hexagon:2611.00279v3

ANOTHER PROOF OF QUANTUM LAURENT POSITIVITY, II

QIYUE TANG

Abstract. We give a further combinatorial proof of Laurent positivity for skew-symmetric quantum cluster algebras.

Contents

1. Introduction 2
1.1. Main result 4
2. Quantum tori and seeds 5
2.1. Based monomials 5
2.2. Compatibility and mutation 6
3. Polynomial modules on fans 7
3.1. Stalks, sections, and minimal extensions 7
3.2. The algebraic fan theorems used below 8
4. Graph cones and local degree bounds 9
4.1. The flow subdivision and its reduced ring 10
4.2. The cographic quotient 11
4.3. Faces and contractions 12
5. Equivariant graph characters 14
5.1. Canonical actions and the cellular formula 14
5.2. The cut-algebra character 16
5.3. Two partition identities 17
5.4. The contraction recursion 18
5.5. Independent permutations of opposite arrows 19
6. Transportation polytopes and primitive characters 21
6.1. Symmetric functions and cycle indices 21
6.2. Thin supports and their factorization 22
6.3. The positive-flow core 24
6.4. The balanced-partition identity 26
7. Positive quantum strings 29
7.1. Elementary and complete strings 29
7.2. Finite reordering 31

Date: November 10, 2026.

2020 *Mathematics Subject Classification.* Primary 13F60; Secondary 05E10, 52B20.

Key words and phrases. Quantum cluster algebra, Laurent positivity, scattering diagram, graph, combinatorial intersection cohomology.

8. Admissible lists and positive reordering 32
8.1. Energy bounds 32
8.2. Auxiliary degrees and compatible atom lists 33
8.3. The local positive reordering theorem 34
9. Finite wall systems and positive coefficients 34
9.1. The total element and its incoming prescription 35
9.2. Positive crossing formulas 37
9.3. Incoming coefficients and a decreasing recursion 39
10. Change of roots and equality of Laurent coefficients 40
10.1. A uniform bound along the mutation direction 40
10.2. Comparison in a common rational ring 42
10.3. A monomial with one incoming coefficient 44
11. Compatible quantum seeds 44
11.1. Central coefficients and an integral root model 45
11.2. The exact exchange relation 46
11.3. Transport of an actual cluster monomial 47
11.4. Specialization and several components 48
References 49

1. Introduction

Fomin and Zelevinsky introduced cluster algebras as an algebraic framework for total positivity and canonical bases [9]. Starting from a seed, repeated exchange relations produce a mutation pattern. Their Laurent phenomenon [9, Theorem 3.1] states that every cluster variable is a Laurent polynomial in the variables of every seed. Immediately after that theorem they conjectured nonnegativity of its coefficients. The exchange relations are subtraction-free, but positivity of the simplified Laurent coefficients requires a separate argument.

Berenstein and Zelevinsky defined quantum cluster algebras by coupling an exchange matrix with a compatible alternating form [2]. The variables of a quantum seed commute up to powers of a parameter, and the associated quantum torus has a distinguished basis of normalized monomials. Their quantum Laurent phenomenon [2, Corollary 5.2] gives Laurent expansions with coefficients in $\mathbb{Z}[v^{\pm 1}]$, where v is the square root of the usual quantum parameter. Quantum positivity asserts that these coefficients belong to $\mathbb{N}[v^{\pm 1}]$, where $\mathbb{N} = \{0, 1, 2, \dots\}$. Specialization at $v = 1$ gives the corresponding classical statement.

Several early approaches made positivity explicit in particular classes. Sherman and Zelevinsky constructed positive bases for rank-two cluster algebras of finite and affine types [26]. Musiker, Schiffler and Williams obtained expansion formulas for cluster algebras from surfaces [18]. Lee and Schiffler proved classical positivity for all skew-symmetric cluster algebras [15]. Gross, Hacking, Keel and Kontsevich developed theta functions and

canonical bases using scattering diagrams [10]. In their construction, broken lines record changes of monomials across walls, and positivity of wall functions controls the signs in Laurent expansions. Huang extended surface expansion formulas to quantum cluster algebras, proving positivity for unpunctured surfaces [12], and obtained further orbifold formulas [13]. Burcroff and Lee constructed quantum-weighted bijections between compatible pairs and broken lines in type A_2 and the Kronecker case [4].

Representation theory supplied graded interpretations of quantum coefficients. Nakajima related cluster structures to graded quiver varieties [20]. Qin expressed quantum cluster variables through Serre polynomials of quiver Grassmannians and proved positivity for expansions in acyclic seeds [22]. Kimura and Qin proved quantum positivity throughout mutation classes containing an acyclic seed [14]. These results identify coefficients with graded dimensions and make the dependence on the quantum parameter part of the positivity problem.

Donaldson–Thomas theory provided a method for treating general mutation sequences. Nagao related cluster transformations to wall crossing for quiver representations [19]. Efimov expressed quantum cluster variables using vanishing cycles and reduced their positivity to purity statements [8]. Davison, Maulik, Schürmann and Szendrői established the required purity for quivers admitting a nondegenerate potential homogeneous of positive degree [7]. Davison subsequently proved quantum Laurent positivity for every skew-symmetric quantum cluster algebra [5, Theorem 2.4]. His proof uses purity of vanishing-cycle cohomology and also establishes Lefschetz and parity properties of the Laurent coefficients.

Davison and Mandel gave a second approach through quantum scattering diagrams and theta bases [6]. Their local positivity theorem is obtained from Donaldson–Thomas integrality and positivity, and their scattering construction propagates it to arbitrary rank. Their quantum theta bases have positive multiplication constants and contain cluster monomials. Laurent positivity is one consequence. Positive rank-two quantum dilogarithm factorizations also appear in Reineke's work [23], through quiver-moduli intersection cohomology.

The graph constructions used below have several precedents. Hausel and Sturmfels described Lawrence toric cohomology by cographic Stanley–Reisner rings modulo cycle relations [11, Theorem 8.3]. Reineke, Rhoades and Tewari related graphical zonotopal algebras, their graded automorphism-group representations and symmetric-quiver Donaldson–Thomas invariants [24]. Szenes and Trapeznikova gave acyclic-graph models, partition recursions and cut-product presentations for type-A toric intersection cohomology [27]. Thin quiver representations and partition formulas also occur in the abelianization results of Reineke, Stoppa and Weist [25] and Mozgovoy and Reineke [17]. In particular, the latter derive Harder–Narasimhan products and exponential formulas for abelian primitive invariants. Meinhardt and Reineke identify Hodge-theoretic Donaldson–Thomas invariants of zero-potential quivers with

the compactly supported intersection cohomology of the closure of the stable locus, for stability conditions whose Euler form is symmetric on dimension vectors of each fixed slope [16, Theorem 1.1].

The fan modules belong to the combinatorial intersection cohomology theory of Barthel, Brasselet, Fieseler and Kaup [1] and Bressler and Lunts [3]. They are constructed recursively from graded polynomial modules on cones. We state the decomposition, cellular and duality results used below. For a symmetric graph cone, a flow subdivision and the cographic quotient give a strict bound on the degrees of its local minimal-extension module. Graph character identities then identify that module, including the actions obtained by independently permuting the two sets of opposite arrows.

Our local coefficient space is the reduced minimal-extension module of a complete transportation fan. A semistable support splits into its maximal positive-flow core and an acyclic graph between balanced components. Summing over invariant partitions identifies the full plethystic logarithm of the factorization series with the Frobenius characteristic of this finite graded representation. The identity holds at every permutation and includes all block cycles and their return maps. These terms determine the signed Adams substitutions and the multiple-cover coefficients. Schur multiplicities then give the positive quantum factors. This character identity and its factorization application are the local steps of the proof.

The earlier paper [28] studies the local factorization through polynomial shuffle quotients, a signed coproduct and primitive strings. The present argument obtains the local multiplicities from the graph and fan modules above. The incoming-coefficient transport and change of roots follow [21, Section 4], as applied to quantum cluster mutation in [28]. We develop this formal framework below and use the graph and fan calculation as its positivity input.

1.1. **Main result.** Let $\widetilde{B}$ be an $m \times n$ integer matrix with skew-symmetric principal part, and let Λ be an integer alternating matrix satisfying

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

We use the normalization $M(a)M(b) = v^{\Lambda(a,b)}M(a+b)$ for the based monomials in a toric frame. A normalized cluster monomial in a seed M' is $M'(a)$ with $a_i \geq 0$ at mutable indices; its frozen exponents may be arbitrary integers.

Theorem 1.1 (Quantum Laurent positivity)**.** *For every compatible quantum seed with skew-symmetric principal exchange matrix, every seed M' obtained by a finite sequence of mutations, and every normalized cluster monomial $M'(a)$, the expansion in the initial toric frame*

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

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

Rational compatible forms are included after a change of parameter. Choose $r \in \mathbb{Z}_{>0}$ such that $r\Lambda$ is integral. With $w = v^{1/r}$ and the integral form $r\Lambda$, the coefficients belong to $\mathbb{N}[w^{\pm 1}]$.

The theorem gives Laurent positivity for arbitrary skew-symmetric mutation classes, including frozen variables and the general positive diagonal compatibility matrix in the Berenstein–Zelevinsky definition. The local factorization is first constructed in a completion with coefficient field $\mathbb{Q}((t^{-1}))$. A coefficient recursion propagates nonnegativity, and an exact mutation identity identifies the transported sections with cluster monomials. A central coefficient extension compares this integral model with the given compatible quantum seed. Finally, the quantum Laurent phenomenon supplies finite Laurent support.

2. Quantum tori and seeds

Lattice vectors are columns. For an integer a, write $[a]_+ = \max(a, 0)$; this notation applies coordinatewise to a vector. All frozen variables are invertible.

2.1. **Based monomials.** Let L be a lattice and let $\Lambda : L \times L \to \mathbb{Z}$ be alternating. The based quantum torus $\mathcal{T}_\Lambda$ is the free $\mathbb{Z}[v^{\pm 1}]$-module with basis $\{X^a : a \in L\}$ and multiplication

$$X^a X^b = v^{\Lambda(a,b)} X^{a+b}. \tag{2.1}$$

Bilinearity gives associativity, and X^{-a} is the inverse of X^a. If $e_1, \dots, e_m$ is a basis and $X_i = X^{e_i}$, then

$$X^a = v^{-\sum_{i<j} a_i a_j \Lambda(e_i,e_j)} X_1^{a_1} \cdots X_m^{a_m}, \qquad X_i X_j = v^{2\Lambda(e_i,e_j)} X_j X_i. \tag{2.2}$$

The torus is an Ore domain and embeds in its skew field of fractions $\mathcal{F}$ over $\mathbb{Q}(v)$; see [2, Section 4.1 and Appendix]. A *toric frame* is a map $M : \mathbb{Z}^m \to \mathcal{F}$ obtained from the based monomials by a lattice isomorphism and a skew-field automorphism. Thus, for an integer alternating form Λ_M,

$$M(a)M(b) = v^{\Lambda_M(a,b)} M(a + b),$$

and its monomials form a basis of the corresponding Laurent torus.

2.2. **Compatibility and mutation.** Fix $n \le m$. An extended exchange matrix is an integer matrix $\widetilde{B} = (b_{ij}) \in \mathrm{Mat}_{m\times n}(\mathbb{Z})$. Its principal part $B = (b_{ij})_{1\le i,j\le n}$ will always be skew-symmetric. A *compatible quantum seed* $(M, \widetilde{B})$ satisfies

$$\widetilde{B}^{\,T}\Lambda_M = (D \;\; 0), \qquad D = \mathrm{diag}(d_1, \dots, d_n), \quad d_i > 0. \tag{2.3}$$

The indices $1, \dots, n$ are mutable; the others are frozen.

Mutation in direction $k \le n$ changes the exchange matrix to

$$b'_{ij} = \begin{cases} -b_{ij}, & i = k \text{ or } j = k, \\ b_{ij} + [b_{ik}]_+ b_{kj} + b_{ik}[-b_{kj}]_+, & i \ne k,\ j \ne k. \end{cases} \tag{2.4}$$

If $\widetilde{b}_k$ is the kth column of $\widetilde{B}$, the new frame is determined by

$$\begin{aligned} M'(e_i) &= M(e_i) \quad (i \ne k), \\ M'(e_k) &= M(-e_k + [\widetilde{b}_k]_+) + M(-e_k + [-\widetilde{b}_k]_+). \end{aligned} \tag{2.5}$$

Define the integral matrix E by

$$Ee_k = -e_k + [\widetilde{b}_k]_+, \qquad Ee_i = e_i \quad (i \ne k). \tag{2.6}$$

Then $\Lambda_{M'} = E^T \Lambda_M E$. The entries with one index k are the same whether $[\widetilde{b}_k]_+$ or $[-\widetilde{b}_k]_+$ is used, since $\Lambda_M(\widetilde{b}_k, e_i) = d_k\delta_{ki}$. This also shows that the two summands in (2.5) have the same commutation factor with $M(e_i)$ for $i \ne k$. To check independence, let $\mathcal{D}_0$ be the division subfield generated by $\mathbb{Q}(v)$ and the $M(e_i)$ with $i \ne k$. Writing $X_k = M(e_k)$, the exchange expression has the form $M'(e_k) = X_k^{-1}R$ with $0 \ne R \in \mathcal{D}_0$. For every $r \in \mathbb{Z}$, the power $M'(e_k)^r$ is X_k^{-r} times a nonzero element of $\mathcal{D}_0$. Distinct powers are therefore linearly independent over $\mathcal{D}_0$, and $X_k = R\, M'(e_k)^{-1}$ shows that the fraction field is unchanged. The new generators thus define a toric frame, whose arbitrary based monomials are obtained by (2.2) with M' in place of X. These definitions are the conventions of [2, Sections 3–4].

The quantum cluster algebra is generated by all mutable variables obtained by finite sequences of these mutations, together with the invertible frozen variables, over $\mathbb{Z}[v^{\pm 1}]$. A *normalized cluster monomial* is $M'(a)$ in some seed, where $a_i \ge 0$ for $i \le n$ and $a_i \in \mathbb{Z}$ for $i > n$.

Theorem 2.1 (Quantum Laurent phenomenon, [2])**.** *Every normalized cluster monomial has a finite expansion in the based monomials of every seed, with coefficients in* $\mathbb{Z}[v^{\pm 1}]$.

Lemma 2.2. *The compatibility constants* d_i *are constant on each connected component of the graph with vertices* $1, \dots, n$ *and edges* $\{i, j\}$ *for* $b_{ij} \ne 0$.

Proof. Multiplication of (2.3) by $\widetilde{B}$ gives $\widetilde{B}^T\Lambda_M\widetilde{B} = DB$. The left side is alternating. Since $B^T = -B$, this implies $DB = BD$, or $(d_i - d_j)b_{ij} = 0$ for every i, j. Equality propagates along each edge path. □

Example 2.3. Take $B = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$ and $\Lambda_M = B$, so $B^T\Lambda_M = I_2$. Mutation at 1 gives

$$M'(e_1) = M(-e_1) + M(-e_1 + e_2).$$

The coefficient of each based monomial is 1. In ordinary ordered variables the second summand is $vX_1^{-1}X_2$, by (2.2).

3. Polynomial modules on fans

All vector spaces and polynomial rings in this section are over $\mathbb{Q}$. The polynomial-module constructions below are defined over $\mathbb{Q}$; their stated exactness, freeness and duality properties may equivalently be checked after the faithful scalar extension $\mathbb{Q} \subset \mathbb{R}$. A cone in a finite-dimensional vector space W means a strongly convex rational polyhedral cone in $W_{\mathbb{R}}$. A fan Δ is a finite collection of cones, closed under taking faces, such that the intersection of two cones is a face of each. We write $\tau \preceq \sigma$ for the face relation, $W_\sigma = \operatorname{span}_{\mathbb{Q}}(\sigma \cap W)$, and $\partial\sigma$ for the fan of proper faces of σ. The fan is complete if its support is $W_{\mathbb{R}}$. When its support is a full-dimensional convex cone, its boundary fan consists of the cones contained in the topological boundary of the support. The boundary of a complete fan is empty. A fan is pure of dimension d if every maximal cone has dimension d.

3.1. Stalks, sections, and minimal extensions.

Put

$$A_W = \operatorname{Sym}(W^*), \qquad A_\sigma = \operatorname{Sym}(W_\sigma^*).$$

A linear form has degree one throughout the fan calculation. Thus the variable q records half the usual cohomological degree. Restriction of linear forms defines a surjection $A_\sigma \to A_\tau$ whenever $\tau \preceq \sigma$.

A sheaf of graded polynomial modules on Δ is a collection of finitely generated graded A_σ-modules $\mathcal{F}_\sigma$, together with compatible degree-preserving restriction maps $\mathcal{F}_\sigma \to \mathcal{F}_\tau$. For a subfan Γ, its sections are the compatible tuples

$$\mathcal{F}(\Gamma) = \left\{ (f_\sigma) \in \prod_{\sigma \in \Gamma} \mathcal{F}_\sigma : f_\sigma|_\tau = f_\tau \text{ if } \tau \preceq \sigma \right\}.$$

For $\Gamma' \subseteq \Gamma$, set $\mathcal{F}(\Gamma, \Gamma') = \ker(\mathcal{F}(\Gamma) \to \mathcal{F}(\Gamma'))$. On the fan of faces of one cone, $\mathcal{F}(\sigma)$ is its stalk $\mathcal{F}_\sigma$.

The sheaf is *pure* if every stalk is a finite free graded A_σ-module and every restriction $\mathcal{F}_\sigma \to \mathcal{F}(\partial\sigma)$ is surjective. The latter condition implies surjectivity of restriction to every subfan: extend a section one cone at a time, in increasing dimension. For a graded A_σ-module M, write

$$\overline{M} = M/A_\sigma^{>0}M. \tag{3.1}$$

If M is a module of sections on a full-dimensional fan in W, the reduction is taken by $A_W^{>0}$.

Definition 3.1. A normalized minimal-extension sheaf $\mathcal{L}$ on Δ is a pure sheaf such that $\mathcal{L}_0 = \mathbb{Q}$ in degree zero and, for every nonzero cone σ, restriction induces an isomorphism

$$\overline{\mathcal{L}_\sigma} \xrightarrow{\sim} \mathcal{L}(\partial\sigma)/A_\sigma^{>0}\mathcal{L}(\partial\sigma). \tag{3.2}$$

This definition gives an explicit recursive construction. Having defined $\mathcal{L}$ on the proper faces of σ, choose homogeneous lifts of a basis of the quotient on the right of (3.2). Take the free A_σ-module on these lifts and map it onto the boundary module. Graded Nakayama's lemma gives surjectivity, and the resulting map is a minimal free cover. Induction constructs $\mathcal{L}$ on the whole fan. Lifting maps between these free covers shows that two such sheaves are isomorphic by an isomorphism that is the identity at the zero cone.

For a simplicial cone the polynomial sheaf $\mathcal{A}$, with $\mathcal{A}_\sigma = A_\sigma$, is minimal. Choose its ray coordinates. Inclusion–exclusion over the coordinate facets extends every compatible boundary tuple to a polynomial. A polynomial vanishing on every facet is divisible by the product of those coordinates, so the restriction kernel lies in the positive-degree ideal. It follows by induction that $\mathcal{A}$ is a minimal extension on a simplicial fan. Its sections are the piecewise polynomial functions on the fan.

3.2. **The algebraic fan theorems used below.** We use the following facts from the polynomial-module theory of fans. They are the decomposition, cellular, and duality theorems of [1, Theorems 2.3–2.4, Corollary 2.5, Theorems 3.8–3.9, 3.14 and 4.3, Corollary 3.15, and Theorem 5.3]; see also [3, Theorems 4.7 and 5.5 and Proposition 7.8].

Theorem 3.2 (Algebraic fan theorems)**.** *Let Δ be a pure d-dimensional fan in W, where $\dim W = d$. Assume that Δ is complete or has full-dimensional convex support, and let $\mathcal{L}$ be a normalized minimal-extension sheaf.*

(i) *The modules $\mathcal{L}(\Delta)$ and $\mathcal{L}(\Delta, \partial\Delta)$ are finite free graded A_W-modules.*

(ii) *If $\pi : \widetilde{\Delta} \to \Delta$ is a subdivision and $\widetilde{\mathcal{L}}$ is pure, its direct image $(\pi_*\widetilde{\mathcal{L}})_\sigma = \widetilde{\mathcal{L}}(\pi^{-1}\sigma)$ is pure. In particular, the direct image of the polynomial sheaf of a simplicial subdivision has $\mathcal{L}$ as a graded direct summand.*

(iii) *The augmented cellular complexes of $\mathcal{L}$ are exact. The complex with terms over all cones resolves $\mathcal{L}(\Delta, \partial\Delta)$; the complex with terms over cones not in $\partial\Delta$ resolves $\mathcal{L}(\Delta)$. A cone τ of codimension c contributes its stalk tensored with the orientation line of W/W_τ, in cellular degree c.*

(iv) *There is an intersection product and an evaluation map giving a perfect graded A_W-bilinear pairing*

$$\mathcal{L}(\Delta) \otimes_{A_W} \mathcal{L}(\Delta, \partial\Delta) \longrightarrow A_W[-d]. \tag{3.3}$$

Here the generator of $A_W[-d]$ has degree d. The product is constructed recursively by free-module lifts and complements; every choice permitted in that construction gives a perfect pairing.

The freeness in part (i) is often called quasiconvexity. We will apply the theorem only to the complete and convex-support cases explicitly listed. In part (iii), choices of orientations give the usual signed restriction maps; the orientation lines retain their natural transformation under a linear automorphism of the fan.

We record the direct-summand construction because reduction of that summand will supply the local degree bound. Suppose $\mathcal{F}$ is pure and $\mathcal{F}_0 = \mathbb{Q}$. Inductively construct maps $\alpha : \mathcal{L} \to \mathcal{F}$ and $\beta : \mathcal{F} \to \mathcal{L}$ with $\beta\alpha = 1$. At a cone σ, freeness of the source and surjectivity onto the boundary permit both maps to be lifted from the proper faces. On reduction, the composite $\beta_\sigma\alpha_\sigma$ is the identity: its boundary restriction is the identity, and (3.2) identifies the reduced stalk with the reduced boundary. A degree-zero endomorphism of a finite free graded module that is invertible on reduction is invertible. To see this directly, order a homogeneous basis by degree. Its matrix is block triangular, and its diagonal blocks are the constant invertible reduced blocks. Replacing β_σ by $(\beta_\sigma\alpha_\sigma)^{-1}\beta_\sigma$ preserves its boundary restriction and makes the composite the identity. This constructs a splitting on the whole fan. Consequently,

$$\overline{\mathcal{L}_\sigma} \text{ is a graded direct summand of } \overline{\mathcal{F}_\sigma}. \tag{3.4}$$

Lemma 3.3 (Smooth factors)**.** *If σ is a cone and $\eta = \mathbb{R}^u_{\geq 0}$ is a simplicial orthant, then*

$$\overline{\mathcal{L}}_{\sigma\times\eta} \cong \overline{\mathcal{L}}_\sigma.$$

For a chosen compatible action on $\mathcal{L}_\sigma$, the product construction carries the induced action of automorphisms of σ and permutations of the orthant coordinates. The latter act trivially on the reduced factor.

Proof. It suffices to adjoin one coordinate z. On a product face put $\mathcal{F}_{\tau\times\mathbb{R}_{\geq 0}} = \mathcal{L}_\tau[z]$ and $\mathcal{F}_{\tau\times\{0\}} = \mathcal{L}_\tau$. A boundary section of $\tau \times \mathbb{R}_{\geq 0}$ is a pair

$$(a(z), b) \in \mathcal{L}(\partial\tau)[z] \times \mathcal{L}_\tau, \qquad a(0) = b|_{\partial\tau}.$$

Write $a(z) - b|_{\partial\tau} = zc(z)$ and lift each coefficient of c to $\mathcal{L}_\tau$. Then $b + z\widetilde{c}(z)$ lifts the pair. The kernel of this boundary restriction is $z\mathcal{L}(\tau, \partial\tau)[z]$. It lies in the positive-degree ideal times $\mathcal{L}_\tau[z]$, so the free cover is minimal. Thus $\mathcal{F}$ is a normalized minimal extension. Reduction removes z. Iteration proves the assertion, including the stated actions. □

4. Graph cones and local degree bounds

Let G be a connected loopless multigraph with vertex set V, edge set E, $s = |V|$, and $R = |E|$. Replace each edge by two opposite arrows and denote the arrow set by $\vec{E}$. The incidence map and its transpose are

$$\partial : \mathbb{Q}^{\vec{E}} \longrightarrow \mathbb{Q}^V_0, \qquad \partial(e_{i\to j}) = e_i - e_j, \qquad \partial^T : \mathbb{Q}^V/\mathbb{Q}\mathbf{1} \longrightarrow \mathbb{Q}^{\vec{E}},$$

where $\mathbb{Q}^V_0$ is the subspace with coordinate sum zero. Set

$$N_G = \mathbb{Q}^{\vec{E}}/\operatorname{im}\partial^T, \qquad \sigma_G = \operatorname{image}(\mathbb{R}^{\vec{E}}_{\geq 0} \to (N_G)_\mathbb{R}). \tag{4.1}$$

The dual cone is the nonnegative circulation cone

$$C_G = \{x \in \mathbb{R}^{\vec{E}}_{\geq 0} : \partial x = 0\}, \qquad d_G = \dim \sigma_G = 2R - s + 1. \tag{4.2}$$

The quotient map shows that σ_G spans N_G. A potential whose differences are nonnegative on both directions of every edge is constant, so $\operatorname{im} \partial^T \cap \mathbb{R}^{\vec{E}}_{\geq 0} = \{0\}$. Equivalently, the circulation that is one on every arrow is strictly positive and defines a strictly positive functional on σ_G. Thus σ_G is strongly convex.

4.1. **The flow subdivision and its reduced ring.** Choose $\theta \in \mathbb{Q}^V_0$ such that $\theta(S) := \sum_{i \in S} \theta_i \neq 0$ for every nonempty proper $S \subset V$. Consider

$$P_\theta = \{x \in \mathbb{R}^{\vec{E}}_{\geq 0} : \partial x = \theta\}. \tag{4.3}$$

This polyhedron is nonempty: along paths in the connected bidirected graph, route the positive supplies to the negative supplies. Adding a sufficiently large strictly positive circulation produces a point with all coordinates positive. Its affine dimension is therefore d_G.

Lemma 4.1. *Every vertex of P_θ has a spanning-tree support. Its minimizing normal fan Σ_θ is a simplicial subdivision of σ_G. Its ring of piecewise polynomials, reduced by global linear forms, is*

$$A_G = \frac{\operatorname{Sym}(\mathbb{Q}^V/\mathbb{Q}\mathbf{1})}{\left(\prod_{e \in \delta(S)} (x_{h(e)} - x_{t(e)}) : \varnothing \neq S \neq V\right)}. \tag{4.4}$$

Here x_i denotes the image of the ith coordinate vector, an orientation of each edge is chosen, $t(e)$ and $h(e)$ are its tail and head, and $\delta(S)$ is the undirected cut between S and its complement.

Proof. If a positive support contains an undirected circuit, including a pair of opposite arrows, a signed circuit flow gives a nonzero vector in $\ker \partial$ supported there. Small positive and negative multiples preserve nonnegativity, so the point is not a vertex. A disconnected feasible support has a component S with $\theta(S) = 0$, which the choice of θ excludes. A vertex support is consequently a spanning tree, with $s - 1$ positive coordinates. The other $2R - s + 1$ coordinate constraints are independent on the affine flow space, because a prescribed flow on the complement of a spanning tree extends uniquely to a circulation. Hence every vertex is simple.

A cost in N_G has a finite minimum on P_θ exactly when it is nonnegative on its recession cone C_G. The minimizing normal fan therefore has support σ_G. Simplicity makes that fan simplicial.

For $I \subseteq \vec{E}$, the coordinate faces $x_a = 0$, $a \in I$, meet precisely when the arrows $J = \vec{E} \setminus I$ support a θ-flow. The flow-feasibility criterion is

$$\theta(S) \leq 0 \quad \text{whenever no arrow of } J \text{ leaves } S. \tag{4.5}$$

Necessity follows by summing the incidence equations over S. For sufficiency, Farkas' lemma would otherwise give a vertex potential y nondecreasing along

every arrow of J and with $\sum_i y_i\theta_i > 0$. Subtract its minimum and write y as a nonnegative linear combination of characteristic functions of its upper level sets. Each upper level set has no outgoing arrow of J, contradicting (4.5).

Write $\delta^+(S) = \{i \to j : i \in S,\ j \notin S\}$. The Stanley–Reisner ideal of Σ_θ is

$$I_\theta = \left(\prod_{a \in \delta^+(S)} z_a : \theta(S) > 0 \right). \tag{4.6}$$

A coordinate face that is empty is represented by a ghost variable and the corresponding linear monomial belongs to I_θ. Every nonempty coordinate face contains a vertex. The active coordinate constraints there are independent, so this face is a facet and distinct non-ghost arrows give distinct rays of the normal fan. For each non-ghost arrow a, let z_a be the piecewise linear function taking value one on the ray generator given by the class of e_a and zero on the other ray generators. On each simplicial cone these functions are linear coordinates. Compatibility on common faces gives a basis of monomials whose ray sets lie in a cone; all other monomials vanish. Thus the piecewise polynomial ring is $\mathbb{Q}[z_a : a \in \vec{E}]/I_\theta$. The global linear forms are $N_G^* = \ker \partial$; a vector $\ell = (\ell_a)$ maps to $\sum_a \ell_a z_a$. Dividing by them identifies $z_{i\to j}$ with $x_i - x_j$. In particular, the variables for opposite arrows become negatives. Of S and $V \setminus S$, exactly one has positive θ-sum. Their outgoing-cut products become the same cut product up to sign. Thus (4.6) becomes the ideal in (4.4). □

This cut-product presentation appears in [11, Theorem 8.3] and [27, Theorem 6.3]. The graph action will be defined directly on the quotient; the subdivision Σ_θ is used for the degree bound.

4.2. **The cographic quotient.** Orient one copy of each undirected edge and put

$$K_G = \ker(\partial : \mathbb{Q}^E \to \mathbb{Q}_0^V), \qquad r_G = R - s + 1.$$

The cographic independence complex consists of the sets $J \subseteq E$ for which $G \setminus J$ is connected. Its Stanley–Reisner ring is

$$S_G = \mathbb{Q}[z_e : e \in E]/(z_J : G \setminus J \text{ is disconnected}), \qquad z_J = \prod_{e \in J} z_e. \tag{4.7}$$

Its maximal faces are the complements of spanning trees, and their cardinality is r_G.

Lemma 4.2. *A basis of K_G, viewed as linear forms in S_G, is a linear system of parameters and a regular sequence, and*

$$S_G/(K_G) \cong A_G. \tag{4.8}$$

Consequently A_G is finite-dimensional and

$$\deg_q \operatorname{Hilb}(A_G; q) \le R - s + 1. \tag{4.9}$$

Every automorphism of G acts naturally on A_G as a graded representation. Parallel-edge permutations fixing the vertices act trivially.

Proof. For any spanning tree T, projection gives an isomorphism $K_G \to \mathbb{Q}^{E\setminus T}$. Indeed, arbitrary values on the non-tree edges have a uniquely determined correction on the tree that makes the divergence zero. Thus the restrictions of a basis of K_G to every maximal face of the complex are a basis of all its linear coordinates.

Order the edges, and order the maximal faces $B_1, \dots, B_N$ lexicographically. For $i < j$, let e be the least edge in $B_i \triangle B_j$; then $e \in B_i \setminus B_j$. Write $T_i = E \setminus B_i$ and $T_j = E \setminus B_j$. Deleting e from T_j gives a cut crossed by an edge $f \in T_i \setminus T_j = B_j \setminus B_i$. Thus $B' = B_j \setminus \{f\} \cup \{e\}$ is a maximal face. Since $e < f$, it precedes B_j, and

$$B_i \cap B_j \subseteq B' \cap B_j = B_j \setminus \{f\}.$$

Hence the intersection of the simplex on B_j with the preceding simplices is a union of its facets.

Let Δ_i be the complex on $B_1, \dots, B_i$, and put

$$R_i = \{e \in B_i : B_i \setminus \{e\} \in \Delta_{i-1}\} \quad (i > 1), \qquad R_1 = \varnothing.$$

The faces introduced at step i are exactly the sets J with $R_i \subseteq J \subseteq B_i$. Consequently, for $i > 1$ the restriction of Stanley–Reisner rings has kernel

$$\ker(\mathbb{Q}[\Delta_i] \longrightarrow \mathbb{Q}[\Delta_{i-1}]) = z_{R_i}\mathbb{Q}[z_e : e \in B_i].$$

Let $A = \mathrm{Sym}(K_G)$ act by the linear forms above. Its restriction to each $\mathbb{Q}[z_e : e \in B_i]$ is an isomorphism. The first ring $\mathbb{Q}[\Delta_1]$ is therefore free over A, as is every displayed kernel. Inductively the restriction sequence splits over A, so S_G is finite free over A. The parameters are thus a regular sequence. Summing monomials by their positive supports gives $\mathrm{Hilb}(S_G; q) = \sum_J (q/(1-q))^{|J|}$, and reduction by the r_G parameters gives

$$\mathrm{Hilb}(A_G; q) = \sum_{J:\, G\setminus J \text{ connected}} q^{|J|}(1-q)^{r_G - |J|}, \tag{4.10}$$

and the degree of this polynomial is at most r_G.

The quotient of the linear span of the variables by K_G is the incidence image $\mathbb{Q}_0^V$, identified with $\mathbb{Q}^V/\mathbb{Q}\mathbf{1}$. Under this identification $z_{i\to j}$ maps to $x_i - x_j$. The monomial ideal in (4.7) is generated by products over cuts: deleting J disconnects G if and only if J contains a nonempty cut. This proves (4.8).

An automorphism acts on the oriented edge space by signed permutation, and preserves both the monomial ideal and K_G. It therefore acts on the quotient. In (4.4) this action is the induced permutation of the vertex differences. This proves the last assertion. □

4.3. **Faces and contractions.** An arrow support $K \subseteq \vec{E}$ is *totally cyclic* if each of its arrows belongs to a directed cycle. Equivalently, it is the positive support of a nonnegative circulation. The empty support is allowed. Write $\pi(K)$ for its strongly connected components, including isolated vertices.

Lemma 4.3. *The faces of σ_G are indexed by totally cyclic supports K. The face corresponding to K is generated by the coordinate classes outside K and has a canonical product decomposition*

$$\tau_K \cong \sigma_{G/\pi(K)} \times \mathbb{R}_{\geq 0}^{u_K}, \qquad u_K = 2|E_{\mathrm{int}}| - |K|. \tag{4.11}$$

Here E_{int} consists of the undirected edges internal to the blocks of $\pi(K)$; loops are deleted from the contracted graph. Moreover,

$$M_K = \ker(\partial : \mathbb{Q}^K \to \mathbb{Q}^V), \qquad \dim M_K = |K| - s + |\pi(K)|, \tag{4.12}$$

is dual to $N_G/(N_G)_{\tau_K}$. All these descriptions are equivariant under automorphisms of the directed configuration preserving K.

Proof. Every face of the nonnegative circulation cone is specified by its positive support. A positive circulation is a sum of directed-cycle flows, so those supports are exactly the totally cyclic ones. Duality of faces then identifies the corresponding face of σ_G with the cone generated by coordinates outside K. The linear relations between these coordinates form

$$\operatorname{im} \partial^T \cap \mathbb{Q}^{\vec{E}\setminus K}.$$

The image of a potential under ∂^T lies in this intersection precisely when its differences vanish on every arrow of K, or equivalently when the potential is constant on each block of $\pi(K)$. Such a potential has zero difference on every unused internal arrow. On the cross arrows its differences are exactly the incidence relations of the contracted graph. The internal unused coordinates are consequently independent orthant coordinates, proving (4.11).

The tangent space to the dual face is the circulation space supported on K, namely M_K. Incidence has rank $s - |\pi(K)|$ on that support, which gives (4.12). The description uses only coordinate subspaces, incidence, and the intrinsic component partition; hence it is equivariant. In the contracted directed graph, an automorphism may permute the two opposite parallel-arrow sets independently. □

Proposition 4.4 (Strict local degree bound)**.** *For $s > 1$,*

$$\deg_q \overline{\mathcal{L}}_{\sigma_G} \leq R - s + 1 < \frac{\dim \sigma_G}{2}. \tag{4.13}$$

For every nonzero cone τ that is a face of σ_G or a product of such a face with a smooth orthant, $\deg_q \overline{\mathcal{L}}_\tau < \dim(\tau)/2$.

Proof. Push the polynomial sheaf of Σ_θ forward to the face fan of σ_G. Its stalk at σ_G is the ring of piecewise polynomials on Σ_θ. By Theorem 3.2(ii) and (3.4), its reduction contains $\overline{\mathcal{L}}_{\sigma_G}$ as a graded direct summand. Lemma 4.1 identifies that reduction with A_G. Lemma 4.2 therefore bounds its degree by $R - s + 1$. The strict inequality follows from

$$2(R - s + 1) < 2R - s + 1 \qquad (s > 1).$$

Apply Lemma 4.3 to each face, and use Lemma 3.3 to remove its orthant factor. A contracted graph with one vertex has zero cone and reduced

module $\mathbb{Q}$. A positive-dimensional smooth orthant has the same degree-zero reduction, which also satisfies the strict bound. □

5. Equivariant graph characters

5.1. **Canonical actions and the cellular formula.** Let a finite group act linearly on a fan. It acts on every polynomial stalk by pullback. A compatible action on a sheaf is semilinear with respect to these stalk actions. Such actions on minimal extensions can be constructed recursively: choose a cone in each orbit, make the minimal free cover equivariant under its stabilizer, and transport it around the orbit. The equivariant cover is obtained by averaging the homogeneous lifting maps.

We will use the strict degree condition

$$\deg_q \overline{\mathcal{L}}_\tau < \frac{\dim \tau}{2} \qquad (\tau \neq 0). \tag{5.1}$$

Proposition 4.4 establishes it for the graph cones and their smooth products. Under (5.1), the normalized minimal extension and its action are canonical. Here is the relevant rigidity argument. Suppose two morphisms agree on the proper faces of a cone τ. Their difference maps $\mathcal{L}_\tau$ into $\mathcal{L}(\tau, \partial\tau)$. The source has free generators in degrees less than $\dim \tau/2$. By (3.3), the target has free generators in the complementary degrees, all greater than $\dim \tau/2$. A degree-zero map between them is zero. Starting at the zero cone proves uniqueness of a normalized morphism. Existence follows from the free-cover construction. In particular, the normalized isomorphisms attached to group elements obey the group law. This is the rigidity criterion of [1, Remark 1.8(ii),(iii)] with the degree convention fixed above.

For a finite graded representation $V = \bigoplus_j V_j$ define

$$\operatorname{Tr}(g \mid V; q) = \sum_j \operatorname{Tr}(g \mid V_j) q^j.$$

For a cone fixed by g, put $L_\tau(g; q) = \operatorname{Tr}(g \mid \overline{\mathcal{L}}_\tau; q)$. On a complete or convex-support fan in W write

$$R_\Delta(g; q) = \operatorname{Tr}(g \mid \overline{\mathcal{L}(\Delta)}; q), \qquad R^!_\Delta(g; q) = \operatorname{Tr}(g \mid \overline{\mathcal{L}(\Delta, \partial\Delta)}; q).$$

Proposition 5.1 (Equivariant cellular formula)**.** *Let Δ be a complete or convex-support fan of dimension $d = \dim W$. Let g belong to a finite group of fan automorphisms, with a compatible action on $\mathcal{L}$. Then*

$$R^!_\Delta(g; q) = \sum_{\substack{\tau \in \Delta \\ g\tau = \tau}} \det(q - g \mid W/W_\tau) L_\tau(g; q). \tag{5.2}$$

For R_Δ the corresponding sum is over the cones outside the boundary fan. In particular the two sums coincide for a complete fan.

Proof. Take the equivariant cellular complex in Theorem 3.2(iii) and reduce it by the global parameters using the Koszul complex. Because the resolved module is free over A_W, its higher Koszul homology is zero. For a fixed cone

τ, an equivariant homogeneous lift of its reduced stalk gives $\mathcal{L}_\tau \cong A_\tau \otimes \overline{\mathcal{L}}_\tau$. Put $U = W/W_\tau$ and $c = \dim U$. The missing global parameters form U^*, so the Koszul Euler character of this term is $\det(1 - qg^{-1} \mid U)L_\tau(g;q)$. Its cellular degree supplies $(-1)^c$, and its orientation line supplies $\det(g \mid U)$. Since g has finite order, that determinant is the orientation sign. The identity

$$(-1)^c \det(g \mid U) \det(1 - qg^{-1} \mid U) = \det(q - g \mid U) \tag{5.3}$$

proves the contribution in (5.2). Summands in a nontrivial orbit of cones have zero trace. The absolute cellular resolution omits the boundary cones, giving the second assertion. □

Proposition 5.2 (Equivariant duality)**.** *Let Δ be complete or have full-dimensional convex support in W, and put $d = \dim W$. Let a finite group act on Δ. Assume (5.1) for every nonzero cone and use the normalized action on $\mathcal{L}$. For every group element g,*

$$R^!_\Delta(g;q) = q^d R_\Delta(g^{-1};q^{-1}). \tag{5.4}$$

In particular, for a complete fan R_Δ is the character of an actual finite graded representation concentrated in degrees $0, \dots, d$ and satisfies this reciprocal identity.

Proof. In (5.2), the zero cone contributes $\det(q - g \mid W)$, with leading coefficient one in degree d. Every nonzero cone contributes in degrees at most

$$d - \dim \tau + \deg L_\tau < d - \tfrac{1}{2} \dim \tau < d.$$

Thus the top reduced relative degree is d, its dimension is one, and every g acts there as the identity. There are no higher reduced degrees. Consequently the degree-zero evaluation in (3.3) is unique up to a scalar and is equivariant: it is determined by the image of this single top generator.

Start the intersection product with ordinary multiplication on $\mathcal{L}_0 = \mathbb{Q}$. At a nonzero cone τ, average a relative interior vector over the cone stabilizer. Its span is a fixed line. Let $B \subset A_\tau$ be the polynomial subring on linear forms that vanish on this line. Projection along the line identifies $\partial\tau$ with a complete fan, so its section module is free over B. Form

$$F_\tau = A_\tau \otimes_B \mathcal{L}(\partial\tau), \qquad \beta : F_\tau \longrightarrow \mathcal{L}(\partial\tau).$$

Lift the restriction $r : \mathcal{L}_\tau \to \mathcal{L}(\partial\tau)$ to a map $\alpha : \mathcal{L}_\tau \to F_\tau$ and average that lift. It still satisfies $\beta\alpha = r$. On reduction $\overline{r}$ is an isomorphism, so $\overline{\alpha}$ is injective and α is a split embedding of graded free modules.

Choose a stabilizer-invariant complement to $\operatorname{im}\overline{\alpha}$ in $\overline{F}_\tau$, lying in $\ker\overline{\beta}$, and choose an equivariant homogeneous lift of this generating representation. For a lifted generator f, its image $\beta(f)$ belongs to $A_\tau^{>0}\mathcal{L}(\partial\tau)$. Surjectivity of r gives a lift $x \in A_\tau^{>0}\mathcal{L}_\tau$ with $r(x) = \beta(f)$. Choose these lifts as a homogeneous linear map on the chosen complement and average that map over the stabilizer. The equality $r(x) = \beta(f)$ is preserved, so the correction is equivariant. Replace f by $f - \alpha(x)$. The new lift lies in $\ker\beta$ and has the

same reduction. These lifts form an equivariant free complement K with $F_\tau = \alpha(\mathcal{L}_\tau) \oplus K$ and $K \subseteq \ker \beta$. Extend the already constructed boundary product to F_τ and project it along K to $\mathcal{L}_\tau$. Repeating by cone orbits gives an equivariant product. These are the lift and complement choices in [1, Section 5.1]; the perfect-pairing assertion of Theorem 3.2(iv) applies to them.

The resulting pairing, followed by the equivariant evaluation, is an invariant perfect pairing. Reducing it by global parameters identifies relative degree j with the dual of absolute degree $d - j$. This gives (5.4). All stalk generators and polynomial rings have nonnegative degrees, so a complete fan has no negative reduced degree. The preceding top-degree calculation excludes degrees above d. □

For a finite group acting on a rational representation, the character at g^{-1} equals that at g. One can average a positive-definite inner product to identify the real representation with its dual. We will therefore replace g^{-1} by g in reciprocal character identities below.

5.2. **The cut-algebra character.** First let g be an automorphism of the undirected multigraph G. Choose orientations of its edges. For an edge orbit O, let $\ell_O = |O|$ and let $\eta_O \in \{1, -1\}$ be the product of the orientation reversals along that orbit. Define

$$Z_G(g; q) = \mathrm{Tr}(g \mid A_G; q).$$

For a disconnected graph we use the convention $Z_G = 0$, and for a single vertex with no edges put $Z_G = 1$.

Lemma 5.3. *For a connected graph,*

$$Z_G(g; q) = \frac{1 - q}{\prod_{C \in V/g}(1 - q^{|C|})} \sum_{\substack{H \subseteq E,\ gH = H \\ (V,H) \text{ connected}}} \prod_{O \subseteq H} (1 - \eta_O q^{\ell_O}) \prod_{O \not\subseteq H} \eta_O q^{\ell_O}. \tag{5.5}$$

Proof. A monomial in S_G is fixed by the edge permutation only if its exponents are constant along every edge orbit. An orbit occurring in its positive support contributes

$$\sum_{a \geq 1} \eta_O^a q^{a\ell_O} = \frac{\eta_O q^{\ell_O}}{1 - \eta_O q^{\ell_O}}.$$

The positive supports in S_G are exactly the J with connected complement. The regular Koszul quotient by K_G multiplies this trace series by $\det(1 - qg \mid K_G)$. The incidence exact sequence

$$0 \longrightarrow K_G \longrightarrow \mathbb{Q}^E \xrightarrow{\partial} \mathbb{Q}^V \longrightarrow \mathbb{Q} \longrightarrow 0$$

gives

$$\det(1 - qg \mid K_G) = \frac{(1 - q) \prod_{O \in E/g}(1 - \eta_O q^{\ell_O})}{\prod_{C \in V/g}(1 - q^{|C|})}.$$

Multiply the support contributions by this determinant and set $H = E \setminus J$. The factors on H and its complement are exactly those in (5.5). □

5.3. **Two partition identities.** For an invariant arrow support D in the doubled graph, give it the weight

$$w_g(D;q) = \det(q - g \mid \mathbb{Q}^D) = \prod_{O \in D/g} (q^{|O|} - 1).$$

Let

$$\mathcal{D}_G(g;q) = \sum_{\substack{gD=D \\ D \text{ acyclic}}} w_g(D;q),$$

$$D_V(g;q) = \det(q - g \mid \mathbb{Q}^V) = \prod_{C \in V/g} (q^{|C|} - 1).$$

Here acyclic means that the arrow support has no directed cycle.

Let ρ be a g-invariant partition of V. Write $R_{\mathrm{cross}}(\rho)$ for the number of undirected edges joining different blocks. Order the blocks and orient these cross edges in that order. Let $\epsilon_\rho(g)$ be the product of the signs of the orientation reversals induced by g. Changing the chosen orientations changes each reversal sign at both ends of an edge orbit, so their product is unchanged. For an orbit of blocks choose a representative C and write ℓ_C for its length. The automorphism induced on $G[C]$ is g^{ℓ_C}.

Lemma 5.4. *The following identities hold:*

$$\frac{\mathcal{D}_G(g;q)}{D_V(g;q)} = \sum_{g\rho=\rho} \epsilon_\rho(g) q^{R_{\mathrm{cross}}(\rho)} \prod_{C \in \rho/g} \frac{Z_{G[C]}(g^{\ell_C}; q^{\ell_C})}{q^{\ell_C} - 1}, \tag{5.6}$$

$$\frac{q^{2R}}{D_V(g;q)} = \sum_{g\rho=\rho} \epsilon_\rho(g) q^{R_{\mathrm{cross}}(\rho)} \prod_{C \in \rho/g} \frac{q^{\ell_C d_C} Z_{G[C]}(g^{\ell_C}; q^{-\ell_C})}{q^{\ell_C} - 1}, \tag{5.7}$$

Here $d_C = 2|E(G[C])| - |C| + 1$. Terms containing a disconnected induced block are zero.

Proof. Substitute (5.5) into the right side of (5.6) and multiply by $D_V(g;q)$. Choosing a connected spanning subgraph in every block amounts to choosing an arbitrary g-invariant spanning subgraph $H \subseteq E$, whose component partition is ρ. Let $c_V = |V/g|$ and let c_H be the number of g-orbits of components of H. Its resulting summand is

$$(-1)^{c_V + c_H} \prod_{O \subseteq H} (1 - \eta_O q^{\ell_O}) \prod_{O \nsubseteq H} \eta_O q^{\ell_O}. \tag{5.8}$$

The signs from the internal incidence quotients give $(-1)^{c_V+c_H}$; the remaining cross-edge signs are exactly $\epsilon_\rho(g)$. This accounts for all orientation signs in the substitution.

Collapse each vertex orbit to one vertex. The number of components of the image of H is c_H. An edge orbit that becomes a loop can be included

or excluded without changing c_H, and its two weights sum to $(1-\eta_O q^{\ell_O})+\eta_O q^{\ell_O}=1$. Every remaining edge orbit joins two distinct vertex orbits and can be oriented from one of them to the other, giving $\eta_O=1$. Thus the sum of (5.8) is the weighted acyclic-orientation sum on the quotient multigraph.

To verify the last assertion explicitly, for a graph with vertex set of size v, the number $a(J)$ of acyclic orientations of an edge set J is

$$a(J)=\sum_{H\subseteq J}(-1)^{|H|+v+c(H)}.$$

Both sides satisfy deletion–contraction, are one on an edgeless graph, and vanish in the presence of a loop. For the left side, deletion–contraction follows by adding one graphical hyperplane: the newly divided chambers correspond to the chambers of the contracted arrangement. Multiplying this identity by $\prod_{e\in J}(x_e-1)$ and summing over J gives

$$\sum_J a(J)\prod_{e\in J}(x_e-1)=\sum_H(-1)^{v+c(H)}\prod_{e\in H}(1-x_e)\prod_{e\notin H}x_e.$$

Apply it with $x_e=q^{\ell_e}$. An invariant directed support upstairs is acyclic exactly when its quotient orientation support is acyclic. A quotient directed cycle lifts to a directed path whose endpoint is a translate of its starting vertex; repeating the path after suitable powers of g closes it to a directed cycle. In particular an arrow orbit becoming a quotient loop cannot occur in an invariant acyclic support. This proves (5.6).

For (5.7), reciprocal substitution in (5.5) gives

$$\frac{q^{d_G}Z_G(g;q^{-1})}{q-1}=\frac{q^R}{D_V(g;q)}\sum_{\substack{gH=H\\(V,H)\ \text{connected}}}\prod_{O\subseteq H}(q^{\ell_O}-\eta_O)\prod_{O\not\subseteq H}\eta_O. \tag{5.9}$$

Assemble the connected components as before. For each edge orbit, the two choices sum to $(q^{\ell_O}-\eta_O)+\eta_O=q^{\ell_O}$. Their product is q^R, and the outside normalization contributes the other q^R. The vertex determinants and block-orbit denominators are those of (5.7). □

5.4. **The contraction recursion.** For a g-invariant totally cyclic support K, the exact sequence

$$0\longrightarrow M_K\longrightarrow\mathbb{Q}^K\xrightarrow{\partial}\mathbb{Q}^V\longrightarrow\mathbb{Q}^{\pi(K)}\longrightarrow 0$$

gives

$$\det(q-g\mid M_K)=\frac{\det(q-g\mid\mathbb{Q}^K)\det(q-g\mid\mathbb{Q}^{\pi(K)})}{D_V(g;q)}. \tag{5.10}$$

Proposition 5.5. *For every connected loopless multigraph,*

$$q^{d_G}Z_G(g;q^{-1})=\sum_{\substack{gK=K\\K\ \text{totally cyclic}}}\det(q-g\mid M_K)Z_{G/\pi(K)}(g;q). \tag{5.11}$$

Proof. Denote the right side of (5.11) by $T_G(g;q)$. The weighted sum of all invariant directed supports is q^{2R}, because each arrow orbit contributes $1+(q^\ell-1)=q^\ell$. Every directed support has a unique totally cyclic core K, formed by its arrows internal to its strongly connected components. Its remaining arrows form an acyclic support on the contracted components. Conversely, adjoining any acyclic cross support to such a core gives a support with that core. Hence

$$\begin{aligned}\frac{q^{2R}}{D_V(g;q)} &= \sum_K \frac{\det(q-g \mid \mathbb{Q}^K)}{D_V(g;q)}\mathcal{D}_{G/\pi(K)}(g;q)\\ &= \sum_K \det(q-g \mid M_K)\frac{\mathcal{D}_{G/\pi(K)}(g;q)}{\det(q-g \mid \mathbb{Q}^{\pi(K)})}.\end{aligned}$$

Apply (5.6) to each contracted graph. Its partitions are the partitions ρ of V coarser than $\pi(K)$. For fixed ρ, the choices of K separate inside its blocks. On a block orbit of length ℓ_C, the determinant on a cyclically permuted sum of spaces is

$$\det(q-g \mid \textstyle\bigoplus_{j=0}^{\ell_C-1} g^j M) = \det(q^{\ell_C}-g^{\ell_C} \mid M).$$

Regrouping therefore gives the exact identity

$$\frac{q^{2R}}{D_V(g;q)} = \sum_{g\rho=\rho} \epsilon_\rho(g) q^{R_{\mathrm{cross}}(\rho)} \prod_{C\in\rho/g} \frac{T_{G[C]}(g^{\ell_C};q^{\ell_C})}{q^{\ell_C}-1}. \tag{5.12}$$

For a disconnected induced block, its contraction is disconnected for every core, and both the corresponding T and Z are zero.

Induct on $|V|$, the one-vertex case being immediate. Every proper block in (5.12) has fewer vertices, so its T equals the corresponding reciprocal $q^{d_C}Z(q^{-1})$ by induction. Compare (5.12) with (5.7). All terms except the one-block partition cancel. That partition has denominator $q-1$ and cross factor one. Its equality is exactly (5.11). □

5.5. **Independent permutations of opposite arrows.** The contraction of a source–sink configuration need not retain a chosen matching between its opposite parallel-arrow sets. We therefore allow an automorphism $\widehat{g}$ of the directed doubled configuration inducing a vertex permutation g, with arbitrary independent bijections on its directed parallel-arrow sets. The cut algebra A_G is acted on through the vertex permutation, so $Z_G(g;q)$ is defined independently of a choice of such a lift.

Lemma 5.6. *Identities* (5.6), (5.7), *and* (5.11) *remain valid for* $\widehat{g}$. *In the support sums and determinants the actual arrow action of* $\widehat{g}$ *is used; in every* Z *factor the induced vertex action is used.*

Proof. Fix an orbit of ordered pairs of distinct vertices under g, of length ℓ, and suppose there are r parallel arrows over each pair. The return permutation $\widehat{g}^\ell$ on those r arrows has cycle lengths $d_1,\dots,d_h$. Invariant

supports in this orbit are unions of its arrow cycles. The sum of the weights of all nonempty such supports is

$$\prod_{j=1}^{h}(1+(q^{\ell d_j}-1))-1=q^{\ell r}-1. \tag{5.13}$$

Acyclicity depends only on which ordered vertex pairs have nonempty support. Thus the complete acyclic-support sum is independent of the return permutations. It can be evaluated with any lift that preserves a pairing of opposite arrows. For such a lift (5.6) is already proved. Its right side depends only on the weighted vertex permutation: the cut characters, numbers of cross edges, and cross-orientation sign are all unchanged by permuting parallel labels. This proves (5.6) for $\widehat{g}$.

The unrestricted support sum is still q^{2R}, independently of the arrow cycles, and the right side of (5.7) has the same vertex-only dependence. The core decomposition and the exact sequence (5.10) are equivariant for the actual arrow action. Therefore the proof of (5.12), followed by the same induction, proves (5.11) for $\widehat{g}$. □

Theorem 5.7 (Local graph character)**.** *Let G be a connected loopless multigraph and let $\widehat{g}$ be an automorphism of its directed doubled configuration, inducing the vertex permutation g. Then the canonical normalized minimal-extension sheaf on the face fan of σ_G satisfies*

$$\operatorname{Tr}(\widehat{g}\mid\overline{\mathcal{L}}_{\sigma_G};q)=Z_G(g;q)=\operatorname{Tr}(g\mid A_G;q). \tag{5.14}$$

The same identification applies to a smooth product after removing its orthant factor.

Proof. Write $L_G(\widehat{g};q)$ for the character on the left. By Proposition 4.4, all cones in the face fan satisfy (5.1), so its canonical action and Proposition 5.2 apply. The cellular formula and Lemma 4.3 give

$$q^{d_G}L_G(\widehat{g};q^{-1})=\sum_{\substack{\widehat{g}K=K\\ K\ \text{totally cyclic}}}\det(q-\widehat{g}\mid M_K)L_{G/\pi(K)}(\widehat{g};q). \tag{5.15}$$

The quotient normal representation in the cellular formula is dual to M_K; their characteristic polynomials agree because these are finite rational representations. Smooth coordinates contribute no reduced stalk factor, by Lemma 3.3.

Induct on the dimension of the cone. The zero cone has both characters equal to one. A nonempty totally cyclic support contains a directed cycle, so $\dim M_K>0$ and τ_K is a proper face. Its nontrivial contracted graph factor therefore has smaller dimension. By induction its local character is the corresponding Z character. The empty support contributes L_G itself in (5.15) and contributes Z_G in (5.11). Subtracting these two recursions gives

$$q^{d_G}f(q^{-1})=f(q),\qquad f(q)=L_G(\widehat{g};q)-Z_G(g;q).$$

Both summands of f are concentrated in nonnegative degrees strictly less than $d_G/2$, by Proposition 4.4 and Lemma 4.2. Its reciprocal is concentrated strictly above $d_G/2$. Hence the equality forces $f = 0$. The smooth-product assertion follows again from Lemma 3.3. □

6. Transportation polytopes and primitive characters

6.1. **Symmetric functions and cycle indices.** Let U and V be independent alphabets. Put

$$\mathcal{S}(U) = \mathbb{Q}[p_1(U), p_2(U), \ldots]$$

and define $\mathcal{S}(V)$ similarly. These rings are graded by $\deg p_j(U) = \deg p_j(V) = j$. For a partition $\lambda = (1^{m_1} 2^{m_2} \cdots)$, put

$$p_\lambda(U) = \prod_j p_j(U)^{m_j}, \qquad z_\lambda = \prod_j j^{m_j} m_j!.$$

If f is a class function on $S_a \times S_b$ with values in $\mathbb{Q}(t)$, its Frobenius characteristic is

$$\mathrm{Frob}(f) = \sum_{\lambda \vdash a,\ \mu \vdash b} \frac{f(\lambda, \mu)}{z_\lambda z_\mu} p_\lambda(U) p_\mu(V). \tag{6.1}$$

Equivalently, this is the sum over all $(\sigma, \tau) \in S_a \times S_b$, divided by $a!b!$. The same definition, with one factor for each alphabet, applies to a product of symmetric groups.

For a finite graded rational $S_a \times S_b$-representation $W = \bigoplus_{j \geq 0} W_j$, its Frobenius characteristic has the expansion

$$\mathrm{Frob}(W; q) = \sum_{j, \lambda, \mu} \dim \mathrm{Hom}_{S_a \times S_b}(S^\lambda \boxtimes S^\mu, W_j)\, q^j s_\lambda(U) s_\mu(V). \tag{6.2}$$

Here S^λ is the irreducible representation indexed by λ, and s_λ is its Schur function. In particular all multiplicities in (6.2) are nonnegative integers. We use also the tableau definition $s_\lambda(U) = \sum_T u^{\mathrm{wt}(T)}$, where T runs over semistandard tableaux of shape λ. It shows that Schur-positive functions remain coefficientwise nonnegative on expansion into monomials in the alphabet letters.

On symmetric functions with coefficients in $\mathbb{Q}(t)$ define

$$\psi_\ell(t) = (-1)^{\ell - 1} t^\ell, \qquad \psi_\ell(p_j(U)) = p_{\ell j}(U), \qquad \psi_\ell(p_j(V)) = p_{\ell j}(V). \tag{6.3}$$

The rule for t is the signed convention of [16, Section 6.2]. Any auxiliary variable recording a positive degree satisfies $\psi_\ell(z) = z^\ell$. These operations obey $\psi_\ell \psi_m = \psi_{\ell m}$; for t this follows from

$$(-1)^{m-1} (-1)^{m(\ell-1)} = (-1)^{\ell m - 1}.$$

In a completion by a positive degree, set

$$\mathrm{Exp}_\psi(f) = \exp\left(\sum_{\ell \geq 1} \frac{\psi_\ell(f)}{\ell} \right), \qquad \mathrm{Log}_\psi(A) = \sum_{\ell \geq 1} \frac{\mu(\ell)}{\ell} \psi_\ell(\log A). \tag{6.4}$$

Here μ is the arithmetic Möbius function. The constant term of f is zero and that of A is one. Möbius inversion, using $\sum_{d|r}\mu(d)=\delta_{r1}$, proves that these two operations are inverse. In this section they are applied in a commutative algebra attached to a single ray.

6.2. **Thin supports and their factorization.** Let I and J be finite sets of sizes a and b. After choosing standard labels, we also use a,b as subscripts in place of I,J. Fix an integer $k\ge 1$ and put

$$\mathcal{E}=I\times J\times\{1,\dots,k\}.$$

An element (i,j,c) is an arrow $i\to j$. Write $n=a+b$ and $E=kab$. The group $S_I\times S_J$ acts on the vertices and arrows, preserving the last arrow label. For an arrow support $H\subseteq\mathcal{E}$, a vertex subset $S\subseteq I\sqcup J$ is *successor-closed* if no arrow of H leaves S. Put

$$A(S)=|S\cap I|,\quad B(S)=|S\cap J|,\quad \mu(S)=\frac{A(S)}{A(S)+B(S)}$$

for $S\neq\varnothing$. When $I\sqcup J\neq\varnothing$, call H *semistable* if $\mu(S)\le\mu(I\sqcup J)$ for every nonempty successor-closed S. The unique support on the empty vertex set is declared semistable.

When $a,b>0$, let $\partial:\mathbb{Q}^{\mathcal{E}}\to\mathbb{Q}^{I\sqcup J}$ send an arrow $i\to j$ to e_i-e_j, and set

$$\theta_i=b\quad(i\in I),\qquad \theta_j=-a\quad(j\in J),\qquad P_{I,J;k}=\{x\ge 0:\partial x=\theta\}. \tag{6.5}$$

The flow feasibility criterion used for graph cones shows that H is semistable precisely when there is a point of $P_{I,J;k}$ supported on H: for a successor-closed S the condition is

$$\theta(S)=bA(S)-aB(S)\le 0.$$

The point with every arrow coordinate $1/k$ is strictly positive. The equations at the sources bound every coordinate, so $P_{I,J;k}$ is a bounded polytope of dimension

$$D=E-n+1. \tag{6.6}$$

The strictly positive point identifies its affine span equivariantly with $W=\ker\partial$. Its minimizing normal fan $\Delta_{I,J;k}$ is a complete fan in W^*.

For $g=(\sigma,\tau)\in S_I\times S_J$, write $\chi_{I,J}(g)=\operatorname{sgn}(\sigma)\operatorname{sgn}(\tau)$ and define

$$\mathcal{F}_{I,J;k}(g;t)=\chi_{I,J}(g)\,t^{n-E}\,\frac{\displaystyle\sum_{gH=H\ \text{semistable}}\det(t^2-g\mid\mathbb{Q}^H)}{\det(t^2-g\mid\mathbb{Q}^{I\sqcup J})}. \tag{6.7}$$

For $a=0$ or $b=0$ there are no arrows. We use the same formula with its single empty support, and put $\mathcal{F}_{\varnothing,\varnothing;k}=1$. The transportation fan and the dimension (6.6) will only be used when both sets are nonempty.

Lemma 6.1 (The ordered support filtration)**.** *Every support H has a unique ordered partition*

$$I\sqcup J=C_1\sqcup\cdots\sqcup C_s$$

such that $C_1 \sqcup \cdots \sqcup C_r$ is successor-closed for each r, the induced support on each C_r is semistable, and

$$\mu(C_1) > \mu(C_2) > \cdots > \mu(C_s).$$

Every automorphism of H preserving I and J fixes each C_r individually.

Proof. If $I \sqcup J = \varnothing$, take the empty partition. Otherwise, choose a nonempty successor-closed subset of maximal slope. The union of two such subsets is again of maximal slope. Indeed, their intersection is successor-closed, and additivity of A and $A + B$ in the inclusion–exclusion identity forces the slope of their union to be at least the maximal slope. It cannot be larger by maximality; if the intersection is nonempty, its slope must equal the same value as well. There is therefore a unique largest subset C_1 of maximal slope. It is semistable by its definition.

Apply the construction to the induced support on the complement. If its largest slope were at least $\mu(C_1)$, the inverse image of a subset of that slope, together with C_1, would be successor-closed in H. Its slope would either exceed the maximal slope or equal it on a larger subset. Both possibilities are excluded. Repetition gives the required partition.

For uniqueness, suppose a partition with the stated properties is given, and let S be a nonempty successor-closed subset. Each nonempty intersection $S \cap C_r$ is successor-closed in the induced support on C_r. Semistability therefore bounds its slope by $\mu(C_r)$. The slope of S is the weighted average of these slopes, and hence is at most $\mu(C_1)$. Equality implies $S \subseteq C_1$, because the remaining block slopes are strictly smaller. Thus C_1 is exactly the largest maximal-slope subset constructed above. Induction proves uniqueness of all blocks. An automorphism preserves this construction. □

Let Ω be an integral alternating form on a lattice, and let α, β be lattice vectors with $\Omega(\alpha, \beta) = k > 0$, and use the multiplication

$$X^u X^v = t^{\Omega(u,v)} X^{u+v}.$$

All alphabet letters are central. Define the axis series

$$\mathcal{E}(U; X^\alpha) = \prod_{u \in U} \prod_{j \geq 0} (1 + u\, t^{-1-2j} X^\alpha). \tag{6.8}$$

It may equivalently be defined by

$$\log \mathcal{E}(U; X^\alpha) = \sum_{\ell \geq 1} \frac{(-1)^{\ell-1} p_\ell(U)}{\ell(t^\ell - t^{-\ell})} X^{\ell\alpha}.$$

On an axis, formula (6.7) gives a factor $(-1)^{\ell-1}/(t^\ell - t^{-\ell})$ for a vertex cycle of length ℓ. The Frobenius characteristic therefore recovers exactly (6.8).

For coprime $p, r > 0$, put

$$\mathcal{A}_{p,r}(U, V, z; t) = 1 + \sum_{m \geq 1} z^m \operatorname{Frob}(\mathcal{F}_{mp,mr;k}(\,\cdot\,;t)). \tag{6.9}$$

Proposition 6.2 (Support factorization)**.** *In the completion by the sum of the α and β coordinates,*

$$\mathcal{E}(U;X^{\alpha})\mathcal{E}(V;X^{\beta}) = \mathcal{E}(V;X^{\beta})\left(\prod_{p/r \text{ increasing}} \mathcal{A}_{p,r}(U,V,X^{p\alpha+r\beta};t)\right)\mathcal{E}(U;X^{\alpha}). \quad (6.10)$$

Only coprime $p, r > 0$ occur in the product.

Proof. Set $q = t^2$. For a g-invariant support, an arrow orbit of length ℓ contributes $q^{\ell} - 1$ to $\det(q - g \mid \mathbb{Q}^H)$. Summing over every support, without a semistability condition, gives

$$\sum_{gH=H} \det(q - g \mid \mathbb{Q}^H) = \prod_{O\in\mathcal{E}/g} (1 + (q^{|O|} - 1)) = q^E. \quad (6.11)$$

Now group these supports by the partition in Lemma 6.1. Its blocks are individually g-stable. An arrow from an earlier block to a later block must be absent. Arrows in the other direction are unrestricted, and their invariant supports sum to q to the power of the number of these arrows. Inside each block the support is semistable. Conversely these conditions produce exactly the supports with the prescribed filtration, by uniqueness in Lemma 6.1.

List the blocks in the reverse order. For two blocks with charges (a_1, b_1) and (a_2, b_2), written left to right, there are ka_1b_2 allowed cross arrows. Their contribution, including the part of the normalization t^{n-E} involving these blocks, is

$$q^{ka_1b_2}t^{-k(a_1b_2+a_2b_1)} = t^{k(a_1b_2-a_2b_1)}. \quad (6.12)$$

This is $t^{\Omega(a_1\alpha+b_1\beta,\, a_2\alpha+b_2\beta)}$. The sign character and the vertex determinant split over the individually stable blocks. Thus the normalized sum over a fixed filtration is precisely the corresponding based-torus product of its semistable block sums.

Frobenius characteristics turn induction over disjoint label subsets into multiplication. In detail, summing over all choices of blocks of given source and sink sizes introduces the binomial coefficients counting those choices; division by $a!b!$ in (6.1) leaves the product of the block normalizations. The block slopes in reverse filtration order are increasing. Hence the support sum becomes the right side of (6.10). On the left, the two axis normalizations multiply to $\chi_{I,J}(g)t^n/\det(q-g \mid \mathbb{Q}^{I\sqcup J})$. The quantum multiplication contributes t^{kab}. By (6.11), this is the normalized sum over all supports. Equality of each homogeneous Frobenius characteristic proves the proposition. □

6.3. **The positive-flow core.** Assume $a, b > 0$ throughout this subsection. A support K will be called a *positive-flow support* if a point of $P_{I,J;k}$ has support exactly K. Its connected components refer to the underlying undirected graph. Such a component C is balanced:

$$b\,|C \cap I| = a\,|C \cap J|. \quad (6.13)$$

The components partition all vertices, because every source and every sink has a nonzero prescribed supply. If C is balanced, put $a_C = |C \cap I|$ and $b_C = |C \cap J|$. Both are positive, and $a/a_C = b/b_C =: s_C > 0$. Thus the restriction of the global supply to C is s_C times the normalized supply $(b_C, -a_C)$ for its own transportation polytope. Multiplication of a flow by s_C identifies their positive supports and their normal fans.

Lemma 6.3 (Core decomposition)**.** *A semistable support H has a unique maximal positive-flow support $K \subseteq H$. Every arrow of H whose endpoints lie in one component of K belongs to K. The remaining arrows form a directed acyclic support between the components of K.*

Conversely, the union of positive-flow supports on balanced blocks with any directed acyclic support between those blocks is semistable, and its maximal positive-flow support is the given union.

Proof. Let K be the union of the supports of all feasible flows supported on H. Since H is finite, averaging finitely many witnessing flows gives one feasible flow positive on every arrow of K. Thus K itself is a positive-flow support and is uniquely maximal.

Suppose an arrow $i \to j$ of H has endpoints in one component of K. Choose an undirected path in K from j to i. Add a small positive amount on the arrow $i \to j$ and add signed amounts on the path to cancel the resulting divergence. All path coordinates were strictly positive, so a sufficiently small amount leaves every coordinate nonnegative. This produces a feasible flow positive on the additional arrow, which must belong to K.

If the remaining arrows contained a directed cycle after the components of K were contracted, place the same small positive flow on its cross arrows. Correct the divergence inside each component along undirected positive-flow paths. Again positivity is preserved for a sufficiently small amount. This would enlarge K, a contradiction.

For the converse, extend the prescribed positive flows by zero on the cross arrows. The result is feasible. Any other feasible flow induces a nonnegative circulation on the contracted cross support, since every block is balanced. A directed acyclic graph has no nonzero nonnegative circulation: a source has no incoming flow and hence no outgoing flow, and induction deletes all vertices. Therefore every cross coordinate remains zero. □

Let

$$\mathcal{R}_{I,J;k} = \overline{\mathcal{L}(\Delta_{I,J;k})}, \qquad R_{I,J;k}(g;q) = \mathrm{Tr}(g \mid \mathcal{R}_{I,J;k}; q).$$

The modules $\mathcal{R}_{I,J;k}$ are finite graded representations constructed by the preceding fan theory. Relabeling gives an $S_a \times S_b$-module $\mathcal{R}_{a,b;k}$. If C is a balanced vertex block, abbreviate the corresponding objects to $\mathcal{R}_C$ and R_C, and put $\chi_C = \chi_{C\cap I, C\cap J}$.

For a positive-flow support K with $gK = K$, let $\pi(K)$ be its component partition and set

$$M_K = \ker(\partial : \mathbb{Q}^K \longrightarrow \mathbb{Q}^{I \sqcup J}).$$

The incidence exact sequence gives

$$\frac{\det(q-g \mid \mathbb{Q}^K)}{\det(q-g \mid \mathbb{Q}^{I\sqcup J})} = \frac{\det(q-g \mid M_K)}{\det(q-g \mid \mathbb{Q}^{\pi(K)})}. \tag{6.14}$$

The face of $P_{I,J;k}$ whose relative interior has support K has tangent space M_K. Its normal cone is generated by the coordinate classes outside K in $\mathbb{Q}^{\mathcal{E}}/\operatorname{im}\partial^T$. The relations between those generators are exactly $\operatorname{im}\partial^T \cap \mathbb{Q}^{\mathcal{E}\setminus K}$. The image of a vertex potential under ∂^T lies in this intersection precisely when the potential is constant on every connected component of K. Its differences therefore vanish on every unused internal arrow, whose coordinate gives an independent smooth orthant factor. On the cross arrows its relations are the incidence relations of the contracted graph. For components C, D of K, the numbers of arrows in the two cross directions are equal:

$$k|C\cap I|\,|D\cap J| = k|D\cap I|\,|C\cap J|.$$

Let $G_{\pi(K)}$ be this symmetric contracted graph, with the indicated undirected multiplicities. Every normal cone therefore satisfies (5.1); the normalized action and Proposition 5.2 apply to $\Delta_{I,J;k}$. The cellular formula of Proposition 5.1 and the local character identification of Theorem 5.7 therefore apply. The quotient normal space is dual to the face tangent space M_K. Their characteristic polynomials agree for a finite rational group action, so they give

$$R_{I,J;k}(g;q) = \sum_{\substack{gK=K\\ K \text{ a positive-flow support}}} \det(q-g \mid M_K)\, Z_{G_{\pi(K)}}(g;q). \tag{6.15}$$

The determinant retains the actual arrow action. In the Z factor only the induced vertex action is needed, by the independent-arrow part of Theorem 5.7.

6.4. **The balanced-partition identity.** Call a partition ρ of $I\sqcup J$ *balanced* if every block satisfies (6.13). For a balanced partition set

$$R_{\mathrm{cr}}(\rho) = \sum_{\{C,D\}\subset\rho} k|C\cap I|\,|D\cap J|.$$

For each unordered pair the displayed number is independent of the order of C, D. It is half the number of original arrows between different blocks. If $g\rho=\rho$, let $\epsilon_\rho(g)$ be the cross-orientation character of the symmetric graph on the blocks: order the blocks, orient every cross edge in that order, and take the product of the signs of the orientation reversals under g. Changing the order conjugates the signed permutation matrix and leaves this character unchanged. In a block orbit $\mathcal{O}\in\rho/g$, choose a representative C and write $\ell_C=|\mathcal{O}|$.

Proposition 6.4 (Balanced partitions)**.** *For $a, b > 0$ one has the following identity of rational class functions:*

$$\begin{aligned}&\frac{\displaystyle\sum_{\substack{gH=H\\ H\ \text{semistable}}} \det(q-g \mid \mathbb{Q}^H)}{\det(q-g \mid \mathbb{Q}^{I\sqcup J})}\\ &\quad= \sum_{\substack{g\rho=\rho\\ \rho\ \text{balanced}}} \epsilon_\rho(g) q^{R_{\mathrm{cr}}(\rho)} \prod_{\mathcal{O}\in\rho/g} \frac{R_C(g^{\ell_C}; q^{\ell_C})}{q^{\ell_C}-1}.\end{aligned} \tag{6.16}$$

Proof. Apply Lemma 6.3. For a fixed invariant core K, summing over the extra acyclic cross support and using (6.14) gives

$$\sum_{\substack{gK=K\\ K\ \text{positive-flow}}} \det(q-g \mid M_K) \frac{\mathcal{D}_{G_{\pi(K)}}(g;q)}{\det(q-g \mid \mathbb{Q}^{\pi(K)})}. \tag{6.17}$$

Here $\mathcal{D}$ is the weighted invariant acyclic-support sum of the preceding graph section. It is evaluated on the actual permutations of the directed arrows. Its partition formula, (5.6), applies to these permutations.

A partition of the vertices of $G_{\pi(K)}$ groups components of K into larger blocks. Their unions form a balanced partition ρ of the original vertices. The cross multiplicity between two such unions is

$$\sum_{\substack{C\subset C'\\ D\subset D'}} k|C\cap I|\,|D\cap J| = k|C'\cap I|\,|D'\cap J|.$$

Consequently the cross power $q^{R_{\mathrm{cr}}(\rho)}$ and the orientation character $\epsilon_\rho(g)$ depend only on ρ, not on the component refinement given by K.

Fix ρ and a block orbit of length ℓ with representative C. The restrictions of K to the blocks of this orbit are determined by its restriction K_C to C, which is fixed by g^ℓ. For a vector space U and a cyclic permutation of ℓ copies whose return map is h, direct block elimination gives

$$\det(q-g \mid U\oplus gU\oplus\cdots\oplus g^{\ell-1}U) = \det(q^\ell - h \mid U). \tag{6.18}$$

Apply this to the tangent spaces of the components of K_C. The factor attached to this block orbit in (6.17), after inserting the partition formula, is

$$\frac{1}{q^\ell-1} \sum_{\substack{g^\ell K_C=K_C\\ K_C\ \text{positive-flow}}} \det(q^\ell - g^\ell \mid M_{K_C})\, Z_{G_{\pi(K_C)}}(g^\ell; q^\ell).$$

By (6.15), this is $R_C(g^\ell; q^\ell)/(q^\ell-1)$. Multiplication over all block orbits, followed by summation over ρ, proves (6.16). □

Theorem 6.5 (The full primitive character)**.** *For $a, b > 0$ set $D = kab - a - b + 1$ and define*

$$P_{a,b;k}(U,V;t) = t^{-D}\,\mathrm{Frob}\Big((\mathrm{sgn}_{S_a}\boxtimes\mathrm{sgn}_{S_b})\otimes\mathcal{R}_{a,b;k}; t^2\Big). \tag{6.19}$$

For every coprime $p, r > 0$,

$$(t - t^{-1}) \operatorname{Log}_{\psi} \mathcal{A}_{p,r}(U, V, z; t) = \sum_{m \geq 1} z^m P_{mp,mr;k}(U, V; t). \tag{6.20}$$

Each $P_{a,b;k}$ is a finite Schur-positive Laurent polynomial. Its t-exponents lie in $[-D, D]$ and are congruent to D modulo 2.

Proof. Multiply (6.16) by $\chi_{I,J}(g)t^{n-E}$ and set $q = t^2$. For a block C, put

$$n_C = |C|, \quad E_C = k|C \cap I|\,|C \cap J|, \quad D_C = E_C - n_C + 1.$$

The equality

$$n = \sum_{\mathcal{O}} \ell_C n_C, \qquad E = \sum_{\mathcal{O}} \ell_C E_C + 2R_{\mathrm{cr}}(\rho)$$

cancels all cross powers:

$$t^{n-E} q^{R_{\mathrm{cr}}(\rho)} = \prod_{\mathcal{O}} t^{\ell_C(n_C - E_C)}. \tag{6.21}$$

The vertex permutation signs satisfy

$$\chi_{I,J}(g) = \prod_{\mathcal{O}} (-1)^{(\ell_C - 1)n_C} \chi_C(g^{\ell_C}). \tag{6.22}$$

Indeed, moving ℓ_C copies of a set of n_C vertices cyclically has determinant $(-1)^{(\ell_C-1)n_C}$, after which the return permutation acts on the representative block.

Order the blocks consecutively by their g-orbits. No orientation reversal occurs between two different block orbits. Within an orbit of length ℓ_C, each pair of blocks has cross multiplicity E_C. The sign of the block cycle is $(-1)^{\ell_C - 1}$. It follows that

$$\epsilon_\rho(g) = \prod_{\mathcal{O}} (-1)^{(\ell_C - 1)E_C}. \tag{6.23}$$

Combining (6.21)–(6.23), the factor for one block orbit is

$$(-1)^{(\ell_C-1)(n_C-E_C)} t^{\ell_C(n_C-E_C)} \chi_C(g^{\ell_C}) \frac{R_C(g^{\ell_C}; t^{2\ell_C})}{t^{2\ell_C} - 1}. \tag{6.24}$$

Put $h = g^{\ell_C}|_C$. The displayed factor is the ℓ_Cth substitution (6.3) of the single-block weight

$$\chi_C(h)\, \frac{t^{n_C - E_C} R_C(h; t^2)}{t^2 - 1} = \frac{\chi_C(h) t^{-D_C} R_C(h; t^2)}{t - t^{-1}}.$$

To take Frobenius characteristics, consider a cycle of ℓ blocks, each with a source and b sink labels. There are $(\ell a)!(\ell b)!/(a!b!)^\ell$ ordered block decompositions. For a fixed return permutation h on the first block, the first $\ell - 1$ block bijections have $(a!b!)^{\ell-1}$ choices; the last is determined by h. Dividing by the ℓ choices of cyclic starting block and by the Frobenius factorials gives

$$\frac{(\ell a)!(\ell b)!}{(a!b!)^\ell} \frac{(a!b!)^{\ell-1}}{\ell(\ell a)!(\ell b)!} = \frac{1}{\ell\, a!b!}.$$

A j-cycle of h becomes an ℓj-cycle of the global permutation, giving $p_j \mapsto p_{\ell j}$ in each alphabet. Unordered collections of m_ℓ block cycles contribute $1/m_\ell!$. Fix coprime $p, r > 0$ and let

$$f = \frac{\sum_{m\geq 1} z^m P_{mp,mr;k}}{t - t^{-1}}$$

be the Frobenius series of the single-block weights on this ray. The total contribution of its block cycles is

$$\prod_{\ell\geq 1} \sum_{m_\ell \geq 0} \frac{1}{m_\ell!} \left(\frac{\psi_\ell(f)}{\ell}\right)^{m_\ell} = \exp\left(\sum_{\ell\geq 1} \frac{\psi_\ell(f)}{\ell}\right).$$

The weight of each cycle is (6.24), so (6.16) reads

$$\mathcal{A}_{p,r} = \mathrm{Exp}_\psi(f).$$

Applying (6.4) proves (6.20).

The representation $\mathcal{R}_{a,b;k}$ is an actual finite graded representation. Tensoring with either sign representation leaves it an actual representation, so (6.2) proves Schur positivity. Its grading is nonnegative. Complete-fan duality (Proposition 5.2) identifies degrees j and $D - j$ dually, so no degree above D occurs. The shift t^{-D} and substitution $q = t^2$ give precisely the exponent bounds and parity stated in the theorem. □

7. Positive quantum strings

7.1. **Elementary and complete strings.** Let L be a lattice with integral alternating form Ω, and let $C \subset L \otimes_{\mathbb{Z}} \mathbb{R}$ be a pointed rational cone. Choose an integral linear charge κ strictly positive on $(C \cap L) \setminus \{0\}$. Work over $\mathbb{K} = \mathbb{Q}((t^{-1}))$ in the κ-completion of the algebra $\bigoplus_{a\in C\cap L} \mathbb{K}X^a$, with multiplication

$$X^u X^v = t^{\Omega(u,v)} X^{u+v}.$$

For a root $p \in (C \cap L) \setminus \{0\}$, an integer h, and a central variable u, define

$$\mathsf{E}_{p,h}(u) = \prod_{j\geq 0} (1 + u\, t^{h-2j} X^p), \qquad \mathsf{H}_{p,h}(u) = \prod_{j\geq 0} (1 - u\, t^{h-2j} X^p)^{-1}. \tag{7.1}$$

Their coefficients converge as Laurent series in t^{-1}. The argument $u = 1$ is omitted. The ordinary logarithms, in the commutative subalgebra of the ray $\mathbb{R}_{\geq 0} p$, are

$$\begin{aligned} \log \mathsf{E}_{p,h}(u) &= \sum_{\ell\geq 1} \frac{(-1)^{\ell-1} u^\ell t^{\ell(h+1)}}{\ell(t^\ell - t^{-\ell})} X^{\ell p}, \\ \log \mathsf{H}_{p,h}(u) &= \sum_{\ell\geq 1} \frac{u^\ell t^{\ell(h+1)}}{\ell(t^\ell - t^{-\ell})} X^{\ell p}. \end{aligned} \tag{7.2}$$

We call a factor in (7.1) an *atom*. Its type is $\varepsilon = 1$ for E and $\varepsilon = 0$ for H.

Let ω be the usual involution on symmetric functions:

$$\omega(p_j) = (-1)^{j-1} p_j, \qquad \omega(s_\lambda) = s_{\lambda'}.$$

Applying it in an input alphabet changes the corresponding E axis series into an H axis series, by (7.2).

Lemma 7.1. *Suppose f is homogeneous of total degree K in a collection of alphabets on which ω acts. Then*

$$\omega\psi_\ell(f) = (-1)^{(\ell-1)K}\psi_\ell(\omega f). \tag{7.3}$$

Proof. For a product of power sums of total degree K, with r power-sum factors, the sign on the left is $(-1)^{\ell K-r}$. The sign contributed by ω on the right is $(-1)^{K-r}$. Their ratio is $(-1)^{(\ell-1)K}$. Power-sum monomials form a basis, proving the assertion. □

Proposition 7.2 (Two-atom factorization)**.** *Let $\Omega(\alpha,\beta) = k > 0$. For arbitrary integers h_α, h_β and arbitrary choices of E or H type, the product of the corresponding atoms on α,β factors, in the opposite angular order, into the same two boundary atoms and positive integral multiplicities of atoms on roots $a\alpha + b\beta$, $a, b > 0$. At each bounded root degree only finitely many output atoms occur.*

An atom obtained using a copies of the first input and b copies of the second has parameter satisfying

$$h_{\text{out}} \le ah_\alpha + bh_\beta + kab. \tag{7.4}$$

Its type is

$$\varepsilon_{\text{out}} \equiv kab + a\varepsilon_\alpha + b\varepsilon_\beta \pmod 2. \tag{7.5}$$

Proof. First give the two inputs independent alphabets and use the normalized parameter -1. Proposition 6.2 identifies the ray factors, and Theorem 6.5 identifies their full plethystic logarithms with Schur-positive Laurent polynomials. Apply ω to each alphabet of H type. It is an automorphism of the coefficient ring, so it preserves the ordered factorization identity.

Expand the resulting primitive polynomial into Schur functions, and then into alphabet monomials. Consider one term t^rM of multiplicity one, where M includes the torus monomial on the output ray. Let K be its total degree in the H-marked alphabets. In the ordinary logarithm of the ray factor, its contribution is

$$\sum_{\ell\ge1} \frac{(-1)^{(\ell-1)(r-1+K)}t^{\ell r}M^\ell}{\ell(t^\ell - t^{-\ell})}. \tag{7.6}$$

Indeed, the substitution of $t^r/(t-t^{-1})$ contributes $(-1)^{(\ell-1)(r-1)}$; moving ω through ψ_ℓ contributes the additional sign in Lemma 7.1. By (7.2), expression (7.6) is the logarithm of an E atom if $r-1+K$ is odd, and of an H atom if it is even. Its parameter is $r-1$.

All multiplicities are nonnegative: ω conjugates Schur partitions, and the monomial coefficients of a Schur function are nonnegative. Now specialize each input alphabet to the single central letter $u_\alpha t^{h_\alpha+1}$ or $u_\beta t^{h_\beta+1}$, as appropriate. The central colors u_α, u_β may then be specialized to arbitrary central color monomials, including monomials sharing variables; setting them

to one gives the uncolored statement. These are ordinary specializations of the established identity. For a term of bidegree (a, b) it changes the output parameter to

$$h_{\mathrm{out}} = r - 1 + a(h_\alpha + 1) + b(h_\beta + 1). \tag{7.7}$$

The factor type already determined by (7.6) is unchanged by this specialization.

The dimension of the relevant transportation polytope is $D = kab - a - b + 1$. Theorem 6.5 gives $-D \le r \le D$ and $r \equiv D \pmod 2$. Inserting the upper bound in (7.7) proves (7.4). Moreover

$$K = a(1 - \varepsilon_\alpha) + b(1 - \varepsilon_\beta).$$

Thus $r - 1 + K \equiv kab + a\varepsilon_\alpha + b\varepsilon_\beta$, which proves (7.5). Each fixed bidegree uses a finite representation and a finite symmetric function. There are finitely many bidegrees in a bounded root degree, so the output there is finite. □

When $\Omega(\alpha, \beta) = 0$, all factors on the two rays commute and they may be interchanged without producing interior atoms. For positive pairing, $\Gamma = \mathbb{Z}\alpha + \mathbb{Z}\beta$ is a rank-two sublattice of L, and the identity holds in its based torus. A coprime pair (p, r) gives a primitive ray generator in Γ; the vector $p\alpha + r\beta$ may be nonprimitive in L. All pairings use the actual form Ω. Commuting factors on the same ambient ray are grouped together.

7.2. **Finite reordering.** Suppose C has dimension two and the restriction of Ω is nonzero. Order its rays by $\prec$ so that $\Omega(a, b) < 0$ whenever the ray of a precedes the distinct ray of b. An adjacent pair in the reverse order therefore has positive pairing.

Lemma 7.3 (Finite reordering)**.** *For every positive integer B, every finite word of atoms with roots in C can be rewritten, modulo roots of charge greater than B, as a finite positive word whose rays are in $\prec$ order.*

Proof. Discard roots of charge greater than B. For an adjacent inverted pair of roots α, β, put $s = \kappa(\alpha) + \kappa(\beta)$. If $s \le B$, use Proposition 7.2 and discard its terms above the truncation. The two boundary atoms are reversed; each new interior atom has root $a\alpha + b\beta$ with $a, b \ge 1$, and hence has charge at least s. If $s > B$, interchange the atoms directly: their commutator has charge at least s, so this is the same operation in the truncated torus.

For $2 \le j \le B$, let I_j count inverted pairs of atom occurrences whose charges sum to j. Let I_{all} count all inverted pairs in the current finite word. A replacement with $s \le B$ leaves every I_j for $j < s$ unchanged and lowers I_s by exactly one. All other old atoms lie before or after both replaced positions, so their relative orders with the two old boundary atoms do not change. A pair involving a new interior atom has charge sum greater than s. A swap with $s > B$ leaves all I_j, $j \le B$, unchanged and lowers I_{all} by one. Consequently the finite tuple

$$(I_2, I_3, \ldots, I_B, I_{\mathrm{all}})$$

strictly decreases lexicographically at every operation. Lexicographic order on a finite product of copies of $\mathbb{N}$ is well-founded, so only finitely many operations occur. The resulting word has no adjacent inversion and is ordered. Every operation is an identity in the truncated torus and preserves nonnegative multiplicities. □

Lemma 7.4 (Uniqueness of ray factors)**.** *An element with constant term one in the completed positive torus has at most one factorization into factors with constant term one, supported on individual rays, in a fixed angular order.*

Proof. Induct on the positive charge. At the coefficient of X^p, a coefficient of that charge in a ray factor can occur only in the factor for $\mathbb{R}_{\geq 0}p$, and it occurs with coefficient one. Every contribution involving two or more nonconstant terms uses only smaller positive charges. After those lower coefficients have been determined, subtraction therefore determines the coefficient of X^p uniquely. □

8. Admissible lists and positive reordering

8.1. **Energy bounds.** A list of atom occurrences with roots $p_i \in (C \cap L)\setminus\{0\}$ and parameters $h_i \in \mathbb{Z}$ is called *admissible* if, for every B, the parameters with $\kappa(p_i) \leq B$ are bounded above, and for each integer J only finitely many such occurrences have $h_i \geq J$. An occurrence is counted with its multiplicity. A *positive packet* on a ray is the product of the atoms of an admissible list on that ray. Infinite products are evaluated at bounded root degree, using the t^{-1}-adic topology on each coefficient.

Lemma 8.1. *Fix a charge bound B and let M_B bound $|\Omega(p,q)|$ for positive roots p, q of charge at most B. An output atom of charge at most B obtained by finite reordering has an input-leaf history with at most B leaves and satisfies*

$$h_{\mathrm{out}} \leq \sum_{\text{leaves } i} h_i + M_B B^2. \tag{8.1}$$

If

$$K_0 = \max(0, \sup_{\kappa(p_i)\leq B} h_i),$$

then every active input occurrence in an output with $h_{\mathrm{out}} \geq J$ satisfies

$$h_i \geq J - (B-1)K_0 - M_B B^2. \tag{8.2}$$

All these output parameters are bounded above by $BK_0 + M_B B^2$.

Proof. An input occurrence is a one-leaf history. In a use of Proposition 7.2, an output with root $a\alpha + b\beta$ has a copies of the history of the first atom and b copies of that of the second. All leaf roots add to the output root. Their positive integral charges show that their number is at most B.

The extra term in (7.4) is $ab\Omega(\alpha,\beta)$. Expand it by bilinearity into pairings of leaves lying in the two different input groups. Along the resulting history tree, a pair of leaf occurrences is counted at most once. Replacing each

pairing by its absolute value bounds the total extra contribution by $M_B B^2$. This proves (8.1). There are at most $B-1$ other leaves, each with parameter at most the nonnegative number K_0. Isolating any chosen h_i in (8.1) gives (8.2). The upper bound follows by bounding all leaves by K_0. □

8.2. **Auxiliary degrees and compatible atom lists.** Index the input occurrences by a set I, write ε_i for the type of the ith occurrence, and attach an independent central variable u_i to it. For a finitely supported $\mathbf{d} = (d_i) \in \mathbb{N}^{(I)}$, put

$$p(\mathbf{d}) = \sum_i d_i p_i, \qquad \varepsilon(\mathbf{d}) = \sum_i \varepsilon_i d_i + \sum_{i<j} \Omega(p_i, p_j) d_i d_j \quad \text{in } \mathbb{Z}/2\mathbb{Z}. \tag{8.3}$$

The order used in the second sum is immaterial modulo 2. Bilinearity gives

$$\varepsilon(\mathbf{d}+\mathbf{e}) - \varepsilon(\mathbf{d}) - \varepsilon(\mathbf{e}) = \Omega(p(\mathbf{d}), p(\mathbf{e})) \pmod 2. \tag{8.4}$$

Also $\varepsilon(a\mathbf{d}) = a\varepsilon(\mathbf{d})$ modulo 2.

Lemma 8.2 (Multiplicity compatibility)**.** *After finite reordering, every output atom of auxiliary degree* $\mathbf{d}$ *has type* $\varepsilon(\mathbf{d})$*. For each* $\mathbf{d}$ *and parameter* h*, its grouped multiplicity is independent of the sequence of adjacent swaps. These multiplicities are compatible with setting any subset of the input variables to zero.*

Proof. The assertion about type holds for the input occurrences. Suppose two atoms of auxiliary degrees $\mathbf{d}_A, \mathbf{d}_B$ are reordered. Formula (7.5) gives type

$$ab\,\Omega(p(\mathbf{d}_A), p(\mathbf{d}_B)) + a\varepsilon(\mathbf{d}_A) + b\varepsilon(\mathbf{d}_B).$$

By (8.4), this is $\varepsilon(a\mathbf{d}_A + b\mathbf{d}_B)$. Induction proves the type assertion throughout the reordering.

The ordered ray factor is unique by Lemma 7.4, with the u_i retained. Take its ordinary logarithm. At auxiliary degree $\mathbf{d}$, every term from a multiple cover of order $\ell \geq 2$ has degree $\mathbf{d}/\ell$, and hence smaller total auxiliary degree. Subtract these contributions, which are already determined inductively. The remainder at $u^{\mathbf{d}} X^{p(\mathbf{d})}$ is

$$\frac{1}{t - t^{-1}} \sum_h n_{\mathbf{d},h} t^{h+1}. \tag{8.5}$$

There are finitely many summands at a fixed degree in the reordering of a finite word. The Laurent monomials t^{h+1} are linearly independent over $\mathbb{Q}$, so the remainder uniquely determines all $n_{\mathbf{d},h}$. The type has already been fixed by (8.3); thus these numbers determine the grouped atom list itself.

If an input variable is set to zero, precisely the output degrees involving that variable vanish. The initial product specializes to the product of the remaining occurrences, and ray-factor uniqueness and (8.5) show that every surviving multiplicity is unchanged. □

8.3. **The local positive reordering theorem.**

Theorem 8.3 (Positive reordering)**.** *Let C be a pointed rational two-dimensional cone in a lattice with integral alternating form Ω. If its restriction to* $\operatorname{span}(C)$ *is nonzero, choose the angular order $\prec$ so that an inverted pair of distinct rays has positive pairing. Every finite word of positive* E *and* H *atoms with roots in C has, in that angular order, a factorization into positive packets. The same statement holds for admissible lists of input occurrences. The resulting packets are admissible, and at each bounded charge and each fixed lower parameter bound their multiplicities are determined by finitely many input occurrences.*

When Ω vanishes on $\operatorname{span}(C)$*, the assertion holds in either order by commutativity.*

Proof. For a finite input word, Lemma 7.3 gives the factorization at each charge bound. Retain independent input variables. The uniqueness of ray factors and Lemma 8.2 show that the grouped multiplicities at different charge bounds agree wherever both are defined. Their limit is the required positive factorization.

For an infinite admissible list, fix a finitely supported auxiliary degree $\mathbf{d}$. Perform the finite construction using only the input occurrences in its support, in their original order. By Lemma 8.2, adjoining any other input occurrences does not change the result at degree $\mathbf{d}$. This defines all output multiplicities unambiguously.

Fix now a charge bound B and a lower output parameter bound J. By Lemma 8.1, an output of these sizes uses only inputs satisfying (8.2). There are finitely many such input occurrences by admissibility. There are finitely many auxiliary degrees of charge at most B supported on that finite set, and each such degree has finitely many parameters and multiplicities. Consequently only finitely many output atoms have charge at most B and parameter at least J. Their parameters have the uniform upper bound in Lemma 8.1. The output list is therefore admissible.

Coefficientwise convergence follows from the same bounds. At a fixed root degree, the number of nonconstant choices from factors is bounded by that degree. Quantum multiplication changes the energy by a bounded sum of root pairings. Parameters are bounded above; hence an arbitrarily negative parameter cannot contribute to a fixed energy, because all other contributions have a fixed upper bound. Among the finitely many remaining factors, the string indices j in (7.1) are also bounded at a fixed energy. Thus every coefficient is a finite sum of nonnegative integers, and the series identities follow by passing to finite charge and energy truncations. □

9. Finite wall systems and positive coefficients

We use the incoming-coefficient and change-of-root framework of [21, 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 $L^\vee_\mathbb{R} = \mathrm{Hom}_\mathbb{R}(L_\mathbb{R},\mathbb{R})$, and put $\mathbb{K} = \mathbb{Q}((t^{-1}))$. Thus an element of $\mathbb{K}$ is a series $\sum_{j\le j_0} a_j t^j$ with rational coefficients and an upper bound on its exponents. We write $\mathbb{N}((t^{-1}))$ for the subsemiring in which every a_j is a nonnegative integer. Multiplication is $X^aX^b = t^{\Omega(a,b)}X^{a+b}$. Fix a finite set $C \subset L$ contained in a lattice basis, and write

$$C^+ = \sum_{c\in C} \mathbb{N}c, \qquad \left|\sum_{c\in C} r_c c\right|_C = \sum_{c\in C} r_c, \qquad a \le_C b \iff b-a \in C^+. \tag{9.1}$$

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.

Use the atoms (7.1) with roots in $C^+ \setminus \{0\}$. Admissibility is measured by the charge $|a|_C$. Theorem 8.3 applies after restriction to the two-dimensional normal planes used below. Products are interpreted at bounded root degree and coefficientwise in t^{-1}.

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

$$V(a) = \Omega(-,a) \in L^\vee_\mathbb{R}, \qquad D_p = \mathsf{E}_{p,-1} = \prod_{j\ge 0}(1 + t^{-1-2j}X^p). \tag{9.2}$$

9.1. **The total element and its incoming prescription.** The group of units with positive root support is

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

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 9.1 (Three-factor decomposition)**.** *Every $g \in G_C$ has a unique factorization*

$$g = g_+(h)g_0(h)g_-(h), \qquad g_\sigma(h) \in G_{C,\sigma}(h). \tag{9.3}$$

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

$$F(h) = \mathrm{Ad}(g_-(h)g^{-1})F_-. \tag{9.4}$$

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 (9.3). Uniqueness identifies its middle factor with the prescribed wall element.

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

$$[X^r] \log g_0(V(r)) = \begin{cases} [X^r] \log D_p, & r = jp,\ p \in C,\ j \geq 1, \\ 0, & \text{otherwise.} \end{cases} \tag{9.5}$$

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,

$$\tau_r \log(UW) = \tau_r \log U + \tau_r \log W. \tag{9.6}$$

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 (9.6) 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.

9.2. **Positive crossing formulas.**

Lemma 9.2 (Positive crossing)**.** *Let $U_{a,k}$ be either $\mathsf{E}_{a,k}$ or $\mathsf{H}_{a,k}$, 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*

$$k - |b| + 2h, \qquad 1 \leq h \leq |b|. \tag{9.7}$$

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,

$$\mathrm{Ad}(D_p)^{\operatorname{sgn}\Omega(p,m)} X^m = \sum_{j=0}^{q} e_j(t^{1-q}, t^{3-q}, \ldots, t^{q-1}) X^{m+jp}, \qquad q = |\Omega(p, m)|. \tag{9.8}$$

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

$$\mathrm{Ad}(U_{a,k})X^m = X^m \frac{U_{a,k}(t^{2b}X^a)}{U_{a,k}(X^a)}. \tag{9.9}$$

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 (9.7), up to its order. All coefficients are nonnegative, and the constant selection has coefficient one. Taking $k=-1$ and $f(Z) = 1+Z$ gives (9.8). In particular, pairing one gives precisely $X^m + X^{m+p}$. □

Theorem 9.3 (Positive wall elements)**.** *At every generic point of a wall, and to every specified root order, its element is an admissible product of E/H atoms 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 9.1 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 8.3 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 (9.5) 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 9.2 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 $\mathbb{N}((t^{-1}))$. □

9.3. **Incoming coefficients and a decreasing recursion.** A *section* is a family of chamber values related by the consistent wall operations. Transporting a value supported in a finite union $U = \bigcup_{j=1}^{s}(m_j + C^+)$ defines a section. Transport and inverse transport add only positive roots, so every chamber value retains this support bound. For $w \in L$, define

$$A_w(F) = [X^w]F(V(w) + \epsilon), \tag{9.10}$$

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.

At the needed finite order, choose h in the target chamber so that the line $h + sV(w)$ avoids every joint. Such a choice exists because the forbidden points lie in finitely many proper linear subspaces. Follow the line 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

$$[X^w]F(h) = A_w(F) + \sum_i \sum_{\substack{j>0 \\ w-ja_i \in U}} J_i(w - ja_i, j)[X^{w-ja_i}]F(I_i). \tag{9.11}$$

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 9.4 (Propagation of incoming coefficients)**.** *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 (9.11), 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 $\mathbb{K}$. □

10. Change of roots and equality of Laurent coefficients

Fix $p \in C$ and put

$$\begin{aligned} &\alpha_c = \Omega(p,c), \qquad c' = c + [-\alpha_c]_+ p \quad (c \neq p), \\ &C' = \{-p\} \cup \{c' : c \in C \setminus \{p\}\}. \end{aligned} \tag{10.1}$$

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

$$S(w) = w - \Omega(p,w)p, \qquad S^\vee h = h \circ S^{-1}. \tag{10.2}$$

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

$$T(h) = \begin{cases} h, & h(p) \geq 0, \\ S^\vee h, & h(p) \leq 0. \end{cases} \tag{10.3}$$

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

10.1. **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

$$M_\sigma(u) = \sum_{c \neq p} [-\sigma \alpha_c]_+ u_c. \tag{10.4}$$

Lemma 10.1 (Finite mutation fibers)**.** *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*

$$0 \leq k \leq M_\sigma(u). \tag{10.5}$$

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

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

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 (9.8) 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,

$$L_\sigma(kp + u) = \sum_{c \neq p} u_c c' + (M_\sigma(u) - k)(-p). \tag{10.6}$$

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

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

Theorem 10.2 (Mutation covariance)**.** *The covector map (10.3), 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

$$F_- = \mathrm{Ad}(D_p)F'_-. \tag{10.7}$$

Their actual transported coefficients satisfy

$$F'(T(h)) = L_\sigma F(h), \qquad \sigma = \operatorname{sgn} h(p), \tag{10.8}$$

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

$$A_w(F) = A_{\mu_p(w)}(F'), \qquad \mu_p(w) = w + [-\Omega(p,w)]_+ p, \tag{10.9}$$

and μ_p is a bijection of L.

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

$$S \circ \mathrm{Ad}(D_p) = \mathrm{Ad}(D_{-p})^{-1}. \tag{10.10}$$

For $q = \Omega(p,m) \geq 0$, put $b_{q,j} = e_j(t^{1-q}, t^{3-q}, \dots, t^{q-1})$. The letters are invariant under inversion and have product one, giving $b_{q,j} = b_{q,q-j}$. Formula (9.8), 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 (10.10) 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

$$X^\ell f(Z) = f(t^{-2\Omega(p,\ell)}Z)X^\ell. \tag{10.11}$$

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 10.1 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 (9.9). 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 (10.6) 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 (10.10). 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 (10.6) 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 (9.6). A portion with no old wall preimage has identity element. The remaining pure incoming element is D_{-p}. The uniqueness of (9.5) 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 (10.10) and (10.7) 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 10.1 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 (10.9). Its label map preserves $\Omega(p,w)$, and hence its inverse subtracts $[-\Omega(p,w)]_+ p$. It is a bijection. □

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

Lemma 10.3 (A single incoming coefficient)**.** *Suppose $\Omega(c,m) \le 0$ for every $c \in C$. The section whose all-negative value is X^m satisfies*

$$A_w(F) = \begin{cases} 1, & w = m, \\ 0, & w \neq m. \end{cases} \tag{10.12}$$

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 $\Omega(r,w) > 0$. Since Ω is alternating, the hypothesis gives

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

No such contribution exists. At the far end the target coefficient is the incoming coefficient, proving (10.12). □

11. Compatible quantum seeds

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

11.1. **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

$$\widehat{B} = \begin{pmatrix} \widetilde{B} \\ I_n \end{pmatrix}, \qquad \widehat{\Lambda} = \begin{pmatrix} \Lambda & 0 \\ 0 & 0 \end{pmatrix}, \qquad p_i = -\widehat{b}_i. \tag{11.1}$$

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

$$\Omega_0 = \begin{pmatrix} 0 & -\Pi^T \\ \Pi & -B \end{pmatrix} \quad \text{on } \widehat{L}. \tag{11.2}$$

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

$$\begin{aligned} \widehat{\Lambda}(p_i, a) &= -d\, a_i^{\mathrm{old}}, \qquad \widehat{\Lambda}(p_i, p_j) = d b_{ij}, \\ \Omega_0(p_i, a) &= -a_i^{\mathrm{old}}, \qquad \Omega_0(p_i, p_j) = b_{ij}. \end{aligned} \tag{11.3}$$

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

$$(\widehat{\Lambda} - d\Omega_0)(r, a) = 0 \qquad (r \in \mathsf{N}, a \in \widehat{L}). \tag{11.4}$$

Lemma 11.1 (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 10.2.*

Proof. For $r \in \mathsf{N}$ and $a \in \widehat{L}$, equation (11.4) 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. The field $\mathbb{K}_v$ contains the coefficients of every

actual BZ Laurent expression. After $t = v^d$, the positive crossing formulas and (9.11) preserve $\mathbb{N}((v^{-1}))$.

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 9.2. 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 9. We then apply this root bimodule map and specialize in the actual central torus (11.1).

11.2. **The exact exchange relation.** At the initial extended seed, set $Ee_k = -e_k + [\widehat{b}_k]_+$ and $Ee_j = e_j$ for $j \neq k$. 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 $D_{p_k} = \mathsf{E}_{p_k,-1}$ therefore satisfies

$$\begin{aligned} \mathrm{Ad}(D_{p_k})X^g &= X^g + X^{g+p_k} \\ &= X^{-e_k+[\widehat{b}_k]_+} + X^{-e_k+[-\widehat{b}_k]_+}. \end{aligned} \tag{11.5}$$

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 11.1 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

$$\widehat{B}' = E\widehat{B}E_0^T, \qquad \widehat{\Lambda}' = E^T\widehat{\Lambda}E. \tag{11.6}$$

For $i, j \neq k$, the (i, j) entry of the first product is

$$\widehat{b}_{ij} + [\widehat{b}_{ik}]_+\widehat{b}_{kj} + \widehat{b}_{ik}[-\widehat{b}_{kj}]_+,$$

which is the exchange-matrix mutation rule; entries in row or column k change sign. Identifying the new exponent lattice with the old one by E, its

negative exchange columns become

$$-E\widehat{b}'_k = -p_k, \qquad -E\widehat{b}'_j = p_j + [-b_{kj}]_+ p_k \quad (j \neq k). \tag{11.7}$$

This is the root change in Theorem 10.2, 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 $E_0^2 = 1$,

$$(\widehat{B}')^T \widehat{\Lambda}' = dE_0\Pi_0 E = d\Pi_0. \tag{11.8}$$

11.3. **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 (11.6) and (11.8) successively in seed coordinates gives $\widehat{B}_j^T \widehat{\Lambda}_j = d\Pi_0$ and makes C_j a basis of N. Since $p_{j,i} \in \mathsf{N}$, (11.4) yields

$$\Omega_0(p_{j,i}, J_j a) = d^{-1}\widehat{\Lambda}(p_{j,i}, J_j a) = d^{-1}\widehat{\Lambda}_j(-\widehat{b}_{j,i}, a) = -a_i. \tag{11.9}$$

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 (11.5) gives the rational identity

$$P_j = \mathrm{Ad}(D_p) P_{j+1}, \qquad p \in C_j \text{ the root mutated at step } j. \tag{11.10}$$

The two finite values are the all-negative initial values for the respective wall systems. Lemma 11.1 and Theorem 10.2 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 (11.9). Lemma 10.3 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 9.4, all coefficients of P_0 belong to $\mathbb{N}((v^{-1}))$. They were already elements of $\mathbb{Z}[v^{\pm 1}]$; hence

$$\mathbb{Z}[v^{\pm 1}] \cap \mathbb{N}((v^{-1})) = \mathbb{N}[v^{\pm 1}] \tag{11.11}$$

proves positivity in the central extension.

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

$$\mathrm{ev}_Y : \mathcal{T}_{\widehat{\Lambda}} \longrightarrow \mathcal{T}_{\Lambda}, \qquad X^{(b,c)} \longmapsto X^b, \qquad Y_i \longmapsto 1. \tag{11.12}$$

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

$$\begin{aligned} \widehat{M}'(e_k)\widehat{M}(e_k) = {} & v^{\widehat{\Lambda}(g_+,e_k)}\widehat{M}([\widehat{b}_k]_+) \\ & + v^{\widehat{\Lambda}(g_-,e_k)}\widehat{M}([-\widehat{b}_k]_+). \end{aligned} \tag{11.13}$$

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. Its mutation matrix fixes the coordinate vectors of all these variables, so their pairings and centering factors are unchanged. 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. □

References

[1] G. Barthel, J.-P. Brasselet, K.-H. Fieseler and L. Kaup, *Combinatorial intersection cohomology for fans*, Tohoku Math. J. (2) **54** (2002), 1–41; arXiv:math/0002181v1.

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

[3] P. Bressler and V. A. Lunts, *Intersection cohomology on nonrational polytopes*, Compos. Math. **135** (2003), 245–278; arXiv:math/0002006v1.

[4] A. Burcroff and K. Lee, *Broken lines and compatible pairs for rank 2 quantum cluster algebras*, Selecta Math. (N.S.) **32** (2026), article 81. doi:10.1007/s00029-026-01189-z.

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

[6] B. Davison and T. Mandel, *Strong positivity for quantum theta bases of quantum cluster algebras*, Invent. Math. **226** (2021), 725–843; arXiv:1910.12915v3.

[7] 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.

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

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

[10] M. Gross, P. Hacking, S. Keel and M. Kontsevich, *Canonical bases for cluster algebras*, J. Amer. Math. Soc. **31** (2018), 497–608; arXiv:1411.1394v2.

[11] T. Hausel and B. Sturmfels, *Toric hyperkähler varieties*, Doc. Math. **7** (2002), 495–534. doi:10.4171/DM/130.

[12] M. Huang, *An expansion formula for quantum cluster algebras from unpunctured triangulated surfaces*, Selecta Math. (N.S.) **28** (2022), article 21. doi:10.1007/s00029-021-00750-2.

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

[14] 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.

[15] K. Lee and R. Schiffler, *Positivity for cluster algebras*, Ann. of Math. (2) **182** (2015), 73–125; arXiv:1306.2415v3.

[16] S. Meinhardt and M. Reineke, *Donaldson–Thomas invariants versus intersection cohomology of quiver moduli*, J. Reine Angew. Math. **754** (2019), 143–178. doi:10.1515/crelle-2017-0010.

[17] S. Mozgovoy and M. Reineke, *Abelian quiver invariants and marginal wall-crossing*, Lett. Math. Phys. **104** (2014), 495–525. doi:10.1007/s11005-013-0671-0.

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

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

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

[21] OpenAI, *Elementary positivity of chromatic quasisymmetric functions*, September 24, 2026. OpenAI Math Release, commit `adc7f1241b42`.
[22] F. Qin, *Quantum cluster variables via Serre polynomials*, J. Reine Angew. Math. **668** (2012), 149–190. doi:10.1515/CRELLE.2011.129.
[23] M. Reineke, *Wild quantum dilogarithm identities*, Ann. Represent. Theory **1** (2024), no. 3, 385–391. doi:10.5802/art.12.
[24] M. Reineke, B. Rhoades and V. Tewari, *Zonotopal algebras, orbit harmonics, and Donaldson–Thomas invariants of symmetric quivers*, Int. Math. Res. Not. IMRN **2023**, no. 23, 20169–20210. doi:10.1093/imrn/rnad033.
[25] M. Reineke, J. Stoppa and T. Weist, *MPS degeneration formula for quiver moduli and refined GW/Kronecker correspondence*, Geom. Topol. **16** (2012), 2097–2134. doi:10.2140/gt.2012.16.2097.
[26] 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.
[27] A. Szenes and O. Trapeznikova, *Intersection cohomology of type-A toric varieties*, Algebr. Comb. **8** (2025), no. 2, 575–596. doi:10.5802/alco.417.
[28] Q. Tang, *Another proof of Laurent positivity for skew-symmetric quantum cluster algebras*, November 2026; Hexagon 2611.00239v1.

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