hexagon:2611.00238v2

ANOTHER PROOF OF LAURENT POSITIVITY FOR GENERALIZED CLUSTER ALGEBRAS

QIYUE TANG

Abstract. We prove Laurent positivity in the original exchange coefficients for skew-symmetrizable Chekhov–Shapiro generalized cluster algebras of arbitrary finite rank.

Contents

1. Introduction 1
2. Generalized cluster patterns 3
3. Finite Gale algebras 5
4. Cut degenerations and chamber algebras 7
5. Separating the core and attached factors 11
6. Finite zero schemes and Young characters 16
7. A polynomial Harder–Narasimhan decomposition 18
8. The attached algebra and the two-input wall function 21
9. Positive formal wall operations 26
10. Mutation and finite Laurent fibers 31
11. Laurent positivity 36
12. Finite algebras for the coefficients 37
References 41

1. Introduction

1.1. **The positivity problem.** Fomin and Zelevinsky introduced cluster algebras in [9], motivated by total positivity and dual canonical bases. A cluster is a transcendence basis, and mutation replaces one of its variables by a rational expression determined by a binomial exchange relation. Their Laurent phenomenon states that every cluster variable is a Laurent polynomial in the variables of any cluster. They also conjectured that the coefficients of these expansions are nonnegative integer combinations of elements of the coefficient group. The caterpillar argument of [10, Theorem 2.1] gives a criterion for the Laurent property in polynomial exchange patterns. In [11], Fomin and Zelevinsky developed principal coefficients, F-polynomials, and separation formulas, organizing the dependence of cluster variables on the coefficient data.

Rank-two formulas provided explicit positive expansions and positive bases. Lee, Li, and Zelevinsky [16] constructed the greedy basis using compatible pairs on Dyck paths. Lee and Schiffler [17] proved Laurent positivity for all skew-symmetric cluster algebras. Gross, Hacking, Keel, and Kontsevich [15] established the skew-symmetrizable case through cluster scattering diagrams and theta functions. Their construction identifies cluster monomials with theta functions and derives Laurent positivity from positive wall functions.

Date: November 9, 2026.
2020 *Mathematics Subject Classification.* Primary 13F60; Secondary 05E05, 05E10, 16G20.
Key words and phrases. Generalized cluster algebra, Laurent positivity, scattering diagram, finite algebra, Young permutation representation, Harder–Narasimhan factorization.

Chekhov and Shapiro introduced generalized cluster algebras in a 2011 preprint, published as [7], to describe Teichmüller spaces of surfaces with orbifold points. The orbifold mutations lead to quadratic exchange polynomials; their algebraic definition allows homogeneous exchange polynomials of arbitrary prescribed degrees. The coefficient tuples mutate with the exchange matrix. They proved the generalized Laurent phenomenon in Theorem 2.5 and established positivity for the motivating orbifold transformations. Conjecture 5.1 asks for positivity for reciprocal exchange polynomials with positive coefficients. We use the corrected third arXiv version of their paper.

The rank-two and coefficient theories developed alongside this conjecture. Rupel [25] extended greedy bases to generalized rank-two algebras with monic palindromic exchange polynomials with positive coefficients, using compatible gradings on maximal Dyck paths. Nakanishi [19] established formulas in terms of c-vectors, g-vectors, and F-polynomials under normalization and quasi-reciprocity conditions. Nakanishi and Rupel [20] related these data to left and right ordinary companion cluster algebras and recovered the companion variables by specializations of generalized F-polynomials.

Cheung, Kelley, and Musiker [8] constructed scattering diagrams and theta functions for reciprocal generalized cluster algebras. Mou [18] developed generalized scattering diagrams by factoring the exchange polynomials into binomials and proved positivity in the factor parameters. The original exchange coefficients give a different positivity question. If

$$P(z) = \prod_i (1 + \alpha_i z),$$

then those coefficients are the elementary symmetric functions $e_s(\alpha)$. For example,

$$\alpha_1^2 + \alpha_2^2 = e_1(\alpha)^2 - 2e_2(\alpha).$$

Thus a positive expression in the factor parameters can have a negative coefficient when written in the original exchange generators.

Burcroff, Lee, and Mou [4, 5] resolved this coefficient positivity problem. Their result was announced in 2024; the PNAS article and the detailed proof appeared in 2025. They introduced tight gradings on Dyck paths and used them to count rank-two wall coefficients in the original polynomial coefficients. Their two proofs [5, Sections 5.4–5.5] use shadowed gradings and the one-bending property of broken lines; the second compares the shadowed-grading and broken-line recursions directly. They then proved positivity for several incoming walls, passed to arbitrary rank, and obtained Laurent positivity and positive theta-function structure constants. Their Laurent theorem [5, Theorem 7.9 and Section 7.2] covers arbitrary finite rank, skew-symmetrizable exchange matrices, frozen variables, and nonreciprocal exchange polynomials. Theorem 2.1 below gives another proof of this Laurent-positivity statement.

Quiver moduli provide a related description of the local wall coefficients. Reineke and Weist [23] expressed tropical-vertex functions through Euler characteristics of stable bipartite quiver moduli. Burcroff, Lee, Mou, Musiker, and Reineke [6] combined these formulas with tight gradings to study coefficient formulas and Weyl symmetry. On the structural side, Ramos and Whiting [22] realized generalized cluster algebras of geometric type as subquotients of ordinary cluster algebras. Akagi and Nakanishi [1] extended the comparison to arbitrary semifields and composite cluster patterns; their Remark 6.5 discusses the distinction between positivity in factor parameters and in exchange coefficients. In the quantum setting, Ye [27] constructed scattering diagrams for reciprocal data with a specified compatibility normalization and proved positivity in the factor parameters.

1.2. **The attached algebra.** Our local construction uses a finite algebra with a particle-permutation action. Fix coprime integers $a, b > 0$, positive integers p, q with $pa - qb = 1$, and

$k \geq 1$. In the variables $a_1, \dots, a_{ka}, b_1, \dots, b_{kb}$, consider the system with parameter T:

$$a_i^p \prod_{j=1}^{kb} (a_i - b_j) = T^b, \qquad (-b_j)^q \prod_{i=1}^{ka} (a_i - b_j) = T^a. \tag{1.1}$$

The group $G_k = S_{ka} \times S_{kb}$ permutes the particles of each color. The horizontal solution algebra, completed at $T = 0$, is finite free over $\mathbb{C}[[T]]$. Its core consists of the factors whose coordinates tend to zero. Its attached factor consists of those whose coordinates have nonzero limits; let $\mathcal{D}_k^{p,q}$ be its special fiber. Each selected point contributes its full Artin local algebra. Section 3 treats the same construction with any positive integer multiplicity on the mixed factors.

For finite alphabets α and β, let $f_{a,b}(U)$ be the two-input wall function in direction (a, b), with formal ray variable U and input polynomials $\prod_i (1 + \alpha_i z)$ and $\prod_j (1 + \beta_j z)$. We prove

$$[U^k] f_{a,b}(U) = \omega_\alpha \omega_\beta \ \mathrm{ch}(\mathcal{D}_k^{p,q}). \tag{1.2}$$

Here ch is the Frobenius characteristic, and ω is the involution of the ring of symmetric functions sending the complete homogeneous function h_j to the elementary function e_j; apply it before evaluating in the finite alphabets. The equality holds for every positive pair (p, q) satisfying $pa - qb = 1$ and every k. An equivariant deformation shows that $\mathcal{D}_k^{p,q}$ is a sum of Young permutation representations. Consequently

$$[U^k] f_{a,b}(U) = \sum_{\lambda \vdash ka,\, \mu \vdash kb} n_{\lambda\mu} e_\lambda(\alpha) e_\mu(\beta), \qquad n_{\lambda\mu} \in \mathbb{Z}_{\geq 0}, \tag{1.3}$$

where the indices are partitions and $e_\lambda = \prod_i e_{\lambda_i}$. This gives positivity in the exchange coefficients at every multiple of the primitive direction. The integers $n_{\lambda\mu}$ are the universal coefficients counted by tight gradings in [6, Theorem 4.10]. Equation (1.2) derives them from the attached factor by an equivariant algebraic construction.

Products of linear forms are classical objects of Gale duality [2]. Toric quantum Stanley–Reisner presentations and finite factors selected by valuations occur in González–Woodward [14]. The construction here identifies the full symmetric-group character of the attached factor of (1.1) with a wall coefficient. A deformation separating the core and attached particles produces an induced-character identity, uniformly over the positive Bézout pairs. This identity is compared with a framed Harder–Narasimhan factorization to prove (1.2).

The proof includes a cut degeneration, the required finite-algebra comparisons, and the equivariant deformation. Rimmer studied regular equivariant deformations and their permutation representations [24, Corollary 6.2.8 and Section 8.1]. Here polynomial Hermite interpolation constructs the deformation and controls the directions transverse to the coordinate-equality strata. The thin-quiver quotients and alternation have antecedents in Franzen [12], and the polynomial factorization is a form of the Chow–Hall factorization of Franzen and Reineke [13]. We adapt the separated-center proof and formal transport arguments of [21, Sections 3–4]. The passage from two inputs to arbitrary rank follows the scattering-diagram reduction of [5, Section 9]; we give its labeled induction and mutation comparison with integral ambient normals. The generalized Laurent phenomenon [7], convex-polytope shellability [3], and standard completion and regular-sequence results [26] supply the background facts.

Finally, Section 12 constructs finite rational algebras of dimensions $n_{\lambda\mu}$ and gives a finite construction for every integer coefficient of a universal Laurent expansion associated with a prescribed mutation word.

2. Generalized cluster patterns

We write $\mathbb{N} = \mathbb{Z}_{\geq 0}$ and $[c]_+ = \max(c, 0)$. All lattices are free abelian groups of finite rank. The coefficient group $\mathbb{P}$ is a torsion-free abelian multiplicative group; hence $\mathbb{Z}[\mathbb{P}]$ is an integral domain.

Its positive semiring $\mathbb{N}[\mathbb{P}]$ consists of finite sums of elements of $\mathbb{P}$ with nonnegative integral coefficients. Frozen variables are invertible throughout and can equivalently be incorporated into this group.

2.1. **Degree-separated exchange matrices.** Fix n mutable and $N-n$ frozen indices. A seed consists of algebraically independent variables $x_1,\dots,x_N$, an integral $N\times n$ matrix $\widetilde{B}=(b_{ij})$, positive degrees $r_1,\dots,r_n$, and coefficient tuples $q_i=(q_{i,0},\dots,q_{i,r_i})$ in $\mathbb{P}$. The principal $n\times n$ part B is skew-symmetrizable:

$$DB=-B^tD,\qquad D=\operatorname{diag}(\delta_1,\dots,\delta_n),\quad \delta_i\in\mathbb{Z}_{>0}.$$

The matrix is degree-separated: the usual physical exchange matrix is $B\operatorname{diag}(r_i)$. This convention avoids divisibility conditions on its columns. If a physical matrix and divisible degrees are given first, division of the i-th column by r_i gives the present matrix. Skew-symmetrizability is preserved by this change: a symmetrizer D_{phys} is replaced by $D_{\mathrm{phys}}\operatorname{diag}(r_i)^{-1}$, followed by multiplication by a common denominator.

For a mutable index k, let

$$M_k^+=\prod_{i=1}^N x_i^{[b_{ik}]_+},\qquad M_k^-=\prod_{i=1}^N x_i^{[-b_{ik}]_+}.$$

Mutation replaces x_k by the rational function determined by

$$x_kx_k'=\sum_{s=0}^{r_k}q_{k,s}(M_k^+)^{r_k-s}(M_k^-)^s,\qquad x_i'=x_i\quad(i\neq k).\tag{2.1}$$

The matrix changes by

$$b_{ij}'=\begin{cases}-b_{ij}, & i=k\text{ or }j=k,\\ b_{ij}+r_k\bigl([b_{ik}]_+b_{kj}+b_{ik}[-b_{kj}]_+\bigr), & i,j\neq k.\end{cases}\tag{2.2}$$

Frozen rows obey the same rule. The degrees and symmetrizer remain fixed.

2.2. **Coefficients and representatives.** Put $\bar q_{i,s}=q_{i,s}/q_{i,0}$. The coefficient rule is

$$\begin{aligned}q_{k,s}'&=q_{k,r_k-s},\\ \bar q_{j,s}'&=\begin{cases}\bar q_{j,s}q_{k,0}^{-sb_{kj}}, & b_{kj}\geq 0,\\ \bar q_{j,s}q_{k,r_k}^{-sb_{kj}}, & b_{kj}\leq 0,\end{cases}\qquad(j\neq k).\end{aligned}\tag{2.3}$$

For $j\neq k$, the common scale $q_{j,0}'\in\mathbb{P}$ is free. Equations (2.1)–(2.3) are the Chekhov–Shapiro rules, with the order of each coefficient tuple chosen to agree with (2.1). Mutation reverses the exchange polynomial at the mutated index, including for nonreciprocal coefficient tuples.

A pattern assigns seeds to the vertices of the regular labeled n-valent tree, with adjacent seeds related by the corresponding mutation. After one seed is fixed, choose the common scales along every edge oriented away from it. The *canonical representative* used below takes $q_{j,0}'=q_{j,0}$ for $j\neq k$ on those edges. All other permitted choices are related by the coefficient rescalings proved in Section 10.

The generalized cluster algebra is the $\mathbb{Z}[\mathbb{P}]$-subalgebra of $\operatorname{Frac}(\mathbb{Z}[\mathbb{P}])(x_1,\dots,x_N)$ generated by the mutable variables at all seeds, with the fixed frozen variables and their inverses included in the coefficient ring. A cluster monomial has nonnegative exponents in the mutable variables of one seed and arbitrary integral exponents in the frozen variables.

Theorem 2.1 (Laurent positivity)**.** *In a Chekhov–Shapiro generalized cluster pattern of arbitrary finite rank, every cluster monomial has a Laurent expansion in the variables of any seed with coefficients in* $\mathbb{N}[\mathbb{P}]$.
Take the exchange coefficients of a chosen seed to be independent symbols, with invertible endpoints, and use the canonical representative rooted at that seed. Every such Laurent expansion has coefficients in

$$R_+ = \mathbb{N}[q_{i,0}^{\pm 1}, q_{i,r_i}^{\pm 1}, q_{i,s}\ (1 \le i \le n,\ 0 < s < r_i)]. \tag{2.4}$$

The first assertion follows from the universal second assertion and the rescaling argument below. If the endpoints of the reference exchange polynomials equal one, (2.4) is the polynomial semiring in their intermediate coefficients. An arbitrary permitted normalization may introduce inverses of intermediate coefficients; the invariant assertion for such a representative is positivity in $\mathbb{N}[\mathbb{P}]$.

We use the following Laurent theorem.

Theorem 2.2 (Chekhov–Shapiro Laurent phenomenon, [7, Theorem 2.5])**.** *Every variable in a generalized cluster pattern is a Laurent polynomial over* $\mathbb{Z}[\mathbb{P}]$ *in the variables of any seed.*

The theorem applies to the coefficient rule (2.3), including nonreciprocal tuples. Its free-abelian coefficient-group formulation in [7] gives the statement above: a finite mutation word involves finitely many coefficient elements and common scales, whose subgroup in $\mathbb{P}$ is finitely generated and torsion-free, hence free abelian. To include arbitrary frozen rows, choose for each frozen index f a positive δ_f divisible by every mutable δ_i. Add frozen columns with

$$b_{if} = -\delta_f b_{fi}/\delta_i \quad (i \le n), \qquad b_{fg} = 0 \quad (f, g > n).$$

The resulting square matrix is integral and skew-symmetrizable by $\operatorname{diag}(\delta_1, \dots, \delta_N)$. Assign degree one to the added columns and apply the Laurent theorem along the prescribed mutable directions. The restriction is the original rectangular pattern, proving the stated formulation.

Remark 2.3. Zero intermediate coefficients, or intermediate coefficients that are positive sums in another coefficient group, are obtained by specializing the universal polynomial expression after Theorem 2.1 is proved. Endpoints are specialized to units. Since no intermediate coefficient is inverted in (2.4), these specializations preserve both the exchange relation and positivity.

3. Finite Gale algebras

Fix relatively prime positive integers a, b, an integer $m \ge 1$, and positive integers p, q such that

$$pa - qb = 1. \tag{3.1}$$

Such a positive pair exists, since one can add a sufficiently large multiple of (b, a) to any integral solution. Throughout the next three sections these integers are fixed. Set $G_0 = \{1\}$. For $k \ge 1$, put

$$A = ka, \qquad B_0 = kb, \qquad G_k = S_A \times S_{B_0}.$$

The two symmetric groups act by permuting the coordinates of their respective colors. Consider the square polynomial system

$$a_i^p \prod_{j=1}^{B_0} (a_i - b_j)^m = T^b \qquad (1 \le i \le A), \tag{3.2}$$

$$(-b_j)^q \prod_{i=1}^{A} (a_i - b_j)^m = T^a \qquad (1 \le j \le B_0). \tag{3.3}$$

All factors selected from this system will be full scheme-theoretic factors. In particular, a multiple solution contributes its entire Artin local algebra.

3.1. **A numerical identity and the valuation criterion.** Set

$$N_A = p + mkb, \qquad N_B = q + mka, \qquad M_k = \max\{kaN_A, kbN_B\}.$$

For integers $0 \le u \le ka$, $0 \le v \le kb$, not both zero, write

$$\Delta = bu - av, \qquad D_0 = N_A u - N_B v.$$

Equation (3.1) gives

$$bD_0 = N_A\Delta + v. \tag{3.4}$$

Consequently D_0 has the sign of Δ whenever $\Delta \neq 0$. Indeed, the only case requiring an estimate is $\Delta < 0$, when $N_A\Delta + v \le -N_A + kb < 0$. If $\Delta = 0$, relative primality gives $(u, v) = \ell(a, b)$ for an integer $1 \le \ell \le k$, and then $D_0 = \ell$. Thus D_0 is nonzero for every nonempty choice of particles, and $|D_0| \le M_k$ by the bounds on u and v.

First replace the right sides of (3.2)–(3.3) by independent nonzero parameters. This gives the polynomial moment map

$$(\mathbf{a}, \mathbf{b}) \longmapsto \left(a_i^p \prod_j (a_i - b_j)^m, (-b_j)^q \prod_i (a_i - b_j)^m \right)_{i,j}.$$

Its inverse image of the target torus is finite over that torus. To see this, apply the valuative criterion. Let R be a discrete valuation ring whose target parameters are units, and suppose a solution is given over its fraction field. If the smallest coordinate valuation ν were negative, divide all coordinates by a coordinate of valuation ν and choose one nonzero residue value of the resulting tuple. Let its cluster contain u particles of the first color and v of the second. Multiply the first-color equations in this cluster and divide by the second-color equations in it. Every internal mixed difference cancels. A difference from a cluster coordinate to a coordinate outside the cluster has valuation exactly ν. The valuation of the quotient is therefore

$$\nu\{pu - qv + mu(B_0 - v) - mv(A - u)\} = \nu D_0.$$

The right side of the quotient is a unit. This contradicts $D_0 \neq 0$. Hence all coordinates are integral over R. Their equations have unit right sides, so the extension remains in the inverse image of the target torus. Uniqueness follows from separatedness. The DVR criterion applies to this finite-type morphism over a Noetherian base; it is proper, and an affine proper morphism is finite.

For the central system, fix an extension of the T-adic valuation to $\overline{\mathbb{C}((T))}$, normalized by $\mathrm{val}(T) = 1$. For a DVR specialization the same calculation, with that DVR's valuation, yields

$$\nu D_0 = \mathrm{val}(T)\Delta. \tag{3.5}$$

The sign information in (3.4) excludes $\nu < 0$. For the normalized central valuation, if all coordinates have positive valuation, apply the argument to their smallest valuation. Then $\Delta \neq 0$, and the two nonzero integers Δ, D_0 have the same sign. Thus

$$\nu = \frac{|\Delta|}{|D_0|} \ge \frac{1}{M_k}. \tag{3.6}$$

Every mixed difference then has valuation at least ν.

3.2. **The possible nonzero centers.** Consider a cluster with common nonzero residue c, containing u first-color and v second-color particles. In the quotient of its equations the internal mixed differences cancel, and the framing and external factors are units. Thus $0 = \mathrm{val}(T)(bu - av)$, so $bu - av = 0$. The cluster contains $\ell(a, b)$ particles. Every nonzero cluster is balanced in this sense; the zero cluster is balanced as well, because the total charge is $k(a, b)$.

The framing contribution to the limiting quotient for this cluster is $(-1)^{qb\ell}c^{\ell}$. Its interaction with an external balanced cluster of charge $h(a,b)$ contributes $(-1)^{mab\ell h}$. Hence

$$c^{\ell} = (-1)^{qb\ell + mab\ell(k-\ell)} = \sigma_k^{\ell}, \qquad \sigma_k = (-1)^{qb+mab(k-1)}.$$

The last equality uses the evenness of $\ell(\ell-1)$. Define

$$P_k(z) = \prod_{h=1}^{k} (z^h - \sigma_k^h). \tag{3.7}$$

Every nonzero center is a root of P_k, and $P_k(0) \neq 0$.

Proposition 3.1. *The horizontal closure of (3.2)–(3.3), completed at $T = 0$, is a finite free $\mathbb{C}[[T]]$-algebra. It has G_k-stable direct factors consisting respectively of the solutions for which all particle coordinates have positive valuation and those for which all have valuation zero. Denote the special-fiber algebras of the full family and these two factors by*

$$\mathcal{B}_k^{p,q}, \qquad \mathcal{C}_k^{p,q}, \qquad \mathcal{D}_k^{p,q}.$$

Their G_k-characters agree with those of their generic fibers. We set $\mathcal{B}_0^{p,q} = \mathcal{C}_0^{p,q} = \mathcal{D}_0^{p,q} = \mathbb{C}$.

Proof. By horizontal closure we mean that the equation ideal is saturated with respect to T. Hence T is a nonzerodivisor, every minimal prime avoids T, and every special support point is a specialization of the generic locus. The generic fiber is finite by the moment-map argument. Equation (3.5) excludes a point at infinity of the projective horizontal closure above $T = 0$: any such point would give, after taking a curve through it and normalizing, a solution with a coordinate of negative valuation. Equation (3.7) then shows that the special support is finite. The projective closure is consequently quasi-finite near this fiber and is finite there. Its affine horizontal closure has the same germ. Flatness of Noetherian completion preserves the injectivity of multiplication by T, giving a finite torsion-free, hence free, module over the DVR $\mathbb{C}[[T]]$.

The special support is contained in the finite set of tuples whose coordinates belong to $\{0\} \cup \{P_k = 0\}$. The complete algebra therefore splits into the full factors supported at these tuples. This splitting can equivalently be obtained by lifting the idempotents of its special fiber. The all-zero tuples define the core factor, and the tuples with no zero coordinate define the attached factor. These conditions are G_k-invariant. The other tuples give the remaining direct factor. Because a finite group in characteristic zero has central isotypic projectors, every isotypic part of a finite free equivariant $\mathbb{C}[[T]]$-module is free with the same rank in its special and generic fibers. This proves the character assertion. □

4. Cut degenerations and chamber algebras

Gale duality for complete intersections uses products of linear forms [2]. Positive-valuation factors obtained by completion and parameter inversion also occur in the toric construction of González–Woodward [14].

Let $Q = (V, E)$ be a finite connected directed graph without oriented cycles; multiple arrows are permitted. Its incidence map is

$$\partial : \mathbb{Z}^E \longrightarrow \mathbb{Z}_0^V, \qquad e \longmapsto e_{\mathrm{head}(e)} - e_{\mathrm{tail}(e)},$$

where $\mathbb{Z}_0^V$ consists of the vectors with coordinate sum zero. The potential lattice and its quotient are

$$L = \left\{ (d_{\mathrm{head}(e)} - d_{\mathrm{tail}(e)})_{e \in E} : d \in \mathbb{Z}^V \right\}, \qquad N = \mathbb{Z}^E / L. \tag{4.1}$$

The coboundary $\partial^* : \mathbb{Z}^V/\mathbb{Z}\mathbf{1} \to \mathbb{Z}^E$ is dual to ∂ and has image L. A reduced incidence matrix of a spanning tree has determinant ± 1. It follows that L is saturated and N is a free lattice of rank $n = |E| - |V| + 1$.

Write ν_e for the image of e in N. Choose a basis $\eta_1, \dots, \eta_n$ of N^* and put

$$\ell_j(\mathbf{w}) = \sum_{e \in E} \langle \eta_j, \nu_e \rangle w_e. \tag{4.2}$$

The common zero space of these forms is $L_{\mathbb{C}}$. Equivalently, $w_e = h_{\mathrm{head}(e)} - h_{\mathrm{tail}(e)}$ for vertex coordinates h, determined up to a common constant. The ideal generated by the ℓ_j is the ideal of $L_{\mathbb{C}}$ and is preserved by graph automorphisms.

Fix a rational vector $\theta \in \mathbb{Q}_0^V$ such that $\theta(S) := \sum_{v \in S} \theta_v \neq 0$ for every nonempty proper subset $S \subset V$. Assume there is a strictly positive rational arrow weight ϖ with $\partial \varpi = \theta$. For a subset S, let $\mathrm{Out}(S)$ be the set of arrows leaving it and define

$$E_S(\mathbf{w}) = \prod_{e \in \mathrm{Out}(S)} w_e, \qquad J_\theta = (E_S : \theta(S) < 0).$$

4.1. **The exact initial ideal.** Let $I_L \subset \mathbb{C}[\mathbf{w}]$ be the lattice ideal generated by $\mathbf{w}^{l+} - \mathbf{w}^{l-}$ for $l \in L$, with disjoint positive and negative supports. It is also the kernel of $\mathbb{C}[\mathbf{w}] \to \mathbb{C}[N]$, $w_e \mapsto X^{\nu_e}$. Use the maximum-weight convention for initial terms.

Lemma 4.1. *One has* $\mathrm{in}_\varpi(I_L) = J_\theta$.

Proof. For the potential $d = 1_S$, the positive part of its coboundary is the incoming cut and the negative part is the outgoing cut. If $\theta(S) < 0$, the outgoing monomial has greater weight; hence $J_\theta \subseteq \mathrm{in}_\varpi(I_L)$.

For the opposite inclusion, shift an integral potential by a constant and decompose it into its nonempty proper threshold sets:

$$d = \sum_{h=\min d+1}^{\max d} 1_{S_h}, \qquad S_h = \{v : d_v \geq h\}.$$

The decomposition of $l = \partial^* d$ is conformal: on each arrow all nonzero summands have the same sign. If $\theta \cdot d > 0$, some S_h has positive weight, and its incoming cut, which is the outgoing cut of its negative-weight complement, divides $\mathbf{w}^{l+}$. The case $\theta \cdot d < 0$ follows by changing d to $-d$. If $\theta \cdot d = 0$ and $l \neq 0$, the nonzero threshold weights include both signs. Thus both monomials of the tied binomial belong to J_θ.

Finally, group a polynomial in I_L by its N-degree. The sum of its coefficients in each group is zero. Each maximum-weight monomial in such a nonzero group has another monomial of no greater weight in the same group. Their quotient difference is a monomial multiple of a lattice binomial; the strict or tied argument places the chosen leading monomial in J_θ. The assertion follows for every polynomial in I_L. □

Choose $e \geq 1$ such that $e\varpi$ and $e\theta$ are integral, and write $T = t^e$. For $l = \partial^* d$ with $\theta \cdot d \geq 0$, form

$$\mathbf{w}^{l+} - t^{e\theta \cdot d} \mathbf{w}^{l-}. \tag{4.3}$$

Reverse the two terms when necessary, and retain a tied binomial unchanged. The ideal of these homogenizations, denoted by I_θ^h, is the weight-homogenization of I_L. Indeed, grouping by N-degree and subtracting a minimum-weight monomial in each group expresses a general homogenization as a polynomial combination of the binomial ones, with nonnegative powers of t. The resulting Rees algebra is flat over $\mathbb{C}[t]$ and has special fiber $\mathbb{C}[\mathbf{w}]/J_\theta$. The standard monomials of a refinement of the weight order form a $\mathbb{C}[t]$-basis.

4.2. **A finite basis for the cut quotient.** The flow polytope

$$\mathcal{P}_\theta = \{u \in \mathbb{R}^E : u_e \geq 0, \partial u = \theta\}$$

has dimension n, since it contains the strictly positive point ϖ. It is bounded: a recession direction would be a nonnegative circulation, whereas successive removal of sources in an acyclic graph shows that such a circulation is zero. At a vertex, the positive support has no undirected cycle, since a signed cycle would allow a small perturbation in both directions. It is connected, since a component S would otherwise have $\theta(S) = 0$. Its positive support is therefore a spanning tree T_0.

Restriction gives an integral isomorphism

$$\ker \partial \longrightarrow \mathbb{Z}^{E \setminus T_0}.$$

One prescribes the complementary coordinates and solves for the tree coordinates using its unimodular reduced incidence matrix. Locally at the vertex, all tree coordinates remain strictly positive. The local feasible cone is consequently the positive orthant in the coordinates $E \setminus T_0$. These zero coordinates define distinct genuine facets, so $\mathcal{P}_\theta$ is simple. A coordinate which can vanish defines a facet, because its zero face contains a vertex. A remaining coordinate never vanishes and is treated as a ghost vertex of the dual complex.

The nonface ideal of this complex is exactly J_θ. If a prescribed zero-arrow set I is infeasible, rational Farkas separation gives a potential d with

$$d_{\mathrm{head}(e)} - d_{\mathrm{tail}(e)} \geq 0 \quad (e \notin I), \qquad \theta \cdot d < 0.$$

After a shift and multiplication by a positive integer, decompose d into threshold indicators. An allowed arrow cannot leave a threshold set. Since the sum of their θ-weights is negative, one threshold S satisfies $\theta(S) < 0$ and $\mathrm{Out}(S) \subseteq I$. Conversely, if all arrows leaving such an S vanish, then a feasible flow would satisfy $\theta(S) = \mathrm{inflow}(S) \geq 0$, a contradiction.

For $n > 0$, translate the polytope by an interior point before taking its relative polar. Its dual simplicial boundary complex, with the ghost variables killed, is thus the complex of J_θ. Its facets are the sets $E \setminus T_0$ for feasible spanning trees. A convex polytope boundary admits a shelling [3]. Let $F_1, \ldots, F_t$ be the shelling facets and $R_j \subseteq F_j$ their restriction faces. The new faces at step j are exactly the interval $[R_j, F_j]$. Write $\mathbf{w}_R = \prod_{e \in R} w_e$ for $R \subseteq E$.

On each facet the coefficient matrix of (4.2) is invertible, by the tree-complement isomorphism: under $N^* = \ker \partial$, this matrix restricts a basis of circulations to the edges outside T_0, so its determinant is ± 1. Set $A_0 = \mathbb{C}[\ell_1, \ldots, \ell_n]$. This is a polynomial subring: restriction to any facet makes its generators independent linear coordinates. Let Δ_j be the complex generated by the first j facets, and let $\mathbb{C}[\Delta_j]$ denote its face ring. The kernel of

$$\mathbb{C}[\Delta_j] \longrightarrow \mathbb{C}[\Delta_{j-1}]$$

is $\mathbf{w}_{R_j}\mathbb{C}[w_e : e \in F_j]$. Indeed, these are exactly its new monomials; every variable outside F_j annihilates this kernel, since a face containing R_j could belong neither to an older facet nor outside the new facet. The kernel is a free rank-one A_0-module generated by $\mathbf{w}_{R_j}$. Beginning with the first facet and lifting the preceding monomial basis at each step gives

$$\mathbb{C}[\mathbf{w}]/J_\theta = \bigoplus_{j=1}^{t} A_0\, \mathbf{w}_{R_j} \quad \text{as an } A_0\text{-module.} \tag{4.4}$$

Independence follows by reducing a relation to the preceding partial complex and then using the free rank-one kernel. In particular the ℓ_j form a regular sequence, and

$$H_\theta = \mathbb{C}[\mathbf{w}]/(J_\theta, \ell_1, \ldots, \ell_n) \tag{4.5}$$

is finite dimensional. If $n = 0$, the graph is a tree and the feasible polytope is a point. Deleting any edge gives a negative-weight cut whose sole outgoing arrow is that edge. Thus every w_e lies in J_θ, and $H_\theta = \mathbb{C}$.

4.3. **The complete positive factor.**

Proposition 4.2. *Let a finite group G act by automorphisms of Q preserving θ. For a sufficiently divisible ramification $T = t^e$, there is a finite free equivariant $\mathbb{C}[[t]]$-algebra with special fiber H_θ whose generic fiber is exactly the full factor of the Gale system*

$$\ell_j(\mathbf{w}) = 0, \qquad \mathbf{w}^{\partial^* d} = t^{e\theta\cdot d} \quad (d \in \mathbb{Z}^V/\mathbb{Z}\mathbf{1}) \tag{4.6}$$

on which all arrow coordinates have positive valuation. Its special and generic fibers have the same ungraded G-character.

Proof. Average ϖ over G and enlarge e to clear its denominators. This preserves positivity and the incidence sum, so the Rees construction is equivariant. The Rees ring is Noetherian and flat over $\mathbb{C}[t]$. Complete it at the maximal ideal $(t, \mathbf{w})$ and then divide by the ℓ_j; call the result S_θ. The completion is Noetherian and local, and flatness of completion preserves the regularity of t. Its quotient modulo t is the completion of $\mathbb{C}[\mathbf{w}]/J_\theta$ at the arrow ideal. Equation (4.4) and flatness of this completion show that $\ell_1, \dots, \ell_n$ remains regular there. Thus $(t, \ell_1, \dots, \ell_n)$ is a regular sequence, all of whose elements lie in the maximal ideal. Such a sequence in a Noetherian local ring can be permuted, so t remains a nonzerodivisor after the linear section [26].

The arrow ideal of $S_\theta/tS_\theta = H_\theta$ is nilpotent. Thus the $(t, \mathbf{w})$-adic and t-adic topologies on S_θ are cofinal. Lift the basis in (4.4). Successive reduction modulo t expresses every element of S_θ as a $\mathbb{C}[[t]]$-linear combination of these lifts. A relation among them has all its coefficients divisible by t; cancel t and repeat. Separatedness proves independence. Hence S_θ is finite free.

Put

$$A' = \mathbb{C}[[t]][\mathbf{w}]/(I_\theta^h, \ell_1, \dots, \ell_n).$$

Its $(t, \mathbf{w})$-adic completion is S_θ. The completion map is flat and is surjective, because the finite basis lifts and all their $\mathbb{C}[[t]]$-coefficients already lie in A'. Its finitely generated kernel J satisfies $J/J^2 = 0$ and is generated by an idempotent. Consequently

$$A' \simeq S_\theta \times A'_{\text{other}}. \tag{4.7}$$

After inverting t, this is a direct-factor decomposition of the polynomial Gale algebra. If $E \neq \varnothing$, acyclicity supplies an integral potential strictly increasing on every arrow. Its lattice relation sets a product involving every arrow coordinate equal to a nonzero power of t, so all arrow coordinates are invertible. If $E = \varnothing$, the graph has one vertex and both generic algebras are $\mathbb{C}((t))$. Thus the generic relations are precisely (4.6).

On S_θ each arrow is topologically nilpotent, so every residue point of its generic fiber has positive arrow valuations. Conversely, take a full Artin local factor of the generic system with this property, after a finite complete valued-field extension. Its arrow multiplication operators are integral over the valuation ring: each is a positive-valuation scalar plus a nilpotent operator. The ring generated by these operators is finite over the valuation ring, and they are topologically nilpotent. The polynomial homomorphism is therefore continuous for $(t, \mathbf{w})$ and factors through the completion. The idempotent decomposition therefore contains this entire Artin factor. For any $\overline{\mathbb{C}((t))}$-point with positive arrow valuations, evaluation in a finite complete field of definition is continuous and therefore factors through S_θ. Since $S_\theta[t^{-1}]$ is finite and is a direct factor of the generic polynomial algebra, this point is isolated. This characterizes the selected factor exactly. Finally, apply the isotypic projectors to the finite free family to obtain the character equality. □

4.4. **The two neighboring chambers.** Apply the construction to the graph with vertices $A_1, \dots, A_{ka}$, $B_1, \dots, B_{kb}$, and ∞, and arrows

$$B_j \longrightarrow A_i \ (m \text{ copies}), \qquad \infty \longrightarrow A_i \ (p \text{ copies}), \qquad B_j \longrightarrow \infty \ (q \text{ copies}).$$

Set the potential at ∞ equal to zero. The arrow coordinates are then $a_i - b_j$, a_i, and $-b_j$. Define

$$\theta_0 = (b \text{ on } A_i,\ -a \text{ on } B_j,\ 0 \text{ at } \infty), \quad \kappa = (N_A \text{ on } A_i,\ -N_B \text{ on } B_j,\ -k \text{ at } \infty).$$

Let $\theta_\delta = \theta_0 - \delta\kappa$ and $\delta_\pm = \pm(2M_k+1)^{-1}$. A positive lift of θ_0 has weights $1/(2mk)$ on the mixed arrows, $b/(2p)$ on the arrows from the frame, and $a/(2q)$ on the arrows to the frame. The bounds $M_k \geq mk$, $M_k \geq p/b$, and $M_k \geq q/a$ show that subtracting $\delta_\pm$ from every weight keeps them positive and changes their incidence sum to $\theta_{\delta_\pm}$.

These two characters are generic. It suffices to consider a subset not containing ∞, with particle counts (u, v). Its weight is $\Delta - \delta D_0$. If $\Delta \neq 0$, then $|\delta D_0| < 1/2$, so its sign is that of Δ. If $\Delta = 0$, then $D_0 = \ell > 0$ and its weight is nonzero. Complements handle the subsets containing the frame. Write $H_- = H_{\theta_{\delta_-}}$ and $H_+ = H_{\theta_{\delta_+}}$.

Corollary 4.3. *There are isomorphisms of ungraded G_k-representations*

$$\mathcal{B}_k^{p,q} \simeq H_-, \qquad \mathcal{C}_k^{p,q} \simeq H_+. \tag{4.8}$$

Proof. Choose a common ramification $T = t^e$ clearing the rational exponents for both chambers. The central valuation then has $\mathrm{val}(t) = 1/e$. Let $\mathbf{w}^{(0)}$ be the arrow coordinates of the θ_0 system, and set $w_\epsilon^{(\delta)} = T^{-\delta} w_\epsilon^{(0)} = t^{-e\delta} w_\epsilon^{(0)}$ for $\epsilon \in E$. For $l = \partial^* d$, the equality $\sum_{\epsilon\in E} l_\epsilon = \kappa \cdot d$ gives

$$(\mathbf{w}^{(\delta)})^l = T^{-\delta\kappa\cdot d}(\mathbf{w}^{(0)})^l = T^{(\theta_0 - \delta\kappa)\cdot d} = t^{e\theta_\delta \cdot d}.$$

The homogeneous linear relations are preserved, so this scaling is an isomorphism of the generic polynomial Gale algebras for θ_0 and θ_δ. Proposition 4.2 selects the factors whose arrow valuations in the θ_0 system exceed δ. For δ_- it selects every solution, because every particle, hence every arrow, has nonnegative valuation. For δ_+ it selects precisely the core: the frame arrows require every particle to have positive valuation, and then (3.6) makes every arrow valuation greater than δ_+. The reverse implication follows again from the frame arrows. The finite free models in Propositions 3.1 and 4.2 transport these generic characters to their special fibers, giving (4.8). □

5. Separating the core and attached factors

A square deformation separates the core and attached equations.

5.1. **A completion lemma over a one-dimensional generic base.** Let K be a characteristic-zero field, and let $F_1, \ldots, F_N, h \in K[s, T, z_1, \ldots, z_N]$. Set

$$A_0 = K[s, T, \mathbf{z}, h^{-1}]/((F_1, \ldots, F_N) : T^\infty).$$

Suppose A_0/TA_0 is finite over $K[s]$. Put

$$R = K[s][[T]], \qquad S = R[T^{-1}], \qquad A' = A_0 \otimes_{K[s,T]} R, \qquad B = \varprojlim_j A_0/T^j A_0.$$

The ring here is $K[s][[T]]$, in which the degrees in s of successive coefficients may be unbounded.

Lemma 5.1. *The algebra B is finite over R, and $A' \to B$ is a surjective flat map selecting a direct factor. After inverting T, $B[T^{-1}]$ is finite locally free over the connected regular one-dimensional ring S. If a finite group acts on the equations, preserves the localization defining A_0, and fixes the parameters, its characters on the fibers of $B[T^{-1}]$ at $s = 0$ and $s = 1$ agree.*

Proof. The ring R is the T-adic completion of $K[s, T]$, and $A'/T^jA' \simeq A_0/T^jA_0$ for every $j \geq 1$. Thus B is the T-adic completion of the Noetherian algebra A'. Lift finitely many $K[s]$-module generators of A_0/TA_0. Successive reduction modulo T expresses every element of B as an R-linear combination of these lifts. They lie in A', so the completion map $A' \to B$ is surjective. It is flat by Noetherian completion. As in (4.7), its kernel is a finitely generated idempotent ideal.

T belongs to the Jacobson radical of R, since $1 - Tf$ has a power-series inverse for every $f \in R$. Every maximal ideal is therefore $(T, g(s))$ for an irreducible $g \in K[s]$. It is generated by two elements and contains the prime chain $(0) \subset (T) \subset (T, g(s))$, so its local ring is regular of dimension two. Thus R is regular, and every prime avoiding T has height at most one. The localization S is a regular integral domain of dimension one (the prime (s) survives), and is connected. On inverting T, the saturation disappears. Thus $B[T^{-1}]$ is a direct factor of the coordinate algebra of the original square system,

$$S[\mathbf{z}, h^{-1}]/(F_1, \dots, F_N).$$

Let $\mathfrak{q}$ be a maximal ideal of $B[T^{-1}]$ and put $\mathfrak{p} = \mathfrak{q} \cap S$. Finiteness makes $\mathfrak{p}$ maximal and the residue extension finite. If $\mathfrak{P}$ is the corresponding prime of $S[\mathbf{z}, h^{-1}]$, then

$$D = (S[\mathbf{z}, h^{-1}])_{\mathfrak{P}}$$

is a regular local ring of dimension $N + 1$: the base $S_{\mathfrak{p}}$ is a DVR, and the closed fiber is an N-variable polynomial ring localized at a maximal ideal with finite residue extension. Localizing the direct factor identifies its local ring with $D/(F_1, \dots, F_N)$. Let π be a uniformizer of $S_{\mathfrak{p}}$. Finiteness over S makes $D/(F_1, \dots, F_N, \pi)$ an Artinian local ring; it is nonzero because all these elements belong to the maximal ideal of D. Thus $F_1, \dots, F_N, \pi$ is a system of parameters in the regular local ring D, hence a regular sequence. In particular, π is a nonzerodivisor on $D/(F_1, \dots, F_N)$. Applying this at every maximal point proves torsion-freeness over S. A finite torsion-free module over each DVR $S_{\mathfrak{p}}$ is free, proving local freeness.

Choose a finite splitting-field extension K'/K. Its base change is $K'[s][[T]][T^{-1}]$, which is a domain. The finite-group isotypic projectors give finite projective modules of constant rank on this connected base. Coefficientwise evaluation $s = c$ sends S to $K((T))$. Thus the ranks, and hence the characters at $c = 0, 1$, are equal. The equality descends from the splitting field. □

5.2. **The coupled equations and their selected factor.** Fix $r, l \geq 0$ with $r + l = k$. Give $r(a, b)$ particles the core labels x_i, y_j and the remaining $l(a, b)$ particles the attached labels u_α, v_β. The group preserving these labels is $G_r \times G_l$. Introduce s and consider

$$x_i^p \prod_j (x_i - y_j)^m \prod_\beta (sx_i - v_\beta)^m = T^b, \tag{5.1}$$

$$(-y_j)^q \prod_i (x_i - y_j)^m \prod_\alpha (u_\alpha - sy_j)^m = T^a, \tag{5.2}$$

$$u_\alpha^p \prod_\beta (u_\alpha - v_\beta)^m \prod_j (u_\alpha - sy_j)^m = T^b, \tag{5.3}$$

$$(-v_\beta)^q \prod_\alpha (u_\alpha - v_\beta)^m \prod_i (sx_i - v_\beta)^m = T^a. \tag{5.4}$$

Each mixed factor is changed in both of its endpoint equations. At $s = 1$ this is the original system. At $s = 0$ the attached system has parameters

$$(p', q') = (p + mrb, q + mra), \qquad p'a - q'b = 1. \tag{5.5}$$

Its nonzero-center sign is unchanged:

$$(-1)^{bq' + mab(l-1)} = (-1)^{bq + mab(k-1)} = \sigma_k.$$

In particular, the attached sector at the decoupled endpoint is the intrinsic attached sector for this shifted framing.

Let I be the equation ideal and localize its T-saturation at

$$h = \left(\prod_{\alpha} u_\alpha \prod_{\beta} v_\beta\right)\left(\prod_{i,\beta}(sx_i - v_\beta)\prod_{\alpha,j}(u_\alpha - sy_j)\right)\prod_i P_k(x_i)\prod_j P_k(y_j). \tag{5.6}$$

An empty product is one. The closed points of the special support of this localized algebra satisfy

$$x_i = y_j = 0, \qquad u_\alpha, v_\beta \in \{P_k = 0\}. \tag{5.7}$$

Indeed, T is a nonzerodivisor after saturation, so every closed special point lies in the closure of the locus $T \neq 0$. Taking a curve through that point in a horizontal irreducible component and normalizing produces a DVR specialization. All affine coordinates and s are integral, and h is a unit. Partition the nonzero residues separately in the core and attached families. For any one cluster, multiply its first-color equations and divide by its second-color equations. Internal differences cancel and all other factors are units, including the mixed factors selected by h. The cluster is therefore balanced. All attached coordinates are nonzero, and the remaining zero core is also balanced.

If a core cluster of charge $i(a,b)$ has center c and an attached cluster of charge $j(a,b)$ has center d, their combined contribution to the relevant quotient is

$$\frac{(sc-d)^{mabij}}{(d-sc)^{mabij}} = (-1)^{mabij}.$$

Its reciprocal is the same sign. Interactions within one family give the identical sign. Thus the calculation of (3.7) applies unchanged: every nonzero center lies in $\{P_k = 0\}$, independently of s. The factors $P_k(x_i), P_k(y_j)$ in (5.6) exclude such centers from the core, proving (5.7).

The special algebra is finitely generated over $\mathbb{C}$. By the Nullstellensatz, (5.7) holds modulo its nilradical. On this reduced support, the denominator h is independent of s and takes only finitely many nonzero constant values: the mixed factors reduce to $-v_\beta$ and u_α, and the core factors to $P_k(0)$. Thus every particle coordinate and h^{-1} satisfies a monic polynomial with constant coefficients modulo the nilradical, and is integral there over $\mathbb{C}[s]$. The special algebra is Noetherian, so its nilradical is nilpotent. Raising a monic equation modulo that nilradical to a sufficiently high power proves integrality in the full algebra. Since it has finitely many generators, it is finite over $\mathbb{C}[s]$.

Lemma 5.1 now applies. After inverting T, its fibers at $s = 0, 1$ are the full sectors with positive-valued core and zero-valued attached coordinates. Indeed, saturation disappears after T is inverted. At either endpoint, take a full Artin factor with these valuations, after a finite complete valued-field extension. Every particle multiplication operator is a scalar of nonnegative valuation plus a nilpotent operator, and hence satisfies a power of a monic linear polynomial over the valuation ring. The denominators in (5.6) have unit residues, so their inverse operators are integral as well. The commutative operator algebra they generate over the valuation ring is finite and complete. The homomorphism from the localized saturated algebra is therefore T-adically continuous and extends to B, placing the entire Artin factor in its idempotent summand. Conversely, finiteness of B over R makes every generic endpoint residue point integral over the valuation ring; (5.7) forces exactly the stated core and attached valuations. Exactness of Noetherian completion on finite modules identifies $B/(s-c)B$ with the completion of $A_0/(s-c)A_0$. Thus the endpoint fibers select precisely these full local factors, and Lemma 5.1 gives equality of their characters.

5.3. **Independent unit parameters and the tensor product.** At the decoupled endpoint, the core right sides are $u_A T^b$ and $u_B T^a$, where

$$u_A = V^{-m}, \qquad u_B = U^{-m}, \qquad U = \prod_\alpha u_\alpha, \quad V = \prod_\beta (-v_\beta). \tag{5.8}$$

These are two G_l-invariant units in the full attached algebra. Replace them first by independent unit parameters.

For the general graph above and an integer $h \geq 1$, let $R_0 = \mathbb{C}[u_1^{\pm 1}, \dots, u_h^{\pm 1}]$ and let $\chi : L \to R_0^*$ be a lattice character with Laurent monomial values. Assume G fixes the unit parameters and preserves χ. Replace (4.3) by

$$\mathbf{w}^{l+} - \chi(l) t^{e\theta \cdot d} \mathbf{w}^{l-}.$$

Since the lattice $N = \mathbb{Z}^E/L$ is free, χ extends to a character of $\mathbb{Z}^E$. A diagonal change by these arrow units identifies the twisted lattice/Rees ideal with the scalar extension of the untwisted one. Its initial ideal is still $J_\theta R_0[\mathbf{w}]$. Leave the residual linear forms in the original $\mathbf{w}$ coordinates. The special section is exactly $H_\theta \otimes R_0$.

Let $I^h_{\theta,\chi}$ denote the twisted Rees ideal. Complete $R_0[t, \mathbf{w}]/I^h_{\theta,\chi}$ at $(t, \mathbf{w})$ and then divide by the original linear forms; denote the result by $S_{\theta,\chi}$. The completion is Noetherian, and the image of $(t, \mathbf{w})$ lies in its Jacobson radical. Untwisting and flatness of completion show that t is regular before the section. Modulo t, the ring is the arrow-adic completion of $R_0[\mathbf{w}]/J_\theta$. Scalar extension of (4.4) and flatness of this completion make the ℓ_j regular there. Every maximal ideal contains t and the ℓ_j, so local permutability of regular sequences makes t regular after the section. Its quotient modulo t is $H_\theta \otimes_{\mathbb{C}} R_0$, whose arrow ideal is nilpotent. Thus the completion is also t-adic. Lift a $\mathbb{C}$-basis of H_θ; successive reduction modulo t proves generation over $R_0[[t]]$, and cancellation of t in a relation, followed by separatedness, proves independence. Averaging the lift of the basis map $H_\theta \to S_{\theta,\chi}$ over the finite group preserves its reduction modulo t. The same lifting argument therefore gives

$$S_{\theta,\chi} \simeq H_\theta \otimes_{\mathbb{C}} R_0[[t]] \quad \text{as equivariant modules.} \tag{5.9}$$

The completion map from $R_0[[t]][\mathbf{w}]/(I^h_{\theta,\chi}, \ell_1, \dots, \ell_n)$ is surjective, since the basis lifts and their coefficients already lie in this algebra. It is flat by Noetherian completion, so its finitely generated kernel is generated by an idempotent. Both this product decomposition and (5.9) survive arbitrary base change.

Choose a common ramification $T = t^e$ and put $\mathcal{O} = \mathbb{C}[[t]]$, $K = \mathbb{C}((t))$, and $R_0 = \mathbb{C}[u_A^{\pm 1}, u_B^{\pm 1}]$. For a core of charge $r(a, b)$ with $r \geq 1$, use its positive chamber. The relevant character is

$$\chi(\partial^* d) = u_A^{\sum_i (d_{A_i} - d_\infty)} u_B^{-\sum_j (d_{B_j} - d_\infty)}.$$

The inverse exponent at the second-color vertices comes from the head-minus-tail incidence convention. At a valued specialization where u_A, u_B are valuation-ring units, their valuations are zero, so the cluster quotient and the numerical gap (3.6) remain valid. The scaling in Corollary 4.3 then identifies the generic positive factor with the generic core of this two-unit system. The full-factor argument in Proposition 4.2 applies to these unit specializations as well. Indeed, specialize the parameters in the units of a finite complete valuation-ring extension. In a full local factor with positive arrow valuations, the arrow multiplication operators generate a finite algebra over that valuation ring and are topologically nilpotent. Evaluation therefore extends continuously to the universal completion. Conversely, the arrows in this completion remain topologically nilpotent after base change. Its specialized idempotent consequently selects exactly the full positive local factors. Denote the universal positive completion by $\mathcal{U}$; for $r = 0$,

put $\mathcal{U} = R_0[[t]]$. Equation (5.9) and Corollary 4.3 give

$$\mathcal{U} \simeq \mathcal{C}_r^{p,q} \otimes_{\mathbb{C}} R_0[[t]] \qquad \text{as equivariant modules.}$$

Let $\mathcal{D}$ be the finite free attached model for the shifted framing (5.5), and set $D_K = \mathcal{D}[t^{-1}]$. For $l = 0$, take $\mathcal{D} = \mathcal{O}$, so $U = V = 1$. Its particle coordinates have nonzero residues, so they and the products U, V are units in $\mathcal{D}$. Completeness of this finite $\mathcal{O}$-algebra makes the substitution (5.8) a continuous homomorphism $R_0[[t]] \to \mathcal{D}$. Form

$$\mathcal{E} = \mathcal{U} \otimes_{R_0[[t]]} \mathcal{D}.$$

This is finite free over $\mathcal{O}$, with

$$\mathcal{E} \simeq \mathcal{C}_r^{p,q} \otimes_{\mathbb{C}} \mathcal{D} \quad \text{as } (G_r \times G_l)\text{-equivariant } \mathcal{O}\text{-modules.} \tag{5.10}$$

The substituted units are G_l-invariant, so G_l fixes the universal core basis and the action in (5.10) is the tensor-product action.

Now invert t. Let A_K be the coordinate algebra over D_K of the core equations (5.1)–(5.2) at $s = 0$. The decoupled sector selected above is a direct factor of A_K. The universal product decomposition commutes with base change to D_K; the central-to-positive-chamber scaling therefore makes $\mathcal{E}_K = \mathcal{E}[t^{-1}]$ another direct factor of A_K. The algebra A_K is finite over D_K: its right sides are units, so it is a base change of the finite moment map over its target torus. Let $\mathfrak{n}$ be the nilradical of D_K. The quotient $D_K/\mathfrak{n}$ is a product of finite extensions of K, and the images of u_A, u_B have valuation zero in each field because they come from units of $\mathcal{D}$: each unit and its inverse are integral over $\mathcal{O}$, so their nonnegative valuations sum to zero. Over each such field, the unit-parameter full-factor argument above and the valuation gap identify the base change of $\mathcal{E}_K$ with the full central core factor of the decoupled equations. Thus the idempotents defining these two factors agree modulo $\mathfrak{n}A_K$. That ideal is nilpotent, and idempotents in a commutative ring lift uniquely across nilpotent ideals. Indeed, if two idempotents e, f agree modulo a nilpotent ideal, then $e(1-f)$ and $f(1-e)$ are nilpotent idempotents. They vanish, giving $e = f$. The two factors therefore agree over D_K, including all attached and core nilpotents.

The finite free equivariant model $\mathcal{D}$ has the same G_l-character as its special fiber $\mathcal{D}_l^{p+mrb,q+mra}$. Applying isotypic projectors to (5.10) gives, for $(g_r, g_l) \in G_r \times G_l$,

$$\operatorname{Tr}((g_r, g_l) \mid \mathcal{E}_K) = \operatorname{Tr}(g_r \mid \mathcal{C}_r^{p,q}) \operatorname{Tr}(g_l \mid \mathcal{D}_l^{p+mrb,q+mra}).$$

Lemma 5.1 transports this character to the coupled sector at $s = 1$. That sector is a direct factor of the finite free horizontal model in Proposition 3.1, so its generic and special characters agree.

Theorem 5.2. *For every $k \geq 0$ and every positive pair satisfying (3.1), there is an isomorphism of ungraded G_k-representations*

$$\mathcal{B}_k^{p,q} \simeq \bigoplus_{r+l=k} \operatorname{Ind}_{G_r \times G_l}^{G_k} \left(\mathcal{C}_r^{p,q} \otimes \mathcal{D}_l^{p+mrb,q+mra} \right). \tag{5.11}$$

Proof. For $k = 0$, both sides are the trivial representation. Assume $k \geq 1$. In an original solution every coordinate has nonnegative valuation, and every nonzero center cluster has balanced charge. Hence the labels with zero residue, namely the positive-valuation core, have counts $r(a, b)$ for some r, with the remaining labels of counts $l(a, b)$. The resulting full algebra factors are indexed by these subsets of the two label sets. For one fixed subset their character is unchanged by (5.1)–(5.4), and the decoupled calculation gives the tensor product above. The stabilizer of the chosen subsets is $G_r \times G_l$. Summing over their G_k-orbit is induction, and summing over r exhausts every full local factor. Equality of characters implies the stated representation isomorphism in characteristic zero. □

6. Finite zero schemes and Young characters

All representations in this section are ungraded. For a partition $\lambda = (\lambda_1, \dots, \lambda_r)$ of n, write $S_\lambda = S_{\lambda_1} \times \cdots \times S_{\lambda_r} \subseteq S_n$. We call $\operatorname{Ind}_{S_\lambda}^{S_n} \mathbf{1}$ a Young permutation module. For a product of symmetric groups, the same terminology means the external tensor product of such modules.

Lemma 6.1. *Let K be an algebraically closed field of characteristic zero, let*

$$V = \bigoplus_{c=1}^{t} K^{n_c}, \qquad G = \prod_{c=1}^{t} S_{n_c},$$

and let G act by permuting the coordinates in each summand. Suppose that $F : V \longrightarrow V$ is a G-equivariant polynomial map. If Z is a G-stable finite union of isolated local factors of $F^{-1}(0)$, with their full Artin scheme structures, then

$$\mathcal{O}(Z) \cong \bigoplus_{\lambda^{(c)} \vdash n_c} m_{\lambda^{(1)},\dots,\lambda^{(t)}} \operatorname{Ind}_{\prod_c S_{\lambda^{(c)}}}^{G} \mathbf{1}, \qquad m_{\lambda^{(1)},\dots,\lambda^{(t)}} \in \mathbb{N}. \tag{6.1}$$

Proof. Put $N = \sum_c n_c$. If $N = 0$, the selected scheme is empty or a point, giving the zero or trivial representation. Assume $N > 0$. Introduce N parameters $t_{c,j}$ and the polynomials

$$P_c(z) = \sum_{j=0}^{n_c - 1} t_{c,j} z^j.$$

The equations

$$F_{c,i}(x) + P_c(x_{c,i}) = 0 \qquad (1 \le c \le t,\ 1 \le i \le n_c) \tag{6.2}$$

are equivariant, with trivial G-action on the parameter space $\mathcal{T} = \mathbb{A}_K^N$. We prove first that a nonempty Zariski open subset of $\mathcal{T}$ gives finite reduced zero schemes.

Stratify V by the equality partitions of the coordinates within each color. On one such stratum, the distinct values in color c are $z_{c,1}, \dots, z_{c,d_c}$, with multiplicities $m_{c,1}, \dots, m_{c,d_c}$. Equalities between different colors do not change this partition. The stratum is smooth of dimension $d = \sum_c d_c$, and the stabilizer of each of its points is $H = \prod_{c,l} S_{m_{c,l}}$.

Keep one equation from (6.2) for each equality block. The other equations in that block agree by equivariance. Evaluation of a polynomial of degree at most $n_c - 1$ at the d_c distinct values is surjective. Thus the derivative of these d equations in the parameter directions is surjective everywhere. Their incidence variety in the stratum times $\mathcal{T}$ is smooth of dimension N. On each component dominating $\mathcal{T}$, its projection is generically finite and generically étale. Its non-étale locus has dimension at most $N - 1$. A nondominating component also has image closure of dimension at most $N - 1$.

Let R be the Jacobian of the block equations in the d block-constant directions. The tangent equation of the incidence is $R\,\delta z + E_z\,\delta t = 0$, where E_z is the surjective evaluation map. The projection differential is invertible exactly when R is invertible. Consequently the parameters admitting a solution on this stratum with singular R are contained in a proper closed subset of $\mathcal{T}$.

At each point of the stratum, the full Jacobian is an endomorphism of the H-module V. This module decomposes into the block-constant subspace and one sum-zero standard module for each repeated equality block. The standard modules belong to distinct factors of H and cannot mix. On the standard module of a block of size $m_{c,l} > 1$, the Jacobian is a scalar $\lambda_{c,l}$. Explicitly, within the block it is the difference of the common diagonal and common off-diagonal matrix entries; matrix blocks joining distinct equality blocks have constant entries and kill sum-zero vectors. Hence

$$\det \operatorname{Jac}(F + P) = \det R \prod_{c,l : m_{c,l} > 1} \lambda_{c,l}^{m_{c,l} - 1}. \tag{6.3}$$

If r_c blocks of color c are repeated, then $d_c + r_c \le n_c$. Hermite interpolation therefore prescribes independently the values of P_c at every block and its first derivative at every repeated block. Indeed the corresponding kernel consists of polynomials divisible by

$$\prod_{m_{c,l}=1} (z - z_{c,l}) \prod_{m_{c,l}>1} (z - z_{c,l})^2,$$

of degree $d_c + r_c$. Variation of the derivative changes $\lambda_{c,l}$ while leaving the block equations fixed. Inside the incidence variety, each condition $\lambda_{c,l} = 0$ is therefore a divisor, if nonempty. Its image closure has dimension at most $N - 1$.

There are finitely many equality partitions. Removing all the bad image closures just described leaves a nonempty open subset $\mathcal{T}^\circ$. At a parameter in this subset, every zero has invertible full Jacobian by (6.3). Its affine zero scheme is consequently reduced and zero-dimensional, hence finite.

Choose a parameter vector whose line through the origin is not contained in $\mathcal{T} \setminus \mathcal{T}^\circ$, and write the resulting perturbation as

$$F^s_{c,i}(x) = F_{c,i}(x) + s\, p_c(x_{c,i}).$$

Over $K((s))$ its complete zero scheme is finite and geometrically reduced. To isolate Z, take its idempotent in $K[x]/(F)$ and average a polynomial lift over G. The resulting invariant polynomial g satisfies $(K[x]/(F))[g^{-1}] = \mathcal{O}(Z)$. Set

$$A = K[[s]][x, g^{-1}]/(F^s), \qquad \widehat{A} = \varprojlim_j A/s^j A.$$

Its special fiber is exactly $\mathcal{O}(Z)$. Lifts of a finite basis of that fiber generate $\widehat{A}$ over $K[[s]]$: successively reduce an element modulo $s, s^2, \dots$ and collect the resulting coefficients as power series.

At each point x_0 of Z, the $N + 1$ elements $F^s_1, \dots, F^s_N, s$ form a parameter system in the regular local ring $K[[s, x - x_0]]$, because their common quotient is the original finite local algebra. They form a regular sequence. In particular s is a nonzerodivisor after quotienting by the N deformed equations. The s-adic completion is the product of these completed local factors: the inverse image of the radical of A/sA has a power contained in sA, and the Chinese remainder theorem identifies the corresponding cofinal completions. It follows that $\widehat{A}$ is finite torsion-free, hence finite free, over the DVR $K[[s]]$.

The completion map $A \to \widehat{A}$ is flat and is also surjective, since the chosen finite lifts and all their $K[[s]]$-coefficients already lie in A. Its kernel J is finitely generated and satisfies $J/J^2 = 0$. A finitely generated idempotent ideal is generated by an idempotent, so $\widehat{A}$ is a direct factor of A. After inverting s, it is therefore a reduced direct factor of the full deformed polynomial zero algebra.

The group G acts over $K[[s]]$. For every irreducible G-module E, the multiplicity module $\mathrm{Hom}_G(E, \widehat{A})$ is a direct summand of a finite free $K[[s]]$-module and is free. Its special and generic ranks agree. Over an algebraic closure of $K((s))$, the generic algebra is the function algebra of a finite G-set of coordinate tuples. The stabilizer of each tuple is the product of the symmetric groups on its equal-coordinate blocks. Decomposing this set into orbits and using equality of characters between the two fibers proves (6.1). □

Regular equivariant deformations and their permutation representations were studied in [24, Corollary 6.2.8 and Section 8.1]. The proof above gives polynomial perturbations with separate degree bounds for each color and retains the selected full local factors.

For an alphabet α, put $p_r(\alpha) = \sum_i \alpha_i^r$ for $r \ge 1$. The complete homogeneous symmetric function $h_j(\alpha)$ is the sum of all monomials of degree j, and the elementary symmetric function $e_j(\alpha)$ is the sum of the squarefree monomials of degree j, with $h_0 = e_0 = 1$. For a partition λ,

write

$$p_\lambda = \prod_i p_{\lambda_i}, \qquad h_\lambda = \prod_i h_{\lambda_i}, \qquad e_\lambda = \prod_i e_{\lambda_i}.$$

For a finite S_n-representation E, its Frobenius characteristic is

$$\operatorname{ch} E = \frac{1}{n!} \sum_{\sigma \in S_n} \operatorname{tr}(\sigma \mid E)\, p_{\lambda(\sigma)},$$

where $\lambda(\sigma)$ records the cycle lengths of σ. We use the analogous two-alphabet characteristic for $S_A \times S_B$. It sends a Young permutation module to $h_\lambda h_\mu$. Indeed

$$\sum_{n \geq 0} h_n z^n = \prod_i (1 - \alpha_i z)^{-1} = \exp\left(\sum_{r \geq 1} \frac{p_r z^r}{r} \right)$$

identifies h_n with the cycle-index average of the trivial S_n-character. For induced external tensor products, regrouping the character average by the preserved ordered packets makes the Frobenius characteristics multiply. These two observations give the asserted Young-module formula. Let $\iota = \omega_\alpha \omega_\beta$ be the involution in the two alphabets determined by $\omega(h_j) = e_j$. Equivalently, $\omega(p_j) = (-1)^{j-1} p_j$, so it is tensoring the representation with the two sign representations. For a product W of symmetric groups, let sgn_W be the product of their sign characters. For a W-representation E, write

$$E^{\operatorname{sgn}_W} = \{u \in E : wu = \operatorname{sgn}_W(w) u \text{ for every } w \in W\}$$

for its sign-isotypic subspace. Since sgn_W is one dimensional, $\dim E^{\operatorname{sgn}_W} = \dim \operatorname{Hom}_W(\operatorname{sgn}_W, E)$.

Apply Lemma 6.1 to the attached factor of the Gale equations over an algebraic closure of $\mathbb{C}((T))$. The source and the target have the same natural particle action of $G_k = S_{ka} \times S_{kb}$. Finite flatness in Proposition 3.1 then transfers its character to the special fiber. We obtain

$$\operatorname{ch} \mathcal{D}_k^{p,q} = \sum_{\lambda \vdash ka,\, \mu \vdash kb} n_{\lambda\mu}^{p,q}\, h_\lambda(\alpha) h_\mu(\beta), \qquad n_{\lambda\mu}^{p,q} \in \mathbb{N}. \tag{6.4}$$

7. A polynomial Harder–Narasimhan decomposition

The polynomial factorization below is the shuffle form of the Chow–Hall Harder–Narasimhan decomposition and destabilizing-product quotient of [13, Theorems 6.2 and 8.1]. Its separated-center proof is adapted from [21, Sections 3.1–3.2].

Let I be a finite set, and let $a_{ij} \in \mathbb{N}$ be an arrow matrix, allowing loops. For $d \in \mathbb{N}^I$ put

$$S_d = \mathbb{Q}[x_{i,\alpha} : i \in I,\ 1 \leq \alpha \leq d_i]^{\prod_i S_{d_i}}, \qquad \chi(d,e) = \sum_i d_i e_i - \sum_{i,j} a_{ij} d_i e_j.$$

Every variable has polynomial degree one. If $f \in S_d$ and $g \in S_e$, define

$$f * g = \sum_{\text{shuffles}} f(x) g(y) \prod_{i,j} \prod_{\alpha=1}^{d_i} \prod_{\beta=1}^{e_j} (y_{j,\beta} - x_{i,\alpha})^{a_{ij} - \delta_{ij}}. \tag{7.1}$$

There is one summand for each partition of the variables of every color into the two ordered packets of sizes d_i, e_i. Only equal-color collisions can create poles, and each possible pole is simple. Pair the summands assigning the two colliding variables to opposite packets. Their residues cancel. Clearing the product of the distinct collision factors shows that (7.1) is a polynomial. Expanding three ordered packets proves associativity. Its degree is

$$\deg(f * g) = \deg f + \deg g - \chi(d,e) \tag{7.2}$$

for homogeneous inputs whose product is nonzero.

For the zero dimension vector, $S_0 = \mathbb{Q}$, with its usual unit. Choose additive maps $c, \eta : \mathbb{Z}^I \to \mathbb{Q}$, with $c(e_i) > 0$ for every coordinate vector, and set $\mu(d) = \eta(d)/c(d)$ for $d \neq 0$. For $d \neq 0$, define

$$I_d^\mu = \sum_{\substack{u+v=d,\ u,v\neq 0 \\ \mu(u)>\mu(d)}} S_u * S_v, \qquad B_d^\mu = S_d/I_d^\mu. \tag{7.3}$$

The sum is an ordinary homogeneous ideal of S_d. Indeed a symmetric multiplier, restricted to the two packets, is a finite sum of tensors in $S_u \otimes S_v$, each of which can be absorbed into the shuffle inputs. Common translation of all particle coordinates preserves I_d^μ, since every kernel factor is a difference.

Theorem 7.1. *Choose homogeneous linear sections $B_d^\mu \to S_d$ for $d \neq 0$. Ordered shuffle multiplication gives an isomorphism of vector spaces*

$$\bigoplus_{\substack{d^1+\cdots+d^s=d \\ \mu(d^1)>\cdots>\mu(d^s)}} B_{d^1}^\mu \otimes \cdots \otimes B_{d^s}^\mu \xrightarrow{\sim} S_d. \tag{7.4}$$

All parts d^i in this sum are nonzero; the empty list represents the unit when $d = 0$. A tensor of polynomial degrees $j_1, \ldots, j_s$ on the left has degree $\sum_i j_i - \sum_{i<l} \chi(d^i, d^l)$ on the right.

Proof. Surjectivity follows by induction on $|d| = \sum_i d_i$. By (7.3), any polynomial is a section representative plus a sum of products with $\mu(u) > \mu(d) > \mu(v)$. Apply the induction hypothesis to the two smaller packets. Draw each resulting list as a polygonal path with successive increments $(c(d^i), \eta(d^i))$. Whenever adjacent slopes are nondecreasing, combine their product and decompose it by the lower-particle induction. A decreasing list lies above the chord of its endpoints, and that chord lies above the old two-segment path. Thus each replacement weakly raises the whole path. If the path is unchanged, the old slopes were equal and the replacement has one segment, so the number of segments decreases.

The two packets being combined never make up the entire list. The original unstable cut lies strictly above the total endpoint chord, and every subsequent path remains above that cut. A whole list consisting of two nondecreasing slopes would lie below the total chord, a contradiction. Thus lower-particle induction is valid at every replacement. There are only finitely many ordered lists of nonzero subvectors summing to d. The path order, followed by decreasing segment count in the equality case, proves termination. Its terminal lists have strictly decreasing slopes, proving surjectivity.

For injectivity, fix a decreasing target list $R = (r_1, \ldots, r_q)$. Partition the variables into these packets, translate all variables in packet j by an independent center z_j, and pass to the quotient $B_{r_j}^\mu$ in every packet. This gives a polynomial map

$$S_d \longrightarrow \left(\bigotimes_j B_{r_j}^\mu \right) [z_1, \ldots, z_q]. \tag{7.5}$$

To analyze its value on a rational shuffle expression, extend scalars to $\mathbb{Q}(z_1, \ldots, z_q)$ and complete in nonnegative internal polynomial degree. A difference between distinct target packets has nonzero constant term $z_j - z_i$ and is invertible in this completion. The original polynomial map still takes values in the uncompleted tensor quotient.

Consider a decreasing source list $D = (d_1, \ldots, d_p)$. Group its shuffle assignments by the full intersection dimension vectors d_{ij} of source packet i with target packet j. The row sums are d_i and column sums are r_j. Taylor coefficients of the source polynomials and the cross-target kernels are symmetric within every intersection packet. In each internal degree they belong to $\bigotimes_{i,j} S_{d_{ij}}$. Expand them into finite sums of pure tensors and absorb these tensors into the

intersection-packet inputs. Inside target j, the remaining sum is the iterated shuffle product in source order.

Write $u_{hj} = \sum_{i \le h} d_{ij}$. If $u_{hj} \neq 0$ and $\mu(u_{hj}) > \mu(r_j)$, then $r_j - u_{hj} \neq 0$. Associativity groups the internal product into its prefix and suffix; it lies in $I^{\mu}_{r_j}$ and vanishes. This is legitimate coefficient by coefficient because that subspace is an ordinary ideal. Consequently every surviving assignment satisfies

$$\eta(u_{hj}) \le \mu(r_j) c(u_{hj}), \qquad 0 \le c(u_{hj}) \le c(r_j). \tag{7.6}$$

At fixed total charge C, maximize $\sum_j \mu(r_j) t_j$ under $0 \le t_j \le c(r_j)$ and $\sum_j t_j = C$. If $j < l$, $t_j < c(r_j)$, and $t_l > 0$, moving $\min\{c(r_j) - t_j, t_l\}$ from t_l to t_j strictly increases the sum. Iterating fills the higher-slope packets first, giving the unique maximum, the height of the target polygon at horizontal coordinate C. Apply this bound with $t_j = c(u_{hj})$. Summing (7.6) therefore puts every source breakpoint below the target polygon. Concavity of the target puts the entire source path below it. Thus a source polygon not pointwise at most the target polygon has zero restriction.

If two polygon functions agree, their strictly decreasing slopes and segment lengths agree in order. At the first breakpoint, equality in the capacity bound forces the first target packet to be full and all later packets to be empty. Since $0 \le d_{ij} \le r_j$ componentwise and c is strictly positive on nonzero nonnegative vectors, full or zero capacity forces the corresponding full or zero dimension vector. Repeating shows $d_i = r_i$ and that the only surviving intersection assignment is diagonal. For this assignment, (7.5) is the tensor of translated section representatives multiplied by

$$K_R = \prod_{1 \le i < j \le q} \prod_{u,v \in I} \prod_{\alpha=1}^{(r_i)_u} \prod_{\beta=1}^{(r_j)_v} (z_j - z_i + x^{(j)}_{v,\beta} - x^{(i)}_{u,\alpha})^{a_{uv} - \delta_{uv}}. \tag{7.7}$$

This kernel is separately symmetric in every complete target packet and has invertible scalar constant term $\prod_{1 \le i < j \le q} (z_j - z_i)^{-\chi(r_i, r_j)}$. It is therefore a unit on the completed tensor quotient, even when that quotient has zero divisors. The unique diagonal assignment contributes with coefficient one.

Take a finite relation between the images of the summands of (7.4). Choose a pointwise-minimal source polygon among the nonzero terms and use its full list as the target. All incomparable or larger polygons vanish. A strictly smaller one is absent by minimality; equal polygons with a different full list also vanish. Cancel (7.7) from the remaining term. The tensor of translated section representatives is then zero. Its vanishing in the completion implies its vanishing before completion, since a nonnegatively graded vector space injects into the product of its graded pieces. Common translation is an automorphism of each polynomial quotient, so inverse translation kills the original tensor coefficient. Remove this list and repeat. This proves injectivity. The grading statement follows from (7.2). □

For a nonnegatively graded vector space V with finite-dimensional graded pieces, write $\mathrm{Hilb}(V; t) = \sum_{j \ge 0} (\dim V_j) t^j$, taking dimensions over its stated ground field. The elementary symmetric generators of S_d give

$$\mathrm{Hilb}(S_d; t) = \prod_{i \in I} \prod_{j=1}^{d_i} (1 - t^j)^{-1}.$$

Theorem 7.1 gives, for $d \neq 0$,

$$\mathrm{Hilb}(B^{\mu}_d; t) = \mathrm{Hilb}(S_d; t) - \sum_{\substack{d^1 + \cdots + d^s = d,\ s \ge 2 \\ d^i \neq 0,\ \mu(d^1) > \cdots > \mu(d^s)}} t^{-\sum_{i<l} \chi(d^i, d^l)} \prod_{i=1}^{s} \mathrm{Hilb}(B^{\mu}_{d^i}; t).$$

Every part in this finite sum has smaller total dimension. Induction therefore proves that $\mathrm{Hilb}(B_d^\mu;t)$ is rational.

Put $\langle d,e\rangle=\chi(e,d)-\chi(d,e)$, and introduce based monomials with multiplication

$$X^dX^e=v^{\langle d,e\rangle}X^{d+e}. \tag{7.8}$$

All series are completed by nonnegative dimension vectors; each fixed dimension involves finitely many decompositions. Coefficients belong to $\mathbb{Q}(v)$. Define the full series and the slope factors by

$$\mathcal{S}=\sum_d v^{\chi(d,d)}\,\mathrm{Hilb}(S_d;v^2)X^d,\qquad H_\rho=1+\sum_{\substack{d\neq0\\ \mu(d)=\rho}} v^{\chi(d,d)}\,\mathrm{Hilb}(B_d^\mu;v^2)X^d.$$

Theorem 7.1 gives $\mathcal{S}=\prod_\rho^{\searrow}H_\rho$, with factors in decreasing slope order. Indeed

$$\begin{aligned}&\chi(d+e,d+e)+2(i+j-\chi(d,e))\\ &\qquad=\chi(d,d)+2i+\chi(e,e)+2j+\langle d,e\rangle,\end{aligned}$$

which is precisely the weight prescribed by (7.8). For any fixed total-dimension cutoff, an ordered factorization by prescribed slopes is unique: the unknown term of each smallest unresolved dimension occurs linearly in its one prescribed factor, and all other terms use smaller dimensions.

8. The attached algebra and the two-input wall function

Fix an integer $m\geq1$, finite alphabets $\alpha=(\alpha_1,\dots,\alpha_s)$ and $\beta=(\beta_1,\dots,\beta_t)$ and disjoint color sets $I_A=\{A_1,\dots,A_s\}$ and $I_B=\{B_1,\dots,B_t\}$. Take the quiver on $I=I_A\sqcup I_B$ with m arrows from each B-color to each A-color. For its dimension vectors put $A(d)=\sum_{i\in I_A}d_i$ and $B(d)=\sum_{j\in I_B}d_j$. Its Euler form and skew pairing are

$$\chi(d,e)=\sum_i d_ie_i-mB(d)A(e),\qquad \langle d,e\rangle=m\big(B(d)A(e)-A(d)B(e)\big). \tag{8.1}$$

The skew form depends only on the two aggregate dimensions. Define the commuting tags by

$$\alpha^{d_A}=\prod_{i=1}^s\alpha_i^{d_{A_i}},\qquad \beta^{d_B}=\prod_{j=1}^t\beta_j^{d_{B_j}}.$$

They retain the color information when the based monomials are projected to the rank-two aggregate lattice.

8.1. **Central factors and their two framed chambers.** Fix coprime $a,b>0$, and write $\gamma=(a,b)$. The central character is $\theta_0=bA-aB$; use slope $-\theta_0/(A+B)$ on the unframed quiver. Let

$$H(U;v)=1+\sum_{k\geq1}\sum_{(A(d),B(d))=k(a,b)} v^{\chi(d,d)}\,\mathrm{Hilb}(B_d^\mu;v^2)\,\alpha^{d_A}\beta^{d_B}U^k. \tag{8.2}$$

All monomials in this central sector commute. The series does not involve a choice of framing.

Choose $p,q>0$ with $pa-qb=1$. Add a vertex ∞, p arrows from ∞ to each A-color, and q arrows from every B-color to ∞. Restrict to framing dimension one.

On framed dimension vectors set

$$\tau=A+B+d_\infty,\qquad \phi=pA-qB,\qquad \mu_\pm=\frac{-\theta_0\pm\epsilon\phi}{\tau}.$$

Within any fixed finite dimension cutoff, choose a rational $\epsilon>0$ small enough that different old slopes $-\theta_0/\tau$ keep their order. Group the new ordered factors by their old slopes. Uniqueness of the old factorization identifies each such group with its old factor. In the old central sector, an unframed dimension $k\gamma$ has $\phi/\tau=1/(a+b)$, whereas a framed dimension $(l\gamma,1)$ has $\phi/\tau=l/(l(a+b)+1)<1/(a+b)$. Hence the positive refinement puts the unframed factor before

the framed factors, and the negative refinement puts it after them. In framing dimension one there can be only one framed factor in a product. If W is the framing-degree-one part of the old central factor, we obtain

$$W = H\,T_{\text{after}} = T_{\text{before}}\,H. \tag{8.3}$$

We verify that these are the chambers used in Corollary 4.3. In total dimension $(k\gamma, 1)$, a noncentral subset has a nonzero integer value under θ_0; its sign is unchanged for a sufficiently small refinement. A balanced unframed subset $l\gamma$ has a larger μ_+-slope than the whole framed dimension, and is forbidden by (7.3). The balanced subset has κ-value $N_A la - N_B lb = l(pa - qb) = l$. This gives exactly the negative value $-\delta l$ of the after character $\theta_0 - \delta\kappa$, with $\delta > 0$. For μ_- the sign is reversed, giving the before character. A framed proper subset has an unframed complement, and their character values are opposite. Thus every forbidden split agrees. A common small refinement exists in each finite cutoff, and enlarging the cutoff changes none of the earlier forbidden splits. Equation (8.3) is therefore a compatible identity of full formal series.

8.2. **The thin and colored polynomial quotients.** The thin-quiver quotient and Vandermonde comparison below are algebraic versions of the presentation and alternation in [12, Sections 3.1–3.2, Lemma 3.8]. The cut products and the linear forms (4.2) have integer coefficients. Hence the cut algebra has the rational form

$$H_{\theta,\mathbb{Q}} = \mathbb{Q}[\mathbf{w}]/(J_\theta, \ell_1, \dots, \ell_n), \qquad H_\theta = H_{\theta,\mathbb{Q}} \otimes_{\mathbb{Q}} \mathbb{C}.$$

The particle actions preserve this form. Scalar extension preserves graded dimensions and commutes with the sign projectors $|W|^{-1}\sum_{w\in W} \operatorname{sgn}_W(w)w$, so it also preserves sign-isotypic dimensions. Fix a colored dimension $(d, 1)$ with $(A(d), B(d)) = k(a, b)$, $k \geq 1$, and replace each color i by d_i labeled particles. Let P be the polynomial ring in these particles and the frame variable x_∞, let $W_d = \prod_i S_{d_i}$, and put

$$R = P^{W_d}, \qquad \Delta_d = \prod_i \prod_{r<l} (x_{i,l} - x_{i,r}), \qquad \zeta(d) = \sum_i \binom{d_i}{2}.$$

For the chosen chamber, let J be the ideal in P generated by all forbidden outgoing cut products

$$E_S = \prod_{e:S\to S^c} (x_{\text{head}(e)} - x_{\text{tail}(e)}).$$

In thin dimension each split is a labeled subset, so its sole shuffle term has kernel E_S and product image $E_S P$. Thus the thin quotient of (7.3) is P/J. Subtracting the common frame coordinate identifies it with $H_{\theta,\mathbb{Q}}[x_\infty]$; setting $x_\infty = 0$ gives $H_{\theta,\mathbb{Q}}$.

Let I denote the corresponding colored ideal in R. For a forbidden split of $(d, 1)$, let e, f be its unframed parts and fix its standard particle subset, with the frame in the prescribed packet. Alternating an arbitrary multiple of E_S first within the two packets produces $E_S\Delta_e\Delta_f$ times an arbitrary element of $P^{W_e\times W_f}$. Write $\Delta_d = \Delta_e\Delta_f D_{e,f}$, where $D_{e,f}$ is the product of the same-color differences between the two packets. Alternating over the remaining shuffles then gives the subspace spanned by

$$\Delta_d \sum_{\sigma\in W_d/(W_e\times W_f)} \sigma\left(\frac{E_S h}{D_{e,f}}\right), \qquad h \in P^{W_e\times W_f}.$$

The sum is exactly the shuffle expression for that split. Normalizing the alternation operators only multiplies it by nonzero rational scalars and does not change the image. Summing over the forbidden splits yields

$$J \cap P^{\operatorname{sgn}_{W_d}} = \Delta_d I, \qquad R/I \xrightarrow[\sim]{\cdot\Delta_d} (P/J)^{\operatorname{sgn}_{W_d}}. \tag{8.4}$$

The equality also proves surjectivity, since every alternating polynomial is divisible by Δ_d with invariant quotient. Exactness of averaging in characteristic zero permits alternation after passing to the quotient.

The frame variable does not occur in Δ_d; all cut products depend only on coordinate differences. Hence (8.4) commutes with setting $x_\infty = 0$ after subtracting it from every other coordinate. The resulting colored quotient is finite, because its sign space is contained in the finite cut algebra. Write the two Hilbert polynomials as $P_{\text{before}}(d;t)$ and $P_{\text{after}}(d;t)$, corresponding respectively to H_- and H_+. Then

$$\begin{aligned} \mathrm{Hilb}(H_-^{\mathrm{sgn}_{W_d}}; v^2) &= v^{2\zeta(d)} P_{\text{before}}(d; v^2), \\ \mathrm{Hilb}(H_+^{\mathrm{sgn}_{W_d}}; v^2) &= v^{2\zeta(d)} P_{\text{after}}(d; v^2). \end{aligned} \tag{8.5}$$

For the pure frame set $P_{\text{before}}(0;t) = P_{\text{after}}(0;t) = 1$. For fixed totals A, B, the thin unframed Euler number is $\chi_{\text{thin}} = A + B - mAB$, and

$$\chi(d,d) - \chi_{\text{thin}} = 2\zeta(d). \tag{8.6}$$

Thus the Vandermonde degree exactly matches the difference of the two Euler normalizations.

8.3. **The framing scalar and regularity at one.** Let $\delta = e_\infty$, $z = X^\delta$, and let f, g be the vectors of arrow counts from and to the frame. Thus $f \cdot d = pA(d)$ and $g \cdot d = qB(d)$. Use the subscript ext for the Euler form and skew pairing of the framed quiver. Its Euler number is

$$\chi_{\text{ext}}(d+\delta, d+\delta) = \chi(d,d) + 1 - (f+g)\cdot d.$$

The common frame coordinate contributes Hilbert series $(1-v^2)^{-1}$. After removing the common scalar $v/(1-v^2)$, the coefficient in either chamber is its Hilbert polynomial multiplied by $v^{\chi(d,d)-(f+g)\cdot d}$. Moreover

$$\langle \delta, d\rangle_{\text{ext}} = (f-g)\cdot d, \qquad X^{d+\delta} = v^{-(f-g)\cdot d} z X^d.$$

Moving the frame monomial to the left therefore changes $f+g$ to $2f$. The normalized framed series are zA_{before} and zA_{after}, where

$$\begin{aligned} A_{\text{before}}(U;v) &= \sum_{k\ge 0} \sum_{(A(d),B(d))=k(a,b)} v^{\chi(d,d)-2pA(d)} \\ &\qquad \cdot P_{\text{before}}(d; v^2)\, \alpha^{d_A} \beta^{d_B} U^k, \\ A_{\text{after}}(U;v) &= \sum_{k\ge 0} \sum_{(A(d),B(d))=k(a,b)} v^{\chi(d,d)-2pA(d)} \\ &\qquad \cdot P_{\text{after}}(d; v^2)\, \alpha^{d_A} \beta^{d_B} U^k. \end{aligned} \tag{8.7}$$

Each constant coefficient is one, contributed by the pure frame. Every coefficient is a Laurent polynomial in v with integral tag coefficients.

For an unframed dimension d with $(A(d), B(d)) = k\gamma$, one has $\langle \delta, d\rangle_{\text{ext}} = pA(d) - qB(d) = k$. Hence $X^d z = v^{-2k} z X^d$, giving $Uz = v^{-2} zU$. Cancel $v/(1-v^2)$ and z in (8.3). The remaining central monomials commute, so the result is

$$A_{\text{before}}(U;v) = A_{\text{after}}(U;v) F_{\text{H}}(U;v), \qquad F_{\text{H}}(U;v) = \frac{H(v^{-2}U;v)}{H(U;v)}. \tag{8.8}$$

Because A_{after} has constant coefficient one, its formal inverse also has Laurent-polynomial coefficients. Thus every coefficient of F_{H} is a Laurent polynomial, and we can define

$$F_{\text{orig}}(U) = F_{\text{H}}(U;v)\big|_{v=1}. \tag{8.9}$$

Both H and F_{orig} are independent of the chosen positive Bézout pair.

8.4. **Recovering the whole character.** For any finite-dimensional representation E of $G = S_A \times S_B$, the coefficient of an individual monomial $\alpha^{d_A}\beta^{d_B}$ in $\operatorname{ch} E$ is $\dim E^{W_d}$. To see this directly from the Frobenius formula, expand each power sum as a sum of alphabet powers. A contribution assigns one color to each cycle. There are $[G : W_d]$ label colorings with multiplicity vector d, each with stabilizer conjugate to W_d. Regrouping the average by these colorings gives

$$[\alpha^{d_A}\beta^{d_B}] \operatorname{ch} E = \frac{[G : W_d]}{|G|} \sum_{w \in W_d} \operatorname{tr}(w \mid E) = \dim E^{W_d}.$$

Tensoring with the two sign representations gives

$$[\alpha^{d_A}\beta^{d_B}] \, \iota \operatorname{ch} E = \dim \operatorname{Hom}_{W_d}(\operatorname{sgn}_{W_d}, E). \tag{8.10}$$

Permuting the colors permutes the tag variables without changing the shuffle construction. Adjoining a color and then setting its tag to zero retains exactly the dimensions having zero component at that color, and recovers the original construction. Thus the coefficient of U^k in F_{orig} is a compatible symmetric polynomial of bidegree (ka, kb). It determines a unique element of the ring of symmetric functions in two alphabets: alphabets of sizes at least ka, kb already distinguish every monomial symmetric function of this bidegree. We apply ι in this stable ring, before specializing to any finite alphabet.

Define the two ordinary character series

$$\mathcal{B}^{p,q}(U) = \sum_{k \geq 0} \operatorname{ch} \mathcal{B}_k^{p,q} U^k, \qquad \mathcal{C}^{p,q}(U) = \sum_{k \geq 0} \operatorname{ch} \mathcal{C}_k^{p,q} U^k.$$

At $v = 1$, the grading powers in (8.7) disappear. Equations (8.5)–(8.10) and Corollary 4.3 identify its before and after coefficients with $\iota\mathcal{B}^{p,q}$ and $\iota\mathcal{C}^{p,q}$. Consequently (8.8) becomes

$$\iota\mathcal{B}^{p,q} = (\iota\mathcal{C}^{p,q}) F_{\mathrm{orig}}, \qquad \mathcal{B}^{p,q} = \mathcal{C}^{p,q} F_{\mathrm{Y}}, \quad F_{\mathrm{Y}} = \iota F_{\mathrm{orig}}. \tag{8.11}$$

The equality holds in the completed ring of symmetric functions in both alphabets, because (8.10) identifies every monomial coefficient in each of the sufficiently large alphabets just described.

Proposition 8.1. *For every $k \geq 0$ and every positive Bézout pair $pa - qb = 1$,*

$$\operatorname{ch} \mathcal{D}_k^{p,q} = [U^k] F_{\mathrm{Y}}, \qquad [U^k] F_{\mathrm{orig}} = \iota \operatorname{ch} \mathcal{D}_k^{p,q}. \tag{8.12}$$

In particular the attached character is independent of (p, q) and

$$[U^k] F_{\mathrm{orig}} = \sum_{\lambda \vdash ka,\, \mu \vdash kb} n_{\lambda\mu}\, e_\lambda(\alpha) e_\mu(\beta), \qquad n_{\lambda\mu} \in \mathbb{N}. \tag{8.13}$$

Proof. Induct on k, with the assertion quantified over all positive pairs (p, q) at each stage. The case $k = 0$ is the one-dimensional empty particle algebra. Frobenius characteristic sends induction of external tensor products to multiplication, so Theorem 5.2 gives

$$\operatorname{ch} \mathcal{B}_k^{p,q} = \sum_{r+l=k} (\operatorname{ch} \mathcal{C}_r^{p,q})(\operatorname{ch} \mathcal{D}_l^{p+mrb,q+mra}).$$

Whenever $r > 0$, we have $l < k$ and $(p + mrb)a - (q + mra)b = 1$. The induction hypothesis therefore replaces that attached character by $[U^l] F_{\mathrm{Y}}$. Comparing with the coefficient of U^k in (8.11) leaves only $r = 0$ and proves (8.12). Finally apply ι to (6.4) to obtain (8.13). □

8.5. **Identification with formal two-input reordering.** On the aggregate lattice, for $w = (R, S)$ set $K_w = m(aS - bR)$. By (8.1), $\langle \gamma, w \rangle = -K_w$. Conjugation by the central factor is exactly

$$X^w \longmapsto X^w \frac{H(v^{-2K_w} U; v)}{H(U; v)}. \tag{8.14}$$

For $K > 0$, the ratio telescopes as

$$\frac{H(v^{-2K}U; v)}{H(U; v)} = \prod_{j=0}^{K-1} F_{\mathrm{H}}(v^{-2j}U; v).$$

For negative K it is the inverse of the corresponding finite product, and for $K = 0$ it is one. Coefficientwise regularity already proved in (8.8) permits specialization at $v = 1$. Substituting $U = x^a y^b$ gives the classical operation

$$x^R y^S \longmapsto x^R y^S F_{\mathrm{orig}}(x^a y^b)^{m(aS-bR)}. \tag{8.15}$$

An arrowless single-color axis has series

$$E_v(z) = \sum_{n\geq 0} \frac{v^{n^2} z^n}{\prod_{j=1}^{n}(1 - v^{2j})}.$$

Direct comparison of coefficients gives $E_v(v^{-2}z) = (1 + v^{-1}z)E_v(z)$. The two aggregate axis factors are

$$\prod_{i=1}^{s} E_v(\alpha_i X^{(1,0)}), \qquad \prod_{j=1}^{t} E_v(\beta_j X^{(0,1)}).$$

The same conjugation calculation gives their classical actions

$$(x, y) \longmapsto (x, yP(x)^m), \qquad (x, y) \longmapsto (xQ(y)^{-m}, y), \tag{8.16}$$

where $P(z) = \prod_i(1 + \alpha_i z)$ and $Q(z) = \prod_j(1 + \beta_j z)$.

The full polynomial series is the A-axis factor followed by the B-axis factor. This can be checked in the shuffle algebra: a split into all A-particles first and all B-particles second has kernel one, so its product is the full polynomial space. For a slope increasing from the A-axis to the B-axis, Theorem 7.1 refactors the same full series into the B-axis, the ordered interior central factors, and the A-axis. Taking conjugations and the classical limits above gives precisely the formal two-input reordering of (8.16). Uniqueness can also be seen directly degree by degree: at the first unresolved degree an interior monomial belongs to exactly one ray, and its correction appears linearly.

Theorem 8.2 (Two-input positivity)**.** *Let*

$$P(z) = 1 + c_1 z + \cdots + c_r z^r, \qquad Q(z) = 1 + d_1 z + \cdots + d_s z^s$$

have independent coefficients. In the formal reordering of $(x, y) \mapsto (x, yP(x))$ *and* $(x, y) \mapsto (xQ(y)^{-1}, y)$*, write the factor on each primitive interior ray* (a, b) *in the form*

$$x^R y^S \longmapsto x^R y^S f_{a,b}(x^a y^b)^{aS-bR}, \qquad f_{a,b}(0) = 1.$$

Every coefficient of $f_{a,b}$ *belongs to* $\mathbb{N}[c_1, \ldots, c_r, d_1, \ldots, d_s]$*. This includes every multiple of the primitive ray.*

Proof. Set $m = 1$ in (8.15). Because $\gcd(a, b) = 1$, its normal $(R, S) \mapsto aS - bR$ is primitive. Thus $f_{a,b} = F_{\mathrm{orig}}$ with the indicated ray variable. The coefficients c_i, d_j are the elementary symmetric functions $e_i(\alpha), e_j(\beta)$. Formula (8.13) is a polynomial with nonnegative integer coefficients in those generators. They are algebraically independent for alphabets of sizes r, s; elementary functions beyond these sizes specialize to zero. The resulting equality therefore proves the assertion over the universal coefficient ring and under every positive specialization of its coefficients. □

9. Positive formal wall operations

We now pass from Theorem 8.2 to arbitrary rank. We use finite quotients of a graded Lie algebra throughout. The several-input mechanism is established in [5, Section 9.2]; our formal transport is adapted from [21, Section 4]. We give the details for the present coefficient convention and primitive ambient normals.

9.1. **The completed Hamiltonian group.** Let M be a lattice, let $\Omega : M \times M \to \mathbb{Z}$ be an alternating form, and let $C = \{p_1, \dots, p_n\}$ be part of a basis of M. Write

$$C^+ = \sum_i \mathbb{N} p_i, \qquad \left| \sum_i a_i p_i \right| = \sum_i a_i.$$

Assume that the span of C meets $\operatorname{rad}\Omega$ only at zero. Let R_+ be a subsemiring of a characteristic-zero domain R, and put $K = \operatorname{Frac}(R)$. The completed Lie algebra

$$\mathfrak{g}_C = \prod_{a \in C^+ \setminus \{0\}} K X^a, \qquad [X^a, X^b] = \Omega(a,b) X^{a+b} \tag{9.1}$$

is filtered by root degree. Its quotient in degrees at most d is nilpotent and finite-dimensional over K. We denote the corresponding Baker–Campbell–Hausdorff group by G_C, understood as the inverse limit of these nilpotent groups. It acts on monomials by Hamiltonian derivations,

$$\{X^a, X^w\} = \Omega(a,w) X^{a+w}.$$

The action is faithful: a nonzero coefficient at the first root a of a Hamiltonian is detected on some X^w, since $\Omega(a,-) \neq 0$. All exponential actions are well-defined on series whose supports are contained in a finite union of translates of C^+.

If a is primitive on a positive root ray, write $\Omega(a,-) = \delta_a n_a$, where $\delta_a > 0$ and $n_a \in M^*$ is primitive. The ray Hamiltonian $H = \sum_{j \geq 1} h_j X^{ja}$ acts by

$$\exp(\{H,-\})(X^w) = X^w f(X^a)^{n_a(w)}, \qquad f(Z) = \exp\left(\delta_a \sum_{j \geq 1} j h_j Z^j \right). \tag{9.2}$$

We call f the wall function with primitive normal. Indeed, $\{H, X^w\}/X^w$ is $n_a(w)\delta_a \sum_j j h_j X^{ja}$, and the derivation annihilates every X^{ja}. Exponentiating proves the formula. Positivity will concern the coefficients of the wall function f.

Lemma 9.1 (Three-factor decomposition)**.** *For $h \in M^*_{\mathbb{R}}$, every $g \in G_C$ has a unique factorization*

$$g = g_+(h) g_0(h) g_-(h),$$

where the logarithms of the three factors are supported on roots with positive, zero, and negative h-value, respectively. These factorizations are compatible with truncation and with refinement under a sufficiently small displacement of h.

Proof. The three indicated spaces are Lie subalgebras, since signs are preserved by addition within each one. Work first through degree d. Suppose that the factors have been found below degree s. For a root a of degree s, its unknown logarithmic coefficient occurs with coefficient one in exactly one of the three factors. Every other contribution to the coefficient of X^a is a commutator involving roots of smaller degree. Thus the coefficient is determined uniquely. This proves existence, uniqueness, and compatibility for increasing d.

For the refinement assertion, take h' close enough to h that the signs of all nonzero h-values through degree d are unchanged. Factor $g_0(h)$ according to the signs of h'. Its positive factor can be multiplied into $g_+(h)$, and its negative factor into $g_-(h)$; the three resulting factors have

precisely the required signs. Uniqueness identifies them with the factors for h'. Passing through the compatible finite quotients proves the assertion. □

In a finite quotient, the finitely many hyperplanes $a^\perp$ divide $M^*_{\mathbb{R}}$ into cells. At a generic point of a codimension-one cell, all roots with zero value lie on one ray, so $g_0(h)$ has a wall function. A crossing from its positive side to its negative side acts by this middle factor; the reverse crossing acts by its inverse. Lemma 9.1 shows that the action from the all-positive chamber to a generic h is $g_-(h)$. Consequently path products depend only on their endpoints. This is the consistency property used below.

Conversely, a consistent central system of such wall factors is recovered from its total product g from the all-positive chamber to the all-negative chamber. Indeed, choose a covector η positive on C, and use the path $h+t\eta$, with t decreasing from a large positive value to a large negative value. The roots crossed before, at, and after $t=0$ have negative, zero, and positive h-value. Grouping these factors gives the three-factor decomposition of g. The total product is independent of h by consistency, and uniqueness recovers each middle factor. This argument is made in every finite quotient. Thus uniqueness of g below is also uniqueness of the wall system, up to subdividing its cells.

Put $V(a)=\Omega(-,a)\in M^*_{\mathbb{R}}$. The point $V(a)$ lies in $a^\perp$. It specifies the incoming direction of a wall on the ray of a. The following observation deals with the fact that this point need not be generic.

Lemma 9.2 (Incoming prescription)**.** *For each i, choose a polynomial $Q_i(Z)$ with constant term one, and write*

$$\Omega(p_i,-)=\delta_i n_i, \qquad \log Q_i(Z)=\sum_{j\ge 1} c_{ij}Z^j,$$

where n_i is primitive and $\delta_i>0$. There is a unique $g\in G_C$ for which

$$[X^a]\log g_0(V(a))=\begin{cases} c_{ij}/(\delta_i j), & a=jp_i,\ j\ge 1,\\ 0, & a \text{ lies on no simple ray.}\end{cases} \tag{9.3}$$

At a small generic displacement within $a^\perp$ of the incoming direction $V(a)$, these are the incoming coefficients.

Proof. The coefficient of X^a in $\log g$ enters the left side of (9.3) with coefficient one. All remaining terms involve smaller roots, by the proof of Lemma 9.1. Thus (9.3) determines $\log g$ recursively.

Let $\lambda_a(H)=[X^a]H$. On the Lie subalgebra supported in $\ker V(a)$, this functional annihilates every bracket. In fact, if $b+c=a$ and $\Omega(b,a)=0$, then $\Omega(b,c)=\Omega(b,a-b)=0$. The same argument applied to the last bracket in a nested bracket shows that

$$\lambda_a(\mathrm{BCH}(H_1,H_2))=\lambda_a(H_1)+\lambda_a(H_2)$$

for H_1,H_2 in this subalgebra. Refine $g_0(V(a))$ at a small generic displacement in $a^\perp$. The factors of strict sign cannot contain X^a, and the displayed additivity proves that its coefficient in the middle factor is unchanged. The same argument applies to every positive multiple of a. Hence the incoming simple wall has function Q_i by (9.2), and all nonsimple incoming functions are one. □

9.2. **Two inputs on an ambient lattice.** We apply the two-input theorem while retaining the integral normals of M.

Lemma 9.3. *Suppose that*

$$\phi_i(X^w)=X^w P_i(X^{r_i})^{n_i(w)} \qquad (i=1,2),$$

where $n_i \in M^*$, $n_i(r_i) = 0$, the n_i are positive scalar multiples of $\Omega(r_i, -)$, and the polynomials P_i have constant term one. If their coefficients belong to R_+, then all functions in their ordered two-input factorization, when written with primitive integral normals on M, have coefficients in R_+.

Proof. The commuting case is immediate. Otherwise, after relabeling, put

$$u = n_1(r_2) > 0, \qquad v = -n_2(r_1) > 0, \qquad e_1 = r_1/v, \quad e_2 = r_2/u.$$

In the enlarged lattice $M' = M + \mathbb{Z}e_1 + \mathbb{Z}e_2$, both n_1 and n_2 remain integral: their values on e_1, e_2 are $(0, 1)$ and $(-1, 0)$. On $L = \mathbb{Z}e_1 + \mathbb{Z}e_2$, the two operations are the standard two-input operations for

$$P_1(x^v), \qquad P_2(y^u).$$

Theorem 8.2, with absent coefficients specialized to zero, therefore gives positive output functions. On a primitive L-ray $r = ae_1 + be_2$, its ambient normal is $an_1 + bn_2$. Indeed, write $n_i = \lambda_i \Omega(r_i, -)$ with $\lambda_i > 0$. Then

$$\lambda_1 v = \lambda_2 u, \qquad an_1 + bn_2 = (\lambda_1 v)\Omega(ae_1 + be_2, -),$$

which gives the required direction and orientation.

The factorization is also valid on all ambient monomials. To see this, write $\partial_n(X^w) = n(w)X^w$; then

$$[X^r\partial_n, X^s\partial_m] = X^{r+s}(n(s)\partial_m - m(r)\partial_n).$$

Every covector generated here lies in the span of n_1, n_2, whose restriction to L is injective. Thus an identity detected on L is an identity of these derivations on M', and hence of their exponential actions.

The logarithms of the inputs have support in $\mathbb{N}r_1 + \mathbb{N}r_2$. Brackets, exponentials, and the recursive ordered factorization preserve this support monoid. Equation (9.2) shows that the output functions preserve it as well. Their nonconstant monomials therefore belong to M. If $\rho r \in M$, the values $n_1(\rho r) = \rho b$ and $n_2(\rho r) = -\rho a$ are integers. Coprimality of a, b gives $\rho \in \mathbb{Z}$, so

$$M \cap \mathbb{Q}r = \mathbb{Z}(hr) \qquad \text{for some } h \in \mathbb{Z}_{>0}.$$

The support condition therefore writes the output function as $f(Z) = g(Z^h)$, where $Z = X^r$ and g has the same nonzero coefficients as f. If $an_1 + bn_2 = tn_{\mathrm{prim}}$ on M, then $t \in \mathbb{Z}_{>0}$, and the action on an ambient monomial is

$$X^w \longmapsto X^w\big(g(X^{hr})^t\big)^{n_{\mathrm{prim}}(w)}.$$

Thus the function in the primitive ambient ray variable, with primitive integral normal, is g^t. It has coefficients in R_+. □

9.3. **Several inputs and higher rank.**

Proposition 9.4. *Let a finite collection of incoming polynomial line functions have roots in a common pointed rank-two cone. The lines pass through the origin, and the functions use the primitive integral normals of the ambient lattice. If their constant terms are one and their coefficients belong to R_+, then they have a consistent completion whose additional outgoing ray functions have coefficients in R_+. The assertion holds coefficientwise for every positive integer weight filtration with finite pieces.*

Proof. We retain the ambient normals, so every two-input operation below is covered by Lemma 9.3. Commuting inputs already form a consistent diagram. Otherwise work in the dual of the two-dimensional root span, whose alternating pairing is nondegenerate.

Give each nonconstant original input monomial an independent coefficient c_s and a label z_s. The label monoid is $F = \mathbb{N}^S$, with positive integer weights w_s, and the geometric exponents define a map $\rho : F \to L$ to the root lattice. A labeled monomial $z^\alpha X^{\rho(\alpha)}$ has weight $\sum_s w_s\alpha_s$. Different labels may have the same geometric image. All operations preserve this grading, whose bounded pieces are finite. Pointedness gives $\rho(\alpha) \neq 0$ for $\alpha \neq 0$.

For full-line inputs at one point, the additional outgoing correction functions are determined recursively. At the first unresolved weight, the logarithm of the small-circle product is a finite sum of labeled Hamiltonian monomials. Each monomial has one prescribed outgoing ray. Its correction coefficient enters the circle product linearly with a nonzero primitive-normal conversion factor, so a unique choice cancels that defect. These corrections are central modulo the next weight. This proves existence, uniqueness, and rational-polynomial dependence on the input coefficients. To count input occurrences, multiply all c_s by ϵ. Modulo ϵ^2, all operations commute and the opposite crossings of each full line cancel. Every nonconstant additional correction therefore contains at least two input occurrences. The original full lines remain unchanged: if an input line has function P and its combined outgoing function is F_0, this construction determines the additional factor F_0/P.

Uniformly over all finite labeled data and positive integer weights, we prove the following assertions simultaneously through weight d:

(i) the additional central correction functions have nonnegative integral polynomial coefficients in the independent c_s;

(ii) after any fixed rational translation of the original lines with distinct parallel lines and no triple original intersection, there is a finite diagram, consistent through weight d, consisting of those unchanged full lines and additional outgoing rays with such coefficients. Each additional ray starts in weight at least two.

Combine rays with the same base point and direction, and read overlapping collinear germs by multiplying their functions. The joints are the ray bases and intersections of nonparallel supports. Both assertions in weight one follow from the first-order calculation above.

Assume both assertions through weight d, and take a translated diagram furnished by the second assertion. At a joint q, use as local full-line inputs the original lines through q and the correction rays whose relative interiors contain q. Rays based at q supply the outgoing germs. Consistency through weight d and local uniqueness identify their functions through that weight with the additional completion of the through-going inputs. Consequently local recomputation preserves every assigned lower coefficient.

An additional coefficient of weight $d+1$ uses at least two positive-weight input occurrences, so only input monomials of weight at most d can contribute. Decompose these retained functions into homogeneous labeled monomials and give each monomial a fresh local label of its original total weight. Its scalar coefficient is a nonnegative polynomial in the c_s, by the affine induction hypothesis. In particular, the calculation depends only on the already assigned coefficients of every through-going ray.

Consider a local label multidegree α of weight $d+1$. If it uses a label of weight at least two, then

$$\sum_s \alpha_s = d + 1 - \sum_s (w_s - 1)\alpha_s \leq d.$$

Give all local labels weight one, retaining their geometric exponents and normals. The central induction assertion through weight d makes this same coefficient a nonnegative integral polynomial in the local input coefficients. Substituting their positive expressions preserves positivity.

If every used label has weight one, it comes from an original line, since every generated ray starts in weight at least two. At most two original lines pass through q. Set all other local labels

to zero; label homogeneity preserves the specified coefficient. Lemma 9.3 gives its positivity. This argument also covers different weights for different monomials of one original polynomial.

At each old joint, assign these weight $d+1$ coefficients to its outgoing corrections, combining equal directions at that base. There are finitely many joints and finitely many monomials at this weight, so only finitely many rays are added. Each modification has weight $d+1$ and is central modulo weight $d+2$. At an old joint, modifications from other bases are through-going central factors whose opposite crossings cancel. At every new intersection away from the bases the same cancellation applies. Thus every joint loop is trivial through weight $d+1$. In a homotopy of a generic path, crossing a smooth part of a wall inserts an operation and its inverse; joint loops are the remaining changes. Every closed path in the finite planar arrangement therefore has product one. This proves the affine assertion through weight $d+1$.

To obtain the central assertion, dilate this finite affine diagram by a positive scalar tending to zero. A fixed circle enclosing all joints has product one. Every original line tends to its prescribed central line, and each correction ray tends to the outgoing half-ray in direction $-V(\rho(\alpha))$. The limiting crossing order changes only among factors of the same geometric direction, which commute. Translation and positive dilation preserve roots and primitive ambient normals. Coincident outgoing corrections consequently use the same primitive normal, and their functions multiply. These products have nonnegative integral coefficients. The unchanged original lines remain separate, including in directions that receive additional corrections. The limiting central diagram is consistent; local uniqueness identifies its corrections with the required ones.

This completes the simultaneous induction. Central uniqueness makes the truncations compatible. Finally specialize c_s to the given coefficients and put $z_s = 1$. Pointedness makes this specialization finite at each geometric exponent, proving the proposition. □

Theorem 9.5. *For the incoming prescription of Lemma 9.2, assume that the coefficients of every Q_i lie in R_+. Every primitive-normal wall function of the resulting consistent system has coefficients in R_+.*

Proof. We argue in finite root degree. Fix a positive ray a, a generic cell of $a^\perp$, and a degree bound d. Within $a^\perp$, follow a generically displaced line parallel to $V(a)$, from the incoming direction to that cell. Only finitely many hyperplanes occur through degree d. A hyperplane $b^\perp$ with $\Omega(b,a)=0$ is parallel to the line and may be avoided. A generic displacement avoids intersections with three linearly independent normals. Every joint where the factor on $a^\perp$ changes therefore has a two-dimensional root span and a nonzero alternating pairing. By refinement in Lemma 9.1, the change is an ordered rank-two factorization.

Induct on d, and then along the finitely many joints of this line. At each joint take the germ on the $+V(b)$-side as the incoming germ for an incident root ray b. For the followed ray a, this is the preceding segment, whose function is known through degree d. The other incoming germs are known through degree $d-1$. An unknown degree d term on another ray enters that degree only linearly and on the same ray; every interaction has larger degree. Set their remaining coefficients to zero and apply Proposition 9.4, retaining the primitive normals of M. The outgoing function on the followed ray is its incoming function multiplied by the additional positive correction. Local uniqueness identifies this product through degree d with the outgoing function, independently of the zero extensions. It is positive. The initial incoming function is Q_i on a simple ray and one otherwise. Compatibility of truncations completes the induction. □

9.4. **A coefficient recursion.** A transported expression means one fixed finite Laurent polynomial in one chamber, with its values in other chambers obtained by the consistent path operations. These values may be infinite series, but their supports remain in the same finite

union $\mathcal{S} = \bigcup_{\nu=1}^{t}(m_\nu + C^+)$. For $w \in M$, define its incoming coefficient by

$$A_w(F) = [X^w]F(V(w) + \varepsilon), \tag{9.4}$$

where a sufficiently small generic displacement is taken at the finite order needed to compute this coefficient. This number does not depend on the displacement. Indeed, a wall separating two such choices has root a with $\Omega(a, w) = 0$. Its possible source monomials $w - ja$ have the same zero pairing with a, and hence cannot change the coefficient of X^w.

Proposition 9.6. *Assume that the coefficients of the Q_i belong to R_+. If all incoming coefficients of a transported expression also belong to R_+, then every coefficient of every one of its values belongs to R_+.*

Proof. For a generic target h, follow $h + sV(w)$ from large s down to zero. All walls are conical, so positive rescaling identifies the initial chamber with that of $V(w) + s^{-1}h$. Its X^w-coefficient is therefore $A_w(F)$. At a crossing with primitive root a, the sign of the crossing is the sign of $\Omega(a, w)$. The source monomial $w - ja$ has exactly the same pairing. By (9.2), the operation on that source consequently involves a nonnegative integral power of the positive wall function of Theorem 9.5. Write $J_i(w - ja_i, j) \in R_+$ for its coefficient adding ja_i, with $j > 0$, at the i-th crossing. Explicitly, if the wall function and its primitive normal are f_i and n_i, then

$$J_i(w - ja_i, j) = [T^j]f_i(T)^{|n_i(w)|}.$$

The coefficient adding zero is one. Taking successive differences gives

$$[X^w]F(h) = A_w(F) + \sum_i \sum_{j>0} J_i(w - ja_i, j)[X^{w-ja_i}]F(h_i), \tag{9.5}$$

where h_i is the chamber immediately before that crossing in this oriented path.

Every source on the right is strictly smaller than w in the addition order of C^+. The interval $\{u \in \mathcal{S} : w - u \in C^+\}$ is finite, because C is part of a lattice basis. It also bounds the roots and crossings needed in (9.5). Induct on this finite partially ordered set, simultaneously for all target chambers. For a minimal element the formula is its incoming coefficient; each subsequent element is a sum of products in R_+. Terms outside $\mathcal{S}$ vanish. This proves the proposition and the finiteness of the recursive calculation for each specified coefficient. □

10. Mutation and finite Laurent fibers

10.1. **The lattice attached to a generalized seed.** Let $\widetilde{B}$ have N rows and n mutable columns, and let $D = \operatorname{diag}(\delta_1, \dots, \delta_n)$ symmetrize its principal part B. Add n auxiliary frozen rows I_n; write the resulting matrix as $\widehat{B}$, and its auxiliary variables as $Y_1, \dots, Y_n$. On $M = \mathbb{Z}^N \oplus \mathbb{Z}^n$, let $\pi : \mathbb{Z}^N \to \mathbb{Z}^n$ select the mutable coordinates and set

$$p_i = -\widehat{b}_i, \qquad n_i(w) = -w_i, \qquad \Omega = \begin{pmatrix} 0 & -\pi^t D \\ D\pi & -DB \end{pmatrix}. \tag{10.1}$$

The matrix is alternating, and direct multiplication gives

$$\Omega(p_i, w) = \delta_i n_i(w), \qquad n_i(p_j) = b_{ij}, \qquad \Omega\left(\sum_i a_i p_i, e_j\right) = -\delta_j a_j \quad (1 \le j \le n). \tag{10.2}$$

The auxiliary coordinates of the p_i form $-I_n$; thus they extend to a lattice basis. The last identity in (10.2) proves the required faithfulness on the root span, with arbitrary frozen rows.

For the exchange polynomial $P_i(Z) = \sum_{s=0}^{r_i} q_{i,s} Z^s$, put $Q_i = P_i/q_{i,0}$. The simple incoming operation is

$$\phi_i(X^w) = X^w Q_i(X^{p_i})^{n_i(w)}.$$

The preceding section constructs the corresponding consistent system. We next compare it with the system of a mutated seed. The support estimate below uses the degree of Q_i.

10.2. **Finiteness in one mutation direction.** Fix $p \in C$, write $\Omega(p,-) = \delta n$ with $\delta > 0$, and let its incoming polynomial Q have degree r, constant term one, and nonzero leading coefficient. For $u = \sum_{c \in C\setminus\{p\}} u_c c \neq 0$, with $u_c \in \mathbb{N}$, define

$$|u|_{\neq p} = \sum_{c \neq p} u_c, \qquad M_\sigma(u) = r \sum_{c \neq p} [-\sigma n(c)]_+ u_c \quad (\sigma \in \{+1,-1\}).$$

Lemma 10.1 (Finite fibers)**.** *On the open half-space* $\operatorname{sgn} h(p) = \sigma$*, every nonpure wall Hamiltonian exponent* $kp + u$ *satisfies*

$$0 \leq k \leq M_\sigma(u). \tag{10.3}$$

Every nonpure wall operation and its inverse consequently has finite Laurent support in the p*-direction at each fixed non-*p *degree when applied to a monomial.*

Proof. The pure wall has function Q on every generic part of $p^\perp$. A sum of positive roots lies on $\mathbb{N}p$ precisely when every summand does, and Hamiltonians on this ray commute. Thus its coefficients remain those of the incoming pure wall.

We induct on $|u|_{\neq p}$, verifying a particular k first in an arbitrary sufficiently large full-root truncation. If $n(u) \neq 0$, call $\sigma = \operatorname{sgn} n(u)$ the upstream side. For a point of the wall $(kp+u)^\perp$ on that side, the line towards its incoming end has p-value

$$h(p) + s\Omega(p, kp+u) = h(p) + s\delta n(u) \quad (s \geq 0),$$

so it does not cross $p^\perp$. At the incoming end only simple rays can have nonzero data; for a nonpure one, $k = 0$.

At any joint off $p^\perp$, all active roots have positive non-p degree. A nonlinear term of the current non-p degree is made from strictly smaller positive non-p degrees, while a term of the current degree enters only linearly on its own ray. The set specified by (10.3) is closed under addition on a fixed side. Therefore the coefficient recursion for the local factorization preserves it. Following the finitely many joints proves the upstream bound for the specified k, and hence uniformly for all k.

At a generic joint on $p^\perp$, the non-p projections of the other roots are positive multiples of a common direction. Their nonzero n-values have the same sign. Move the pure factor across their ordered product. The needed conjugation on Hamiltonians is Poisson pullback, and on a monomial it is

$$X^{kp+w} \longmapsto X^{kp+w} Q(X^p)^{|n(w)|}. \tag{10.4}$$

It increases the p-exponent by at most $r|n(w)|$. For these roots,

$$M_{-\sigma}(w) - M_\sigma(w) = r|n(w)|.$$

The conjugated product therefore obeys the downstream bound. Ordered refactorization cannot produce a first term outside an additive set containing the support of the product: at such a first term the product coefficient is zero, and every contribution already determined from smaller roots also has permitted support. This proves the downstream bound. Subsequent off-cut joints are handled by the same induction as before. If $n(u) = 0$, the incoming line is parallel to the pure cut, the two numerical bounds coincide, and the argument works separately on each side.

Finally, an exponential or inverse at positive non-p degree uses only finitely many decompositions at a fixed such degree. Applying the just proved bound to each factor gives finite p-support. The pure Hamiltonian may have infinitely many p-terms; its action is the explicitly rational map given by the polynomial Q. □

10.3. **Reversal of an exchange polynomial.** Write

$$P(Z) = A + \sum_{j=1}^{r-1} q_j Z^j + BZ^r, \qquad Q = P/A, \qquad \tau = B/A,$$

and define

$$Q^\vee(Z) = \tau^{-1} Z^r Q(Z^{-1}), \qquad S(w) = w - rn(w)p, \qquad S_\tau(X^w) = \tau^{-n(w)} X^{S(w)}. \tag{10.5}$$

Since $n(p) = 0$, S is an integral unimodular shear. It preserves Ω: the two cross terms in $\Omega(Su, Sv)$ cancel because $\Omega(p, w) = \delta n(w)$. The scalar in S_τ is a multiplicative character, so S_τ is a Poisson automorphism of the rational torus. If $\phi_{p,Q}(X^w) = X^w Q(X^p)^{n(w)}$, then

$$S_\tau \phi_{p,Q} = \phi^{-1}_{-p,Q^\vee}. \tag{10.6}$$

Indeed, the primitive normal of $-p$ is $-n$, and both sides send X^w to

$$X^w \big(\tau^{-1} X^{-rp} Q(X^p)\big)^{n(w)}.$$

This is a rational identity for every integer $n(w)$.

Replace the simple-root set by

$$C' = \{-p\} \cup \{c' = c + r[-n(c)]_+ p : c \in C \setminus \{p\}\}. \tag{10.7}$$

It is again part of a lattice basis. Let $L_+ = \mathrm{id}$ and $L_- = S$. For a root satisfying (10.3), direct calculation gives

$$L_\sigma(kp + u) = \sum_{c \neq p} u_c c' + (M_\sigma(u) - k)(-p) \in (C')^+. \tag{10.8}$$

Put $b_0 = r \max_{c \neq p} |n(c)|$, with $b_0 = 0$ if there is no other root. A pushed root of degree at most d requires an old root of degree at most $(b_0 + 1)d$, because $k \leq b_0 |u|_{\neq p}$ and its new degree is at least $|u|_{\neq p}$. Thus these changes of cone make sense in compatible finite truncations.

Lemma 10.2 (The mutated formal system)**.** *Use the identity on the half-space $h(p) > 0$ and S_τ on $h(p) < 0$, with the corresponding dual linear change of the base point. Replace the pure factor by $(-p, Q^\vee)$. The result is the consistent system for the cone C' with incoming functions*

$$Q^\vee \text{ on } -p, \qquad \begin{cases} Q_c(X^{c'}), & n(c) \geq 0, \\ Q_c(\tau^{-n(c)} X^{c'}), & n(c) \leq 0 \end{cases} \text{ on } c'. \tag{10.9}$$

All other incoming coefficients vanish.

Proof. Choose a lattice complement to p containing the other simple roots, and put $Z = X^p$. Work over $K(Z)$, completing only in the positive monoid generated by the other roots modulo p; allow finite unions of support translates for transported expressions. At a fixed projected exponent the coefficient is a rational function of Z. Lemma 10.1 makes all nonpure operations and their inverses well-defined in this ring. The pure operation is rational in Z, and S_τ fixes Z and the projected exponents.

Old consistency holds in this common ring. To verify it, expand $K(Z)$ injectively into $K((Z))$. At a fixed non-p degree the nonpure roots are finite by (10.3); a further prescribed coefficient of Z is tested in a finite full-root truncation. Old consistency holds there. Hence a rational difference of two path products has zero Laurent expansion coefficientwise and is zero. Equation (10.6) now gives consistency of the pushed system in the common rational ring. Expanding at $Z^{-1} = 0$, with (10.8), places it in the C'-completion.

It remains to check the incoming data, including their vanishing away from simple roots. Symplecticity gives

$$V(Sw) = V(w) \circ S^{-1}, \qquad n(Sw) = n(w).$$

An incoming direction for c therefore lies on the identity piece if $n(c) > 0$, and on the S_τ piece if $n(c) < 0$. This produces exactly the two functions on c' in (10.9). The pure incoming function is $Q^\vee$ by (10.6). An effective root a' not on a new simple ray has no incoming coefficient. To check this explicitly, suppose first that $n(a') > 0$. Its incoming direction lies on the identity piece, so its inverse image is a'. An old simple root on this piece has $n(c) > 0$, and is precisely the new simple root $c' = c$. Thus a' has an inverse image on no old simple ray, or has no inverse image in the old positive support. The coefficient vanishes in either case. If $n(a') < 0$, use the inverse image $S^{-1}a'$. The sign of its n-value is unchanged. The old simple roots on this piece map exactly to the new simple roots $c' = S(c)$, so the same argument applies. This also checks every positive multiple of the root.

Directions with $n(a') = 0$ lie along the cut. Both pieces fix their exponents and their coefficient characters. At the possibly nongeneric incoming point, the coefficient functional for a' annihilates brackets supported in $\ker V(a')$, as in the proof of Lemma 9.2. Refinement to a small displacement on either side therefore preserves that coefficient. An inverse image on an old simple ray lies on a new simple ray as well; all remaining coefficients are zero. Uniqueness in that lemma identifies the whole pushed system. □

10.4. **Incoming coefficients across a mutation.**

Proposition 10.3. *Let F_{old} and F_{new} be finite Laurent polynomials, taken as values in the all-negative chambers of the systems for C and C', respectively. Suppose, in the common rational torus, that*

$$F_{\mathrm{old}} = \phi_{p,Q}(F_{\mathrm{new}}).$$

Then

$$\mu(w) = w + r[-n(w)]_+ p, \qquad A'_{\mu(w)}(F_{\mathrm{new}}) = \tau^{[-n(w)]_+} A_w(F_{\mathrm{old}}). \tag{10.10}$$

The label map μ is a bijection.

Proof. Compare the equality in the common rational ring of the preceding proof. Choose $h(c) < -1$ for $c \neq p$, and $|h(p)| < (b_0 + 1)^{-1}$. For every supported nonpure root $kp + u$, the estimate $k \le b_0|u|_{\neq p}$ shows that $h(kp + u) < 0$. A short path across the pure cut here meets no other wall. The old value on its positive side is the finite polynomial F_{new}. On its negative side the pushed value is

$$S_\tau(F_{\mathrm{old}}) = \phi^{-1}_{-p,Q^\vee}(F_{\mathrm{new}}).$$

Thus these are the values of the new transported expression on the two sides.

Continue to any required point within one open half-space. Only nonpure operations are crossed. By Lemma 10.1, at each fixed non-p exponent they take a finite Laurent p-fiber to a finite Laurent p-fiber. Therefore both descriptions read the same finite Laurent polynomial in Z. Its coefficients are independent of which of the two completions is used.

If $n(w) > 0$, the incoming direction of w lies on the identity piece, so its label and coefficient remain unchanged. If $n(w) < 0$, its label is $S(w)$ and its coefficient is multiplied by $\tau^{-n(w)}$. When $n(w) = 0$, the incoming direction lies on the cut. The pure operation fixes every potential source $w - jp$, since $n(w - jp) = 0$. Every wall through that incoming point also preserves the coefficient of X^w, by the argument for (9.4). We may therefore compute it at a small displacement on either side. Both piecewise maps fix its label and coefficient. These cases give (10.10). Since $n(\mu(w)) = n(w)$, the inverse label map is $v \mapsto v - r[-n(v)]_+ p$. □

10.5. **The complete generalized mutation rule.** We verify that the formal mutation just constructed is the mutation of Section 2, including its coefficient rule. Fix k, and put $A = q_{k,0}$, $B_* = q_{k,r_k}$, and $\tau = B_*/A$. Define

$$Ee_k = -e_k + r_k[\widehat{b}_k]_+, \qquad Ee_i = e_i \quad (i \neq k), \qquad E_A(X^w) = A^{w_k} X^{Ew}. \tag{10.11}$$

Here frozen basis vectors are included among the fixed vectors. The diagonal entry $b_{kk} = 0$ gives $E^2 = 1$. Evaluation on each generator proves

$$\mu_k^* = \phi_{p_k,Q_k} E_A. \tag{10.12}$$

For the k-th generator, the right side is

$$\frac{A}{x_k}(M_k^+)^{r_k} Q_k(M_k^-/M_k^+),$$

which is precisely (2.1); the other generators are fixed.

Let F be the $n \times n$ identity matrix except for

$$F_{kk} = -1, \qquad F_{kj} = r_k[-b_{kj}]_+ \quad (j \neq k).$$

The matrix mutation (2.2) reads

$$\widehat{B}' = E\widehat{B}F. \tag{10.13}$$

Its new roots, pulled back by E, are consequently

$$Ep'_k = -p_k, \qquad Ep'_j = p_j + r_k[-b_{kj}]_+ p_k \quad (j \neq k).$$

For $j \neq k$, put $\alpha = b_{kj}$. The coefficient character in $E_A(X^{p'_j})$ is A^α, since the k-th coordinate of p'_j is α. Recalling (2.3), we have

$$\bar{q}'_{j,s} = \begin{cases} \bar{q}_{j,s} A^{-s\alpha}, & \alpha \geq 0, \\ \bar{q}_{j,s} B_*^{-s\alpha}, & \alpha \leq 0, \end{cases} \qquad \bar{q}_{j,s} = q_{j,s}/q_{j,0}.$$

After substitution of $A^\alpha X^{Ep'_j}$, this is the old input function if $\alpha \geq 0$, and the old function evaluated at $\tau^{-\alpha} X^{Ep'_j}$ if $\alpha \leq 0$. At k, reversal of the coefficient tuple and division by its new constant term gives $Q_k^\vee$. These are exactly (10.9).

The new form is $\Omega' = E^t \Omega E$. For $i \neq k$, computation from (10.13) gives

$$\begin{aligned} \Omega'(p'_i, w) &= \Omega(p_i + r_k[-b_{ki}]_+ p_k, Ew) \\ &= \delta_i n_i(w) + r_k(\delta_i[b_{ik}]_+ - \delta_k[-b_{ki}]_+) n_k(w). \end{aligned}$$

The last coefficient vanishes since DB is alternating. For $i = k$, both the root and the corresponding coordinate change sign. Hence the new data still satisfy (10.2). This proves compatibility for arbitrary symmetrizers. It also proves that the construction can be iterated with all frozen rows retained.

This monomial change also identifies incoming coefficients. The equality $\Omega' = E^t \Omega E$ carries an incoming direction to the corresponding incoming direction in the pulled-back system. For an exponent w in the new seed, E_A carries its coefficient to A^{w_k} times that coefficient at Ew. Thus the monomial change is a bijection on incoming labels, with endpoint-monomial multipliers. Combining this observation with Proposition 10.3 gives the full incoming-coefficient transport for mutation.

10.6. **Arbitrary coefficient representatives.** The freedom in the common coefficient scales is accounted for by the following elementary rescaling statement.

Lemma 10.4. *Two representatives rooted at the same seed differ, at every seed, by rescaling each cluster variable by an element of the coefficient group. The same holds for their cluster monomials.*

Proof. Denote the two representatives by superscripts a and c, and suppose at one seed that, for coefficient-group elements t_i,

$$x_i^{\mathrm{a}} = t_i x_i^{\mathrm{c}}, \qquad \frac{\bar{q}_{j,s}^{\mathrm{a}}}{\bar{q}_{j,s}^{\mathrm{c}}} = \prod_i t_i^{sb_{ij}}. \tag{10.14}$$

The normalized exchange functions then agree after substitution. Writing $\lambda_k = q_{k,0}^{\mathrm{a}}/q_{k,0}^{\mathrm{c}}$, the exchange relation gives

$$t_i' = t_i \quad (i \neq k), \qquad t_k' = \lambda_k t_k^{-1} \prod_i t_i^{r_k[b_{ik}]_+}. \tag{10.15}$$

We verify that (10.14) persists. For $j \neq k$, put $\alpha = b_{kj}$. Using (10.13),

$$\prod_i (t_i')^{sb_{ij}'} = \left(\prod_i t_i^{sb_{ij}}\right) \lambda_k^{-s\alpha} \prod_i t_i^{sr_k b_{ik}[-\alpha]_+}. \tag{10.16}$$

If $\alpha \geq 0$, the last product is one and the constant-endpoint coefficient rule gives the right side of (10.16). If $\alpha \leq 0$, use

$$\frac{q_{k,r_k}^{\mathrm{a}}}{q_{k,r_k}^{\mathrm{c}}} = \lambda_k \prod_i t_i^{r_k b_{ik}};$$

the leading-endpoint rule gives exactly the same right side. At k, reversing the coefficient tuple gives

$$\frac{\bar{q}_{k,s}'^{\,\mathrm{a}}}{\bar{q}_{k,s}'^{\,\mathrm{c}}} = \frac{\bar{q}_{k,r_k-s}^{\mathrm{a}}/\bar{q}_{k,r_k-s}^{\mathrm{c}}}{q_{k,r_k}^{\mathrm{a}}/q_{k,r_k}^{\mathrm{c}}} = \prod_i t_i^{-sb_{ik}} = \prod_i (t_i')^{sb_{ik}'},$$

because $b_{kk} = 0$ and the k-th column changes sign. Thus the invariant propagates, independently of the allowed new uniform coefficient scalings.

Starting with the same reference seed, take $t_i = 1$. Every variable of an arbitrary representative is therefore a coefficient-group monomial times the variable of the canonical representative. The same holds for cluster monomials. This proves the lemma. □

It follows that positivity in the universal canonical representative implies positivity in $\mathbb{N}[\mathbb{P}]$ for all representatives: specialize the universal coefficient symbols to their elements of $\mathbb{P}$ and multiply by the rescaling monomial.

11. Laurent positivity

Lemma 11.1. *In its own seed, a cluster monomial cX^m, with $m_i \geq 0$ in mutable coordinates and with c a coefficient-group monomial, has incoming coefficient c at m and zero at every other exponent.*

Proof. By (10.2), $\Omega(p_i, m) = -\delta_i m_i \leq 0$. All transported values have support in $m + C^+$. Fix $w = m + r$ with $0 \neq r \in C^+$, and start in the all-negative chamber at a generic h_-. Follow $h_- + sV(w)$, $s \geq 0$, towards the incoming direction of w. A crossed root a must satisfy $\Omega(a, w) > 0$, since its value starts negative and can change sign only with positive slope. Any nonidentity contribution from m to w would require

$$r = \sum_j t_j a_j, \qquad t_j > 0,$$

with roots crossed by this path. It would follow that $\Omega(r, w) > 0$. But

$$\Omega(r, w) = \Omega(r, m + r) = \Omega(r, m) \leq 0,$$

a contradiction. The calculation may be made in the finite interval between m and w. The coefficient at m remains c, because a sum of nonzero roots in C^+ cannot be zero. □

Proof of Theorem 2.1. Choose the reference seed and a finite mutation word leading to the seed of the given cluster monomial. Add the auxiliary frozen rows and variables used in (10.1). The generalized Laurent phenomenon, Theorem 2.2, supplies a finite Laurent expression for the monomial in every intermediate seed.

In the terminal seed, write the cluster monomial as X^m, including its frozen factor in m. Lemma 11.1 gives the single incoming coefficient one at m. For each adjacent pair, the monomial map E_A places their expressions in one lattice, and (10.12) gives the rational relation required in Proposition 10.3. The incoming labels are carried by its bijection, and the scalar factors are Laurent monomials in the two endpoint coefficients. Thus, going back along the word, the incoming data in the reference seed still consist of a single positive coefficient. Theorem 9.5 and Proposition 9.6 prove positivity of every coefficient in the reference-seed Laurent expression.

The coefficient semiring is the fixed one in (2.4). In the canonical representative, $q'_{j,0} = q_{j,0}$ for $j \neq k$. Reversal interchanges the endpoint coefficients at the mutated vertex; the other endpoint rules use only endpoint monomials. Every intermediate coefficient is either an initial interior coefficient, possibly with its index reversed, multiplied by an endpoint Laurent monomial, or is itself an endpoint Laurent monomial. The same is true of all scalar factors in (10.10). Consequently every positive operation above takes place in that fixed semiring. If the initial endpoints are one, no initial interior coefficient is inverted.

Finally set $Y_1 = \cdots = Y_n = 1$. This is a homomorphism on the finite Laurent ring already obtained, and it recovers the original mutation word and monomial. Coefficients of coincident monomials are added, preserving positivity. This proves the canonical-representative assertion. Lemma 10.4, followed by specialization to $\mathbb{P}$, proves the assertion for every allowed coefficient representative. □

12. Finite algebras for the coefficients

We realize the coefficients in Proposition 8.1 as dimensions of explicitly constructed finite algebras. We work over $\mathbb{Q}$; dimensions are preserved by finite extensions used to split points. We allow the zero algebra, of dimension zero.

12.1. **A finite algebra for an elementary coefficient.** Fix a, b, m, k, p, q as in Section 3, and put $A = ka$, $B_0 = kb$, and $M = 1 + \max(a, b)$. Choose polynomials P_A, P_B with rational coefficients, of degrees at most $A - 1$ and $B_0 - 1$, respectively. Perturb the Gale equations to

$$\begin{aligned} a_i^p \prod_j (a_i - b_j)^m - T^b + T^M P_A(a_i) &= 0, \\ (-b_j)^q \prod_i (a_i - b_j)^m - T^a + T^M P_B(b_j) &= 0. \end{aligned} \tag{12.1}$$

The coefficients can be chosen so that the full zero scheme over $\overline{\mathbb{Q}(T)}$ is finite and reduced. Indeed, over $\mathbb{Q}(T)$ the factor T^M is a unit, so the same value and first-derivative interpolation used for the Young deformation applies. The nonempty open condition on the coefficient parameters contains a rational point.

The attached factor of (12.1) has the same ungraded G_k-representation as $\mathcal{D}_k^{p,q}$. Insert a parameter s before both perturbations, localize at all particle coordinates, take T-saturation, and then complete in T. At a special point, normalize a curve from the locus $T \neq 0$. All coordinates are integral on that curve. The right sides of the original product equations now have the form

$$T^b(1 - sT^{M-b}P_A(a_i)), \qquad T^a(1 - sT^{M-a}P_B(b_j)).$$

The parenthesized factors have residue one. Multiplication of the equations in any equal-center cluster therefore gives exactly the balance and center equations proved in Section 3. Every

nonzero center lies in the same fixed finite root set of

$$P_k(z)=\prod_{h=1}^{k}(z^h-\sigma_k^h),\qquad \sigma_k=(-1)^{bq+mab(k-1)}.$$

The inverted coordinates take only nonzero values there. Their inverses are consequently integral on the reduced special support, and then integral on the whole special algebra by nilpotence of its nilradical. The special algebra is finite over $\mathbb{Q}[s]$.

Lemma 5.1 makes the completed algebra a finite direct factor over $\mathbb{Q}[s][[T]]$. After T is inverted, it is finite locally free over the connected regular one-dimensional ring $\mathbb{Q}[s][[T]][T^{-1}]$, so its isotypic ranks are constant.

Its endpoint fibers are the complete attached factors. Indeed, after inverting T, the equations are the displayed square system with the perturbations multiplied by s. At $s=0,1$, consider a full Artin local solution whose reduced particle coordinates have valuation zero. The coordinates and their inverses are integral over a finite extension of $\mathbb{Q}[[T]]$: a monic equation for a reduced coordinate becomes a monic equation in the local algebra after taking a power that kills its nilradical. Thus the whole local solution extends to the localized completion. Conversely, the center calculation gives exactly these valuations on the selected factor. Completion commutes with the endpoint quotients, and the direct factor retains every local Artin length. The fibers at $s=0$ and $s=1$ therefore have the same character.

For partitions $\lambda=(\lambda_1,\dots,\lambda_r)\vdash A$ and $\mu=(\mu_1,\dots,\mu_t)\vdash B_0$, fix labeled blocks of these sizes and give them coordinates $u_1,\dots,u_r$ and $v_1,\dots,v_t$. Let $I_{\lambda\mu}$ be the ideal in $\mathbb{Q}[T,u,v]$ generated by

$$\begin{aligned} u_i^p\prod_{j=1}^{t}(u_i-v_j)^{m\mu_j}-T^b+T^MP_A(u_i)&=0,\quad 1\le i\le r,\\ (-v_j)^q\prod_{i=1}^{r}(u_i-v_j)^{m\lambda_i}-T^a+T^MP_B(v_j)&=0,\quad 1\le j\le t. \end{aligned}\tag{12.2}$$

Set

$$\Delta=\prod_{i<i'}(u_i-u_{i'})\prod_{j<j'}(v_j-v_{j'}),\qquad Z=\prod_i u_i\prod_j v_j,$$

and define, in the order displayed,

$$J_{\lambda\mu}=I_{\lambda\mu}:(T\Delta)^{\infty},\qquad K_{\lambda\mu}=(J_{\lambda\mu}\bmod T):Z^{\infty},\qquad R_{\lambda\mu}=\mathbb{Q}[u,v]/K_{\lambda\mu}.\tag{12.3}$$

The first saturation selects the exact equality pattern on the generic fiber and takes its full horizontal closure. The second selects nonzero centers. Distinct generic blocks can collide at $T=0$; taking their horizontal closure before selecting nonzero centers retains the resulting lengths.

Let $m_d(\lambda)$ denote the multiplicity of the part d in λ, and define

$$W_{\lambda\mu}=\prod_d S_{m_d(\lambda)}\times\prod_d S_{m_d(\mu)}.$$

It permutes blocks of equal size and acts on $R_{\lambda\mu}$.

Proposition 12.1. *The algebra $R_{\lambda\mu}$ is finite, and the coefficient of $e_\lambda(\alpha)e_\mu(\beta)$ in the local wall coefficient of charge $k(a,b)$ is*

$$n_{\lambda\mu}=\dim_{\mathbb{Q}}R_{\lambda\mu}^{W_{\lambda\mu}}.\tag{12.4}$$

Proof. Over a splitting field of $\mathbb{Q}((T))$, the perturbed attached factor is a finite reduced G_k-set. Each G_k-orbit whose point stabilizers are conjugate to $S_\lambda\times S_\mu$ meets the chosen exact labeled block stratum in one orbit of the normalizer quotient $N_{G_k}(S_\lambda\times S_\mu)/(S_\lambda\times S_\mu)=W_{\lambda\mu}$. This

action is free on the exact block stratum. Thus its invariant function space has dimension one for every such G_k-orbit.

After the first saturation in (12.3), T is a nonzerodivisor. Complete

$$\mathbb{Q}[T, u, v, Z^{-1}]/J_{\lambda\mu}$$

in T. Its special fiber is the localization of $\mathbb{Q}[u, v]/(J_{\lambda\mu} \bmod T)$ at Z. The center calculation makes its selected support finite, with Z nonzero at every point. The image in this finite-dimensional localization has injective, hence bijective, multiplication by Z. Thus localization equals the quotient by Z-torsion, namely

$$\mathbb{Q}[u, v]/((J_{\lambda\mu} \bmod T) : Z^{\infty}) = R_{\lambda\mu}.$$

The generator-lifting argument of Lemma 5.1 makes the completion finite over $\mathbb{Q}[[T]]$. Completion preserves the nonzerodivisor T, so this finite module is torsion-free over the discrete valuation ring $\mathbb{Q}[[T]]$, hence free. Characteristic-zero averaging makes the invariant submodule a direct summand; it commutes with base change and has the same rank in the two fibers, including all nonreduced special lengths. Therefore its special dimension equals the number of generic G_k-orbits of the indicated Young type. The character identity in Proposition 8.1 identifies that number with $n_{\lambda\mu}$. □

The invariant algebra itself has explicit equations. For each group of equal λ-parts adjoin variables for the elementary symmetric functions of the corresponding u_i; do the same for the equal μ-parts and the v_j. Eliminate u, v from $K_{\lambda\mu}$ together with these defining equations. The resulting quotient in the new variables is $R_{\lambda\mu}^{W_{\lambda\mu}}$. To see surjectivity, average any lift of an invariant residue class. The invariant polynomial ring for each permutation block is generated by its elementary symmetric functions. Elimination gives exactly the kernel.

12.2. **Effective choices.** The generic perturbations can be chosen by a terminating rational procedure. Enumerate integer coefficient tuples for P_A, P_B. For each tuple, use a Gröbner basis over $\mathbb{Q}(T)$ to test

$$(F_1, \dots, F_{A+B_0}, \det \operatorname{Jac} F) = (1),$$

where F denotes the left sides of (12.1). A successful test says that every affine zero has invertible Jacobian, so the zero scheme is finite and reduced. The nonempty generic open set contains an integer tuple: after clearing denominators, a nonzero polynomial in T and the parameters has a nonzero coefficient when viewed as a polynomial in T, and the integer lattice is Zariski dense. Enumeration therefore terminates. Saturation, elimination, finite standard monomial bases, invariant subspaces, and their multiplication tables are then finite rational computations.

12.3. **A finite formula for a whole Laurent expansion.** Fix a mutation word and a cluster monomial in its final seed, using the universal canonical representative of Theorem 2.1. Append the auxiliary frozen variables $Y_1, \dots, Y_n$ used in the global proof. Let $\widehat{F}$ be the resulting cluster monomial, expressed in the initial variables. The initial roots p_i have auxiliary coordinates $(p_i)_{Y_j} = -\delta_{ij}$. Transporting the unique incoming monomial backwards determines an exponent g and an endpoint coefficient monomial η. For a mutation at k with new exponent u, use the current old-seed data and the map E from (10.11). The inverse update is

$$g_{\mathrm{old}} = Eu - r_k[-u_k]_+ p_k, \qquad \eta_{\mathrm{old}} = \eta_{\mathrm{new}} q_{k,0}^{[u_k]_+} q_{k,r_k}^{-[-u_k]_+}. \tag{12.5}$$

It follows directly by inverting the incoming-coefficient identity in Section 10; in particular all factors in η are endpoint monomials.

The positive recurrence implies that every transported support lies in $g + \sum_i \mathbb{N} p_i$. Start with the finite Laurent expression supplied by Theorem 2.2. Its transports have support in a finite

union of translates of the free positive root cone. In each finite interval below a target exponent outside $g+\sum_i \mathbb{N}p_i$, the incoming coefficient and every predecessor outside that cone are zero. Induction in the recurrence proves the assertion.

Compute the lowest valuation of the final rational expression in each Y_j directly from the exchange word. All these expressions are subtraction-free in independent positive variables. Lowest terms of summands therefore cannot cancel. Let $\widehat{b}_{ik}$ denote the current auxiliary extended exchange matrix. If ℓ_i are the current Y_j-valuations of all seed variables, including the original frozen variables and the auxiliary Y's, mutation at k gives

$$\ell'_k = -\ell_k + r_k \min\left\{\sum_i [\widehat{b}_{ik}]_+\ell_i, \sum_i [-\widehat{b}_{ik}]_+\ell_i\right\}. \tag{12.6}$$

Endpoint coefficients have Y_j-valuation zero, and the minimum over the exchange sum is attained at an endpoint, which proves the formula. Let L_j be the resulting valuation of the final cluster monomial. By its Laurent finiteness this is the minimum Y_j-exponent.

An exponent $g+\sum_i k_i p_i$ in its support consequently satisfies

$$0 \le k_j \le K_j := g_{Y_j} - L_j. \tag{12.7}$$

This gives a finite box using only the matrix and coefficient mutations, the incoming update, and the valuation recursion.

For $1 \le i \le n$ and $1 \le s \le r_i$, let $\bar{q}_{i,s}$ denote the initial coefficient ratios defined in Section 2. Write $\mathbf{k} = (k_1, \dots, k_n)$, $\mathbf{K} = (K_1, \dots, K_n)$, and $\beta = (\beta_{i,s})_{1\le i\le n,\, 1\le s\le r_i}$, with componentwise inequalities for these indices. There are explicitly constructible finite $\mathbb{Q}$-algebras $\mathcal{A}_{\mathbf{k},\beta}$ such that

$$\widehat{F} = \eta \sum_{0\le \mathbf{k}\le \mathbf{K}} X^{g+\sum_i k_i p_i} \sum_{\substack{\beta_{i,s}\ge 0 \\ \sum_s s\beta_{i,s} = k_i \ (1\le i\le n)}} \dim_{\mathbb{Q}} \mathcal{A}_{\mathbf{k},\beta} \prod_{i,s} \bar{q}_{i,s}^{\beta_{i,s}}. \tag{12.8}$$

Keep every initial coefficient separately labeled. Construct the wall operations only through total root degree $\sum_j K_j$. Replace each two-input coefficient by the algebra of Proposition 12.1. Products of coefficients are realized by tensor products, and finite sums by direct products of finite algebras; the coefficients zero and one are represented by the zero algebra and $\mathbb{Q}$. Integral powers of a wall function use ordered slot choices, with one tensor product for each choice. The finite labeled perturbation construction and the positive recurrence use precisely these operations. The wall construction proceeds through finitely many root degrees; at each degree it crosses finitely many generic joints, with the local positive expressions constructed by the uniform induction on label weight in Proposition 9.4. Once these wall coefficients are fixed, the recurrence of Proposition 9.6 strictly decreases root degree. These three finite procedures construct $\mathcal{A}_{\mathbf{k},\beta}$ in finitely many steps.

All geometric choices can be made over $\mathbb{Q}$. At each finite stage, the directions and supports are rational and finite in number. Choose rational original-line translations with distinct parallel lines and no triple original intersection, and compute subsequent joints by rational linear algebra. The local construction includes multiple generated intersections and overlapping rays. In higher rank, choose rational generic points and displacements in the required cells, avoiding the finitely many forbidden proper linear subspaces. Enumerating rational candidates and testing these conditions by exact linear algebra terminates.

The displayed weight constraint follows because the initial coefficient $\bar{q}_{i,s}$ has root weight sp_i.

Finally set $Y_1 = \cdots = Y_n = 1$ in the finite expression (12.8). If several lifted exponents project to the same old exponent, sum their finitely many coefficient algebras by a direct product. Thus column dependencies, nonprimitive columns, and zero projected roots are covered. The rational

algorithm constructs the coefficient algebras over the universal independent parameters; positive specialization then gives the coefficient-group statement.

References

[1] R. Akagi and T. Nakanishi, *Relation between generalized and ordinary cluster algebras*, preprint, arXiv:2512.21062v2, 2026. https://arxiv.org/abs/2512.21062v2.

[2] F. Bihan and F. Sottile, *Gale duality for complete intersections*, Ann. Inst. Fourier (Grenoble) **58** (2008), no. 3, 877–891. doi:10.5802/aif.2372.

[3] H. Bruggesser and P. Mani, *Shellable decompositions of cells and spheres*, Math. Scand. **29** (1971), 197–205. doi:10.7146/math.scand.a-11045.

[4] A. Burcroff, K. Lee, and L. Mou, *Scattering diagrams, tight gradings, and generalized positivity*, Proc. Natl. Acad. Sci. USA **122** (2025), no. 18, e2422893122. doi:10.1073/pnas.2422893122. Announcement preprint: https://arxiv.org/abs/2409.15235v1, 2024.

[5] A. Burcroff, K. Lee, and L. Mou, *Positivity of generalized cluster scattering diagrams*, preprint, arXiv:2503.03719v1, 2025. https://arxiv.org/abs/2503.03719v1.

[6] A. Burcroff, K. Lee, L. Mou, G. Musiker, and M. Reineke, *Cluster scattering diagrams via quiver moduli and tight gradings*, J. Lond. Math. Soc. (2) **114** (2026), no. 2, e70642. doi:10.1112/jlms.70642.

[7] L. Chekhov and M. Shapiro, *Teichmüller spaces of Riemann surfaces with orbifold points of arbitrary order and cluster variables*, Int. Math. Res. Not. IMRN (2014), no. 10, 2746–2772. doi:10.1093/imrn/rnt016. Corrected preprint: https://arxiv.org/abs/1111.3963v3.

[8] M.-W. Cheung, E. Kelley, and G. Musiker, *Cluster scattering diagrams and theta functions for reciprocal generalized cluster algebras*, Ann. Comb. **27** (2023), no. 3, 615–691. doi:10.1007/s00026-022-00623-1.

[9] S. Fomin and A. Zelevinsky, *Cluster algebras I: Foundations*, J. Amer. Math. Soc. **15** (2002), no. 2, 497–529. doi:10.1090/S0894-0347-01-00385-X.

[10] S. Fomin and A. Zelevinsky, *The Laurent phenomenon*, Adv. in Appl. Math. **28** (2002), no. 2, 119–144. doi:10.1006/aama.2001.0770.

[11] S. Fomin and A. Zelevinsky, *Cluster algebras IV: Coefficients*, Compos. Math. **143** (2007), no. 1, 112–164. doi:10.1112/S0010437X06002521.

[12] H. Franzen, *Chow rings of fine quiver moduli are tautologically presented*, Math. Z. **279** (2015), no. 3–4, 1197–1223. doi:10.1007/s00209-015-1409-8.

[13] H. Franzen and M. Reineke, *Semistable Chow–Hall algebras of quivers and quantized Donaldson–Thomas invariants*, Algebra Number Theory **12** (2018), no. 5, 1001–1025. doi:10.2140/ant.2018.12.1001. https://arxiv.org/abs/1512.03748v2.

[14] E. González and C. T. Woodward, *Quantum cohomology and toric minimal model programs*, Adv. Math. **353** (2019), 591–646. doi:10.1016/j.aim.2019.07.004. https://arxiv.org/abs/1207.3253v9.

[15] M. Gross, P. Hacking, S. Keel, and M. Kontsevich, *Canonical bases for cluster algebras*, J. Amer. Math. Soc. **31** (2018), no. 2, 497–608. doi:10.1090/jams/890.

[16] K. Lee, L. Li, and A. Zelevinsky, *Greedy elements in rank 2 cluster algebras*, Selecta Math. (N.S.) **20** (2014), no. 1, 57–82. https://arxiv.org/abs/1208.2391v3.

[17] K. Lee and R. Schiffler, *Positivity for cluster algebras*, Ann. of Math. (2) **182** (2015), no. 1, 73–125. doi:10.4007/annals.2015.182.1.2.

[18] L. Mou, *Scattering diagrams for generalized cluster algebras*, Algebra Number Theory **18** (2024), no. 12, 2179–2246. doi:10.2140/ant.2024.18.2179.

[19] T. Nakanishi, *Structure of seeds in generalized cluster algebras*, Pacific J. Math. **277** (2015), no. 1, 201–218. doi:10.2140/pjm.2015.277.201.

[20] T. Nakanishi and D. Rupel, *Companion cluster algebras to a generalized cluster algebra*, Travaux Math. **24** (2016), 129–149. https://arxiv.org/abs/1504.06758v1.

[21] OpenAI, *Elementary positivity of chromatic quasisymmetric functions*, preprint, September 24, 2026. OpenAI mathematical preprints. Version at repository commit `adc7f1241b42e322a6451854ab7e4b4c146bf78a`.

[22] R. Ramos and D. Whiting, *Generalized cluster algebras are subquotients of cluster algebras*, preprint, arXiv:2504.20931v3, 2025. https://arxiv.org/abs/2504.20931v3.

[23] M. Reineke and T. Weist, *Refined GW/Kronecker correspondence*, Math. Ann. **355** (2013), no. 1, 17–56. doi:10.1007/s00208-012-0778-0.

[24] C. R. Rimmer, *Finite polynomial maps and G-variant map germs*, Ph.D. thesis, University of Liverpool, December 1995. https://livrepository.liverpool.ac.uk/3175815/.

[25] D. Rupel, *Greedy bases in rank 2 generalized cluster algebras*, preprint, arXiv:1309.2567v1, 2013. https://arxiv.org/abs/1309.2567v1.

[26] The Stacks Project Authors, *The Stacks Project*, 2026, Tags 0CM1, 0BNH, 00LJ, 02JN, 00EH, 05KK, 0208, and 00J9. https://stacks.math.columbia.edu.

[27] H. Ye, *Quantum generalized cluster algebras from scattering diagrams: Positivity and mutation*, preprint, arXiv:2611.00991v1, 2026. https://arxiv.org/abs/2611.00991v1.

Zhili College, Tsinghua University, Beijing 100084, China

Email address: tangqy24@mails.tsinghua.edu.cn