# FORESTS, ORIENTATIONS, AND TRUNCATIONS OF PLÜCKER MAPS

SANAY SEHGAL

Abstract. Over a field of characteristic two, we study truncations of the Plücker map of the Grassmannian cone of two-dimensional planes, obtained by reducing the two-by-two minors modulo monomial ideals that contain the squares of one row of variables. Identifying monomials in the Plücker coordinates with multigraphs, we show that a monomial survives the truncation exactly when its graph is a forest in which every vertex has bounded valency. We determine a minimal generating set for the kernel. It consists of the Plücker relations, the cycles, the stars, and one exceptional family: rainbows in the even case and spiders in the odd case. In particular, the kernel is infinitely generated. We give a basis of the image of every valency class, indexed by forests, and compute its rank, obtaining Catalan and ballot numbers. As a consequence, for Frobenius truncations the kernels are GL-stable ideals that are not GL-finitely generated.

We use the following conventions throughout: $R$ is either $\mathbb{Z}$ or a field $\Bbbk$, $\mathbb{N} = \{0, 1, \dots\}$, and $\mathfrak{S}_Y$ is the group of bijections of a set $Y$.

Let $\ell \in \mathbb{N}_{\geq 1}$ and $n \in \mathbb{N} \cup \{\infty\}$, and let $X = (x_{k,i})$ be the $\ell \times n$ matrix of variables, with rows indexed by $1 \leq k \leq \ell$ and columns by $1 \leq i \leq n$. For a subset $I = \{i_1 < \cdots < i_\ell\}$ of the columns, let $X_I$ denote the $\ell \times \ell$ submatrix of $X$ on the columns in $I$. The classical *Plücker embedding* sends the Grassmannian $\mathrm{Gr}(\ell, n)$ of $\ell$-dimensional subspaces of $\Bbbk^n$ into the projective space $\mathbb{P}(\bigwedge^\ell \Bbbk^n)$, by sending the row space of a full rank $\ell \times n$ matrix to the vector of its maximal minors. In coordinates it is the ring map

$$\Phi^R \colon R[m_I : I \subseteq [n],\ |I| = \ell] \longrightarrow R[x_{k,i}], \qquad \Phi^R(m_I) = \det X_I,$$

which we call the *Plücker map*. A monomial of $R[m_I]$ is a finite multiset of $\ell$-subsets of the columns, that is, an $\ell$-uniform multi-hypergraph on $[n]$, so the kernel of $\Phi^R$ and of its truncations below is a collection of sums of hypergraphs.

Let $q = (q_1, \dots, q_\ell)$ be a vector of *truncation exponents*, with each $q_k$ a positive integer or $\infty$, and let

$$T_q = \left(x_{k,i}^{q_k} : 1 \leq k \leq \ell,\ 1 \leq i \leq n\right)$$

be the ideal that bounds the exponent in row $k$ by $q_k$; we take $q_k = \infty$ to mean that row $k$ is not truncated. Let $p_q^R \colon R[x_{k,i}] \to R[x_{k,i}]/T_q$ be the quotient map. The *truncated Plücker map* is the composite

$$t_q^R \;=\; p_q^R \circ \Phi^R \colon R[m_I] \longrightarrow R[x_{k,i}]/T_q.$$

**Question.** *Given $R$, $\ell$ and $q$, what is a minimal generating set of* $\ker(t_q^R)$*?*

In this paper we answer the question when $R = \Bbbk$ is a field of characteristic two, $\ell = 2$, $n = \infty$, and $\min(q_1, q_2) = 2$. Exchanging the rows, we take $q = (2, r)$ with $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$. For $\ell = 2$ we write $x_i = x_{1,i}$, $y_i = x_{2,i}$, $m_{ij} = m_{\{i,j\}}$, $\Bbbk[x, y] = \Bbbk[x_i, y_i : i \geq 1]$ and $\Phi = \Phi^{\Bbbk}$, so that $\Phi(m_{ij}) = x_i y_j - x_j y_i$, and we write $\mathfrak{m}_r = T_{(2,r)} = (x_i^2, y_i^r)$, so that $\mathfrak{m}_\infty = (x_i^2)$, $p_r = p^{\Bbbk}_{(2,r)}$ and $t_{2,r} = t^{\Bbbk}_{(2,r)} = p_r \circ \Phi$.

Let $\mathcal{P}$ be the set of Plücker relations, let $\mathcal{C}$ be the set of cycle monomials and $\mathcal{Z}_{r+1}$ the star monomials (Subsection 2.1), with $\mathcal{Z}_\infty = \emptyset$, let $\mathcal{D}$ be the set of cut relations (Definition 5.3), and let $\mathcal{R}_r$ and $\mathcal{S}_r$ be the rainbow and spider relations respectively (Subsection 6.1). We write $t_{2,\mathrm{even}}$ and $t_{2,\mathrm{odd}}$ for $t_{2,r}$ at even and at odd $r$.

---

*Time*: Autumn 2026.

*Topic*: Algebraic combinatorics; combinatorial commutative algebra.

**Theorem A.** *Let* $\mathrm{char}(\Bbbk) = 2$ *and* $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$. *Then*

$$\ker(t_{2,r}) = (\mathcal{P},\ \mathcal{C},\ \mathcal{Z}_{r+1},\ \mathcal{D}).$$

*In terms of minimal generators,*

$$\ker(t_{2,\infty}) = (\mathcal{P}, \mathcal{C}), \qquad \ker(t_{2,\mathrm{even}}) = (\mathcal{P}, \mathcal{C}, \mathcal{Z}_{r+1}, \mathcal{R}_r), \qquad \ker(t_{2,\mathrm{odd}}) = (\mathcal{P}, \mathcal{C}, \mathcal{Z}_{r+1}, \mathcal{S}_r),$$

*where* $\mathcal{C}$, $\mathcal{R}_r$ *and* $\mathcal{S}_r$ *are taken one per class. In particular, none of these ideals is generated by finitely many* $\mathfrak{S}_\infty$*-orbits.*

**Theorem B.** *Let* $\mathrm{char}(\Bbbk) = 2$ *and consider a class of* $r$*-forests with* $n$ *vertices,* $k$ *trees, and* $u$ *vertices of valency* $r$. *The rank of the image of the class under* $t_{2,r}$ *is*

$$B(n,k) \quad \text{if } r = \infty,$$

$$\binom{n-u-2}{k-1} - \binom{n-u-2}{k-3} \quad \text{if } r \text{ is even}, \qquad \binom{n-u-1}{k-1} - \binom{n-u-1}{k-u-2} \quad \text{if } r \text{ is odd}.$$

*In particular we recover the Catalan number* $C_k$ *at* $r = 2$, *obtained by taking* $u = n - 2k$ *in the even formula, and* $B(n,k)$ *is the ballot number obtained by taking* $u = 0$ *in the even formula.*

**Acknowledgments.** The author thanks Karthik Ganapathy for proposing the problem of determining the kernel of the truncated Plücker map and for helpful discussions throughout the writing of the paper, Steven Sam for useful conversations, and Roger Van Peski for advice on exposition. Finally, the author is grateful to his family.

**Support.** This work was carried out while the author was a summer researcher at the University of California, San Diego, and an undergraduate student at Columbia University. The author was supported by a Columbia Research Fellowship.

## 1. Introduction

1.1. **Context.** The kernel of the Plücker map is the basic example of an algebra presented by straightening (Section 3). Its image, the algebra generated by the 2×2 minors, is the homogeneous coordinate ring of the Grassmannian of two-dimensional planes, and Hodge's straightening law [24] says that its kernel is generated by the set $\mathcal{P}$ of Plücker relations (Proposition 3.9). Doubilet, Rota and Stein [11] turned straightening into combinatorics, and De Concini, Eisenbud and Procesi made it an axiom: the algebra of minors is an algebra with straightening law [7, 3], whose standard monomials are indexed by semistandard tableaux and counted by Kostka numbers [13].

Monomials in the Plücker coordinates are graphs, and this places the problem among the algebras of graphs of combinatorial commutative algebra [28, 21, 36]. Sending $m_{ij}$ to the monomial $t_i t_j$ instead of the minor gives the edge subring of the complete graph, the toric ring of the second hypersimplex studied by De Loera, Sturmfels and Thomas [8], whose relations are the binomials of even closed walks [35, 36]. Taking the minors over the edges of one fixed graph gives its binomial edge ideal, introduced by Herzog, Hibi, Hreinsdóttir, Kahle and Rauh [22] and by Ohtani [30]. Postnikov and Shapiro [31], Ardila and Postnikov [2], and Holtz and Ron [25] study algebras whose bases are indexed by the trees and forests of a fixed graph. In this paper the graphs vary, the object is the algebra the minors generate, and the question is which sums of graphs vanish.

The truncations arise from positive characteristic. When $r$ is a power of two, $\mathfrak{m}_r$ is the sum of the Frobenius powers $(x_i)^{[2]}$ and $(y_i)^{[r]}$, and $\mathfrak{m}_\infty = (x_i)^{[2]}$; we call these the *Frobenius truncations*. Two works study the target rings. Casabella and Yu [4] work in $\Bbbk[x,y]/\mathfrak{m}_2$ and compute the $\mathrm{GL}_n$-characters of the pieces of the ideal generated by the minors, and Gao, Raicu and VandeBogert [18] pose conjectures on the rings $\Bbbk[x,y]/(x_i^p, y_i^p)$ under the hypothesis $a - b \geq p - 1$ on the bidegree, which excludes the diagonal $a = b$ where the image of $t_{2,r}$ lives. Neither work considers the kernel, its generators, bases of the image indexed by graphs, or the truncations with $r > 2$. All of these are new here: the presentation of Theorem A and the forest bases of Theorem B are new at every $r$,

including $r = 2$, where the character formula of [4] determines the dimensions but gives neither a basis nor any relations.

The final motivation is noetherianity. Identifying $\Bbbk[m_{ij}]$ with $\mathrm{Sym}(\bigwedge^2 \mathbf{V})$, where $\mathbf{V} = \Bbbk^\infty$ (Subsection 8.1), Nagpal, Sam and Snowden [29] proved that this algebra is noetherian for ideals stable under $\mathrm{GL}_\infty$ in characteristic zero, and Ganapathy [16] proved that it is not in characteristic two, by a chain of ideals generated by Plücker relations and cycles. For the Frobenius truncations, $\mathfrak{m}_r$ is stable under linear changes of the columns (Lemma 8.3), so their kernels are GL-stable, and the minimal generators of Theorem A occur in every degree.

**Theorem C.** *Let* $\mathrm{char}(\Bbbk) = 2$ *and let* $r$ *be a power of* 2 *or* $r = \infty$*, so that* $t_{2,r}$ *is a Frobenius truncation. Then* $\ker(t_{2,r})$ *is a* GL*-stable ideal of* $\mathrm{Sym}(\bigwedge^2 \mathbf{V})$ *which is not* GL*-finitely generated. In particular* $\mathrm{Sym}(\bigwedge^2 \mathbf{V})$ *is not* GL*-noetherian.*

Ganapathy [16] proved the last statement by exhibiting an ascending chain of GL-stable ideals, and the preimage in $\mathrm{Sym}(\bigwedge^2 \mathbf{V})$ of the union of his chain is $(\mathcal{P}, C) = \ker(t_{2,\infty})$. Our proof is independent of his and goes through the minimal generating set of Theorem A, which determines exactly which generators witness the failure. For $r = 2^t$ the kernels are GL-stable ideals that have not been studied before, and their minimal generators include the stars and rainbows.

1.2. **Outline of the proofs.** We consider multigraphs on finite subsets of $\mathbb{N}_{\geq 1}$, without loops, in which a repeated edge is a repeated element of the edge multiset. A multigraph $G$ corresponds to the monic monomial $m_G = \prod_e m_e$ of $\Bbbk[m_{ij}] = \Bbbk[m_{ij} : 0 < i < j]$, obtained by sending the edge $\{i, j\}$ to the variable $m_{ij}$, so that multiplying monomials is taking the union of edge multisets and dividing is taking a subgraph. Figure 1 shows the three kinds of graph that occur as monomial generators of the kernels.

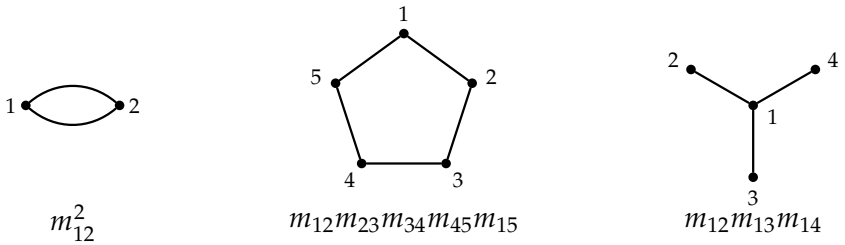



FIGURE 1. The doubled edge $C_2$, the cycle $C_5$ and the star $\mathcal{Z}_3$, with their monomials

The proofs of Theorems A and B follow one pattern. All relations split into relations inside single classes, the graphs with the same valency at every vertex (Proposition 2.2), and the terms of $\Phi(m_G)$ are indexed by the orientations of $G$ (Subsection 2.2); after setting $y_i = 1$ this is the orientation expansion of the graph polynomial of Alon and Tarsi [1]. In each class we cut the graphs down in three steps, each modding out by one family of relations, and then show that what remains is independent by reading off a graph invariant.

(1) *Modding out by monomials excludes subgraphs.* An orientation survives the truncation exactly when every vertex has at most one outgoing edge and at most $r - 1$ incoming edges. So a graph dies exactly when it contains a cycle, including the doubled edge, or a star $\mathcal{Z}_{r+1}$, and the survivors are the *r-forests*, the forests of maximum valency at most $r$ (Proposition 4.5).

(2) *Modding out by Plücker relations leaves cross-free graphs.* A Plücker relation is the sum of the three perfect matchings of four vertices, so it rewrites a crossing into a disjoint and a nested pair, and every graph is a sum of cross-free graphs of its class (Proposition 3.2). Applied to $r$-forests this leaves the cross-free $r$-forests (Proposition 5.5).

(3) *Modding out by the exceptional relations chooses representatives.* The surviving orientations of an $r$-forest direct each tree toward one vertex of valency less than $r$, so its image is recorded by its root selections (Proposition 4.6), and for two trees by the cut of one root set in a complete graph (Proposition 5.4). A *spider* identifies two forests with the same root sets by exchanging the centers of two stars, and a *rainbow* cancels root sets paired by a perfect matching, which needs an even number of leaves; this is the parity in Theorem A. Modulo these we keep one representative for each image, for instance one linear forest for each cross-free endpoint pairing at $r = 2$ (Proposition 6.4).

(4) *A graph invariant shows independence.* Each representative carries an invariant read off from one term of its image: the opener, the multiset of left endpoints of its edges, for $\Phi$; the root maxima at $r = \infty$ and at even $r$; and a finer odd label at odd $r$ (Definition 7.7). That term has coefficient one and occurs only in representatives on one side of it in dominance order (Propositions 7.1 and 3.6), and distinct representatives have distinct invariants (Propositions 3.4 and 7.5). So the representatives are independent (Proposition 3.5).

Steps (1) to (3) span every class modulo the proposed relations, which for finite $r$ is proved by moving branches between vertices and inserting vertices of valency two (Subsection 7.4), and step (4) shows that nothing more is needed, which proves Theorem A. Counting the invariants proves Theorem B: cross-free perfect matchings give the Catalan numbers [34], a reflection argument on lattice paths gives the ballot numbers, and partial matchings give the odd formula. The rank of a class is the value at its valency vector of the multigraded Hilbert function [28, Chapter 8] of $\Bbbk[m_{ij}]/\ker(t_{2,r})$, so these counts are the Hilbert functions of the truncated algebras, as the two-row Kostka numbers [13] are for $\Bbbk[m_{ij}]/(\mathcal{P})$; at $\Phi$ the cross-free and nesting-free graphs are the noncrossing and nonnesting graphs of de Mier [9], and the nesting-free graphs are the standard monomials (Propositions 3.12 and 3.13). Minimality of each generator is a count on its class that every relation inherited from a smaller class fails (Subsection 6.3). Finally Section 8 places the maps among GL-algebras and proves Theorem C.

The sections follow these steps. Section 2 sets up classes and orientations, Section 3 carries out step (2) and step (4) for $\Phi$, Section 4 carries out step (1) and records images by root selections, Section 5 gives the cut criterion, Section 6 gives the exceptional relations and proves minimality, and Section 7 carries out step (4) for the truncations and proves generation.

## 2. Graph algebras

2.1. **Multigraphs.** Throughout the paper let $G$ be a multigraph, as in Section 1, with edge multiset $E$ and vertex set $V = \{v_1, \dots, v_n\} \subseteq \mathbb{N}_{\geq 1}$, the set of endpoints of its edges, with $v_1 < \cdots < v_n$. The *cycle* $C_n$ is the graph on $n$ vertices with edges $\{1,2\}, \{2,3\}, \dots, \{n-1,n\}, \{1,n\}$, so that $C_2$ is the doubled edge. The *star* $\mathcal{Z}_n$ is the graph on $n+1$ vertices in which one vertex, the *center*, is joined to each of the other $n$.

Throughout, let $\mathfrak{n}$ be an ideal of $\Bbbk[x,y]$ and let $\tau_{\mathfrak{n}} = \overline{p}_{\mathfrak{n}} \circ \Phi$ where $\overline{p}_{\mathfrak{n}} : \Bbbk[x,y] \to \Bbbk[x,y]/\mathfrak{n}$ is the natural projection modulo the ideal $\mathfrak{n}$. We recover $t_{2,r}$ when $\mathfrak{n} = (x_i^2, y_i^r)$ and $\Phi$ when $\mathfrak{n} = (0)$.

The *valency* of a vertex $v$ is the number of edges which connect to $v$ with multiedges counted with multiplicity. We define the *valency vector* of a graph as

$$\mathcal{W}(G) = (\mathrm{val}(v_1), \dots, \mathrm{val}(v_n)).$$

The number of edges in a graph we denote by $d(G)$.

**Definition 2.1.** We define the *class* of the graph $G$ as

$$\mathcal{O}(G) = \{\tilde{G} : V_G = V_{\tilde{G}} \text{ and } \mathcal{W}(G) = \mathcal{W}(\tilde{G})\}$$

to be the set of graphs on the vertex set $V$ which have the same valency vector as $G$.

The *components* of $G$ are its maximal *connected* subgraphs and they partition the edge multiset: we denote $G = \bigsqcup_{i=1}^{q} G^{(i)}$ for the partition into its components. Then $m_G = \prod_{i=1}^{q} m_{G^{(i)}}$ and $\tau_{\mathfrak{n}}(m_G) = \prod_{i=1}^{q} \tau_{\mathfrak{n}}(m_{G^{(i)}})$.

We define a *transversal* of $G$ to be $\{T(e) : e \in E\}$ where $T : E \to V$ and $T(\{i,j\}) \in \{i,j\}$. Hence to each edge copy, we assign one of the vertices it lies on. The number of transversals is $2^{d(G)}$ since each edge is assigned independently.

Let $\varphi : \mathbb{N}_{\geq 1} \to \mathbb{N}_{\geq 1}$ be injective. Then we define $\varphi(G)$ as the relabeled graph formed by replacing all $v$ with $\varphi(v)$ for all vertices (and edges). For an ideal $I \subseteq \Bbbk[m_{ij}]$, we denote $G \in I$ to indicate that $m_G \in I$ for any labeling of $G$. Similarly, we denote $H \mid G$ to indicate that $H$ is a subgraph of $G$, meaning for any $G$, we have that $m_{\varphi(H)} \mid m_G$ for some $\varphi$.

When determining kernel relations (not monomials) in $\Phi$ or $t_{2,r}$, the question is which sums of graphs are in the kernel across all possible graphs. Concretely, we can fix some $G_1$ and ask about sums $m_{G_1} + \cdots + m_{G_n}$ which are in the kernel. In this subsection we grade $\Bbbk[m_{ij}]$ and $\Bbbk[x,y]$ to refine this question.

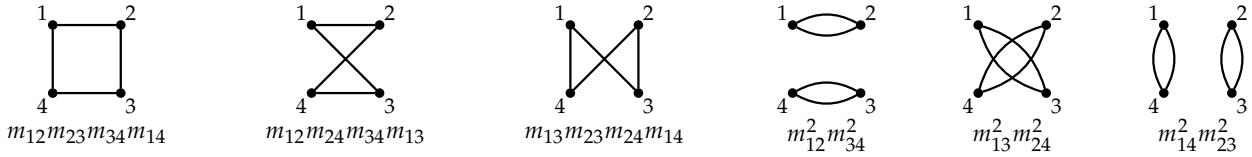



FIGURE 2. The valency class $O(C_4)$ of $(2,2,2,2)$

Let $\mathbf{d} \in \mathbb{N}^{\mathbb{N}}$ be a valency vector with finitely many nonzero entries. We write $O(\mathbf{d})$ for the class of multigraphs with valency vector $\mathbf{d}$, which is finite since its members have $\frac{1}{2}\sum_v d_v$ edges on a fixed vertex set, and we write $\Bbbk[m_{ij}]_{\mathbf{d}}$ for the span of the monomials $m_G$ with $G \in O(\mathbf{d})$, so that $\dim_{\Bbbk} \Bbbk[m_{ij}]_{\mathbf{d}} = |O(\mathbf{d})|$. Adding edges adds valencies, so $\Bbbk[m_{ij}]_{\mathbf{d}} \cdot \Bbbk[m_{ij}]_{\mathbf{b}} \subseteq \Bbbk[m_{ij}]_{\mathbf{d}+\mathbf{b}}$. Similarly, for a monomial $\mu = \prod_v x_v^{a_v} y_v^{b_v}$ of $\Bbbk[x,y]$ we define its *valency vector* by $\mathrm{val}_v(\mu) = a_v + b_v$, and we write $\Bbbk[x,y]_{\mathbf{d}}$ for the span of the monomials with valency vector $\mathbf{d}$.

Each factor $\Phi(m_{ij}) = x_i y_j - x_j y_i$ has both of its terms of valency one at $i$ and at $j$ and zero elsewhere, so every term of $\Phi(m_G)$ has valency vector $\mathcal{W}(G)$. Reducing modulo an ideal generated by monomials deletes some terms and keeps the others, so

$$\tau_{\mathfrak{n}}(\Bbbk[m_{ij}]_{\mathbf{d}}) \subseteq \Bbbk[x,y]_{\mathbf{d}}$$

whenever $\mathfrak{n}$ is generated by monomials, as $(0)$ and $\mathfrak{m}_r$ are. The following proposition is clear.

**Proposition 2.2.** *Let $\mathfrak{n}$ be generated by monomials and let $f = \sum_i c_i m_{H_i} \in \ker(\tau_{\mathfrak{n}})$. Then for every class $O$ the partial sum $\sum_{H_i \in O} c_i m_{H_i}$ lies in* $\ker(\tau_{\mathfrak{n}})$.

In particular a relation with no proper subsum in the kernel is supported in a single class.

This is the mechanism in which we change problems about characterizing maps into finite linear algebra.

2.2. **Orientations.** In this subsection we use directed multigraphs to index coefficients of $\Phi(m_G)$ and derive a formula for the Plücker map in terms of directed graphs.

Consider the *directed multigraph algebra*, which we define to be $\Bbbk[m_{i\to j} : i \neq j]$ and denote by $\Bbbk[m_{i\to j}]$. Throughout let $D$ be a directed multigraph (without directed loops or isolated vertices) with directed multiset $I$ and a finite vertex set $V = \{v_1, \ldots, v_n\}$. Then the assignment

$$D \longleftrightarrow m_D = \prod_{i \in I} m_i$$

is a bijection from set of directed multigraphs on finitely many vertices to the set of monic monomials of $\Bbbk[m_{i\to j}]$ mirroring the undirected case.

To each directed multigraph we can assign its undirected counterpart through the non-injective *forgetting map* $\Gamma : \Bbbk[m_{i\to j}] \to \Bbbk[m_{ij}]$ defined on generators by $\Gamma(m_{i\to j}) = m_{ij}$. Hence $\Gamma$ is simply the map which forgets the direction of the edges of $D$.

**Definition 2.3.** An *orientation* of $G$ is an element of $\Gamma^{-1}(G)$ (which is really $\Gamma^{-1}(m_G)$). In addition, we define the set of orientations of $G$ to be $\Omega(G) = \Gamma^{-1}(G)$.

Let $t$ be the number of edges of $G$ counted *without* multiplicity. Then $|\Omega(G)|$ is the product, over those $t$ edges, of one more than the number of copies of the edge, since in $\Bbbk[m_{i\to j}]$ only the number of copies oriented each way is recorded. Note that if $G$ is a *simple graph* (meaning $G$ has no multiedges), then $|\Omega(G)| = 2^{d(G)}$.

Let $D \in \Omega(G)$ with $v_i \in V$. Then we define the *outvalency* $\mathrm{outval}(v_i)$ to be the number of edges of form $v_i \to v_j$ for some $v_j \in V$ with multiplicity included. Adjacently, we define the *invalency* $\mathrm{inval}(v_i)$ to be the number of edges of form $v_j \to v_i$ for some $v_j \in V$ with multiplicity included again. Hence outvalency is the number of edges oriented away from the vertex and invalency is the number of edges oriented towards the vertex.

**Definition 2.4.** The *directed Plücker map* is defined as $\Psi : \Bbbk[m_{i\to j}] \to \Bbbk[x, y]$ determined by $\Psi(m_{i\to j}) = x_i y_j$.

The *Segre embedding* [20, Lecture 2] is the map which embeds $\mathbb{P}^{n-1} \times \mathbb{P}^{n-1}$ into the projectivization of $\Bbbk^n \otimes \Bbbk^n$ by sending $([v], [w])$ to $[v \otimes w]$. Then the directed Plücker map is the coordinate version of the two dimensional Segre embedding.

The *Plücker splitting map* which sends the multigraph algebra into the directed multigraph algebra is defined as $\Delta : \Bbbk[m_{ij}] \to \Bbbk[m_{i\to j}]$ determined by $\Delta(m_{ij}) = m_{i\to j} - m_{j\to i}$. Naturally, we have that $\Phi = \Psi \circ \Delta$:

$$\begin{array}{ccc} \Bbbk[m_{ij}] & \xrightarrow{\Delta} & \Bbbk[m_{i\to j}] \\ & \searrow^{\Phi} & \downarrow{\scriptstyle \Psi} \\ & & \Bbbk[x, y] \end{array}$$

Let $H$ be an orientation of $G$. We denote an element of $E(H)$ as $i \to j$ (meaning the edge which directs from $i$ to $j$). Then let $\mathrm{Des}(H) = \{i \to j \in H : i > j\}$ be the multiset of *descents* of $H$, the directed edges which direct from a larger vertex to a smaller vertex. Then we define the *sign* of $H$ to be $\varepsilon(H) = (-1)^{|\mathrm{Des}(H)|}$.

Let $e = i \to j$ be an edge in $H$. Then we define $m(e)$ to be the multiplicity of $e$ in the underlying multigraph $\Gamma(H)$ (more specifically in the edge multiset of $\Gamma(H)$). And, we define $\mathrm{Des}_e(H) = \{e \in \Gamma(\mathrm{Des}(H)) : \Gamma(e) = \{i, j\}\}$ to be the multiset of misorientations of $e$.

We define the graph $\mathcal{B}_n$ (the *banana graph*) to be the graph with $n$ edges between two vertices ($n \geq 1$) with edge set $E(\mathcal{B}_n)$. Note that $\mathcal{B}_2 = C_2$. Let $G$ be a multigraph, then we can factor $G$ as disjoint banana graphs $m_{\mathcal{B}_{n_i}}$, giving $m_G = \prod_i m_{\mathcal{B}_{n_i}}$.

**Proposition 2.5.** *For a multigraph $G$ with edge set $E$, taken without multiplicity, we have that*

$$\Phi(m_G) = \sum_{H \in \Omega(G)} \left( \prod_{e \in E} \binom{m(e)}{|\mathrm{Des}_e(H)|} \right) \varepsilon(H) \Psi(m_H) \in \Bbbk[x, y].$$

*Proof.* For the banana $\mathcal{B}_c$ on $a < b$ with $c$ copies,

$$\Phi(m_{ab}^c) = (\Psi(m_{a\to b}) - \Psi(m_{b\to a}))^c = \sum_{k=0}^{c} \binom{c}{k} (-1)^{c-k} (\Psi(m_{a\to b}^k m_{b\to a}^{c-k})) = \sum_{k=0}^{c} \binom{c}{k} (-1)^{c-k} (x_a^k x_b^{c-k} y_a^{c-k} y_b^k)$$

through the binomial theorem. Let $m_G = \prod_{e\in E} m_e^{m(e)}$, then

$$\Phi(m_G) = \Phi\left(\prod_{e\in E} m_e^{m(e)}\right) = \prod_{e\in E}\Phi\left(m_e^{m(e)}\right) = \prod_{e\in E}\left(\sum_{k=0}^{m(e)}\binom{m(e)}{k}(-1)^{m(e)-k}\Psi\left(m_{a_e\to b_e}^{k} m_{b_e\to a_e}^{m(e)-k}\right)\right)$$

$$= \sum_{(k_e)_{e\in E}}\left(\prod_{e\in E}\binom{m(e)}{k_e}(-1)^{m(e)-k_e}\Psi\left(m_{a_e\to b_e}^{k_e} m_{b_e\to a_e}^{m(e)-k_e}\right)\right) = \sum_{H\in\Omega(G)}\left(\prod_{e\in E}\binom{m(e)}{|\operatorname{Des}_e(H)|}\right)\varepsilon(H)\Psi(m_H),$$

where $e = \{a_e, b_e\}$ with $a_e < b_e$ and the sum is over tuples with $0 \le k_e \le m(e)$. Each tuple records how many of the copies of each edge are directed from the smaller vertex to the larger, so it is an orientation $H \in \Omega(G)$. Its misoriented copies at $e$ are the remaining $m(e)-k_e$, so $|\operatorname{Des}_e(H)| = m(e)-k_e$ and $\binom{m(e)}{k_e} = \binom{m(e)}{|\operatorname{Des}_e(H)|}$. Then across all edges we get $\prod_e(-1)^{|\operatorname{Des}_e(H)|} = \varepsilon(H)$ and the result follows. □

Then $|\Gamma^{-1}(G)| = \prod_e(m(e)+1)$ since the choice of $k_e$ is independent across edges. We can specialize to the case of a simple graph. For a *simple* graph $G$, we get that

$$\Phi(m_G) = \sum_{H\in\Omega(G)}\varepsilon(H)\prod_v x_v^{\operatorname{outval}(v)} y_v^{\operatorname{inval}(v)} \in \Bbbk[x, y].$$

Since for an orientation $D$, we have that $\Psi(m_D) = \prod_v x_v^{\operatorname{outval}(v)} y_v^{\operatorname{inval}(v)}$ since each outvalency corresponds to the $x$-variable and each invalency corresponds to the $y$-variable by construction. As $G$ is simple every $m(e) = 1$, so each binomial coefficient in Proposition 2.5 is 1.

If $G$ is simple then $|\Gamma^{-1}(G)| = 2^{d(G)}$ since each choice is independent. Figure 3 expands the path $m_{12}m_{23}$ this way.




FIGURE 3. The four orientations of $m_{12}m_{23}$ and their terms, with descents dashed

2.3. **Cycles.** In this section we generalize Proposition 2.5 by characterizing orientations which produce the same monomial.

A *directed cycle* is an orientation of a cycle in which $\operatorname{inval}(v) = \operatorname{outval}(v)$ for all vertices in the cycle. This means that all edges in the cycle are oriented in the same direction.

**Definition 2.6.** Let $H$ be an orientation of $G$. We define the *cancellation class* of an orientation as

$$[H] = \{H' \in \Omega(G) : \Psi(m_{H'}) = \Psi(m_H)\}.$$

Suppose $H$ has vertex set $V$. If $H', K \in [H]$ meaning $\Psi(m_{H'}) = \Psi(m_K)$ then

$$\prod_{v\in V} x_v^{\operatorname{outval}_K(v)} y_v^{\operatorname{inval}_K(v)} = \prod_{v\in V} x_v^{\operatorname{outval}_{H'}(v)} y_v^{\operatorname{inval}_{H'}(v)},$$

hence by cancellation $\operatorname{outval}_K(v) = \operatorname{outval}_{H'}(v)$ and $\operatorname{inval}_K(v) = \operatorname{inval}_{H'}(v)$ for all $v \in V$.

Let $\mathcal{J}(H)$ be the set of directed subcycles in $H$.

**Definition 2.7.** Let $H$ be an orientation. Then we define the *cycle algebra* of $H$ to be

$$\mathbb{J}(H) = \left\{\bigcup_{a\in A} a : A \subseteq \mathcal{J}(H),\ a \cap a' = \emptyset \text{ for all distinct } a, a' \in A\right\}$$

with the $\emptyset$ counting as a subset.

We can now state the general formula for $\Phi(m_G)$ categorized by orientation type. Setting $y_i = 1$ turns $\Phi(m_G)$ into the graph polynomial of Alon and Tarsi [1], whose coefficients are signed counts of orientations; the formula below refines that count by keeping the invalencies and grouping orientations into cancellation classes.

**Proposition 2.8.** *Let $G$ be a simple graph. Then*

$$\Phi(m_G) = \sum_{[H]} \varepsilon(H) c(H) \Psi(m_H)$$

*where*

$$c(H) = \sum_{J \in \mathbb{J}(H)} \begin{cases} 1 \text{ if } |J| \text{ is even}, \\ -1 \text{ if } |J| \text{ is odd} \end{cases}$$

*with $|J|$ the number of edges of $J$ and $|\emptyset| = 0$ being counted as even.*

*Proof.* Let $J \in \mathbb{J}(H)$ and define $J(H)$ to be the orientation $H$ with the cycles of $J$ reversed. Then the map sending $J$ to $J(H)$ is a bijection

$$\mathbb{J}(H) \longleftrightarrow [H].$$

Let $K \in [H]$, then define $A$ to be the set of edges on which $H$ and $K$ differ, oriented as in $H$. We have that $\mathrm{outval}_A(v) = \mathrm{inval}_A(v)$ for all $v \in V$ since

$$\mathrm{outval}_{H\setminus A}(v) + \mathrm{outval}_A(v) = \mathrm{outval}_H(v) = \mathrm{outval}_K(v) = \mathrm{outval}_{K\setminus A}(v) + \mathrm{inval}_A(v)$$

and $H\setminus A = K\setminus A$. Then we can traverse the edges of $A$ in their direction and at each vertex continue, so $A$ is a disjoint union of directed cycles. Therefore $A \in \mathbb{J}(H)$ and $A(H) = K$, giving surjectivity.

If $J(H) = J'(H)$ then $J$ and $J'$ are both exactly the set of edges on which $H$ and $J(H)$ disagree, so $J = J'$. Hence the map is injective, and with surjectivity it is a bijection.

Reversing a directed cycle of length $s$ changes the number of descending edges by $s$ minus twice the number of descending edges on the cycle, so it multiplies $\varepsilon$ by $(-1)^s$. Hence $\varepsilon(J(H)) = \varepsilon(H)(-1)^{|J|}$, and summing Proposition 2.5 over the class $[H]$ gives the formula. □

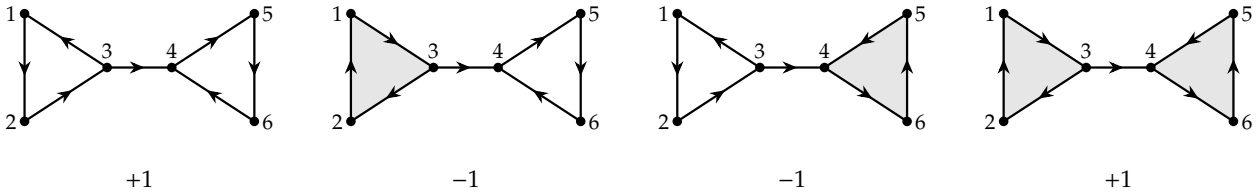



FIGURE 4. A cancellation class $[H]$ with $|\mathbb{J}(H)| = 4$ and $c(H) = 0$, reversed triangles shaded

Figure 4 shows the four orientations obtained from one $H$ by reversing its two disjoint directed triangles, with their signs $\varepsilon$; they share the monomial $x_1x_2x_3^2x_4x_5x_6\, y_1y_2y_3y_4^2y_5y_6$.

For a multigraph the orientations record only how many copies of each edge point each way, since multiedges are powers $m_e^{m(e)}$, so the reversal argument above no longer applies. We collapse each banana to a single edge. Let $\overline{G}$ be the underlying simple graph of $G$, with an edge $e = \{a, b\}$, $a < b$, for each banana $\mathcal{B}_{m(e)}$ of $G$. For an orientation $H$ on $G$ we denote the number of copies of $e$ directed $a \to b$ by $f_e(H)$, so that $|\mathrm{Des}_e(H)| = m(e) - f_e(H)$.

**Definition 2.9.** Let $H$ be an orientation of a multigraph $G$. We define the *circulation set* $\mathrm{Circ}(H)$ to be the set of integer vectors $x = (x_e)$ indexed by the edges of $\overline{G}$ such that

$$-|\mathrm{Des}_e(H)| \le x_e \le f_e(H) \quad \text{for every edge } e,$$

and such that at every vertex $v$ the sum of $x_e$ over the edges $e = \{v, b\}$ with $v < b$ equals the sum of $x_e$ over the edges $e = \{a, v\}$ with $a < v$. We define $|x| = \sum_e x_e$.

**Proposition 2.10.** *Let $G$ be a multigraph. Then*

$$\Phi(m_G) = \sum_{[H]} \varepsilon(H) c(H) \Psi(m_H)$$

*where*

$$c(H) = \sum_{x \in \mathrm{Circ}(H)} (-1)^{|x|} \prod_e \binom{m(e)}{f_e(H) - x_e}.$$

*Proof.* An orientation $K$ on $G$ is determined by the numbers $f_e(K)$. We can define $x_e = f_e(H) - f_e(K)$. At a vertex $v$, the outvalency of $K$ differs from that of $H$ by the sum of $x_e$ over the edges $e = \{v, b\}$ with $v < b$ minus the sum of $x_e$ over the edges $e = \{a, v\}$ with $a < v$. So $K \in [H]$ exactly when $x$ satisfies the vertex condition of Definition 2.9, and the bounds on $x_e$ say exactly that $0 \le f_e(K) \le m(e)$. Hence $K \mapsto x$ is a bijection $[H] \longleftrightarrow \mathrm{Circ}(H)$.

By Proposition 2.5 the coefficient of $\Psi(m_K)$ contributed by $K$ is $\varepsilon(K) \prod_e \binom{m(e)}{|\mathrm{Des}_e(K)|}$. We have $\binom{m(e)}{|\mathrm{Des}_e(K)|} = \binom{m(e)}{f_e(K)} = \binom{m(e)}{f_e(H) - x_e}$ and $|\mathrm{Des}_e(K)| = |\mathrm{Des}_e(H)| + x_e$, so $\varepsilon(K) = \varepsilon(H)(-1)^{|x|}$. Summing over $K \in [H]$ gives the formula. □

When $G$ is simple every $m(e) = 1$, so each $x_e$ is 0 or points in the direction $H$ uses, and the circulations are exactly the sets of edges of $H$ forming a disjoint union of directed cycles. All binomial coefficients are 1, and Proposition 2.10 reduces to Proposition 2.8.

**Proposition 2.11.** *Let $n \ge 2$. Modulo $(x_i^2)$ we have that*

$$\Phi(m_{C_n}) \equiv \left((-1)^{n-1} - 1\right) x_1 \cdots x_n \, y_1 \cdots y_n.$$

*In particular, if $\Bbbk$ has characteristic two then $C_n \in \ker(\tau_{\mathfrak{n}})$ for every ideal $\mathfrak{n}$ containing the $x_i^2$.*

*Proof.* Suppose first that $n \ge 3$, so that $C_n$ is simple. An orientation contributes a monomial divisible by some $x_v^2$ as soon as one vertex has outvalency 2. The outvalencies sum to $n$, so if none is 2 then each is exactly 1 and the orientation is a directed cycle. There are two, both contributing $x_1 \cdots x_n y_1 \cdots y_n$. The orientation $1 \to 2 \to \cdots \to n \to 1$ has the single descending edge $n \to 1$, so its sign is $-1$; the reverse has $n - 1$ descending edges, so its sign is $(-1)^{n-1}$. Their sum is $(-1)^{n-1} - 1$.

For $n = 2$ we have $\Phi(m_{12}^2) = x_1^2 y_2^2 - 2x_1 x_2 y_1 y_2 + x_2^2 y_1^2$, whose only term not divisible by an $x_i^2$ is $-2x_1 x_2 y_1 y_2$. □

## 3. Plücker relation

We now use the graph dictionary of Section 1 throughout our proofs. A *weighted graph list* is a finite collection of graphs with weights in $\Bbbk$; repeated graphs have their weights added. Its *orientation record* at a transversal is the signed sum of the orientations producing that transversal. A relation is a weighted list whose surviving records all have weight zero. In characteristic two the signs disappear, and an unweighted list is counted by parity. Adding a fixed collection of edges to every graph in a list is the graph operation written as multiplication of monomials in Section 1. Thus the kernel and basis statements retain their original notation, but we prove them by local moves, cancellation of records, and counting.

3.1. **Crossings.** In this section we describe a method of combinatorial straightening, in the spirit of Doubilet, Rota and Stein [11], which we will use to prove that $\ker(\Phi) = (\mathcal{P})$. Recall the Plücker relation is

$$m_{ab} m_{cd} - m_{ac} m_{bd} + m_{ad} m_{bc} = 0$$

where $a < b < c < d$. We can reframe this in terms of edges.

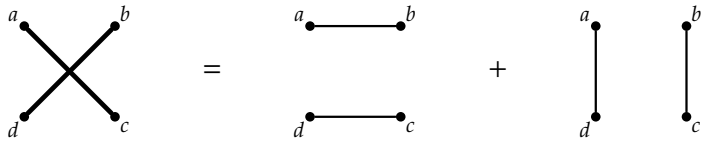



FIGURE 5. The Plücker relation

**Definition 3.1.** Let $G$ be a multigraph with $\{a, b\}$ and $\{c, d\}$ being edges in $G$. We say that the edges $\{a, b\}$ and $\{c, d\}$ are *crossed* if

$$a < c < b < d.$$

Hence in the Plücker relation, the term $m_{ac}m_{bd}$ is crossed.

For a multigraph $G$ we define $\mathcal{F}(G)$ to be the number of crosses in $G$. If $\mathcal{F}(G) = 0$, we call $G$ *cross-free*. We straighten using total edge length first and the number of crossings second; the two measures ensure termination even when the first smoothing changes crossings elsewhere.

**Proposition 3.2.** *Let $G$ be a multigraph. Then $m_G \equiv m_{G_1} + \cdots + m_{G_n}$ modulo $(\mathcal{P})$ with each $G_i$ cross-free and in the class of $G$.*

*Proof.* Give an edge $ij$, with $i < j$, length $j - i$, and sum these lengths over edge copies. Replace a crossing $ac, bd$, where $a < b < c < d$, by the two lists $ab, cd$ and $ad, bc$, retaining every other edge. The first has smaller total length. The second has the same length and fewer crossings: the crossing of these two copies disappears, and the number of crossings with any fixed third copy does not increase. To check the last assertion, place the endpoints of the third copy in the intervals cut out by $a, b, c, d$; a crossing occurs exactly when one endpoint is strictly between the endpoints of the edge being tested and the other is outside, with shared endpoints excluded. The nested pair never contributes more than the crossed pair. This check also allows endpoints equal to $a, b, c, d$ and repeated edge copies.

Each child list therefore decreases the lexicographically ordered pair (total length, number of crossings). These are nonnegative integers and the first coordinate never increases, so repeated replacements terminate. Every replacement preserves the valency at every vertex. The terminal graphs are cross-free, proving the assertion. □

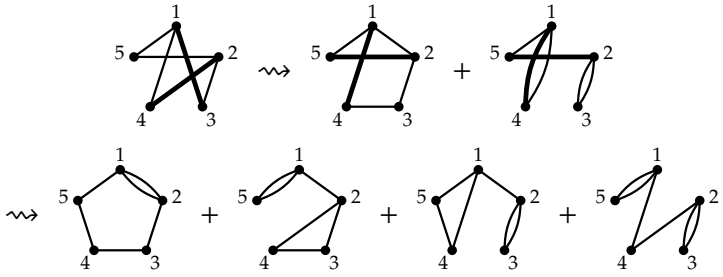



FIGURE 6. The straightening of $K_{2,3}$; the thick crossing is smoothed next

By linearity, every element of $\mathbb{k}[m_{ij}]$ is congruent modulo $(\mathcal{P})$ to a sum of *cross-free monomials* (cross-free graphs).

3.2. **Openers.** We use an invariant to keep track of cross-free graphs from which we show linear independence of cross-free elements.

Let $S, T$ be transversals of $G$ (multisets) with $d = d(G)$. Then we define the dominance order on the set of transversals of $G$. First, we order $S$ and $T$ as $s_1 \leq \cdots \leq s_d$ and $t_1 \leq \cdots \leq t_d$ respectively. We

define $S \preceq T$ if $s_k \leq t_k$ for all $k \in [d]$. We call $\preceq$ the *dominance order* on the set of transversals of $G$. It is the order introduced by Gale [14], usually called the Gale order.

Let $G$ be a multigraph with edge multiset $E$. Then for an edge $e = (i, j) \in E$, we define $\min(e) = i$, the smaller of the two vertices by convention.

**Definition 3.3.** We define the *opener* of $G$ as

$$\mathcal{A}(G) = \{\min(e) : e \in E\}.$$

We say an edge copy $e$ *opens* at $\min e$ and *closes* at $\max e$.

Then the opener is a transversal of $G$, hence $|\mathcal{A}(G)| = d$. We have that $\mathcal{A}(G) \preceq T$ for all transversals $T$ of $G$.

The opener is useful as it indexes cross-free graphs in the same class.

**Proposition 3.4.** *Suppose $G$ and $G'$ are cross-free graphs in the same class. Then if $\mathcal{A}(G) = \mathcal{A}(G')$ we get that $G = G'$.*

*Proof.* For a vertex $u$, we define $\mathrm{Cl}(u) = \{\min(e) : e \text{ is an edge copy of } G \text{ with } \max(e) = u\}$ to be the multiset of openers of the copies of $G$ closing at $u$. Similarly, we define $\mathrm{Op}(u)$ to be the multiset of edge copies of $G$ opening at $u$, so that $|\mathrm{Op}(u)|$ is the multiplicity of $u$ in $\mathcal{A}(G)$.

A multigraph is determined by the family $(\mathrm{Cl}_G(u))_u$, since every edge copy occurs once at its closer. Then we want to show that $\mathrm{Cl}_G(u) = \mathrm{Cl}_{G'}(u)$ for every $u$.

Since $G$ and $G'$ are in the same class, they have the same valencies at each vertex. Then $\mathrm{val}_G(u)$ can be broken into the edges which connect to it as a minimum and the edges which connect to it as maximum, hence $\mathrm{val}_G(u) = |\mathrm{Op}(u)| + |\mathrm{Cl}(u)|$, therefore $|\mathrm{Cl}_G(u)| = |\mathrm{Cl}_{G'}(u)|$ since $\mathcal{A}(G) = \mathcal{A}(G')$.

Suppose that $G \neq G'$. Let $u$ be the smallest vertex which satisfies $\mathrm{Cl}_G(u) \neq \mathrm{Cl}_{G'}(u)$. Then since the multisets have the same size, by counting there exists some vertex $t$ such that the multiplicity of $t$ in $\mathrm{Cl}_G(u)$ is greater than the multiplicity of $t$ in $\mathrm{Cl}_{G'}(u)$. Similarly, there exists some vertex $s$ such that the multiplicity of $s$ in $\mathrm{Cl}_G(u)$ is less than the multiplicity of $s$ in $\mathrm{Cl}_{G'}(u)$, and, exchanging the roles of $G$ and $G'$ if necessary, we may take $t < s$. Then $G$ has an edge copy $\{t, u\}$ with $t < s < u$.

Since $\mathcal{A}(G) = \mathcal{A}(G')$, the same number of copies open at $s$ in $G$ and in $G'$. Then every copy which opens at $s$ closes at some $r$, if $r < u$, then $\mathrm{Cl}_G(r)$ and $\mathrm{Cl}_{G'}(r)$ agree by minimality of $u$. Then $G$ and $G'$ have equally many copies which open at $s$ and close below $u$. At $u$, $G'$ has more, hence $G$ has more copies opening at $s$ and closing above $u$. Therefore $G$ contains the edge copy $\{s, \ell\}$ with $\ell > u$.

Then $G$ contains the copies $\{t, u\}$ and $\{s, \ell\}$ with $t < s < u < \ell$ on four distinct vertices which defines a cross and contradicts $G$ being cross-free. Hence openers determine cross-free graphs in the same class. □

### 3.3. **Independence.**

**Proposition 3.5** (Distinguishing records)**.** *Let $B_1, \ldots, B_s$ be weighted collections of records, and assign to each $B_i$ a distinct record $q_i$ of weight 1. Suppose that the $q_i$ are ordered so that $q_i$ can occur in $B_j$ only when $j \leq i$. Then no nonzero weighted list of the $B_i$ has all record weights zero. The same conclusion holds with the order reversed.*

*Proof.* In a proposed nonzero list choose the least index $i$ with nonzero weight $c_i$. At record $q_i$, every earlier contribution has weight zero and every later contribution is absent. Its total weight is therefore $c_i \neq 0$. For the reversed order start with the greatest index instead. This also gives an explicit recovery procedure: read the first distinguished record, subtract that many copies of its collection, and continue. □

We have that $\Phi(m_G) = \prod_e \pi_e$ over edge copies of $G$ with $\pi_e = x_a y_b - x_b y_a$ for $e = \{a, b\}$, $a < b$. For each edge, we choose a direction which is equivalent to choosing a smaller or larger vertex.

Then, in each monomial produced by a choice, the variable $x_v$ records how many copies were directed out of $v$ and the $x$-exponent vector is the multiset of chosen tails (a transversal of $G$). Within

a class, the $x$-exponents determine the monomial since $\deg(x_v) + \deg(y_v) = \mathrm{val}(v)$. Then we denote $\mu_T$ for the monomial of the $x$-exponent multiset $T$. Hence every monomial occurring in $\Phi(m_G)$ is $\mu_T$ for some transversal $T$ of $G$. The converse fails as Subsection 2.3 details.

**Proposition 3.6.** *Let $G$ be a multigraph. Then $\mu_{\mathcal{A}(G)}$ occurs in $\Phi(m_G)$ with coefficient exactly 1. Moreover, if $G$ and $G'$ are multigraphs in the same class and $\mu_{\mathcal{A}(G)}$ occurs in $\Phi(m_{G'})$, then $\mathcal{A}(G') \preceq \mathcal{A}(G)$.*

*Proof.* Expand $\Phi(m_G) = \prod_e \pi_e$. Each copy contributes one of its two vertices as the $x$-variable, a choice for every copy is a transversal $T$, the monomial produced is $\mu_T$, and its sign is $(-1)^{|L|}$ where $L$ is the set of copies choosing their larger vertex. So $\Phi(m_G)$ is the sum over transversals $T$ of $\mu_T$ weighted by $\sum(-1)^{|L|}$ over the choices realizing $T$.

Take $T = \mathcal{A}(G)$. Every choice satisfies $T(e) \geq \min e$, so $\sum_e T(e) \geq \sum_e \min e$ with equality exactly when $T(e) = \min e$ for every copy. Equal multisets have equal sums, so the only choice realizing $\mathcal{A}(G)$ takes every minimum. Its set $L$ is empty and the coefficient is 1.

For the second statement, let $T$ be any transversal of $G$ and consider $\mathcal{A}(G)$ and $T$ in weakly increasing order as $a_1 \leq \cdots \leq a_d$ and $t_1 \leq \cdots \leq t_d$. Fix $k$ and let $S$ be the $k$ copies realizing $t_1, \ldots, t_k$. Each copy in $S$ has its minimum at most its chosen vertex, so at least $k$ entries of $\mathcal{A}(G)$ are at most $t_k$, whence $a_k \leq t_k$. Thus $\mathcal{A}(G) \preceq T$. Now if $\mu_{\mathcal{A}(G)}$ occurs in $\Phi(m_{G'})$ then $\mathcal{A}(G)$ is a transversal of $G'$, and applying this to $G'$ gives $\mathcal{A}(G') \preceq \mathcal{A}(G)$. □

Note that no cross-free hypothesis appears in Proposition 3.6. Cross-freeness enters only through Proposition 3.4, which is what guarantees that the openers $\mathcal{A}(G)$ are distinct as $G$ ranges over the cross-free members of a class.

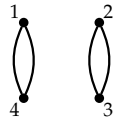
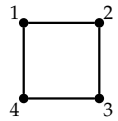
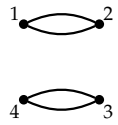
| | {1,1,2,2} | {1,1,2,3} | {1,1,3,3} |
|---|---|---|---|
|  | 1 | −2 | 1 |
|  | 0 | 1 | −1 |
|  | 0 | 0 | 1 |

FIGURE 7. The three cross-free members of the class $(2, 2, 2, 2)$

**Proposition 3.7.** *Let $\mathcal{O}(G)$ be a class and let $G_1, \ldots, G_m$ be the cross-free multigraphs in $\mathcal{O}(G)$. Then $\Phi(m_{G_1}), \ldots, \Phi(m_{G_m})$ are linearly independent.*

*Proof.* Order the distinct openers by any total order extending dominance. In a nonzero weighted list of cross-free graphs, choose the graph $G$ whose opener is first among those with nonzero weight. An orientation of another graph $H$ can produce $\mathcal{A}(G)$ only if $\mathcal{A}(H) \preceq \mathcal{A}(G)$ by Proposition 3.6. No earlier graph has nonzero weight, and no other cross-free graph has this opener by Proposition 3.4. The unique orientation of $G$ producing its opener therefore leaves exactly the weight of $G$ at that record. It cannot cancel. This proves independence by Proposition 3.5. □

**Example 3.8.** Take the class $\mathcal{O}(C_4)$ of Figure 2. Three of its six members are cross-free, and listed so that their openers increase in dominance they are

$$G_1 = m_{14}^2 m_{23}^2, \qquad G_2 = m_{12}m_{23}m_{34}m_{14}, \qquad G_3 = m_{12}^2 m_{34}^2,$$

with openers $\{1,1,2,2\}$, $\{1,1,2,3\}$, $\{1,1,3,3\}$. The table in Figure 7 counts signed orientations with these three records. The first record recovers the weight of $G_1$; after removing it, the second recovers the weight of $G_2$, and then the third recovers the weight of $G_3$.

3.4. **Kostka numbers.** Combining straightening with independence gives a combinatorial characterization for the basis $\Phi$ and we determine ranks for the basis.

The statement of the following proposition is classical: the basis goes back to Kempe [27] (see also [26, Theorem 2.2]), and $\ker(\Phi) = (\mathcal{P})$ is Hodge's straightening law [24]. The proof through openers is new, and it is the argument that we adapt to the truncations in Section 7.

**Proposition 3.9.** *Let $\mathcal{O}(G)$ be a class with valency vector $\mathbf{d}$. Then*

$$\{\Phi(m_H) : H \in \mathcal{O}(G) \text{ is cross-free}\}$$

*is a $\Bbbk$-basis of $\Phi\big(\Bbbk[m_{ij}]_{\mathbf{d}}\big)$. Hence $\ker(\Phi) = (\mathcal{P})$.*

*Proof.* Apply the terminating smoothing procedure to each graph in a weighted list. Each local replacement preserves all signed orientation records: on the four affected vertices the two possible tail choices giving the same record either agree on the two sides or cancel with opposite signs. The unchanged edges can be oriented independently. Thus every list has the same records as a list of cross-free graphs.

If the original list is a relation, so is this terminal list. Its first nonzero opener record would contradict Proposition 3.7; consequently it is the empty weighted list. Every relation has therefore been reduced to zero by Plücker moves. Conversely, each such move preserves the records. This proves both the basis statement and the description of the kernel, for arbitrary weighted lists rather than only for individual graphs. □

Let $\mathbf{d} = (a_1, \ldots, a_n)$ with $a_i > 0$ for all $i \in [n]$. Then $|\mathcal{O}(G)|$ is the coefficient of $x_1^{a_1} \cdots x_n^{a_n}$ in

$$\prod_{1 \le i < j \le n} \frac{1}{1 - x_i x_j} \in \mathbb{Z}[[x_1, \ldots, x_n]].$$

This is since a monomial of $\Bbbk[m_{ij}]$ is determined by the multiplicity $c_{ij} \ge 0$ of each variable $m_{ij}$, and its valency vector has $i$-th entry $\sum_{j \ne i} c_{ij}$. Expanding each factor as $\sum_{c \ge 0} (x_i x_j)^c$ and multiplying, the coefficient of $x_1^{a_1} \cdots x_n^{a_n}$ counts the tuples $(c_{ij})$ with $\sum_{j \ne i} c_{ij} = a_i$ for every $i$. It follows that if $\mathcal{W}(G) = (1, \ldots, 1)$ with $2k$ entries, then $|\mathcal{O}(G)| = (2k-1)!!$.

We now count the basis. The count is naturally stated for a second interleaving condition, and the two agree.

**Definition 3.10.** Let $G$ be a multigraph with $\{a, b\}$ and $\{c, d\}$ being edges in $G$. We say that the edges $\{a, b\}$ and $\{c, d\}$ are *nested* if

$$a < c < d < b.$$

If $G$ has no nested pair of edge copies we call $G$ *nesting-free*, and we denote $\mathcal{N}(G)$ for the set of nesting-free members of $\mathcal{O}(G)$.

A crossing and a nesting are the two ways two edge copies can interleave, so neither condition implies the other: $m_{13}m_{24}$ is nesting-free and not cross-free, while $m_{14}m_{23}$ is cross-free and not nesting-free. The opener identifies the two families.

For graphs without repeated edges, de Mier [9, Theorem 3.3] proves that noncrossing and nonnesting graphs are equinumerous; the following proposition gives a bijection for multigraphs that preserves the opener.

**Proposition 3.11.** *Let $\mathcal{O}(G)$ be a class with $a = \mathcal{W}(G)$ and $d = d(G)$. A multiset $U$ of $d$ vertices is the opener of a cross-free member of $\mathcal{O}(G)$ if and only if it is the opener of a nesting-free member, and in each case that member is unique. Hence the cross-free members of $\mathcal{O}(G)$ are equinumerous with $\mathcal{N}(G)$.*

*Proof.* Let $u_i$ be the multiplicity of $i$ in the prescribed opener $U$, and let $c_i = a_i - u_i$, $U_i = \sum_{j \le i} u_j$ and $S_i = \sum_{j \le i} c_j$. Such data are realizable exactly when $0 \le u_i \le a_i$, $U_n = S_n = d$, and $U_{i-1} \ge S_i$ for every $i$. Indeed, at vertex $i$ one must close $c_i$ previously opened copies before opening its $u_i$ new copies; this condition says exactly that enough old copies are available.

For a cross-free graph close the most recent available openers first (a stack). For a nesting-free graph close the earliest ones first (a queue). The stack rule avoids crossings and the queue rule avoids nestings. Conversely, violating the respective rule forces a crossing or a nesting with a copy closed later. Copies with identical opener labels are indistinguishable and do not affect uniqueness. Thus both rules work for exactly the same data and produce a unique graph. Keeping $U$ and changing stack to queue is the required bijection. □

Figure 8 shows the two rules on eight vertices of valency one, with the openers filled.

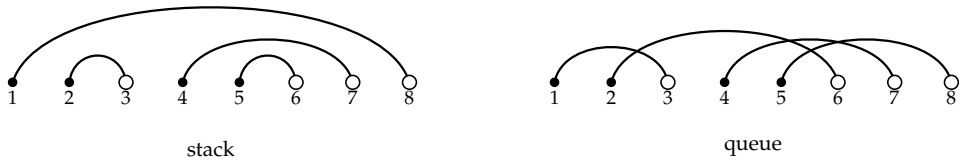



FIGURE 8. The cross-free and the nesting-free member with opener $\{1, 2, 4, 5\}$

That nesting-free graphs index a basis is equivalent to the classical standard monomial basis [24, 11]; the proof below uses only openers.

**Proposition 3.12.** *Suppose $G$ and $G'$ are nesting-free graphs in the same class with $\mathcal{A}(G) = \mathcal{A}(G')$. Then $G = G'$. Consequently Proposition 3.9 holds with nesting-free in place of cross-free.*

*Proof.* The queue construction in Proposition 3.11 proves uniqueness. The opener-record argument of Proposition 3.7 uses only that uniqueness, so it proves independence for nesting-free graphs as well. There are as many of them as cross-free graphs. Since the latter already form a basis, the former independent family of the same size is also a basis. No representation-theoretic identification is needed. □

Let $d \ge 1$. A *tableau* of shape $(d, d)$ is a pair of sequences $u = (u_1, \dots, u_d)$, $v = (v_1, \dots, v_d)$ of positive integers, written as the two rows of a $2 \times d$ array. It is *semistandard* if the rows are weakly increasing and the columns are strictly increasing, that is $u_t < v_t$ for every $t \in [d]$ with $u_1 \le \cdots \le u_d$ and $v_1 \le \cdots \le v_d$. The *content* of a tableau is the vector whose $i$-th entry is the number of times $i$ occurs among the $u_t$ and $v_t$. For $a = (a_1, \dots, a_n)$ we denote $Y(d, d, a)$ as the semistandard tableaux of shape $(d, d)$ and content $a$. The *Kostka number* $K_{(d,d),a}$ is $|Y(d, d, a)|$ [33, Section 7.10], [13].

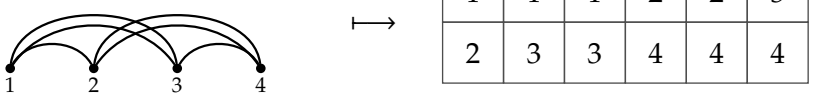



FIGURE 9. The nesting-free multigraph $m_{12}m_{13}^2m_{24}^2m_{34}$ and its tableau

**Proposition 3.13.** *Let $a = \mathcal{W}(G)$ and $d = d(G)$. Then listing the edges of a nesting-free member as the columns of a $2 \times d$ array, openers along the first row and closers along the second, is a bijection*

$$\mathcal{N}(G) \longleftrightarrow Y(d, d, a),$$

*and hence $|\mathcal{N}(G)| = K_{(d,d),a}$.*

*Proof.* Sort the edge columns first by opener and, for equal openers, by closer. If the closer row decreased between two columns with distinct openers, the two edges would be nested. Thus a nesting-free graph gives two weakly increasing rows, and every column is strictly increasing because edges have distinct endpoints. Conversely, a semistandard tableau gives a multigraph of the prescribed valencies, and its increasing closer row excludes a nesting. These constructions undo one another. □

The tuples $(u_i)$ with $U_n = d$ and $U_{i-1} \geq S_i$ appearing in the proof of Proposition 3.11 are therefore in bijection with $Y(d, d, a)$, which gives the standard tuple description of $K_{(d,d),a}$ [33, Section 7.10]. If $\mathcal{W}(G) = (1, \ldots, 1)$ with $2k$ entries, then $|\mathcal{N}(G)| = C_k$, the $k$-th Catalan number [33, Exercise 6.19], and by Proposition 3.11 the cross-free members are $C_k$ in number as well.

## 4. Forests

From now on we assume $\mathrm{char}(\mathbb{k}) = 2$. Then $\varepsilon(H) = 1$ for every orientation, so the sign drops out of Proposition 2.5 and the only coefficients are the binomials.

4.1. **Monomials.** In this subsection we classify the monomials which are in the kernel and describe the monomials spanning the image of $t_{2,r}$.

**Proposition 4.1.** *Let $C_n$ be a cycle. Then $C_n \in \ker(t_{2,r})$ for all $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$.*

*Proof.* A surviving orientation has at most one outgoing edge at each vertex. On a cycle every vertex must then have exactly one outgoing edge, so there are precisely the two cyclic orientations. Reversal pairs them, preserving every vertex's incoming and outgoing counts. The two records cancel in characteristic two. Every other orientation violates the outgoing bound and is discarded. The same argument includes a doubled edge: distinguish its two copies while counting orientations. □

**Proposition 4.2.** *Let $\mathcal{Z}_n$ be a star. Then $\mathcal{Z}_n \in \ker(t_{2,r})$ iff $n \geq r + 1$ for $r \neq \infty$, and $\mathcal{Z}_n \notin \ker(t_{2,\infty})$ for all $n$.*

*Proof.* At the center an orientation with $a$ outgoing edges has $n - a$ incoming edges. It is allowed exactly when $a \leq 1$ and $n - a \leq r - 1$. No such choice exists if $n \geq r + 1$. If $n \leq r$, choose one leaf as root and direct every edge toward it. Different chosen leaves give different records, each occurring once. The center can also be a root when $n \leq r - 1$. These are all the allowed orientations. With no incoming bound, the leaf-root orientations always remain. □

**Definition 4.3.** We define $K(\tau_{\mathfrak{n}}) = \{m \in \ker(\tau_{\mathfrak{n}}) : m \text{ is a monomial}\}$.

Recall that an *r-forest* is a forest in which every vertex has valency at most $r$, and by convention an $\infty$-forest is a forest.

**Proposition 4.4.** *Let $F$ be an $r$-forest with trees $T_1, \ldots, T_k$. Then we have that $F \notin \ker(t_{2,r})$.*

*Proof.* Choose a leaf in each component and direct every edge toward that leaf. A nonroot vertex has one outgoing edge and at most $r - 1$ incoming edges; a root leaf has no outgoing edge and one incoming edge. Thus the record is allowed. It occurs just once: in a tree the chosen sink forces the direction of every edge along the unique path to it. Equivalently, removing a leaf determines its edge direction from its outgoing count, and induction determines all remaining directions. This record cannot cancel. □

**Proposition 4.5.** *We have that* $K(t_{2,r}) = (\mathcal{C}, \mathcal{Z}_{r+1})$ *and hence* $\mathrm{im}(t_{2,r})$ *is spanned by the* $t_{2,r}(m_F)$ *with* $F$ *an* $r$*-forest.*

*Proof.* If a graph contains a cycle, fix that cycle. Among the orientations that meet the outgoing bound, its edges must be cyclically oriented; reverse just those edges. This is a fixed-point-free pairing with identical records, even in the presence of other edges. If the graph contains an $(r+1)$-star, its center cannot satisfy both bounds.

Conversely, a graph containing no cycle is a forest. If its maximum valency exceeds $r$, the incident edges at that vertex contain an $(r+1)$-star. Every remaining graph is an $r$-forest, and its leaf-root record survives by Proposition 4.4. Thus the forbidden subgraphs are exactly the stated ones, and deleting them from a weighted list leaves only $r$-forests. □

The proof of Proposition 4.4 shows more, and we record it since the rest of the paper uses the result.

For $r = 2$ and $F$ a single path the formula specializes to [4, Proposition 2.5]: the image of a path is the product of $x_v y_v$ over its interior vertices with the minor of its endpoints.

**Proposition 4.6.** *Let* $F$ *be an* $r$*-forest with trees* $T_1, \dots, T_k$*. The orientations of* $F$ *whose term survives modulo* $\mathfrak{m}_r$ *are exactly those obtained by choosing one vertex* $z_j$ *of valency at most* $r-1$ *in each tree and directing every edge towards the chosen vertex of its tree. The image is*

$$t_{2,r}(m_F) = \sum_{z} \prod_{v \neq z_1,\dots,z_k} x_v y_v^{\mathrm{val}(v)-1} \prod_{j=1}^{k} y_{z_j}^{\mathrm{val}(z_j)},$$

*the sum being over these choices, and its terms are pairwise distinct monomials with coefficient* 1*.*

*Proof.* An orientation survives exactly when $\mathrm{outval}(v) \leq 1$ and $\mathrm{inval}(v) \leq r-1$ for every $v$. A forest on $N$ vertices with $k$ trees has $N-k$ edges, so the outvalencies sum to $N-k$; if none exceeds 1 then exactly $k$ of them are 0, and a tree with $n_j$ vertices has $n_j - 1$ edges, so each tree contains exactly one vertex $z_j$ of outvalency 0. Every other vertex of $T_j$ has one outgoing edge, and following these edges gives a walk which cannot repeat a vertex, so it ends at $z_j$ and every edge is directed towards $z_j$. A vertex $v \neq z_j$ then has invalency $\mathrm{val}(v) - 1 \leq r-1$ automatically, while $z_j$ has invalency $\mathrm{val}(z_j)$, which is the stated bound. The $x$-support of the term of $z$ is the complement of $\{z_1, \dots, z_k\}$, so distinct choices give distinct monomials. □

For instance the tree of Figure 10 has the blocked vertex 5 of valency 3, and $t_{2,3}$ sends it to a sum of four terms.

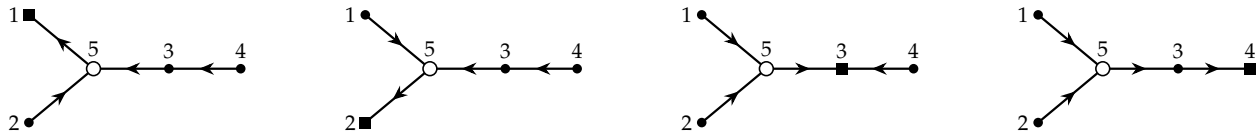



FIGURE 10. The surviving orientations of a tree at $r = 3$, one for each root

We often use notation $(1^{a_1}, 2^{a_2}, \dots)$ for a forest's shape where $a_i$ is the number of vertices of valency $i$.

4.2. **Linear forests.** Two linear forests have the same image only if they lie in the same valency class, so we fix a class. In this subsection we give a combinatorial bijection for the linear forest class which recovers its size. Throughout let $L = P_{n_1,\dots,n_k}$ be a linear forest.

**Definition 4.7.** The *linear forest class* of $L$ is

$$\mathcal{O}_2(L) = \{L' \in \mathcal{O}(L) : L' \text{ is a linear forest}\}.$$

More generally, for $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$ we define $O_r(F)$ to be the set of $r$-forests in $O(F)$, so that the linear forest class is $O_2(L)$.

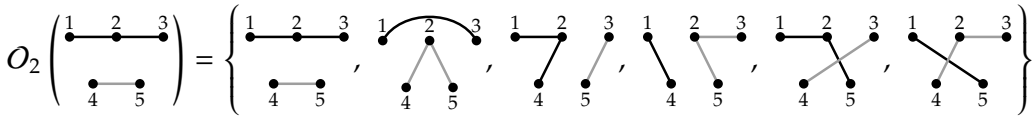



FIGURE 11. The linear forest class for $L = P_{3,2}$.

We use the letter $N$ for the total number of vertices in $L$, which is $N = \sum_{i=1}^{k} n_i$. Define the *end point set* to be $\mathrm{End}(L) = \{v \in L : \mathrm{val}(v) = 1\}$ and *interior point set* $\mathrm{Int}(L) = \{v \in L : \mathrm{val}(v) = 2\}$. Then $|\mathrm{End}(L)| = 2k$ and $|\mathrm{Int}(L)| = N - 2k$.

For linear forests $L, L'$, $L' \in O_2(L)$ iff $\mathrm{End}(L) = \mathrm{End}(L')$ and $\mathrm{Int}(L) = \mathrm{Int}(L')$ since if the endpoint set and interior set match, the valencies match so $L' \in O_2(L)$.

Let $P(L)$ be the perfect matchings of $\mathrm{End}(L)$. Then $|P(L)| = (2k-1)!! = \frac{(2k)!}{2^k\, k!}$. For lists of pairs from a perfect matching or endpoints, we always index its pairs $\{a_1, b_1\}, \ldots, \{a_k, b_k\}$ so that $a_j < b_j$ for all $j \in [k]$ and $a_1 < a_2 < \cdots < a_k$.

**Definition 4.8.** Let $L' \in O_2(L)$ with paths $Q_1, \ldots, Q_k$, where $Q_j$ has endpoints $a_j, b_j$. We define

$$\rho : O_2(L) \longrightarrow P(L), \qquad \rho(L') = \big\{\{a_1, b_1\}, \ldots, \{a_k, b_k\}\big\},$$

to be the *endpoint pairing* which matches a linear forest to its endpoints along paths.

Define $S(L)$ to be the set of ordered $k$-tuples of sequences $s = (s_1, \ldots, s_k)$ of distinct elements of $\mathrm{Int}(L)$ whose underlying sets partition $\mathrm{Int}(L)$. Then we have that $s$ is determined by two independent choices: the lengths $(m_1, \ldots, m_k)$ of its sequences, and the assignment of the elements of $\mathrm{Int}(L)$ to the positions those lengths provide. Hence $|S(L)| = \binom{N-k-1}{k-1}(N-2k)! = \frac{(N-k-1)!}{(k-1)!}$.

**Proposition 4.9.** *We have that $O_2(L)$ is in bijection with $P(L) \times S(L)$, hence*

$$|O_2(L)| = |P(L)| \cdot |S(L)| = (2k-1)!! \cdot \frac{(N-k-1)!}{(k-1)!}.$$

*Proof.* Let $p \in P(L)$ and $s = (s_1, \ldots, s_k) \in S(L)$. Index the pairs of $p$ as $\{a_j, b_j\}$ with $a_j < b_j$ for each $j$ and $a_1 < \cdots < a_k$, and let $L'$ be the disjoint union of the paths which traverse $a_j$ to $s_j$ to $b_j$. Its vertices of valency 1 are the $a_j, b_j$, that is $\mathrm{End}(L)$, and those of valency 2 are the entries of the $s_j$, that is $\mathrm{Int}(L)$; so $L' \in O_2(L)$. This indexing is the standard labeling of $L'$. For any $L' \in O_2(L)$ with paths $Q_1, \ldots, Q_k$ as in Definition 4.8, we define $\iota(L') \in S(L)$ to be the tuple whose $j$-th entry lists the interior vertices of $Q_j$ from $a_j$ to $b_j$, the *filling* of $L'$. Then $\big(\rho(L'), \iota(L')\big) = (p, s)$ by construction.

Conversely, a linear forest in $O_2(L)$ is recovered from its pairing and filling by the same construction, since its $j$-th path is determined by the traversal $a_j$ to $\iota(L')_j$ to $b_j$. The two assignments are therefore mutually inverse. □

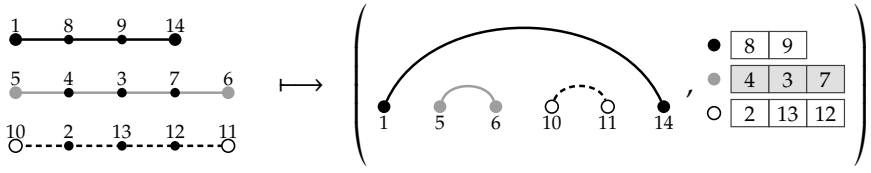



FIGURE 12. The bijection sending a linear forest to its perfect matching and interior

4.3. **Roots.** Fix $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$ and let $F$ be an $r$-forest with trees $T_1, \dots, T_k$ and $N$ vertices. A vertex is a *root* if $\mathrm{val}(v) < r$.

**Definition 4.10.** Let $F$ be an $r$-forest with trees $T_1, \dots, T_k$. The *root family* is the unordered partition $\mathcal{X}_r(F) = \{X_1, \dots, X_k\}$ with $X_i$ being the set of roots in $T_i$. Then we define $\bigcup_i X_i$ to be the *root set* $\mathcal{Y}_r(F)$, which is the underlying set of the root family.

Then a *root selection* is a transversal of the root family, hence $(z_1, \dots, z_k)$ with $z_i \in X_i$ is a root selection for $F$. By Proposition 4.6 the surviving orientations of $F$ are exactly the root selections, so

$$\#\{\text{terms of } t_{2,r}(m_F)\} = \prod_{i=1}^{k} |X_i|,$$

and the image of $F$ depends only on the root family. Being a root is a condition on valency, so the root set is determined by the class, while the root family is not. Let $u$ be the number of vertices of valency $r$ in $F$, then $|\mathcal{Y}_r(F)| = N - u$.

At $r = \infty$ every vertex is a root. At $r = 2$ the 2-forests are the linear forests, the roots are the endpoints, and the root family is the endpoint pairing $\rho$ of Definition 4.8. In general the root family is what replaces the endpoint pairing, and the endpoint pairing is the case in which every block has size two.




FIGURE 13. A root family and a root selection (squares) of a 3-forest

We now count the forests in a class. By Prüfer's theorem [33, Section 5.3] we have that for a tree $T$ on a fixed vertex set with prescribed valencies,

$$|O_\infty(T)| = \frac{(|T|-2)!}{\prod_{v \in T}(\mathrm{val}(v)-1)!},$$

where $O_\infty(T)$ is the set of forests in $O(T)$. Every such forest is a tree, since a forest in the class of $T$ has the same numbers of vertices and edges as $T$, hence exactly one component.

Denote $n_d$ as the number of vertices of $F$ of valency $d$. We define an *ordered split* to be a matrix $(n_{d,j})$ of non-negative integers with $d \geq 1$ and $1 \leq j \leq k$ such that $\sum_j n_{d,j} = n_d$ for all $d$, and we denote $|V_j| = \sum_d n_{d,j}$ for the size of its $j$-th part. The split is *admissible* if $\sum_{d \geq 1}(2-d)n_{d,j} = 2$ for every $j$. Note that this is just forest handshaking: a part carries a tree with the prescribed valencies exactly when its valencies sum to $2(|V_j| - 1)$. We define $\Im(F)$ as the set of admissible ordered splits.

**Proposition 4.11.** *Let $F$ be a forest with $k$ trees. Then*

$$|O_\infty(F)| = \frac{1}{k! \prod_d \big((d-1)!\big)^{n_d}} \sum_{o \in \Im(F)} \prod_{d \geq 1} \binom{n_d}{n_{d,1}, \dots, n_{d,k}} \prod_{j=1}^{k} (|V_j| - 2)!.$$

*Proof.* A forest in $O_\infty(F)$ is determined by the partition of its vertex set into the vertex sets of its trees together with a tree on each part, and conversely any such pair with every part admissible is a forest in $O_\infty(F)$, since the valencies are unchanged.

Count ordered tuples first. An ordered tuple $(V_1, \dots, V_k)$ of disjoint parts covering the vertices, with $n_{d,j}$ vertices of valency $d$ in $V_j$, is chosen by fixing the matrix $(n_{d,j})$ and then, for each $d$ independently, distributing the $n_d$ vertices of valency $d$ among the parts, which is $\binom{n_d}{n_{d,1}, \dots, n_{d,k}}$. The tuple carries a forest exactly when the split is admissible, and then Prüfer counts the trees on $V_j$ as

$(|V_j|-2)!/\prod_{v\in V_j}(\mathrm{val}(v)-1)!$. The denominator depends only on the valencies present, so multiplying over the parts gives $\prod_d((d-1)!)^{n_d}$, the same for every split.

Finally a forest determines its parts as a set and not as a tuple, and the parts are pairwise distinct, being disjoint and nonempty, so each forest arises from exactly $k!$ ordered tuples. □

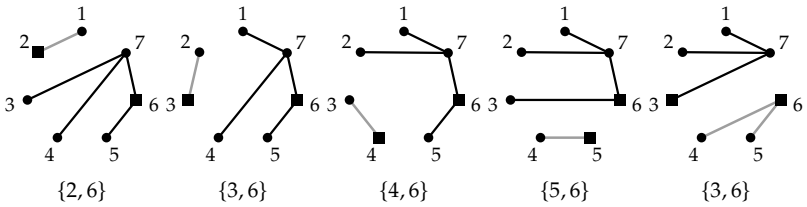



FIGURE 14. The five cross-free 3-forests of the class $(1^5, 2, 3)$

For a linear forest, $n_1 = 2k$ and $n_2 = N - 2k$, we get

$$\frac{1}{k!}\cdot\frac{(2k)!}{2^k}\sum_{m_1+\cdots+m_k=N-2k}\binom{N-2k}{m_1,\ldots,m_k}\prod_j m_j! = (2k-1)!!\cdot\frac{(N-k-1)!}{(k-1)!},$$

since every summand equals $(N-2k)!$ and there are $\binom{N-k-1}{k-1}$ compositions. This recovers Proposition 4.9.

4.4. **Root-selection inventories.** A root selection $R$ in a fixed valency class determines the term

$$\mu_R = \prod_{v\notin R} x_v y_v^{\mathrm{val}(v)-1}\prod_{v\in R} y_v^{\mathrm{val}(v)}.$$

We will keep only the combinatorial record $R$. Write $x_R$ for a record carrying the set $R$, and write $x_v = x_{\{v\}}$. An *inventory* assigns weights in $\Bbbk$ to finitely many records. When the weights are 0 and 1, adding inventories takes symmetric difference. General weights are retained throughout, so the arguments apply over every field of characteristic two, not just $\mathbb{F}_2$.

Juxtaposition of inventories means: choose a record from each, discard choices with a repeated vertex, and record the union with the product of the chosen weights. Add weights when several choices give the same union. Thus

$$x_R x_S = \begin{cases} x_{R\sqcup S} & R\cap S = \emptyset,\\ 0 & \text{otherwise.}\end{cases}$$

This is a rule for counting choices. In particular $L_B = \sum_{v\in B} x_v$ records one choice from $B$. Two independent choices from the same set cancel: equal choices are discarded, and unequal choices are paired by exchanging their order. We use the notation $L_B L_B = 0$ for this pairing.

**Proposition 4.12.** *Let $F$ have eligible-root blocks $X_1,\ldots,X_k$. Its root-selection inventory is $L_{X_1}\cdots L_{X_k}$. Under the record identification $x_R \leftrightarrow \mu_R$, this inventory is $t_{2,r}(m_F)$.*

*Proof.* Choose one vertex from each tree. The blocks are disjoint, so the chosen set determines the choices uniquely. Directing each tree toward its chosen vertex gives the unique surviving orientation with that record, by Proposition 4.6. □

**Definition 4.13.** For a forest $F$ let $P_F$ be its *unrestricted root inventory*: one chooses one vertex from each component, allowing every vertex. For a positive valency vector $\mathbf{d}$, let $\Pi(\mathbf{d})$ be the collection of inventories obtained from weighted lists of forests with these valencies. An unrealizable class contributes only the zero inventory. For $U\subseteq V$ let $\pi_U$ discard every record meeting $U$.

Thus $P_F = \prod_T L_{V(T)}$ is a compact notation for the choice rule. For $U = \{v : d_v = r\}$ the allowed inventory is $\pi_U P_F$. All inventories occurring in the class image are precisely $\pi_U \Pi(\mathbf{d})$.

**Lemma 4.14** (Moving pairs of branches)**.** *For distinct vertices $a, b$ with $d_b \geq 3$,*

$$\Pi(\mathbf{d}) = \Pi(\mathbf{d} + 2\mathbf{e}_a - 2\mathbf{e}_b).$$

*Proof.* If $a, b$ lie in the same tree, move two branches at $b$ not containing $a$ to $a$. This preserves the component vertex sets and hence every root-selection record.

If they lie in different trees, choose three branches at $b$ and call their vertex sets $B_1, B_2, B_3$. Form three forests by moving respectively $B_1 \cup B_2$, $B_1 \cup B_3$, and $B_2 \cup B_3$ to $a$. For two fixed vertices in the union of the affected trees, being in opposite trees is the exclusive-or of their two component-membership indicators. Each vertex in a chosen branch changes its indicator in exactly two of the three moves; every other vertex changes it zero times. The parity of the three separation indicators is therefore its original separation indicator. Thus each two-root record occurs in the three new forests with exactly its original parity. Choices in the other components are unchanged. Their sum has inventory $P_F$ (Figure 15).

This proves one inclusion by explicit moves. The reverse moves are permitted because the new valency of $a$ is at least three. Repeating the moves shows that the inventories depend only on the parities of the positive valencies and their sum. □

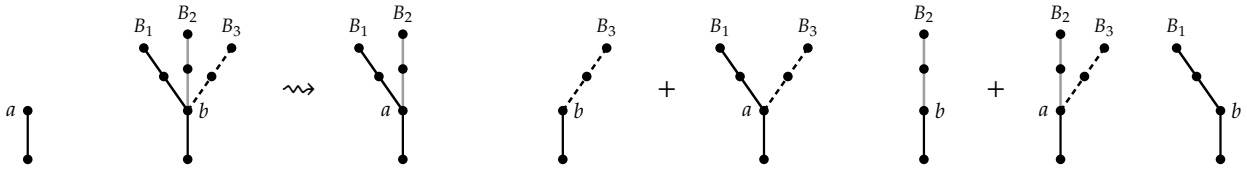



FIGURE 15. Moving two of the three branches $B_1, B_2, B_3$ at $b$ to $a$, in all three ways

Fix an anchor $a \in V$ and a new vertex $v$. Define $D_a$ on a record by retaining it only when it contains $a$, and then deleting $a$. Define $\eta_v(P) = P + x_v D_a P$: retain every record and also make a copy with $a$ replaced by $v$ whenever possible. These are explicit selection operations; no differentiation is involved.

**Lemma 4.15** (Inserting a vertex of valency two)**.** *Every inventory of $\Pi(\mathbf{d}, 2)$ has a unique expression*

$$\eta_v(P) + x_v Q, \qquad P \in \Pi(\mathbf{d}), \quad Q \in \Pi(\mathbf{d} + 2\mathbf{e}_a).$$

*Every expression of this form is attainable. The same statement holds after discarding records meeting any fixed set disjoint from $\{a, v\}$.*

*Proof.* Suppress $v$, joining its two neighbors, to obtain a forest $G$. If it is inserted in component $j$, the old selections remain, and the new selections choose $v$ and one vertex from each other component. Write $P_j$ for the latter choices before adding $v$. Inserting in the component containing $a$ gives exactly $\eta_v(P_G)$.

Compare insertion in another component $j$ with insertion in the anchor component. Their common old selections cancel. The remaining inventory is $x_v(P_j + D_a P_G)$. Delete an edge $bc$ of component $j$ and add $ab, ac$ (Figure 16). The two components become one; choosing its root in either of the old components gives precisely $P_j + D_a P_G$. The new valencies are $\mathbf{d} + 2\mathbf{e}_a$.

Conversely, in any forest with these new valencies the anchor has at least three incident edges. Delete two of them, say $ab, ac$, and add $bc$. This separates its tree into two nontrivial trees. The sum of the two ways of inserting $v$ in those trees recovers $x_v$ times the original inventory. Hence both kinds of terms are attainable.

Discarding records containing $v$ recovers $P$ uniquely. Remove $\eta_v(P)$ from the inventory and erase the now compulsory mark $v$; this recovers $Q$ uniquely. This recovery procedure also proves the last assertion: filtering labels other than $a, v$ commutes with every selection operation just used. □

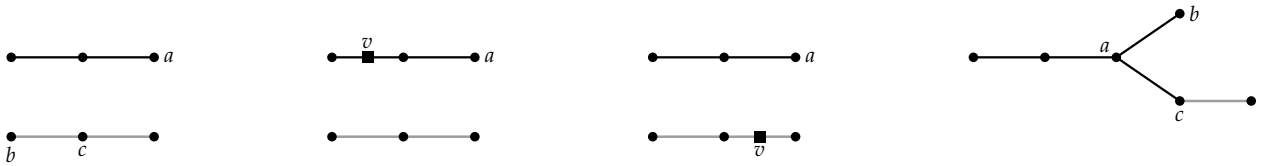



FIGURE 16. A forest $G$, the insertions of $v$ in the anchor component and in another component, and the merged forest

## 5. CUTS

The relations of the next section are all instances of one mechanism, and this section describes it. Throughout fix $r$ and let $F$ be an $r$-forest with root family $\{X_1, \dots, X_k\}$ and root set $\mathcal{Y}_r(F)$ of size $m = N - u$.

5.1. **Cut space.** Let $\mathcal{K}_m$ be the complete graph on the $m$ roots, with edge set $E$ and vertex set $V = \mathcal{Y}_r(F)$. Throughout this section every coefficient that appears is 0 or 1, and we add sets by symmetric difference. We write $\mathbb{F}_2[V]$ for the set of subsets of $V$ and $\mathbb{F}_2[E]$ for the set of subsets of $E$, each with addition $S + T = S \triangle T$, so that a sum of monomials with coefficients in $\mathbb{F}_2$ is the same as a set of monomials.

For $S \subseteq V$ the *cut* of $S$ is

$$\delta(S) = \{e \in E : |e \cap S| = 1\},$$

the set of edges of $\mathcal{K}_m$ with exactly one endpoint in $S$, and we call $\delta\colon \mathbb{F}_2[V] \to \mathbb{F}_2[E]$ the *cut map*. Its image is the *cut space* of $\mathcal{K}_m$ [10, Section 1.9]. An edge meets $S \triangle T$ in exactly one endpoint exactly when it meets exactly one of $S$ and $T$ in exactly one endpoint, so

$$\delta(S \triangle T) = \delta(S) \triangle \delta(T)$$

(see Figure 17). The only subsets with empty cut are $\emptyset$ and $V$, and we write $\ker(\delta) = \{\emptyset, \mathcal{Y}_r(F)\}$ for them. Indeed, if $S$ is neither $\emptyset$ nor $V$ then there are $u \in S$ and $v \notin S$, and $\{u, v\}$ is an edge of $\mathcal{K}_m$ lying in $\delta(S)$, while neither $\emptyset$ nor $V$ separates any pair. Hence two subsets have the same cut exactly when they are equal or complementary, and the cut space has $2^{m-1}$ elements.

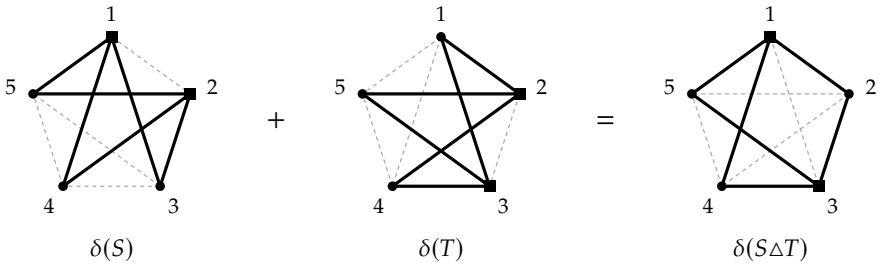



FIGURE 17. The cut map on $\mathcal{K}_5$ with $S = \{1, 2\}$ and $T = \{2, 3\}$

Since the monomials occurring in the image of a class are indexed by the possible root sets of selections (Subsection 4.4), a single $\mathcal{K}_m$ serves the whole class.

5.2. **Groves.** Fix $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$. We call a forest with two trees a *grove*, and an $r$-forest with two trees an *$r$-grove*. Throughout this subsection $G$ is an $r$-grove, we denote $\mathcal{Y} = \mathcal{Y}_r(G)$ for its root set, and we let $m = |\mathcal{Y}|$. Recall that $\mathcal{Y}$ is determined by the class, so a single $\mathcal{K}_m$ is sufficient for all of $\mathcal{O}_r(G)$.

**Definition 5.1.** For a grove $G'$ with eligible-root blocks $X_1, X_2$, its *root split* $\sigma(G')$ is the unordered bipartition $\{X_1, X_2\}$ of $\mathcal{Y}$. When writing a symmetric difference of root splits, choose either block as representative. Changing that choice complements the final subset, so the condition that the result be $\emptyset$ or $\mathcal{Y}$ is independent of every choice.

By Proposition 4.12 the image $t_{2,r}(m_{G'})$ is determined by the root family of $G'$, which for a grove is recorded by $\sigma(G')$. The cut map turns that record into the image.

**Proposition 5.2.** *Let $G'$ be an $r$-grove. Then, identifying $\mu_{\{a,b\}}$ with the edge $\{a, b\}$ of $\mathcal{K}_m$,*

$$t_{2,r}(m_{G'}) = \delta(\sigma(G')).$$

*Proof.* A pair $\{a, b\}$ is a root selection exactly when its vertices lie in different trees. This says exactly that it crosses the root split. Every such pair specifies one choice in each tree and therefore has weight one. No other pair occurs. □

**Definition 5.3.** A *cut relation* in a grove class is a weighted list of groves for which every pair of eligible roots is separated with total weight zero. Write $\mathcal{D}(G)$ for these lists and $\mathcal{D} = \bigcup_G \mathcal{D}(G)$ over all grove classes.

Here is a direct test. Choose one block $X_i$ of each grove $G_i$ in a weighted list $\sum_i c_i G_i$, and let

$$q_v = \sum_{i:\, v \in X_i} c_i.$$

The total separation weight of $\{v, w\}$ is $q_v + q_w$: a block containing both is counted twice and one containing neither is counted zero times. Thus the list is a cut relation exactly when all the $q_v$ are equal. Changing the block chosen for $G_i$ adds $c_i$ to every $q_v$, preserving this test. For weights in $\mathbb{F}_2$ it says that the chosen blocks have symmetric difference $\emptyset$ or $\mathcal{Y}$.

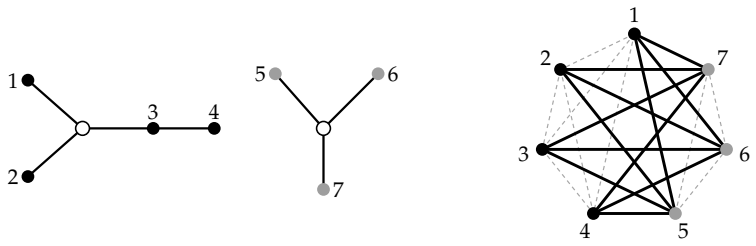



FIGURE 18. Root selections pick one root from each tree

**Proposition 5.4.** *Let $G_1, \ldots, G_n$ be distinct $r$-groves in a class. Then*

$$m_{G_1} + \cdots + m_{G_n} \in \ker(t_{2,r}) \qquad \textit{if and only if} \qquad \sigma(G_1) \triangle \ldots \triangle \sigma(G_n) \in \{\emptyset, \mathcal{Y}\}.$$

*Proof.* For each root pair, count the groves in which its endpoints lie on opposite sides of the chosen split. Its parity is the sum of the two endpoint-membership parities in the symmetric difference of the chosen blocks. All pairs have even count exactly when those membership parities are all equal, that is, when the difference is empty or all of $\mathcal{Y}$. The weighted version is the same test with the $q_v$ above. □

A monomial of the class which is not an $r$-forest lies in $K(t_{2,r})$ by Proposition 4.5, so adding one to a relation changes neither side of Proposition 5.4. Hence the criterion describes $\ker(t_{2,r})$ on the class and not only its forest part.

The cut description also counts the missing coordinates. Fix a root $o$. A cut is determined by its $m-1$ pairs incident to $o$, since for $v, w \neq o$ its value on $vw$ is the sum of its values on $ov$ and $ow$. The latter $m-1$ values can be prescribed freely. Thus the full cut space has dimension $m-1$, and the class image misses at least $\binom{m}{2} - m + 1$ of the pair coordinates. Equivalently, the triangles $ovw$, one for each pair $v, w \neq o$, form a basis of even-degree edge sets: remove all edges not incident to $o$ using their triangles; the even-degree condition then removes the remaining edges too. This gives $\binom{m-1}{2}$ independent parity tests, the usual cycle-space count [10, Section 1.9].

5.3. **Spanning.** We first allow two kinds of graph operations: smooth a crossing by its Plücker replacement, and erase a graph containing a cycle or an $(r+1)$-star. All of these preserve the allowed orientation records.

**Proposition 5.5.** *Every weighted list in a valency class reduces by these operations to a weighted list of cross-free $r$-forests in the same class. Equivalently, the class is spanned by these forests modulo $(\mathcal{P}, C, \mathcal{Z}_{r+1})$.*

*Proof.* First use the terminating smoothing algorithm of Proposition 3.2. Then erase every terminal graph containing a forbidden subgraph. By Proposition 4.5 the graphs left are exactly cross-free $r$-forests. Both steps preserve the records. □

The remaining data are especially simple:

$$F' \longmapsto \mathcal{X}_r(F') \longmapsto \{\text{one root chosen from each block}\}.$$

There are two possible sources of relations. Distinct forests can have the same root family; or different root families can have selection inventories whose weighted sum is zero. We call these the *spider* and *rainbow* mechanisms, respectively. More formally, group a weighted list by root family and replace each group by one chosen representative with its total weight. The differences inside the groups are of the first kind; the remaining relation is of the second kind.

For groves, records are pairs and these relations are exactly the cut relations. For a single tree every eligible vertex is a possible root, so all trees of the class have the same inventory. For three or more trees, records are transversals of several blocks; the same two sources of cancellation remain possible, but pair cuts alone no longer record all selections. Proposition 7.9 will show that relations with these larger records still follow from local grove relations after attaching further edges.

## 6. Relations

Every relation in this section is a cut relation in the sense of Proposition 5.4. There are two families, one for each parity of $r$, and the base cases $r = 2$ and $r = 3$ are binomials. At $r = 2$ the binomial is an identity in the image that also appears in Casabella and Yu [4, Corollary 2.7]; that it is a minimal generator of the kernel, and the families for $r > 2$, are new. The base case $r = 3$ is Proposition 6.2.

6.1. **Rainbows and spiders.** Let $r$ be even, let $b, a_1, \ldots, a_{r+2}$ be distinct, and let $A = \{a_1, \ldots, a_{r+2}\}$. For a pair $\{s, t\} \subseteq A$ we define

$$F_{st} = m_{st} \prod_{w \in A \setminus \{s,t\}} m_{wb},$$

the grove joining $s$ to $t$ and the center $b$ to the remaining $r$ leaves. We define the *rainbow relation at $r$* as

$$\sum_{\{s,t\} \in M} F_{st}$$

for a perfect matching $M$ of $A$, and denote the set of rainbow relations as $\mathcal{R}_r$. Every term lies in the class $(1^{r+2}, r)$, the relation has $(r+2)/2$ terms, and each term has degree $r+1$.

At $r = 2$ let $a, b, c, d, e$ be distinct and take $M = \{\{a, c\}, \{d, e\}\}$. Then the rainbow relation is

$$m_{ab}m_{bc}m_{de} + m_{bd}m_{be}m_{ac}$$

and we call the elements of $\mathcal{R}_2$ *rainbows*. The identity also appears in [4, Corollary 2.7], where it is not considered as a kernel generator. We also define the *degenerate rainbow relation* as

$$m_{ab}m_{bc}m_{cd} + m_{ac}m_{bc}m_{bd}$$

and denote the set of degenerate rainbows as $\mathcal{R}^\circ$.

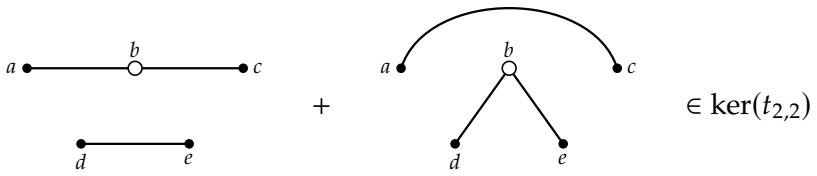



FIGURE 19. The rainbow relation $\mathcal{R}_2$ in the class $(1^4, 2)$

**Proposition 6.1.** *$\mathcal{R}_r \subseteq \ker(t_{2,r})$ for every even $r$. In the class $(r, 1^{r+2})$, a relation with a nonzero common leaf-incidence weight can exist only when $r$ is even. In particular no perfect-matching rainbow exists when $r$ is odd.*

*Proof.* The eligible roots are the leaves. For a fixed pair of leaves $x, y$, the pair is a selection of $F_{st}$ exactly when $st$ contains exactly one of $x, y$. In a perfect matching, either $x, y$ belong to the same matched pair and the count is zero, or they belong to different pairs and the count is two. Every root pair therefore cancels.

More generally let $w_p$ be the total weight of leaf edges incident to $p$ in a list $\sum c_{st}F_{st}$. The root pair $xy$ has weight $w_x + w_y$, so a relation has $w_p = c$ independent of $p$. Counting all leaf incidences gives $(r+2)c = 2\sum c_{st} = 0$. If $r$ is odd this forces $c = 0$. This does not rule out zero-incidence relations: for instance $F_{st} + F_{sx} + F_{tx}$ is a relation for distinct leaves $s, t, x$. □

Now let $r \geq 3$ and let $u, v, a_1, \ldots, a_r, b_1, \ldots, b_r$ be distinct. We define the *spider relation at $r$* as

$$\prod_{i=1}^{r} m_{a_i v} \prod_{i=1}^{r} m_{b_i u} \; + \; \prod_{i=1}^{r} m_{a_i u} \prod_{i=1}^{r} m_{b_i v}$$

and denote the set of spider relations as $\mathcal{S}_r$. Every term lies in the class $(1^{2r}, r, r)$ and is a pair of disjoint copies of $\mathcal{Z}_r$ with the centers exchanged, so the relation is a binomial of degree $2r$. We call the elements of $\mathcal{S}_3$ *spiders*.

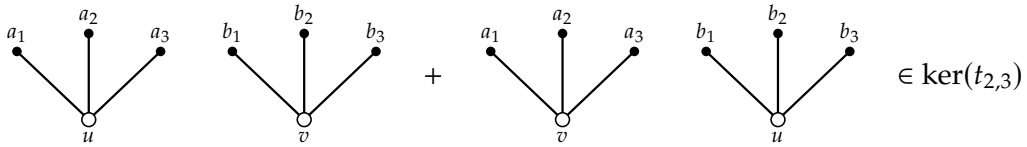



FIGURE 20. The spider relation $\mathcal{S}_3$ in the class $(1^6, 3, 3)$

**Proposition 6.2.** *$\mathcal{S}_r \subseteq \ker(t_{2,r})$ for every $r \geq 3$.*

*Proof.* Each term is an $r$-grove. The centers have valency $r$ and are not roots, the $2r$ leaves are, so $\mathcal{Y}$ is the set of leaves. The trees of the first term have root families $\{a_1, \ldots, a_r\}$ and $\{b_1, \ldots, b_r\}$, and so do the trees of the second, hence both terms have the same image under $\sigma$ and their symmetric difference is $\emptyset$. Proposition 5.4 gives the relation. □

The two families are the two mechanisms of Subsection 5.3. The spider is an identification: two distinct monomials with the same root family, hence the same image. The rainbow is a cancellation: $(r+2)/2$ distinct root families whose cuts sum to zero.

6.2. **Perfect matchings.** At $r = 2$ the rainbow and degenerate rainbow relations act as moves on linear forests. In this subsection we show that the moves act transitively on linear forests with a common endpoint pairing, so the classes of the resulting equivalence are indexed by the perfect matchings $P(L)$. This is what Section 7 uses to index the basis at $r = 2$.

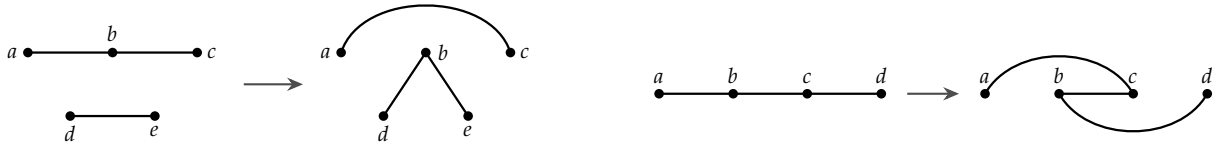
**Definition 6.3.** Let $L$ be a linear forest

- a *rainbow move* corresponds to switching the edges $\{a,b\}, \{b,c\}, \{d,e\}$ with $\{b,d\}, \{b,e\}, \{a,c\}$;
- a *degenerate rainbow move* corresponds to switching the edges $\{a,b\}, \{b,c\}, \{c,d\}$ with $\{a,c\}, \{b,c\}, \{b,d\}$.

If $L'$ is derived from $L$ through a rainbow move or degenerate rainbow move we denote $L \to L'$, and we denote $L \sim L'$ if $L'$ can be derived from $L$ through a series of rainbow and degenerate rainbow moves. We define the *rainbow class* of $L$ to be

$$R(L) = \{L' \in \mathcal{O}_2(L) : L \sim L'\},$$

and we define $\mathrm{Rain}(L)$ to be the set of rainbow classes in $\mathcal{O}_2(L)$.




(A) Regular rainbow move (B) Degenerate rainbow move

FIGURE 21. Rainbow moves

Rainbow moves take linear forests to linear forests, and each move is undone by another, so $\sim$ is an equivalence relation on linear forests.

Rainbow moves preserve the endpoint pairing: by induction on the length of a sequence of moves it suffices to treat a single move $L' \to L''$, and either move takes endpoints to endpoints and leaves them on the same path. So $L' \sim L''$ implies $\rho(L') = \rho(L'')$. The converse is significantly more difficult and is the main result of this subsection. We prepare by making a few definitions. Let $L = P_{n_1,\ldots,n_k}$ be a linear forest and $\rho(L)$ be its perfect matching. We label the paths of $L$ consecutively, so that its $j$-th path has vertex set $\{N_{j-1}+1, \ldots, N_j\}$, where $N_j = n_1 + \cdots + n_j$ and $N_0 = 0$. For $\rho = \{\{a_1,b_1\}, \ldots, \{a_k,b_k\}\} \in P(L)$, indexed as in Subsection 4.2, we define the *standard linear forest* $L_\rho^{\mathrm{id}}$ to be the linear forest whose $j$-th path is $a_j, N_{j-1}+2, N_{j-1}+3, \ldots, N_j - 1, b_j$, so that its interior edges are the $\{c, c+1\}$ with $N_{j-1} + 2 \le c < c+1 \le N_j - 1$. We define the *standard path* $P_\rho^{\mathrm{id}}$ to be the singular case of the standard linear forest.

Then we call a linear forest $L$ *aligned* if $L$ is the disjoint union of $P_i$ such that $P_i \in \mathcal{O}_2(P_\rho^{\mathrm{id}})$ for all $i \in [k]$. Finally, we denote the set of aligned linear forests as $A(L_\rho^{\mathrm{id}})$.

We note the bijection between the elements of $\mathcal{O}_2(P_n)$ and $\mathfrak{S}_{[2,n-1]}$ since the class of paths have fixed endpoints. We use the standard symmetric group $\mathfrak{S}_{n-2} \cong \mathfrak{S}_{[2,n-1]}$, for $w \in \mathfrak{S}_{n-2}$, an *inversion* of $w$ is some $a, b \in [n-2]$ with $a < b$ such that $w(a) > w(b)$. Then let $\mathrm{inv}(w)$ be the number of inversions in the permutation $w$.

The final note we give is that the degenerate rainbow relation, replacing $a-b-c-d$ with $a-c-b-d$, corresponds to the element $(bc) \in \mathfrak{S}_{[2,n-1]}$. Hence these two formulations are equivalent within the conversion between paths and permutations.

**Proposition 6.4.** *Let $L', L'' \in \mathcal{O}_2(L)$. Then $L' \sim L''$ if and only if $\rho(L') = \rho(L'')$.*

**Claim 1**: Let $\rho(L) = \rho(L')$ then $L \sim \hat{L} \sim L'$ with $\hat{L} \in A(L_\rho^{\mathrm{id}})$.




(A) Aligned (B) Not aligned

FIGURE 22. Two linear forests in $O_2(L_\rho^{\mathrm{id}})$ for $\rho = \{\{1,10\},\{6,15\},\{7,11\}\}$, each vertex shaded by the path of $L_\rho^{\mathrm{id}}$ containing it

*Proof.* It suffices to show that $L \sim \hat{L}$. Split $L_\rho^{\mathrm{id}}$ into its paths $P_{n_1}, \ldots, P_{n_k}$ and similarly split $L$ into its paths $Q_{n_1}, \ldots, Q_{n_k}$ with the condition that we order the paths such that $\mathrm{End}(P_{n_i}) = \mathrm{End}(Q_{n_i})$. We define the value

$$\beta(L) = \sum_{j=1}^{k} |\operatorname{Int}(P_{n_j}) \cap \operatorname{Int}(Q_{n_j})|$$

which counts the number of vertices along the paths of $Q_{n_j}$ which are in the correct corresponding path $P_{n_j}$. Let $m = |\operatorname{Int}(L_\rho^{\mathrm{id}})|$, then $\beta(L) \leq m$ with equality exactly when $L$ is aligned. Suppose $\beta(L) < m$. Then some interior vertex $v$ lies on $Q_{n_i}$ while being an interior vertex of $P_{n_j}$ for some $j \neq i$. Let $a$ and $c$ be the two neighbors of $v$ in $Q_{n_i}$. Since $n_j \geq 2$ the path $Q_{n_j}$ has an edge $\{d,e\}$, and the two paths are disjoint, so $v \notin \{d,e\}$ and $\{d,e\} \neq \{a,c\}$. The rainbow move on $a,v,c,d,e$ is therefore available: it deletes $\{a,v\},\{v,c\},\{d,e\}$ and inserts $\{v,d\},\{v,e\},\{a,c\}$, transferring $v$ from the $i$-th path to the $j$-th and leaving every other interior vertex where it was.

The move preserves the endpoint pairing, so the paths of the result are again indexed by $[k]$ in the same way. The $i$-th term of $\beta$ loses nothing, since $v$ was not an interior vertex of $P_{n_i}$, and the $j$-th term gains one, so $\beta$ rises by exactly one. As $m - \beta(L)$ is a non-negative integer, iterating carries $L$ to an aligned linear forest $\hat{L} \in A(L_\rho^{\mathrm{id}})$ in $m - \beta(L)$ moves.

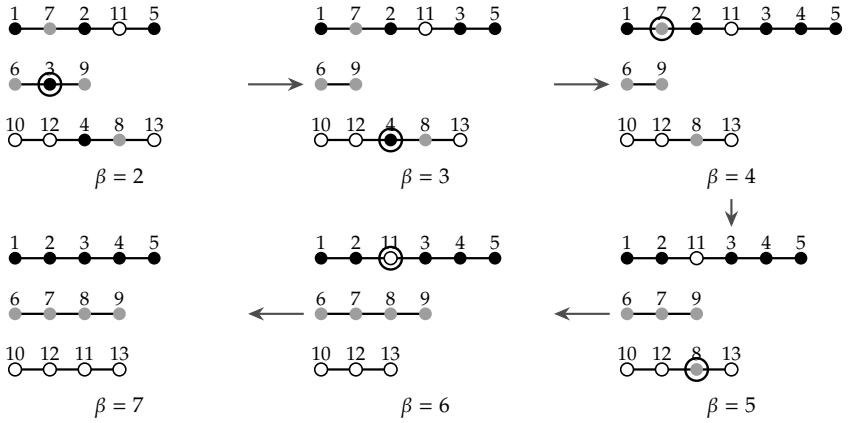



FIGURE 23. The alignment of the linear forest $L = P_{5,4,4}$; the ringed vertex moves next

□

**Claim 2**: For $P \in O_2(P^{\mathrm{id}})$, we have that $P \sim P^{\mathrm{id}}$.

*Proof.* We turn to the symmetric group translation. Define $P = P^w$ for $w \in \mathfrak{S}_{[2,n-1]}$. Then $P = P^{\mathrm{id}}$ exactly when $w$ is the identity, which happens exactly when $\mathrm{inv}(w) = 0$: a permutation with no inversion is increasing, hence the identity, and the converse is immediate.

Suppose $\mathrm{inv}(w) > 0$. Then $w$ has a descent $i$, since a permutation with no descent is increasing. The four vertices $w(i-1), w(i), w(i+1), w(i+2)$ are consecutive on $P$, and the degenerate rainbow move on them interchanges the entries in positions $i$ and $i+1$ and fixes the rest, so it replaces $w$ by $\tau_i w$ where $\tau_i$ is the adjacent transposition. The pair $(i, i+1)$ is an inversion of $w$ and is not one of $\tau_i w$. No other pair changes status: a pair meeting neither position is untouched, and a pair meeting exactly one of them is interchanged with the pair meeting the other, so the number of inversions among those two is the same before and after. Hence $\mathrm{inv}(\tau_i w) = \mathrm{inv}(w) - 1$.

Inducting on $\mathrm{inv}(w)$ carries $P$ to $P^{\mathrm{id}}$ in $\mathrm{inv}(w)$ degenerate moves, and in no fewer degenerate moves, since each degenerate move changes inv by exactly one; rainbow moves which are not degenerate could shorten the sequence.

For an aligned linear forest the same argument applies path by path. A degenerate move replaces a segment $a, b, c, d$ of a path by $a, c, b, d$, so its four vertices lie consecutively on a single path; the move therefore acts on that path alone and leaves the others untouched. Writing an aligned $\hat{L}$ as $\bigsqcup_j P^{w_j}$ and performing these sequences for $j = 1, \ldots, k$ in turn carries $\hat{L}$ to $L_\rho^{\mathrm{id}}$ in $\sum_j \mathrm{inv}(w_j)$ degenerate moves, and in no fewer degenerate moves.

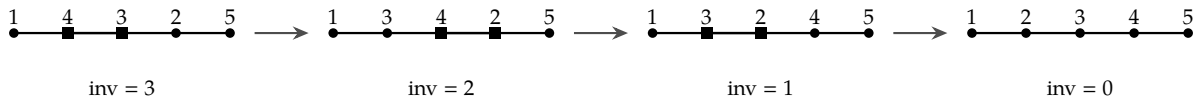



FIGURE 24. Flattening the path $P^w \to P_5^{\mathrm{id}}$; the squared pair is exchanged next

□

**Claim 3**: If $\rho(L') = \rho(L'')$ then $L' \sim L''$, which with the argument above proves Proposition 6.4.

*Proof.* Claims 1 and 2 were carried out in the standard labeling, so it remains to move between matchings. For $\tau \in \mathfrak{S}_{\mathrm{End}(L)}$ define $\hat{\tau}$ to be $\tau$ on $\mathrm{End}(L)$ and the identity on $\mathrm{Int}(L)$, which is a permutation of $V(L)$. Then $\hat{\tau}L' \in \mathcal{O}_2(L)$, since relabeling preserves valencies, hence $\rho(\hat{\tau}L') = \tau(\rho(L'))$. Then $L' \to L''$ iff $\hat{\tau}L' \to \hat{\tau}L''$, since $\hat{\tau}$ is injective. Therefore $L' \sim L''$ iff $\hat{\tau}L' \sim \hat{\tau}L''$.

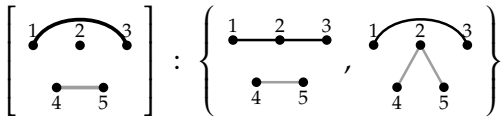
Finally $\mathfrak{S}_{\mathrm{End}(L)}$ acts transitively on $P(L)$. So given $\rho(L') = \rho(L'')$, we can choose $\tau$ carrying that matching to $\rho(L^{\mathrm{id}})$; then Claims 1 and 2 give $\hat{\tau}L' \sim \hat{\tau}L''$, and it follows that $L' \sim L''$. □

The proof also bounds the distance: a linear forest with the standard matching reaches $L^{\mathrm{id}}$ in at most $(m - \beta(\hat{L})) + \sum_j \binom{n_j-2}{2}$ moves, since $\mathrm{inv}(w_j) \le \binom{n_j-2}{2}$. Here inv is the Coxeter length in the adjacent transpositions, so the flattening distance is the Kendall tau distance to the identity, attained at the longest element. This is not the Cayley distance, which counts arbitrary transpositions rather than adjacent ones.

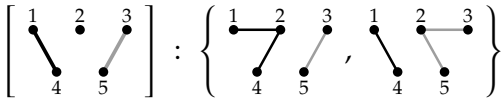
Consequently $\rho$ descends to a map $\mathrm{Rain}(L) \to P(L)$, which is injective by Proposition 6.4 and surjective since $L_\rho^{\mathrm{id}}$ realizes every $\rho$. Hence the rainbow classes of $\mathcal{O}_2(L)$ are in bijection with $P(L)$, and there are $(2k-1)!!$ of them.

In the language of Subsection 5.2, a rainbow class is the set of linear forests with a given root family, and Proposition 6.4 says that the rainbow moves connect it.

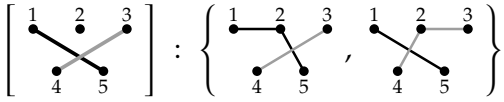
6.3. **Minimal generators.** Theorem A asserts both completeness and minimality of the listed relations. We phrase minimality as an obstruction to graph moves. A relation in a fixed class is a weighted graph list with zero surviving record at every allowed orientation. A relation is *inherited* if it is a sum of lists obtained by adding one edge to every graph of a relation in a class with one fewer edge. Adding several fixed edges is already covered: remove the last edge and regard the other additions as part of the smaller relation. Edge attachment preserves relations for a direct counting reason: split an orientation into its old edges and its new edge, then group by the two resulting



(A) $\rho = \{1,3\},\{4,5\}$



(B) $\rho = \{1,4\},\{3,5\}$



(C) $\rho = \{1,5\},\{3,4\}$

FIGURE 25. Corresponding perfect matchings of the rainbow classes of $O_2(L) = O_2(P_{3,2})$

records. A forbidden old record cannot become allowed by adding incidences; every allowed old record has total weight zero. Thus every resulting allowed record also has total weight zero.

Write $I$ for the collection of all relations, $I_{\mathbf{a}}$ for those in class $\mathbf{a}$, and $\mathrm{Inh}(\mathbf{a})$ for the inherited ones. The graph-algebra dictionary identifies the latter collection with $(\mathfrak{M}I)_{\mathbf{a}}$, but we use its edge-attachment definition in every argument below. An individual relation is *new* if it is not inherited. This is the meaning of an individual minimal generator. A *minimal generating family* must in addition generate every relation and contain no redundant member. Showing that each member is new does not by itself prove the latter assertion; the class normal forms below supply that additional step.

**Lemma 6.5.** *Let $B \subseteq I$ be a set of relations, each in a single class, and let $B_{\mathbf{a}} = B \cap I_{\mathbf{a}}$.*

(1) *If for every $\mathbf{a}$ every relation of the class $\mathbf{a}$ is a combination of elements of $B_{\mathbf{a}}$ plus an inherited relation, then $B$ generates $I$.*
(2) *If moreover for every $\mathbf{a}$ no nonzero combination of elements of $B_{\mathbf{a}}$ is inherited, then no proper subset of $B$ generates $I$, and every generating set of $I$ consisting of relations of single classes has at least $|B_{\mathbf{a}}|$ elements in the class $\mathbf{a}$.*

*Proof.* For (1), induct on the number of edges. The empty graph has a nonzero empty record, so it has no relation. By hypothesis a relation in the current class is a weighted sum of the chosen relations in that class and lists obtained by attaching an edge to smaller relations. The latter are generated by the induction hypothesis. Attaching that edge to each step of their reductions gives the required reduction in the current class.

For (2), consider any generating family $B'$. In a reduction of a relation of class $\mathbf{a}$, every use of a generator with at least one added edge is inherited. Thus, after recording these inherited terms, only generators of $B'_{\mathbf{a}}$ remain. Suppose that this class has fewer members of $B'$ than of $B$. Use the following explicit exchange procedure. Express a first member of $B_{\mathbf{a}}$ using $B'_{\mathbf{a}}$ and inherited lists. Some coefficient of a member of $B'$ must be nonzero, since otherwise that first member is inherited. Solve for that member of $B'$ and replace it by the first member of $B$. The new list still generates the same reductions.

Repeat with a second, third, and subsequent member of $B_{\mathbf{a}}$. At every stage some remaining member of $B'$ must have a nonzero coefficient: otherwise the new member would be a combination of the earlier members of $B$ and inherited lists, contradicting the hypothesis. Each step consumes a different member of $B'$. The procedure therefore cannot accommodate more members of $B$ than

there are members of $B'$. This proves the size bound without a coefficient matrix. The same argument shows directly that omitting a member of $B$ cannot generate it. □

In particular the number of generators required in each class is intrinsic. For the Plücker class, the obstruction is already visible with one edge: orienting an edge toward either endpoint distinguishes it from all other edges. There are no one-edge relations. Consequently an inherited relation has at least three edges, whereas a Plücker relation has two.

Fix $n \geq 2$, a vertex set $V$ with $|V| = n$, and let $\mathbf{a}$ be the valency vector with value 2 on $V$, the valency vector of a cycle through $V$. The connected members of $O(\mathbf{a})$ are exactly the cycles through all of $V$. For a sum $f$ of members of $O(\mathbf{a})$ with coefficients in $\mathbb{k}$, we define its *connected count* $\kappa(f)$ to be the sum of the coefficients of its connected terms.

Let $e = \{i, j\}$ be a pair of vertices of $V$ and let $\mathbf{b} = \mathbf{a} - \mathcal{W}(e)$, so that $\mathbf{b}$ has valency 1 at $i$ and at $j$ and valency 2 elsewhere. Every vertex of a member of $O(\mathbf{b})$ has valency 1 or 2, so such a member is a disjoint union of paths and cycles, and only $i$ and $j$ can be the ends of a path. Hence it is a path from $i$ to $j$ together with disjoint cycles. We call it *spanning* when it contains no cycle, equivalently when its path passes through every vertex.

**Lemma 6.6.** *Let $n \geq 2$ and $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$. Then $m_{C_n}$ is a minimal generator of* $\ker(t_{2,r})$.

*Proof.* Fix an edge $e = ij$ that might have been attached to a smaller relation. After deleting it from the cycle valency class, a graph is an $i$–$j$ path together with disjoint cycles. Among the allowed orientation records, look only at the record of a single tree rooted at $i$. It occurs exactly once in every spanning $i$–$j$ path. It occurs in no graph with a cycle, because cyclic orientations cancel in pairs. Since a smaller relation has zero weight at this record, its spanning paths have total weight zero.

Restoring $ij$ makes a graph connected exactly when that path was spanning. Thus every relation obtained by attaching $ij$ has connected count zero. This holds for every possible attached edge, so it holds for every inherited relation. A cycle through all vertices has connected count one (Figure 26). It is a relation by cycle reversal, and cannot be inherited. □

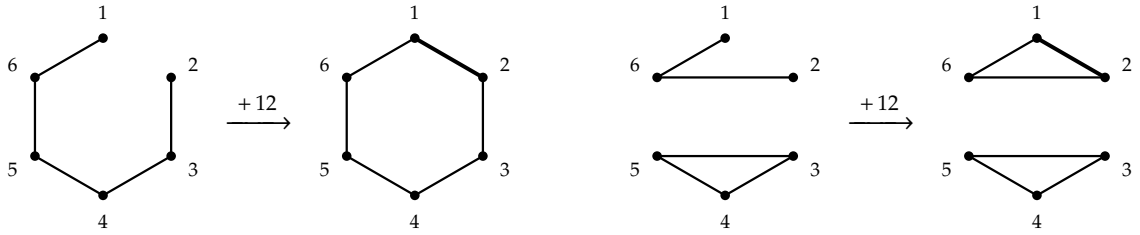



FIGURE 26. Restoring the edge 12 to a spanning path and to a path with a cycle

**Lemma 6.7.** *Let $r \in \mathbb{N}_{\geq 2}$. Then $m_{\mathcal{Z}_{r+1}}$ is a minimal generator of* $\ker(t_{2,r})$.

*Proof.* The valency class $(r+1, 1^{r+1})$ contains exactly one graph: every leaf must be joined to the center. If a proposed last attached edge joins two leaves, the remaining center would still need $r+1$ incidences but only $r-1$ leaf incidences are available. No smaller graph exists. If the attached edge joins the center to a leaf, the smaller class consists of the $r$-star alone. Rooting that star at any of its leaves gives a record of weight one, so the smaller class has no nonzero relation. There is therefore no nonzero inherited relation in the $(r+1)$-star class. The star itself has no allowed orientation and is new. □

For the exceptional families we use weights on the leaves. Let $r$ be even, let $b$ be a vertex and $A$ a set of $r+2$ further vertices, and let $\mathbf{a}$ be the valency vector with value $r$ at $b$ and 1 on $A$. Every

member of $O(\mathbf{a})$ is a grove $F_{st}$, consisting of a leaf edge $\{s,t\} \subseteq A$ and the other leaves joined to $b$. For a sum $f = \sum c_{st}F_{st}$ we define the *weight* of a leaf $p$ to be

$$w_p(f) = \sum_{t \in A\setminus\{p\}} c_{pt},$$

the sum of the coefficients of the members whose leaf edge contains $p$.

**Lemma 6.8.** *Let $r$ be even. A sum $f = \sum c_{st}F_{st}$ is a relation of $\ker(t_{2,r})$ if and only if all of its leaf weights are equal, and every inherited relation of the class* **a** *has all of its leaf weights equal to zero.*

*Proof.* For a root pair $xy$, count the leaf edges containing exactly one of $x, y$. The leaf-incidence counts $w_x, w_y$ count these once; the edge $xy$ itself is counted twice and cancels. The root-pair weight is therefore $w_x + w_y$. All records vanish exactly when all leaf weights are equal.

Now delete a possible last attached edge $e$. If it is a leaf edge, the smaller class has just the surviving $r$-star, so there is no smaller relation. If $e = sb$, then $s$ becomes isolated and can be omitted. The center $b$ has valency $r-1$ and is now an eligible root. The smaller class again consists of a leaf pair and a star, but $b$ belongs to no leaf pair. Its incidence count is therefore zero. The same root-pair count forces every other leaf-incidence count of a smaller relation to be zero as well. Reattaching $sb$ does not change the leaf pair, and adds zero to the incidence count at $s$. Hence every inherited relation has zero weight at every leaf. □

Now let $r \geq 3$ be odd, let $u, v$ be vertices and $A$ a set of $2r$ further vertices, and let **a** be the valency vector with value $r$ at $u$ and at $v$ and 1 on $A$. A forest of $O(\mathbf{a})$ is either two stars or a tree containing the edge $\{u,v\}$ together with a separate leaf edge. For a sum $f$ of members of $O(\mathbf{a})$, the *weight* $w_\alpha(f)$ of a leaf $\alpha$ is the sum of the coefficients of the forests of $f$ in which $\alpha$ does not lie in the component of $v$; members with a cycle are ignored. The leaves outside the component of $v$ are the roots of the other tree, so by Proposition 5.2 and the computation in the proof of Lemma 6.8, a sum $f$ is a relation exactly when all of its leaf weights are equal.

**Lemma 6.9.** *Let $r \geq 3$ be odd. Every inherited relation of the class* **a** *has all of its leaf weights equal to zero.*

*Proof.* Fix a proposed last edge $e$, a leaf $\alpha$, and a graph $H$ in the smaller valency class. Keep isolated vertices as components; their unique root is compulsory. Every forest in this smaller class has three components. Let $c_R(H)$ be 1 if the three vertices of $R$ select one eligible root in each component, and 0 otherwise; set it to 0 for graphs with a cycle. This is a direct test on components. The following four local rules express whether $\alpha$ lies outside the component of $v$ after $e$ is restored:

(1) If $e = \{p,q\}$ with $p, q \in A$, it is $c_{\{p,q,\gamma\}}(H)$ for any $\gamma \in A \setminus \{p,q\}$ when $\alpha \in \{p,q\}$, and zero otherwise.
(2) If $e = \{w,v\}$ with $w \in A$, it is $c_{\{w,v,\alpha\}}(H)$ when $\alpha \neq w$, and zero when $\alpha = w$.
(3) If $e = \{w,u\}$ with $w \in A$, it is $Q(H) + c_{\{w,u,\alpha\}}(H)$ when $\alpha \neq w$, and $Q(H)$ when $\alpha = w$, where

$$Q(H) = \sum_{\substack{|R|=3 \\ w,u\in R}} c_R(H).$$

(4) If $e = \{u,v\}$, it is $c_{\{u,v,\alpha\}}(H)$.

For $\{p,q\}$, the vertices $p, q$ are isolated and the rest form one tree. For $\{w,v\}$, the vertex $w$ is isolated and $v$ is eligible, so $\{w,v,\alpha\}$ selects the components exactly when $\alpha \neq w$ lies outside the component of $v$. For $\{w,u\}$, separate centers give a star at $v$ with $r$ leaves, so $Q(H) = 1$ and the additional coefficient records membership there; joined centers give a third component which is a leaf pair, so $Q(H) = 0$ and that coefficient records membership in the pair. For $\{u,v\}$, separate centers leave a third component which is a leaf pair; joined centers give coefficient zero, and adding $\{u,v\}$ creates a cycle. In every case both sides vanish for a member $H$ with a cycle.

Apply the appropriate rule to every graph of a smaller relation, retaining its weight. Each fixed root triple has total weight zero by the definition of a relation. The rules express the desired leaf count as a sum of exactly these zero triple counts. It follows that every edge attachment has zero count at $\alpha$, and hence so does every inherited relation. The use of odd $r$ is concrete: in case (3), selecting the last root among the $r$ leaves of the separate $v$-star has odd count. □

**Lemma 6.10.** *Let $r$ be even. Then every element of $\mathcal{R}_r$ is a minimal generator of* $\ker(t_{2,r})$. *Let $r$ be odd. Then every element of $\mathcal{S}_r$ is a minimal generator of* $\ker(t_{2,r})$.

*Proof.* A rainbow $\sum_{\{s,t\}\in M} F_{st}$ uses each leaf in exactly one leaf edge, so it has weight 1 at every leaf, and by Lemma 6.8 it is not inherited. The two terms of a spider exchange complementary leaf sets, so a spider has weight 1 at every leaf, and by Lemma 6.9 it is not inherited. Both are relations by Propositions 6.1 and 6.2. □

**Proposition 6.11.** *Let $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$. Every element of the generating set of Theorem A is a minimal generator of* $\ker(t_{2,r})$. *In particular* $\ker(t_{2,r})$ *has a minimal generator in every degree $s \geq 2$.*

*Proof.* The Plücker relations are minimal by the degree count above. The cycles are Lemma 6.6, the star is Lemma 6.7, and the rainbows and spiders are Lemma 6.10. The cycle $C_s$ has degree $s$. □

We next give explicit reduction procedures inside the generator classes. They show that, after inherited moves, one chosen generator and its scalar weight are the only remaining data. This is the step that turns individual counting obstructions into a minimal generating family.

**Lemma 6.12.** *Let $V$ be a finite vertex set with $|V| \geq 2$, and let $\mathbf{a}$ be the valency vector with value 2 on $V$. Then every relation of the class $\mathbf{a}$ is a multiple of any cycle through $V$ plus an inherited relation, and no nonzero multiple of a cycle through $V$ is inherited.*

*Proof.* Every member of $O(\mathbf{a})$ is a disjoint union of cycles covering $V$, some of which may be doubled edges, so it is a relation by Proposition 4.1. A disconnected member is obtained from one of its cycle components, which is already a relation, by attaching the nonempty collection of edges in the other components. It is therefore inherited. The remaining members are the cycles through all of $V$. If $|V| \leq 3$ there is exactly one.

Suppose $|V| \geq 4$, let $H$ be a cycle through $V$, and let $w_1, w_2, w_3, w_4$ be consecutive on $H$. In characteristic two

$$m_{w_1w_2}m_{w_3w_4} + m_{w_1w_3}m_{w_2w_4} + m_{w_1w_4}m_{w_2w_3}$$

is a Plücker relation, and adjoining the other $|V|-2$ edges of $H$ to all three matchings gives an inherited relation. Its three terms are $m_H$, the monomial of the cycle $H'$ obtained from $H$ by exchanging $w_2$ and $w_3$, and a member with two components, namely the doubled edge $\{w_2, w_3\}$ and a cycle through the other vertices (Figure 27). Hence $m_H$ and $m_{H'}$ differ by an inherited relation. Exchanging adjacent vertices connects any two cycles through $V$, since adjacent transpositions generate the symmetric group, so any two cycles through $V$ differ by an inherited relation, which proves the first statement.

By the proof of Lemma 6.6 every inherited relation has connected count zero, while a nonzero multiple $c\, m_H$ of a cycle has connected count $c \neq 0$. □

**Lemma 6.13.** *Let $r$ be even, let $b$ be a vertex and $A$ a set of $r+2$ further vertices, and let $\mathbf{a}$ be the valency vector with value $r$ at $b$ and 1 on $A$. Then every relation of the class $\mathbf{a}$ is a multiple of any rainbow with center $b$ and leaves $A$ plus an inherited relation, and no nonzero multiple of a rainbow is inherited.*

*Proof.* Let $f$ be a relation of the class and let $\xi$ be a rainbow. By Lemma 6.8 all weights of $f$ equal some $c \in \Bbbk$, and all weights of $\xi$ equal 1, so $g = f - c\xi$ has all weights zero. We show that $g$ is inherited.

For distinct $s, t, x \in A$,

$$F_{st} + F_{sx} + F_{tx} = \left(m_{st}m_{xb} + m_{sx}m_{tb} + m_{tx}m_{sb}\right) \prod_{y\in A\setminus\{s,t,x\}} m_{yb}$$

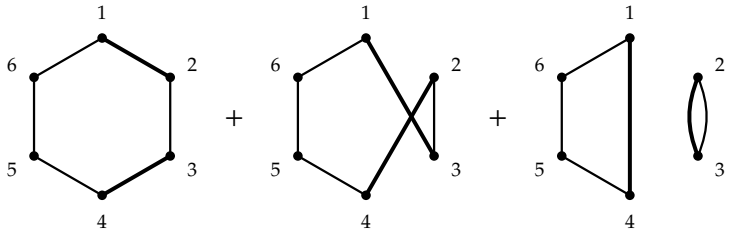



FIGURE 27. The Plücker move on four consecutive vertices $1, 2, 3, 4$ of a cycle

is obtained by applying a Plücker move on $s, t, x, b$ and attaching the other $r-1$ edges. It is inherited (Figure 28), and every leaf occurs in exactly zero or two of its leaf pairs. Fix $s_0 \in A$ and write $g = \sum c_{pq} F_{pq}$. Subtracting $c_{pq}(F_{pq} + F_{s_0 p} + F_{s_0 q})$ for every pair $\{p, q\}$ not containing $s_0$ leaves a sum $g'$ of members $F_{s_0 p}$ only, with all weights still zero. The weight of $g'$ at $p \neq s_0$ is its coefficient of $F_{s_0 p}$, so $g' = 0$ and $g$ is a sum of inherited relations.

For the second statement, a nonzero multiple $c\xi$ has weights $c \neq 0$, while inherited relations have weights zero by Lemma 6.8. □

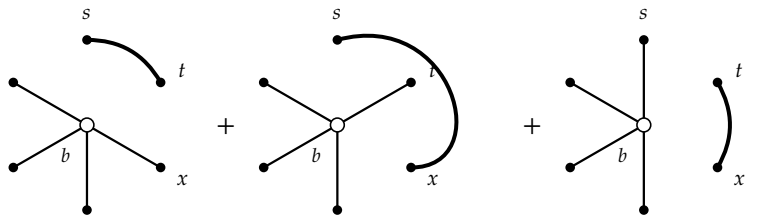



FIGURE 28. The inherited relation $F_{st} + F_{sx} + F_{tx}$ at $r = 4$

**Lemma 6.14.** *Let $r \geq 3$ be odd, let $u, v$ be vertices and $A$ a set of $2r$ further vertices, and let* **a** *be the valency vector with value $r$ at $u$ and at $v$ and $1$ on $A$. Then every relation of the class* **a** *is a multiple of any spider with centers $u, v$ and leaves $A$ plus an inherited relation, and no nonzero multiple of a spider is inherited.*

*Proof.* A cyclic member must contain the doubled edge $uv$, since every other vertex is a leaf. Delete all its other edges; the doubled edge is already a relation, so restoring those edges makes the member inherited. The forests are of two kinds. For $S \subseteq A$ with $|S| = r$ let $Z_S$ join $u$ to $S$ and $v$ to $A \setminus S$. For distinct $p, q \in A$ and $S \subseteq A \setminus \{p, q\}$ with $|S| = r - 1$ let $Z_{pq,S}$ consist of the edge $\{p, q\}$ and the tree with the edge $\{u, v\}$ joining $u$ to $S$ and $v$ to the remaining vertices of $A$. The forest $Z_S$ has weight 1 exactly at the leaves of $S$, and $Z_{pq,S}$ has weight 1 exactly at $p$ and $q$ (Figure 29).

In characteristic two a Plücker relation is the sum of the three perfect matchings of its four vertices. Keep the remaining edges fixed and perform that three-matching move on four vertices. This gives the following three explicit inherited switches.

(1) For $s \in S$ and $t \in A \setminus S$, the relation on $u, v, s, t$ times the edges from $u$ to $S \setminus \{s\}$ and from $v$ to $A \setminus (S \cup \{t\})$ is $Z_S + Z_{(S\setminus\{s\})\cup\{t\}} + Z_{st,S\setminus\{s\}}$.
(2) For $x \in S$, the relation on $p, q, x, u$ times the edges of $Z_{pq,S}$ other than $\{p, q\}$ and $\{u, x\}$ is $Z_{pq,S} + Z_{px,(S\setminus\{x\})\cup\{q\}} + Z_{qx,(S\setminus\{x\})\cup\{p\}}$.
(3) For $x \in S$ and $y \in A \setminus (S \cup \{p, q\})$, the relation on $u, v, x, y$ times the edges of $Z_{pq,S}$ other than $\{u, x\}$ and $\{v, y\}$ is $Z_{pq,S} + Z_{pq,(S\setminus\{x\})\cup\{y\}}$ plus a member containing the doubled edge $\{u, v\}$.

By (3), and since exchanges connect any two splits of $A \setminus \{p, q\}$, the forests $Z_{pq,S}$ with the same leaf edge $\{p, q\}$ differ by inherited relations; let $z_{pq}$ denote any one of them. Fix $p_0 \in A$ and an $r$-subset $S_0 \subseteq A$. For $p, q \neq p_0$ we may choose $S \ni p_0$, and (2) with $x = p_0$ shows that $z_{pq}$ differs

from $z_{pp_0} + z_{qp_0}$ by an inherited relation. By (1), and since exchanges connect any two $r$-subsets of $A$, every $Z_S$ differs from $Z_{S_0}$ plus a sum of the $z_{pq}$ by an inherited relation. Hence every sum of members of the class is a combination of $Z_{S_0}$ and of the forests $z_{p_0x}$ for $x \in A \setminus \{p_0\}$, plus an inherited relation.

Now let $f$ be a relation of the class and let $\zeta$ be a spider. All weights of $f$ equal some $c \in \Bbbk$, and all weights of $\zeta$ equal 1, so $g = f - c\zeta$ has all weights zero. Write $g = c_0 Z_{S_0} + \sum_x c_x z_{p_0x} + h$ with $h$ inherited. By Lemma 6.9 the weights of $h$ are zero, so the weights of $c_0 Z_{S_0} + \sum_x c_x z_{p_0x}$ are zero. Adding up its weights over all leaves counts each $z_{p_0x}$ twice and $Z_{S_0}$ exactly $|S_0| = r$ times, so the total is $c_0 r = c_0$, since $r$ is odd and the characteristic is two. Hence $c_0 = 0$, and then the weight at each $x \neq p_0$ is $c_x = 0$. Therefore $g = h$ is inherited, which proves the first statement.

For the second statement, a nonzero multiple $c\zeta$ has weights $c \neq 0$, while inherited relations have weights zero by Lemma 6.9. □

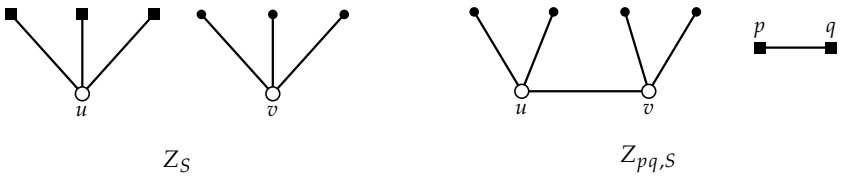



FIGURE 29. The two kinds of forest in the spider class at $r = 3$, with the leaves of weight one as squares

## 7. BASES

At every level we give a reduction procedure and a distinguished record that recovers the weight of each normal form. At $\Phi$ the record was the opener. Here it is the set of largest eligible roots, and at odd $r$ we refine it using marked vertices and partial matchings. Denote the Iverson bracket notation for a set $A$,

$$[a \in A] = \begin{cases} 1 \text{ if } a \in A \\ 0 \text{ if } a \notin A \end{cases}$$

and similarly with the membership switched for $[a \notin A]$.

Let $F$ be an $r$-forest with root family $\{X_1, \ldots, X_k\}$. We define the *root maxima* of $F$ to be

$$\mathcal{M}(F) = \{\max X_1, \ldots, \max X_k\},$$

the largest root of each tree. This is a root selection, so $\mu_{\mathcal{M}(F)}$ is a term of $t_{2,r}(m_F)$. We order root selections by dominance as in Subsection 3.2.

**Proposition 7.1.** *Let $F$ be an $r$-forest. Then $\mu_{\mathcal{M}(F)}$ occurs in $t_{2,r}(m_F)$ with coefficient 1, and if $F'$ is an $r$-forest in the same class with $\mu_{\mathcal{M}(F)}$ occurring in $t_{2,r}(m_{F'})$, then $\mathcal{M}(F) \preceq \mathcal{M}(F')$.*

*Proof.* The first statement is Proposition 4.6, since $\mathcal{M}(F)$ is a root selection and every term has coefficient 1. For the second, $\mathcal{M}(F)$ is the dominance maximum among the root selections of $F$: if $T$ is a selection with entries $t_1 \leq \cdots \leq t_k$ and $\mathcal{M}(F)$ has entries $m_1 \leq \cdots \leq m_k$, then the $k - j + 1$ blocks contributing $t_j, \ldots, t_k$ have maxima at least $t_j$, so at least $k - j + 1$ entries of $\mathcal{M}(F)$ are at least $t_j$, whence $m_j \geq t_j$. If $\mu_{\mathcal{M}(F)}$ occurs in $t_{2,r}(m_{F'})$ then $\mathcal{M}(F)$ is a root selection of $F'$, and applying this to $F'$ gives $\mathcal{M}(F) \preceq \mathcal{M}(F')$. □

Consequently, in any nonzero weighted list of forests with distinct root maxima, choose a greatest root maximum among those carrying nonzero weight, using a total order extending dominance. Only that forest can contribute to its record. Its weight survives, proving independence by Proposition 3.5.

The $x$-support of $\mu_R$ is the complement of $R$, which is why the extremal selection is the maximum here and the minimum at $\Phi$.

7.1. **Even valency.** We give the basis at $r = 2$, where a class of 2-forests is a linear forest class $O_2(L)$ and the root family is the endpoint pairing $\rho$. By Proposition 4.12 the image of $L' \in O_2(L)$ is

$$\chi_\rho = \prod_{j=1}^{k} (x_{a_j} + x_{b_j}), \qquad \rho(L') = \{\{a_j, b_j\}\},$$

which depends only on $\rho(L')$, in agreement with Proposition 6.4. Every $\rho \in P(L)$ is realized, by $L_\rho^{\mathrm{id}}$, so

$$t_{2,2}\big(\Bbbk[m_{ij}]_{\mathcal{W}(L)}\big) = \mathrm{Span}\{\chi_\rho : \rho \in P(L)\}.$$

A perfect matching is a multigraph, so it is cross-free or not in the sense of Definition 3.1, and straightening acts on it directly.

An equivalent identity appears as [4, Corollary 2.8]; we give the short proof in the language of selections.

**Lemma 7.2.** *Let $a < b < c < d$ be endpoints and let $\rho, \rho_1, \rho_2 \in P(L)$ agree away from these four and contain $\{a, c\}, \{b, d\}$, respectively $\{a, b\}, \{c, d\}$, respectively $\{a, d\}, \{b, c\}$. Then $\chi_\rho = \chi_{\rho_1} + \chi_{\rho_2}$.*

*Proof.* Fix the choices at all pairs outside the four displayed vertices. Of their three perfect matchings, any fixed two-element set is a valid one-per-pair selection for exactly two: the excluded matching is the one that pairs its two elements together. Thus, in the sum of the three selection inventories, every record appears twice and cancels. This is the asserted two-term smoothing rule. □

The character formula of [4, Theorem 5.4] determines the dimension of each class image under $t_{2,2}$ but no basis. The following proposition gives an explicit basis of each class image, indexed by cross-free endpoint pairings, and its proof is independent of the character computation.

**Proposition 7.3.** *Let $O_2(L)$ be a linear forest class with $k$ paths. Then $\{\chi_\rho : \rho \in P(L)$ cross-free$\}$ is a basis of the image of the class under $t_{2,2}$, so the images of one linear forest for each cross-free endpoint pairing form a basis, and*

$$\mathrm{rank}(t_{2,2}) = C_k.$$

*Proof.* Apply the smoothing rule of Lemma 7.2, with the decreasing length-and-crossing order of Proposition 3.2. This reduces every pairing to cross-free pairings without changing its selection inventory.

For a cross-free perfect matching, the larger endpoints determine all pairs: each closer joins the most recent unused opener. Distinct cross-free pairings therefore have distinct sets of larger endpoints. In a nonzero weighted list, the greatest such set is a root record contributed by only its own pairing, with weight one. The list cannot vanish. Finally the stack rule identifies these pairings with balanced up-down paths of length $2k$ that never go below zero. Reflection counts them as $\binom{2k}{k} - \binom{2k}{k-1} = C_k$. □

Openers are up-steps and closers are down-steps, as in Figure 30.

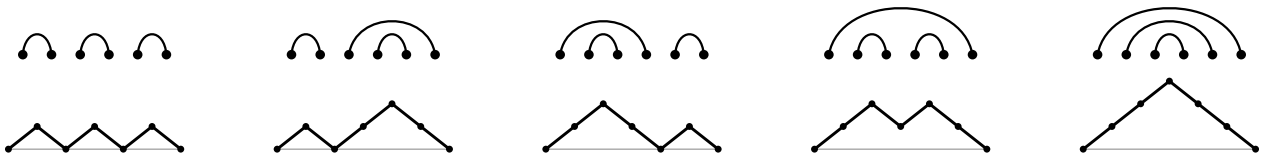


FIGURE 30. The $C_3 = 5$ cross-free perfect matchings of six points and their paths

By Proposition 6.4 the linear forests with a given pairing are exactly one rainbow class, so the basis is one representative per rainbow class with cross-free pairing, which is the $r = 2$ case of the characterization of bases.

For even $r$ we use an auxiliary class in place of the pairing. Let $V$ be the vertex set, $\mathbf{d}$ the valency vector, $U = \{v : d_v = r\}$, $Y = V \setminus U$, and choose a leaf $a \in Y$. We define $\bar{\mathbf{d}} = \mathbf{d}|_Y + |U|(r-2)\mathbf{e}_a$ and order $Y$ by increasing $\bar{d}_v$, breaking ties by the vertex labels. For each forest $H$ with valencies $\bar{\mathbf{d}}$ which is cross-free with respect to this order on $Y$, and not with respect to the vertex labels, let $\widehat{H}$ replace its tree containing $a$ by one chosen tree on the same vertices together with $U$, with the original valencies; leave the other trees unchanged. Existence is checked below. Root maxima are taken in the chosen order.

We also get that the root maxima separate the forests $\widehat{H}$ constructed above. Since the root family of $\widehat{H}$ is the component partition of $H$. Therefore its root maxima determine $H$ by Proposition 7.5, since the auxiliary class is degree-sorted.

**Proposition 7.4.** *Let $r$ be even and let $\mathcal{O}(F)$ be a class of $r$-forests with $n$ vertices, $k$ trees and $u$ vertices of valency $r$. Then the images of the forests $\widehat{H}$ constructed above form a basis of the image of the class under $t_{2,r}$, and*

$$\operatorname{rank}(t_{2,r}) = \binom{n-u-2}{k-1} - \binom{n-u-2}{k-3}.$$

*Proof.* Move $r-2$ branches in pairs from each blocked vertex $b \in U$ to the anchor $a$. Apply Lemma 4.14 at each step. This changes which forests realize an inventory but not the collection of attainable inventories. Every blocked vertex now has valency two. Suppress all these vertices (Figure 31). Each component has leaves outside $U$, so suppression leaves a nontrivial tree. The remaining valencies are $\bar{\mathbf{d}}$, and the surviving root records are unchanged. Conversely, subdividing edges inserts all vertices of $U$ back, without changing any root choice outside $U$.

The degree-sorted cross-free forests $H$ on $Y$ therefore give all normal forms, by Proposition 7.6. To realize such a form in the original class, place all of $U$ in the tree of $H$ containing $a$. The desired positive valencies on this enlarged tree sum to $2(|V(T_a)| + u - 1)$. A Prüfer word in which each vertex occurs $d_v - 1$ times constructs the required tree; choose one such word once and for all. Leave all other components unchanged. The resulting $\widehat{H}$ has exactly the component root blocks of $H$.

Distinct $H$ have distinct root maxima, so the greatest-record argument proves independence for these realizations too. Counting their stack paths by Proposition 7.6 gives $B(n-u, k)$, the stated formula. □

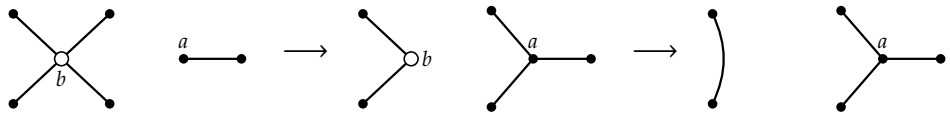



FIGURE 31. Reducing a blocked vertex $b$ at $r = 4$: two branches move to the anchor $a$, and $b$ is suppressed

At $u = n - 2k$ every non-leaf is blocked and the formula is $C_k$, recovering Proposition 7.3.

7.2. **Infinite valency.** At $r = \infty$ every vertex is a root, so the root family is the partition of the vertex set into trees and the root maxima are the largest vertex of each tree. For the normal forms we order the vertices by increasing valency, breaking ties by label. This order affects the chosen representatives, but not the collection of root inventories or its rank.

**Proposition 7.5.** *Let $\mathcal{O}(F)$ be a degree-sorted class of forests. Then the root maxima separate the cross-free forests of $\mathcal{O}(F)$.*

*Proof.* Let the valencies be $d_1 \le \cdots \le d_n$. Suppose $v < w < z$ and $vw, vz$ are edges of a cross-free forest. The component $T$ of $F - v$ containing $w$ excludes $z$, since $F$ has no cycle, and lies between $v, z$: a path leaving this interval would meet $v, z$ or cross $vz$. The only edge from $v$ to $T$ is $vw$. Hence $T$ contains a leaf $t$ of $F$ (choose $w$ if $T = \{w\}$, otherwise a leaf other than $w$), contradicting $1 = d_t \ge d_v \ge 2$.

Thus every vertex has at most one neighbor to its right. Orient edges towards their larger endpoint. A tree on $s$ vertices has $s - 1$ edges and outdegrees at most one, so its maximum is its unique vertex of outdegree zero. Hence $\mathcal{A}(F) = [n] \setminus \mathcal{M}(F)$. The root maxima determine the opener, which determines the cross-free forest by Proposition 3.4. □

**Proposition 7.6.** *Let $O(F)$ be a degree-sorted class of forests with $n$ vertices and $k$ trees. Then the images of the cross-free forests form a basis of the image of the class under $t_{2,\infty}$, and*

$$\operatorname{rank}(t_{2,\infty}) = \binom{n-2}{k-1} - \binom{n-2}{k-3} = B(n,k).$$

*Proof.* Spanning is Proposition 5.5, and independence follows from Propositions 7.1, 7.5 and 3.5. It remains to count the root maxima.

Let $d_1 \le \cdots \le d_n$ and let $\ell$ denote the number of leaves. The degree sum gives $\ell = 2k + \sum_{i>\ell}(d_i - 2) \ge 2k$. Let $R \subseteq [n]$ of size $k$ containing $n$, and let $a_i = [i \notin R]$. At vertex $i$, we can close the $d_i - a_i$ most recent openers, then open at $i$ if $a_i = 1$. The number left open is $h_i = \sum_{j \le i}(2a_j - d_j)$, with $h_n = 0$. This is possible exactly when $h_i \ge a_i$ for all $i$. The graph is cross-free and has no cycle, since a cycle's smallest vertex would open two edges. Thus it is a forest with root maxima $R$, and every cross-free forest arises this way by Proposition 7.5 and the opener construction.

For $i \le \ell$, we have $h_i = i - 2|R \cap [i]|$, so feasibility is the nonnegative up-down path condition: at an up-step the height is at least one, at a down-step at least zero. If $\ell < n$, then for $\ell \le i < n$,

$$h_i = \sum_{j>i}(d_j - 2 + 2[j \in R]) \ge d_n \ge 2, \qquad i - 2|R \cap [i]| \ge \ell - 2(k-1) \ge 2.$$

Both conditions therefore hold beyond the leaves, including the final vertex where $a_n = h_n = 0$. If all vertices are leaves, the construction is the matching path itself. Thus $R \setminus \{n\}$ indexes nonnegative paths of length $n - 1$ with $k - 1$ down-steps. Reflecting through the first visit to $-1$ counts the bad paths by $\binom{n-1}{k-2}$, giving $\operatorname{rank}(t_{2,\infty}) = \binom{n-1}{k-1} - \binom{n-1}{k-2}$, which equals $\binom{n-2}{k-1} - \binom{n-2}{k-3} = B(n,k)$ by Pascal's rule. □

Figure 32 shows three of the $B(8,3) = 14$ forests of the class $(1^6, 2^2)$, with the step at $i < n$ going down exactly when $i$ is a root maximum.

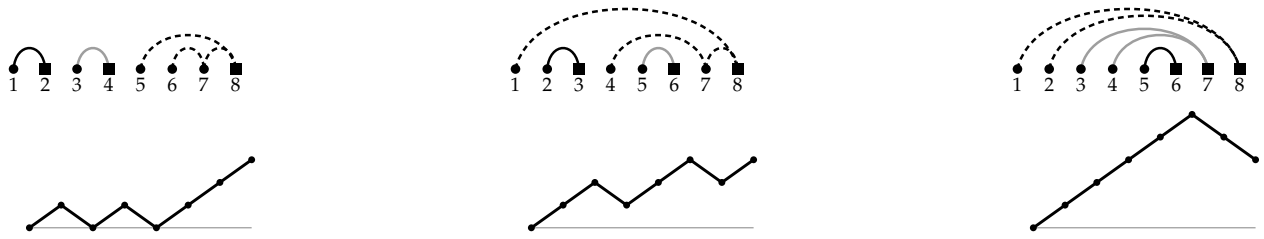



FIGURE 32. Cross-free forests of the class $(1^6, 2^2)$ with $k = 3$, their root maxima, and their paths

7.3. **Odd valency.** At odd $r$ the root maxima are still triangular by Proposition 7.1, so any family of cross-free $r$-forests with distinct root maxima is independent. But this invariant is too coarse: Figure 14 shows five cross-free 3-forests in the class $(1^5, 2, 3)$ which are independent, with only four distinct root maxima. So the basis at odd $r$ is not indexed by the root maxima, and a finer invariant is needed.

Fix a class with $k$ trees and use the root-selection inventories of Subsection 4.4 and the notation $U$, $Y$ introduced before Proposition 7.4. Let $u = |U|$, choose a leaf $a$, let $V_{\mathrm{ev}}$ be the even-valency vertices and define $O = Y \setminus (V_{\mathrm{ev}} \cup \{a\})$. Order $U \sqcup O$ with $U$ first, using increasing vertex labels within each part, and write $O = \{o_1, \dots, o_N\}$ in this order.

**Definition 7.7.** An *odd label* is a pair $(S, C)$, with $S \subseteq V_{\mathrm{ev}}$, $C \subseteq O$, satisfying

$$|S| + |C| = k - 1, \qquad 2|C \cap \{o_1, \dots, o_j\}| \le u + j \quad (0 \le j \le N).$$

Let $\rho_C$ pair each closer in $C$ with the most recent unused noncloser of $U \sqcup O$. Let

$$Q_C = \pi_U \prod_{\{v,w\} \in \rho_C} (x_v + x_w), \qquad b_{S,C} = L_Y x_S Q_C, \qquad \lambda(S, C) = \{a\} \sqcup S \sqcup C.$$

The inventory $b_{S,C}$ has the following selection rule. Mark every vertex of $S$, choose one eligible endpoint from each pair of $\rho_C$, and choose one additional vertex of $Y$. Discard repetitions and count identical resulting sets with their weights. Thus the formula in Definition 7.7 specifies a finite family of choices, not an auxiliary polynomial computation.

Compare records lexicographically by membership, testing $a$ first, then the vertices of $V_{\mathrm{ev}}$, then $o_N, \dots, o_1$; membership ranks ahead of nonmembership. The greatest record is $\lambda(S, C) = \{a\} \sqcup S \sqcup C$. It arises uniquely by choosing $a$ as the extra vertex and every closer from its pair. The pair choices contain no even-valency vertex and no $a$, so its weight is one. Intersecting this record with $V_{\mathrm{ev}}$ and $O$ recovers $S, C$. The greatest record in a nonzero weighted list of these inventories therefore cannot cancel.

**Proposition 7.8.** *Let $r$ be odd and let $\mathcal{O}(F)$ be a class of $r$-forests with $n$ vertices, $k$ trees and $u$ vertices of valency $r$. Then the selection inventories $b_{S,C}$ over its odd labels form a basis of the image of the class under $t_{2,r}$ in root coordinates, and*

$$\operatorname{rank}(t_{2,r}) = \binom{n-u-1}{k-1} - \binom{n-u-1}{k-u-2}.$$

*Proof.* Let $2h$ be the number of odd valencies. Every tree has at least two odd-valency vertices, so $h \ge k$. By moving branches in pairs to $a$, reduce every other odd valency to one and every even valency to two. The anchor then has valency $D = 2h - 2k + 1$, and the vertices of $U$ have valency one; their labels remain forbidden as roots.

First suppose there are no even-valency vertices. A forest is a star at $a$ together with $p = k - 1$ disjoint pairs among the other $M = 2h - 1$ vertices. Its root inventory can be recorded by choosing one endpoint from each pair and one extra vertex from $V$. Extra choices from a pair cancel: the two ways to choose its two different endpoints give the same record, while repeated choices are discarded. The remaining extra choices are exactly the vertices of the star. In the compact notation this inventory is $L_V \prod_{bc \in \rho}(x_b + x_c)$. Records containing $a$ recover all the pair selections uniquely. This recovery remains valid after forbidding $U$, since $a \notin U$.

Straighten a partial matching by the four-point move of Lemma 7.2. If an unmatched vertex $b$ lies under an arc $ac$, replace $ac$ by the two alternatives $ab$ and $bc$. Their endpoint choices have the same parity as those of $ac$, because the choice $b$ occurs twice. This move and the disjoint smoothing decrease total arc length; the nested smoothing preserves length and decreases crossings. The process terminates at cross-free partial matchings having no unmatched vertex under an arc.

Such a matching is determined by its closers: each closer must use the most recent unused noncloser, or else there is a crossing or an unused vertex under its arc. Order the blocked vertices

first. A normal form has no permitted selection exactly when a closer lies in $U$, because its partner then also lies in $U$. Every other normal form has a unique greatest permitted record, its closer set $C$. These records distinguish all survivors. With $N = M - u$, they are counted by

$$\binom{N}{p} - \binom{N}{p-u-1}. \tag{1}$$

Indeed, start a walk at height $u$, read the vertices of $O$, and take a down-step at each closer and an up-step elsewhere. The matching rule is possible exactly when the walk stays nonnegative (Figure 33). Reflecting the initial segment through its first visit to $-1$ counts the excluded walks by the second binomial. The final height is nonnegative since $M \geq 2p$.

Restore the even-valency vertices one at a time by Lemma 4.15. For each such vertex there are two types of inventory: the insertion type $\eta_v(P)$, or a compulsory mark $x_v Q$. Let $S$ be the set of vertices restored by the second type. The remaining all-odd forest has $k - |S|$ components. If $L_0 = L_{Y \setminus V_{\mathrm{ev}}}$, its normal forms are $L_0 Q_C$. An insertion at a vertex outside $S$ adds precisely the records obtained by replacing the extra choice $a$ by that vertex. Restoring all of them therefore gives

$$x_S \left( L_0 + \sum_{v \in V_{\mathrm{ev}} \setminus S} x_v \right) Q_C = L_Y x_S Q_C = b_{S,C}.$$

The equality follows from the selection rule: vertices of $S$ are already marked and cannot be chosen again. The unique recovery in Lemma 4.15 proves that these forms account for every inventory. Their distinguished records prove independence as above.

To count the labels, first choose the $j$ compulsory even vertices, then the $k - 1 - j$ closers. This gives

$$\sum_{j=0}^{|V_{\mathrm{ev}}|} \binom{|V_{\mathrm{ev}}|}{j} \left( \binom{N}{k-1-j} - \binom{N}{k-u-2-j} \right) = \binom{N + |V_{\mathrm{ev}}|}{k-1} - \binom{N + |V_{\mathrm{ev}}|}{k-u-2}.$$

For each of the two sums, choose a subset of a disjoint union of sets of sizes $|V_{\mathrm{ev}}|$ and $N$, and classify it by its size in the first set; this is Vandermonde's counting identity [19, (5.27)]. Out-of-range binomials and ranks with $k \leq 0$ are zero. Since $N + |V_{\mathrm{ev}}| = n - u - 1$, the result follows. □

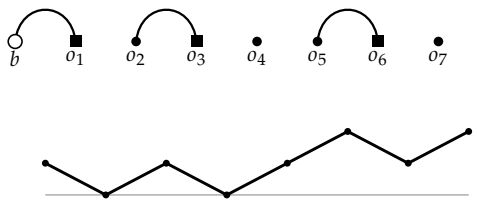



FIGURE 33. A canonical partial matching with $u = 1$, its closers, and its walk

The even and odd formulas can differ for the same $n, k, u$, so the parity split of Theorem B is genuine. Propositions 7.3, 7.4, 7.6 and 7.8 together prove Theorem B.

7.4. **Generation.** We prove completeness by reducing every zero-record list to the empty list using the displayed local relations. Relabeling vertices preserves all orientation counts and all our moves. At $r = \infty$, first order vertices by valency, smooth crossings, and erase cyclic graphs. Proposition 7.6 says that a nonempty terminal list has a nonzero distinguished root record. Thus every relation reduces to zero using only Plücker moves and cycle erasures, proving $\ker(t_{2,\infty}) = (\mathcal{P}, C)$.

For finite $r$, erase graphs with a forbidden subgraph and work in a fixed surviving forest class $\mathbf{d}$. Let $U = \{v : d_v = r\}$ and choose a leaf $a \notin U$ as anchor. A relation can have nonzero unrestricted inventory, but every record in it must contain a blocked vertex after all weights have been

collected. For $W \subseteq U$ write $K_W$ for the attainable inventories that disappear when records meeting $W$ are discarded. This describes $K_W$ by a test on finite sets; the formulas below provide explicit decompositions of such inventories.

**Proposition 7.9.** *Let $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$. Then*

$$\ker(t_{2,r}) = (\mathcal{P},\ \mathcal{C},\ \mathcal{Z}_{r+1},\ \mathcal{D}).$$

*Proof.* The infinite case was proved above. We treat the two finite parities separately, always keeping the original blocked labels fixed during branch transfers.

*Even $r$: isolate one blocked mark at a time.* Move $r-2$ branches from a blocked vertex $b$ to $a$. The vertex $b$ now has valency two, so Lemma 4.15 gives the unique form $\eta_b(P) + x_b Q$. Its records not containing $b$ recover $P$. Consequently

$$K_{\{b\}} = x_b \Pi(\mathbf{d}|_{V\setminus\{b\}} + r\mathbf{e}_a).$$

For several blocked vertices, take any inventory in $K_U$ and use this same decomposition. Filtering the remaining blocked labels makes $P$ zero. Induct on $|U|$ to express $P$ as a sum of attainable inventories each carrying one particular blocked mark $c \neq b$. Applying $\eta_b$ retains that mark, since it replaces only $a$ by $b$. The other part $x_b Q$ already carries $b$. Hence

$$K_U = \sum_{b\in U} K_{\{b\}}.$$

This is an algorithm for sorting the disappearing records by a compulsory blocked mark, with attainability at every step supplied by the insertion lemma.

We realize each term $x_b P_H$ by an attached rainbow. The anchor of $H$ has valency $r+1$, since it was a leaf in the original class. Delete all $r+1$ edges at $a$ in its tree, obtaining $r+2$ pieces, each with a specified attachment vertex (one is $a$). Match these pieces in pairs. For each matched pair, join its attachment vertices to one another and join $b$ to the attachment vertices of all the other pieces. Retain all internal edges of the pieces and all unaffected components. The sum of these forests is a rainbow with those fixed edges attached, in the original valency class (Figure 34).

Count its local two-root records before blocking $b$. A record consisting of $b$ and a vertex $z$ of a piece occurs just once: choose the matched pair containing that piece. A record of two vertices in different matched pairs occurs twice, once for each of their pairs. Two vertices in the same matched pair can never lie in the two different components. Thus all records without $b$ cancel, and precisely the records $\{b, z\}$ remain. With the unaffected choices restored, the inventory is $x_b P_H$, as required.

*Odd $r$: isolate a pair of blocked labels.* Move branches in pairs to $a$ until the other odd vertices have valency one and the even vertices have valency two. First omit the latter. As in Proposition 7.8, all inventories have normal forms consisting of one extra root and the endpoint choices of a canonical partial matching. Order $U$ first. The surviving normal forms have distinct greatest permitted records, so a disappearing inventory is a combination of the normal forms containing a pair entirely in $U$. Every such form already disappears upon forbidding the two labels of that pair. Therefore

$$K_U = \sum_{\{b,c\}\subseteq U} K_{\{b,c\}}.$$

For one fixed pair, order $b, c$ first. A canonical partial matching then vanishes on forbidding $b, c$ exactly when it contains the pair $bc$. Delete this pair. Its root choice was either $b$ or $c$; the extra root can be taken outside the pair because the two ways to choose both $b, c$ cancel. Conversely, attach $bc$ to a normal form on the remaining vertices. This gives exactly

$$K_{\{b,c\}} = (x_b + x_c)\Pi(\mathbf{d}|_{V\setminus\{b,c\}} + (2r-2)\mathbf{e}_a).$$

Here the factor $x_b + x_c$ means append either mark, and collect equal records. The change at the anchor accounts for lowering both blocked valencies from $r$ to one before deleting their pair.

These two decompositions remain valid when the even vertices are restored. Indeed, the insertion lemma recovers the two underlying inventories separately after filtering $U$, so both must disappear. Induction applies to each. Replacing $a$ by a restored vertex, or appending that vertex as a compulsory mark, preserves either of the marks $b, c$ and commutes with the choice between them. Finally undo the paired branch transfers; those preserve the full collection of inventories on both sides. This justifies the decompositions for the original arbitrary valencies, not only for the reduced star classes.

Realize $(x_b + x_c)P_H$ by an attached spider. The anchor has valency $2r-1$. Delete its incident edges to make $2r$ pieces, divide them into two groups of $r$, and attach the groups respectively to $b, c$. Make a second forest with these attachments exchanged. Keep every internal edge and every unaffected component. The two forests form a spider with fixed edges attached and have the original valencies.

The local root record $\{b, c\}$ occurs twice. Two roots in opposite groups occur twice, and two roots in the same group never occur. For any vertex $z$ of a piece, each of $\{b, z\}$ and $\{c, z\}$ occurs once. Thus the sum of the two unrestricted inventories is exactly $(x_b + x_c)P_H$.

*Complete the reduction.* For a given truncated relation $f$, the corresponding construction produces a weighted sum $q$ of attached rainbows or attached spiders with exactly the same unrestricted root inventory. The difference $f - q$ has every unrestricted record zero, so the infinite-case reduction expresses it using Plücker moves and cycle erasures. Restore the forbidden graphs erased at the start; each is generated by its contained cycle or $(r+1)$-star. This proves the asserted generation, as well as the sharper rainbow and spider descriptions. Both exceptional families are cut relations by Propositions 6.1 and 6.2, which gives the stated equality with $\mathcal{D}$. □

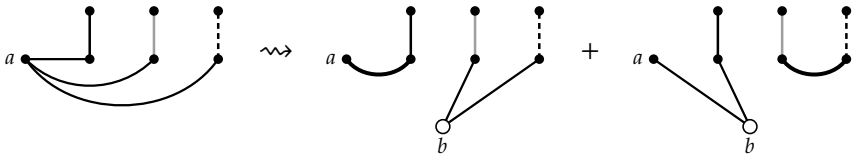



FIGURE 34. An attached rainbow at $r = 2$: two matched pieces are joined, and $b$ is joined to the others

**Proposition 7.10.** *Let $r \in \mathbb{N}_{\geq 2} \cup \{\infty\}$. Choose one cycle through each finite vertex set $V \subseteq \mathbb{N}_{\geq 1}$ with $|V| \geq 2$, and, if $r$ is even, one rainbow for each center $b$ and each set $A$ of $r+2$ leaves, or, if $r$ is odd, one spider for each pair of centers $u, v$ and each set $A$ of $2r$ leaves. Together with $\mathcal{P}$ and $\mathcal{Z}_{r+1}$ these elements form a minimal generating set of* $\ker(t_{2,r})$. *More precisely, in the class of each chosen element every relation is a multiple of that element plus an inherited relation, and no nonzero multiple of it is inherited, while in every other class every relation is inherited. In particular every generating set of* $\ker(t_{2,r})$ *consisting of relations of single classes has an element in the class of each chosen element.*

*Proof.* The preceding reduction expresses every relation by attached Plücker, cycle, star, and exceptional relations. In a fixed class, every use with at least one added edge is inherited. The only possible new terms are therefore generators already in that class. The four kinds of valency vectors are distinct, and the prescribed vertex data distinguish the classes within each kind.

For four vertices of valency one, the graphs are the three perfect matchings. Each fixed root pair is a selection of exactly two of them. A relation thus forces the weights of those two matchings to agree; using the three partitions forces all three weights to agree. The only relation is a scalar multiple of the Plücker relation, and none is inherited from a one-edge class. A star class consists of just its star, and Lemma 6.7 excludes inheritance there.

For cycle, rainbow, and spider classes, the explicit switches of Lemmas 6.12 to 6.14 reduce any relation to one chosen representative with a scalar weight. Its connected count or leaf-incidence count is precisely that scalar. Every inherited list has count zero, so the scalar cannot be removed by inherited moves. All other classes have no new term and consist of inherited relations.

Induction on the number of edges, as in Lemma 6.5, now proves generation by the chosen representatives. Omitting any representative leaves its nonzero count unattainable in its class: smaller relations give count zero, and a relation in a different class cannot be used without adding an edge or changing the valencies. Hence the family is minimal and every generating family of single-class relations must meet each of these classes. □

Together with Proposition 6.11, this proves Theorem A.

## 8. Application to GL-algebras

The graph-counting proofs are complete. We retain the algebraic interpretation here to record all consequences of the original setup. The definitions and the formal passage to quotients in this section are algebraic; the obstruction to finite generation is the connected count already proved above.

8.1. **Stability.** We define GL-algebras from representation stability theory, whose main use is to stabilize the commutative rings, ideals, and maps we discuss. Our reference throughout is Sam and Snowden [32].

Let $\mathfrak{S}_n$ be the group of bijections of $[n]$, regarded as a subgroup of $\mathfrak{S}_{n+1}$ by extending a bijection of $[n]$ to fix $n+1$. The *infinite symmetric group* is

$$\mathfrak{S}_\infty = \bigcup_{n\in\mathbb{N}} \mathfrak{S}_n,$$

the group of bijections of $\mathbb{N}_{\geq 1}$ fixing all but finitely many elements. Similarly let $\mathrm{GL}_n(\Bbbk)$ be the group of invertible $n \times n$ matrices over $\Bbbk$, regarded as a subgroup of $\mathrm{GL}_{n+1}(\Bbbk)$ through the homomorphism sending $A$ to $\left(\begin{smallmatrix} A & 0 \\ 0 & 1 \end{smallmatrix}\right)$. The *infinite general linear group* is

$$\mathrm{GL}_\infty = \bigcup_{n\in\mathbb{N}} \mathrm{GL}_n(\Bbbk),$$

the group of invertible $\mathbb{N}_{\geq 1} \times \mathbb{N}_{\geq 1}$ matrices differing from the identity in all but finitely many entries. By using permutation matrices, $\mathfrak{S}_\infty \subseteq \mathrm{GL}_\infty$.

Recall that we define $\mathbf{V}$ to be $\Bbbk^\infty = \bigcup_n \Bbbk^n$, the standard representation. A representation of $\mathrm{GL}_\infty$ over $\Bbbk$ is *polynomial* if it is a subquotient of a direct sum of tensor powers of $\mathbf{V}$ [15].

**Definition 8.1.** A GL-*algebra* is a commutative unital $\Bbbk$-algebra with an action of $\mathrm{GL}_\infty$ by algebra automorphisms which is polynomial as a representation.

If $W$ is a polynomial representation then $\mathrm{Sym}(W)$ is a GL-algebra with the functorial action [32]. We use two instances. For the direct sum $\mathbf{V} \oplus \mathbf{V}$ with the diagonal action, we have

$$\mathrm{Sym}(\mathbf{V} \oplus \mathbf{V}) \cong \Bbbk[x, y],$$

with $x_i$ and $y_i$ the basis vector $e_i$ of the first and second summand. For the exterior square $\bigwedge^2 \mathbf{V}$, the quotient of $\mathbf{V} \otimes \mathbf{V}$ by the tensors with a repeated factor, with basis $e_i \wedge e_j$ for $i < j$ and action $g \cdot (v \wedge w) = gv \wedge gw$, we have

$$\mathrm{Sym}(\bigwedge^2 \mathbf{V}) \cong \Bbbk[m_{ij}],$$

with $m_{ij} = e_i \wedge e_j$. Hence both algebras in the Plücker map are GL-algebras.

**Proposition 8.2.** *The Plücker map* $\Phi$ *is* $\mathrm{GL}_\infty$*-equivariant, so* $(\mathcal{P})$ *is a* $\mathrm{GL}_\infty$*-stable ideal and* $\mathrm{im}(\Phi)$ *is a* GL*-subalgebra of* $\Bbbk[x, y]$.

*Proof.* Write $g = (g_{pi})$. Think of replacing a vertex label $i$ by a label $p$ with weight $g_{pi}$, independently at every endpoint. For the edge $ij$, expand its two oriented records. Terms with equal new labels cancel. For distinct labels $p < q$, collecting the two assignments of its endpoints gives weight $g_{pi}g_{qj} - g_{qi}g_{pj}$ multiplying the edge record $x_p y_q - x_q y_p$. This is exactly the label-replacement rule on $e_i \wedge e_j$. Replacing labels before orienting an edge or afterward therefore gives the same weighted records. For several edges the choices are independent, so the equality holds for every graph and weighted list. This is equivariance. Zero-record lists and attainable records are preserved, giving the kernel and image assertions. □

The truncation ideal is stable only for the right $r$, and this is the one place the paper's results depend on $r$ being a power of two.

**Lemma 8.3.** *Let* $\mathrm{char}(\Bbbk) = 2$ *and let* $r$ *be a power of* 2 *or* $r = \infty$. *The ideal* $\mathfrak{m}_r$ *is* $\mathrm{GL}_\infty$*-stable. Consequently* $\Bbbk[x, y]/\mathfrak{m}_r$ *inherits the structure of a* GL*-algebra,* $p_r$ *is* $\mathrm{GL}_\infty$*-equivariant,* $t_{2,r}$ *is a map of* GL*-algebras, and* $\ker(t_{2,r})$ *is a* GL*-stable ideal of* $\Bbbk[m_{ij}]$ *containing* $\ker(\Phi)$.

*Proof.* Consider $2^t$ copies of one label and replace each copy independently by a new label, the weight being the product of its replacement weights. Rotate the $2^t$ positions cyclically. Every orbit has power-of-two size. An orbit of size one is a constant assignment; every nonconstant orbit has even size and cancels in characteristic two. The weight and the resulting unordered record are constant on an orbit. Only constant assignments remain. Applied to two copies of $x_i$ and $r = 2^t$ copies of $y_i$ (Figure 35), this proves the familiar identities

$$g \cdot x_i^2 = \sum_p g_{pi}^2 x_p^2, \qquad g \cdot y_i^r = \sum_p g_{pi}^r y_p^r.$$

Thus a forbidden repeated-label record remains a sum of forbidden records after any label replacement. At $r = \infty$ only the first identity is needed. The truncation is stable, and the induced maps and their kernels have the stated equivariance properties. □

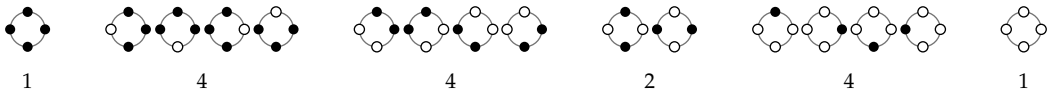



FIGURE 35. The rotation orbits of the label choices for four copies with two labels, and their sizes

A *twisted commutative algebra* is a graded $\Bbbk$-algebra $A = \bigoplus_{n \geq 0} A_n$ with an action of $\mathfrak{S}_n$ on each $A_n$, such that multiplication is equivariant and commutative up to the block-swapping permutation. In characteristic zero, Schur–Weyl duality identifies twisted commutative algebras with polynomial GL-algebras [32], and this is how results stated for twisted commutative algebras apply to the algebras of this paper. In positive characteristic the identification fails [29, Section 1.6], and we work with GL-algebras throughout.

8.2. **Noetherianity.** We recall the standard results on the noetherianity of GL-algebras, focusing on the ones appearing in the map $t_{2,r}$.

Let $A$ be a GL-algebra. An ideal $I \subseteq A$ is GL-*stable* if $g \cdot f \in I$ for all $g \in \mathrm{GL}_\infty$ and $f \in I$, and is GL-*finitely generated* if there are $f_1, \ldots, f_s \in I$ whose orbits $\mathrm{GL}_\infty \cdot f_1, \ldots, \mathrm{GL}_\infty \cdot f_s$ generate $I$ as an ordinary ideal.

**Definition 8.4.** A GL-algebra $A$ is $\mathfrak{S}_\infty$*-noetherian* if it satisfies the ascending chain condition on $\mathfrak{S}_\infty$-stable ideals, GL-*noetherian* if it satisfies the ascending chain condition on GL-stable ideals, equivalently if every GL-stable ideal is GL-finitely generated, and *topologically* GL-*noetherian* if it satisfies the ascending chain condition on GL-stable radical ideals.

**Lemma 8.5.** *Let $A$ be a* GL*-algebra. If $A$ is $\mathfrak{S}_\infty$-noetherian then it is* GL*-noetherian, and if $A$ is* GL*-noetherian then it is topologically* GL*-noetherian. If $I \subseteq A$ is* GL*-stable then the* GL*-stable ideals of $A/I$ correspond, preserving inclusion, to those of $A$ containing $I$.*

*Proof.* Since $\mathfrak{S}_\infty \subseteq \mathrm{GL}_\infty$, every GL-stable ideal is $\mathfrak{S}_\infty$-stable, so an ascending chain of GL-stable ideals is one of $\mathfrak{S}_\infty$-stable ideals and the first implication follows. The second holds because radical ideals are among the ideals. For the equivalence in the definition, given the ascending chain condition a GL-stable $I$ has a maximal subideal generated by finitely many orbits, and adjoining the orbit of any element of $I$ outside it contradicts maximality. Finally the quotient map $A \to A/I$ is equivariant and surjective, so $J \mapsto J/I$ is an inclusion preserving bijection from the GL-stable ideals of $A$ containing $I$ to those of $A/I$, with inverse the preimage. □

In particular a quotient of a GL-noetherian algebra by a GL-stable ideal is GL-noetherian. The algebra $\Bbbk[x, y]/\mathfrak{m}_r$ is well behaved.

The first statement of the following is due to Cohen [5, 6]; see also Hillar and Sullivant [23, Theorem 1.1].

**Proposition 8.6.** *Over any field $\Bbbk[x, y]$ is $\mathfrak{S}_\infty$-noetherian, hence* GL*-noetherian, and so is $\Bbbk[x, y]/\mathfrak{m}_r$ whenever $r$ is a power of 2 or $r = \infty$.*

*Proof.* The first statement is Cohen's theorem, and Lemma 8.5 gives GL-noetherianity. For the quotient, $\mathfrak{m}_r$ is GL-stable for these $r$ by Lemma 8.3, and Lemma 8.5 passes noetherianity to the quotient. □

The algebra $\Bbbk[m_{ij}]$ is more delicate. Nagpal, Sam and Snowden [29] proved that in characteristic zero $\Bbbk[m_{ij}]$ is GL-noetherian, as is every quotient by a GL-stable ideal, and Draisma [12] proved that $\Bbbk[m_{ij}]$ is topologically GL-noetherian over any infinite field. In characteristic two the algebraic statement fails.

The following is due to Ganapathy [16, Theorem 1.3].

**Proposition 8.7.** *Suppose* $\mathrm{char}(\Bbbk) = 2$*. Then the* GL*-algebra* $\mathrm{Sym}(\bigwedge^2 \mathbf{V})$ *is not* GL*-noetherian.*

Ganapathy's proof exhibits an infinite ascending chain of GL-stable ideals, and non-noetherian GL-algebras have since been found in every positive characteristic [17]. The following proposition gives an independent proof: our kernels are explicit GL-stable ideals which are not GL-finitely generated, and their minimal generators, determined in Theorem A, show exactly where finite generation fails.

**Proposition 8.8.** *Let $r = 2^t$ or $r = \infty$. Then* $\ker(t_{2,r})$ *is not* GL*-finitely generated. In particular $\Bbbk[m_{ij}]$ is not* GL*-noetherian.*

*Proof.* The kernel is stable by Lemma 8.3. Label replacement changes neither the number of edges nor its degree. Finitely many orbit generators would consequently have a common upper bound $B$ on their numbers of edges. We may separate their finitely many degree components, which are again relations; the resulting orbit family still generates.

Choose a cycle with more than $B$ edges. Any expression for it using those orbit generators must attach at least one edge to every term. Separate this expression into valency classes, as allowed by Proposition 2.2. In the cycle's class it is a sum of inherited relations. But every such list has connected count zero, whereas the cycle has connected count one by Lemma 6.6. This contradiction rules out a finite orbit-generating family.

The same edge-count argument uses only that relabeling preserves relations. It therefore also rules out finitely many $\mathfrak{S}_\infty$-orbit generators for every finite $r \geq 2$ and for $r = \infty$. □

By Lemma 8.3 the ideal $\ker(t_{2,r})$ is GL-stable, so this proves Theorem C.

**Large Language Models.** The author used Claude (Anthropic), through Claude Code and the Claude app, for the following work: Claude Opus wrote code to test the invariants of Sections 3 and 7 and to compute the relations of Section 6 for $2 < r < \infty$. Opus was also used to discuss ideas and proof strategies throughout the progression of the paper. Opus helped with editing, proposing expository edits, and checking bibliographic data. Finally, Opus conducted retrospective literature searches, which did not affect the development of the results, and found the related work of Alon and Tarsi [1] (Section 2) and of Casabella and Yu [4] (Sections 4, 6 and 7). All mathematical content was discovered, developed, and checked by the author, who takes full responsibility for the paper.

## References


[1] N. Alon and M. Tarsi, *Colorings and orientations of graphs*, Combinatorica **12** (1992), no. 2, 125–134.
[2] F. Ardila and A. Postnikov, *Combinatorics and geometry of power ideals*, Trans. Amer. Math. Soc. **362** (2010), no. 8, 4357–4384.
[3] W. Bruns and U. Vetter, *Determinantal Rings*, Lecture Notes in Mathematics **1327**, Springer-Verlag, 1988.
[4] L. Casabella and T. Yu, *Standard monomial theory modulo Frobenius in characteristic two*, J. Pure Appl. Algebra **229** (2025), no. 9, Paper No. 108051.
[5] D. E. Cohen, *On the laws of a metabelian variety*, J. Algebra **5** (1967), 267–273.
[6] D. E. Cohen, *Closure relations, Buchberger's algorithm, and polynomials in infinitely many variables*, in: Computation Theory and Logic, Lecture Notes in Comput. Sci. **270**, Springer, 1987, 78–87.
[7] C. De Concini, D. Eisenbud, and C. Procesi, *Hodge algebras*, Astérisque **91**, Société Mathématique de France, 1982.
[8] J. A. De Loera, B. Sturmfels, and R. R. Thomas, *Gröbner bases and triangulations of the second hypersimplex*, Combinatorica **15** (1995), 409–424.
[9] A. de Mier, *k-noncrossing and k-nonnesting graphs and fillings of Ferrers diagrams*, Combinatorica **27** (2007), no. 6, 699–720.
[10] R. Diestel, *Graph Theory*, 5th ed., Graduate Texts in Mathematics **173**, Springer, 2017.
[11] P. Doubilet, G.-C. Rota, and J. Stein, *On the foundations of combinatorial theory IX: Combinatorial methods in invariant theory*, Studies in Appl. Math. **53** (1974), 185–216.
[12] J. Draisma, *Topological Noetherianity of polynomial functors*, J. Amer. Math. Soc. **32** (2019), no. 3, 691–707.
[13] W. Fulton, *Young Tableaux*, London Mathematical Society Student Texts **35**, Cambridge University Press, 1997.
[14] D. Gale, *Optimal assignments in an ordered set: an application of matroid theory*, J. Combinatorial Theory **4** (1968), 176–180.
[15] K. Ganapathy, GL-*algebras in positive characteristic I: the exterior algebra*, Selecta Math. (N.S.) **30** (2024), no. 4, Paper No. 63.
[16] K. Ganapathy, *Non-noetherian* GL-*algebras in characteristic two*, preprint (2024), `arXiv:2408.04630`.
[17] K. Ganapathy, *Ideal-theoretic non-noetherianity of polynomial functors in positive characteristic*, preprint (2025), `arXiv:2506.15134`.
[18] Z. Gao, C. Raicu, and K. VandeBogert, *Some questions arising from the study of cohomology on flag varieties*, preprint (2023), `arXiv:2308.04334`.
[19] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science*, 2nd ed., Addison-Wesley, 1994.
[20] J. Harris, *Algebraic Geometry: A First Course*, Graduate Texts in Mathematics **133**, Springer-Verlag, 1992.
[21] J. Herzog and T. Hibi, *Monomial Ideals*, Graduate Texts in Mathematics **260**, Springer-Verlag, 2011.
[22] J. Herzog, T. Hibi, F. Hreinsdóttir, T. Kahle, and J. Rauh, *Binomial edge ideals and conditional independence statements*, Adv. in Appl. Math. **45** (2010), no. 3, 317–333.
[23] C. J. Hillar and S. Sullivant, *Finite Gröbner bases in infinite dimensional polynomial rings and applications*, Adv. Math. **229** (2012), no. 1, 1–25.
[24] W. V. D. Hodge, *Some enumerative results in the theory of forms*, Proc. Cambridge Philos. Soc. **39** (1943), 22–30.
[25] O. Holtz and A. Ron, *Zonotopal algebra*, Adv. Math. **227** (2011), no. 2, 847–894.
[26] B. Howard, J. Millson, A. Snowden, and R. Vakil, *The ideal of relations for the ring of invariants of n points on the line*, J. Eur. Math. Soc. (JEMS) **14** (2012), no. 1, 1–60.
[27] A. B. Kempe, *On regular difference terms*, Proc. London Math. Soc. **25** (1894), 343–359.
[28] E. Miller and B. Sturmfels, *Combinatorial Commutative Algebra*, Graduate Texts in Mathematics **227**, Springer-Verlag, 2005.
[29] R. Nagpal, S. V. Sam, and A. Snowden, *Noetherianity of some degree two twisted commutative algebras*, Selecta Math. (N.S.) **22** (2016), no. 2, 913–937.
[30] M. Ohtani, *Graphs and ideals generated by some* 2*-minors*, Comm. Algebra **39** (2011), no. 3, 905–917.

[31] A. Postnikov and B. Shapiro, *Trees, parking functions, syzygies, and deformations of monomial ideals*, Trans. Amer. Math. Soc. **356** (2004), no. 8, 3109–3142.
[32] S. V. Sam and A. Snowden, *Introduction to twisted commutative algebras*, preprint (2012), `arXiv:1209.5122`.
[33] R. P. Stanley, *Enumerative Combinatorics, Volume 2*, Cambridge Studies in Advanced Mathematics **62**, Cambridge University Press, 1999.
[34] R. P. Stanley, *Catalan Numbers*, Cambridge University Press, 2015.
[35] R. H. Villarreal, *Rees algebras of edge ideals*, Comm. Algebra **23** (1995), no. 9, 3513–3524.
[36] R. H. Villarreal, *Monomial Algebras*, 2nd ed., Monographs and Research Notes in Mathematics, CRC Press, 2015.

Department of Mathematics, Columbia University, New York, NY
*Email address*: `ss7675@columbia.edu`
*URL*: `https://sanay-s.com`