uniqueness of the Hamming and the MacDonald codes*, J. Combin. Theory **1** (1966), no. 1, 96–104.

[5] P. Brändén and L. Saud-Maia-Leite, *Totally nonnegative matrices, chain enumeration and zeros of polynomials*, Adv. Math. **487** (2026), Paper No. 110760; arXiv:2412.06595v3 (11 December 2025). Numbered references above are to that version.

[6] P. Brändén and L. Saud-Maia-Leite, *On chain polynomials of geometric lattices*, arXiv:2508.13810v3 (11 December 2025). Numbered references above are to that version; versions 1 and 2 carry a misprint in Conjecture 4.3, see [13].

[7] P. Brändén and L. Vecchi, *Chow polynomials of uniform matroids are real-rooted*, arXiv:2501.07364.

[8] P. Brändén and L. Vecchi, *Chow polynomials of totally nonnegative matrices and posets*, arXiv:2509.17852v3 (10 November 2025). Numbered references above are to that version.

[9] T. Brylawski, *Intersection theory for embeddings of matroids into uniform geometries*, Stud. Appl. Math. **61** (1979), no. 3, 211–244. The theorem number cited above is the one [13, §8.2] give for the statement that a lower rank uniform geometric lattice is upper rank uniform; we have not consulted the original.

[10] R. Cheng and S. Liu, *Kazhdan–Lusztig polynomials of matroids need not be unimodal*, arXiv:2607.24186v2 (28 July 2026). Numbered references above are to that version.

[11] J. Chun, N. Hayashida, E. Partida and Z. Wang, *The Chow polynomial of the noncrossing partition lattice*, arXiv:2608.03806v1 (4 August 2026). The numbered reference above is to that version, which is the only one.

[12] B. Coron, L. Ferroni and S. Li, *Matroid analogues of Gal's conjecture*, arXiv:2604.04550v3 (10 July 2026). All numbered references above are to this version.

[13] B. Coron, L. Ferroni and S. Li, *Chow polynomials of rank-uniform labeled posets*, arXiv:2511.13819v2 (18 December 2025). Numbered references above are to that version; the numbering of its introduction, from which all of them come, is the same in version 1.

[14] E.-M. Feichtner and S. Yuzvinsky, *Chow rings of toric varieties defined by atomic lattices*, Invent. Math. **155** (2004), no. 3, 515–536.

[15] L. Ferroni, J. P. Matherne, M. Stevens and L. Vecchi, *Hilbert–Poincaré series of matroid Chow rings and intersection cohomology*, Adv. Math. **449** (2024), Paper No. 109733; arXiv:2212.03190v3. Numbered references above are to that version; versions 1 and 2 cite the Ferroni–Schröter statement inside Conjecture 1.6 as Conjecture 10.19 instead.

[16] L. Ferroni, J. P. Matherne and L. Vecchi, *Chow functions for partially ordered sets*, arXiv:2411.04070v3 (5 May 2026). Numbered references above are to that version; Theorem 3.10 carries the same number in version 2.

[17] L. Ferroni and B. Schröter, *Valuative invariants for large classes of matroids*, J. Lond. Math. Soc. (2) **110** (2024), no. 3, Paper No. e12984; arXiv:2208.04893v4. The conjecture cited above carries the number 8.18 in versions 3 and 4 and the number 10.19 in versions 1 and 2.

[18] Ś. R. Gal, *Real root conjecture fails for five- and higher-dimensional spheres*, Discrete Comput. Geom. **34** (2005), no. 2, 269–284.

[19] T. Hameister, S. Rao and C. Simpson, *Chow rings of vector space matroids*, J. Comb. **12** (2021), no. 1, 55–83; arXiv:1802.04241.

[20] G. O. H. Katona, *A theorem of finite sets*, in: Theory of Graphs (Tihany, 1966), Academic Press, 1968.

[21] D. C. Kurtz, *A sufficient condition for all the roots of a polynomial to be real*, Amer. Math. Monthly **99** (1992), 259–263.

[22] J. B. Kruskal, *The number of simplices in a complex*, in: Mathematical Optimization Techniques, Univ. California Press, 1963.

[23] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Comput. Sci., Springer, 2021.

[24] The mathlib Community, *The Lean mathematical library*, in: Proc. 9th ACM SIGPLAN Int. Conf. on Certified Programs and Proofs (CPP 2020), ACM, 2020.

[25] E. Nevo and T. K. Petersen, *On γ-vectors satisfying the Kruskal–Katona inequalities*, Discrete Comput. Geom. **45** (2011), no. 3, 503–521; arXiv:0909.0694.

[26] R. P. Stanley, *Combinatorics and Commutative Algebra*, 2nd ed., Birkhäuser, 1996.

[27] M. Stevens, *Real-rootedness conjectures in matroid theory*, Bachelor thesis, Stanford University, 8 June 2021, `https://mcs0042.github.io/bachelor.pdf`. The two conjectures cited above are stated there as joint with June Huh.