hexagon:2611.00239v3

ANOTHER PROOF OF LAURENT POSITIVITY FOR SKEW-SYMMETRIC QUANTUM CLUSTER ALGEBRAS

QIYUE TANG

Abstract. We give another proof of Laurent positivity for skew-symmetric quantum cluster algebras.

Contents

1. Introduction 2
1.1. The positivity theorem 3
1.2. Finite coefficient spaces 4
1.3. The proof 4
2. Quantum seeds and Laurent expressions 5
2.1. Based tori and normalization 5
2.2. Compatibility and mutation 5
2.3. The abstract parameter 7
3. Polynomial shuffle algebras and ordered decomposition 7
3.1. The shuffle product and its quotients 7
3.2. Separated restrictions 9
3.3. The Hilbert-series identity 10
4. A regular coproduct and primitive strings 11
4.1. The separated-center filtration 11
4.2. Product symbols 12
4.3. Construction of the coproduct 13
4.4. Signed PBW and strings 14
5. Positive factors and effective bounds 16
5.1. Completion of the finite-input proof 16
5.2. Admissible infinite inputs 16
5.3. Finite linear-algebra realization 17
6. Finite wall systems and positive coefficients 17
6.1. The total element and its incoming prescription 18
6.2. Positive crossing formulas 19
6.3. Incoming coefficients and a decreasing recursion 20
7. Change of roots and equality of Laurent coefficients 21
7.1. A uniform bound along the mutation direction 22
7.2. Comparison in a common rational ring 23
7.3. A monomial with one incoming coefficient 25
8. From root transport to arbitrary quantum seeds 25
8.1. Central coefficients and an integral root model 25

Date: November 9, 2026.
2020 *Mathematics Subject Classification.* Primary 13F60; Secondary 16G20, 16T30.
Key words and phrases. Quantum cluster algebra, Laurent positivity, polynomial shuffle algebra, primitive element, scattering diagram.

8.2. The exact exchange relation 26
8.3. Transport of an actual cluster monomial 27
8.4. Specialization and several components 27
9. Finite graded spaces for all coefficients 28
9.1. The incoming label and exact auxiliary valuations 29
9.2. A finite graph of local reorderings 30
9.3. A global energy bound 31
9.4. The finite polynomial matrices 32
9.5. Decorated histories and specialization 33
References 35

1. Introduction

Fomin and Zelevinsky introduced cluster algebras to study the coordinate changes underlying total positivity and canonical bases [8]. A seed consists of a set of algebraically independent variables and exchange data; mutation replaces one variable by the quotient of a sum of two monomials by that variable. Their Laurent phenomenon [8, Theorem 3.1] states that every variable obtained by repeated mutation is a Laurent polynomial in every seed. Immediately after that theorem they conjectured that the Laurent coefficients are nonnegative. Repeated exchange relations give rational expressions involving sums, products and inverses. The Laurent phenomenon reduces the denominators to monomials, while the positivity conjecture concerns the coefficients in the resulting Laurent expansions.

Berenstein and Zelevinsky introduced quantum cluster algebras in [1]. Their compatible pairs couple an exchange matrix with an integral alternating form. The variables in a quantum seed commute up to powers of a parameter, and centered monomials form a distinguished basis of its quantum torus. They proved the quantum Laurent phenomenon [1, Corollary 5.2] and constructed the bar involution [1, Section 6]. Write $\mathbb{N} = \{0, 1, 2, \ldots\}$ and $v = q^{1/2}$. Quantum positivity asks that the coefficients in the centered basis belong to $\mathbb{N}[v^{\pm 1}]$. Specialization at $v = 1$ gives classical positivity.

Early results connected positivity with explicit bases and expansion formulas. Sherman and Zelevinsky constructed positive canonical bases for rank-two cluster algebras of finite and affine types [21]. Musiker, Schiffler and Williams gave Laurent expansion formulas for cluster algebras from triangulated surfaces, including punctured surfaces, and proved positivity in that setting [15]. Lee and Schiffler proved positivity for every skew-symmetric cluster algebra [14]. Gross, Hacking, Keel and Kontsevich constructed classical theta functions by broken lines in scattering diagrams and proved positivity in their theory of canonical bases [10]. In this description, a Laurent coefficient is obtained by recording the monomial changes along a broken line, and positive wall functions give positive contributions to the expansion.

Geometric methods also gave graded interpretations of quantum coefficients. Nakajima related cluster monomials to simple modules and perverse sheaves on graded quiver varieties, obtaining classical positivity for cluster algebras with a bipartite seed [17]. Qin expressed quantum cluster monomials and quantum F-polynomials in terms of Serre polynomials of quiver Grassmannians of rigid modules; this proved quantum positivity for expansions in acyclic seeds [19]. Kimura and Qin subsequently used graded quiver varieties and deformed Grothendieck rings to prove quantum positivity throughout the mutation class of an acyclic seed [12].

Donaldson–Thomas theory supplied a way to study arbitrary mutation sequences. Nagao related compositions of cluster transformations to wall crossing and quiver Grassmannians [16]. Efimov expressed quantum cluster variables through vanishing-cycle cohomology and reduced

positivity to a purity statement; he proved positivity and a Lefschetz property when the initial or final seed is acyclic [6]. Davison, Maulik, Schürmann and Szendrői proved the required purity for positively graded nondegenerate potentials [4]. Davison then proved quantum Laurent positivity for every skew-symmetric quantum cluster algebra [2, Theorem 2.4]. His purity theorem for vanishing-cycle cohomology supplies the nonnegative coefficients [2, Theorem 5.2 and Section 6]. The result also gives a Lefschetz property and a parity constraint on the coefficients.

Davison and Mandel developed quantum scattering diagrams and quantum theta bases with positive multiplication [3]. Their positivity for local wall factors [3, Theorem 2.15] follows from cohomological DT integrality and positivity of DT invariants [3, Theorems 6.6 and 6.15]. Cluster monomials occur among the theta functions, so this also proves quantum Laurent positivity. Kontsevich and Soibelman gave a polynomial presentation of cohomological Hall algebras [13]. For symmetric quivers with zero potential, Efimov proved that a sign twist makes the Hall algebra free supercommutative and obtained a quantum-dilogarithm factorization with nonnegative multiplicities [7, Theorem 1.1 and Corollary 4.1]. Semistable Hall quotients [9] and the PBW and integrality theorems of cohomological Donaldson–Thomas theory [5] describe related structures behind these factorizations.

Huang gave combinatorial quantum Laurent expansion formulas for orbifolds without orbifold points of weight two, including punctured cases [11, Section 6].

The dependence on frozen coefficients and quantization has also been studied separately. Tran introduced quantum F-polynomials and their recursion [22]. Qin used these polynomials to compare compatible quantizations and coefficient patterns. For a common positive integer compatibility constant δ, Qin obtains the coefficients of the quantum cluster-variable expansions from the unit-compatibility case by substituting $q^{\delta/2}$ for $q^{1/2}$ in the coefficient polynomials [20, arXiv v3, Theorem 3.0.2, formula (3)]. His correction technique compares more general compatible realizations [20, arXiv v3, Section 9]. Ye used a central lattice extension to separate exchange-polynomial roots in generalized quantum scattering diagrams, and combined it with Davison–Mandel positivity and a positive cylinder projection [23]. That construction gives positivity in the adjoined polynomial roots.

The proof below applies the polynomial shuffle and primitive-string construction of OpenAI [18, Section 3] to quantum cluster mutation. Independent-center restrictions on polynomial shuffle quotients define a regular signed coproduct. A direct PBW argument and polynomial-string decomposition express the elementary factor multiplicities as dimensions of primitive starting spaces. The formal wall transport, including the change of roots, is developed in [18, Section 4]; finite energy slices and decorated histories appear in [18, Section 6]. We reconstruct the local and transport arguments, give their comparison with arbitrary integral compatible BZ seeds, and construct finite graded spaces for all Laurent coefficients. The comparison uses a central coefficient extension and an integral auxiliary form agreeing with the original quantization on root actions.

1.1. **The positivity theorem.** Let $\widetilde{B}$ be an $m \times n$ integer matrix whose principal $n \times n$ part B is skew-symmetric. Let Λ be an integer alternating $m \times m$ matrix satisfying

$$\widetilde{B}^T \Lambda = (D \;\; 0), \qquad D = \operatorname{diag}(d_1, \dots, d_n), \qquad d_i > 0. \tag{1.1}$$

The first n variables are mutable and the remaining variables are frozen. We invert the frozen variables. If M is a toric frame, our normalization is

$$M(a)M(b) = v^{\Lambda(a,b)} M(a+b). \tag{1.2}$$

A normalized cluster monomial in a seed M' means $M'(a)$ with $a_i \geq 0$ for $i \leq n$; the frozen exponents may be any integers.

Theorem 1.1 (Quantum Laurent positivity)**.** *For every compatible seed satisfying* (1.1)*, every seed M' obtained by a finite sequence of mutations, and every normalized cluster monomial $M'(a)$, the expansion*

$$M'(a) = \sum_{b\in\mathbb{Z}^m} c_b(v)M(b)$$

is finite and satisfies $c_b(v) \in \mathbb{N}[v^{\pm 1}]$ for every b. The same assertion holds with any seed in the mutation class as the expansion seed.

Theorem 1.1 proves Laurent positivity for all skew-symmetric mutation classes, with the general positive diagonal matrix D of the BZ definition. Davison uses $\widetilde{B}^T\Lambda = (I_n \;\; 0)$ in [2, Section 2.1]. Our comparison treats the compatibility constant on each mutable component and retains all frozen pairings.

The quantum Laurent phenomenon [1, Corollary 5.2], recalled in Theorem 2.1, supplies finite Laurent support and integral coefficients. The polynomial reordering and wall-transport statements used to prove nonnegativity are established below.

1.2. **Finite coefficient spaces.** The construction also gives an interpretation of individual Laurent coefficients by finite-dimensional vector spaces. These spaces are defined by polynomial maps, kernels, quotients and finite choices of wall-crossing terms.

Theorem 1.2 (Finite coefficient spaces)**.** *Fix the integer seed data, a finite mutation word, and the exponent of a normalized cluster monomial at the last seed. There is a terminating construction, using rational linear algebra and integer arithmetic, of finite-dimensional $\mathbb{Z}$-graded rational vector spaces $\mathcal{V}_b = \bigoplus_E \mathcal{V}_{b,E}$ such that*

$$M'(a) = \sum_{b\in\mathbb{Z}^m} \sum_{E\in\mathbb{Z}} \dim_{\mathbb{Q}} \mathcal{V}_{b,E}\; v^E M(b). \tag{1.3}$$

Only finitely many $\mathcal{V}_{b,E}$ are nonzero. The construction produces bounds on all exponents b and all energies E before the coefficient spaces are computed.

The spaces depend on graded complements, rational bases and generic paths; their graded dimensions are determined by (1.3).

1.3. **The proof.** Sections 3–5 prove positive reordering of elementary quantum products. An ordered Harder–Narasimhan decomposition identifies each ray factor with the Hilbert series of a polynomial shuffle quotient. The primitive decomposition gives its elementary factors with nonnegative integral multiplicities.

Sections 6–7 pass from positive factors to Laurent coefficients. Incoming coefficients determine a section by a triangular recurrence with positive terms. Under mutation, a bound on each non-pure wall in the mutated-root direction makes the change of completion a comparison of finite Laurent fibers. A terminal normalized cluster monomial has one incoming coefficient, equal to one; the change of roots transports this datum to the expansion seed.

Section 8 compares this construction with BZ mutation. Adjoining central frozen variables makes the exchange-root lattice an integral direct summand. On a component of compatibility constant d, an integral auxiliary form gives the same left and right root actions after $t = v^d$. The centered two-term exchange relation identifies the formal and BZ mutations. Specializing the auxiliary variables to one in the finite Laurent algebra proves Theorem 1.1.

Section 9 proves Theorem 1.2. Exact auxiliary-variable valuations give a finite box of exponents. Local parameter bounds give an upper quantum energy, and bar invariance gives the lower energy. Backward parameter bounds then reduce the primitive spaces and the incoming recurrence to finite rational matrices and finite sets of decorated histories.

2. Quantum seeds and Laurent expressions

2.1. **Based tori and normalization.** For an integer alternating form Λ on $\mathbb{Z}^m$, the based quantum torus is the free $\mathbb{Z}[v^{\pm 1}]$-module

$$\mathcal{T}_\Lambda = \bigoplus_{a\in\mathbb{Z}^m} \mathbb{Z}[v^{\pm 1}]X^a, \qquad X^aX^b = v^{\Lambda(a,b)}X^{a+b}.$$

Bilinearity proves associativity. Its identity is X^0, and $(X^a)^{-1} = X^{-a}$. Let $e_1,\dots,e_m$ be the standard basis of $\mathbb{Z}^m$ and put $X_i = X^{e_i}$. Then

$$X_iX_j = v^{2\Lambda_{ij}}X_jX_i, \qquad X^a = v^{-\sum_{i<j} a_ia_j\Lambda_{ij}}X_1^{a_1}\cdots X_m^{a_m}. \tag{2.1}$$

The second identity is valid for arbitrary integer exponents. We call these basis elements *centered monomials*.

Over $\mathbb{Q}(v)$ this torus is an iterated skew Laurent polynomial ring, so it has an Ore division ring of fractions. A toric frame M in such a division ring is a copy of a based torus with basis $M(a)$ satisfying (1.2); its images generate the ambient division ring. Whenever a seed is fixed, we identify $M(a)$ with X^a.

2.2. **Compatibility and mutation.** A quantum seed consists of a toric frame M and an integer matrix $\widetilde{B}$ satisfying (1.1) for its form Λ. Write $\widetilde{b}_k$ for the kth column. For a scalar r, write $[r]_+ = \max(r,0)$; positive parts of vectors are taken coordinatewise. Mutation in $k \le n$ replaces the generators by

$$\begin{aligned} &M'(e_j) = M(e_j) \quad (j\neq k),\\ &M'(e_k) = M(-e_k + [\widetilde{b}_k]_+) + M(-e_k + [-\widetilde{b}_k]_+). \end{aligned} \tag{2.2}$$

The exchange matrix changes according to

$$\widetilde{b}'_{ij} = \begin{cases} -\widetilde{b}_{ij}, & i=k \text{ or } j=k,\\ \widetilde{b}_{ij} + [\widetilde{b}_{ik}]_+b_{kj} + \widetilde{b}_{ik}[-b_{kj}]_+, & i\neq k,\ j\neq k. \end{cases} \tag{2.3}$$

The form of the new frame is $\Lambda' = E^T\Lambda E$, where

$$Ee_j = e_j\ (j\neq k), \qquad Ee_k = -e_k + [\widetilde{b}_k]_+. \tag{2.4}$$

Let E_0 be the principal $n\times n$ block of E. Skew-symmetry and (2.3) give

$$\widetilde{B}' = E\widetilde{B}E_0^T, \qquad B' = E_0BE_0^T,$$

so B' is skew-symmetric. The matrix $\widetilde{B}^T\Lambda\widetilde{B} = DB$ is alternating; taking transposes gives $DB = BD$, or $(d_i-d_j)b_{ij} = 0$. The definition of E then gives $E_0D = DE_0$, $E^2 = I_m$ and $E_0^2 = I_n$. Hence

$$\begin{aligned} (\widetilde{B}')^T\Lambda' &= E_0\widetilde{B}^T(E^T)^2\Lambda E\\ &= E_0(D\ \ 0)E = (E_0DE_0\ \ 0) = (D\ \ 0). \end{aligned}$$

The normalized monomial $M'(a)$ is defined from these generators by (2.1), using Λ'.

Let $g_+ = -e_k + [\widetilde{b}_k]_+$ and $g_- = -e_k + [-\widetilde{b}_k]_+$. For $j\neq k$,

$$\Lambda(g_+,e_j) - \Lambda(g_-,e_j) = \Lambda(\widetilde{b}_k,e_j) = d_k\delta_{kj} = 0.$$

Thus the two summands in $M'(e_k)$ have the same commutation factor with $M(e_j)$. That factor is the one prescribed by $E^T\Lambda E$. Let $\mathcal{D}$ be the division subring generated over $\mathbb{Q}(v)$ by the other coordinates, and put $x = M(e_k)$, $x' = M'(e_k)$. The exchange relation has the form $x' = Px^{-1}$ with $0\neq P\in\mathcal{D}$. Thus $x = (x')^{-1}P$, so the new coordinates generate the same division ring. For every $r\in\mathbb{Z}$, $(x')^r = P_rx^{-r}$ with $P_r\in\mathcal{D}^\times$. The distinct powers of x are linearly independent over $\mathcal{D}$ in the skew Laurent polynomial ring. A finite relation between the new monomials, grouped

by the exponent r of x', therefore has zero coefficient in each group. Linear independence of the based monomials in the other coordinates makes every scalar coefficient zero. This proves that the normalized monomials define a toric frame.

Mutation is involutive. To verify the scalar factors, write $b_+ = [\widetilde{b}_k]_+$ and $b_- = [-\widetilde{b}_k]_+$. Their kth entries vanish, and compatibility gives

$$\Lambda(b_+, b_-) = \Lambda(\widetilde{b}_k, b_-) = 0, \qquad \Lambda'(e_k, b_\pm) = \Lambda(g_+, b_\pm) = -\Lambda(e_k, b_\pm).$$

The second mutation reverses column k. If its new coordinate is x'', its exchange relation, multiplied on the right by x', is

$$x''x' = \sum_{\epsilon\in\{+,-\}} v^{-\Lambda'(e_k,b_\epsilon)} M'(b_\epsilon) = \sum_{\epsilon\in\{+,-\}} v^{\Lambda(e_k,b_\epsilon)} M(b_\epsilon) = xx'.$$

Here $M'(b_\epsilon) = M(b_\epsilon)$ because these exponents have zero kth coordinate. Cancelling x' gives $x'' = x$. These are the mutation conventions of [1, Section 4].

All frozen generators remain fixed. A normalized cluster monomial has nonnegative exponents on the mutable generators and arbitrary exponents on these fixed frozen generators. An unnormalized ordered monomial differs from it by a power of v.

Theorem 2.1 (Quantum Laurent phenomenon, Berenstein–Zelevinsky)**.** *For any two seeds in the same quantum mutation class, every cluster variable of either seed belongs to the based quantum torus of the other. In particular, every normalized cluster monomial has a finite expansion with coefficients in $\mathbb{Z}[v^{\pm 1}]$ in the based torus of every seed.*

The first assertion is [1, Corollary 5.2]. The assertion for monomials follows by multiplication, and the frozen inverse powers are already basis monomials in every seed. We use the theorem for integrality and finite Laurent support.

The mutable graph has vertices $1, \dots, n$ and an edge between distinct i, j exactly when $b_{ij} \neq 0$.

Lemma 2.2. *The compatibility condition has the following consequences.*

(i) *The columns of $\widetilde{B}$ are linearly independent over $\mathbb{Q}$.*
(ii) *$DB = BD$. Consequently $d_i = d_j$ whenever i and j belong to the same connected component of the mutable graph.*
(iii) *Distinct mutable components stay disconnected under mutation, and a mutation in one component does not change the exchange columns of another component.*

Proof. The right side of (1.1) has rank n, proving the first assertion. The relation $DB = BD$ proved above says $(d_i - d_j)b_{ij} = 0$ for all i, j. Equality of the constants then propagates along every edge. For the last assertion, if k lies in one component and j in another, $b_{kj} = 0$. Every correction to column j in (2.3) is therefore zero, including all its frozen entries. The mutable cross-component entries remain zero, so induction proves the assertion for an arbitrary mutation word. □

The following seed illustrates the integral lattice and frozen pairings treated by the central extension in Section 8.

Example 2.3. With one mutable variable and three frozen variables, take

$$\widetilde{B} = \begin{pmatrix} 0 \\ 2 \\ 0 \\ 0 \end{pmatrix}, \qquad \Lambda = \begin{pmatrix} 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix}.$$

Then $\widetilde{B}^T\Lambda = (2\ 0\ 0\ 0)$. The column is nonprimitive and generates the nonsaturated sublattice $2\mathbb{Z}e_2$, while the frozen–frozen pairing $\Lambda(e_3, e_4) = 1$ is not divisible by the compatibility constant 2. The exchange relation is $M'(e_1) = M(-e_1 + 2e_2) + M(-e_1)$.

2.3. **The abstract parameter.** The polynomial and wall arguments will use an integral alternating form Ω on a lattice L and an independent symbol t. Their coefficient field is

$$\mathbb{K} = \mathbb{Q}((t^{-1})), \qquad X^a X^b = t^{\Omega(a,b)} X^{a+b}.$$

An element of $\mathbb{K}$ has a bounded upper t-degree, and may have an infinite negative tail. Write $\mathbb{N}((t^{-1}))$ for the subsemiring whose coefficients are nonnegative integers. An integer exponent of t will be called its *energy*. All limits will be taken first at bounded positive root degree, with explicit energy bounds for any remaining infinite products. On a mutable component with compatibility constant d, the comparison with the original quantum seed will set $t = v^d$.

3. Polynomial shuffle algebras and ordered decomposition

We give the polynomial construction of the positive reordering theorem of [18, Section 3], whose multiplicities will be used at two-dimensional intersections of walls.

Let L be a lattice with an integral alternating form Ω, and write $\mathbb{K} = \mathbb{Q}((t^{-1}))$. Let $\mathcal{C} \subset L$ be an additive monoid contained in a strictly convex cone in a real plane. Choose an additive integral charge $c : \mathcal{C} \to \mathbb{N}$ which is positive on $\mathcal{C} \setminus \{0\}$ and has finite bounded sublevel sets. We use the multiplication

$$X^a X^b = t^{\Omega(a,b)} X^{a+b}$$

and complete with respect to charge. For $p \in \mathcal{C} \setminus \{0\}$ and $k \in \mathbb{Z}$, put

$$E_{p,k} = \prod_{j \geq 0} (1 + t^{k-2j} X^p), \qquad H_{p,k} = \prod_{j \geq 0} (1 - t^{k-2j} X^p)^{-1}. \tag{3.1}$$

An occurrence of a unit consists of its root p, parameter k, and type E or H. Repeated occurrences are distinguished.

Definition 3.1. A list of unit occurrences is *admissible* if, for every charge bound B, the parameters of its roots of charge at most B are bounded above, and only finitely many such occurrences have parameter at least J, for every $J \in \mathbb{Z}$.

Theorem 3.2 (Positive reordering)**.** *Suppose an admissible product of the units (3.1) is ordered by rays, with $\Omega(p_i, p_j) > 0$ whenever the ray of p_i precedes the distinct ray of p_j. Its factorization in the opposite ray order has, on every ray, an admissible product of E and H units with nonnegative integral multiplicities. At each exponent, the coefficients belong to $\mathbb{N}((t^{-1}))$. Every coefficient at a fixed exponent and fixed power of t is a finite sum.*

The assertion holds after any charge truncation. If the input list above any parameter threshold and below any charge bound is effectively given, and a bound on the relevant pairings is computable, every bounded part of the output unit list is computable by finite rational linear algebra.

Ray factorization itself is unique over $\mathbb{K}$: at the next charge, the coefficient of X^p not already supplied by smaller charges is assigned to the ray of p. A product of this new term with a positive-charge term has larger charge.

3.1. **The shuffle product and its quotients.** First take a finite list, indexed by $1, \dots, r$ in input ray order. Indices on the same ray may occur in any order. Define

$$a_{ii} = \begin{cases} 0 & \text{for type } E, \\ 1 & \text{for type } H, \end{cases} \qquad a_{ji} = [\Omega(p_i, p_j)]_+ \quad (i \neq j). \tag{3.2}$$

For dimension vectors $d, e \in \mathbb{N}^r$, set

$$p(d) = \sum_i d_i p_i, \qquad \chi(d,e) = \sum_i d_i e_i - \sum_{i,j} a_{ij} d_i e_j, \tag{3.3}$$

$$s(d) = \tfrac{1}{2}\chi(d,d), \qquad u(d) = \sum_i d_i(k_i + 1 - a_{ii}). \tag{3.4}$$

Directly from (3.2),

$$\chi(d,e) - \chi(e,d) = \Omega(p(d), p(e)). \tag{3.5}$$

Extend c linearly to the span of the cone. Choose a rational linear functional η on the root span such that $\eta(p)/c(p)$ increases in the input ray order, and define $\mu(d) = \eta(p(d))/c(p(d))$ for $d \neq 0$. Such an η exists by choosing a rational affine coordinate on the line $c = 1$ with the prescribed orientation; for a single ray take $\eta = 0$. In particular, $\chi(d,e) = \chi(e,d)$ when $\mu(d) = \mu(e)$.

Introduce a pack of variables of size d_i for each index i, and let

$$S_d = \mathbb{Q}[x_{i,\alpha} : 1 \le i \le r,\ 1 \le \alpha \le d_i]^{\prod_i \mathfrak{S}_{d_i}}, \qquad S_0 = \mathbb{Q}.$$

Every variable has ordinary degree one. For $f \in S_d$ and $g \in S_e$ define

$$f * g = \sum_{\text{pack shuffles}} f(x)g(y)K_{d,e}(x,y), \qquad K_{d,e}(x,y) = \prod_{i,j} \prod_{\substack{\alpha \le d_i \\ \beta \le e_j}} (y_{j,\beta} - x_{i,\alpha})^{a_{ij} - \delta_{ij}}. \tag{3.6}$$

Each shuffle selects the variables belonging to the first input in each output pack, and each pack shuffle occurs once. This is the polynomial shuffle presentation of the Hall product [13, Section 2.4].

Lemma 3.3. *The product (3.6) is polynomial and associative. It changes ordinary degree by $-\chi(d,e)$. For $d \neq 0$, the subspace*

$$I_d = \sum_{\substack{a+b=d, a, b \neq 0 \\ \mu(a) > \mu(d)}} S_a * S_b \quad \textit{is an ordinary homogeneous ideal of } S_d. \tag{3.7}$$

Put $B_d = S_d/I_d$ and $B_0 = \mathbb{Q}$. The shuffle product induces $B_d \otimes B_e \to B_{d+e}$ on each common slope. It preserves the sum of the shifted degrees

$$\deg_{\mathrm{sh}} f = \deg_{\mathrm{pol}} f + s(d). \tag{3.8}$$

Common translation and common differentiation preserve I_d.

Proof. Only differences between variables of the same pack can occur in a denominator, and each has exponent at most one. Along a fixed collision hyperplane, interchange the two colliding variables in a shuffle. The corresponding residues are opposite: all other factors have equal restrictions to that hyperplane. Thus the common numerator is divisible by its defining difference. The distinct differences are relatively prime irreducible polynomials, so all possible poles cancel. Both parenthesizations of a threefold product enumerate the same final splittings and use each of the three cross-kernels once. This proves associativity. The degree of $K_{d,e}$ is $-\chi(d,e)$.

Let $h \in S_d$, and restrict it to packs of dimensions a, b. This restriction is separately symmetric, hence is a finite sum $\sum_\nu h_{\nu,1} \otimes h_{\nu,2} \in S_a \otimes S_b$. Its invariance under output shuffles gives

$$h(f * g) = \sum_\nu (h_{\nu,1} f) * (h_{\nu,2} g).$$

Every summand $S_a * S_b$ of (3.7) is therefore an ordinary S_d-submodule. The degree formula gives homogeneity.

Suppose d, e have slope θ. A destabilizing first factor in S_d remains the first factor after multiplication by S_e. If $e = a + b$ with $\mu(a) > \theta$, associativity rewrites $S_d * (S_a * S_b) = (S_d * S_a) * S_b$,

whose first dimension $d + a$ also has slope greater than θ. This proves descent in both arguments. On that slope, $s(d + e) = s(d) + s(e) + \chi(d, e)$, which compensates the shuffle degree. Finally, every kernel depends on differences of variables. Simultaneous translation consequently preserves each destabilizing image and its sum. Differentiating the translation identity proves the assertion for common differentiation. □

3.2. **Separated restrictions.** For an ordered list $P = (r_1, \dots, r_q)$ of nonzero vectors with $\sum_j r_j = d$, divide each pack into labeled target groups of these dimensions. Add an independent center z_j to every variable in target group j, and take the quotient in each group. This defines

$$\widetilde{R}^z_P : S_d \longrightarrow \left(\bigotimes_{j=1}^{q} B_{r_j} \right) [z_1, \dots, z_q]. \tag{3.9}$$

Pack symmetry makes the choice of variables within each pack irrelevant. Translation is an automorphism of $B_d[z]$, with inverse translation by $-z$.

To restrict a shuffle expression, divide every source group among the target groups. Over $\mathbb{Q}(z_1, \dots, z_q)$, expand kernels between distinct targets in the internal variables. In each fixed internal degree their coefficients are finite sums of polynomial tensors, separately symmetric in the source–target intersection packs. Absorb these tensors into the source factors before taking the target shuffles. Kernels within a target are left unchanged. If a source prefix has nonzero dimension b in a target of dimension a, with $\mu(b) > \mu(a)$, that target product belongs to I_a and vanishes. The homogeneous quotient ideals make these Taylor calculations valid in the internal-degree completion. Extension to the independent-center field is faithful; a zero Taylor restriction therefore implies a zero original polynomial restriction.

Lemma 3.4 (Ordered decomposition)**.** *Choose homogeneous vector-space sections $B_a \to S_a$ of the quotient maps. Ordered shuffle multiplication induces an isomorphism*

$$\bigoplus_{\substack{d^1+\cdots+d^q=d,\ d^j\neq 0 \\ \mu(d^1)>\cdots>\mu(d^q)}} B_{d^1} \otimes \cdots \otimes B_{d^q} \longrightarrow S_d, \tag{3.10}$$

with the ordinary degree shifts supplied by the shuffle kernels. For $d = 0$ the source is the empty tensor $\mathbb{Q}$.

Proof. We first prove surjectivity by induction on $|d|_1 = \sum_i d_i$. After removing the chosen lift of B_d, an element of S_d is a sum of destabilizing twofold products. Each factor has smaller dimension and may be expanded by induction. Associate to a list of factors the polygon whose successive vertices are

$$\left(c(p(d^1 + \cdots + d^j)), \eta(p(d^1 + \cdots + d^j)) \right).$$

If two adjacent slopes are nondecreasing, merge their factors and expand the product by induction. The old two-segment polygon lies below its chord, whereas every replacement has decreasing slopes and lies above that chord. Hence the polygon weakly rises. Equality requires a single collinear replacement, so the number of segments decreases. The initial destabilizing cut puts the polygon strictly above the total chord somewhere; all subsequent replacements preserve this property. Consequently the merged pair cannot be the whole list: a two-segment list with nondecreasing slopes lies below the total chord. Thus each application of induction really has smaller total dimension. There are only finitely many ordered lists of nonzero subvectors of d, so the process terminates with strictly decreasing slopes.

For injectivity, restrict a source product to a target list $R = (r_1, \dots, r_q)$ with strictly decreasing slopes. In a surviving assignment, every nonempty part of a source prefix inside target j has slope at most $\mu(r_j)$. If that part has charge b_j, its height is at most $b_j \mu(r_j)$; an empty part has charge

and height zero. At fixed total charge, the sum of these bounds is maximized by filling the targets in decreasing slope order. The source polygon therefore lies below the target polygon at each source corner, and hence everywhere by concavity of the target polygon. If the polygons coincide, equality at their successive corners forces complete filling of the first target, then the second, and so on. Positive charge excludes a nonzero unfilled subdimension of charge zero. Thus the complete dimension-vector lists agree, even if distinct input occurrences have the same projected root.

In a finite relation among (3.10), choose a nonzero summand whose polygon is minimal for pointwise comparison, and use its list as the target. A surviving source polygon must be smaller, or equal with the same complete list. Minimality leaves only the chosen list. Its restriction is the tensor of the translated quotient classes multiplied by the cross-target kernels. These kernels are separately symmetric in the complete target packs and have nonzero scalar constant term over the center field. They are invertible in the internal-degree completion of the quotient tensor. Canceling them, and then undoing translation, shows that the chosen tensor is zero, a contradiction. □

Lemma 3.5. *If all entries of P have the slope of d, then (3.9) descends to a map R^z_P on B_d. These maps compose under refinement, with successive centers added. When a product on a common slope is restricted, only assignments whose nonempty source–target subgroups have that slope survive.*

Proof. A destabilizing first source group has slope greater than the common target slope. Its nonempty pieces in the targets cannot all have slope at most that slope, since their charge-weighted average is the source slope. At least one target is killed by its instability ideal. This proves descent. Refinement is composition of substitutions, and therefore still holds after descent. For the first source factor already on the common slope, a nonempty off-slope piece forces another piece above the slope by the same weighted-average argument. That target vanishes. In all remaining terms the products descend in their input subgroups by Lemma 3.3; the Taylor coefficients can be absorbed because the quotient ideals are ordinary ideals. □

3.3. **The Hilbert-series identity.** For an ordinarily graded space $V = \bigoplus_{\ell\geq 0} V_\ell$ with finite-dimensional pieces, write

$$\mathrm{Hilb}_V(z) = \sum_{\ell\geq 0} \dim_{\mathbb{Q}} V_\ell \, z^\ell.$$

The pack-symmetric polynomial ring has Hilbert series

$$\mathrm{Hilb}_{S_d}(z) = \prod_i \prod_{j=1}^{d_i} (1 - z^j)^{-1}.$$

Indeed, monomial orbits in a pack of size b are indexed by $\lambda_1 \geq \cdots \geq \lambda_b \geq 0$. With $\lambda_{b+1} = 0$, their differences $\nu_j = \lambda_j - \lambda_{j+1}$ satisfy $\sum_j \lambda_j = \sum_j j\nu_j$ and range independently over $\mathbb{N}$. From (3.1), the coefficient of X^{bp} in $E_{p,k}$ and $H_{p,k}$ is, respectively,

$$\frac{t^{bk-b(b-1)}}{\prod_{j=1}^b (1-t^{-2j})}, \qquad \frac{t^{bk}}{\prod_{j=1}^b (1-t^{-2j})}.$$

For E, subtract $0, 1, \ldots, b-1$ from the strictly increasing chosen indices; for H the index list is already weakly increasing. Thus the full ordered input product equals

$$\sum_{d\in\mathbb{N}^r} X^{p(d)} t^{u(d)-\chi(d,d)} \, \mathrm{Hilb}_{S_d}(t^{-2}). \tag{3.11}$$

The exponent is

$$u(d)-\chi(d,d)=\sum_i\bigl(d_ik_i-(1-a_{ii})d_i(d_i-1)\bigr)+\sum_{i<j}d_id_j\Omega(p_i,p_j).$$

The diagonal summands give the two single-occurrence formulas, and the second sum is the quantum multiplication shift in the input order. Occurrences on the same ray commute. The identity holds for every complete dimension vector; at bounded charge there are finitely many such vectors, whose contributions are then collected at each projected root.

For two factors of dimensions a,b in (3.10), the ordinary degree changes by $-\chi(a,b)$. Their energy relative to separate normalization consequently changes by

$$2\chi(a,b)-\chi(a,b)-\chi(b,a)=\Omega(p(a),p(b)).$$

Lemma 3.4 therefore identifies the factor on a ray of slope θ in the reversed product as

$$1+\sum_{\mu(d)=\theta}X^{p(d)}t^{u(d)-\chi(d,d)}\operatorname{Hilb}_{B_d}(t^{-2}). \tag{3.12}$$

An element of original shifted degree h contributes energy $u(d)-2h$.

4. A regular coproduct and primitive strings

Fix a slope θ. All nonzero dimensions in the remainder of this section have slope θ. In particular χ is symmetric on these dimensions. For a list $P=(r_1,\dots,r_q)$ write $s(P)=\sum_j s(r_j)$.

4.1. **The separated-center filtration.** Give the centers auxiliary degree zero. Define a decreasing filtration of B_d, indexed by $w\in\frac12\mathbb{Z}$, by

$$f\in F^wB_d\quad\Longleftrightarrow\quad R^z_Pf\text{ has internal degree at least }w-s(P)\text{ for every }P. \tag{4.1}$$

Each restriction preserves total polynomial degree in the internal variables and centers. Components of distinct total degree cannot cancel in the internal-degree tests, so F^wB_d is homogeneous for the original polynomial degree. The one-entry list is included. For fixed d, there are finitely many lists; their shift sums give a lower bound for the filtration. In ordinary degree ℓ, the one-entry test gives upper bound $\ell+s(d)$. Thus the filtration is finite in each ordinary degree. Set $\operatorname{gr}^w B_d=F^wB_d/F^{w+1/2}B_d$ and give $B_0=\mathbb{Q}$ weight zero. Passing to $\operatorname{gr}B_d$ preserves every original shifted graded dimension.

For $[f]\in\operatorname{gr}^wB_d$ define

$$L_P[f]=\prod_{i<j}(z_j-z_i)^{\chi(r_i,r_j)}(R^z_Pf)_w, \tag{4.2}$$

where the subscript selects internal degree $w-s(P)$; a negative or nonintegral degree selects zero. These symbols initially have rational center coefficients.

Lemma 4.1. *All symbols (4.2) are polynomial in their centers, and together they detect* gr^wB_d. *For the tensor filtration*

$$F^w(B_a\otimes B_b)=\sum_{u+v\ge w}F^uB_a\otimes F^vB_b, \tag{4.3}$$

an element $f\in B_a\otimes B_b$ belongs to $F^w(B_a\otimes B_b)$ if and only if, for every pair of lists R,S summing to a,b, $(R^{z^R}_R\otimes R^{z^S}_S)f$ has total internal degree at least $w-s(R)-s(S)$, where all centers are mutually independent. For a zero dimension take the empty list, with $s(\varnothing)=0$ and $R_\varnothing=L_\varnothing=\mathrm{id}_{\mathbb{Q}}$. The symbols $L_R\otimes L_S$ detect the associated graded.

Proof. Merge two target groups a, b in P. The shift sum increases by $\chi(a,b)$. Refine that merged group again, counting the relative center $\delta = z_b - z_a$ as an internal variable of degree one. The coarse filtration test gives degree at least $w - s(P) - \chi(a,b)$. To contribute fine internal degree $w - s(P)$, a term must therefore contain δ to degree at least $-\chi(a,b)$. When $\chi(a,b) < 0$, this is precisely the divisibility needed for the corresponding denominator in (4.2). Perform this comparison for every pair. In a rational vector-space basis of the target quotient tensor, each scalar polynomial coefficient is divisible by each required diagonal power. The differences for distinct unordered pairs are relatively prime irreducibles in the center polynomial ring, so their product divides that coefficient. This proves polynomiality.

If every leading restriction is zero, (4.1) places the representative in $F^{w+1/2}B_d$. Multiplication by the nonzero scalar prefactors cannot alter this conclusion after extension to the independent-center field. This proves detection.

The tensor tests hold on the right-hand side of (4.3). For the converse, choose filtration complements in the finite-dimensional spaces of each original degree. If a tensor violated the proposed bound, take its lowest nonzero total weight below w. The tensor product of the detecting families on its two factors is injective. Different factor weights (u, v) have different internal bidegrees $(u - s(R), v - s(S))$ and cannot cancel. Some tensor restriction thus detects the violating term. The same argument proves joint detection on every associated tensor grade. Whenever scalar parameters are adjoined, all testing centers are taken independent of those parameters. Accordingly the detection applies coefficientwise to actual rational vectors and to their scalar extensions. □

For example, with one particle $B_d = \mathbb{Q}[x]$ and $R^z_{(d)}f = f(x+z)$. Every x^ℓ has auxiliary weight $s(d)$, but its original shifted degree is $\ell + s(d)$; its symbol is z^ℓ. Thus the auxiliary weight and the degree recording energy are distinct.

4.2. **Product symbols.** For a target list $R = (r_1, \dots, r_q)$, let $*^{\otimes R}$ mean: bring the two factors assigned to each target together, with no additional sign, and use its internal shuffle product. Zero entries in a symbol list are omitted, with units reinserted in their target positions.

Lemma 4.2. *The common-slope product preserves the sum of filtration weights. For $f \in \mathrm{gr}^u B_d$, $g \in \mathrm{gr}^v B_e$, its symbols are*

$$L_R(f * g) = \sum_{\substack{d_i + e_i = r_i \ (1 \le i \le q) \\ \sum_i d_i = d,\ \sum_i e_i = e}} (-1)^{\sum_{i>j} \chi(d_i, e_j)} \, *^{\otimes R} \, (L_{(d_i)} f \otimes L_{(e_i)} g). \tag{4.4}$$

Only dimensions on the fixed slope occur in nonzero terms.

Proof. Lemma 3.5 leaves only the indicated assignments. Expand their cross-target kernels at separated centers. In target i the shuffle lowers polynomial degree by $\chi(d_i, e_i)$, which is compensated by $s(r_i) - s(d_i) - s(e_i) = \chi(d_i, e_i)$. Every nonconstant internal Taylor term raises internal degree. This proves preservation of the filtration bound.

On the leading grade, the cross-target kernel for $i < j$ contributes

$$(z_j - z_i)^{-\chi(d_i, e_j)} (z_i - z_j)^{-\chi(d_j, e_i)}.$$

Multiplication by the target prefactor in (4.2) cancels the cross terms in $\chi(r_i, r_j)$. What remains is the product of the two source prefactors and the sign $(-1)^{\chi(d_j, e_i)}$. Multiplying over all $i < j$ gives (4.4). □

4.3. **Construction of the coproduct.** Let $d = a + b$ be a proper splitting on the slope, and let $f \in F^w B_d$ have ordinary degree ℓ. At $Z = \infty$ expand

$$K_{a,b}(x + Z, y)^{-1} f(x + Z, y) = \sum_j C_j(f) Z^j. \tag{4.5}$$

Every C_j is an actual separately symmetric polynomial tensor: a specified power of Z involves only finitely many terms from the finitely many integer powers of linear factors. The restrictions descend, and their target ideals are ordinary ideals, so C_j induces a map from B_d to $B_a \otimes B_b$. Homogeneity gives

$$\deg_{\mathrm{pol}} C_j(f) = \ell + \chi(a, b) - j, \qquad C_j(f) = 0 \quad \text{if } j > \ell + \chi(a, b). \tag{4.6}$$

Lemma 4.3. *Each C_j preserves the tensor filtration. On the associated graded, all its negative-j components vanish. The assignment*

$$\Delta_{a,b}[f] = [C_0(f)] \tag{4.7}$$

is therefore a well-defined map $\operatorname{gr} B_d \to \operatorname{gr} B_a \otimes \operatorname{gr} B_b$, preserving the original shifted degree and the total auxiliary weight. It is characterized by

$$(L_R \otimes L_S)\Delta_{a,b}[f] = L_{R,S}[f], \tag{4.8}$$

where R, S are arbitrary refinements of a, b, with independent centers. The components with a zero dimension are the identity tensored with the unit.

Proof. Refine a, b into lists $R = (r_i)$ and $S = (s_j)$. The internal-degree-zero term of the inverse kernel is

$$\prod_{i,j} (z_j^S - z_i^R - Z)^{\chi(r_i, s_j)}.$$

All its other internal terms have positive degree. The restriction of f to the concatenated list satisfies the lower weight bound w. Thus every refined restriction of each $C_j(f)$ satisfies the same tensor lower bound. Lemma 4.1 implies $C_j(f) \in F^w(B_a \otimes B_b)$ in the actual tensor space. For fixed j, (4.6) bounds the internal degrees contributing to extraction, so coefficient extraction commutes with these refinements and quotient maps.

On weight w, only the displayed zero-degree inverse kernel remains. Multiplying by the two separate symbol prefactors supplies exactly the cross-list factors, and gives

$$(L_R \otimes L_S)\left(\sum_j [C_j(f)] Z^j\right) = L_{R,S}[f](z^R + Z, z^S). \tag{4.9}$$

The right side is a polynomial in the independent centers and Z. For $j < 0$, tensor detection therefore gives $[C_j(f)] = 0$ before any specialization of Z. Equation (4.6) leaves only finitely many nonnegative powers on this grade. Their polynomial may now be evaluated at $Z = 0$, defining (4.7) and yielding (4.8). Replacing f by a representative in $F^{w+1/2}$ raises every tensor coefficient by that weight, so the map is independent of its representative. Finally,

$$\ell + \chi(a, b) + s(a) + s(b) = \ell + s(d)$$

proves preservation of the original shifted degree. □

Set

$$\mathcal{H} = \mathbb{Q} \oplus \bigoplus_{\mu(d) = \theta} \operatorname{gr} B_d, \qquad \varepsilon(a, b) = (-1)^{\chi(a,b)}.$$

On the chosen slope, ε is a symmetric bilinear interchange sign. On tensor products we use

$$(x \otimes y)(x' \otimes y') = \varepsilon(\deg y, \deg x')(x * x') \otimes (y * y'), \tag{4.10}$$

where these degrees are dimension vectors.

Proposition 4.4. *With $\Delta = \sum_{a+b=d} \Delta_{a,b}$, the space $\mathcal{H}$ is a connected dimension-graded bialgebra. Its coproduct is cocommutative with interchange sign ε.*

Proof. Both iterated coproducts, tested against three lists of symbols, give the symbol of their concatenation by (4.8). Joint tensor detection proves coassociativity. Interchanging two lists changes their scalar prefactor by $(-1)^{\chi(a,b)}$ and permutes the underlying restriction. Detection proves signed cocommutativity. The empty-list restrictions give the counit and unit identities.

For multiplicativity use (4.4). If the two inputs contribute dimensions d_R, d_S and e_R, e_S to two test lists, the signs for the concatenated list differ from the two separate signs by $(-1)^{\chi(d_S,e_R)}$. This is exactly the factor in (4.10). Both sides of multiplicativity consequently have identical symbols, hence are equal. Lastly $\mathcal{H}_0 = \mathbb{Q}$, and every proper splitting strictly reduces $|d|_1$, proving connectedness. □

4.4. **Signed PBW and strings.** Let $\mathcal{P} = \mathrm{Prim}(\mathcal{H})$ denote the subspace of positive dimension annihilated by all proper coproducts. Equation (4.8) and tensor detection identify it with the common kernel of all symbols for lists having at least two nonzero entries. Indeed such a list can be divided into two nonempty lists; the converse follows by testing each proper coproduct against all its refinements.

Lemma 4.5 (Signed vector-space PBW)**.** *Symmetrized multiplication gives an isomorphism of vector spaces preserving dimension, original shifted degree, and auxiliary weight:*

$$\mathrm{Sym}_{\varepsilon}(\mathcal{P}) \xrightarrow{\sim} \mathcal{H}. \tag{4.11}$$

Here exchanging homogeneous vectors of dimensions a, b in the signed symmetric space multiplies the tensor by $\varepsilon(a, b)$.

Proof. For maps $f, g : \mathcal{H} \to A$ into an algebra with multiplication m_A, their convolution is $f \star g = m_A(f \otimes g)\Delta$. The augmentation is projection to $\mathcal{H}_0 = \mathbb{Q}$. Write e for the augmentation followed by the unit, $J = \mathrm{id} - e$, and

$$\mathcal{E} = \log_\star(\mathrm{id}) = \sum_{q \geq 1} \frac{(-1)^{q-1}}{q} J^{\star q}.$$

On dimension d, a nonzero term has at most $|d|_1$ positive-dimensional factors, so this is a finite operator. In the braided tensor target put $\alpha(h) = h \otimes 1$ and $\beta(h) = 1 \otimes h$. Then $\alpha \star \beta = \Delta$, and signed cocommutativity also gives $\beta \star \alpha = \Delta$. Since Δ is multiplicative,

$$\Delta\mathcal{E} = \log_\star(\Delta) = \log_\star(\alpha) + \log_\star(\beta) = \mathcal{E} \otimes 1 + 1 \otimes \mathcal{E}.$$

Thus $\mathcal{E}$ has primitive image. The convolution exponential identity

$$\mathrm{id} = e + \sum_{q \geq 1} \frac{1}{q!} \mathcal{E}^{\star q}$$

and the signed permutation invariance of the iterated coproduct show that signed symmetrized products of primitives span $\mathcal{H}$.

In characteristic zero the signed symmetric tensors are identified with invariant tensors by averaging over signed permutations. In a finite relation among their products, choose the largest tensor length q. Apply the iterated coproduct with q factors, projecting each factor to positive dimension. Products of fewer than q primitives vanish. For a product of q primitives the result is the sum of its $q!$ signed permutations, which on the invariant tensor is $q!$ times that tensor. The length-q term of the relation is therefore zero. Descending on q proves injectivity. All maps used preserve the three stated gradings. □

For a homogeneous basis of $\mathcal{P}$, a basis vector can be repeated in a signed symmetric tensor exactly when its self-sign is $+1$. Thus (4.11) describes polynomial generators for even self-sign and exterior generators for odd self-sign. In the original algebra, an odd primitive x satisfies

$$\Delta(x * x) = (x * x) \otimes 1 + 1 \otimes (x * x),$$

because its two mixed coproduct terms cancel. Thus its square is an even primitive.

Lemma 4.6 (Primitive strings)**.** *For every nonzero d, the space $\mathcal{P}_d$ is a direct sum of polynomial strings with original shifted degrees*

$$h, h+1, h+2, \dots, \qquad h - s(d) \in \mathbb{N}.$$

Each graded piece is finite dimensional, and every vector in these strings has self-sign $(-1)^{\chi(d,d)}$.

Proof. On S_d consider the operators

$$D_d = \sum_{i,\alpha} \frac{\partial}{\partial x_{i,\alpha}}, \qquad T_d f = \left(\frac{1}{|d|_1} \sum_{i,\alpha} x_{i,\alpha} \right) f. \tag{4.12}$$

They descend to B_d by translation invariance and ordinary ideality, respectively. We use the same symbols for the induced operators on B_d and its associated graded. On a restriction to $P = (r_1, \dots, r_q)$,

$$R_P^z(D_d f) = \sum_j \partial_{z_j} R_P^z f,$$
$$R_P^z(T_d f) = \left(\frac{\sum_j |r_j|_1 z_j}{|d|_1} + \frac{\sum_{j,i,\alpha} x_{j,i,\alpha}}{|d|_1} \right) R_P^z f.$$

Both preserve the filtration lower bound. In the second formula the internal-variable summand raises internal degree and disappears on the given associated grade. The symbol prefactor depends only on center differences, so its common center derivative is zero. Hence

$$L_P(D_d f) = \sum_j \partial_{z_j} L_P(f), \qquad L_P(T_d f) = \frac{\sum_j |r_j|_1 z_j}{|d|_1} L_P(f). \tag{4.13}$$

These operators preserve the common kernel of all proper symbols, and therefore preserve $\mathcal{P}_d$.

Their commutator is $[D_d, T_d] = 1$. Moreover D_d is locally nilpotent and T_d raises the original shifted degree by one. Define the finite-on-each-vector operator

$$\pi_d = \sum_{j \geq 0} \frac{(-1)^j}{j!} T_d^j D_d^j.$$

The commutator gives

$$D_d \pi_d = \sum_{j \geq 0} \frac{(-1)^j}{j!} T_d^j D_d^{j+1} + \sum_{j \geq 1} \frac{(-1)^j}{(j-1)!} T_d^{j-1} D_d^j = 0,$$
$$\sum_{r \geq 0} \frac{T_d^r}{r!} \pi_d D_d^r = \sum_{m \geq 0} T_d^m D_d^m \sum_{j=0}^{m} \frac{(-1)^j}{(m-j)! j!} = \mathrm{id}\,.$$

All sums are finite on each vector; the last equality is the binomial identity. Hence

$$f = \sum_{j \geq 0} \frac{T_d^j}{j!} \pi_d(D_d^j f).$$

A homogeneous basis of $\mathcal{P}_d \cap \ker D_d$ and its T_d-iterates thus span $\mathcal{P}_d$. They are independent: in a finite relation $\sum_{j=0}^{q} T_d^j b_j = 0$ with $D_d b_j = 0$, applying D_d^q gives $q! b_q = 0$, and downward induction eliminates the other terms. Therefore

$$\mathcal{P}_d \simeq \mathbb{Q}[T_d] \otimes (\mathcal{P}_d \cap \ker D_d). \tag{4.14}$$

Every component comes from nonnegative ordinary polynomial degree plus $s(d)$. The starting degrees have the claimed form. Finite ordinary polynomial spaces and the finite filtration give finite graded dimensions. The self-sign depends only on d. □

5. Positive factors and effective bounds

5.1. **Completion of the finite-input proof.** For a primitive string starting in ordinary degree ℓ and auxiliary weight w, its original shifted degree is $h = \ell + s(d)$. Its successive energies in (3.12) are $u(d) - 2h, u(d) - 2h - 2, \ldots$. Lemma 4.5 and Lemma 4.6 consequently make this string contribute

$$\begin{cases} E_{p(d),\, u(d)-2h}, & \chi(d,d) \text{ odd}, \\ H_{p(d),\, u(d)-2h}, & \chi(d,d) \text{ even}. \end{cases} \tag{5.1}$$

The integer multiplicity of this occurrence is

$$\dim_{\mathbb{Q}} (\mathcal{P}_d \cap \ker D_d)_{\ell,w}. \tag{5.2}$$

Summing over auxiliary weights recovers the original Hilbert series. Thus (3.12) is a positive product of the required units. Combining this with Lemma 3.4 proves Theorem 3.2 for finite input lists.

5.2. **Admissible infinite inputs.** For an infinite list give every occurrence its own index and use dimension vectors of finite support. All constructions for one fixed d involve only its support: the polynomial rings, instability ideals, restrictions, filtrations, primitive spaces, and differentiation kernels are unchanged by adding unused indices.

Fix a charge bound $B \geq 1$. Choose $K \geq 0$ bounding above all input parameters with root charge at most B, and choose $A \geq 1$ bounding the absolute pairings of roots of charge at most B. Such roots form a finite set. If $c(p(d)) \leq B$, then $|d|_1 \leq B$ and

$$\chi(d,d) \geq -AB^2, \qquad u(d) \leq \sum_i d_i k_i + B.$$

Since $h \geq s(d)$, the output parameter $k = u(d) - 2h$ satisfies

$$k \leq \sum_i d_i k_i + B + AB^2 \leq BK + B + AB^2. \tag{5.3}$$

If $k \geq J$, single out one of the active copies of an occurrence i. There are at most $B - 1$ remaining copies, each with parameter at most K. Hence every active occurrence satisfies

$$k_i \geq J - (B-1)K - B - AB^2. \tag{5.4}$$

Admissibility leaves finitely many indices in this interval and thus finitely many possible d. For each such vector, $h \leq (u(d) - J)/2$ bounds the ordinary starting degree. The spaces in (5.2) are finite dimensional, and only finitely many auxiliary weights occur in a fixed ordinary degree. This proves admissibility of the output list.

Consider now a coefficient with root charge at most B and energy e. For either the input or the output list, let $K' \geq 0$ bound the parameters and let $A' \geq 0$ bound the absolute root pairings

at this charge. List the $D \le B$ selected monomial copies in their multiplication order, with roots p_ν, parameters k_ν, and string indices $j_\nu \ge 0$. Their energy is

$$e = \sum_{\nu=1}^{D} (k_\nu - 2j_\nu) + \sum_{1 \le \mu < \nu \le D} \Omega(p_\mu, p_\nu). \tag{5.5}$$

The multiplication shift satisfies

$$\left| \sum_{\mu<\nu} \Omega(p_\mu, p_\nu) \right| \le A' \binom{D}{2} \le A' B^2.$$

The j_ν belonging to one E occurrence are distinct; for H repetitions are allowed. In either case,

$$e \le BK' + A'B^2, \tag{5.6}$$

$$k_\nu \ge e - (B-1)K' - A'B^2, \qquad \sum_\nu j_\nu \le \frac{BK' + A'B^2 - e}{2}. \tag{5.7}$$

Only finitely many unit occurrences and index choices contribute. Each coefficient is therefore a finite nonnegative integer, and its upper energy is bounded. The finite-list identities pass coefficientwise to the infinite list, proving the remaining algebraic assertions of Theorem 3.2.

5.3. **Finite linear-algebra realization.** Suppose the bounded-charge input list above a specified parameter threshold can be listed, and a bound A is supplied. The maximum of zero and the parameters in the list above threshold zero gives K. For an output threshold J, equation (5.4) gives a finite list of possible indices. Enumerate the dimension vectors supported there with charge at most B. The equation $k = u(d) - \chi(d,d) - 2\ell$ bounds every relevant ordinary degree.

In a fixed degree use monomial symmetric functions as bases of S_d. The instability test is the rational comparison $\eta(p(a))c(p(d)) > \eta(p(d))c(p(a))$. To form the instability ideal in output degree ℓ, only pairs of input degrees with sum $\ell + \chi(a,b)$ occur; these are nonnegative integers and hence give a finite computation. Shuffle multiplication and quotient maps are rational matrices. Each separated restriction is a polynomial substitution with finite ordinary degree, so the filtration tests and their associated graded spaces are kernels and quotients of finite rational matrices. The common kernels of the proper symbols, followed by the kernel of D_d, give (5.2). Equation (5.7) makes the expansion of every requested product coefficient finite.

In the later wall applications all roots lie in a free positive cone with known generators $c_1, \dots, c_n$, and charge is the sum of their coordinates. One can then take

$$A_B = \max\left\{1,\, B^2 \max_{i,j} |\Omega(c_i, c_j)|\right\}.$$

This also bounds pairings in the monoids cut out by two-dimensional joint planes.

6. Finite wall systems and positive coefficients

We use the wall construction and change of roots of [18, Section 4].

Let L be a lattice with an integral alternating form Ω. Write $L^\vee = \mathrm{Hom}(L, \mathbb{Z})$, $L_\mathbb{R} = L \otimes_\mathbb{Z} \mathbb{R}$ and $L_\mathbb{R}^\vee = \mathrm{Hom}_\mathbb{R}(L_\mathbb{R}, \mathbb{R})$, and put $\mathbb{K} = \mathbb{Q}((t^{-1}))$. Multiplication is $X^a X^b = t^{\Omega(a,b)} X^{a+b}$. Fix a finite set $C \subset L$ contained in a lattice basis, and write

$$C^+ = \sum_{c \in C} \mathbb{N} c, \qquad \left| \sum_{c \in C} r_c c \right|_C = \sum_{c \in C} r_c, \qquad a \le_C b \iff b - a \in C^+. \tag{6.1}$$

The nonzero elements of C^+ are called roots. Every interval for $\le_C$ is finite. Our completed modules consist of series whose support is contained in a finite union of translates of C^+.

Multiplication of two such series is defined coefficientwise: for a fixed output exponent, only finitely many pairs of exponents from their prescribed translates can contribute.

We may enlarge L to $L \oplus L^\vee$, with form

$$\Omega((a,f),(b,g)) = \Omega(a,b) + f(b) - g(a).$$

Pairing with all $(0,g)$ forces $a = 0$ for a vector (a,f) in the radical; pairing then with all $(b,0)$ forces $f = 0$. Thus the form is nondegenerate, and its restriction to L is the given form. When the original form is degenerate, we perform the wall construction in this extension. Wall factors have exponents in C^+, so a section starting in the original lattice stays there. Pairings between original exponents are unchanged, and all resulting identities restrict to the original torus. In this construction we continue to write L and Ω for the enlarged lattice and form when the extension is used. Set

$$V(a) = \Omega(-,a) \in L^\vee_{\mathbb{R}}, \qquad D_p = E_{p,-1} = \prod_{j\geq 0}(1 + t^{-1-2j}X^p). \tag{6.2}$$

6.1. **The total element and its incoming prescription.** The group of positive units is

$$G_C = 1 + \prod_{0\neq r\in C^+} \mathbb{K}X^r.$$

Inversion, logarithms and exponentials are well-defined in root degree. For $h \in L^\vee_{\mathbb{R}}$, let $G_{C,+}(h)$, $G_{C,0}(h)$ and $G_{C,-}(h)$ be the subgroups supported respectively on roots with positive, zero and negative h-value.

Lemma 6.1. *Every $g \in G_C$ has a unique factorization*

$$g = g_+(h)g_0(h)g_-(h), \qquad g_\sigma(h) \in G_{C,\sigma}(h). \tag{6.3}$$

At any bounded root degree the factors are constant on the cells of a finite hyperplane arrangement. On perturbing h, they are refined by factoring the middle element according to the perturbation.

Proof. Suppose the coefficients of all three factors below degree d are known. For a root r of degree d, precisely one of the three factors may contain X^r. Its unknown coefficient occurs in the product with coefficient one; all other contributions use smaller positive roots. Subtracting these known contributions determines it. Induction proves existence and uniqueness. There are finitely many roots of degree at most d; their hyperplanes therefore give the stated finite arrangement. Near a fixed h, strict signs do not change at that order. Factoring its middle element by the new signs and invoking uniqueness proves refinement. □

At a fixed root order, a *chamber* is a connected component of the complement of the relevant hyperplanes. A *wall* is a codimension-one cell on which the middle factor is nonidentity, and a *joint* lies on the intersection of at least two independent wall hyperplanes. We write $\mathrm{Ad}(u)(f) = ufu^{-1}$ for conjugation. The all-positive and all-negative chambers are given respectively by $h(c) > 0$ and $h(c) < 0$ for every $c \in C$.

Write $r^\perp = \{h \in L^\vee_{\mathbb{R}} : h(r) = 0\}$. At a generic point of $r^\perp$, the middle factor is supported on the positive ray through r. We regard it as the wall element there. A crossing from positive to negative evaluation applies $\mathrm{Ad}(g_0)$, and a crossing in the other direction applies its inverse. Earlier crossing elements occur on the right in an ordered product. If I and II are the chambers before and after a positive-to-negative crossing, refinement gives

$$g_-(II) = g_0 g_-(I).$$

Thus transport from an all-positive chamber to h is $\mathrm{Ad}(g_-(h))$. It depends only on the endpoint. If the value F_- is specified in the all-negative chamber, its transported value is consequently

$$F(h) = \mathrm{Ad}(g_-(h)g^{-1})F_-. \tag{6.4}$$

All these statements are interpreted at the finite root order needed for a requested coefficient.

Conversely, a consistent system of wall elements with this positive support comes from a total element. Take its ordered product g from an all-positive to an all-negative chamber. At the chosen finite order, join an all-positive point to a generic wall point h_0, and h_0 to an all-negative point, by generic straight segments. Replace their common endpoint by a short transverse crossing of the wall. The first segment crosses only roots negative at h_0, the short crossing contributes the wall element, and the last segment crosses only roots positive at h_0. Its product is a factorization of the form (6.3). Uniqueness identifies its middle factor with the prescribed wall element.

For every $0 \neq r \in C^+$, prescribe

$$[X^r] \log g_0(V(r)) = \begin{cases} [X^r] \log D_p, & r = jp,\ p \in C,\ j \geq 1, \\ 0, & \text{otherwise.} \end{cases} \tag{6.5}$$

These equations also have a unique solution. Indeed, in $\log g = \sum_r \ell_r X^r$, the coefficient ℓ_r contributes linearly with coefficient one to the middle logarithm at $V(r)$, because $V(r)(r) = 0$. All its other contributions involve roots of smaller degree. The equations therefore determine the ℓ_r successively.

The incoming end of the ray r is the positive $V(r)$ direction. On the subalgebra supported in $\ker V(r)$ the functional $\tau_r(B) = [X^r]B$ is a trace. In fact, if $a + b = r$ and $\Omega(a, r) = 0$, then $\Omega(a, b) = 0$, so the two orders of multiplication have the same X^r-coefficient. For positive units U, W in this subalgebra,

$$\tau_r \log(UW) = \tau_r \log U + \tau_r \log W. \tag{6.6}$$

To prove this, introduce a central variable z by replacing X^a with $z^{|a|_C} X^a$. At each fixed root order, expansion of $\log(1 + A)$ and cyclicity combine the differentiated terms of each power of A:

$$\tau_r \frac{d}{dz} \log U(z) = \tau_r\big(U'(z)U(z)^{-1}\big).$$

The product rule and cyclicity then give

$$\tau_r\big((UW)'(UW)^{-1}\big) = \tau_r(U'U^{-1}) + \tau_r(W'W^{-1}).$$

The logarithms have zero constant term in z. Integrating and setting $z = 1$ proves (6.6) coefficientwise in root degree. Apply this to a small generic perturbation of $V(r)$ within $r^\perp$. In the refined factorization, a strictly positive or strictly negative factor cannot contain X^r in its logarithm. Hence the middle factor retains its prescribed coefficient. Repeating this for every positive multiple on the ray shows that its incoming wall element is D_p on a simple ray and the identity on every other ray.

6.2. **Positive crossing formulas.**

Lemma 6.2. *Let $U_{a,k}$ be either $E_{a,k}$ or $H_{a,k}$, and let $b = \Omega(a, m) \in \mathbb{Z}$. If $b \neq 0$, the coefficients of $\mathrm{Ad}(U_{a,k})^{\operatorname{sgn} b} X^m$ are nonnegative. Its coefficient at X^m is one. In centered notation its factors have slots of energies*

$$k - |b| + 2h, \qquad 1 \leq h \leq |b|. \tag{6.7}$$

An E-slot can be chosen at most once; an H-slot can be chosen any nonnegative number of times. If $b = 0$, the adjoint action is the identity. For the simple unit, more precisely,

$$\mathrm{Ad}(D_p)^{\operatorname{sgn}\Omega(p,m)} X^m = \sum_{j=0}^{q} e_j(t^{1-q}, t^{3-q}, \ldots, t^{q-1}) X^{m+jp}, \qquad q = |\Omega(p, m)|. \tag{6.8}$$

Here e_j denotes the elementary symmetric polynomial of degree j, with $e_0 = 1$.

Proof. Write $U_{a,k}(Z) = \prod_{j\geq 0} f(t^{k-2j}Z)$, where $f(Z) = 1+Z$ or $f(Z) = (1-Z)^{-1}$. The commutation relation gives

$$\mathrm{Ad}(U_{a,k})X^m = X^m \frac{U_{a,k}(t^{2b}X^a)}{U_{a,k}(X^a)}. \tag{6.9}$$

For $b > 0$ the quotient is $\prod_{h=1}^{b} f(t^{k+2h}X^a)$. For $b < 0$, its reciprocal is $\prod_{h=b+1}^{0} f(t^{k+2h}X^a)$. Multiplying a selected copy of X^a on the right of X^m contributes the centering factor t^{-b}. In either case this gives exactly the list (6.7), up to its order. All coefficients are nonnegative, and the constant selection has coefficient one. Taking $k = -1$ and $f(Z) = 1+Z$ gives (6.8). In particular, pairing one gives precisely $X^m + X^{m+p}$. □

Theorem 6.3. *At every generic point of a wall, and to every specified root order, its element is an admissible product of E/H units on that ray with nonnegative integral multiplicities. A monomial crossing in the direction $-V(m)$ has coefficients in $\mathbb{N}((t^{-1}))$.*

Proof. Fix a degree bound d and proceed by induction in d. For a ray $\mathbb{R}_{>0}r$, follow a generic line in $r^{\perp}$ parallel to $V(r)$ from its incoming end to the desired cell. At this finite order it meets finitely many joints. If s is not proportional to r and $\Omega(s,r) = 0$, then $s^{\perp}$ is parallel to that line and can be avoided by its offset. The offset also avoids intersections of three independent normals. Every joint where the factor changes is therefore governed by a two-dimensional normal plane on which Ω is nonzero.

Order its positive rays so that the pairing of an earlier ray with a later ray is positive. At the incoming direction $V(a)$ of one of these rays, earlier normals evaluate positively and later ones negatively. Refinement in Lemma 6.1 consequently identifies the incoming local product with the product in this order, and the outgoing factors with its factorization in reverse order. The monoid obtained by intersecting this plane with C^+ is pointed, with positive integral charge $a \mapsto |a|_C$ and finitely many roots of bounded charge. Theorem 3.2 applies.

At degree d, the factor on the ray being followed is known on the preceding segment, and other rays are needed only through degree $d-1$. A degree-d term contributes at that degree only linearly, on its own ray; every interaction with another positive root has larger degree. Choose any admissible extensions of the known factorizations, apply positive reordering, and use uniqueness of the truncated ray factorization. This determines the required factor through degree d. Starting with the incoming factors supplied by (6.5) proves the assertion successively along the line.

For a crossing in direction $-V(m)$, positive pairing with its positive normal gives a positive-to-negative crossing and hence an adjoint; negative pairing gives an inverse adjoint. Every unit argument is a positive multiple of that normal, so Lemma 6.2 proves nonnegativity. At bounded root degree, only a bounded number of nonconstant copies can be selected, and their interaction pairings are bounded. The admissible unit parameters have a common upper bound. For a fixed output energy, none of the selected parameters can tend to $-\infty$, since the remaining bounded number of parameters and pairing shifts cannot compensate for it. Admissibility then leaves only finitely many occurrences; their string indices are bounded as well. Thus every energy coefficient is a finite nonnegative integer, and the upper energy is bounded. This is precisely membership in $\mathbb{N}((t^{-1}))$. □

6.3. **Incoming coefficients and a decreasing recursion.** A *section* is a family of chamber values related by the consistent wall operations. We will use sections obtained by transporting a finite Laurent polynomial. More generally the same construction works with a value supported in a finite union $U = \bigcup_{j=1}^{s}(m_j + C^+)$. Transport and inverse transport add only positive roots, so every chamber value retains this support bound. Define

$$A_w(F) = [X^w]F(V(w)+\epsilon), \tag{6.10}$$

where $\epsilon \in L^\vee_{\mathbb{R}}$ is a sufficiently small generic perturbation in the finite arrangement needed for this coefficient. The coefficient is independent of this choice: a separating wall has normal a with $\Omega(a, w) = 0$, and every potential predecessor $w - ja$ also pairs to zero with a. Its wall action therefore fixes all monomials that could contribute to X^w.

Fix a generic target point h. At the needed finite order, follow $h + sV(w)$ from large s down to zero. Write its successive crossing maps as $W_i : I_i \longrightarrow I_{i+1}$ and let a_i be the primitive positive normal of the ith wall. Put

$$J_i(u, j) = [X^{u+ja_i}]W_i(X^u).$$

The initial coefficient on this path is $A_w(F)$. Every crossing has the direction of $-V(w)$; also $\Omega(a_i, w-ja_i) = \Omega(a_i, w)$. The preceding theorem gives $J_i(w-ja_i, j) \in \mathbb{N}((t^{-1}))$, with $J_i(u, 0) = 1$. Subtracting successive X^w-coefficients and telescoping gives

$$[X^w]F(h) = A_w(F) + \sum_i \sum_{\substack{j>0 \\ w-ja_i \in U}} J_i(w - ja_i, j)[X^{w-ja_i}]F(I_i). \tag{6.11}$$

The inner exponents lie in the finite set $U \cap (w - C^+)$. Only finitely many roots can join two members of this set, so the outer sum is finite as well.

Theorem 6.4. *If all incoming coefficients of a section belong to $\mathbb{N}((t^{-1}))$, then every coefficient of every chamber value belongs to $\mathbb{N}((t^{-1}))$.*

Proof. For a target w, induct on the finite partially ordered set $U \cap (w - C^+)$, simultaneously in all chambers. In (6.11), every predecessor is strictly smaller than the target. The incoming term is nonnegative by assumption, and all other terms are products of a nonnegative jump coefficient with a coefficient already treated by the induction. The finiteness and energy bounds established above justify the sums and products in $\mathbb{K}$. □

The recursion is also effective. For example,

$$[X^{jp}]D_p = \frac{t^{-j^2}}{\prod_{i=1}^{j}(1 - t^{-2i})}. \tag{6.12}$$

This follows by selecting j distinct string positions and subtracting $0, 1, \ldots, j-1$ from them. Thus the incoming coefficients of the wall elements are rational functions of t. Finite triangular factorization computes any bounded-order wall data. Alternatively, positive factorizations are computed using the finite shuffle constructions of Theorem 3.2. Once a support interval and an energy are specified, (6.11) has finitely many strictly decreasing dependencies.

7. Change of roots and equality of Laurent coefficients

Fix $p \in C$ and put

$$\alpha_c = \Omega(p, c), \qquad c' = c + [-\alpha_c]_+ p \quad (c \neq p), \qquad C' = \{-p\} \cup \{c' : c \in C \setminus \{p\}\}. \tag{7.1}$$

The roots C' are again part of a lattice basis. The symplectic shear and its action on covectors are

$$S(w) = w - \Omega(p, w)p, \qquad S^\vee h = h \circ S^{-1}. \tag{7.2}$$

Indeed, expanding $\Omega(Sa, Sb)$ makes the two cross terms cancel. Since $\Omega(p, S(w)) = \Omega(p, w)$, the inverse is $S^{-1}(w) = w + \Omega(p, w)p$. Let $L_+ = \mathrm{id}$, $L_- = S$, and define the piecewise linear covector map

$$T(h) = \begin{cases} h, & h(p) \geq 0, \\ S^\vee h, & h(p) \leq 0. \end{cases} \tag{7.3}$$

The two definitions agree on $p^\perp$.

7.1. **A uniform bound along the mutation direction.** Write a non-pure positive root uniquely as

$$r = kp + u, \qquad u = \sum_{c \neq p} u_c c \neq 0, \qquad k, u_c \in \mathbb{N},$$

and set $|u| = \sum_{c \neq p} u_c$. For $\sigma \in \{+1, -1\}$ define

$$M_\sigma(u) = \sum_{c \neq p} [-\sigma \alpha_c]_+ u_c. \tag{7.4}$$

Lemma 7.1. *On the open side* $\operatorname{sgn} h(p) = \sigma$*, every non-pure root* $kp + u$ *in the support of a wall element or its logarithm satisfies*

$$0 \leq k \leq M_\sigma(u). \tag{7.5}$$

The pure wall on $p^\perp$ *has element* D_p *wherever no other wall direction is present.*

Proof. Projection onto the face $\mathbb{N}p$ is an algebra homomorphism: a sum involving a nonzero non-p component cannot be pure. At $p^\perp$, the pure part of the middle factor is therefore the pure part of the total element. The incoming prescription identifies it with D_p.

We induct on $|u|$. For each fixed k, work in a finite full-root truncation containing $kp + u$; the bound below is independent of that truncation. Let $b = \Omega(p, u)$. If $b \neq 0$, call the side $\sigma = \operatorname{sgn} b$ upstream. The path in $(kp + u)^\perp$ from its incoming end to a point on that side stays off the pure cut, since its forward tail has evaluation $h(p) + sb$ on p. If $b = 0$, the path is parallel to the cut and the same reasoning applies on either side. At the incoming end, the only nonidentity data are simple, and their non-pure roots have $k = 0$.

At a joint away from the cut, all normals have positive non-p degree. A new nonlinear term of non-p degree $|u|$ uses only terms of strictly smaller non-p degree. A term of the current degree can enter linearly only on its own ray and is carried from the preceding segment. By induction all smaller terms have the stated bound. Since M_σ is additive, the support set

$$\{kp + u : 0 \leq k \leq M_\sigma(u)\}$$

is closed under addition. The coefficient recursion for ray factorization therefore preserves it. This proves the upstream bound for each fixed k, through finitely many joints at the necessary full order. Its numerical bound is independent of that order. We may now conclude that the entire upstream p-fiber is finite.

For $b \neq 0$, continue across a generic joint on $p^\perp \cap u^\perp$. In its two-dimensional normal plane, the non-pure projections are positive multiples of the same ray modulo p. Their pairings with p all have sign σ. The upstream bounds already proved allow this genericity to be imposed in a finite full-root arrangement for all non-p degrees currently needed. Separate the pure factor from the incoming product. For $\sigma = +1$ it is first, so passage across it conjugates the non-pure product by D_p; for $\sigma = -1$ it is last, so conjugation is by D_p^{-1}.

Formula (6.8) adds between zero and $|\Omega(p, w)|$ copies of p to a monomial with non-p part w. On the projected ray under consideration,

$$M_{-\sigma}(w) = M_\sigma(w) + |\Omega(p, w)|.$$

The conjugated product thus satisfies the downstream bound. Refactorization preserves that additive support set: for a first putative exponent outside it, the product coefficient is zero and all subtracted terms are sums of already allowed exponents, so its factor coefficient must also be zero. At the current non-p degree, the cut calculation uses smaller degrees and current upstream coefficients. Subsequent off-cut joints use the same smaller-degree induction. Finally, logarithms and exponentials preserve an additive positive support set, proving both assertions about support. □

In particular, non-cut wall elements have constant term one in non-p degree, and their inverses have finite p-fibers too. The same is true after either shear, with finite Laurent fibers. Furthermore,

$$L_\sigma(kp+u) = \sum_{c\neq p} u_c c' + (M_\sigma(u) - k)(-p). \tag{7.6}$$

For $\sigma = +1$ this follows by substituting the definitions of c'; for $\sigma = -1$ use $M_+(u) + \Omega(p,u) = M_-(u)$. Thus supported roots map into C'^+. Put

$$B = \max(\{0\} \cup \{|\alpha_c| : c \neq p\}).$$

If a pushed root has new degree at most d, then $|u| \leq d$ and $k \leq B|u|$, so its old degree is at most $(B+1)d$. Every bounded new-order operation is consequently determined by finite old-order data.

7.2. **Comparison in a common rational ring.**

Theorem 7.2. *The covector map (7.3), with exponent map L_σ on the open side of sign σ, carries the wall system for C to that for C', after replacing the pure wall by D_{-p}.*

Let F and F' be sections specified by finite Laurent polynomials in the respective all-negative chambers. Suppose their initial values satisfy the rational identity

$$F_- = \mathrm{Ad}(D_p) F'_-. \tag{7.7}$$

Their actual transported coefficients satisfy

$$F'(T(h)) = L_\sigma F(h), \qquad \sigma = \operatorname{sgn} h(p), \tag{7.8}$$

where the equality is taken in each finite Laurent p-fiber. In particular,

$$A_w(F) = A_{\mu_p(w)}(F'), \qquad \mu_p(w) = w + [-\Omega(p,w)]_+ p, \tag{7.9}$$

and μ_p is a bijection of L.

Proof. First we prove the identity at the pure cut:

$$S \circ \mathrm{Ad}(D_p) = \mathrm{Ad}(D_{-p})^{-1}. \tag{7.10}$$

For $q = \Omega(p,m) \geq 0$, put $b_{q,j} = e_j(t^{1-q}, t^{3-q}, \ldots, t^{q-1})$. The letters are invariant under inversion and have product one, giving $b_{q,j} = b_{q,q-j}$. Formula (6.8), followed by S, sends X^m to $\sum_{j=0}^{q} b_{q,j} X^{m-(q-j)p}$. This equals $\mathrm{Ad}(D_{-p})^{-1} X^m$. For negative pairing, apply the proved equality to X^{-m} and take inverses. Both sides act by algebra homomorphisms on rational monomial expressions, so (7.10) follows.

Set $Z = X^p$ and choose a lattice complement $L = \mathbb{Z}p \oplus L_0$ containing $C \setminus \{p\}$. Allow coefficients in $\mathbb{K}(Z)$ on the monomials X^ℓ, $\ell \in L_0$, and complete only in the free positive monoid generated by the images of $C \setminus \{p\}$. Section values may be supported in finitely many translates of this monoid. The multiplication rule, in addition to the centered torus multiplication, is

$$X^\ell f(Z) = f(t^{-2\Omega(p,\ell)} Z) X^\ell. \tag{7.11}$$

Every specified non-p coefficient is therefore a rational function of Z, and its products use finitely many decompositions of that non-p exponent.

Lemma 7.1 places each non-cut wall element and its inverse in this ring, with finite Laurent support in Z at every non-p degree. Pure conjugation fixes Z and multiplies each X^ℓ by a rational function of Z, by (6.9). The shear is the identity modulo p. Hence all operations in the comparison act on the same ring.

Old consistency holds there as well. The expansion map $\mathbb{K}(Z) \hookrightarrow \mathbb{K}((Z))$ at zero is injective. On a monomial, the expansion of an old crossing agrees with its action in the old positive completion. At bounded non-p degree, the non-cut operations are determined by finitely many old roots, by the fiber bound. To compare a further coefficient in Z, choose the additional finite full-root order

needed for it. Old path independence applies at that order. Thus a difference of two rational path operations has zero expansion at zero in every non-p coefficient, and injectivity makes the difference zero in the common ring. The path maps are continuous algebra homomorphisms; their equality on monomials extends through the localization and then the non-p completion.

Push the non-cut elements by L_σ and insert D_{-p} on the pure cut. Equation (7.6) and the following degree bound place them in the new completion, with finite data at every new order. In a loop, the linear identifications on successive open sides cancel; at a cut crossing the remaining discrepancy is precisely corrected by (7.10). The pushed actions are therefore consistent in the common rational ring. Expansion at $Z^{-1} = 0$ gives the new-cone actions. Indeed, replacing c by c' changes a fixed projected coefficient only by a Laurent monomial in Z and a power of t, and (7.6) ensures that every increment is new-cone positive. The finite-order bound makes this expansion legitimate. Consistency of actions implies consistency of elements after the nondegenerate enlargement: if a logarithm has a first nonzero term λX^r, choose m with $\Omega(r, m) \neq 0$; its commutator with X^m is nonzero at that first degree.

We next identify the incoming prescription. A symplectic map A satisfies $V(Ar) = A^\vee V(r)$. Thus, on an eventual open side, the incoming direction of a pushed ray is the image of its old incoming direction. For a simple root $c \neq p$, the relevant side has sign α_c, and its image is exactly $c + [-\alpha_c]_+ p = c'$; its unit becomes $D_{c'}$. Other non-pure incoming logarithmic coefficients are zero. If the incoming direction is parallel to the cut, its exponent pairs to zero with p, and both linear pieces agree there. Either generic offset gives the same coefficient by (6.6). A portion with no old wall preimage has identity element. The remaining pure incoming element is D_{-p}. The uniqueness of (6.5) identifies the pushed system with the system for C'.

Choose a point on the cut with $h(c) < -1$ for every $c \neq p$, and take $h(p)$ small on either side, with $|h(p)| < (B + 1)^{-1}$. Such covectors exist since C is part of a basis. For every non-pure supported root,

$$h(kp + u) \leq B|u|\,|h(p)| - |u| < 0.$$

The short crossing between these points meets only the pure wall. The negative side is in the old all-negative chamber, and the positive side is in the new all-negative chamber, since $h(-p) < 0$ and

$$h(c') = h(c) + [-\alpha_c]_+ h(p) < -1 + \frac{B}{B+1} < 0.$$

On the positive side the old section takes the value $\mathrm{Ad}(D_p)^{-1} F_- = F'_-$, a finite Laurent polynomial by hypothesis. On the negative side the pushed old value is SF_-, while (7.10) and (7.7) give

$$SF_- = \mathrm{Ad}(D_{-p})^{-1} F'_-.$$

This value is also finite.

To reach any other chamber on one open side, transport entirely within that half-space. Only non-cut walls are crossed. Their elements and inverses have finite Laurent p-fibers. For a fixed non-p root-degree bound N, Lemma 7.1 leaves only roots of old degree at most $(B+1)N$. Thus a specified non-p coefficient requires only finitely many decompositions and crossings. Starting from either of the finite values just obtained therefore preserves finite Laurent support in each fiber. Applying L_σ preserves it as well. Each compared fiber is the same finite Laurent polynomial, so its actual coefficients agree independently of whether we expand at zero or infinity.

Finally, $V(w)$ lies on the side of sign $\Omega(p, w)$. For positive pairing the exponent map is the identity, and for negative pairing it is $S(w) = w - \Omega(p, w)p$. For zero pairing both pieces fix w, and incoming coefficients are independent of the generic offset. This proves (7.9). Its label map preserves $\Omega(p, w)$, and hence its inverse subtracts $[-\Omega(p, w)]_+ p$. It is a bijection. □

7.3. **A monomial with one incoming coefficient.**

Lemma 7.3. *Suppose $\Omega(c,m)\le 0$ for every $c\in C$. The section whose all-negative value is X^m satisfies*

$$A_w(F)=\begin{cases}1, & w=m,\\ 0, & w\neq m.\end{cases} \tag{7.12}$$

Proof. Every wall action and its inverse add only positive roots. Thus the support is contained in $m+C^+$, and the coefficient at X^m is one in every chamber. Consider $w=m+r$ with $0\neq r\in C^+$. In the finite truncation needed for this coefficient, start from a generic all-negative point h_- and follow $h_-+sV(w)$ for $s\ge 0$. At every crossed wall its positive normal a satisfies $\Omega(a,w)>0$: its evaluation starts negative and increases to zero. Any nonidentity contribution from X^m to X^w would therefore give

$$r=\sum_j q_j a_j, \qquad q_j>0, \qquad \Omega(a_j,w)>0,$$

so $\Omega(r,w)>0$. Alternatingness and the hypothesis instead give

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

No such contribution exists. At the far end the target coefficient is the incoming coefficient, proving (7.12). □

8. From root transport to arbitrary quantum seeds

For the seed data of Theorem 1.1, adjoin central coefficients so that the exchange-root lattice is an integral direct summand.

8.1. **Central coefficients and an integral root model.** For the moment suppose that all current mutable vertices have the same compatibility constant d. This includes a single mutable component with the other components held fixed and regarded as frozen. Add n frozen variables $Y_1,\dots,Y_n$ which commute with every original variable and with each other. On $\widehat{L}=\mathbb{Z}^m\oplus\mathbb{Z}^n$ put

$$\widehat{B}=\begin{pmatrix}\widetilde{B}\\ I_n\end{pmatrix}, \qquad \widehat{\Lambda}=\begin{pmatrix}\Lambda & 0\\ 0 & 0\end{pmatrix}, \qquad p_i=-\widehat{b}_i. \tag{8.1}$$

Thus $p_i=(-\widetilde{b}_i,-e_i)$, where the second coordinate is in the new frozen lattice. The extended pair is compatible with the same d. The original coordinate basis together with $p_1,\dots,p_n$ is a basis of $\widehat{L}$: the change-of-basis matrix has diagonal blocks $I_m,-I_n$. In particular

$$\mathsf{N}=\bigoplus_{i=1}^{n}\mathbb{Z}p_i$$

is a direct summand. All subsequent root mutations are integral basis changes within this same lattice N.

Let $\Pi:\mathbb{Z}^m\longrightarrow\mathbb{Z}^n$ be the projection to the mutable coordinates. Define the integral alternating form

$$\Omega_0=\begin{pmatrix}0 & -\Pi^T\\ \Pi & -B\end{pmatrix} \quad \text{on } \widehat{L}. \tag{8.2}$$

For an ambient vector $a=(a^{\mathrm{old}},a^Y)$, compatibility and skew-symmetry of B give

$$\begin{aligned}\widehat{\Lambda}(p_i,a)&=-d\,a_i^{\mathrm{old}}, & \widehat{\Lambda}(p_i,p_j)&=db_{ij},\\ \Omega_0(p_i,a)&=-a_i^{\mathrm{old}}, & \Omega_0(p_i,p_j)&=b_{ij}.\end{aligned} \tag{8.3}$$

Indeed, $p_i^T\Omega_0 = (-e_i^T\Pi, \widetilde{b}_i^T\Pi^T + e_i^T B)$, and the second block is zero because the transpose of column i of B is the negative of its row i. Consequently

$$(\widehat{\Lambda} - d\Omega_0)(r, a) = 0 \qquad (r \in \mathsf{N}, a \in \widehat{L}). \tag{8.4}$$

Lemma 8.1 (Comparison of root actions)**.** *Set $t = v^d$ and extend scalars from $\mathbb{Q}((t^{-1}))$ to $\mathbb{K}_v = \mathbb{Q}((v^{-1}))$. Matching centered monomials by their labels intertwines left and right multiplication by every root series, and hence every root-wall conjugation, between the model with form Ω_0 and the actual model with form $\widehat{\Lambda}$. It also intertwines the rational pure-wall operations used in Theorem 7.2.*

Proof. For $r \in \mathsf{N}$ and $a \in \widehat{L}$, equation (8.4) gives

$$t^{\Omega_0(r,a)} = v^{\widehat{\Lambda}(r,a)}, \qquad t^{\Omega_0(a,r)} = v^{\widehat{\Lambda}(a,r)}.$$

Thus the label matching is an isomorphism of bimodules for the two identified root subalgebras, on each coset of N. It intertwines finite products and all coefficientwise root multiplication, inverses of root units and conjugations in the positive root completion. We make this comparison over $\mathbb{K}_v$, which contains the coefficients of every actual BZ Laurent expression.

For the rational operation at a simple root p, choose $\widehat{L} = \mathbb{Z}p \oplus L'$; this is possible because p is a member of a basis of the direct summand N. Write $Z = X^p$. Localizing the nonzero elements of $\mathbb{K}_v[Z^{\pm 1}]$ gives a module with summands $\mathbb{K}_v(Z)X^\ell$, $\ell \in L'$. The commutation rule in both models is

$$ZX^\ell = v^{2\widehat{\Lambda}(p,\ell)}X^\ell Z = t^{2\Omega_0(p,\ell)}X^\ell Z.$$

The localization is legitimate: conjugating a nonzero Laurent polynomial in Z by any monomial scales Z by a nonzero scalar and therefore preserves the multiplicative set. For a finite sum of monomials, the product of the finitely many scaled denominators gives a common left or right denominator, proving the Ore condition. Multiplication by a root monomial is likewise intertwined after this localization. The pure adjoint and its inverse send each monomial to a rational function of Z times that monomial, by the finite-ratio calculation in Lemma 6.2. Their formulas consequently agree under the label matching. The same comparison applies after completion in the remaining positive root directions. □

The wall identities for Ω_0 restrict to $\widehat{L}$ from the nondegenerate enlargement of Section 6. We then apply this root bimodule map and specialize in the actual central torus (8.1).

8.2. **The exact exchange relation.** At the initial extended seed, define E by the convention (2.4), with $\widehat{B}$ in place of $\widetilde{B}$. Thus $E^2 = 1$. If $g = Ee_k$, then $\Omega_0(p_k, g) = 1$ and $\Omega_0(p_k, Ee_j) = 0$ for $j \neq k$. The simple wall $D_{p_k} = E_{p_k,-1}$ therefore satisfies

$$\begin{aligned} \mathrm{Ad}(D_{p_k})X^g &= X^g + X^{g+p_k} \\ &= X^{-e_k+[\widehat{b}_k]_+} + X^{-e_k+[-\widehat{b}_k]_+}. \end{aligned} \tag{8.5}$$

Both coefficients are exactly one. In the uncentered ratio the single factor has coefficient t; multiplication of X^g by X^{p_k} has coefficient t^{-1}, so these powers cancel. All other coordinate generators are fixed by this adjoint. In the actual extended torus, Lemma 8.1 proves the identical assertion after $t = v^d$. Since the new generators have form $E^T\widehat{\Lambda}E$, the identity extends from coordinate generators to every centered monomial, including those with negative exponents.

Let E_0 be the $n \times n$ mutable block of E. Matrix mutation is

$$\widehat{B}' = E\widehat{B}E_0^T, \qquad \widehat{\Lambda}' = E^T\widehat{\Lambda}E. \tag{8.6}$$

Identifying the new exponent lattice with the old one by E, its negative exchange columns become

$$-E\widehat{b}'_k = -p_k, \qquad -E\widehat{b}'_j = p_j + [-b_{kj}]_+p_k \quad (j \neq k). \tag{8.7}$$

This is the root change in Theorem 7.2, since $\Omega_0(p_k, p_j) = b_{kj}$, and is an integral basis change in N.

Let $\Pi_0 : \widehat{L} \longrightarrow \mathbb{Z}^n$ be the mutable-coordinate projection. Since $\Pi_0 E = E_0 \Pi_0$ and $E_0^2 = 1$,

$$(\widehat{B}')^T \widehat{\Lambda}' = dE_0\Pi_0 E = d\Pi_0. \tag{8.8}$$

8.3. **Transport of an actual cluster monomial.** Fix a mutation word of length s and a normalized cluster monomial at its last seed. Let $E^{(j)}$ be the full mutation matrix from the exponent lattice of seed $j+1$ to that of seed j, and put $J_0 = 1$, $J_{j+1} = J_j E^{(j)}$. The map J_j places the exponent lattice of seed j in the fixed initial lattice. Write $\widehat{B}_j$ for its exchange matrix and set

$$\begin{aligned} \widehat{\Lambda}_j &= J_j^T \widehat{\Lambda} J_j, \quad \Omega_j = J_j^T \Omega_0 J_j, \\ p_{j,i} &= -J_j \widehat{b}_{j,i}, \quad C_j = \{p_{j,i} : 1 \le i \le n\}. \end{aligned}$$

Applying (8.6) and (8.8) successively in seed coordinates gives $\widehat{B}_j^T \widehat{\Lambda}_j = d\Pi_0$ and makes C_j a basis of N. Since $p_{j,i} \in \mathsf{N}$, (8.4) yields

$$\Omega_0(p_{j,i}, J_j a) = d^{-1}\widehat{\Lambda}(p_{j,i}, J_j a) = d^{-1}\widehat{\Lambda}_j(-\widehat{b}_{j,i}, a) = -a_i. \tag{8.9}$$

Thus the exchange and root-change calculations apply at every stage with the current forms Ω_j and $\widehat{\Lambda}_j$.

Write the chosen cluster monomial in seed j and transport its monomial labels by J_j; call the result P_j. Each P_j is an actual finite based Laurent polynomial over $\mathbb{Z}[v^{\pm 1}]$, by Theorem 2.1. Equation (8.5) gives the rational identity

$$P_j = \mathrm{Ad}(D_p) P_{j+1}, \qquad p \in C_j \text{ the root mutated at step } j. \tag{8.10}$$

The two finite values are the all-negative initial values for the respective wall systems. Lemma 8.1 and Theorem 7.2 thus identify their actual incoming coefficients by a bijection of labels.

At the last seed, write the chosen normalized monomial as $M_s(a)$. Then $P_s = X^{J_s a}$ and $\Omega_0(p_{s,i}, J_s a) = -a_i \le 0$ by (8.9). Lemma 7.3 therefore says that its incoming data consist of a single coefficient equal to one. Applying the label bijections backwards proves the same assertion for P_0. By Theorem 6.4, all coefficients of P_0 belong to $\mathbb{N}((v^{-1}))$. They were already elements of $\mathbb{Z}[v^{\pm 1}]$; hence

$$\mathbb{Z}[v^{\pm 1}] \cap \mathbb{N}((v^{-1})) = \mathbb{N}[v^{\pm 1}] \tag{8.11}$$

proves positivity in the central extension.

8.4. **Specialization and several components.** The actual initial Laurent torus has the homomorphism

$$\mathrm{ev}_Y : \mathcal{T}_{\widehat{\Lambda}} \longrightarrow \mathcal{T}_{\Lambda}, \qquad X^{(b,c)} \longmapsto X^b, \qquad Y_i \longmapsto 1. \tag{8.12}$$

It is a homomorphism because the added directions are in the radical of $\widehat{\Lambda}$. Its application to a finite positive Laurent polynomial preserves positivity.

Deletion of the auxiliary exchange rows commutes with matrix mutation. If E is the mutation matrix for the original current seed and its form is Λ, the extended mutation matrix has the block form

$$\widehat{E} = \begin{pmatrix} E & 0 \\ R & I_n \end{pmatrix}, \qquad \widehat{E}^T \begin{pmatrix} \Lambda & 0 \\ 0 & 0 \end{pmatrix} \widehat{E} = \begin{pmatrix} E^T \Lambda E & 0 \\ 0 & 0 \end{pmatrix},$$

where R is the $n \times m$ matrix with $R_{ij} = \delta_{jk}[\widehat{b}_{m+i,k}]_+$. Thus at every seed the extended form has the original form as its upper-left block and is zero on the auxiliary directions.

We prove by induction that evaluation sends every extended cluster variable to its original counterpart. Let $\widehat{M}$ and M be the current frames, with forms $\widehat{\Lambda}$ and Λ in this calculation, and write $g_{\pm} = -e_k + [\pm\widehat{b}_k]_+$. Multiplying the extended exchange relation on the right by $\widehat{M}(e_k)$ gives

$$\begin{aligned}\widehat{M}'(e_k)\widehat{M}(e_k) = v^{\widehat{\Lambda}(g_+,e_k)}\widehat{M}([\widehat{b}_k]_+) \\ + v^{\widehat{\Lambda}(g_-,e_k)}\widehat{M}([-\widehat{b}_k]_+).\end{aligned} \tag{8.13}$$

All factors lie in the finite Laurent algebra of the initial extended seed. The right side uses nonnegative mutable powers and frozen units. By the induction hypothesis and the block-form calculation, its evaluation is the original exchange expression multiplied on the right by $M(e_k)$. Hence

$$\mathrm{ev}_Y(\widehat{M}'(e_k))M(e_k) = M'(e_k)M(e_k).$$

Right cancellation of the nonzero element $M(e_k)$ proves the induction. The same block-form calculation preserves the centering factors, so evaluation also sends every normalized cluster monomial to its original counterpart.

Proof of Theorem 1.1. If $n = 0$, every cluster monomial is already an initial based monomial, and the assertion is immediate. When the mutable compatibility constants are equal, the preceding argument proves the theorem. In general they are constant on each mutable connected component, by Lemma 2.2. Fix one component and hold the other mutable variables frozen. The preceding proof applies to this smaller mutable set. A mutation in another component changes neither its exchange columns nor the variables that occur in its exchange relation. Every cluster variable from this component in an arbitrary mutation word is therefore the same rational expression as the one obtained from the subsequence of mutations in that component. All cluster variables are consequently positive in the chosen initial seed. A normalized cluster monomial is a power of v times an ordered product of these variables and frozen powers. Based torus multiplication multiplies coefficients and introduces only powers of v, so it preserves positivity. This proves the assertion for monomials. Finally the initial seed was arbitrary, which proves the assertion relative to every seed. □

For the finite construction, add one central variable for every mutable vertex and use the single form (8.2) for the full principal matrix B. For a component c, let N_c be the span of its roots and d_c its compatibility constant. Then

$$\widehat{\Lambda}(r,a) = d_c\Omega_0(r,a) \quad (r \in \mathsf{N}_c, a \in \widehat{L}), \qquad \Omega_0(\mathsf{N}_c, \mathsf{N}_{c'}) = 0 \quad (c \neq c'). \tag{8.14}$$

Run the root system of component c with parameter $t_c = v^{d_c}$. The root subalgebras of distinct components commute, and their wall actions commute on every ambient monomial. Their product therefore defines the simultaneous system, with coefficients in $\mathbb{Q}((v^{-1}))$ and with the same cone and incoming-label rules. Local refactorization leaves commuting component factors unchanged. The simple cut in one component also commutes with every other component action. Thus the proofs of the incoming recurrence and of change of chart apply component by component to this product system. Apply it directly to the whole normalized terminal monomial, whose incoming coefficient is one. All scalar factors are then those of the original centered monomial and its root actions.

9. Finite graded spaces for all coefficients

We retain the central extension of Section 8 and the finite matrix construction of [18, Sections 3 and 6]. We compute all exponent and energy bounds from the seed, mutation word and terminal monomial exponent.

9.1. **The incoming label and exact auxiliary valuations.** Let $\widehat{F}$ be the lifted normalized cluster monomial. At a mutation in direction k from seed j, use the current coordinates $p_k = -\widehat{b}_{j,k}$ and the pulled-back form $\Omega_j = J_j^T \Omega_0 J_j$. If w is the incoming label in the new coordinates, the old label is

$$g_{\mathrm{old}} = Ew - [-w_k]_+ p_k. \tag{9.1}$$

By (8.9), $\Omega_j(p_k, Ew) = w_k$. The map $z \mapsto z + [-\Omega_j(p_k, z)]_+ p_k$ preserves the pairing with p_k, so its inverse subtracts this same multiple of p_k. Applying (9.1) backwards to the terminal exponent computes $g \in \widehat{L}$ in the initial coordinates. The formula holds componentwise in (8.14). From now on, p_i denotes the initial root.

Write

$$\iota : \mathbb{Z}^n \longrightarrow \widehat{L}, \qquad \iota(k) = \sum_i k_i p_i, \qquad C^+ = \iota(\mathbb{N}^n).$$

The incoming recurrence gives

$$\operatorname{supp} \widehat{F} \subset g + C^+. \tag{9.2}$$

By Laurent finiteness, the section with initial value $\widehat{F}$ is supported in a finite union $\mathcal{U}$ of translates of C^+. If $w \notin g + C^+$, every predecessor $w - ja$ in the incoming recurrence is also outside $g + C^+$. Its incoming datum is zero. Induction over the finite partially ordered set $\mathcal{U} \cap (w - C^+)$, simultaneously in all chambers, therefore makes every such coefficient zero. This proves (9.2).

The smallest auxiliary exponents are computed using an ordered series division ring.

Lemma 9.1 (Positive leading coefficients)**.** *Let $\mathcal{D}$ be an ordered division ring and let σ be an order-preserving automorphism. The skew Laurent-series division ring $\mathcal{D}((Z;\sigma))$, with $Za = \sigma(a)Z$ and exponents bounded below, is ordered by its first nonzero coefficient. An integral based quantum torus over $\mathbb{Q}(v)$ embeds into an iterated division ring of this kind in which every centered monomial is positive. For any specified central coordinate Y, the last series variable can be chosen to be Y.*

Proof. If two positive series have first terms $a_p Z^p$ and $b_q Z^q$, the first coefficient of their product is $a_p \sigma^p(b_q) > 0$. Addition preserves positivity: at different first exponents the earlier coefficient survives, and at equal exponents the two positive coefficients add. These rules also give trichotomy. A nonzero series is its invertible first monomial times $1 + h$, where h has strictly positive Z-valuation. The geometric series for $(1 + h)^{-1}$ is defined coefficientwise and has constant term one. This proves both the division assertion and the positivity of inverses of positive series.

Order $\mathbb{Q}(v)$ by the sign at positive infinity. In particular, every integer power of v is positive. Adjoin the coordinate variables successively as skew Laurent-series variables. At each step the required automorphism scales earlier coordinate variables by central powers of v. It extends coefficientwise to their series division rings and preserves the first exponent. By induction it preserves the sign of the recursively leading coefficient, since each scaling factor is a power of v. It is thus order-preserving. The quantum commutation relations give the embedding. It extends to the Ore division ring of fractions, since every nonzero torus element has become invertible. Centered monomials differ from ordered products of coordinate powers by a power of v, which is positive in this order. A central coordinate may be adjoined last, with identity automorphism. □

Fix j and place Y_j last in this construction. The resulting division ring is $\mathcal{D}((Y_j))$, where $\mathcal{D}$ is ordered. For a nonzero series, define

$$\nu_{Y_j}\left(\sum_r f_r Y_j^r\right) = \min\{r : f_r \neq 0\}.$$

Every cluster variable is positive in this ring: mutation uses sums, products, inverses and integer powers of v. Each of the two exchange summands is positive. If their Y_j-valuations agree, their leading coefficients in $\mathcal{D}$ are positive and their sum is nonzero.

Consequently the valuations $\ell_i = \nu_{Y_j}(X_i)$ of the current frame coordinates satisfy the exact recurrence

$$\ell'_k = -\ell_k + \min\left\{\sum_i [\widehat{b}_{ik}]_+\ell_i,\ \sum_i [-\widehat{b}_{ik}]_+\ell_i\right\}, \qquad \ell'_i = \ell_i \quad (i \neq k). \tag{9.3}$$

Initially the old coordinate valuations are zero and $\nu_{Y_j}(Y_h) = \delta_{jh}$. If the terminal exponent is a, set

$$L_j = \sum_i a_i \ell_i^{\text{terminal}}. \tag{9.4}$$

The centering scalar has valuation zero, so $L_j = \nu_{Y_j}(\widehat{F})$, also when frozen exponents are negative. Since $\widehat{F}$ is an actual finite Laurent polynomial, L_j is exactly the smallest Y_j-exponent among its nonzero terms. At any fixed Y_j-exponent the remaining based monomials are linearly independent in the residue quantum torus, which proves this last assertion.

For $w = g + \iota(k)$, the auxiliary coordinates are $w_{Y_i} = g_{Y_i} - k_i$. Combining (9.2)–(9.4) gives a complete finite box:

$$0 \le k_i \le K_i := g_{Y_i} - L_i, \qquad N = \sum_i K_i, \qquad \mathcal{I} = \{g + \iota(k) : 0 \le k_i \le K_i\}. \tag{9.5}$$

All K_i are nonnegative, since $\widehat{F} \neq 0$. Both their values and the uniqueness of the label k follow before any coefficients are computed. If $N = 0$, the incoming recurrence gives $\widehat{F} = X^g$; its only coefficient space is $\mathbb{Q}$ in degree zero. We henceforth suppose $N \ge 1$.

9.2. **A finite graph of local reorderings.** At positive root degree at most N there are finitely many roots, wall hyperplanes and cells. The proof of Theorem 6.3 constructs the required factorization on each ray and cell by following a generic line and applying Theorem 3.2 at its joints. Retain each chosen line traversal as an ordered list of nodes and join each query to its predecessor computation. These applications form a finite directed acyclic graph, ordered first by root degree and then by position along the line on the ray being constructed. Queries to other rays use smaller root degree. Indeed, a term of the current degree on a different ray could affect the target ray only through a nonlinear product, which has larger degree. Queries on the current ray refer only to earlier points of that same line. This proves acyclicity.

All these choices are effective over $\mathbb{Q}$. Enumerate rational points and lines with the prescribed rational directions until they avoid the finitely many forbidden hyperplanes and intersections in the truncated arrangement. Each choice is tested by integer linear algebra. The input leaves are the simple units $E_{p_i,-1}$. For a component, choose an integer

$$A \ge \max\left(1,\ N^2 \max_{i,j} |\Omega_0(p_i, p_j)|\right).$$

It bounds the pairings between all relevant roots. If the parameters of the elementary units entering a node are at most $K \ge 0$, the effective bound in Theorem 3.2 is

$$k_{\text{out}} \le NK + N + AN^2. \tag{9.6}$$

Propagating this bound from $K = 0$ at the leaves gives an integer $K_{*,c} \ge 0$ bounding every output parameter in component c.

There is a complementary finite procedure for any lower output threshold L. If a node is asked for parameters $k_{\text{out}} \ge L$, an input occurrence that participates in such an output must have

$$k_i \ge L - (N-1)K - N - AN^2. \tag{9.7}$$

The ordinary polynomial degree ℓ of the start of any required primitive string satisfies

$$0 \le \ell \le \left\lfloor \frac{N(K+1) + AN^2 - L}{2} \right\rfloor. \tag{9.8}$$

The degree estimate follows from the exact output parameter $u(d)-\chi(d,d)-2\ell$ and the inequalities $u(d) \le N(K+1)$ and $\chi(d,d) \ge -AN^2$. To obtain (9.7), isolate one active input copy in the inequality $k_{\mathrm{out}} \le \sum_i d_i k_i + N + AN^2$; at most $N-1$ other copies remain. These are the estimates proved with Theorem 3.2. A negative upper bound in (9.8) means that the node contributes nothing. Apply (9.7) recursively to predecessor nodes. Acyclicity makes the recursion terminate, and each retained input list is finite. Each occurrence keeps its own index even if several occurrences have the same root, parameter and type.

9.3. **A global energy bound.** In the simultaneous component system, set

$$\kappa_* = \max_c d_c K_{*,c}, \qquad T_* = \max_{\substack{a \in C^+ \cap \mathsf{N}_c,\ 0<|a|\le N \\ w \in \mathcal{I},\ c}} |\widehat{\Lambda}(a,w)|, \qquad U = N(\kappa_* + T_*). \tag{9.9}$$

Here $|a|$ is the sum of its coefficients in the initial positive root basis. The maximum is over finite sets. For example, it is enough to use the explicit upper bound

$$T_* \le N \max_i |\widehat{\Lambda}(p_i, g)| + N^2 \max_{i,j} |\widehat{\Lambda}(p_i, p_j)|.$$

The local parameter k in component c is measured in powers of $t_c = v^{d_c}$. Lemma 6.2 gives the exact v-energies of its centered slots at a source exponent w:

$$d_c k - |\widehat{\Lambda}(a,w)| + 2d_c h, \qquad 1 \le h \le |\widehat{\Lambda}(a,w)|/d_c. \tag{9.10}$$

The slot number is an integer by (8.14). At zero pairing the crossing contributes no nonconstant term. For an E unit each slot is chosen zero or one times; for an H unit its multiplicity is any nonnegative integer. Root-degree truncation bounds the latter multiplicities. Roots on one ray pair to zero, so the energies in (9.10) add under slot selection.

Expand the incoming recurrence as histories of positive jumps from g. Suppose a history ending at $w \in \mathcal{I}$ selects q nonconstant copies, with roots a_ν, parameters k_ν and components c_ν. Its prefix labels $w_\nu = g + \sum_{\mu=1}^{\nu} a_\mu$ lie in $\mathcal{I}$. Since each $|a_\nu| \ge 1$,

$$q \le \sum_{\nu=1}^{q} |a_\nu| = |w - g| \le N.$$

Let δ_ν be the slot shift in (9.10), evaluated at $w_{\nu-1}$. Then $|\delta_\nu| \le T_*$ and $d_{c_\nu} k_\nu \le \kappa_*$. The total energy satisfies

$$E = \sum_{\nu=1}^{q} (d_{c_\nu} k_\nu + \delta_\nu) \le q(\kappa_* + T_*) \le U. \tag{9.11}$$

The exact crossing and incoming identities express each Laurent coefficient of $\widehat{F}$ as the sum of these histories. Admissibility makes each specified energy coefficient finite, so (9.11) bounds every coefficient above by U.

Bar invariance supplies the other half of the bound. On the based quantum torus the anti-involution

$$\overline{v} = v^{-1}, \qquad \overline{X^a} = X^a, \qquad \overline{PQ} = \overline{Q}\,\overline{P} \tag{9.12}$$

extends to its division ring. If the coordinates of a toric frame are fixed, every centered monomial of that frame is fixed: reversing the ordered product contributes the inverse of the change in the centering power. The statement includes negative powers. Each mutated variable is a sum of two centered monomials, so induction along the mutation word gives $\overline{\widehat{F}} = \widehat{F}$. Write its finite initial Laurent expansion as $\widehat{F} = \sum_w c_w(v) X^w$. Linear independence of the based monomials yields

$$c_w(v) = c_w(v^{-1}).$$

Thus

$$[v^E]c_w(v) = 0 \qquad \text{unless } -U \le E \le U. \tag{9.13}$$

In a history of final energy $E \ge -U$, isolate a selected copy with parameter k in component c. There are at most $N-1$ other parameters, each with v-weight at most κ_*, and the sum of all slot shifts is at most NT_*. Thus

$$-U \le E \le d_c k + (N-1)\kappa_* + NT_*,$$

which gives

$$d_c k \ge L_* := -U - (N-1)\kappa_* - NT_*. \tag{9.14}$$

Query component c only for $k \ge \lceil L_*/d_c \rceil$, and propagate these queries backwards using (9.7). Equations (9.8) then give finite bounds for every local matrix calculation.

Every retained prefix has energy at least

$$N(L_* - T_*). \tag{9.15}$$

Indeed, each selected copy has energy at least $L_* - T_*$, this number is nonpositive, and there are at most N copies. Thus (9.15) and U bound all intermediate energies used in the enumeration.

9.4. **The finite polynomial matrices.** At a local node, use its finite occurrence list. For each dimension vector d of root charge at most N, use monomial symmetric functions in each variable pack as a basis of the space S_d. In the required ordinary degrees, the unstable ideal is the span of the explicit shuffle images

$$S_a * S_b, \qquad a + b = d, \quad a, b \ne 0, \quad \mu(a) > \mu(d).$$

In output degree ℓ, the input degrees sum to $\ell + \chi(a, b)$. Since $\chi(a, b) \le \sum_i a_i b_i \le N^2$, the necessary matrix domains have ordinary degrees at most the bound (9.8) plus N^2. Taking the quotient of these rational matrices gives B_d in all required degrees.

Write ω for the auxiliary weight. The independently centered restrictions define the finite filtration $F^\omega B_d$. In ordinary degree ℓ, the possible weights lie between

$$\min_P \sum_{r \in P} s(r) \quad \text{and} \quad \ell + s(d),$$

in half-integer steps, where P runs through the finitely many ordered decompositions of d on the slope. Restriction does not increase internal polynomial degree. Thus the filtration and its associated graded require only finite matrix operations.

The primitive subspace $\mathcal{P}_d$ is the common kernel of the proper coproduct symbols. To compute this kernel, one may test the leading restrictions before multiplying them by their nonzero rational prefactors in the independent centers. Such a prefactor is invertible over the rational function field of the centers and does not change a kernel. Expanding the remaining expression in center monomials and comparing its quotient-vector coefficients gives finite rational matrices in the stated degree range. The centers remain independent throughout this calculation.

Finally impose the kernel of $D_d = \sum_{i,\alpha} \partial/\partial x_{i,\alpha}$, whose target has ordinary degree $\ell - 1$. The resulting actual multiplicity spaces are

$$W_{d,\ell,\omega} = (\mathcal{P}_d \cap \ker D_d)_{\ell,\omega}. \tag{9.16}$$

By Lemma 4.6, a basis vector of this space gives one output unit with root $p(d)$ and parameter

$$k = u(d) - \chi(d, d) - 2\ell.$$

Its type is E if $\chi(d, d)$ is odd and H if it is even. Choose an ordered basis by rational row reduction, with a fixed ordering of all input monomials. This gives the individual occurrence labels needed at subsequent nodes. The full dimension vector and the separate occurrence labels are retained after projection to roots.

9.5. **Decorated histories and specialization.** Fix the paths used in (6.11) for every target exponent $w \in \mathcal{I}$ and chamber I of the truncated arrangement. The indices i, normals a_i and source chambers I_i below refer to the fixed path for (w, I). Retain the unit parameters in (9.14) and the propagated predecessor ranges. For a summand (i, j) of that recurrence, put $u = w - ja_i$ and let $\mathsf{S}_i(u, j; e)$ be the set of nonempty slot selections contributing exponent ja_i and energy e to the jump coefficient $J_i(u, j)$. Each selected occurrence is labeled by the chosen basis of (9.16). An E slot permits a subset of these labels and an H slot permits a multiset. Their energies are given by (9.10).

For $w \in \mathcal{I}$ and $N(L_* - T_*) \le E \le U$, define the history sets recursively by

$$\mathsf{H}(w, I; E) = \mathsf{H}_0(w; E) \;\sqcup \bigsqcup_{\substack{i,\ j>0,\ e\in\mathbb{Z} \\ u=w-ja_i\in\mathcal{I}}} \mathsf{H}(u, I_i; E - e) \times \mathsf{S}_i(u, j; e), \tag{9.17}$$

where I_i is the source chamber of the crossing and

$$\mathsf{H}_0(w; E) = \begin{cases} \{\varnothing\}, & w = g,\ E = 0, \\ \varnothing, & \text{otherwise.} \end{cases}$$

The element $\varnothing$ in the first case denotes the empty history. Each alternative in (9.17) retains its label (i, j); the product concatenates the earlier history and its slot selection. Set $\mathsf{H}(w, I; E) = \varnothing$ if $w \notin \mathcal{I}$ or $E \notin [N(L_* - T_*), U]$. Since $j > 0$, each recursive call strictly decreases root degree. The recursion is therefore finite. The exponent and chamber sets, unit lists and slot sets are finite, and every history contains at most N copies.

Let I_- be the all-negative chamber. Define $\widehat{\mathcal{V}}_{k,E}$ as the free rational vector space on $\mathsf{H}(g + \iota(k), I_-; E)$. Thus the empty history gives $\widehat{\mathcal{V}}_{0,0} = \mathbb{Q}$. Products and disjoint unions in (9.17) give tensor products and direct sums of these spaces. The slot choices give the exterior and symmetric powers of the multiplicity spaces for E and H, respectively.

At each local node, the primitive-string factorization identifies the dimensions of (9.16) with the elementary factor multiplicities. At each edge, Lemma 6.2 identifies slot enumeration with the exact centered jump coefficient. Induction in root degree in the incoming recurrence expresses the actual coefficient as the sum over decorated histories. Equations (9.14) and (9.15) retain every history in the final window, so

$$\widehat{F} = \sum_{\substack{k\in\mathbb{N}^n \\ k_i\le K_i\ (1\le i\le n)}} \sum_{E=-U}^{U} \dim_{\mathbb{Q}} \widehat{\mathcal{V}}_{k,E}\, v^E X^{g+\iota(k)}. \tag{9.18}$$

Equation (9.13) gives vanishing outside this window. The component actions in (8.14) act on the whole centered monomial, retaining its original normalization in (9.18).

Proof of Theorem 1.2. Let $\mathrm{pr} : \widehat{L} \longrightarrow \mathbb{Z}^m$ delete the auxiliary coordinates. For $-U \le E \le U$, define

$$\mathcal{V}_{b,E} = \bigoplus_{\substack{k\in\mathbb{N}^n,\ k_i\le K_i\ \forall i \\ \mathrm{pr}(g+\iota(k))=b}} \widehat{\mathcal{V}}_{k,E}, \qquad \mathcal{V}_b = \bigoplus_{E=-U}^{U} \mathcal{V}_{b,E}. \tag{9.19}$$

Set $\mathcal{V}_{b,E} = 0$ for $|E| > U$. The sums are finite and their summands are constructed by rational linear algebra and finite choices. Apply the homomorphism (8.12) to (9.18). The specialization argument in Section 8 identifies the left side with the required original cluster monomial. The right side is exactly (1.3). The label recursion, tropical valuations, finite reordering graph and energy bounds are computed before the coefficient spaces, and every subsequent operation terminates. If the mutable set is empty, take the single initial monomial space $\mathbb{Q}$ in energy zero. This proves the theorem. □

Example 9.2. For the seed of Example 2.3, adjoin the single central coordinate Y and mutate the mutable variable once. In the extended exponent lattice,

$$p = -2e_2 - e_Y, \qquad g = -e_1 + 2e_2 + e_Y.$$

The inverse incoming-label rule gives g, and the tropical recurrence gives $L_Y = 0$. Hence $K_Y = 1$ and the full exponent box is $\{g, g+p\}$. The component constant is $d = 2$ and $\widehat{\Lambda}(p, g) = 2$. The simple unit has parameter $k = -1$, so its sole slot in (9.10) has energy $2(-1) - 2 + 2 \cdot 2 = 0$. The empty history and the history selecting this slot give, respectively,

$$\widehat{\mathcal{V}}_{0,0} = \mathbb{Q}, \qquad \widehat{\mathcal{V}}_{1,0} = \mathbb{Q}, \qquad \widehat{F} = X^g + X^{g+p}.$$

All other spaces vanish. Specializing $Y = 1$ gives the two coefficient spaces for $M(-e_1 + 2e_2) + M(-e_1)$. The original nonprimitive exchange column and the frozen pairing $\Lambda(e_3, e_4) = 1$ have both been retained.

References

[1] A. Berenstein and A. Zelevinsky, *Quantum cluster algebras*, Adv. Math. **195** (2005), 405–455; arXiv:math/0404446v2.

[2] B. Davison, *Positivity for quantum cluster algebras*, Ann. of Math. (2) **187** (2018), 157–219. doi:10.4007/annals.2018.187.1.3.

[3] B. Davison and T. Mandel, *Strong positivity for quantum theta bases of quantum cluster algebras*, Invent. Math. **226** (2021), 725–843. doi:10.1007/s00222-021-01061-1.

[4] B. Davison, D. Maulik, J. Schürmann and B. Szendrői, *Purity for graded potentials and quantum cluster positivity*, Compos. Math. **151** (2015), 1913–1944. doi:10.1112/S0010437X15007332.

[5] B. Davison and S. Meinhardt, *Cohomological Donaldson–Thomas theory of a quiver with potential and quantum enveloping algebras*, Invent. Math. **221** (2020), 777–871. doi:10.1007/s00222-020-00961-y.

[6] A. I. Efimov, *Quantum cluster variables via vanishing cycles*, preprint (2011), arXiv:1112.3601v2.

[7] A. I. Efimov, *Cohomological Hall algebra of a symmetric quiver*, Compos. Math. **148** (2012), 1133–1146. doi:10.1112/S0010437X12000152.

[8] S. Fomin and A. Zelevinsky, *Cluster algebras I: Foundations*, J. Amer. Math. Soc. **15** (2002), 497–529; arXiv:math/0104151.

[9] H. Franzen and M. Reineke, *Semi-stable Chow–Hall algebras of quivers and quantized Donaldson–Thomas invariants*, Algebra Number Theory **12** (2018), 1001–1025. doi:10.2140/ant.2018.12.1001.

[10] M. Gross, P. Hacking, S. Keel and M. Kontsevich, *Canonical bases for cluster algebras*, J. Amer. Math. Soc. **31** (2018), 497–608; arXiv:1411.1394.

[11] M. Huang, *Positivity for quantum cluster algebras from orbifolds*, preprint (2024), arXiv:2406.03362v1.

[12] Y. Kimura and F. Qin, *Graded quiver varieties, quantum cluster algebras and dual canonical basis*, Adv. Math. **262** (2014), 261–312. doi:10.1016/j.aim.2014.05.014.

[13] M. Kontsevich and Y. Soibelman, *Cohomological Hall algebra, exponential Hodge structures and motivic Donaldson–Thomas invariants*, Commun. Number Theory Phys. **5** (2011), 231–352; arXiv:1006.2706v2.

[14] K. Lee and R. Schiffler, *Positivity for cluster algebras*, Ann. of Math. (2) **182** (2015), 73–125. doi:10.4007/annals.2015.182.1.2.

[15] G. Musiker, R. Schiffler and L. Williams, *Positivity for cluster algebras from surfaces*, Adv. Math. **227** (2011), 2241–2308; arXiv:0906.0748.

[16] K. Nagao, *Donaldson–Thomas theory and cluster algebras*, Duke Math. J. **162** (2013), 1313–1367. doi:10.1215/00127094-2142753.

[17] H. Nakajima, *Quiver varieties and cluster algebras*, Kyoto J. Math. **51** (2011), 71–126. doi:10.1215/0023608X-2010-021.

[18] OpenAI, *Elementary positivity of chromatic quasisymmetric functions*, September 24, 2026, Sections 3–4 and 6. OpenAI Math Release, fixed repository commit `adc7f1241b42`.

[19] F. Qin, *Quantum cluster variables via Serre polynomials*, J. Reine Angew. Math. **668** (2012), 149–190. doi:10.1515/CRELLE.2011.129.

[20] F. Qin, *t-Analog of q-characters, bases of quantum cluster algebras, and a correction technique*, Int. Math. Res. Not. IMRN **2014**, no. 22, 6175–6232. doi:10.1093/imrn/rnt115. arXiv:1207.6604v3.

[21] P. Sherman and A. Zelevinsky, *Positivity and canonical bases in rank 2 cluster algebras of finite and affine types*, Mosc. Math. J. **4** (2004), 947–974; arXiv:math/0307082.

[22] T. Tran, *F-polynomials in quantum cluster algebras*, Algebr. Represent. Theory **14** (2011), 1025–1061. doi:10.1007/s10468-010-9226-6.

[23] H. Ye, *Quantum generalized cluster algebras from scattering diagrams: positivity and mutation*, preprint (2026), arXiv:2611.00991v1.

Zhili College, Tsinghua University, Beijing 100084, China
Email address: `tangqy24@mails.tsinghua.edu.cn`