hexagon:2611.00790v1

CRYSTAL STRUCTURES ON THETA BASES

JIARUI FEI

Abstract. We construct crystal operators on the regular theta functions of skew-symmetrizable cluster algebras with polynomial frozen variables. A frozen vertex that becomes a sink under mutation defines a locally nilpotent derivation of the upper cluster algebra. These derivations satisfy the Serre relations. Under mutable reddening, the regular theta basis of the middle algebra is perfect for both the sink and the source derivations. In each seed the g-vectors of its elements are the integral points of a rational polyhedral cone.

A tropical criterion detects the sink and source operators that form Chevalley pairs. Such pairs give integrable Kac–Moody actions and, under reddening, normal crystals on the regular theta basis. In the skew-symmetric case with mutable reddening and finite-dimensional Jacobian algebra, we compute the order of the theta function attached to a supported presentation weight along each frozen divisor. We illustrate this theory on various examples.

Contents

Introduction 1
1. Theta functions and coefficients 5
2. Canonical boundary derivations and degrees 7
3. Regular theta cones and tropical crystals 9
4. Tropical DT theory 14
5. Primitive theta lifts and changes of coefficients 19
6. Exact boundary pairs and Lie actions 22
7. Perfect theta bases and crystal structures 27
8. Cartan type and geometric applications 32
9. Theta F-pairings and maximal adjunction 38
10. Maximal operators and string coordinates 43
11. Pointed biperfect bases and transitions 46
Declaration of generative AI and AI-assisted technologies 48
References 49

Introduction

A frozen variable in a cluster algebra determines a divisor. When mutation makes its vertex a sink, it also determines a derivation. We study how these derivations act on theta functions and when they are the simple-root operators of an integrable Lie algebra action. In [12] the corresponding operators were constructed from boundary representations of quivers with potential. Here the construction uses the valued exchange matrix and extends to skew-symmetrizable seeds.

The natural domain for the crystal is the regular middle algebra. Its basis consists of the finite theta functions that extend across the frozen divisors. The derivations act on the entire upper cluster algebra, and the crystal describes a basis of that algebra whenever upper theta spanning holds. In the skew-symmetric setting, Proposition 4.9 gives a criterion, in terms of

2020 *Mathematics Subject Classification.* Primary 13F60; Secondary 05E10, 17B37, 14T20.

Key words and phrases. Theta basis, crystal structure, biperfect basis, skew-symmetrizable cluster algebra, Lagrangian Grassmannian, tropical DT transformation.

The author was supported in part by the National Natural Science Foundation of China (Nos. 12571038).

frozen vanishing orders, for the supported generic characters of [12, Remark 7.8] to span the regular middle algebra.

0.1. **Tropical crystals.** Let $V = U \sqcup F$ be the vertex set of an ice seed, with mutable vertices U and frozen vertices F. We use the row exchange convention of [12]: a positive integral diagonal matrix $D = \operatorname{diag}(d_a)$ satisfies $BD = -(BD)^{\mathsf{T}}$. The extended exchange matrix in the usual column convention is $\widetilde{B} = -B_{U,V}^{\mathsf{T}}$, whose principal part has left symmetrizer D_U.

Over a field $\Bbbk$ of characteristic zero, we work with the upper cluster algebra in which the frozen variables are not inverted,

$$\mathcal{U}^+ = \bigcap_t \Bbbk[x_u(t)^{\pm 1} : u \in U][x_f : f \in F],$$

and write $\mathcal{U}^{\mathrm{loc}}$ for the upper algebra obtained by inverting them as well. The intersections range over all seeds reached by mutation. We use the theta functions of [21], defined by specializing auxiliary principal coefficients to 1 while leaving the original frozen variables unchanged. A theta function is *finite* if its principal-coefficient lift is a Laurent polynomial, and *regular* if it belongs to $\mathcal{U}^+$. An integral tropical point has a g-vector $g(t)$ in each seed t. When the seed is understood we write g for this vector and ϑ_g for the corresponding theta function. In these coordinates, $\mathcal{B}_\theta^+$ denotes the set of g-vectors of regular finite theta functions.

A frozen vertex i is *sink-reachable* if mutations lead to a seed t_i with $b_{iu} \leqslant 0$ for all $u \in U$. Such a seed is an *optimizing seed* for i. In this seed define

$$M_i = x^{m_i} = \prod_{u \in U} x_u^{-b_{iu}}, \qquad v_i = m_i - e_i.$$

The monomial M_i does not depend on the optimizing seed (Lemma 2.1). A frozen vertex is *source-reachable* when the opposite inequalities hold in some seed; its corresponding monomial is denoted by $M_j^\star$. For a chosen set I of sink-reachable vertices, write a_{ij} for the ordinary x_i-degree of M_j in an i-optimizing seed. Section 2 proves that this degree is independent of the optimizing seed. The matrix $C = (c_{ij})$, defined by $c_{ii} = 2$ and $c_{ij} = -a_{ij}$ for $i \neq j$, is a symmetrizable generalized Cartan matrix. Once the derivation R_i is introduced, a_{ij} is also its nilpotence degree on M_j.

Theorem A (Tropical theta crystals)**.** *Suppose the mutable cluster pattern admits a reddening sequence. In each seed, $\mathcal{B}_\theta^+$ consists of the integral points of a rational polyhedral cone $\mathcal{C}^+$. For every sink-reachable i, tropical transport to an optimizing seed defines partial maps*

$$r_i(g) = g + v_i \quad (g_i > 0), \qquad l_i(g) = g - v_i \quad (g - v_i \in \mathcal{C}^+),$$

where an undefined value is denoted by 0_{cr}. These maps do not depend on the optimizing seed. They are partial inverses, and $\rho_i(g) = g_i$ is the number of times r_i can be applied to g. In particular, maximal raising is the retraction $g \mapsto g + g_i v_i$ onto $\mathcal{C}^+ \cap \{g_i = 0\}$.

If compatible integral gradings are chosen so that the degree of r_i in the jth grading is c_{ji}, these operators define an upper seminormal crystal of type C. The opposite seed gives the corresponding operators for source-reachable vertices.

Theta reciprocity expresses the inequalities of the regular cone in terms of Newton polytopes of dual theta functions [8]. We prove that these polytopes have integral vertices and use their inequalities to show that tropical raising preserves regularity. Upper seminormality records the length of each raising string. Before one chooses exact source partners, an inverse string may be infinite. Lemma 3.11 characterizes the existence of compatible rational gradings, and Proposition 3.12 gives sufficient conditions for their existence. This construction is the theta analogue of the crystal on supported presentation weights in [12, Section 5].

0.2. **Boundary derivations and theta perfectness.** The algebraic counterpart of tropical raising is the *canonical boundary derivation* $R_i = M_i\partial_{x_i}$ in an i-optimizing seed. Source-reachable vertices similarly give $R_j^\star = M_j^\star\partial_{x_j}$. We write $\mathfrak{n}_+(C)$ for the Lie algebra generated by the positive Chevalley generators for C, and $U(\mathfrak{n}_+(C))$ for its universal enveloping algebra. Define the *regular middle theta algebra* by $\mathrm{Mid}^+ = \mathrm{Mid}^{\mathrm{loc}} \cap \mathcal{U}^+$, where $\mathrm{Mid}^{\mathrm{loc}}$ is the span of finite specialized theta functions.

Theorem B (Boundary actions and perfectness)**.** *The derivations R_i and $R_j^\star$ are canonical and locally nilpotent on $\mathcal{U}^+$. For any set I of sink-reachable frozen vertices, the assignments $e_i \mapsto R_i$ satisfy the Serre relations for C and give an action of $\mathfrak{n}_+(C)$ by derivations. Equivalently, $\mathcal{U}^+$ is a $U(\mathfrak{n}_+(C))$-module algebra. These conclusions require neither reddening nor theta spanning.*

If the mutable pattern admits a reddening sequence, both types of boundary derivations preserve Mid^+*. Its regular theta basis is normalized perfect for each of them. In an i-optimizing seed, the leading terms satisfy*

$$R_i\vartheta_g = g_i\vartheta_{r_i(g)} + \sum_{h_i \leqslant g_i - 2} c_h\vartheta_h, \qquad R_i^{(g_i)}\vartheta_g = \vartheta_{g+g_iv_i},$$

for $g_i > 0$, with $R_i^{(a)} = R_i^a/a!$.

The Lie action extends the positive-part construction of [12, Theorem 8.3] to skew-symmetrizable exchange data. Its Serre relations follow from canonical boundary degrees and Cao's F-invariant [6]. Perfectness uses the highest frozen face of a theta function and Qin's freezing theorem [35]. Here *normalized perfect* means that the leading coefficient is the nilpotence degree, as in the theorem. The source and sink derivations give two perfect structures; they are not assumed to be paired or to commute.

0.3. **Tropical DT and exact boundary pairs.** The positive-part action does not by itself give the negative Chevalley generators. We obtain them when an i-sink derivation can be matched with a source derivation at a frozen vertex k. In an i-optimizing seed, this *exact boundary pairing* is expressed by

$$g(M_k^\star) = e_i + e_k - m_i.$$

To recognize this identity without first finding both optimizing seeds, we use the *mutable tropical Donaldson–Thomas transformation* T and the Newton polytopes of principal-coefficient F-polynomials.

For every frozen vertex j, write $a_j = B_{j,U}$, the mutable part of its exchange row. For a mutable projective weight b, let F_b be the principal-coefficient F-polynomial of the theta function with mutable g-vector $-b$. When it has a componentwise greatest exponent f_b, define

$$\mathsf{E}(a,b) = \max_{\nu \in \operatorname{supp} F_b} (-aD_U\nu^{\mathsf{T}}), \qquad \mathsf{H}(a,b) = aD_Uf_b^{\mathsf{T}} + \mathsf{E}(a,b).$$

These expressions are support functions of the Newton polytopes of F_b and its reversal; in general neither is symmetric in its arguments. Under reddening, F_b is finite and f_b exists for every integral projective weight. We orient T by $T^{-1}b = -b - f_bB_{UU}$.

Theorem C (Tropical criterion for exact pairs)**.** *Under reddening, $\mathsf{E}(a,b) = \mathsf{H}(b,Ta)$ for all integral mutable projective weights a, b. A sink-reachable vertex i and a source-reachable vertex k form an exact boundary pair if and only if*

$$T^{-1}a_i = Ta_k, \qquad \mathsf{H}(a_j,a_i) + \mathsf{H}(T^{-1}a_i, Ta_j) = d_j(\delta_{ij} + \delta_{kj}) \quad (j \in F).$$

The first equality compares mutable projective weights. The remaining equalities recover the frozen coordinates, including those at vertices unrelated to either boundary. They give a purely tropical counterpart of [12, Lemma 4.29]. We also obtain a criterion using only the two boundary F-polynomials, which does not require reddening (Proposition 4.2).

0.4. **Integrable actions and normal crystals.** For an exact system, choose a source partner $\sigma(i)$ for each selected sink vertex i, and set $L_i = R^\star_{\sigma(i)}$ and $H_i = [R_i, L_i]$. We use *normal crystal* in the usual sense for symmetrizable Kac–Moody algebras: the restriction to every finite-type Levi subalgebra is a disjoint union of finite-dimensional highest-weight crystals. This does not require the entire crystal to be of finite type or to be a union of affine highest-weight crystals.

Theorem D (Integrable actions)**.** *Suppose every selected sink vertex has an exact source partner. Then σ is injective, $d_i = d_{\sigma(i)}$, and the operators R_i, L_i, H_i define an integrable action of the derived Kac–Moody algebra $\mathfrak{g}'(C)$ on $\mathcal{U}^+$. The H_i are integral compatible gradings, and every connected component of C is of finite or affine type.*

Under reddening, this action preserves Mid^+ *and its regular theta basis gives a normal crystal. The same conclusion holds on $\mathcal{U}^+$ without reddening whenever theta functions with finite principal-coefficient lifts form a basis of $\mathcal{U}^+$.*

The commutators themselves supply the Cartan gradings. Rank-one perfectness and restriction to finite-type Levi subalgebras then give normality, as in the generic-basis setting of [12, Section 8.4]. The restriction on Cartan type comes from a different argument. A Lie algebra of derivations generated by finitely many elements has polynomial growth on a field of finite transcendence degree (Lemma 8.1). Exact pairing is an additional condition on the frozen exchange data, even in the presence of reddening.

0.5. **Maximal operators, basis changes, and applications.** Fix a seed and assume mutable reddening. For each $p \in \mathbb{Z}^U$ there is a unique regular theta function whose g-vector has mutable part p and whose order along every frozen divisor is zero. These functions form a basis of the middle algebra over the polynomial ring in the frozen variables (Proposition 5.1). In the skew-symmetric setting of [12], assume also that the full Jacobian algebra is finite dimensional. Extending a general mutable decorated representation of weight $-p$ by zero gives a second choice of frozen coordinates. Proposition 4.9 computes the difference between the two choices at each frozen vertex. The difference vanishes at every vertex that admits an optimizing seed. The proof uses Qin's unitriangular comparison [34, Theorems 5.1.3 and 5.2.2] in a principal extension, so the extended exchange matrix of the original seed may have deficient rank.

Maximal raising sends all cluster monomials of a given seed into a single cluster (Theorem 10.4). For the theta valuation pairing, we determine the cone of directions on which maximal raising admits an adjoint. A single additional cluster monomial, together with the seed variables, tests membership in this cone (Theorem 9.8).

On the base affine space, an opposite cluster automorphism realizes the pinned Lusztig involution (Theorem 8.6). On the Plücker coordinate ring of $\mathrm{LGr}(n, 2n)$, the additional affine color acts by the negative highest-root derivation up to a pinning scalar. Together with the finite colors it gives a level-zero $C_n^{(1)}$-crystal (Proposition 8.8).

0.6. **Conventions.** We write g-vectors as rows and valuation directions as columns, and include zero in $\mathbb{N}$. Unless otherwise stated, mutation is at a mutable vertex; mutation at a frozen vertex in an enlarged pattern will be called *auxiliary mutation*. An undefined crystal operator takes the value 0_{cr}. The presentation weights of [12] correspond to $\delta = -g$. We follow its notation: $R_i, R_j^\star$ for derivations, $r_i, l_i, r_j^\star$ for crystal operators, and ρ_i, λ_i for their string functions. For an exact pair, $L_i = R^\star_{\sigma(i)}$ and $l_i = r^\star_{\sigma(i)}$. The symbols E, H denote the tropical support functions defined above; $\mathcal{T}_\gamma$ denotes a change of coefficient seed.

We invoke *upper theta spanning* only when asserting that $\mathrm{Mid}^+ = \mathcal{U}^+$: it means that specializations of theta functions with finite principal-coefficient lifts span $\mathcal{U}^{\mathrm{loc}}$. Valuative independence then identifies the polynomial-frozen subbasis.

1. Theta functions and coefficients

As in the generic-character setting of [12, Section 7], we distinguish the polynomial-frozen upper cluster algebra from its localization at the frozen variables. Here theta functions must also be distinguished according to finiteness and regularity. We use the row exchange matrix B and symmetrizer D of the introduction. For the chiral-dual lattice construction, choose a positive integer ℓ divisible by every d_a. The extended matrix in the column convention is $\widetilde{B} = (-B^{\mathsf{T}})_{V,U}$. Mutation at $k \in U$ gives

$$x_k x_k' = \prod_{a \in V} x_a^{[b_{ka}]_+} + \prod_{a \in V} x_a^{[-b_{ka}]_+}.$$

The frozen weights d_f belong to the seed datum. We follow the standard mutation rules of [17, 18], with the row convention specified in the introduction. Entries of B between two frozen indices do not affect the exchange relations.

For $F_0 \subseteq F$, let

$$\mathcal{U}_{F_0} = \bigcap_t \Bbbk[x_u(t)^{\pm 1} : u \in U][x_f^{\pm 1} : f \in F \backslash F_0][x_f : f \in F_0].$$

In particular, $\mathcal{U}^+ = \mathcal{U}_F$ and $\mathcal{U}^{\mathrm{loc}} = \mathcal{U}_\varnothing$. We write $\mathcal{A}^{\mathrm{loc}}$ for the corresponding cluster algebra, generated by all mutable cluster variables with the frozen variables inverted. A *reddening sequence* is a mutation sequence whose terminal principal-coefficient c-matrix is a negative permutation matrix; see [26]. Tropical duality [32] identifies this condition with the injective reachability used in [34].

We construct theta functions from principal coefficients and specialize the auxiliary coefficients to 1. We use *finite theta function* in the sense fixed in the introduction: its principal-coefficient lift is a Laurent polynomial. The original frozen variables are not specialized, and remain polynomial in $\mathcal{U}^+$. We use the positivity, pointedness, and mutation covariance of theta functions from [21]. When the mutable pattern admits a reddening sequence, every integral tropical point indexes a finite principal theta function.

Lemma 1.1. *Suppose the finite theta functions form a basis of $\mathcal{U}^{\mathrm{loc}}$. Then*

$$\Theta_{F_0} = \{\vartheta_g : \mathrm{ord}_{x_f}(\vartheta_g) \geqslant 0 \text{ for all } f \in F_0\}$$

is a basis of $\mathcal{U}_{F_0}$.

Proof. For each frozen index f, the valuation $\nu_f = \mathrm{ord}_{x_f}$ is unchanged by mutation. Indeed, at least one exchange monomial is independent of x_f, and the exchange polynomial remains nonzero modulo x_f. Mutation therefore induces a birational change of coordinates at the generic point of the divisor $x_f = 0$. It follows that

$$\mathcal{U}_{F_0} = \{u \in \mathcal{U}^{\mathrm{loc}} : \nu_f(u) \geqslant 0 \text{ for all } f \in F_0\}.$$

By valuative independence after coefficient specialization [8, Theorem 3.10], a finite theta expansion satisfies $\nu_f(\sum_g c_g \vartheta_g) = \min_{c_g \neq 0} \nu_f(\vartheta_g)$. Thus an expansion belongs to $\mathcal{U}_{F_0}$ precisely when each of its terms does. The isolated zero-column case is covered by Remark 1.2. □

Remark 1.2. A zero extended exchange column is an isolated mutable vertex and has exchange relation $xx' = 2$. Its theta functions are $\vartheta_a = 2^{\max(-a,0)} x^a$. Since the characteristic is zero, these nonzero scalar multiples of Laurent monomials satisfy valuative independence.

For the remaining scattering diagram, the exchange map is $P = \widetilde{B}$. In the notation of [8, Lemma 4.3], the skew form on coefficient degrees is $B^\bullet = (Q^\bullet)^{\mathsf{T}} P$, where $Q^\bullet e_u = (d_u/\ell) e_u^*$. A noninitial wall cannot have zero x-direction: if its final bracket combines degrees v_1, v_2 with $P(v_1 + v_2) = 0$, then $B^\bullet(v_1, v_2) = B^\bullet(v_1, v_1 + v_2) = 0$, so the bracket vanishes. Thus the diagram satisfies the hypothesis for coefficient specialization in [8, Assumption 3.9 and Theorem 3.10]. Tensoring with the isolated Laurent factors proves the assertion for all extended exchange matrices.

We record two sufficient conditions for theta spanning.

Proposition 1.3. *Suppose $\widetilde{B}$ has full column rank and the mutable pattern admits a reddening sequence. Then the theta functions form a basis of $\mathcal{U}^{\mathrm{loc}}$, and Θ_{F_0} is a basis of $\mathcal{U}_{F_0}$ for every $F_0 \subseteq F$.*

Proof. The localized assertion is [34, Theorem A.1.5], which applies the large cluster complex theorem of [21] under the full-rank convention. Lemma 1.1 gives the second assertion. □

Proposition 1.4. *The conclusion of Proposition 1.3 also holds if full column rank is replaced by the existence, in some seed, of a real linear functional ℓ_{cone} such that $\ell_{\mathrm{cone}}(\widetilde{B}e_k) > 0 \quad (k \in U)$. Equivalently,* $\ker \widetilde{B} \cap \mathbb{R}^U_{\geqslant 0} = \{0\}$.

Proof. This is the positive-cone case of the convexity condition in [21, Theorem 0.3(7)]. Reddening gives a large cluster complex, so [21, Proposition 0.7] gives the theta basis of $\mathcal{U}^{\mathrm{loc}}$. We then apply Lemma 1.1. The equivalence is the theorem of alternatives for a finite family of vectors. Full column rank implies the condition by assigning value 1 to each independent exchange column. □

1.1. **Opposite theta functions.** The *opposite seed* has the same variables and frozen indices and the exchange matrix $-B$. We identify its rational function field with the original one by fixing the initial variables.

Proposition 1.5. *If the mutable pattern admits a reddening sequence, the original and opposite sets of finite theta functions coincide. This holds in arbitrary extended rank, with theta functions defined by specialization from principal coefficients.*

Proof. In full extended rank this is [36, Proposition 4.2]. Reddening supplies its injective-reachability hypothesis by tropical duality. The proof uses the invertible DT twist and its action on finite theta functions [27, Theorems 5.2 and 6.2(2)]; its Laurent polynomial identities hold over every field of characteristic zero. For arbitrary rank, adjoin independent principal frozen coefficients. The resulting matrix has full column rank and the same mutable principal part, so the equality applies to this extension and its opposite. Every theta function in the extension is an auxiliary frozen Laurent monomial times a theta function whose g-vector has zero auxiliary coordinates. Specialization at 1 therefore maps the two theta sets onto the original and opposite theta sets, respectively. Their equality descends. □

Lemma 1.6. *Suppose the mutable pattern admits reddening. A frozen vertex is sink-reachable if and only if it is source-reachable. More precisely, if t is a sink optimizer, its predecessor $t[-1]$ is a source optimizer. If t is a source optimizer, its successor $t[1]$ is a sink optimizer.*

Proof. Write $a(t) = B_{i,U}(t)$. Matrix mutation shows that $-a(t)$ transforms as a mutable g-vector for the column exchange matrix $-B^{\mathsf{T}}_{U,U}$. Work in the principal-coefficient pattern, since the mutable g-vector transition is independent of the frozen rows. By injective reachability, the mutable g-vectors of the variables of $t[1]$, measured at t, are the negatives of the standard basis vectors up to permutation. Likewise those of the variables of t, measured at $t[-1]$, are the negatives of the standard basis vectors up to permutation [34, Definition 2.3.1 and Proposition 2.3.3].

If $a(t) \leqslant 0$, then $-a(t)$ is the degree of a monomial in the mutable variables of t. Its mutable degree at $t[-1]$ is therefore nonpositive. By tropical transport this degree is $-a(t[-1])$, so $t[-1]$ is a source optimizer. If $a(t) \geqslant 0$, choose a monomial in the mutable variables of $t[1]$ with degree $-a(t)$ at t. Its degree at $t[1]$ is its nonnegative exponent vector, giving $a(t[1]) \leqslant 0$ as required. □

1.2. **Degrees and codegrees.** The *codegree* of ϑ_g is its pointed g-vector for the opposite seed. In deficient rank, both degree and codegree are computed from the principal-coefficient lift. Let $\mathcal{F}_g$ denote the associated F-polynomial computed from the principal-coefficient theta lift. With our row convention, $\widehat{y}_u = x^{-B_{u,V}}$ and $\vartheta_g = x^g \mathcal{F}_g(\widehat{y})$.

Proposition 1.7. *Under reddening, $\mathcal{F}_g$ has a componentwise largest exponent $f(g) \in \mathbb{N}^U$, with coefficient 1. The codegree of ϑ_g and its F-polynomial in the opposite seed satisfy*

$$\eta_\theta(g) = \operatorname{codeg}(\vartheta_g) = g - f(g)B_{U,V}, \qquad \mathcal{F}^{\mathrm{op}}_{\eta_\theta(g)}(y) = y^{f(g)}\mathcal{F}_g(y^{-1}).$$

These identities remain valid after geometric coefficient specialization.

Proof. First suppose the extended exchange matrix has full rank. Proposition 1.5 gives

$$x^g\mathcal{F}_g(x^{-B_{U,V}}) = x^h G(x^{B_{U,V}}), \qquad \mathcal{F}_g(0) = G(0) = 1.$$

The monomial x^h on the right occurs on the left, so $h = g - fB_{U,V}$ for an exponent f of $\mathcal{F}_g$. Injectivity of the exchange map matches each exponent e of $\mathcal{F}_g$ with an exponent $f - e$ of G. Thus $e \leqslant f$ componentwise, and the coefficient of y^f equals the constant coefficient of G, namely 1. This proves both formulas.

In arbitrary rank, apply this argument to the principal lift $x^g\mathcal{F}_g(x^{-B_{U,V}}y)$. Its opposite expression is $x^h y^k G(x^{B_{U,V}}y^{-1})$. The augmented exchange map is injective, giving $k = f$, $h = g - fB_{U,V}$, and the same polynomial identity. Specializing $y = 1$ completes the proof. □

Consequently, the localized DT twist of [27] acts on g-vectors by

$$\mathrm{DT}_\theta(g) = -\eta_\theta(g) = -g + f(g)B_{U,V}.$$

The twist takes the opposite theta function indexed by m to the original theta function indexed by $-m$. On mutable g-vectors, this is inverse tropical transport to the injective seed, with the normalization of [27]. We use this convention in Section 4.

2. Canonical boundary derivations and degrees

The derivations in this section are the exchange-theoretic analogues of those attached to boundary representations in [12, Sections 4 and 8.1]. We construct their monomial multipliers directly from optimizing seeds. In a sink optimizing seed for a frozen vertex i, define

$$M_i = \prod_{u\in U} x_u^{-b_{iu}}, \qquad \zeta_i = (b_{if})_{f\in F}, \qquad Z_i^+ = x^{\zeta_i}, \qquad R_i = M_i\partial_{x_i}.$$

We extend ζ_i by zero on U. A *source optimizing seed* satisfies the reverse inequalities. We first show that these monomials and derivations do not depend on the optimizing seed.

2.1. Canonical monomials and the degree formula.

Lemma 2.1. *If i has a sink optimizing seed, both M_i and ζ_i are independent of its choice. In every seed t,*

$$g(M_i; t) = -\operatorname{row}_i B(t) + \zeta_i.$$

Proof. The identity holds in the chosen optimizing seed. The matrix mutation rule shows that its right side transforms by tropical g-mutation for $-B^{\mathsf{T}}$; adding the frozen vector ζ_i does not change the mutable coordinate controlling the mutation. In any other optimizing seed, the mutable coordinates are nonnegative. The cluster monomial property of theta functions [21] therefore gives $M_i = \prod_{u\in U} x_u^{-b_{iu}} \prod_{f\in F} x_f^{\zeta_{if}-b_{if}}$. Every mutable cluster variable has order zero at each frozen divisor, as in the proof of Lemma 1.1. Taking these orders gives $\zeta_{if} = b_{if}$ and proves the claim. □

For mutable cluster monomials P, Q, let g_P^t, g_Q^t be their mutable g-vectors, and let F_P^t, F_Q^t be their F-polynomials computed with principal coefficients. For a polynomial F and a column vector q, write $F[q] = \max_{e\in\operatorname{supp} F} eq$. Cao's *$F$-invariant* is

$$\mathsf{F}_D(P, Q) = F_P^t[D_U(g_Q^t)^{\mathsf{T}}] + F_Q^t[D_U(g_P^t)^{\mathsf{T}}].$$

It is seed-independent, symmetric, and nonnegative. It vanishes precisely when PQ is a cluster monomial [6, Theorems 4.10 and 4.19, Proposition 4.30]. The principal-coefficient construction makes the invariant independent of the original geometric coefficients.

In an i-sink optimizing seed write $M_i = x^m$, with $m_u = -b_{iu} \geqslant 0$, and let $\rho_i(P)$ be the degree of P in x_i.

Proposition 2.2. *For every mutable cluster monomial P, $d_i\rho_i(P) = \mathsf{F}_D(M_i, P)$.*

Proof. In the separation formula $P = x^g F_P(\widehat{y})$, a term with coefficient exponent e has x_i-exponent $g_i - d_i^{-1}\langle mD_U, e\rangle$, since $\widetilde{b}_{iu} = -b_{ui} = -(d_u/d_i)m_u$. All these exponents are at most g_i, and the constant term of F_P realizes equality. Their minimum is zero because P has order zero along $x_i = 0$. Positivity therefore gives

$$\rho_i(P) = g_i = \frac{1}{d_i} F_P[D_U m^{\mathsf{T}}].$$

Since $F_{M_i} = 1$ and $g_{M_i} = m$, we have $F_P[D_U m^{\mathsf{T}}] = \mathsf{F}_D(M_i, P)$, proving the formula. □

Corollary 2.3. *The rational derivation R_i is independent of the optimizing seed. A mutable cluster monomial lies in its kernel precisely when it is compatible with M_i.*

Proof. By Lemma 2.1, the same M_i is a cluster monomial in every optimizing seed. Each mutable variable in a second optimizing seed is compatible with M_i, so the degree formula shows that R_i annihilates it. On frozen variables both derivations have value $\delta_{ia}M_i$. They therefore agree on a transcendence basis. The kernel assertion follows from the degree formula and Cao's compatibility criterion. □

2.2. **Regularity and local nilpotence.**

Theorem 2.4. *The derivation R_i preserves $\mathcal{U}^+$ and is locally nilpotent. For a source-reachable vertex j, the derivation $R_j^\star = M_j^\star \partial_{x_j}$, where $M_j^\star = \prod_{u\in U} x_u^{b_{ju}}$ in a source optimizing seed, is likewise canonical and locally nilpotent on $\mathcal{U}^+$.*

Proof. Suppose first that the extended exchange matrix has full column rank. Work in an i-sink optimizing seed. The derivation preserves its Laurent ring with polynomial frozen variables. Consider the adjacent seed at $k \in U$. If $b_{ik} = 0$, it kills both x_k and x_k'. Otherwise the exchange polynomial has the form $P = Ax_i^q + C$, where $q > 0$ and A, C are independent of x_i. We have $R_i(x_k') = qAx_i^{q-1}\frac{M_i}{x_k}$. The exponent of x_k in M_i is $-b_{ik} \geqslant 1$. Substituting $x_k = P/x_k'$ therefore gives an element of the adjacent Laurent ring, still polynomial in the frozen variables. The same holds for $R_i(x_i) = M_i$; all other generators are killed. Thus R_i preserves the initial and adjacent seed rings.

Full rank gives total coprimeness, and the upper-bound theorem identifies their fully localized intersection with $\mathcal{U}^{\mathrm{loc}}$ [2, Corollary 1.9]. Moreover, invariance of frozen valuations gives, for any seed t,

$$\mathcal{U}^+ = \mathcal{U}^{\mathrm{loc}} \cap \Bbbk[x_u(t)^{\pm 1} : u \in U][x_f : f \in F].$$

The derivation preserves both rings on the right when t is the optimizing seed, and hence preserves $\mathcal{U}^+$.

For arbitrary rank, start at the optimizing seed and adjoin frozen variables using

$$\widehat{B} = \begin{pmatrix} B & I \\ -I & 0 \end{pmatrix}, \qquad \widehat{D} = \mathrm{diag}(D, D).$$

The extended mutable matrix has full column rank, so the lifted derivation preserves the lifted upper algebra. For any seed t, lift each coordinate $x_a(t)$ to the corresponding cluster or frozen variable $\widehat{x}_a(t)$. Its derivative belongs to the lifted upper algebra. Specializing the auxiliary variables to 1 shows that $R_i(x_a(t))$ belongs to the seed ring at t, with the original frozen variables polynomial. The derivation then preserves that entire ring, including inverses of its mutable coordinates. Intersecting over t proves the assertion. This argument requires specialization only for the images of seed coordinates.

Finally, every element of $\mathcal{U}^+$ is polynomial in x_i in the optimizing seed, and M_i is independent of x_i. Consequently a sufficiently high power of R_i kills each element. The source case follows by replacing B with $-B$. □

2.3. **Cartan and mixed degrees.** For sink-reachable vertices i, j, let $a_{ij} = \rho_i(M_j)$. The degree formula gives

$$d_i a_{ij} = d_j a_{ji} = \mathsf{F}_D(M_i, M_j).$$

Thus $c_{ii} = 2$ and $c_{ij} = -a_{ij}$ for $i \neq j$ define a symmetrizable generalized Cartan matrix, with symmetrizer $\operatorname{diag}(d_i)$ up to a common positive factor. We will also use a formula involving the frozen vectors ζ_i. Evaluate Lemma 2.1 in an i-sink optimizing seed. The degree of M_j is its ith pointed coordinate, and $b_{ij} = \zeta_{ij}$ there. Hence

$$a_{ij} = -b_{ji} + \zeta_{ji} = \zeta_{ji} + \frac{d_j}{d_i}\zeta_{ij} \qquad (i \neq j).$$

For a source-reachable frozen vertex j, define $\rho_j^\star(P) = \deg_{x_j} P$ in a source optimizing seed. This is the nilpotence degree for the canonical source derivation, not the minimum frozen valuation. The separation formula now has increasing x_j-exponents. Its constant term gives the minimum zero, and its maximum gives $\rho_j^\star(P) = d_j^{-1} F_P[D_U g(M_j^\star)_U^{\mathsf{T}}]$ in that seed. Therefore

$$d_j \rho_j^\star(P) = \mathsf{F}_D(M_j^\star, P), \qquad d_i \rho_i(M_j^\star) = d_j \rho_j^\star(M_i) = \mathsf{F}_D(M_i, M_j^\star).$$

In particular, the mixed degrees vanish precisely when M_i and $M_j^\star$ are compatible.

The boundary derivations depend on the valued frozen data, not only on the mutable principal matrix. Exchange entries between frozen vertices play no role in mutation or in these derivations. The symmetrizing weights at frozen vertices, however, affect the coefficient entries and hence can change M_i and $M_j^\star$.

3. Regular theta cones and tropical crystals

In [12, Section 5], the tropical crystal is carried by μ-supported presentation weights. Here we use the g-vectors of regular finite theta functions instead. Throughout this section, we assume reddening. Then every integral g-vector has a finite principal theta function. We construct the tropical operators on the g-vectors of regular theta functions. The canonical derivations enter only when we compare different optimizing seeds.

3.1. **Theta reciprocity and the regular cone.** Let M be the exponent lattice, with basis e_a, and let $N = M^*$. For theta reciprocity we use the *chiral dual* cluster datum with lattices

$$N^\bullet = \bigoplus_{a \in V} \mathbb{Z}\frac{d_a}{\ell} e_a^*, \qquad M^\bullet = \bigoplus_{a \in V} \mathbb{Z}\frac{\ell}{d_a} e_a.$$

In the seed-datum notation of [8], p_u denotes an exchange direction and q_u its coefficient direction. The original datum has $p_u = \widetilde{B} e_u$, $q_u = e_u^*$, $q_u^\bullet = (d_u/\ell) e_u^*$, and $p_u^\bullet = (\ell/d_u) p_u$. The dual exponents are $q_u^\bullet$ and its normals are $p_u^\bullet$. These normals lie in $M^\bullet$ because $-d_a b_{ua} = d_u b_{au}$. Our choice of ℓ ensures that every frozen valuation vector e_f^* belongs to $N^\bullet$. Write $\vartheta_n^\vee$ for a dual theta function. Theta reciprocity [8, Corollary 3.2, Claim 4.16, and Theorem 5.19] gives

$$\mathcal{K}(g, n) = \operatorname{val}_n(\vartheta_g) = \operatorname{val}_g(\vartheta_n^\vee)$$

whenever the values are finite. All valuations are minimum exponents; we use the same convention for theta functions of a seed in which some mutable vertices have been frozen.

Lemma 3.1. *Every integral point of $N^\bullet$ indexes a finite Laurent theta function on the localized chiral-dual cluster variety.*

Proof. We treat isolated zero exchange columns separately. Their rank-one factors are treated in Remark 1.2: the original factor is $2^{(-a)_+} x^a$, the dual factor is a Laurent monomial, and their pairing is linear.

Use $(d_a/\ell)e_a^*$ as the integral coordinate basis of $N^\bullet$. The dual exponent and normal matrices are $(E_U, B_{V,U})$, where E_U is the coordinate inclusion. The injective lattice map $\iota : z \longmapsto (B_{V,U}^{\mathsf{T}} z, z)$ defines a linear morphism to the datum

$$\widehat{P} = \begin{pmatrix} B_{UU}^{\mathsf{T}} \\ E_U \end{pmatrix}, \qquad \widehat{Q} = \begin{pmatrix} I_U \\ 0 \end{pmatrix},$$

since $\iota E_U = \widehat{P}$ and $\iota^{\mathsf{T}} \widehat{Q} = B_{V,U}$. This is a principal $\mathcal{A}$-datum with mutable matrix opposite to $\widetilde{B}_{U,U}$, together with frozen Laurent factors. Its weight matrix is ℓD_U^{-1}, since $D_U B_{UU}^{\mathsf{T}}$ is skew-symmetric. Reddening and the principal-coefficient case of Proposition 1.5 give a finite theta function for every integral point of the target lattice. Inverting the auxiliary frozen coordinates identifies the coefficient matrices $-I_U$ and I_U.

By the linear-morphism theorem and its positive-chamber formulation [8, Theorem 4.9 and Section 4.4], the image of $\vartheta_z^\vee$ is the target theta function indexed by ιz. This formulation also covers a source datum without a full-dimensional positive chamber. Injectivity of ι prevents distinct source exponents from merging, and positivity prevents cancellation. Finiteness of the target theta function therefore implies finiteness of $\vartheta_z^\vee$. □

Corollary 3.2. *In every seed, the g-vectors indexing regular theta functions are the lattice points of a rational polyhedral cone. More precisely, if* $\mathcal{S}_f^\vee = \operatorname{supp}(\vartheta^\vee_{e_f^*})$, *then*

$$\mathcal{B}_\theta^+ = M \cap \mathcal{C}^+, \qquad \mathcal{C}^+ = \bigcap_{f \in F} \bigcap_{a \in \mathcal{S}_f^\vee} \{g : \langle g, a \rangle \geqslant 0\}.$$

Proof. Since $e_f^* \in N^\bullet$, the preceding lemma makes $\mathcal{S}_f^\vee$ finite. Reciprocity identifies the frozen valuation of ϑ_g with $\min_{a \in \mathcal{S}_f^\vee} \langle g, a \rangle$. By the frozen-valuation characterization in the proof of Lemma 1.1, nonnegativity of these valuations for all $f \in F$ is equivalent to regularity in $\mathcal{U}^+$. The supports lie in the rational lattice $N^\bullet$, so the cone is rational polyhedral. □

3.2. **Boundary polytopes and closure.** For $f \in F$, write

$$\vartheta^\vee_{e_f^*} = \sum_{c \in A_f} a_c z^{e_f^* + c}, \qquad a_c > 0, \qquad A_f \subseteq \bigoplus_{u \in U} \frac{d_u}{\ell} \mathbb{Z}_{\geqslant 0} e_u^*.$$

The set A_f contains zero and has no frozen coordinates. We call $\mathcal{P}_f = \operatorname{conv}(A_f)$ the *boundary polytope* at f. For a mutable g-vector p, the least frozen exponent needed for regularity at f is

$$\beta_f(p) = -\operatorname{val}_{e_f^*}(\vartheta_{(p,0)}) = \max_{c \in \mathcal{P}_f}(-p \cdot c).$$

Lemma 3.3. *Every vertex of* $\mathcal{P}_f$ *belongs to* $\mathbb{Z}^U_{\geqslant 0}$. *Consequently the regular cone is defined by the integral inequalities*

$$g_f + g_U \cdot c \geqslant 0 \qquad (f \in F,\ c \in \operatorname{Vert}(\mathcal{P}_f)).$$

Proof. For every $w \in \mathbb{Z}^U$, reciprocity gives an integral support value

$$h_{\mathcal{P}_f}(w) = -\operatorname{val}_{e_f^*}(\vartheta_{(-w,0)}) \in \mathbb{Z}.$$

Choose an integral vector w which strictly exposes a given vertex c. For sufficiently large N, each $Nw + e_u$ still strictly exposes c. Thus $c_u = h_{\mathcal{P}_f}(Nw + e_u) - N h_{\mathcal{P}_f}(w) \in \mathbb{Z}$. Nonnegativity follows from the definition of A_f. It suffices to test a linear inequality on the vertices, giving the asserted description of the cone. □

Theorem 3.4. *In an i-sink optimizing seed, let* $m_U = -B_{i,U} \geqslant 0$, $m_F = 0$, *and* $v = m - e_i$. *For every* $g \in \mathcal{B}_\theta^+$,

$$g + sv \in \mathcal{B}_\theta^+ \qquad (0 \leqslant s \leqslant g_i,\ s \in \mathbb{Z}).$$

At $s = g_i$ *the ith coordinate is zero. Translation by v is injective where* $g_i > 0$, *with partial inverse given by translation by* $-v$ *whenever the resulting g-vector is regular.*

Proof. We use the inequalities defining the regular cone. Let $J = \{u \in U : m_u > 0\}$, and consider the chiral-dual seed obtained by freezing these mutable vertices. For any remaining mutable vertex u, we have $\langle e_i^*, p_u^\bullet \rangle = -(\ell/d_i)m_u = 0$. The same pairing vanishes on every normal created by scattering, because these normals are nonnegative integral combinations of the remaining initial normals. Thus e_i^* belongs to the lineality space of the frozen scattering diagram, where its theta function is the monomial $z^{e_i^*}$.

Freezing has a concrete effect on the original dual theta function. Before coefficient specialization, quotient the scattering Lie algebra by the ideal generated by degrees that are positive at one of the vertices in J. The resulting consistent scattering diagram is the one for the frozen seed; on broken lines the quotient removes exactly the terms whose bending degrees have a positive coordinate in J. Since the dual bend directions $(d_u/\ell)e_u^*$ are independent, specializing the auxiliary coefficients cannot identify a removed term with a surviving term. The frozen theta function at e_i^* is a monomial, so the only exponent $c \in A_i$ with $c_u = 0$ for every $u \in J$ is $c = 0$.

Consequently, $m \cdot c > 0$ for every nonzero exponent $c \in A_i$. By convexity the same is true at every nonzero vertex of $\mathcal{P}_i$. Lemma 3.3 proves that these vertices are integral; hence $m \cdot c \geqslant 1$. This integral lower bound is the point needed for a single raising step.

For each $f \in F$ and each vertex c of $\mathcal{P}_f$, regularity is expressed by $g_f + g_U \cdot c \geqslant 0$. Replacing g by $g + v$ changes the left side by

$$v_f + v_U \cdot c = m \cdot c - \delta_{if}.$$

When $f \neq i$ this change is nonnegative. When $f = i$ and $c \neq 0$, it is also nonnegative by the preceding bound. The remaining case is $f = i$, $c = 0$, whose inequality is $g_i \geqslant 0$. A raising step changes it to $g_i - 1 \geqslant 0$, valid exactly when $g_i > 0$.

Thus one raising step preserves all defining inequalities. Repeating the argument while the ith coordinate is positive proves regularity of $g + sv$ for $0 \leqslant s \leqslant g_i$. The final assertions follow because translation by a fixed vector is injective, with inverse translation defined precisely when the resulting g-vector belongs to the regular cone. □

Corollary 3.5. *The map $P_i(g) = g + g_i v$ is a linear retraction of $\mathcal{C}^+$ onto $\mathcal{C}^+ \cap \{g_i = 0\}$. For $g \in \mathcal{C}^+$, the segment $g + tv$, $0 \leqslant t \leqslant g_i$, lies in $\mathcal{C}^+$.*

Proof. All defining inequalities except $g_i \geqslant 0$ are nondecreasing in the direction v. The remaining coordinate is $g_i - t$. Thus $P_i(g)$ lies on the face $g_i = 0$, which P_i fixes pointwise. □

Proposition 3.6. *Let $\mathcal{N} = \{c + e_f^* : f \in F,\ c \in \mathrm{Vert}(\mathcal{P}_f)\}$ and $a_n = \langle v, n \rangle$. For a regular integral theta g-vector g, the number of successive inverse translations $g \mapsto g - v$ is*

$$\sup\{t \in \mathbb{Z}_{\geqslant 0} : g - tv \in \mathcal{C}^+\} = \min_{\substack{n \in \mathcal{N} \\ a_n > 0}} \left\lfloor \frac{\langle g, n \rangle}{a_n} \right\rfloor,$$

where the minimum of the empty set is $+\infty$.

Proof. At $g - tv$, the normal e_i^* imposes only $g_i + t \geqslant 0$. Normals with $a_n = 0$ give unchanged inequalities; every other normal has $a_n > 0$ and gives $t \leqslant \langle g, n \rangle / a_n$. Taking the integral points on this interval proves the formula. □

3.3. **Intrinsic boundary operators.** A *tropical point* consists of integer coordinates in every seed, related by tropical mutation. Write Φ_w for the piecewise integral-linear change of coordinates along a mutation word w. At a single mutation μ_k, the transformation on theta g-vectors is

$$g'_k = -g_k, \qquad g'_a = g_a + [b_{ka}]_+ g_k - b_{ka}[g_k]_+ \quad (a \neq k).$$

The matrix mutates simultaneously. Thus Φ_w is an invertible piecewise integral-linear map, and theta covariance identifies the regular domains in its source and target seeds. We write 0_{cr} when a crystal operator is undefined, whereas the lattice origin 0 indexes the unit.

There is a useful regularity test when source optimizing seeds are available.

Lemma 3.7. *If w_f reaches a source optimizing seed for f, then* $\mathrm{ord}_{x_f}(\vartheta_g) = (\Phi_{w_f} g)_f$. *If such a word exists for every $f \in F$, then*

$$\mathcal{B}_\theta^+(t) = \{g \in \mathbb{Z}^V : (\Phi_{w_f} g)_f \geqslant 0 \text{ for all } f \in F\}.$$

Proof. In the source seed, skew-symmetrizability gives $\widetilde{b}_{fu} = (d_u/d_f) b_{fu} \geqslant 0$. All exponents in the pointed expansion differ from $\Phi_{w_f} g$ by $-\nu B_{U,V}$, $\nu \geqslant 0$, and its leading term occurs. The least x_f-exponent is therefore $(\Phi_{w_f} g)_f$. □

For a sink-reachable vertex i, choose an optimizing word $\boldsymbol{\mu}_i^+$ and write $\Phi_i = \Phi_{\boldsymbol{\mu}_i^+}$. In the resulting seed let $M_i = x^{m_i}$, with $(m_i)_u = -b_{iu}$ and $(m_i)_f = 0$, and let $v_i = m_i - e_i$. For a regular theta g-vector g, write $a = \Phi_i g$ and define

$$\rho_i(g) = a_i, \qquad r_i(g) = \begin{cases} \Phi_i^{-1}(a + v_i), & a_i > 0, \\ 0_{\mathrm{cr}}, & a_i = 0. \end{cases}$$

We set $l_i(g) = \Phi_i^{-1}(a - v_i)$ when the resulting theta g-vector is regular, and $l_i(g) = 0_{\mathrm{cr}}$ otherwise.

Theorem 3.8. *The maps r_i, l_i and the function ρ_i are independent of the optimizing seed and its mutation word. They commute with seed changes, r_i is an injective partial map, and l_i is its partial inverse. Moreover,*

$$\rho_i(g) = \max\{n \geqslant 0 : r_i^n(g) \neq 0_{\mathrm{cr}}\}.$$

The function ρ_i decreases by one under r_i and increases by one under l_i, whenever the corresponding neighbor exists.

Proof. The zero vertex of $\mathcal{P}_i$ gives $a_i \geqslant 0$. Theorem 3.4 shows that $a + v_i$ is regular when $a_i > 0$. Translation is injective and decreases the ith coordinate by one, proving the string-length assertion in the chosen seed. If $a - v_i$ lies in the regular cone, its ith coordinate is positive and r_i sends it to a. Every preimage of a has this form. Thus the stated regularity test gives the partial inverse.

To compare optimizing seeds, let $n = a_i$ and $p = \Phi_i^{-1}(a + n v_i)$. Freeze $J = \{u : (m_i)_u > 0\}$ in the optimizing seed, and use a superscript 0 for the frozen theta functions. A bend of positive degree on J strictly lowers the x_i-exponent. Since ϑ_p is regular and its pointed ith coordinate is zero, no such bend occurs. Hence $\vartheta_p = \vartheta_p^0$. The vector v_i is supported on vertices now frozen, so frozen translation gives $\vartheta_p = x^{n v_i} \vartheta_g^0$. Freezing selects the highest x_i-degree part of ϑ_g, and therefore

$$\deg_{R_i} \vartheta_g = n, \qquad R_i^{(n)} \vartheta_g = \vartheta_p, \qquad R_i^{(n)} = R_i^n / n!.$$

Here R_i is canonical by Theorem 2.4. Thus n and p are independent of the optimizing seed; for p we also use injectivity of the theta parametrization after specialization [8].

On each stratum of constant ρ_i, maximal raising in color i is injective, since in one optimizing seed it is translation by $\rho_i v_i$. The two candidates for $r_i(g)$ both have $\rho_i = n - 1$, and maximal raising in color i takes both to p. Injectivity on this stratum shows that they coincide. This proves the assertions for r_i, and those for l_i follow from its characterization as the partial inverse. □

Corollary 3.9. *The maximal raising map is*

$$P_i(g) = r_i^{\rho_i(g)} g = \Phi_i^{-1}(a + a_i v_i), \qquad a = \Phi_i g.$$

It is an idempotent map onto $\rho_i^{-1}(0)$, and $R_i^{(\rho_i(g))} \vartheta_g = \vartheta_{P_i g}$.

Proof. Iterating in one optimizing seed gives the formula. The derivative identity is part of the preceding proof. □

3.4. **Compatible weights.** Let I be a family of sink-reachable frozen vertices, and let C be its Cartan matrix. In a reference seed set

$$q_i = g(M_i) - e_i, \qquad L_{\mathbb{Q}} = \mathbb{Q}^V / \operatorname{row}_{\mathbb{Q}}(B_{U,V}).$$

A *compatible grading* is a linear functional on $L_{\mathbb{Q}}$; equivalently, it vanishes on all mutable exchange rows. A family $(w_j)_{j\in I}$ is *adapted* if $w_j(q_i) = c_{ji}$. We define $\mathrm{wt}_j(g) = w_j(g)$ and $\lambda_j(g) = \rho_j(g) + \mathrm{wt}_j(g)$.

Corollary 3.10. *Adapted integral weights make $(r_i, l_i, \rho_i, \lambda_i, \mathrm{wt}_i)_{i\in I}$ an upper seminormal crystal with coroot-coordinate weights on $\mathcal{B}_\theta^+$. Adapted rational weights give the rational weight convention of* [12].

Proof. All exponents of a theta function differ by exchange vectors, so every compatible grading makes it homogeneous. The degree of R_i is $w_j(q_i) = c_{ji}$, hence $\mathrm{wt}_j(r_i g) - \mathrm{wt}_j(g) = c_{ji}$. The inverse-arrow and string-length axioms follow from Theorem 3.8. The definition of λ_i supplies the remaining weight identity and gives its increment $-1 + c_{ii} = 1$ under r_i. □

Lemma 3.11. *Let $Q : \mathbb{Q}^I \to L_{\mathbb{Q}}$ send e_i to $[q_i]$. Adapted rational weights exist if and only if* $\ker Q \subseteq \ker C$. *When they exist, they form an affine space over* $\mathrm{Hom}_{\mathbb{Q}}(L_{\mathbb{Q}}/\operatorname{im} Q, \mathbb{Q}^I)$.

Proof. An adapted grading is a map $W : L_{\mathbb{Q}} \to \mathbb{Q}^I$ with $WQ = C$. The kernel inclusion is exactly the condition that this prescription be well-defined on $\operatorname{im} Q$. Extend it linearly to $L_{\mathbb{Q}}$; two extensions differ by a map through the stated quotient. □

Proposition 3.12. *For any finite family I of sink-reachable boundaries,* $\ker Q$ *is totally isotropic for the symmetric form $D_I C$. Adapted rational weights therefore exist when $D_I C$ is positive semidefinite, in particular when every component of C has finite or affine type. In finite type the classes $[q_i]$ are linearly independent. Adapted rational weights also exist when $|I| \leqslant 2$. These assertions do not require reddening.*

Proof. By Section 2, $g(M_i) = -\operatorname{row}_i B + \zeta_i$, $\zeta_{ii} = 0$, and $a_{ij} = \zeta_{ji} + (d_j/d_i)\zeta_{ij}$ for $i \neq j$. If $t = (t_i) \in \ker Q$, choose v_u with $\sum_i t_i q_i = \sum_u v_u \operatorname{row}_u B$. Let h have mutable coordinates v_u, frozen coordinates t_i on I, and zero elsewhere. Then $hB = \sum_i t_i(\zeta_i - e_i)$. Skew-symmetry of BD gives

$$0 = hBDh^{\mathsf{T}} = \sum_{i,j\in I} d_j t_i t_j (\zeta_{ij} - \delta_{ij}) = -\frac{1}{2} t D_I C t^{\mathsf{T}}.$$

Polarization proves total isotropy. For a positive semidefinite form every isotropic vector lies in the radical, giving $\ker Q \subseteq \ker C$. Positive definiteness gives $\ker Q = 0$. We then apply Lemma 3.11.

For two connected vertices, write $C = \left(\begin{smallmatrix} 2 & -a \\ -b & 2 \end{smallmatrix}\right)$, where $a, b > 0$. The case $ab \leqslant 4$ is positive semidefinite. For $ab > 4$, a nonzero rational isotropic vector would imply that $ab(ab-4) = (ab-2)^2 - 4$ is an integer square. This number lies strictly between $(ab-3)^2$ and $(ab-2)^2$, a contradiction. The disconnected case is positive definite. □

The integral question reduces to a divisibility condition. Let $G = \mathbb{Z}^V / \operatorname{sat}_{\mathbb{Z}} \operatorname{row}_{\mathbb{Z}}(B_{U,V})$ and $L = \sum_i \mathbb{Z}[q_i] \subseteq G$. When rational adapted weights exist, the assignment $[q_i] \mapsto Ce_i$ defines a map $L \to \mathbb{Z}^I$. It extends to an integral grading precisely when its rational extension is integral on $\operatorname{sat}_G L$, a direct summand of G. Smith normal form therefore gives a divisibility criterion.

Without an exact source partner, the formal function $\lambda_i = \rho_i + \mathrm{wt}_i$ need not be nonnegative or equal the length of the inverse string. The actual inverse-string length is given by Proposition 3.6, which can be infinite.

3.5. **The opposite crystal.** Apply the construction to $-B$. Its sink-reachable vertices are the source-reachable vertices of B. Proposition 1.5 identifies the regular theta sets, and Proposition 1.7 relates their g-vectors. This gives the starred operators $r_j^\star$ and their string functions $\rho_j^\star$ on $\mathcal{B}_\theta^+$. Its raising length is the pointed coordinate in an opposite optimizing seed. This is generally different from the minimum frozen valuation computed in Lemma 3.7. For exact boundary pairs, Section 7.3 identifies these two families as raising and lowering operators of a seminormal theta crystal.

4. Tropical DT theory

The exact boundary pairs defined below are the seed-theoretic analogues of the τ-exact pairs of [12, Section 4]. We express their existence through the mutable Donaldson–Thomas transformation and principal-coefficient F-polynomials.

Fix a reference seed. Recall that BD is skew-symmetric, and for $j \in F$ let $a_j = B_{j,U}$. Here $a_j = B_{j,U}$ is the restriction of the frozen exchange row at j to the mutable indices. At mutation μ_k it changes by

$$a'_k = -a_k, \qquad a'_u = a_u + [b_{ku}]_+[a_k]_+ - [-b_{ku}]_+[-a_k]_+ \quad (u \neq k).$$

This is the tropical mutation rule for projective weights, the negatives of mutable theta g-vectors. It applies to every frozen index j, regardless of reachability.

4.1. **Tropical F-polynomials and frozen coordinates.** Let P be a mutable cluster monomial with mutable pointed degree $-b$. We call b a *reachable projective weight* and denote its principal-coefficient F-polynomial by F_b, with exponent convention $p^*(\nu) = -\nu B_{U,V}$. Multiplication by a frozen Laurent monomial does not change F_b. The polynomial F_b has a componentwise largest exponent f_b; its constant and largest coefficients are both 1. This fact holds without reddening. Indeed, P is a cluster monomial for both B and $-B$, and their separation formulas in a principal extension have the form

$$P = x^g \sum_{\nu \geqslant 0} c_\nu x^{-\nu B} = x^\eta \sum_{\mu \geqslant 0} d_\mu x^{\mu B}, \qquad c_0 = d_0 = 1.$$

The injective principal exchange map gives $\eta = g - f_b B$ and matches $\mu = f_b - \nu \geqslant 0$, exactly as in Proposition 1.7.

For a projective weight b with finite F_b, and any $a \in \mathbb{Z}^U$, define

$$\mathsf{E}(a,b) = \max_{\nu \in \operatorname{supp} F_b} (-aD_U\nu^{\mathsf{T}}), \qquad \mathsf{H}(a,b) = aD_U f_b^{\mathsf{T}} + \mathsf{E}(a,b).$$

The formula for $\mathsf{E}(a,b)$ is defined whenever F_b is finite; the formula for $\mathsf{H}(a,b)$ also requires its componentwise largest exponent. For reachable projective weights, both expressions are nonnegative: $0 \in \operatorname{supp} F_b$ and $f_b \in \operatorname{supp} F_b$. The quantity $\mathsf{E}(a,b)$ evaluates the support function of $-\operatorname{Newt}(F_b)$ at aD_U. Likewise, $\mathsf{H}(a,b)$ is the support function of $f_b - \operatorname{Newt}(F_b)$. Neither expression is generally symmetric in a and b. We write $\operatorname{codeg} P = g(P) - f_b B_{U,V}$ for the opposite pointed degree. Under reddening, Proposition 1.7 provides these expressions for every integral mutable weight.

Proposition 4.1. *For a mutable cluster monomial P of mutable pointed degree $-b$ and every $j \in F$,*

$$g(P)_j = \frac{\mathsf{E}(a_j,b)}{d_j}, \qquad \operatorname{codeg}(P)_j = \frac{\mathsf{H}(a_j,b)}{d_j}.$$

In particular, both quotients are integers.

Proof. The frozen valuation of P at j is zero, since it is zero in a seed where P is a mutable monomial. In the separation formula, the term indexed by ν has exponent

$$g(P)_j - \sum_{u \in U} \nu_u b_{uj} = g(P)_j + d_j^{-1} a_j D_U \nu^{\mathsf{T}}.$$

Its minimum is zero, giving the first identity. Positivity makes the calculation valid even when coefficient specialization identifies formal exponents. Subtracting $f_b B_{U,V}$ from the pointed degree gives the second identity. □

Let i be sink-reachable and k source-reachable. Define $b_k = -g(M_k^\star)_U$ and $c_i = \operatorname{codeg}(M_i)_U$. The boundary monomial formula and its opposite give $g(M_i)_U = -a_i$ and $\operatorname{codeg}(M_k^\star)_U = a_k$ in every seed. Hence

$$c_i = -a_i - f_{a_i} B_{UU}, \qquad a_k = -b_k - f_{b_k} B_{UU}.$$

We call (i,k) an *exact boundary pair* if $\operatorname{codeg}(M_i) + g(M_k^\star) = e_i + e_k$.

Proposition 4.2. *The pair (i,k) is exact if and only if $c_i = b_k$ and*

$$\mathsf{H}(a_j, a_i) + \mathsf{E}(a_j, b_k) = d_j(\delta_{ij} + \delta_{kj}) \qquad (j \in F).$$

The criterion is invariant under mutation and applies in arbitrary extended rank, without a reddening hypothesis.

Proof. The negative of an ordinary degree and an opposite degree obey the same tropical mutation rule. Translation by a frozen vector commutes with this rule, because it changes no mutable coordinate. Thus the defining exactness equation is mutation-invariant. In an i-sink optimizing seed, $M_i = x^{m_i}$ has codegree m_i, so the equation is $g(M_k^\star) = e_i + e_k - m_i$. Its mutable coordinates give $c_i = b_k$, while Proposition 4.1 gives the stated frozen equations. □

4.2. **The mutable DT transformation.** Assume now that the mutable pattern admits reddening. Let $C_\theta(b) = -b - f_b B_{UU}$ and $T = C_\theta^{-1}$. By Proposition 1.7, passing from a theta function's g-vector to its codegree is a bijection. Hence C_θ is a bijection on mutable projective weights. In particular,

$$T^{-1}(b) = -b - f_b B_{UU}, \qquad T(a) = -a \quad (a \geqslant 0).$$

The second formula follows from $F_{-a} = 1$. In pointed coordinates the localized DT twist acts by $g \mapsto -\operatorname{codeg}(\vartheta_g)$; in projective coordinates $b = -g$, this is T^{-1}. We call T the *mutable DT transformation* in the inverse orientation. The principal extension gives this convention in arbitrary rank.

The boundary identities above become $b_k = Ta_k$ and $c_i = T^{-1}a_i$. Exactness therefore requires the mutable matching $T^{-1}a_i = Ta_k$, together with all frozen equations in Proposition 4.2. We next express those equations using H alone.

Theorem 4.3. *Under reddening, these support functions satisfy*

$$\mathsf{E}(a,b) = \mathsf{H}(b, Ta) \qquad (a, b \in \mathbb{Z}^U).$$

Proof. We work in a principal coefficient extension, where the exchange map has full column rank and the auxiliary coefficients are invertible. In this proof, V and B refer to the enlarged vertex set and exchange matrix. The advantage of the extension is that every coefficient exponent is recorded independently; the calculation at the end involves only principal-coefficient F-polynomials of the mutable pattern and hence descends to the original coefficients.

Fix $a, b \in \mathbb{Z}^U$, and write $c = Ta$. Choose finite theta functions L, P in the principal extension whose full pointed degrees are $-\widetilde{b}$ and $-\widetilde{c}$, with mutable negative degrees b, c, respectively. Let $\widetilde{a} = \operatorname{codeg} P$. By the definition of the DT transformation, its mutable part is a, and the degree–codegree identity gives $\widetilde{a} = -\widetilde{c} - f_c B_{U,V}$.

The principal extension has an integral skew matrix Γ with the compatibility relations

$$B_{U,V}\Gamma = -(D_U \mid 0), \qquad \Gamma B_{U,V}^{\mathsf{T}} = (D_U \mid 0)^{\mathsf{T}}.$$

One may use the explicit construction in the proof of Proposition 9.3. Multiplying the preceding codegree identity by Γ gives

$$\widetilde{b}\Gamma\widetilde{a}^{\mathsf{T}} = \widetilde{c}\Gamma\widetilde{b}^{\mathsf{T}} - bD_U f_c^{\mathsf{T}}.$$

This accounts for the linear correction that distinguishes H from E.

We next apply theta reciprocity in the principal extension. To check its hypotheses, let $P_0 = -B_{U,V}^{\mathsf{T}}$ be the seed exchange map and let $Q_0 : \mathbb{Z}^U \hookrightarrow \mathbb{Z}^V$ be the coordinate inclusion. Set $Q_0^\bullet = Q_0 D_U$, $D_0 = D_U^{-1}$, and $\Lambda = -\Gamma$. The compatibility relations above become

$$\Lambda P_0 = Q_0^\bullet, \qquad \Lambda^{\mathsf{T}} P_0 = -Q_0^\bullet,$$

so Λ is a rational Λ-structure in the sense of [8, Section 5.1].

Apply the same-datum reciprocity statement [8, Claim 5.7 and Theorem 5.19] to the tropical directions $m_1 = -\widetilde{a}^{\mathsf{T}}$ and $m_2 = -\widetilde{b}^{\mathsf{T}}$. It identifies the valuation of L at $\Gamma\widetilde{a}^{\mathsf{T}}$ with the valuation of the theta function indexed by $-\widetilde{a}$ in the negative scattering chamber at $-\Gamma\widetilde{b}^{\mathsf{T}}$. The two chambers here belong to the same scattering diagram. Reflection through the origin takes broken lines in the negative chamber to broken lines in the positive chamber for the opposite seed. By the codegree formula of Proposition 1.7, the negative-chamber expression is $\nu(P)$, where $\nu(x^m) = x^{-m}$. Thus reciprocity yields

$$\mathrm{val}_{\Gamma\widetilde{a}^{\mathsf{T}}}(L) = \mathrm{val}_{\Gamma\widetilde{b}^{\mathsf{T}}}(P).$$

All theta functions involved are finite. The comparison of broken-line valuations with ordinary theta valuations at these integral tropical directions follows from [8, Corollary 3.2].

The pointed expansions of L and P now express the two minimum valuations as

$$-\widetilde{b}\Gamma\widetilde{a}^{\mathsf{T}} - \mathsf{E}(a,b) \quad \text{and} \quad -\widetilde{c}\Gamma\widetilde{b}^{\mathsf{T}} - \mathsf{E}(b,c).$$

Their equality, together with the linear correction calculated above, gives

$$\mathsf{E}(a,b) = bD_U f_c^{\mathsf{T}} + \mathsf{E}(b,c) = \mathsf{H}(b,c).$$

Finally, $c = Ta$ and the right side depends only on principal-coefficient F-polynomials of the mutable pattern. We can therefore specialize the auxiliary coefficients to 1 without losing the identity, proving the assertion for the original exchange datum. □

4.3. **The boundary criterion.**

Corollary 4.4. *Under reddening, (i,k) is an exact boundary pair if and only if $T^{-1}a_i = Ta_k$ and*

$$\mathsf{H}(a_j, a_i) + \mathsf{H}(T^{-1}a_i, Ta_j) = d_j(\delta_{ij} + \delta_{kj}) \qquad (j \in F).$$

Proof. Apply Theorem 4.3 to $\mathsf{E}(a_j, Ta_k)$ in Proposition 4.2, and use $Ta_k = T^{-1}a_i$. □

For $j \notin \{i,k\}$, both terms are nonnegative and must vanish separately. Hence the frozen coordinates impose conditions that are not detected by the equality of mutable weights. The categorical interpretation is the weighted counterpart of [12, Lemma 4.29].

4.4. **Adding frozen vertices.**

Fix the mutable pattern, D_U, and a pair (i,k) satisfying $c_i = b_k$. Write $F_i = F_{a_i}$, $G_k = F_{b_k}$, and $f_i = f_{a_i}$. Define the *admissible coefficient cone* by

$$\mathcal{C}_{i,k} = \{a \in \mathbb{R}^U : aD_U(f_i - \nu)^{\mathsf{T}} \leqslant 0 \ (\nu \in \operatorname{supp} F_i), \quad aD_U\mu^{\mathsf{T}} \geqslant 0 \ (\mu \in \operatorname{supp} G_k)\}.$$

It is a closed rational polyhedral cone determined by the two principal-coefficient F-polynomials.

Proposition 4.5. *Suppose (i,k) is exact. Add frozen vertices with compatible integral exchange entries and positive symmetrizing weights, leaving the existing entries unchanged. The pair remains exact if and only if every new coefficient vector a_j belongs to $\mathcal{C}_{i,k}$. For a collection of exact pairs, the corresponding condition is membership in the intersection of their admissible coefficient cones. Reddening is not required.*

Proof. The optimizing words still work after the extension: mutation of a coefficient vector depends only on that vector and the mutable principal matrix. The mutable projective weights of the two boundaries, their principal-coefficient F-polynomials, and the existing frozen-coordinate equations are unchanged. For a new frozen vertex with vector a, the remaining equation is $\mathsf{H}(a, a_i) + \mathsf{E}(a, b_k) = 0$. Both terms are nonnegative, and

$$\mathsf{H}(a, a_i) = \max_{\nu \in \operatorname{supp} F_i} aD_U(f_i - \nu)^{\mathsf{T}}, \qquad \mathsf{E}(a, b_k) = \max_{\mu \in \operatorname{supp} G_k} (-aD_U\mu^{\mathsf{T}}).$$

Both vanish precisely for $a \in \mathcal{C}_{i,k}$. Proposition 4.2 proves the claim for one pair and then for any prescribed collection. □

Every integral $a \in \mathcal{C}_{i,k}$ gives an integral extension: choose $d_j = 1$, set $B_{j,U} = a$ and $B_{U,j} = -D_U a^{\mathsf{T}}$, and set the new entries between frozen vertices to zero. This preserves skew-symmetry of BD with integral positive symmetrizing weights.

4.5. **Categorical comparison.** We first assume the skew-symmetric quiver-with-potential hypotheses of [12], including its reference seed without arrows between frozen vertices and its full mutation data. The underlying mutation and presentation theories are [9, 10]; their relation to tropical F-polynomials is developed in [11]. Let i be reachable and j dually reachable. We use the sign convention $\delta = -g$ for projective presentation weights and write $\epsilon_j^\star$ for the weight of $E_j^\star$.

Lemma 4.6. *In every seed, $\epsilon_j^\star = e_j - g(M_j^\star)$.*

Proof. In a source optimizing seed this is the boundary weight formula of [12, Section 4]. Its two sides transform by the same projective-weight mutation rule, giving the formula in every seed. □

Proposition 4.7. *Under these categorical hypotheses, (i, j) is an exact boundary pair if and only if $\tau^{-1}E_i \cong E_j^\star$.*

Proof. In an i-sink optimizing seed, [12, Section 4] gives $\delta(\tau^{-1}E_i) = m_i - e_i$. The preceding lemma gives $\delta(E_j^\star) = e_j - g(M_j^\star)$. Equality of these weights is exactly $g(M_j^\star) = e_i + e_j - m_i$. Since both objects are rigid, equality of presentation weights is equivalent to isomorphism [9]; compare [12, Lemma 4.29]. □

For skew-symmetrizable matrices, good Jacobi-finite nondegenerate ramified quivers with potentials provide a related categorical model [15]. When the first decorated representation mutates to pure decoration, [15, Proposition B.1] expresses its Hom and E-invariants by tropical F-polynomials. The generic formulas of [15, Remark B.2] may require passing to a positive multiple of the presentation weight. These results do not by themselves identify the full and mutable boundary objects required for the exact-pair comparison above.

4.6. **Supported presentation weights and regular theta functions.** We compare regular theta functions with the supported generic characters of [12]. Assume a skew-symmetric ice quiver with potential $(\Delta, \mathcal{S})$ as in [12, Construction 2.5], with finite-dimensional full Jacobian algebra and reddening for the mutable pattern. Work over $\mathbb{C}$ and take $D = I$. Let $\mathcal{T} = -\operatorname{trop}(\Delta, \mathcal{S})$ be the negatives of the μ-supported presentation weights, and let V_{gen} be the span of their generic characters. We use the projected dominance order $\lhd$ of Definition 11.1.

Qin's unitriangular comparison of pointed bases [34, Theorems 5.1.3 and 5.2.2] concerns the localized algebra. Here the frozen variables are not inverted: their vanishing orders determine which theta functions are regular. The basis theorem of [15, Theorem 8.4] concerns inverted frozen variables and does not replace this calculation.

Lemma 4.8. *If $g = -\delta \in \mathcal{T}$, then ϑ_g is regular and*

$$C_{\mathrm{gen}}(\delta) = \vartheta_g + \sum_{\substack{h \in \mathcal{B}_\theta^+ \\ h \lhd g}} c_{g,h}\vartheta_h$$

is a finite unitriangular expansion. Moreover, $\eta_\theta(-\delta) = -\check{\delta}$.

Proof. Adjoin frozen sources $z_u \to u$ for all mutable u, as in [12, proof of Proposition 8.1]. This adds an identity block to the exchange matrix. The radical of the projective at z_u is the old projective at u, so the framed Jacobian algebra remains finite dimensional. A presentation of weight $(\delta, 0)$ has the same cokernel extended by zero, hence the same quotient Grassmannians. Its generic character is regular and pointed at $(g, 0)$ by [12, Lemma 7.6].

Qin's full-rank comparison now applies. Its theta terms are regular by Proposition 1.3. Specializing the auxiliary sources to 1 preserves regularity; by Lemma 11.2, the projected g-vectors of all nonleading terms remain strictly below g. This proves the expansion. Qin's support-dimension formula [34, Propositions 3.4.8 and 5.2.4] identifies the largest mutable F-exponent with $(\dim \operatorname{coker} \delta)|_U$. Combining Proposition 1.7 with [12, (1.6)] gives the opposite-weight identity. □

Fix a mutable presentation weight b. For each frozen vertex j, let $\mathcal{E}_j^\mu$ be its extracted mutable decorated object and write $q_j(b) = \mathrm{e}_{J_\mu}(\mathcal{E}_j^\mu, b)$, with the second presentation general of weight b. Its extension by zero has supported weight $\delta^0 = (b, -q(b))$. Recall $a_j = B_{j,U}$.

Proposition 4.9. *For every frozen vertex j,*

$$\operatorname{ord}_{x_j=0}(\vartheta_{-\delta^0}) = q_j(b) - \mathsf{E}(a_j, b) \geqslant 0.$$

The order vanishes when j admits a sink or source optimizer, and at every frozen vertex when b is reachable. For fixed mutable part $-b$, the allowable frozen coordinates are $q(b) + \mathbb{N}^F$ for $\mathcal{T}$ and $\beta_B(-b) + \mathbb{N}^F$ for $\mathcal{B}_\theta^+$, where $\beta_B(-b)_j = \mathsf{E}(a_j, b)$. Consequently $V_{\mathrm{gen}} = \mathrm{Mid}^+$ if and only if $q_j(b) = \mathsf{E}(a_j, b)$ for all b, j. In particular, the supported generic characters form a basis of Mid^+ when every frozen vertex has an optimizer.

Proof. Supported weights above b are obtained from δ^0 by negative frozen decorations [12, Definition 6.7 and Lemma 6.9]; their frozen coordinates form $q(b) + \mathbb{N}^F$. By Proposition 5.1, the corresponding regular theta coordinates are $\beta_B(-b) + \mathbb{N}^F$. Substitute $-\delta^0 = (-b, q(b))$ into (5.1). The resulting order is nonnegative by Lemma 4.8.

If j has an optimizer, Lemma 1.6 allows us to choose a source optimizer. Mutation of the mutable object and its decoration introduces no frozen representation spaces, and its presentation weight follows the mutation rule [12, Theorem 2.2(4)]. At the chosen source there is no incoming space at j. Thus both projective Betti numbers at j vanish by [12, (1.3)–(1.5)], so the jth g-coordinate is zero. Lemma 3.7 identifies this coordinate with the order along $x_j = 0$. If b is reachable, choose a seed where its general object is purely negative. The extension by zero remains purely negative, and $\vartheta_{-\delta^0}$ becomes a mutable cluster monomial; all its frozen orders therefore vanish.

Finally, the lemma's expansions are unitriangular for a strict order with finite lower sets (Lemma 11.2). Hence the generic characters and the ϑ_h with $h \notin \mathcal{T}$ together form a basis of Mid^+. Their spans agree precisely when the two sets of frozen coordinates agree over each $-b$, which is the stated equality criterion. □

Corollary 4.10. *For every μ-supported weight δ, $g = -\delta$ intertwines the raising operators of [12] with the theta raising operators and preserves upper string lengths, for reachable and dually reachable colors. If $\mathcal{T} = \mathcal{B}_\theta^+$, the regular theta basis of Mid^+ is BK-biperfect for the two compatible gradings. If it spans $\mathcal{U}^+$, the theta assertion of [12, Conjecture 8.14] follows. Both conditions hold if the supported generic characters span $\mathcal{U}^+$.*

Proof. In an optimizing seed, the raising derivation is $M_i\partial_{x_i}$, the string length is $-\delta_i$, and the raising step changes δ by $e_i - m_i$; see [12, Lemmas 5.3–5.4 and 6.2]. The formulas agree with the theta operators under $g = -\delta$ and commute with mutation. For a starred color use the opposite-weight identity in Lemma 4.8. Normalized perfectness follows from Theorem 7.2; the ordinary and opposite expansions give the two gradings. If $V_{\mathrm{gen}} = \mathcal{U}^+$, then $V_{\mathrm{gen}} \subseteq \mathrm{Mid}^+ \subseteq \mathcal{U}^+$ forces equality, and the proposition gives $\mathcal{T} = \mathcal{B}^+_\theta$. □

5. Primitive theta lifts and changes of coefficients

The mutation behavior of the boundary weights parallels the formulas in [12, Sections 4 and 6]. We then study the effect of changing the polynomial frozen coefficients.

We first work on the *finite theta locus* $\Theta_\mu \subseteq \mathbb{Z}^U$, consisting of those mutable g-vectors p for which $\vartheta_{(p,0)}$ has a finite principal-coefficient lift. Multiplication by a frozen monomial does not affect this finiteness property. Reddening will be needed only for the polyhedral description and the comparison with opposite theta functions.

5.1. **Primitive regular theta lifts.** In a fixed seed, write the g-vector of a theta function as $(p, q) \in \mathbb{Z}^U \oplus \mathbb{Z}^F$. Frozen factorization gives $\vartheta_{(p,q)} = x_F^q \vartheta_{(p,0)}$. Let ν_f be the valuation along $x_f = 0$. As in Section 2, mutation preserves this valuation, so an element of $\mathcal{U}^{\mathrm{loc}}$ belongs to $\mathcal{U}^+$ precisely when all its frozen valuations are nonnegative.

Recall that $D = \mathrm{diag}(d_v)$ and $a_f = B_{f,U}$. The principal-coefficient theta F-polynomial associated with $(p, 0)$ is F_{-p}. An exponent ν of this polynomial contributes $a_f D_U \nu^{\mathsf{T}}/d_f$ to the x_f-exponent. Taking the minimum gives

$$\nu_f(\vartheta_{(p,q)}) = q_f - \frac{\mathsf{E}(a_f, -p)}{d_f}. \tag{5.1}$$

Positivity ensures that this minimum is unchanged by specialization of the auxiliary coefficients. Define

$$\beta_B(p)_f = \frac{\mathsf{E}(a_f, -p)}{d_f}, \qquad \chi^B_p = x_F^{\beta_B(p)} \vartheta_{(p,0)}.$$

Equation (5.1) shows that $\beta_B(p)$ is integral, and the constant term of F_{-p} shows that it is nonnegative. We call χ^B_p the *primitive regular theta function* of mutable g-vector p; the adjective refers to its vanishing orders along the original frozen divisors. Among the theta functions with this mutable g-vector, it is the unique one whose order along every frozen divisor is zero.

Proposition 5.1. *The g-vectors of regular theta functions have the unique form $(p, \beta_B(p)+a)$, where $p \in \Theta_\mu$ and $a \in \mathbb{N}^F$. Consequently*

$$\mathrm{Mid}^+(B) = \bigoplus_{p \in \Theta_\mu} \Bbbk[x_F]\chi^B_p.$$

The same decomposition holds for $\mathcal{U}^+(B)$ if theta functions span that algebra. Moreover, for $p, p' \in \Theta_\mu$ and $n \in \mathbb{N}$,

$$\beta_B(p + p') \leqslant \beta_B(p) + \beta_B(p'), \qquad \beta_B(np) = n\beta_B(p).$$

In the g-vector lattice of each seed, the regular finite theta functions are indexed by a saturated additive monoid. For reachable sink and source boundaries i, j, respectively,

$$\deg_{R_i}(x_F^a \chi^B_p) = \deg_{R_i}(\chi^B_p) + a_i, \qquad \deg_{R^\star_j}(x_F^a \chi^B_p) = \deg_{R^\star_j}(\chi^B_p) + a_j.$$

When the theta basis is perfect for these derivations, these degrees are its string lengths.

Proof. The valuation formula gives the stated parametrization of regular theta g-vectors. Frozen factorization then proves spanning over $\Bbbk[x_F]$. A polynomial relation among primitive lifts would be a linear relation among distinct theta functions, so the sum is direct.

The finite theta locus is additive and saturated [21, Theorem 0.3(4)]. The product of two principal theta functions has a positive theta expansion containing the term indexed by the

sum of their g-vectors [8, Example 2.6]. Taking frozen valuations after specialization therefore gives

$$\nu_f(\vartheta_{(p+p',0)}) \geqslant \nu_f(\vartheta_{(p,0)}) + \nu_f(\vartheta_{(p',0)}).$$

This is the subadditivity assertion. The homogeneity of theta valuations [8, Lemma 2.9] gives $\nu_f(\vartheta_{(np,0)}) = n\nu_f(\vartheta_{(p,0)})$, and hence homogeneity of β_B. These two properties, together with saturation of Θ_μ, prove the monoid assertion. Finally, nilpotence degree in an optimizing seed is ordinary degree in its boundary variable. Multiplication by x_F^a increases that degree by the corresponding coordinate of a. □

For a regular theta function with g-vector (p, q), the nonnegative vector $a = q - \beta_B(p)$ consists of its orders of vanishing along the original frozen divisors. We call a its *frozen excess*. Thus this vector has a geometric meaning, independent of the chosen expression as a frozen monomial times χ_p^B.

Corollary 5.2. *Assume reddening. For the boundary polytope $\mathcal{P}_f$ of Section 3.2,*

$$\beta_B(p)_f = \max_{c\in\mathcal{P}_f}(-\langle p, c\rangle) = \beta_f(p).$$

Thus $\beta_B(p)_f$ is a convex integral piecewise-linear function, and the regular theta g-vectors form a saturated, finitely generated additive monoid.

Proof. The identity follows by comparing (5.1) with the definition of β_f. Lemma 3.3 gives its integral linear pieces. Gordan's lemma applied to the cone of Corollary 3.2 proves the monoid assertion. □

The monoid law here is addition of g-vectors in a fixed seed; it need not agree with multiplication of individual theta functions.

5.2. **Seed identifications.** A *direct seed identification* is a field isomorphism Φ that identifies two seeds by a permutation π of their indices, preserving mutable and frozen vertices, exchange data, and symmetrizing weights. An *opposite seed identification* identifies the exchange data with their negatives. Mutation followed by such a relabeling gives examples. Only exchange entries with at least one mutable index are relevant.

Proposition 5.3. *A direct seed identification satisfies*

$$\Phi R_i \Phi^{-1} = R'_{\pi(i)}, \qquad \Phi R_j^\star \Phi^{-1} = R_{\pi(j)}^{\star\prime}.$$

An opposite seed identification interchanges the positive and negative derivations:

$$\Phi R_i \Phi^{-1} = R_{\pi(i)}^{\star\prime}, \qquad \Phi R_j^\star \Phi^{-1} = R'_{\pi(j)}.$$

Both identifications preserve the upper algebras with polynomial frozen variables. A direct identification also preserves the regular middle algebras and their finite theta bases. The latter assertion holds for an opposite identification whenever the opposite theta bases agree, in particular under reddening.

On stable perfect theta bases, the induced maps intertwine the leading operators, their partial inverses, and their string lengths. Compatible gradings are transported with them. Direct identifications preserve exact pairs; opposite identifications send (i, j) to $(\pi(j), \pi(i))$.

Proof. An optimizing seed maps to an optimizer of the same kind in the direct case and of the other kind in the opposite case. Its canonical monomial maps to the corresponding target monomial. The formulas for the derivations follow in that chart and hence in the rational function field.

The map identifies the corresponding Laurent rings and permutes the polynomial frozen variables. Their intersections give the upper algebras. Mutation covariance gives the direct theta assertion; Proposition 1.7 gives the opposite assertion. Frozen valuations select the regular subbases. Conjugation preserves nilpotence degrees and the unique leading term of a perfectness expansion, so it intertwines the crystal maps. It also preserves the finite normalization identities for exact pairs and their commutators. □

Under reddening, let η_θ be characterized by $\vartheta_g^B = \vartheta_{\eta_\theta(g)}^{-B}$. If an opposite identification returns to the same valued ice pattern and has tropical map σ, its permutation of the original theta g-vectors is $\kappa_\sigma = \eta_\theta^{-1}\sigma$. Its order depends on the chosen identification.

5.3. **Relative changes of coefficients.** Let $\gamma = (\boldsymbol{\mu}, \pi)$ consist of a mutation sequence and a mutable permutation returning the principal matrix to itself. Extend π by the identity on F and write $B^\gamma = \pi\boldsymbol{\mu}(B)$. We regard this as a new seed with independent initial coordinates y. Its principal matrix agrees with that of B, whereas its coefficient entries may differ. The inverse of mutation pullback is denoted by $\mathcal{T}_\gamma$; thus $\mathcal{T}_\gamma(x_f) = y_f$ for $f \in F$. We identify the two polynomial coefficient rings by these equalities.

Let γ_g be the tropical change of g-vectors and γ_μ its mutable part. For a g-vector (p, q), the mutable coordinates of $\gamma_g(p, q)$ are $\gamma_\mu(p)$, which depends only on p and the principal pattern.

Proposition 5.4. *The relative change of coefficients gives isomorphisms*

$$\mathcal{T}_\gamma : \mathcal{U}^+(B) \longrightarrow \mathcal{U}^+(B^\gamma), \qquad \mathcal{T}_\gamma : \mathrm{Mid}^+(B) \longrightarrow \mathrm{Mid}^+(B^\gamma).$$

They intertwine the canonical boundary derivations with frozen indices fixed. On primitive lifts and frozen excess they satisfy

$$\mathcal{T}_\gamma(\chi_p^B) = \chi_{\gamma_\mu(p)}^{B^\gamma}, \qquad (p, a) \longmapsto (\gamma_\mu(p), a).$$

Equivalently,

$$\gamma_g(p, \beta_B(p) + a) = \big(\gamma_\mu(p), \beta_{B^\gamma}(\gamma_\mu(p)) + a\big).$$

These isomorphisms transport exact-pair actions and, whenever the theta bases are stable and perfect, their crystals and compatible gradings. No reddening assumption is needed.

Proof. Apply Proposition 5.3 to the original pattern with its initial seed changed to the terminal seed of γ. The image $\mathcal{T}_\gamma(\chi_p^B)$ has mutable g-vector $\gamma_\mu(p)$ and order zero along every frozen divisor. Uniqueness of the primitive lift identifies it with $\chi_{\gamma_\mu(p)}^{B^\gamma}$. Multiplication by x_F^a gives the remaining formulas.

For completeness, a compatible grading changes under mutation at k by

$$w_k' = -w_k + \sum_v [\widetilde{b}_{vk}]_+ w_v, \qquad w_v' = w_v \quad (v \neq k).$$

Compatibility makes the two exchange monomials homogeneous of the same weight. The formula therefore preserves integral gradings and the values of these gradings on tropical points. □

Relative changes compose as cluster transformations between the corresponding coefficient patterns. The terminal matrix alone need not determine the resulting field map.

A reddening sequence with its standard mutable relabeling returns the principal matrix and hence gives such a transformation; see [27, Section 2.12]. With the projective convention for T used in Section 4, its mutable map is $\gamma_\mu(p) = -T(-p)$. Thus it acts on primitive coordinates by $(p, a) \mapsto (-T(-p), a)$, and the inverse uses T^{-1}.

The localized theta DT transformation is instead $\mathrm{DT}_\theta(g) = -\eta_\theta(g)$. By Proposition 1.7, its mutable part is $p \mapsto -T^{-1}(-p) = \gamma_\mu^{-1}(p)$. Its realization by a DT-type variation inverts the original frozen variables. In contrast, $\mathcal{T}_\gamma$ fixes the polynomial coefficient ring and changes the coefficient pattern.

More generally, an invertible integral variation followed by mutation permutes the corresponding finite theta family [27, Theorem 6.2]. It identifies the regular subfamilies precisely when the transformation and its inverse preserve the nonnegative frozen valuations. If it also conjugates the boundary derivations to the prescribed target derivations, Proposition 5.3 gives crystal intertwining. We discuss opposite symmetries for exact boundary pairs in Section 7.3.

6. Exact boundary pairs and Lie actions

This section gives the theta counterpart of the Chevalley-relation construction in [12, Section 8]; exact boundary pairs replace τ-exact boundary representations.

We begin with the commutator formulas for canonical boundary derivations. The exact-pair condition then gives the normalization required by the Chevalley relations.

6.1. **Hamiltonian identities and Serre relations.** Suppose first that B is invertible. The skew-symmetric matrix $\Lambda = B^{-1}D$ defines a log-canonical Poisson bracket by $\{x^p, x^q\} = p\Lambda q^{\mathsf{T}} x^{p+q}$. Since $B\Lambda = D$, this bracket is compatible with mutation. For a source boundary j, let $\zeta_j^\star = -B_{j,F}(t_j^*)$, extended by zero on U, and $Z_j^- = x^{\zeta_j^\star}$. The sink corrections ζ_i and Z_i^+ are as in Section 2.

Lemma 6.1. *The boundary derivations satisfy*

$$R_i = -\frac{Z_i^+}{d_i x_i}\{M_i/Z_i^+, -\}, \qquad R_j^\star = \frac{Z_j^-}{d_j x_j}\{M_j^\star/Z_j^-, -\}.$$

If $a_{ij} = \deg_{R_i} M_j$ for distinct sink boundaries, then

$$d_i x_i R_i M_j + d_j x_j R_j M_i = d_i a_{ij} M_i M_j. \tag{6.1}$$

For a sink boundary i and a source boundary j,

$$d_i x_i R_i M_j^\star = d_j x_j R_j^\star M_i. \tag{6.2}$$

The last two identities hold for arbitrary exchange rank.

Proof. In an i-sink optimizing seed, $M_i/Z_i^+ = x^{-B_{i,V}}$, so $\{M_i/Z_i^+, x_a\} = -\delta_{ia} d_i (M_i/Z_i^+) x_a$. This proves the first formula on the seed coordinates. The source formula follows in the same way with the opposite sign.

Let K_i be the compatible Euler derivation with weight vector $\zeta_i\Lambda$. Expanding the first formula gives $d_i x_i R_i f = M_i K_i f - \{M_i, f\}$. The degree formula of Section 2 and the homogeneity of M_j give

$$a_{ij} = \zeta_{ji} + \frac{d_j}{d_i}\zeta_{ij}, \qquad K_i M_j = \left(d_j \zeta_{ij} + \zeta_i \Lambda \zeta_j^{\mathsf{T}}\right) M_j.$$

Adding the formulas for (i, M_j) and (j, M_i) cancels both skew-symmetric terms and proves (6.1).

The source formula has the form $d_j x_j R_j^\star f = \{M_j^\star, f\} - M_j^\star K_j^\star f$, where $K_j^\star$ has weights $\zeta_j^\star \Lambda$. Combining it with the sink formula gives

$$d_i x_i R_i M_j^\star - d_j x_j R_j^\star M_i = \left(d_i \zeta_{ji}^\star - d_j \zeta_{ij}\right) M_i M_j^\star.$$

In a j-source optimizing seed, the pointed expansion of M_i has only nonnegative corrections to its x_j-exponent. Its frozen valuation is zero, so its leading x_j-exponent is zero. The formula $g(M_i) = -B_{i,V} + \zeta_i$ then gives $\zeta_{ij} = b_{ij}(t_j^*)$. Skew-symmetrizability implies $d_j \zeta_{ij} = d_i \zeta_{ji}^\star$, proving (6.2).

For singular B, apply the identities after the invertible coefficient extension used in Section 2 and specialize its auxiliary frozen variables to one. □

Proposition 6.2. *For distinct reachable sink boundaries, the Serre relation $(\operatorname{ad} R_i)^{a_{ij}+1} R_j = 0$ is sharp. The matrix with diagonal entries 2 and off-diagonal entries $-a_{ij}$ is a generalized Cartan matrix, symmetrized by $\operatorname{diag}(d_i)$. The same statements hold for the source derivations.*

For a sink boundary i and a source boundary j, let $b = \deg_{R_i} M_j^\star$. If $i \neq j$, then $[R_i, R_j^\star] = 0$ for $b = 0$, and the sharp vanishing exponent of $(\operatorname{ad} R_i)^n R_j^\star$ is $b + 2$ for $b > 0$. If $i = j$, the two derivations commute for $b \leqslant 1$, and the sharp exponent is $b + 1$ for $b \geqslant 2$.

Proof. Work in an i-sink optimizing seed. The multiplier M_i has x_i-degree zero. In (6.1), the terms of highest x_i-degree cancel, so $\deg_{x_i} R_j M_i \leqslant a_{ij} - 1$. The Hamiltonian formula also gives $\deg_{x_i} R_j x_k \leqslant a_{ij}$ for $k \neq i$, while $R_j x_i = 0$. Since $R_i^2 x_k = 0$ for every seed coordinate,

$$(\operatorname{ad} R_i)^n R_j(x_k) = R_i^n R_j(x_k) - n R_i^{n-1} R_j R_i(x_k).$$

These degree bounds prove vanishing at $n = a_{ij}+1$. Evaluation at x_j gives $(\operatorname{ad} R_i)^{a_{ij}} R_j(x_j) = R_i^{a_{ij}} M_j \neq 0$. The degree formula gives $d_i a_{ij} = d_j a_{ji}$, which proves the Cartan assertions. Changing B to $-B$ gives the source assertions.

For $i \neq j$, the mixed identity shows that $R_j^\star M_i = 0$ when $b = 0$, and otherwise $\deg_{x_i} R_j^\star M_i = b$. The source Hamiltonian formula gives $\deg_{x_i} R_j^\star x_k \leqslant b$ for $k \neq i$. The same commutator calculation proves the mixed vanishing bounds. When $b > 0$, sharpness follows from

$$(\operatorname{ad} R_i)^{b+1} R_j^\star(x_i) = -(b+1) R_i^b R_j^\star M_i \neq 0.$$

If $i = j$, equation (6.2) gives $R_i M_i^\star = R_i^\star M_i$, and the source Hamiltonian formula gives $\deg_{x_i} R_i^\star x_k \leqslant b - 1$ for $k \neq i$. Moreover, $(\operatorname{ad} R_i)^n R_i^\star(x_i) = (1-n) R_i^n M_i^\star$. This proves commutativity for $b = 0, 1$ and sharpness of the exponent $b+1$ for $b \geqslant 2$. □

By the Serre presentation of $\mathfrak{n}_+(C)$ [22], the assignment $e_i \mapsto R_i$ defines a Lie algebra homomorphism $\mathfrak{n}_+(C) \to \operatorname{Der}_{\Bbbk}(\mathcal{U}^+)$. Equivalently, $\mathcal{U}^+$ is a $U(\mathfrak{n}_+(C))$-module algebra, with locally nilpotent simple-root operators. This extends the action of [12, Theorem 8.3] to skew-symmetrizable exchange data. The source derivations define the analogous action for their Cartan matrix. Neither action requires reddening, theta spanning, or exact partners.

For adapted compatible gradings w_j, let H_j be the corresponding Euler derivations. They commute and satisfy $[H_j, R_i] = c_{ji} R_i$, extending the action to the positive Borel subalgebra of the derived Kac–Moody algebra. Exact pairs will supply the negative simple-root operators and the canonical Cartan gradings.

Proposition 6.3. *If every boundary in I is both sink-reachable and source-reachable, the Cartan matrices of $(R_i)_{i\in I}$ and $(R_i^\star)_{i\in I}$ agree.*

Proof. The proof of Lemma 6.1 gives $d_i \zeta_{ji}^\star = d_j \zeta_{ij}$. For $i \neq j$, the source degree formula therefore reads

$$a_{ij}^\star = \zeta_{ji}^\star + \frac{d_j}{d_i}\zeta_{ij}^\star = \frac{d_j}{d_i}\zeta_{ij} + \zeta_{ji} = a_{ij}.$$

Both Cartan matrices have diagonal entries two. □

6.2. **Exact pairs and finite normalization.** Choose an i-sink optimizing seed. Write $M_i = x^{m_i}$ and $v_i = m_i - e_i$. For a reachable source boundary j, its canonical source monomial is $M_j^\star = R_j^\star(x_j)$.

Definition 6.4. The ordered pair (i, j) is *exact* if

$$g(M_j^\star; t_i) = e_i + e_j - m_i. \tag{6.3}$$

Here $g(M_j^\star; t_i)$ is the full g-vector in the column convention $\widetilde{B} = -B_{U,V}^{\mathsf{T}}$.

Proposition 6.5. *Exactness is equivalent to the corresponding identity*

$$R_i M_j^\star = x_j \quad (i \neq j), \qquad R_i^2 M_i^\star = 2M_i \quad (i = j).$$

Proof. Let $U_i = \{k \in U : b_{ik} < 0\}$ in the chosen optimizer. The principal-coefficient expansion of $M_j^\star$ is

$$\widehat{M_j^\star} = x^g \sum_{\nu \in \mathbb{N}^U} c_\nu x^{-\nu B} y^\nu, \qquad c_\nu \in \mathbb{Z}_{\geqslant 0}, \quad c_0 = 1.$$

A term has x_i-exponent $g_i - \sum_k \nu_k b_{ki} \leqslant g_i$, with equality precisely when $\nu_k = 0$ for $k \in U_i$. Qin's freezing theorem [35, Theorems 5.2 and 5.11] identifies their sum with the theta function indexed by g in the seed obtained by freezing U_i. Denote its specialization by q_g. The expansion is finite, and positivity gives $\deg_{x_i} M_j^\star = g_i$ after specialization.

The vector v_i is supported on frozen indices of this new seed. Frozen factorization gives $x^{rv_i}q_g = q_{g+rv_i}$. If $n = g_i$, we therefore have $R_i^n M_j^\star = n! q_{g+nv_i}$. Under (6.3), $n = 1 + \delta_{ij}$, and $g + nv_i$ equals e_j when $i \neq j$ and m_i when $i = j$. The corresponding functions are x_j and M_i, proving the normalization.

Conversely, either normalization implies $g_i = 1 + \delta_{ij}$. The pointed term x^g contributes $n! x^{g+nv_i}$ to the maximal derivative. All contributions have nonnegative coefficients, while the normalized derivative is a single monomial. Its exponent must therefore be $g + nv_i$. Solving for g proves (6.3). This argument applies after any specialization of auxiliary coefficients, even when the exponent map is not injective. □

The normalized identities are independent of the optimizing seed. Thus exactness is a property of the original boundary derivations, not of a chosen seed.

6.3. **The rank-one action.** Let i be a reachable sink boundary and j a reachable source boundary. Abbreviate $R = R_i$, $R^\star = R_j^\star$, $M = M_i$, and $N = M_j^\star$. In an i-sink optimizing seed, $R = M\partial_{x_i}$ and $M = x^m$. The boundary construction gives $RM = R^\star N = 0$ and

$$R(x_f) = \delta_{if} M, \qquad R^\star(x_f) = \delta_{jf} N \qquad (f \in F).$$

We use the mixed identity

$$d_i x_i RN = d_j x_j R^\star M. \tag{6.4}$$

Theorem 6.6. *If (i, j) is exact, then $H = [R_i, R_j^\star]$ is an integral compatible grading and*

$$[H, R_i] = 2R_i, \qquad [H, R_j^\star] = -2R_j^\star.$$

These derivations define an integrable $\mathfrak{sl}_2$ action on $\mathcal{U}^+$. Distinct exact partners satisfy $d_i = d_j$ and $R_j^\star M_i = x_i$.

Conversely, suppose the canonical boundary derivations form a Chevalley triple with a compatible grading satisfying $H(x_i) = (-1 + \delta_{ij})x_i$ and $H(x_j) = (1 - \delta_{ij})x_j$. Then (i, j) is exact.

Proof. Assume first that B is invertible and use $\Lambda = B^{-1}D$. The source Hamiltonian formula supplies a frozen Laurent monomial $G = x^g$ with

$$R^\star = \frac{1}{d_j x_j}\left(\{N, -\} - N\frac{\{G, -\}}{G}\right). \tag{6.5}$$

On a seed coordinate it reads

$$\frac{R^\star(x_k)}{x_k} = \frac{1}{d_j x_j}\left(\frac{\{N, x_k\}}{x_k} - N(g\Lambda)_k\right). \tag{6.6}$$

Suppose $i \neq j$. The normalization $RN = x_j$ gives $N = x_i x_j/M + C$, with C independent of x_i. Write $a = d_i/d_j$ and $q = e_i + e_j - m - g$. Differentiating (6.6) gives $[R, R^\star](x_k) = d_j^{-1}(q\Lambda)_k x_k$ for $k \neq i$. Equation (6.4) gives $R^\star M = ax_i$, so $[R, R^\star](x_i) = -ax_i$. Comparing the x_i-coefficients in the known values of $R^\star(x_f)$, $f \in F$, shows that $d_j^{-1}(q\Lambda)_i = 0$, $d_j^{-1}(q\Lambda)_j = 1$, and all remaining frozen entries vanish. Thus H is diagonal with weights $w = d_j^{-1}q\Lambda - ae_i$.

This grading is compatible. Indeed, $Bw^{\mathsf{T}} = -d_j^{-1}Dq^{\mathsf{T}} - aBe_i^{\mathsf{T}}$. For $k \in U$, we have $q_k = -m_k$ and $m_k = -b_{ik} = (d_i/d_k)b_{ki}$, so the kth entry vanishes. The identities $RM = R^\star N = 0$, $RN = x_j$, and $R^\star M = ax_i$ imply $HM = aM$ and $HN = -N$. Hence $[H, R] = 2aR$. Every diagonal derivation preserves a log-canonical bracket and commutes with the Euler derivation $\{G, -\}/G$. Applying H to (6.5) therefore gives $[H, R^\star] = -2R^\star$. Jacobi now yields $0 = [H, [R, R^\star]] = 2(a - 1)H$. Since $H(x_j) = x_j$, we obtain $a = 1$, proving the distinct case.

For a self pair, write $x = x_i$ and $d = d_i$. The normalization $R^2N = 2M$ gives $N = x^2/M + Ax + C$, where A, C are independent of x. Let $q = 2e_i - m - g$. For $k \in U$, (6.6) takes the form

$$\frac{R^\star(x_k)}{x_k} = \frac{1}{d}\left(\frac{x}{M}(q\Lambda)_k + A_k + \frac{C_k}{x}\right),$$

with A_k, C_k independent of x. Polynomiality in the original frozen variable x forces $C_k = 0$. It follows that $[R, R^\star](x_k) = d^{-1}(q\Lambda)_k x_k$. The bracket kills every frozen variable, including x, because $RN = R^\star M$. Comparing the coefficients of x^2/M in $R^\star(x_f)$ gives $d^{-1}(q\Lambda)_i = 1$ and zero at the other frozen entries. Thus H is diagonal with weights $w = d^{-1}q\Lambda - e_i$. The preceding matrix calculation again proves compatibility.

We have $HM = R(R^\star M) = R^2 N = 2M$. Every cluster variable is homogeneous for a compatible grading, so the source monomial N is homogeneous. Its term x^2/M has weight -2, and therefore $HN = -2N$. The formula for R and the source Hamiltonian formula now give $[H, R] = 2R$ and $[H, R^\star] = -2R^\star$.

For general B, use the invertible coefficient extension

$$\widehat{B} = \begin{pmatrix} B & I \\ -I & 0 \end{pmatrix}, \qquad \widehat{D} = \operatorname{diag}(D, D).$$

All added indices are frozen, and their variables z are inverted. The original frozen variables remain polynomial. The lifted operators kill z and specialize to $R, R^\star$ at $z = 1$. In the lifted i-optimizer, Laurent positivity gives nonnegative integral coefficients for $\widehat{N}$. Thus every nonzero term of $\widehat{R}\widehat{N}$ has a positive integral coefficient, and every nonzero term of $\widehat{R}^2\widehat{N}$ has a positive even coefficient. Their specializations are x_j and $2M$, respectively. Consequently, for some integral γ,

$$\widehat{R}\widehat{N} = x_j z^\gamma \quad (i \neq j), \qquad \widehat{R}^2\widehat{N} = 2\widehat{M} z^\gamma \quad (i = j).$$

Replace $\widehat{R^\star}, \widehat{N}$ by $R^{\star\prime} = z^{-\gamma}\widehat{R^\star}$ and $N' = z^{-\gamma}\widehat{N}$. The operators kill $z^{-\gamma}$, so local nilpotence, regularity, and the mixed identity are preserved. The normalization is now exact. With $G' = z^{-\gamma}\widehat{G}$, the product rule gives

$$R^{\star\prime} = \frac{1}{d_j x_j}\left(\{N', -\} - N'\frac{\{G', -\}}{G'}\right).$$

Thus the preceding calculation applies. Its use of homogeneity is valid because N' is a cluster monomial times an auxiliary frozen Laurent monomial. The resulting bracket identities specialize to the original rational function field.

The diagonal bracket in the extension kills each auxiliary variable. Its compatible grading therefore descends to a compatible grading for B. To prove integrality, let f be an eigenvector of H. PBW expresses its cyclic $\mathfrak{sl}_2$ module as the span of $(R^\star)^b R^a f$. Local nilpotence bounds a, and then bounds b for each remaining a, so this module is finite dimensional. Its H-weights are integral. Applying this to the seed coordinates proves that the compatible grading is integral. Applying it to all homogeneous elements proves integrability on $\mathcal{U}^+$.

For the converse with $i \neq j$, evaluate $[R, R^\star] = H$ on x_j: this gives $RN = x_j$. For a self pair, $RN = R^\star M$ and $H(x_i) = 0$ give $HM = 2M$, whence $R^2 N = R(R^\star M) = [R, R^\star]M = 2M$. The preceding proposition proves exactness. □

In particular, the difference of the two nilpotence degrees is already determined by the bracket: if $Hf = mf$ and $f \neq 0$, then $\deg_{R^\star} f - \deg_R f = m$. This follows from the rank-one integrable module structure and does not require a separate degree-balance hypothesis.

Proposition 6.7. *Consider the* Langlands dual exchange datum $B^\vee = -B^{\mathsf{T}} = D^{-1}BD$, *with symmetrizer* $D^\vee = \ell D^{-1}$, *and use the same mutable and frozen indices. The sink and source optimizing sequences are the same for the two exchange data. A pair of frozen vertices is exact for* B *if and only if it is exact for* $B^\vee$. *On any common set of sink-reachable vertices, the boundary Cartan matrices satisfy* $C^\vee = C^{\mathsf{T}}$.

Proof. Mutation commutes with $B \mapsto D^{-1}BD$, and positive diagonal rescaling preserves the signs of the frozen entries. In the column convention, $\widetilde{B}^\vee = D\widetilde{B}D_U^{-1}$. Consequently tropical mutation intertwines the rational linear map $g \mapsto gD$: this follows directly from its coordinate formula and $\min(d_k g_k, 0) = d_k \min(g_k, 0)$.

The canonical sink exponent at i satisfies $m_i^\vee = d_i^{-1} m_i D$. Starting in a source optimizer for j and transporting its canonical monomial back to any reference seed therefore gives

$$g(M_j^{\star,\vee}) = d_j^{-1} g(M_j^\star) D.$$

If (i,j) is exact, Theorem 6.6 gives $d_i = d_j$. Thus its defining identity becomes

$$g(M_j^{\star,\vee}) = d_j^{-1}(e_i + e_j - m_i)D = e_i + e_j - m_i^\vee.$$

This proves exactness for the dual datum. Apply the same argument again to obtain the converse; a common scalar in the symmetrizer is irrelevant.

The identical tropical calculation for M_j gives $\deg_{R_i^\vee} M_j^\vee = (d_i/d_j)\deg_{R_i} M_j$. The symmetrizing identity $d_i c_{ij} = d_j c_{ji}$ then gives $c_{ij}^\vee = c_{ji}$, including the diagonal entries. □

6.4. **The derived Kac–Moody action.** For a single exact pair one can determine the mixed commutators with all other reachable boundary derivations.

Lemma 6.8. *Let (i,k) be an exact sink–source pair. For every sink-reachable h and every source-reachable j, respectively,*

$$\deg_{R_h} M_k^\star = \delta_{hi} + \delta_{hk}, \qquad \deg_{R_j^\star} M_i = \delta_{ji} + \delta_{jk}.$$

Consequently $[R_h, R_k^\star] = 0$ for $h \neq i$, and $[R_i, R_j^\star] = 0$ for $j \neq k$.

Proof. In an i-sink optimizer, put $N = M_k^\star$. Exactness gives $g(N) = e_i + e_k - m_i$. Since $g(N)_U = -m_i \leqslant 0$, the first summand of Cao's invariant $\mathsf{F}_D(M_h, N)$ vanishes, since the constant term of F_{M_h} is one and every exponent is nonnegative. For the second summand, use $g(M_h)_U = -B_{h,U}$. The term of N indexed by $\nu \in \operatorname{supp} F_N$ has x_h-exponent

$$g_h(N) - d_h^{-1} \nu D_U (g(M_h)_U)^{\mathsf{T}}.$$

These exponents have minimum zero, since N is a mutable cluster monomial and its Laurent expansion is positive. The degree formula of Proposition 2.2 now gives

$$\deg_{R_h} N = d_h^{-1} \mathsf{F}_D(M_h, N) = g_h(N) = \delta_{hi} + \delta_{hk}.$$

If $h \neq i$, this degree is zero when $h \neq k$ and is one when $h = k$. The two cases of Proposition 6.2 both give $[R_h, R_k^\star] = 0$. For the second degree formula and commutator apply the same argument to (k,i) in the opposite seed. Indeed, replacing B by $-B$ exchanges the two summands in the exactness equation $\operatorname{codeg}(M_i) + g(M_k^\star) = e_i + e_k$. □

Let I be a set of sink-reachable boundaries and assign an exact source partner $\sigma(i)$ to each $i \in I$. Put $L_i = R_{\sigma(i)}^\star$ and $H_i = [R_i, L_i]$. The sink Cartan matrix has entries $c_{ii} = 2$ and $c_{ij} = -\deg_{R_i} M_j$ for $i \neq j$.

Theorem 6.9. *The map σ is injective, and*

$$e_i \longmapsto R_i, \qquad f_i \longmapsto L_i, \qquad h_i \longmapsto H_i$$

defines an integrable action of the derived Kac–Moody algebra $\mathfrak{g}'(C)$ on $\mathcal{U}^+$ by derivations. Each H_i is an integral compatible grading, and $d_i = d_{\sigma(i)}$. The source Cartan matrix, reindexed by σ, equals C.

Proof. Theorem 6.6 gives the rank-one relations, the compatible integral gradings, and the equality $d_i = d_{\sigma(i)}$. Lemma 6.8 gives $[R_j, L_i] = 0$ for $j \neq i$. If $\sigma(i) = \sigma(j)$ with $i \neq j$, then $H_i = [R_i, L_j] = 0$, contrary to $[H_i, R_i] = 2R_i \neq 0$.

The H_i commute because they are diagonal on each seed. In their optimizing charts, R_j and L_j are monomial partial derivatives. They are therefore eigenvectors of $\operatorname{ad} H_i$, of weights u_{ij} and v_{ij} respectively. Applying $\operatorname{ad} H_i$ to $H_j = [R_j, L_j]$ gives $(u_{ij} + v_{ij})H_j = 0$, so $v_{ij} = -u_{ij}$.

Fix $i \neq j$. In the adjoint action of the ith $\mathfrak{sl}_2$, the vector R_j is killed by $\operatorname{ad} L_i$. Hence

$$(\operatorname{ad} L_i)(\operatorname{ad} R_i)^n R_j = -n(u_{ij} + n - 1)(\operatorname{ad} R_i)^{n-1} R_j.$$

By the sharp sink Serre relation, the first vanishing power is $n = 1 - c_{ij}$. Substitution gives $u_{ij} = c_{ij}$ and therefore $v_{ij} = -c_{ij}$. Dually, L_j is killed by $\operatorname{ad} R_i$ and has weight $-c_{ij}$. At the first vanishing power of $\operatorname{ad} L_i$, the dual rank-one identity gives that power as $1 - c_{ij}$. Such a power exists by the source Serre relation of Proposition 6.2. This identifies the source Cartan matrix and proves all Chevalley–Serre relations. The simple-root derivations are locally nilpotent, and compatibility of the integral gradings gives a simultaneous weight decomposition of $\mathcal{U}^+$. The action is therefore integrable. □

7. Perfect theta bases and crystal structures

The generic-character analogues of perfectness and normality are established in [12, Sections 8.3–8.4]. For theta functions the principal additional issue is to identify a stable algebra with the required basis.

7.1. **The highest frozen face.** We now identify the leading term of a boundary derivative. Tang [36, Proposition 3.2] proves the following statement when the extended exchange matrix has full rank and every theta function is finite. We use specialized valuative independence to allow arbitrary rank and restrict the finiteness assumption to the theta functions in the chosen basis.

Theorem 7.1. *Let A be a subalgebra of the seed function field with a basis $\{\vartheta_g : g \in \mathcal{B}\}$ of theta functions having finite principal-coefficient lifts. Fix $a \in F$ with $\widetilde{b}_{ak} \leqslant 0$ for all $k \in U$, and let $U_a = \{k \in U : \widetilde{b}_{ak} < 0\}$. Suppose $u \in \mathbb{Z}^V$ satisfies $u_a = 0$ and $\operatorname{supp}(u) \cap U \subseteq U_a$. If $\partial = x^u \partial_{x_a}$ preserves A and is locally nilpotent, then its theta basis is normalized perfect.*

More precisely, its nilpotence degree is $\rho_\partial(\vartheta_g) = g_a$, and

$$\ker(\partial^{r+1}|_A) = \operatorname{span}_{\Bbbk}\{\vartheta_g : g \in \mathcal{B},\ g_a \leqslant r\} \qquad (r \geqslant 0).$$

Writing $v = u - e_a$, for $g_a > 0$ we have $g + v \in \mathcal{B}$ and

$$\partial \vartheta_g = g_a \vartheta_{g+v} + \sum_{h_a \leqslant g_a - 2} c_h \vartheta_h.$$

The leading map $g \mapsto g + v$ is injective, and

$$\partial^{g_a} \vartheta_g = g_a! \vartheta_{g + g_a v}, \qquad \partial^{g_a + 1} \vartheta_g = 0.$$

Proof. Since x^u is independent of x_a, local nilpotence implies that each element of A is polynomial in x_a over $\Bbbk(x_v : v \neq a)$. Indeed, a negative lowest power x_a^r would survive every derivative, with nonzero coefficient proportional to $r(r-1)\cdots(r-m+1)$ after m steps. Thus nilpotence degree equals x_a-degree.

For a principal lift, write $\widehat{\vartheta}_g = x^g \sum_{n \in \mathbb{N}^U} c_{g,n} x^{-nB_{U,V}} y^n$. The terms of highest x_a-degree are exactly those with $n_k = 0$ for $k \in U_a$. Qin's freezing theorem [35, Theorem 5.2] identifies their sum with the theta function indexed by g in the seed obtained by freezing U_a. Denote its specialization by q_g. Positivity gives $\deg_{x_a} \vartheta_g = g_a$ and identifies q_g with its highest homogeneous part.

The functions q_g are linearly independent by specialized valuative independence [8, Theorem 3.10]; isolated zero exchange columns are handled as in Remark 1.2. Therefore the terms of maximal g_a in a finite theta sum cannot cancel. This proves the kernel formula.

The vector v is supported on frozen indices of the further frozen seed. Frozen factorization gives $x^v q_g = q_{g+v}$, a finite theta function of the frozen pattern. The highest part of $\partial \vartheta_g$ is consequently $g_a q_{g+v}$. Expand the derivative in the basis of A. The kernel formula excludes terms with $h_a > g_a - 1$, and comparison of highest parts gives $\sum_{h_a = g_a - 1} c_h q_h = g_a q_{g+v}$. Linear independence forces the unique leading term claimed in the formula. Translation by a fixed vector is injective. Iterating the formula proves the maximal derivative identity, since every remainder is killed by the remaining derivatives. □

7.2. **Stability of the regular middle algebra.** Assume reddening in this subsection. By Section 1, every principal theta function is finite, and specialized valuative independence gives

$$\mathrm{Mid}^{\mathrm{loc}} = \bigoplus_{g} \Bbbk\vartheta_g, \qquad \mathrm{Mid}^{+} = \mathrm{Mid}^{\mathrm{loc}} \cap \mathcal{U}^{+} = \bigoplus_{\vartheta_g \in \mathcal{U}^{+}} \Bbbk\vartheta_g.$$

We prove stability by lifting an individual theta function and its derivative to a coefficient extension with full rank.

Theorem 7.2. *Every canonical reachable sink or source derivation preserves* Mid^{+} *and is locally nilpotent there. The regular theta basis is normalized perfect for every canonical boundary derivation under consideration. The action of* $U(\mathfrak{n}_{+}(C))$ *therefore restricts to* Mid^{+}*. For an exact system of pairs, the derived Kac–Moody action on* $\mathcal{U}^{+}$ *restricts to an integrable action on* Mid^{+}*.*

Proof. We use the coefficient extension of full rank

$$\widehat{B} = \begin{pmatrix} B & I \\ -I & 0 \end{pmatrix}, \qquad \widehat{D} = \mathrm{diag}(D, D),$$

with all new indices frozen. We invert their variables z and leave the original frozen variables polynomial. The original principal pattern and optimizing sequences are unchanged. The lifted boundary derivations kill z and specialize to the original derivations at $z = 1$.

Let ϑ_g be a regular theta function and take its finite lift $\widehat{\vartheta}_{(g,0)}$ to this extension. Coefficient functoriality gives $\widehat{\vartheta}_{(g,0)}|_{z=1} = \vartheta_g$. In every seed, positivity shows that specialization preserves the minimum exponent of each original frozen variable. Thus the lift is regular along all original frozen divisors.

By Proposition 1.3 and Lemma 1.1, the lifted upper algebra has a theta basis with this regularity condition. Its boundary derivations preserve the algebra, so $\widehat{R}_i\widehat{\vartheta}_{(g,0)} = \sum_{\widehat{h}} c_{\widehat{h}}\widehat{\vartheta}_{\widehat{h}}$ is a finite expansion in regular theta functions. After $z = 1$, each summand is a regular theta function for the original pattern; repeated g-vectors can be combined. This proves $R_i\,\mathrm{Mid}^{+} \subseteq \mathrm{Mid}^{+}$. The same argument proves source stability.

Theorem 7.1 now gives sink perfectness. For source perfectness, apply it to the opposite seed and use the equality of opposite theta bases from Section 1. Local nilpotence is inherited from $\mathcal{U}^{+}$. For exact pairs, the Chevalley–Serre identities of Theorem 6.9 restrict to the stable algebra. Theta functions are homogeneous for the compatible Cartan gradings, so their integral weight decomposition restricts as well. □

If localized theta functions span $\mathcal{U}^{\mathrm{loc}}$, the regular subbasis spans $\mathcal{U}^{+}$ by Lemma 1.1. This applies, for example, under the full-rank or strict-cone hypotheses of Section 1. The following alternative does not require reddening.

Proposition 7.3. *If* $\mathcal{A}^{\mathrm{loc}} = \mathcal{U}^{\mathrm{loc}}$*, the finite principal-lift theta functions form a basis of* $\mathcal{U}^{\mathrm{loc}}$*, and their regular subfamily forms a basis of* $\mathcal{U}^{+}$*. Every reachable sink derivation is normalized perfect for this basis.*

Proof. The middle algebra contains all cluster monomials [21, Theorem 0.3]. Thus $\mathcal{A}^{\mathrm{loc}} \subseteq \mathrm{Mid}^{\mathrm{loc}} \subseteq \mathcal{U}^{\mathrm{loc}}$, and the assumed equality gives spanning. Specialized valuative independence gives linear independence [8, Theorem 3.10]. Lemma 1.1 selects the regular subbasis. Theorems 2.4 and 7.1 give its sink perfectness. □

In particular, local acyclicity gives this conclusion by [31, Theorem 2].

Question. Does a reddening sequence imply that the regular finite theta functions form a basis of $\mathcal{U}^{+}$ for every geometric coefficient system with polynomial frozen variables?

The next argument gives perfectness for exact source partners on the same basis, without comparing opposite theta functions.

7.3. **Normal theta crystals.** Let $A \subseteq \mathcal{U}^+$ have a basis of finite theta functions, and suppose it is stable under the exact-pair operators $E_i = R_i$ and $F_i = R^\star_{\sigma(i)}$. The compatible gradings $H_i = [E_i, F_i]$ make the theta basis homogeneous: exponents in a pointed expansion differ by mutable exchange rows, which these gradings annihilate. For a nonzero vector a, write $\rho_i(a) = \max\{n : E_i^n a \neq 0\}$ and $\lambda_i(a) = \max\{n : F_i^n a \neq 0\}$.

Lemma 7.4. *Let (E, F, H) act integrably on a vector space over a field of characteristic zero. An H-homogeneous normalized E-perfect basis is also normalized F-perfect. If $Hb = mb$, then $\lambda(b) - \rho(b) = m$, and the two leading maps are partial inverses.*

Proof. The finite-dimensional assertion is the rank-one perfectness lemma of [13, Lemma 5.13]. We explain why the same argument applies to an arbitrary integrable module and how the coefficients are determined.

By complete reducibility for integrable $\mathfrak{sl}_2$-modules, the module is a direct sum of finite-dimensional simple modules. Fix an H-weight m. In a simple summand of highest weight $N \geqslant |m|$, the raising and lowering lengths at weight m are $(N - m)/2$ and $(N + m)/2$, respectively. In particular their difference is m.

Let $V_m^{\leqslant N}$ be the sum of the weight-m subspaces belonging to simple summands of highest weight at most N. The preceding formula identifies $V_m^{\leqslant N}$ with the kernel of the corresponding power of E on the weight-m space. Since the basis is E-perfect, its elements of weight m and raising length at most $(N - m)/2$ form a basis of this filtration step. Thus the given basis induces a basis on every successive isotypic quotient.

On the quotient of highest weight N, the map E between consecutive weight spaces is an isomorphism whenever the source is not highest. Normalized E-perfectness says that its matrix between the induced bases has exactly one nonzero entry in each column, with coefficient ρ. Because the map is an isomorphism, these leading arrows are bijections between the two quotient bases. Its inverse determines a unique partial lowering arrow.

For a basis element of weight m, write $\rho = (N - m)/2$ and $\lambda = (N + m)/2$ on its highest nonzero isotypic quotient. There the standard $\mathfrak{sl}_2$ identity is $EF = (\rho + 1)\lambda$ on the weight-m space. The normalized E-coefficient on the preceding weight space is $\rho + 1$; hence the coefficient of F along the inverse arrow is λ. Contributions from lower isotypic quotients have lowering length at most $\lambda - 2$, giving exactly the normalized F-perfectness expansion. The two arrows are partial inverses and $\lambda - \rho = m$.

Although infinitely many simple summands may occur, every vector is supported on finitely many of them. All filtration and quotient calculations above therefore take place in finite sums, and no finite-dimensionality assumption on the whole module is needed. □

Write r_i, l_i for the leading E_i, F_i maps and define $\mathrm{wt}(b) = \sum_i \mathrm{wt}_i(b)\bar{\varpi}_i$, where $H_i\vartheta_b = \mathrm{wt}_i(b)\vartheta_b$. Let $\bar{\alpha}_j = \sum_i c_{ij}\bar{\varpi}_i$. The Cartan brackets and Lemma 7.4 give the crystal axioms with these coroot-coordinate weights. On affine components these are classical weights, with the null root mapped to zero, as in [12, Section 4.4]. The two string functions are their actual iteration lengths, so the crystal is *seminormal*. In an i-sink optimizer,

$$g(r_i b) = g(b) + m_i - e_i \quad \text{if } g_i(b) > 0, \qquad \rho_i(b) = g_i(b).$$

Since the rational derivations are independent of the seed, so are their leading maps.

We use *normal crystal* in the standard Kac–Moody sense: for every subset of simple roots of finite type, the restricted crystal is a disjoint union of finite-dimensional highest-weight crystals; see [12, Section 8.4] and [5, 25]. Normality is a condition on all finite-type Levi restrictions; it makes no assertion that an affine crystal is itself a union of affine highest-weight crystals.

Corollary 7.5. *The theta crystal is normal: for every $J \subseteq I$ of finite Cartan type, its restriction to the colors in J is a disjoint union of crystals of finite-dimensional simple $\mathfrak{g}(C_J)$-modules.*

Proof. The restricted integrable module is a direct sum of finite-dimensional simple modules. Apply the perfect-basis comparison theorem [3, Main Theorem 5.55] as in [12, Section 8.4]; see also [5, 25]. □

Corollary 7.6. *Suppose $\mathcal{U}^+$ has a basis of theta functions with finite principal-coefficient lifts. Reachable sink boundaries with exact source partners make this basis normalized perfect for both collections of derivations. Their leading operators define a normal crystal with integral coroot-coordinate weights. In particular, this conclusion holds when $\mathcal{A}^{\mathrm{loc}} = \mathcal{U}^{\mathrm{loc}}$.*

Proof. Theorem 6.9 gives the integrable action on $\mathcal{U}^+$. Theorem 7.1 gives positive perfectness, and Lemma 7.4 gives negative perfectness and inverse maps. Apply Corollary 7.5 and, for the final assertion, Proposition 7.3. □

Corollary 7.7. *Assume reddening and exact source partners for the sink boundaries in I. Then $\mathcal{B}_\theta^+$ is a normal crystal. Its raising maps are r_i, and its lowering maps are $l_i = r^\star_{\sigma(i)}$. Its integral coroot gradings are the boundary commutators.*

Proof. Apply the preceding rank-one and finite-type arguments to the stable regular middle algebra of Theorem 7.2. □

Write $a = \Phi_{\boldsymbol{\mu}_i^+}(g)$ for the g-vector in an i-sink optimizer. The complete string is

$$r_i^n g = \Phi^{-1}_{\boldsymbol{\mu}_i^+}(a + nv_i), \qquad l_i^n g = \Phi^{-1}_{\boldsymbol{\mu}_i^+}(a - nv_i),$$

for $0 \leqslant n \leqslant \rho_i(g)$ and $0 \leqslant n \leqslant \lambda_i(g)$, respectively. The boundary criterion in Section 4.3 determines when the source maps provide these inverses.

7.4. **Compatible weights and Weyl symmetries.** Coroot gradings alone need not yield a full weight-lattice realization with linearly independent simple roots. To decide whether compatible gradings can complete them to a full realization, fix a seed and let

$$q_i = g(M_i) - e_i, \qquad L_{\mathbb{Q}} = \mathbb{Q}^V / \operatorname{row}_{\mathbb{Q}}(B_{U,V}).$$

The classes of q_i are the degrees of the root derivations modulo the exchange relations. Compatible gradings are linear forms on $L_{\mathbb{Q}}$.

Proposition 7.8. *The coroot grading extends by compatible linear gradings to a full realization with independent simple roots if and only if the classes of q_i are linearly independent in $L_{\mathbb{Q}}$. Equivalently,*

$$\operatorname{rank}\left(B_{U,V}\$q_i\right)_{i\in I}\big) = \operatorname{rank}(B_{U,V}) + |I|.$$

When this holds, $|I| - \operatorname{rank} C$ additional integral gradings suffice, with a suitable integral weight lattice.

Proof. On the span of the independent classes q_i, the existing coroot grading has matrix C. Choose rational forms separating its kernel, extend them to $L_{\mathbb{Q}}$, and clear denominators. Together with the coroot forms, they send the q_i to independent integral vectors, which serve as the simple roots. The dimension is $2|I| - \operatorname{rank} C$. Conversely, a relation among the q_i would give the same relation among simple roots in every compatible realization. If the weight lattice is prescribed in advance, integrality may require an additional primitivity condition. □

Proposition 4.5 describes the coefficient vectors that preserve a prescribed system of exact boundary pairs.

Question. Given a simultaneous system of exact boundary pairs, can one adjoin finitely many frozen vertices so that all the pairs remain exact and the simple-root degree classes become linearly independent in the extended quotient $L_{\mathbb{Q}}$? If so, what is the minimum number of additional frozen vertices?

7.4.1. *String reflections and Tits automorphisms.* The *string reflection* is

$$s_i b = \begin{cases} l_i^{\mathrm{wt}_i(b)} b, & \mathrm{wt}_i(b) \geqslant 0, \\ r_i^{-\mathrm{wt}_i(b)} b, & \mathrm{wt}_i(b) < 0. \end{cases}$$

It is an involution and sends weight ν to $\nu - \nu(h_i)\bar{\alpha}_i$. In an i-optimizer, $g(s_i b) = g(b) - \mathrm{wt}_i(b)(m_i - e_i)$. Normality on finite-type rank-two restrictions gives the finite braid relations [28, Corollaries 4.9 and 4.11]; all other rank-two pairs have infinite Coxeter exponent. Thus the reflections define a Weyl group action on theta g-vectors, compatible with mutation.

Local nilpotence also defines algebra automorphisms $\dot{s}_i = \exp(E_i)\exp(-F_i)\exp(E_i)$. For h in the integral coroot lattice and a weight-ν vector a, define $\Theta_h(a) = (-1)^{\nu(h)}a$. The standard Tits relations [22, 23] give

$$\dot{s}_i^2 = \Theta_{h_i}, \qquad \dot{s}_i \Theta_h \dot{s}_i^{-1} = \Theta_{h - \alpha_i(h)h_i}, \qquad \Theta_h \Theta_{h'} = \Theta_{h+h'}, \quad \Theta_{2h} = 1,$$

together with the braid relations. This is the Tits extension action of [12, Theorem 8.16], applied to the present integrable module algebra.

Write $A^{(i)}_{\leqslant N}$ for the isotypic filtration in Lemma 7.4 and set $N_i(b) = \rho_i(b) + \lambda_i(b)$. Its leading action is

$$\dot{s}_i \vartheta_b \equiv (-1)^{\lambda_i(b)} \vartheta_{s_i b} \pmod{A^{(i)}_{\leqslant N_i(b)-1}}. \tag{7.1}$$

This is the perfect-basis compatibility in [12, Corollary 8.17]. To fix the sign convention, identify an associated-graded string with $X^{N-q}Y^q$, where $q = \rho_i$, $E_i = X\partial_Y$, and $F_i = Y\partial_X$. The lift sends X to $-Y$ and Y to X, giving the sign $(-1)^{N-q} = (-1)^{\lambda_i}$. Lower filtration terms may be present in the action on individual theta functions.

Question. Under what conditions does the Tits automorphism $\dot{s}_i$ send every regular theta function to a theta function up to sign, rather than merely agree with it modulo lower terms of the isotypic filtration?

7.4.2. *Opposite symmetries and Lusztig involutions.* Let J be of finite type, and write $w_{0,J}\alpha_i = -\alpha_{i^*}$. If $\mathcal{H}_\lambda$ denotes the J-highest elements of restricted J-weight λ, normality identifies the restricted crystal with $\coprod_\lambda B_J(\lambda) \times \mathcal{H}_\lambda$. Let ξ_λ be the Lusztig involution on $B_J(\lambda)$, characterized by exchanging highest and lowest elements and satisfying $\xi_\lambda r_i = l_{i^*}\xi_\lambda$ [28, Section 2.4].

Proposition 7.9. *Suppose an opposite symmetry induces a bijection κ with $\kappa r_i = l_{i^*}\kappa$ and $\kappa l_i = r_{i^*}\kappa$ for $i \in J$. Then there are permutations p_λ of $\mathcal{H}_\lambda$ such that*

$$\kappa(b, h) = (\xi_\lambda(b), p_\lambda(h)).$$

Consequently κ agrees with the Lusztig involution on each component if and only if it fixes every component, and $\kappa^2 = \mathrm{Id}$ precisely when each p_λ is involutive.

Proof. The intertwining relations map highest elements to lowest elements of components with the same highest weight. They determine the map on each connected highest-weight crystal uniquely, giving the stated factorization. □

Proposition 5.3 gives these intertwining relations when the symmetry matches sink color i with the exact source partner of i^*, and matches the source direction back to the corresponding sink direction. A direct relative change of coefficients instead intertwines raising maps with raising maps.

7.4.3. *Simultaneous kernels.*

Corollary 7.10. *Let J be a finite family of perfect boundary derivations D_j on Mid^+, with string lengths $\rho_j(g)$. For integers $n_j \geqslant 0$,*

$$\bigcap_{j \in J} \ker D_j^{n_j+1} = \bigoplus_{\rho_j(g) \leqslant n_j \ (j \in J)} \Bbbk \vartheta_g.$$

In particular, the common invariant algebra has the theta basis with all $\rho_j(g) = 0$. For positive operators of a finite-type exact subsystem, these are its highest-weight theta elements. The assertions also hold in $\mathcal{U}^+$ when the theta functions span it.

Proof. Apply the simultaneous kernel-filtration argument of [12, Corollary 8.11]: each kernel is a coordinate subspace in the same perfect basis. □

8. Cartan type and geometric applications

8.1. **Growth of derivation algebras.** Let $C = (c_{ij})_{i,j\in I}$ be a symmetrizable generalized Cartan matrix, and let D_I be a positive rational diagonal matrix such that D_IC is symmetric. We extend the growth argument of [12, Proposition 8.7] from Laurent algebras to fields of finite transcendence degree.

Lemma 8.1. *Let $K/\Bbbk$ have characteristic zero and finite transcendence degree d. If $\mathfrak{d}$ is generated by finitely many $\Bbbk$-derivations of K, then its filtration by Lie word length satisfies*

$$\dim_{\Bbbk} \mathfrak{d}_{\leqslant n} = O(n^d).$$

In particular, $\mathrm{Der}_{\Bbbk}(K)$ *contains no free Lie algebra on two generators.*

Proof. Let $D_1, \dots, D_r$ generate $\mathfrak{d}$, and choose a transcendence basis $x_1, \dots, x_d$. The field

$$L = \Bbbk(x_1, \dots, x_d, D_a(x_j) : 1 \leqslant a \leqslant r,\ 1 \leqslant j \leqslant d)$$

is finite over $\Bbbk(x_1, \dots, x_d)$ and is preserved by every D_a. Indeed, differentiating the separable minimal polynomial of an element of L expresses its derivative in L. A derivation vanishing on L vanishes on the algebraic extension K/L, so restriction is injective on $\mathfrak{d}$.

Write $L = \Bbbk(y_1, \dots, y_s)$ and choose a common denominator q for the finitely many $D_a(y_b)$. Then all D_a preserve $R = \Bbbk[y_1, \dots, y_s, q^{-1}]$. For its degree filtration there is a constant c with $D_a(R_{\leqslant m}) \subseteq R_{\leqslant m+c}$. A Lie word of length at most n therefore sends each algebra generator to $R_{\leqslant 1+cn}$. Evaluation on the generators is injective, giving

$$\dim_{\Bbbk} \mathfrak{d}_{\leqslant n} \leqslant (s+1) \dim_{\Bbbk} R_{\leqslant 1+cn} = O(n^d).$$

The last bound is Hilbert growth for a finitely generated algebra of dimension d. A free Lie algebra on two generators has exponential growth [7, Proposition 4.3]. □

Theorem 8.2. *Let A be a commutative domain over a characteristic-zero field, with $\mathrm{trdeg}_{\Bbbk} \mathrm{Frac}(A) < \infty$. Suppose derivations E_i, F_i, H_i satisfy the Chevalley–Serre relations for C, and every connected component acts nontrivially. Then D_IC is positive semidefinite.*

Proof. If $E_i = 0$, then $H_i = [E_i, F_i] = 0$, and $[H_i, E_j] = c_{ij}E_j$ forces $E_j = 0$ at every adjacent vertex. This propagates through the component. The corresponding H_j and F_j also vanish. Thus the hypothesis implies $E_i \neq 0$ for all i.

Extend the derivations to $K = \mathrm{Frac}(A)$. If D_IC is not positive semidefinite, a negative principal minor has a connected factor with negative determinant. It gives a connected, nonsingular, indefinite principal submatrix C_J. The restricted operators define $\phi : \mathfrak{g}_{\Bbbk}(C_J) \to \mathrm{Der}_{\Bbbk}(K)$. A relation $\sum_{i\in J} a_iH_i = 0$ implies $C_J^{\mathsf{T}}a = 0$ by bracketing with the nonzero E_j. Thus ϕ is injective on the Cartan subalgebra. The ideal-intersection property for Kac–Moody algebras then makes ϕ injective [7, Section 2.1].

The Perron eigenvalue of $2I - C_J$ is greater than two. Approximating a positive Perron eigenvector by a rational vector and clearing denominators gives $v \in \mathbb{Z}^J_{>0}$ with $C_Jv < 0$. The fundamental imaginary-root criterion [22] shows that $\beta = \sum_{i\in J} v_i\alpha_i$ is an imaginary root, with $(\beta, \beta) = v^{\mathsf{T}}D_JC_Jv < 0$. By [7, Corollary 4.8], $\mathfrak{g}_{\mathbb{C}}(C_J)$ contains a free Lie algebra on two generators and has exponential growth in its Chevalley generators. The Chevalley–Serre presentation is defined over $\mathbb{Q}$; dimensions in its finite Lie word filtration are unchanged by extension to any characteristic-zero field. The injective image of ϕ therefore contradicts Lemma 8.1. □

The argument uses neither finite generation of A nor local nilpotence. The hypothesis that each connected component acts nontrivially is essential: an indefinite component could otherwise act by zero.

Corollary 8.3. *Under the hypotheses of Theorem 8.2, every connected component of C has finite or affine type. For $i \neq j$, one has $0 \leqslant c_{ij}c_{ji} \leqslant 4$. Equality with four makes $\{i, j\}$ an entire connected component. The possible nonzero pairs $(-c_{ij}, -c_{ji})$ are listed below, with the last three giving the rank-two affine cases:*

$$(1,1),\ (1,2),\ (2,1),\ (1,3),\ (3,1),\ (1,4),\ (4,1),\ (2,2).$$

Proof. The classification of indecomposable positive semidefinite symmetrizations gives finite or affine type [22]. The two-by-two principal minor gives $c_{ij}c_{ji} \leqslant 4$. If $a = -c_{ij}$ and $b = -c_{ji}$ satisfy $ab = 4$, choose w supported on $\{i, j\}$ with $w_i = a$ and $w_j = 2$. Then $w^{\mathsf{T}} D_I C w = 0$, so positive semidefiniteness implies $Cw = 0$. For $k \notin \{i, j\}$, this reads $ac_{ki} + 2c_{kj} = 0$. Both terms are nonpositive, hence both vanish. The symmetric zero pattern of a Cartan matrix proves the component assertion. □

8.2. **Maximal unipotent subgroups.** Let G be a simply connected complex semisimple group, let T be a maximal torus, and let U_G be the maximal unipotent subgroup associated with a choice of positive roots. Fix a pinning, write e_i for the positive Chevalley generators, and set $x_i(t) = \exp(te_i)$. The simple roots, simple coroots, and fundamental weights are α_i, $\alpha_i^\vee$, and ϖ_i, respectively. The left and right simple-root derivations are

$$D_i^L f(u) = \frac{d}{dt}\Big|_{t=0} f(x_i(t)u), \qquad D_i^R f(u) = \frac{d}{dt}\Big|_{t=0} f(ux_i(t)).$$

We use the conjugation action $(h \cdot f)(u) = f(h^{-1}uh)$ of T. Thus a generalized minor $\Delta_{\gamma,\delta}$ has weight $\delta - \gamma$, and both derivations raise the weight by α_i. Write $\mathcal{B}_\theta(U_G)$ for the regular theta basis of $\mathbb{C}[U_G]$. Its theta basis is biperfect for the left and right root derivations by [36, Theorem 1.1], and the corresponding bicrystal is $B(\infty)$ by [24, Theorem 1.7].

Let w_0 be the longest Weyl-group element, and define i^* by $w_0\alpha_i = -\alpha_{i^*}$. The frozen variables are $A_j = \Delta_{\varpi_j, w_0\varpi_j}$. The derivation D_i^R is the canonical sink derivation at A_{i^*}, while D_i^L is the canonical source derivation at A_i; the pinned coefficients are one [36, Lemmas 5.1 and 5.2]. We index the crystal operators by simple roots, and distinguish these indices from those of the frozen minors.

Proposition 8.4. *The two boundary Cartan matrices are $C = (\langle \alpha_i^\vee, \alpha_j \rangle)_{i,j}$. The coroot coordinates of the conjugation weight give the unique integral compatible grading adapted to either the left or right simple-root derivations.*

Let ι be the pinned antiautomorphism of U_G fixing every $x_i(t)$. Its pullback is an opposite cluster automorphism and permutes the regular theta basis. The induced map on $B(\infty)$ is Kashiwara's involution.

Proof. The left and right regular actions are faithful on the Lie algebra of U_G, so the optimal Serre exponents of their simple-root derivations are $1 - c_{ij}$. Proposition 6.2 therefore identifies the two boundary Cartan matrices with C.

A reduced-word minor seed has $\ell(w_0)$ variables, precisely $\operatorname{rank} G$ of them frozen, and its extended exchange matrix has full rank; see [33, Theorem 8.1.2] and [36, Section 5]. Hence the grading space $L_{\mathbb{Q}}$ of Section 3 has dimension $\operatorname{rank} G$. The conjugation grading is integral and compatible. Its values on the root shifts $[q_i]$ are the columns of C, since D_i^R has degree α_i. The matrix C is invertible, so these root shifts form a basis of $L_{\mathbb{Q}}$. An adapted grading is consequently unique. The same argument applies to the left derivations.

We verify the seed action of ι. For a reduced word $\mathbf{j} = (j_1, \ldots, j_N)$ of w_0, let $v_k = s_{j_1} \cdots s_{j_k}$, with $v_0 = 1$. Its *prefix minor seed* and *suffix minor seed* have coordinates

$$X_k(\mathbf{j}) = \Delta_{\varpi_{j_k}, v_k\varpi_{j_k}}, \qquad Y_k(\mathbf{j}) = \Delta_{v_{k-1}\varpi_{j_k}, w_0\varpi_{j_k}}.$$

For a fixed word they are mutation-equivalent [33, Theorems 8.1.2 and 8.2.2]; changing the reduced word also preserves the mutation class [13, Lemma 3.8]. Set $\mathbf{j}^\dagger = (j_N^*, \dots, j_1^*)$ and $\rho(k) = N + 1 - k$. The standard identity for generalized minors [16, Proposition 2.7] gives

$$\iota^* X_k(\mathbf{j}) = \Delta_{v_k w_0 \varpi_{j_k^*}, w_0 \varpi_{j_k^*}} = Y_{\rho(k)}(\mathbf{j}^\dagger).$$

Indeed, the prefix of length $N-k$ in $\mathbf{j}^\dagger$ is $v_k w_0$. Reversal interchanges previous and next occurrences of each letter. The exchange formulas in [33, Remarks 8.1.4 and A.3.1] therefore give

$$\widetilde{b}^Y_{\rho(a),\rho(k)}(\mathbf{j}^\dagger) = -\widetilde{b}^X_{a,k}(\mathbf{j})$$

for every mutable k and every a. Last occurrences become first occurrences, so frozen indices are preserved. More explicitly, $\iota^* A_j = A_{j^*}$; the symmetrizing weights are also preserved. Thus ι^* is an opposite seed identification.

The minor seeds are injective-reachable [33, Lemma 9.1.3]. Proposition 5.3 now shows that ι^* permutes the regular theta basis. Since $\iota(x_i(t)) = x_i(t)$ and ι reverses multiplication,

$$\iota^* D_i^R = D_i^L \iota^*, \qquad \iota^* D_i^L = D_i^R \iota^*.$$

Consequently the induced involution fixes 1 and exchanges the two crystal structures with the same simple-root color. Under the bicrystal identification with $B(\infty)$, it is Kashiwara's involution [24, Theorem 1.7 and Remark 1.8]. □

If σ denotes the tropical map of this opposite cluster automorphism, its action on the original theta g-vectors is $\kappa_\sigma = \eta_\theta^{-1}\sigma$, as in Proposition 5.3. The proposition applies to every finite semisimple type. For its simply-laced categorical counterpart, see [12, Proposition 10.1].

8.3. **Lusztig's involution on the base affine space.** Let $X = U_G^- \backslash G$, with G acting by right translation. The left torus action gives the dominant-weight grading

$$\mathbb{C}[X] = \bigoplus_{\lambda \in P_+} R_\lambda, \qquad R_\lambda = \{f : f(hg) = \lambda(h) f(g),\ h \in T\} \cong L(\lambda).$$

We call λ the degree of $f \in R_\lambda$ and use $\mathrm{wt}(f)$ for its weight under the right torus action. Minimal monomial lifting of the reduced-word seeds realizes $\mathbb{C}[X]$ as a polynomial-frozen upper cluster algebra [19, arXiv version, Theorem 8.3.2 and Lemma 8.3.4]. The lifted seeds are mutation-equivalent, admit mutable reddening, and give a regular theta basis homogeneous for both gradings [13, Corollary 3.9, Lemma 5.1 and Theorem 5.4].

The linear Lusztig involution Ξ_λ of R_λ exchanges the pinned highest and lowest vectors and satisfies $\Xi_\lambda E_i = F_{i^*} \Xi_\lambda$ and $\Xi_\lambda F_i = E_{i^*} \Xi_\lambda$, where E_i, F_i are the right Chevalley derivations [28, Section 2.4]. Theorem 8.6 realizes this linear map on the theta functions themselves. The two families of frozen minors have the indexing of [12, Section 10.3],

$$x_{p_i} = \Delta_{\varpi_i, \varpi_i}, \qquad x_{q_i} = \Delta_{\varpi_{i^*}, w_0 \varpi_{i^*}}.$$

We first identify the pinning of the two boundary actions.

Proposition 8.5. *In every simple factor of rank at least two, $E_i = R_{q_i}$ and $F_i = R^\star_{p_i}$. The pairs (q_i, p_i) are exact and have the Cartan matrix of that factor. In every finite type the theta basis of R_λ is normalized perfect for the pinned Chevalley derivations and has crystal $B(\lambda)$.*

Proof. Fix a reduced word $\mathbf{j} = (j_1, \dots, j_N)$ for w_0 and write $v_k = s_{j_1} \cdots s_{j_k}$. Its lifted prefix variables are $x_k = \Delta_{\varpi_{j_k}, v_k \varpi_{j_k}}$. Write $s(k)$ for the next occurrence of j_k, with $s(k) = +\infty$ for a last occurrence. If a word starts with i in a simple factor of rank at least two, i occurs again. Otherwise $w_0 \varpi_i = s_i \varpi_i$, which pairs positively with a neighboring coroot. Reversal gives the corresponding assertion for a word ending in i.

Choose first a word ending in i^* and let N^- be its previous occurrence. In the unipotent chart the pinned right derivative is $x_{N^-} \partial_{x_N}$ and the terminal row is $\widetilde{b}_{Nk} = -\delta_{k,N^-}$ [36, Lemma 5.1 and equation (15)]. The two minors have the same degree ϖ_{i^*} and E_i kills every

highest minor x_{p_j}. Hence the formula holds in the lifted chart, where $x_N = x_{q_i}$. The two occurrences have equal symmetrizing weight, so $b_{q_i k} = -\delta_{k,N^-}$ as well. This gives $E_i = R_{q_i}$ with coefficient one.

For the negative derivative choose a word starting with i. Since $v_k^{-1}(-\alpha_i)$ is positive for every $k \geqslant 1$, F_i kills each x_k. On the highest minors, $F_i x_{p_j} = \delta_{ij} x_1$ by the pinned fundamental i-string. To compute the added rows, let a_j be the first occurrence of j in the word. Minimal lifting gives $\widetilde{b}_{p_j k} = -\sum_{r:j_r=j} \widetilde{b}_{rk}$. For each mutable k, the prefix exchange rule yields

$$\widetilde{b}_{p_j k} = \begin{cases} 1 & k = a_j, \\ \langle \alpha_j^\vee, \alpha_{j_k} \rangle & k < a_j < s(k), \\ 0 & \text{otherwise.} \end{cases} \tag{8.1}$$

Indeed, for $j = j_k$ the horizontal entries cancel except at the first occurrence. For $j \neq j_k$, intervals of letter j crossing $(k, s(k))$ contribute in pairs of opposite sign, unless the first occurrence of j lies in $(k, s(k))$. In that case only the crossing at the right endpoint contributes. Taking $j = i$ gives $\widetilde{b}_{p_i k} = \delta_{k1}$ and, since $d_{p_i} = d_1$, $b_{p_i k} = \delta_{k1}$. Thus $F_i = R^\star_{p_i}$.

Now $E_i x_{p_i} = 0$ and $H_i x_{p_i} = x_{p_i}$, whence $E_i F_i x_{p_i} = x_{p_i}$. This is the distinct-partner normalization of Proposition 6.5. Faithfulness of the simple Lie algebra action gives its usual optimal Serre exponents, hence its Cartan matrix. Perfectness and normality follow from the exact-pair theorems, or from the root-operator construction of [13, Corollaries 5.12 and 5.15] with the pinning just checked. On an A_1 factor, the coordinate ring is the polynomial ring on a highest and a lowest coordinate. Its monomials are perfect for the usual $\mathfrak{sl}_2$ derivations [12, Remark 10.6]. The seed in that factor has no mutable vertices. Products give the semisimple case. □

Theorem 8.6. *The algebra involution $\Xi = \bigoplus_\lambda \Xi_\lambda$ is a degree-preserving opposite cluster automorphism. It exchanges x_{p_i} with $x_{q_i *}$. If ξ_λ denotes Lusztig's involution on the degree-λ crystal, then*

$$\Xi(\vartheta_g) = \vartheta_{\xi_\lambda(g)} \qquad (\vartheta_g \in R_\lambda).$$

Proof. For the Cartan product $m : R_\lambda \otimes R_\mu \to R_{\lambda+\mu}$, the maps $\Xi_{\lambda+\mu} m$ and $m(\Xi_\lambda \otimes \Xi_\mu)$ intertwine the same twisted action. The Cartan component has multiplicity one. The two maps agree on the tensor of the highest vectors because the pinned lowest minors multiply with coefficient one. Thus Ξ preserves multiplication. It takes the pinned extremal vector of weight $w\varpi_i$ to that of weight $w_0 w \varpi_i$ [28, Proposition 2.3], giving

$$\Xi(\Delta_{\varpi_i, w\varpi_i}) = \Delta_{\varpi_i, w_0 w \varpi_i}.$$

Reverse $\mathbf{j}$. The prefix of the reversed word of length $N-k$ represents $w_0 v_k$. The reflections lying strictly between k and $s(k)$ fix ϖ_{j_k}, so the image of x_k is the reversed-word coordinate at

$$\pi(k) = \begin{cases} N + 1 - s(k) & s(k) < +\infty, \\ p_{j_k} & s(k) = +\infty. \end{cases}$$

The image of p_i is the last occurrence of i in the reversed word. This gives the asserted exchange of the frozen families.

For the full extended matrix, assign to each word position the interval $[k, s(k)]$ and to p_j the interval $[-\infty, a_j]$. The word exchange rule uses consecutive intervals of the same letter and interlacing intervals of distinct letters [2, Definitions 2.2–2.3]. Formula (8.1) is precisely this rule at an initial frozen interval. The map π sends every interval to its image under the reflection $t \mapsto N + 1 - t$. It reverses the order of consecutive intervals and interchanges the two orders of interlacing. The reflection preserves the ordered pair of simple-root indices for each entry, hence also its Cartan factor in nonsimply-laced type. Thus

$$\widetilde{b}^{\mathbf{j}^{\mathrm{rev}}}_{\pi(a),\pi(k)} = -\widetilde{b}^{\mathbf{j}}_{a,k} \qquad (k \text{ mutable},\ a \text{ arbitrary}).$$

The two lifted seeds belong to the same mutation class, so Ξ is an opposite cluster automorphism. By Proposition 5.3 it permutes the regular theta functions. It preserves R_λ, sends highest to lowest, and interchanges raising of color i with lowering of color i^*. These properties characterize ξ_λ on the connected crystal $B(\lambda)$. □

Let $\mathcal{B}^+_{\theta,\lambda}$ be the set of g-vectors, in the reference seed, of regular theta functions of degree λ. These are the integral points of the bounded degree polyhedron in [13, Theorem 5.4]. Imposing a weight equation singles out the theta basis of its weight space. The perfect-basis tensor product rule [24, Proposition 1.2] gives

$$[L(\lambda)\otimes L(\mu) : L(\nu)] = \#\{g \in \mathcal{B}^+_{\theta,\lambda} : \mathrm{wt}(g) = \nu-\mu,\ \rho_i(g) \leqslant \langle \alpha_i^\vee, \mu\rangle \text{ for every } i\}.$$

The string bounds give finitely many linear inequalities in the reference seed [13, Proposition 5.5]. For the left action on G/U_G, a component of right-T degree λ is $L(\lambda)^*$. The corresponding involution is $\xi_{-w_0\lambda}$, which is ξ_λ in types B_n and C_n.

8.4. **Lagrangian Grassmannians.** Let $G = \mathrm{Sp}_{2n}(\mathbb{C})$, $n \geqslant 2$, and let $X = G/P_n = \mathrm{LGr}(n,2n)$. Its Plücker ring is $\mathcal{R} = \bigoplus_{m\geqslant 0} L(m\varpi_n)^*$. Write b_0 for the section cutting out the Schubert divisor and b_i for the minimal lift of the terminal minor of letter i, $1 \leqslant i \leqslant n$. The sections b_0, b_n are the highest and lowest Plücker coordinates. Indeed b_0 cuts out the complement of the open cell with multiplicity one, and the terminal n-minor on that cell is the opposite Plücker coordinate divided by b_0. Its numerator has order zero along $b_0 = 0$, so it is already the minimal lift. The minimally lifted word seed identifies $\mathcal{R}$ with its polynomial-frozen upper algebra and gives its homogeneous theta basis [13, Theorems 3.4 and 5.4].

Write $\theta = 2e_1$ for the highest root of C_n, and choose root vectors $X_\theta, X_{-\theta}$ with $[X_\theta, X_{-\theta}] = H_\theta$. Let $\mathcal{E}$ denote the Euler derivation for the Plücker grading.

Lemma 8.7. *For $q \geqslant 0$ there is a G-module decomposition*

$$\mathrm{Der}(\mathcal{R})_q \simeq L(q\varpi_n)^* \oplus L(\theta + q\varpi_n)^*.$$

The lowest weight lines of the two summands are generated by $b_n^q\mathcal{E}$ and $b_n^q X_{-\theta}$, respectively. Their highest weight lines are generated by $b_0^q\mathcal{E}$ and $b_0^q X_\theta$.

Proof. The punctured cone is the $\mathbb{G}_m$-bundle of $\mathcal{O}_X(-1)$. Its homogeneous vector fields of degree q are sections of the twisted Atiyah sequence [1]

$$0 \longrightarrow \mathcal{O}_X(q) \longrightarrow \mathcal{A}(q) \longrightarrow T_X(q) \longrightarrow 0.$$

Such fields extend across the vertex because $\mathcal{R}$ is normal and their values on the degree-one generators extend. The tangent representation of the Levi is irreducible with lowest weight $-\theta$. Borel–Weil–Bott gives $H^0(X, T_X(q)) = L(\theta + q\varpi_n)^*$ and $H^1(X, \mathcal{O}_X(q)) = 0$ [4]. Complete reducibility splits the resulting sequence on global sections. Multiplication by an extremal Plücker section gives the asserted vertical lines. The two products with root vectors have nonzero horizontal symbols and extremal weights $\pm(\theta + q\varpi_n)$. □

Proposition 8.8. *All frozen vertices admit sink and source optimizers. Their exact partners are $\sigma(i) = n - i$, $0 \leqslant i \leqslant n$. The resulting boundary action has type $C_n^{(1)}$, preserves Plücker degree, and acts at level zero. Its homogeneous theta basis is normalized perfect. The finite colors in degree m give $B(m\varpi_n)$. For $m > 0$ there is no affine highest element.*

Proof. Put $J = \{1, \dots, n-1\}$ and $w = w_0 w_J$. We use the word

$$(n, n-1, \dots, 1)(n, n-1, \dots, 2)\cdots(n).$$

It sends e_a to $-e_{n+1-a}$ and has length $\ell(w) = n(n+1)/2$. Give each letter $i < n$ weight $d_i = 1$ and letter n weight $d_n = 2$, and set $d_0 = 2$. Use the sign of the word matrix in [13, Section 3], for which the adapted positive simple-root charts have nonpositive active frozen rows.

We record the initial rows at the two ends. Let $N = \ell(w)$, let p be the penultimate position of letter n, and let $s(k)$ be the successor of position k. The only surviving Schubert divisor comes from deleting the first letter. Its column on mutable rows is $(\mathbf{1}_{s(k)=1} - \mathbf{1}_{k=1})_k = -e_1^{\mathsf{T}}$ in column notation [13, Section 5.1]. The terminal minor b_n is at position N. The word row k has last nonzero entry -1 in column $s(k)$ [13, proof of Proposition 3.6]. Thus its entry in column N vanishes unless $k = p$, when it is -1. Since the two mutable neighbors also have weight 2, the two frozen rows on mutable indices are e_1 and e_p, respectively. Both ends are source-optimized in this seed. The word pattern admits reddening [13, Lemma 5.1], so Lemma 1.6 gives sink optimizers for both.

For the finite colors let X_i, Y_i be geometric Chevalley derivations with $[X_i, Y_i] = H_i^{\mathrm{geom}}$. In the chart adapted to X_i, the only frozen variable on which X_i acts nontrivially is b_i. For $i < n$, $w^{-1}\alpha_i = \alpha_{n-i}$, so the negative adapted chart for Y_i has active variable b_{n-i}. The remaining simple root n is a left descent and Y_n has active variable b_0 [13, Theorem 4.1]. These charts give every remaining optimizer. The mutable exponents in that theorem are precisely the exponents of the canonical multipliers, since $d_i = d_{n-i}$. The geometric derivations therefore differ from $R_i, R^\star_{n-i}$, respectively, by a scalar and an ordinary frozen monomial.

In fact these monomials are constant. Every frozen variable has positive Plücker degree. Divisibility by a nonconstant frozen monomial would force the image of $\mathcal{R}_1$ under the geometric derivation into a space of dimension at most one, since $\mathcal{R}_0 = \mathbb{C}$. A long-root operator has rank at least 2^{n-1} on $\mathcal{R}_1$, by the distinct pairs in the extremal orbit $\{(\pm 1, \dots, \pm 1)\}$. For a short simple root $\alpha_i = e_i - e_{i+1}$, take an extremal weight with signs $(+,-)$ in positions $i, i+1$. Its pairing with $\alpha_i^\vee$ is two, so the corresponding root operator has rank at least two on this string alone. We may therefore write $X_i = c_i R_i$ and $Y_i = c_i' R^\star_{n-i}$.

To fix the scalars, put $k = n - i$. When $k \neq i$, the $\mathfrak{sl}_2$-string through b_k has raising length zero and lowering length one, so $X_i Y_i b_k = b_k$. When $k = i$, the weight vector b_i has raising and lowering lengths both equal to one. Its H_i^{geom}-weight is therefore zero, and the rank-one relations give $X_i^2 Y_i b_i = 2 X_i b_i$. Thus, up to the common factor $(c_i c_i')^{-1}$, the canonical operators satisfy the normalization in Proposition 6.5. To determine that factor, expand the cluster monomial $R_k^\star b_k$ in an i-sink optimizer and take its maximal divided R_i-derivative. Its pointed term survives with coefficient one, by the highest frozen face calculation in that proposition. The extended exchange matrix has full mutable rank [13, Section 3], so no other F-exponent contributes to that term. The divided derivative is a scalar times the monomial b_k or $R_i b_i$, hence the scalar is one. This proves $c_i c_i' = 1$ and exactness of all finite pairs.

The only sink still to be paired is 0, and the only unused source is n. Apply Lemma 6.8 to each finite pair, first with the extra sink 0, then dually with the extra source n. It gives

$$[Y_j, R_0] = [X_j, R_n^\star] = 0 \qquad (1 \leqslant j \leqslant n).$$

Coefficient-row mutation is invertible and fixes the zero row. Thus the rows at 0 and n remain nonzero, and their canonical multipliers are nonconstant mutable cluster monomials. Since b_0, b_n have degree one, the two derivations are homogeneous of degrees $q, q' \geqslant 0$. Lemma 8.7 now places R_0 in the span of the two lowest vectors in degree q. Its Euler coefficient vanishes on evaluating at b_n, since $R_0 b_n = X_{-\theta} b_n = 0$. Likewise the Euler coefficient of $R_n^\star$ vanishes on evaluating at b_0. Thus

$$R_0 = c b_n^q X_{-\theta}, \qquad R_n^\star = c' b_0^{q'} X_\theta.$$

In its optimizer the mutable monomial $R_0 b_0$ has order zero along $b_n = 0$. The right side has order at least q there, so $q = 0$. Similarly $q' = 0$. Since $\langle \varpi_n, \theta^\vee \rangle = 1$, $X_{-\theta} X_\theta b_n = b_n$. The same pointed-coefficient calculation for $R_0(R_n^\star b_n) = c c' b_n$ gives $c c' = 1$, proving exactness of $(0, n)$.

The extra root is $-\theta$, and $H_0 = -H_\theta$. The finite roots together with this extra pair give $C_n^{(1)}$ and kill its canonical center. All the operators preserve degree. Perfectness and finite-type normality identify the finite-color crystal on $\mathcal{R}_m = L(m\varpi_n)^*$ with $B(m\varpi_n)$, since

$w_0 = -1$. Its unique highest-weight element corresponds to b_0^m. For $m > 0$, $R_0(b_0^m) \neq 0$, proving the final assertion. □

In rank two the Plücker relation reads $xy = b_1^2 + b_0 b_2$. Writing $E_i = R_i$ and $F_i = R^\star_{2-i}$, the six operators are

i	0	1	2
E_i	$y\partial_{b_0}$	$x\partial_{b_1}$	$y\partial_{b_2}$
F_i	$x\partial_{b_2}$	$y\partial_{b_1}$	$x\partial_{b_0}$

Each derivative is taken in the cluster containing its displayed multiplier. For example, E_0 fixes y, so differentiating the exchange relation gives $E_0 x = b_2$. Hence $E_0(F_0 b_2) = b_2$. The same relation gives $E_1 y = 2b_1$ and thus $E_1^2(F_1 b_1) = 2x$. The Cartan matrix in this order is $\left(\begin{smallmatrix} 2 & -1 & 0 \\ -2 & 2 & -2 \\ 0 & -1 & 2 \end{smallmatrix}\right)$.

Remark 8.9. For a comparison with Kirillov–Reshetikhin crystals one must also identify the affine arrows under a finite-color crystal isomorphism. Cluster structures for the corresponding Cartan algebras are studied in [14].

9. Theta F-pairings and maximal adjunction

The symmetric F-invariant records the Newton-polytope behavior of maximal raising. We also determine when maximal raising admits an adjoint for the valuation pairing $\mathcal{K}$. Principal-coefficient F-polynomials make both constructions available even when the extended exchange matrix has deficient rank; compare the categorical maximal-operator theory of [12, Appendix A].

The first two subsections require only finite principal-coefficient theta functions; the remaining three also assume reddening. We regard theta g-vectors and coefficient exponents as rows, and valuation directions as columns.

9.1. **Principal-coefficient F-polynomials.** At a seed t, adjoin auxiliary frozen variables z_k, $k \in U$, with extended column matrix

$$\widehat{\widetilde{B}}_t = \begin{pmatrix} \widetilde{B}_t \\ I_U \end{pmatrix}.$$

We invert these variables. In square row notation, a compatible extension of the exchange data is

$$\widehat{B}_t = \begin{pmatrix} B_t & -E_U \\ E_U^{\mathsf{T}} & 0 \end{pmatrix}, \qquad \widehat{D} = \operatorname{diag}(D, D_U),$$

where E_U is the inclusion of the mutable coordinates. The *principal-coefficient F-polynomial* of a finite theta function is defined by its expansion in this extension

$$\widehat{\vartheta}^t_{(g,0)} = x^g \mathcal{F}_{t,g}(\widehat{y}_k z_k : k \in U), \qquad \widehat{y}_k = x^{-B_{k,V}(t)}.$$

The auxiliary exponents determine $\mathcal{F}_{t,g}$ uniquely. It has nonnegative integral coefficients and constant term one. This extends the usual F-polynomial for cluster variables [18, Definition 3.3] to finite theta functions using their principal-coefficient expansions [21]. The qualifier "principal-coefficient" distinguishes these polynomials from the F-polynomials of cluster algebras with universal coefficients. In the projective-weight convention of Section 4, it is $\mathcal{F}_{t,g} = F_{-g_U}$.

Lemma 9.1. *Fix a principal extension at one seed. If a finite theta function has g-vector (g_t, a_t) in another seed t, its canonical F-polynomial is $\mathcal{F}_{t,g_t}$. Multiplication by a Laurent monomial in the original or auxiliary frozen variables leaves this polynomial unchanged.*

Proof. At t, the auxiliary exchange matrix is the principal c-matrix C_t. It is integral unimodular by sign coherence and the tropical seed identities [21, 32]. Frozen translation gives $\widehat{\vartheta}_{(g_t,a_t)} = z^{a_t}\widehat{\vartheta}_{(g_t,0)}$. To compare the latter function with the principal extension based

at t, make the monomial change $x^r z^b \longmapsto x^r u^{bC_t^{-\mathsf{T}}}$. It fixes the original variables and takes the incoming function $1 + x^{-B_{k,V}(t)} z^{e_k C_t^{\mathsf{T}}}$ to $1 + x^{-B_{k,V}(t)} u_k$. It also fixes the mutable wall normals. Uniqueness of consistent completion and mutation invariance of cluster scattering diagrams [21] therefore identify the two scattering diagrams and their theta functions. The change identifies their coefficient monoids, so the comparison holds before specialization. Frozen translation proves the last assertion. □

Lemma 9.2. *Every vertex of* $\mathrm{Newt}(\mathcal{F}_{t,g})$ *has coefficient one.*

Proof. By the monicity of the principal cluster scattering diagram, every vertex term of a finite principal theta function has coefficient one by [8, Theorem 3.15]. The map from coefficient exponents to full exponents is the affine injection $\alpha \longmapsto (g + \alpha \widetilde{B}_t^{\mathsf{T}}, \alpha)$. Since this injection preserves the Newton vertices and their coefficients, the claim follows. □

9.2. **The symmetric F-invariant.** Let $N_{t,g} = \mathrm{Newt}(\mathcal{F}_{t,g})$, let $h_{t,g}$ be its support function, and write $w_{t,g} = D_U(g_t)_U^{\mathsf{T}}$. We suppress the fixed symmetrizer D from the notation for the invariant below.

Proposition 9.3. *For theta functions with finite principal-coefficient lifts, the quantity*

$$\mathsf{F}(\vartheta_g, \vartheta_h) = h_{t,g}(w_{t,h}) + h_{t,h}(w_{t,g})$$

is independent of t. It is a symmetric nonnegative integer and is unchanged by frozen Laurent factors in either argument. When the extended exchange matrix has full column rank, it is Cao's F-invariant.

Proof. Lift both functions to the same principal extension. This extension admits a compatible skew form. With $K = (D_U \mid 0_F)$, we use

$$\Gamma = \begin{pmatrix} 0 & -K^{\mathsf{T}} \\ K & \widetilde{B}^{\mathsf{T}} K^{\mathsf{T}} \end{pmatrix}, \qquad (\widetilde{B}^{\mathsf{T}} \mid I_U)\Gamma = (D_U \mid 0_F \mid 0_U).$$

The lower right submatrix is skew-symmetric by skew-symmetrizability. The lifted theta functions are positive and compatibly pointed, so Cao's invariant theorem applies [6, Theorem 4.10]. Lemma 9.1 identifies his canonical polynomials in every seed with the principal-coefficient polynomials defined above. This proves seed independence. Since $0 \in N_{t,g}$, both support values are nonnegative. Integrality and invariance under frozen factors follow from the definition. When $\widetilde{B}$ has full column rank, its exponent map is injective, so the principal-coefficient polynomial is also the canonical polynomial in the original coefficients. □

Two theta functions are *F-compatible* when their F-invariant vanishes. Compatibility alone does not imply that their product is a single theta function.

9.3. **Maximal raising and Newton faces.** Assume reddening for the remainder of the section. Work in an i-sink optimizer, and write

$$m_U = -B_{i,U}, \qquad m_F = 0, \qquad v = m - e_i, \qquad q = D_U m_U^{\mathsf{T}}, \qquad J = \{k \in U : m_k > 0\}.$$

For $g \in \mathcal{B}_\theta^+$, maximal raising gives the i-highest element $P_i g = r_i^{\rho_i(g)} g$ of its string. In this seed we write N_g and h_g for the Newton polytope and its support function, and set $N_g^0 = N_g \cap q^\perp$ and $h_g^0 = h_{N_g^0}$.

Proposition 9.4. *In the optimizer,*

$$P_i g = g + g_i v, \qquad \mathcal{F}_{P_i g} = \mathcal{F}_g|_{Y_k = 0 \ (k \in J)}, \qquad N_{P_i g} = N_g^0.$$

The polynomial identity holds before coefficient specialization.

Proof. Base the principal extension at the optimizer. The lift $\widehat{\vartheta}_{(g,0)}$ is regular at every original frozen divisor by positivity, and at every auxiliary divisor because its auxiliary exponents are nonnegative. The maximal derivative formula of Corollary 3.9, applied in this extension, gives $R_i^{(g_i)}\widehat{\vartheta}_{(g,0)} = \widehat{\vartheta}_{(g+g_iv,0)}$. Skew-symmetrizability shows that a term of coefficient degree α has x_i-exponent $g_i - d_i^{-1}\alpha q$. The divided derivative selects exactly those terms for which $\alpha q = 0$, equivalently $\alpha_k = 0$ for $k \in J$. Comparing auxiliary exponents proves the polynomial identity. The exponent set lies in $\alpha q \geqslant 0$, so its convex hull meets $q^\perp$ in precisely the convex hull of the surviving terms. □

For two regular theta g-vectors, define

$$A_i(g;h) = h_g(w_h) - h_g^0(w_h), \qquad B_i(g;h) = h_h(w_g + g_i q) - h_h(w_g).$$

Corollary 9.5. *Maximal raising satisfies*

$$\begin{aligned} \mathsf{F}(P_i g, P_i h) &= \mathsf{F}(g,h) - A_i(g;h) - A_i(h;g), \\ \mathsf{F}(P_i g, h) &= \mathsf{F}(g,h) - A_i(g;h) + B_i(g;h), \\ 0 \leqslant B_i(g;h) &\leqslant g_i \mathsf{F}(M_i, \vartheta_h), \\ g_i &= \operatorname{ord}_{x_i}(\vartheta_g) + \frac{1}{d_i}\mathsf{F}(M_i, \vartheta_g). \end{aligned}$$

In particular, simultaneous maximal raising does not increase F. *Equality holds exactly when the* w_h*-maximizing face of* N_g *meets* N_g^0 *and the* w_g*-maximizing face of* N_h *meets* N_h^0. *The map* P_i *is idempotent, its image is* F*-compatible with* M_i*, and its fixed points satisfy* $\operatorname{ord}_{x_i}(\vartheta_g) = \mathsf{F}(M_i, \vartheta_g) = 0$. *Every finite composition of simultaneous maximal raising maps preserves* F*-compatibility.*

Proof. We have $w_{P_i g} = w_g + g_i q$ and $N_{P_i g} = N_g^0$. The restricted polytopes annihilate q, giving the first two identities and the equality criterion. Nonnegative coefficient exponents and subadditivity give $0 \leqslant B_i(g;h) \leqslant g_i h_h(q)$. Since M_i is a seed monomial, $h_h(q) = \mathsf{F}(M_i, \vartheta_h)$.

The smallest x_i-exponent of ϑ_g is $g_i - d_i^{-1} h_g(q)$, which proves the last identity. The new i-coordinate and the support value at q are both zero. This proves idempotence and the assertions about the image and fixed points. The final assertion follows by iteration. □

If $\mathsf{F}(M_i, \vartheta_h) = 0$, raising only the first argument also does not increase F. The change is then $A_i(g;h)$. The restricted support function can also be recovered by

$$h_g^0(w) = h_g(w - tq) \qquad \text{for all sufficiently large } t.$$

Indeed, each exponent α outside the face ceases to maximize once $t \geqslant (\alpha w - h_g^0(w))/(\alpha q)$, and there are only finitely many exponents.

An isolated zero exchange column still contributes to the coefficient-independent invariant. Its rank-one factor is $\vartheta_a = 2^{(-a)_+} x^a$, with $\mathcal{F}_a(Y) = (1+Y)^{(-a)_+}$, and its contribution to $\mathsf{F}(a,b)$ is $d_k((-a)_+(b)_+ + (-b)_+(a)_+)$. Mutation negates both g-vectors and preserves this value. Since $m_k = 0$, boundary raising leaves this factor unchanged.

Proposition 9.6. *There is an assignment* Q *to theta functions with finite principal lifts such that*

$$\mathsf{F}(P_i g, h) = \mathsf{F}(g, Qh) \qquad \textit{for all regular } g, h$$

if and only if $M_i = 1$. *In that case,* Q *may be the identity.*

Proof. Suppose $m_k > 0$. The adjacent cluster variable x_k' has principal-coefficient F-polynomial $1 + Y_k$, so $\mathsf{F}(M_i, x_k') = d_k m_k > 0$. But $\mathsf{F}(x_i, u) = 0$ for every possible target u, whereas $P_i(x_i) = M_i$. These two identities contradict the proposed adjunction. If $M_i = 1$, maximal raising removes an independent frozen factor, which leaves F unchanged. □

9.4. **Adjunction for the valuation pairing.** The valuation pairing $\mathcal{K}(g,n) = \mathrm{val}_n(\vartheta_g)$ need not have a global adjoint to maximal raising. We now describe the precise set of dual tropical directions on which such an adjoint exists. In an i-optimizing seed, the maximal-raising map $P_i g = g + g_i v$ is linear, with transpose

$$Q_i n = n + \langle v, n\rangle e_i^*.$$

Both maps are idempotent. This transpose is the only possible candidate for an adjoint: a valuation identity valid for all regular theta functions must in particular hold for each seed variable. The only changed coordinate is $(Q_i n)_i = \langle m, n\rangle$. Moreover, Q_i preserves the chiral-dual lattice $N^\bullet$: in its integral coordinates $z_a = (\ell/d_a)n_a$, it becomes

$$(\widehat{Q}_i z)_a = z_a \ (a \neq i), \qquad (\widehat{Q}_i z)_i = \sum_{k\in U} b_{ki} z_k.$$

Here $d_k m_k / d_i = b_{ki}$.

We write a superscript 0 for theta functions of the seed in which J is frozen; the exponent and valuation lattices are not changed. For every remaining mutable index k, we have $m_k = 0$ and $p_k(i) = 0$. Thus P_i fixes its exponent p_k, and Q_i fixes its normal q_k. Consequently P_i is an endomorphism of the frozen seed datum. The face formula gives $\vartheta_{P_i g} = x^{g_i v}\vartheta^0_g$. Adjunction for linear seed morphisms [8, Proposition 4.15], followed by reciprocity in the frozen datum, gives

$$\mathcal{K}(P_i g, n) = \mathrm{val}_{Q_i n}(\vartheta^0_g) = \mathrm{val}_g(\vartheta^{\vee,0}_{Q_i n}). \tag{9.1}$$

The frozen functions are finite, since freezing selects terms from finite positive expansions.

Let $\mathcal{S}_i = \mathrm{supp}(\vartheta_{-m})$ in the optimizing torus. Define

$$W_i(n) = -\langle m, n\rangle - \mathcal{K}(-m, Q_i n), \qquad \Omega_i = \{n : W_i(n) = 0\}.$$

Since $-m \in \mathcal{S}_i$ and $P_i m = m$, we have

$$W_i(n) = \max_{a\in\mathcal{S}_i}\big(-\langle a + m, Q_i n\rangle\big) \geqslant 0.$$

Lemma 9.7. *The equality $W_i(n) = 0$ holds if and only if $\vartheta^{\vee,0}_{Q_i n} = \vartheta^\vee_{Q_i n}$.*

Proof. For $q = Q_i n$, write the positive dual expansion as

$$\vartheta^\vee_q = \sum_{\alpha\in A_q} c_\alpha z^{q + \sum_{k\in U}\alpha_k \frac{d_k}{\ell} e_k^*}, \qquad c_\alpha > 0.$$

Reciprocity at $(-m, q)$ gives $W_i(n) = \max_{\alpha\in A_q}\sum_{k\in J}\frac{d_k m_k}{\ell}\alpha_k$. Each summand is nonnegative, and $m_k > 0$ exactly for $k \in J$. The maximum vanishes exactly when every contributing degree has $\alpha_J = 0$. These are precisely the degrees that survive freezing. The dual bend exponents are independent, so no deleted degree can coincide with a surviving one. □

Theorem 9.8 (Maximal adjunction)**.** *For $n \in N^\bullet$, there exists $q \in N^\bullet$ satisfying*

$$\mathcal{K}(P_i g, n) = \mathcal{K}(g, q) \qquad \text{for every } g \in \mathcal{B}^+_\theta$$

if and only if $n \in \Omega_i$. The target is uniquely $q = Q_i n$. The domain is the rational polyhedral cone

$$\Omega_i = \bigcap_{a\in\mathcal{S}_i}\{n : \langle a + m, Q_i n\rangle \geqslant 0\}.$$

The map Q_i retracts Ω_i onto $\Omega_i \cap v^\perp$. The same assertions hold for rational valuation directions.

Proof. Suppose first that $W_i(n) = 0$. By Lemma 9.7, freezing does not remove any term from $\vartheta^\vee_{Q_i n}$. Hence the two dual theta functions in (9.1) agree. Theta reciprocity in the original datum now yields $\mathcal{K}(P_i g, n) = \mathcal{K}(g, Q_i n)$ for every regular g.

Conversely, apply a proposed adjunction to the seed variables. Maximal raising fixes x_a for $a \neq i$ and sends x_i to x^m. Thus $q_a = n_a$ for $a \neq i$ and $q_i = \langle m, n\rangle$, which forces $q = Q_i n$.

Choose a nonnegative frozen vector f such that $\vartheta_\gamma = x^f \vartheta_{-m}$ is regular, where $\gamma = -m + f$. Such f exists by finiteness of ϑ_{-m} and invariance of frozen valuations. The vector $-m$ annihilates every normal of the frozen diagram, so $\vartheta^0_\gamma = x^\gamma$. Hence

$$\mathcal{K}(P_i\gamma, n) = -\langle m, n\rangle + \langle f, Q_i n\rangle,$$
$$\mathcal{K}(\gamma, Q_i n) = \mathcal{K}(-m, Q_i n) + \langle f, Q_i n\rangle.$$

The difference between these two valuations is $W_i(n)$. If the proposed adjunction holds for this single regular theta function, the difference must vanish. Hence $W_i(n) = 0$, proving necessity as well as the sufficiency of a single additional test beyond the seed variables.

The finite support formula gives the cone. Idempotence and $\langle m, Q_i n\rangle = \langle m, n\rangle$ imply $W_i(Q_i n) = W_i(n)$; the image of Q_i is $v^\perp$. For rational directions, clear denominators and use homogeneity of valuations. The seed-variable tests still prove uniqueness. □

For a fixed regular theta g-vector g, the difference

$$\Delta_i(g; n) = \mathcal{K}(P_i g, n) - \mathcal{K}(g, Q_i n)$$

is nonnegative. By reciprocity, it is the minimum over the frozen dual support minus the minimum over the full dual support, both evaluated at g. It is piecewise linear and homogeneous in g. The vector $\gamma = -m + f$ constructed in the proof satisfies $\Delta_i(\gamma; n) = W_i(n)$, so Ω_i is exactly the common zero set of these differences.

Corollary 9.9. *It suffices to test maximal adjunction on the seed variables and one regular theta function ϑ_γ, independent of n. This function may be chosen to be a cluster monomial. For a word $P_{i_r} \cdots P_{i_1}$, set $n_r = n$ and $n_{k-1} = Q_{i_k} n_k$. If $n_k \in \Omega_{i_k}$ at every step, then for every $g \in \mathcal{B}^+_\theta$,*

$$\mathcal{K}(P_{i_r} \cdots P_{i_1} g, n) = \mathcal{K}(g, n_0).$$

Proof. Reddening provides a seed with mutable variables I_k of g-vectors $-e_k + c_k$, with c_k frozen. Their pointing monomials occur in regular Laurent expansions, so $c_k \geqslant 0$. Let $c = \sum_k m_k c_k$. Frozen factorization gives $\vartheta_{-m} = x^{-c} \prod_k I_k^{m_k}$. Thus we may choose $f = c + e_i$ and $\vartheta_\gamma = x_i \prod_k I_k^{m_k}$. Its maximal image is $M_i^{c_i} \prod_{a \in F \setminus \{i\}} x_a^{c_a}$. For $q = Q_i n$, the obstruction becomes

$$W_i(n) = \sum_{k \in J} m_k \big(\langle -e_k + c_k, q\rangle - \mathrm{val}_q(I_k)\big).$$

Each term is nonnegative because the pointing monomial occurs. The valuations can be computed from the exchange relations. Successive application of Theorem 9.8 proves the assertion for a word. □

The adjunction characterizes Ω_i and Q_i intrinsically. The regular seed variables separate dual tropical points, so uniqueness identifies the adjoints computed in different optimizing seeds. Let n_i^- be the tropical point with coordinate $-e_i^*$ in an i-sink optimizer. Its valuation on a nonzero element of $\mathcal{U}^+$ is $-\deg_{R_i}$. Indeed, in this chart the nilpotence degree is the ordinary polynomial degree in x_i. Taking differences of degrees on quotients extends the valuation to the function field, independently of the optimizer. The kth string coordinate is therefore $-\mathcal{K}(P_{i_{k-1}} \cdots P_{i_1} g, n_{i_k}^-)$. Under the successive domain conditions of Corollary 9.9, this is a pairing of g with a single transported dual point. Compare the categorical transfer of string data in [12, Appendix A.2].

9.5. **The tropical invariant.** Suppose the original exponent lattice admits a rational skew form Γ with $\widetilde{B}^{\mathsf{T}}\Gamma = (D_U \mid 0_F)$. Cao's *tropical invariant*, in our row convention, is

$$\mathsf{A}_\Gamma(g, h) = g\Gamma h^{\mathsf{T}} + h_g(D_U h_U^{\mathsf{T}}) = -\mathcal{K}(g, -\Gamma h^{\mathsf{T}}).$$

The second equality follows by expanding the valuation in the principal-coefficient F-polynomial. Theorem 9.8 therefore gives, for an integral target h', the criterion

$$\mathsf{A}_\Gamma(P_i g, h) = \mathsf{A}_\Gamma(g, h') \quad \text{for every regular } g$$

$$\iff \quad -\Gamma h^{\mathsf{T}} \in \Omega_i \quad \text{and} \quad \Gamma (h')^{\mathsf{T}} = Q_i \Gamma h^{\mathsf{T}}.$$

If Γ is invertible, the only possible target satisfies $(h')^{\mathsf{T}} = \Gamma^{-1} Q_i \Gamma h^{\mathsf{T}}$. If Γ is singular, the integral solutions are either empty or a coset of the integral kernel of Γ. In either case, clearing denominators and using Smith normal form tests integrality.

10. Maximal operators and string coordinates

10.1. **String data and reconstruction.** In this subsection we assume mutable reddening when using the crystal operators on all regular theta functions, as in Section 3. The piecewise-linear maps attached to fixed optimizing words are defined on all of $\mathbb{R}^V$ in the reference seed.

For a finite word $\mathbf{i} = (i_1, \dots, i_r)$ of sink-reachable colors and $g \in \mathcal{B}_\theta^+$, define

$$g^{(0)} = g, \qquad n_k = \rho_{i_k}(g^{(k-1)}), \qquad g^{(k)} = P_{i_k}(g^{(k-1)}).$$

We call $\mathbf{s_i}(g) = (n_1, \dots, n_r)$ the *string datum* of g for $\mathbf{i}$. In general it must be accompanied by the terminal g-vector $g^{(r)}$.

Lemma 10.1. *The map $\Psi_{\mathbf{i}} : g \longmapsto (\mathbf{s_i}(g), g^{(r)})$ is injective. Its inverse applies $l_{i_r}^{n_r}, \dots, l_{i_1}^{n_1}$ in that order. Moreover, $R_{i_r}^{(n_r)} \cdots R_{i_1}^{(n_1)} \vartheta_g = \vartheta_{g^{(r)}}$.*

Proof. The inverse follows from the partial inversion of r_i and l_i, as in [12, Appendix A.2]. The identity follows by successive applications of Corollary 3.9. □

A string tuple need not determine its initial g-vector without the final g-vector; compare [20].

Proposition 10.2. *Fix optimizing words for the colors in $\mathbf{i}$. Then $\Psi_{\mathbf{i}}$ extends to a homogeneous piecewise integral-linear embedding $\mathbb{R}^V \longrightarrow \mathbb{R}^r \times \mathbb{R}^V$ with a piecewise integral-linear inverse on its image. Under mutable reddening, the image of the real regular cone $\mathcal{C}^+$ is a finite union of rational polyhedral cones. Each linearity cone can be subdivided so that its image contains exactly the images of all its lattice points. On each such cone, the integral image is a saturated finitely generated additive monoid.*

Proof. For one step, write $\Phi_i = \Phi_{\boldsymbol{\mu}_i^+}$ for the chosen tropical transport. The extension sends g to (n, p), where

$$a = \Phi_i(g), \qquad n = a_i, \qquad p = \Phi_i^{-1}(a + n v_i).$$

Its inverse is $(n, p) \mapsto \Phi_i^{-1}(\Phi_i(p) - n v_i)$. These are homogeneous piecewise integral-linear maps with finitely many linearity regions. Composition proves the first assertion.

By Corollary 3.2, $\mathcal{C}^+$ is rational polyhedral. Refine it by the linearity regions of the map and its inverse. On a resulting cone K, let A and L be the integral matrices of the map and inverse. If $y \in A(K)$ is integral, then $Ly \in K$ is integral and $A(Ly) = y$. Thus $A(K \cap \mathbb{Z}^V) = A(K) \cap \mathbb{Z}^{r+|V|}$. Gordan's lemma gives finite generation. The integral points of a rational cone form a saturated monoid in the group they generate. □

10.2. **Cluster monomials.** We next prove that maximal raising preserves cluster monomials with nonnegative frozen exponents. This is the theta analogue of the orthogonalization result in [12, Appendix A.1]. We use Qin's freezing theorem and Laurent regularity. Here an *ordinary cluster monomial* has nonnegative exponents on both the mutable and the original frozen variables of one seed.

Lemma 10.3. *Fix an i-sink seed and let $J = \{u \in U : b_{iu} < 0\}$. For each seed t, there is a seed s, reached from the optimizer by mutations outside J, such that the highest x_i-degree part of every cluster monomial of t is a localized monomial of s. Its exponents at mutable indices outside J are nonnegative.*

Proof. In a principal-coefficient expansion $x^g \sum_{\nu \geqslant 0} c_\nu x^{-\nu B_{U,V}}$, the terms of highest x_i-degree are exactly those with $\nu_J = 0$. Qin's freezing theorem sends this sum to a localized cluster monomial of the J-frozen pattern [35, Proposition 5.5 and Theorem 5.11]. The chamber of t lies in one reachable chamber of the frozen pattern, so a single seed s works for all its monomials [35, Lemmas 5.3 and 5.6]. Theorem 5.11 is formulated using principal-coefficient F-polynomials and also applies in deficient rank. For the common-seed assertion in that case, apply the chamber argument in the principal extension and then specialize. □

Theorem 10.4. *Let $R_i = M_i \partial_{x_i}$ be a canonical sink derivation. If z is an ordinary cluster monomial and $d = \deg_{R_i} z$, then $R_i^{(d)} z$ is an ordinary cluster monomial. For each seed t, all maximal images of its monomials lie in one seed s, reachable from an i-optimizer without mutating its mutable neighbors. On the monomial semigroup of t, maximal raising is multiplicative and has a nonnegative integral exponent matrix.*

These assertions require neither reddening nor theta spanning. They also hold for canonical source derivations.

Proof. In the optimizer, write $M_i = x^{m_i}$ and let $J = \{u : b_{iu} < 0\}$. If q_z is the highest x_i-degree part of z, then $R_i^{(d)} z = x^{d(m_i - e_i)} q_z$. Choose the common seed s from Lemma 10.3. The multiplier is supported on indices frozen in the J-frozen pattern. Thus the right side is a localized monomial $x(s)^c$, with $c_u \geqslant 0$ for $u \in U \backslash J$.

To prove $c_u \geqslant 0$ for $u \in J$, note that mutations outside J leave the ith exchange entries unchanged. Hence $b_{iu}(s) < 0$. The exchange polynomial P_u at u has two different x_i-exponents and is a nonunit in the Laurent ring in the other variables. If $c_u = -r < 0$, substitution $x_u = P_u / x_u'$ would give the nonunit denominator P_u^r and a Laurent monomial numerator in the adjacent seed. This contradicts $R_i^{(d)} z \in \mathcal{U}^+$. Regularity at the original frozen divisors gives $c_f \geqslant 0$ for $f \in F$.

The common seed works for the entire source cluster. Degrees add and highest-degree parts multiply, so maximal raising is multiplicative. Its exponent matrix consists of the nonnegative exponent vectors of the images of the source variables. Reversing the exchange matrix proves the source assertion. □

Corollary 10.5. *For a fixed word of boundary operators and a fixed source cluster, all successive maximal images are cluster monomials. Their string coordinates are integral linear functions of the nonnegative source exponents. A common sequence of target seeds works for the entire source cluster cone.*

Proof. Iterate Theorem 10.4. At each step, degree is additive on the monomial semigroup, and the exponent matrices compose. □

Remark 10.6 (One-seed string coordinates). Let A be a characteristic-zero algebra with regular theta basis, stable under the boundary derivations for a word $\mathbf{i} = (i_1, \dots, i_r)$. Given a character $\epsilon : A \to \Bbbk$, put

$$\Phi_{\mathbf{i}}^*(f) = \epsilon\big(\exp(z_r R_{i_r}) \cdots \exp(z_1 R_{i_1}) f\big).$$

Suppose the variables of a seed t belong to A and satisfy $\Phi_{\mathbf{i}}^*(x_v) = z^{\mathbf{d}_v}$ with $\mathbf{d}_v \in \mathbb{N}^r$. Let M have rows $\mathbf{d}_v$ and assume $-B_{u,V}(t)\mathsf{M} \leqslant_{\mathrm{lex}} 0$ for every mutable u, comparing the first coordinate first. For every $\vartheta_g \in A$ with $\epsilon(\vartheta_{g^{(r)}}) \neq 0$, one then has $\mathbf{s}_{\mathbf{i}}(g) = g(t)\mathsf{M}$. Indeed, successive maximal divided derivatives give the leading exponent $\mathbf{s}_{\mathbf{i}}(g)$ with coefficient $\epsilon(\vartheta_{g^{(r)}})$. In the pointed expansion, every exchange monomial has nonpositive lexicographic exponent after pullback. The constant term of the F-polynomial survives by positivity, so the same leading exponent is $g(t)\mathsf{M}$. Any integral right inverse of M thus recovers $g(t)$ from the string, in particular when M is square and unimodular. For the unipotent factorization charts, see [12, Appendix A.3].

10.3. **Tropical formulas for boundary operators.** Given optimizing mutation sequences, we can compute the boundary monomials and crystal operators by tropical mutation. Let $\boldsymbol{\mu}_i^+$ be a sink word and $\boldsymbol{\mu}_k^-$ a source word. Their terminal exchange matrices determine the exponent vectors m_i and n_k of the boundary monomials.

Proposition 10.7. *For two sink colors $i \neq j$, the Cartan entry is*

$$c_{ij} = -\big(\Phi_{\boldsymbol{\mu}_i^+}\Phi_{\boldsymbol{\mu}_j^+}^{-1} m_j\big)_i.$$

The sink i and source k form an exact pair if and only if

$$\Phi_{\boldsymbol{\mu}_i^+}\Phi_{\boldsymbol{\mu}_k^-}^{-1} n_k = e_i + e_k - m_i.$$

Proof. The first expression is the full g-vector of M_j in the i-sink optimizer. Its ith coordinate is the maximum x_i-degree, hence $\deg_{R_i} M_j = -c_{ij}$. The second expression is the full g-vector of $M_k^\star$ in that optimizer; the assertion is the exactness criterion of Section 4.3. □

For degrees, the tropical exchange relation is different from the tropical mutation rule for g-vectors. If q_a is the maximum degree of x_a, then mutation at k gives

$$q_k' = \max\left\{\sum_a [b_{ka}]_+ q_a, \sum_a [-b_{ka}]_+ q_a\right\} - q_k, \qquad q_a' = q_a \ (a \neq k).$$

Starting with degree vector e_i at the sink optimizer, transport it to the reference seed by this rule and call the result α_i. For an exact partner, similarly transport $e_{\sigma(i)}$ from the source optimizer and call the result β_i.

Corollary 10.8. *The weight of the reference seed variables under $H_i = [R_i, R^\star_{\sigma(i)}]$ is the integral vector $w_i = \beta_i - \alpha_i$.*

Proof. The entries of α_i and β_i are the raising and lowering lengths of the seed variables. Each seed variable is homogeneous for H_i. In a locally finite $\mathfrak{sl}_2$-module, the lowering length minus the raising length of a homogeneous vector is its weight. Indeed, in its finite-dimensional cyclic submodule, both lengths are largest on the summands of largest highest weight. This proves the formula. □

We regard w_i as the compatible linear form $w_i(g) = gw_i^{\mathsf{T}}$ in the reference seed. For $g \in \mathcal{B}_\theta^+$, write $a = \Phi_{\boldsymbol{\mu}_i^+} g$. Then

$$r_i g = \Phi_{\boldsymbol{\mu}_i^+}^{-1}(a + v_i), \qquad P_i g = \Phi_{\boldsymbol{\mu}_i^+}^{-1}(a + a_i v_i),$$

where the first map is defined exactly when $a_i > 0$. Before exact pairing, the inverse candidate is $\Phi_{\boldsymbol{\mu}_i^+}^{-1}(a - v_i)$ and exists exactly when it is regular. The boundary Newton polytopes give finite regularity tests and the inverse-string length through Lemma 3.3 and Proposition 3.6. If every frozen vertex has a source optimizer, Lemma 3.7 gives an alternative test using only tropical transport.

For exact partners, $\rho_i(g) = a_i$ and $\lambda_i(g) = a_i + w_i(g)$. The lowering neighbor exists exactly when $\lambda_i(g) > 0$, and its g-vector is the same candidate for the inverse. The string reflection is $s_i(g) = \Phi_{\boldsymbol{\mu}_i^+}^{-1}\big(a - w_i(g)v_i\big)$. This is the Weyl action of Section 7.4.1. Thus the optimizing sequences determine each raising step and the i-highest element. For exact partners they also determine each lowering step, the i-lowest element $l_i^{\lambda_i(g)} g$, and the reflection of the string. Cartan entries and exactness require no reddening hypothesis; regularity and strings require the hypotheses of the results used to define them. Finding such sequences is a separate reachability problem.

11. Pointed biperfect bases and transitions

Let $A = \mathrm{Mid}^+$, with theta basis $\{\vartheta_g : g \in \mathcal{B}\}$. Fix normalized perfect boundary derivations D_a, indexed by $a \in \mathcal{I}$, together with the weight gradings to be preserved. The indices may correspond to canonical derivations at sink- or source-reachable frozen vertices or to R_i, L_i for an exact system. Write (r_a, l_a, ρ_a) for the corresponding partial maps and upper string lengths. The zero g-vector indexes the unit, whereas 0_{cr} means that a crystal operator is undefined.

The formal criteria for changes of perfect bases are established in [12, Section 9 and Appendix B]. Here we determine the admissible support of a transition matrix after coefficient specialization, allowing deficient extended rank. We then impose linearity over the polynomial coefficient ring. The same results apply to any theta-spanned algebra stable under the chosen operators.

11.1. **Dominance after coefficient specialization.** Assume reddening. We use a principal coefficient extension as our fixed full-rank extension; when the original matrix has full column rank, we may instead use the identity extension. For the principal extension of $L = \mathbb{Z}^V$, write $\widehat{L} = L \oplus \mathbb{Z}^U$ and auxiliary frozen variables z. Frozen factorization and specialization give

$$\widehat{\vartheta}_{(g,c)} = z^c \widehat{\vartheta}_{(g,0)}, \qquad \mathrm{sp}\, \widehat{\vartheta}_{(g,c)} = \vartheta_g.$$

The lifted theta functions form a basis of the localized upper algebra of the extension. Their specializations form the localized middle basis in the original coefficients.

At a seed t, define strict dominance by

$$\widehat{h} \prec_t \widehat{g} \quad \Longleftrightarrow \quad \widehat{h} = \widehat{g} + n\widehat{\widetilde{B}}_t^{\mathsf{T}} \quad \text{for some } n \in \mathbb{N}^U \backslash \{0\}.$$

Write $\widehat{h} \prec \widehat{g}$ if this holds in every seed, in the corresponding tropical coordinates. The set $\widehat{D}_g = \{\widehat{h} : \widehat{h} \prec (g,0)\}$ is finite: it lies in the lower set for an initial seed and an injective-related seed, which is finite by [34, Lemma 5.1.1].

Definition 11.1. For $g, h \in L$, write $h \lhd g$ if $(h,c) \prec (g,0)$ for some $c \in \mathbb{Z}^U$, and define $E_g = \{h \in \mathcal{B} : h \lhd g\}$. A family $(b_g)_{g\in\mathcal{B}}$ *admits compatible pointed lifts* if every b_g is the specialization of an upper element compatibly pointed at $(g,0)$ in the fixed coefficient extension.

Lemma 11.2. *The relation $\lhd$ is a strict partial order with finite lower sets. For a fixed global coefficient extension, it is independent of the choice of initial seed.*

Proof. Finiteness follows by projecting $\widehat{D}_g$. Auxiliary frozen translation commutes with tropical mutation and preserves dominance. Thus comparisons $(h,c) \prec (g,0)$ and $(k,d) \prec (h,0)$ imply $(k, c+d) \prec (h,c) \prec (g,0)$, proving transitivity. If $(g,c) \prec (g,0)$, then $c \neq 0$ by strict dominance in the extension. Translation and transitivity would give $(g,nc) \prec (g,0)$ for every $n \geqslant 1$, contradicting finiteness of $\widehat{D}_g$. This proves irreflexivity.

In a different initial seed t, the tropical point represented by $(g,0)$ has coordinates (g_t, c_t). Subtracting $(0, c_t)$ from each comparison preserves dominance and normalizes the leading auxiliary coordinate to zero. Projection commutes with tropical mutation, so it defines the same relation on the original tropical points. □

Theorem 11.3. *A family in A admits compatible pointed lifts if and only if*

$$b_g = \vartheta_g + \sum_{h\in E_g} T_{g,h}\vartheta_h, \qquad T_{g,h} \in \Bbbk.$$

Every such family is a basis. Its transition matrix and inverse have diagonal one, finite support for each fixed first index, and off-diagonal support on $\lhd$.

If the original extended matrix has full column rank, the identity extension gives the class of all compatibly pointed bases indexed by $\mathcal{B}$.

Proof. Expand a pointed lift in the theta basis of the extension. The finite pointed-decomposition theorem [34, Lemma 3.4.9 and Theorem 5.1.3] gives $\widehat{b}_g = \widehat{\vartheta}_{(g,0)} + \sum_{(h,c)\in\widehat{D}_g} u_{h,c}\widehat{\vartheta}_{(h,c)}$. Each $(h,c) \in \widehat{D}_g$ projects to $h \lhd g$, so $h \neq g$ by Lemma 11.2. After specialization, collecting the coefficients of each ϑ_h gives the required formula. Since $b_g \in A$, independence of the localized theta functions makes all coefficients outside $\mathcal{B}$ vanish.

Conversely, choose for each $h \in E_g$ a witness $c(g,h)$ with $(h, c(g,h)) \prec (g,0)$. Then

$$\widehat{b}_g = \widehat{\vartheta}_{(g,0)} + \sum_{h\in E_g} T_{g,h}\widehat{\vartheta}_{(h,c(g,h))}$$

is a compatible pointed lift. The sum is finite.

Write the transition operator as $1 + N$. Every iterate of N on ϑ_g is supported on E_g, and every nonzero step strictly decreases $\lhd$. Since E_g is finite, N is locally nilpotent. The finite geometric series on each basis vector gives the inverse and proves the basis assertion. □

Let P_g be the image of $\widehat{D}_g$ under projection to the original g-vector lattice.

Proposition 11.4. *A finite theta combination f admits a compatible pointed lift at $(g,0)$ if and only if*

$$f \in \vartheta_g + \bigoplus_{h\in P_g} \Bbbk\vartheta_h.$$

For such f, the space of lifts is affine of dimension $|\widehat{D}_g| - |P_g|$. Moreover, there is a finite set of seeds, depending on g, at which pointedness of an upper element of the extension implies compatible pointedness in every seed.

Proof. The space of lifts is $\widehat{\vartheta}_{(g,0)} + \bigoplus_{\widehat{h}\in\widehat{D}_g} \Bbbk\widehat{\vartheta}_{\widehat{h}}$. Specialization sums the coefficients in each fiber over P_g. The resulting surjective linear map has kernel of dimension $|\widehat{D}_g| - |P_g|$, proving the first two assertions.

Start with the finite lower set for the initial and injective-related seeds. For every candidate outside $\widehat{D}_g$, choose a seed where its strict dominance fails. These finitely many additional seeds suffice by uniqueness of pointed decomposition. This is an existence assertion about the test set; it gives no bound on the required mutation words. □

11.2. **Perfectness and transition parameters.** For $n \geqslant 0$, let $K_{a,n} = \ker D_a^{n+1}$, and set $K_{a,n} = 0$ for $n < 0$. Perfectness gives $K_{a,n} = \operatorname{span}_{\Bbbk}\{\vartheta_g : \rho_a(g) \leqslant n\}$. The derivation induces a map $K_{a,n}/K_{a,n-1} \to K_{a,n-1}/K_{a,n-2}$.

We use the transition criterion of [12, Proposition 9.1 and Appendix B.1]: a linear automorphism fixing the unit and preserving the chosen gradings gives a normalized perfect basis with the same crystal data exactly when it preserves all $K_{a,n}$ and its associated graded maps commute with those of D_a. When at least one color is chosen, finite weight-preserving corrections of strictly smaller ρ_a for every chosen color preserve normalized perfectness [12, Corollary 9.2].

For $g \neq h$, call (g,h) *admissible* if $g \neq 0$, $h \in E_g$, the prescribed weights agree, and $\rho_a(h) \leqslant \rho_a(g)$ for every $a \in \mathcal{I}$. Set $T_{g,g} = 1$ and entries involving 0_{cr} equal to zero.

Proposition 11.5. *A basis in Theorem 11.3 fixes the unit, preserves the chosen gradings, and is normalized perfect for all prescribed crystal data if and only if its off-diagonal coefficients vanish on inadmissible pairs and satisfy*

$$T_{r_a g,h} = T_{g,l_a h} \qquad \text{if } \rho_a(g) > 0, \quad \rho_a(h) = \rho_a(g) - 1. \tag{11.1}$$

Proof. Apply the transition criterion of [12, Theorem 9.10 and Appendix B.1], with the finite support sets E_g from Theorem 11.3. Admissibility expresses preservation of weights and kernel filtrations. Equation (11.1) expresses commutation on the associated graded spaces. □

To solve these equations, form a graph on the off-diagonal pairs in $\mathcal{B}^2$, with edges

$$(g,h) \longleftrightarrow (r_a g, r_a h) \qquad \text{when } \rho_a(g) = \rho_a(h) > 0.$$

Mark every inadmissible pair. Also mark (g,h) if some a satisfies $\rho_a(g) = \rho_a(h)$, $l_a g \neq 0_{\mathrm{cr}}$, and $l_a h = 0_{\mathrm{cr}}$. Let $\mathfrak{C}$ be the connected components with no marked vertex.

Corollary 11.6. *The normalized biperfect bases with compatible pointed lifts and the fixed crystal data and gradings are parametrized by $\prod_{C\in\mathfrak{C}} \Bbbk$. For parameters (t_C), set $T_{g,h} = t_C$ when $(g,h) \in C$, and set every other off-diagonal coefficient to zero. The diagonal coefficients are one, and each basis vector is a finite sum. The theta basis is unique in this class exactly when $\mathfrak{C} = \varnothing$.*

Proof. The graph description of [12, Section 9.3] applies to (11.1): coefficients are constant on connected components and vanish on components containing a marked vertex. Each E_g is finite, so every parameter assignment defines the required basis. □

Corollary 11.7. *For a finite subset G of the theta g-vectors, the possible restrictions $(b_g)_{g\in G}$ form an affine space with one coordinate for each unmarked component meeting a pair whose first entry lies in G. If a group preserves the theta basis, projected order, gradings, and crystal data up to permutation of the operator indices, its equivariant bases have one parameter per orbit on $\mathfrak{C}$. Every localized cluster monomial in A is fixed by every basis with compatible pointed lifts.*

Proof. The first assertion follows from finiteness of the sets E_g. A symmetry preserves graph edges and marks; equivariance makes the parameters constant on its orbits. For the final assertion, lift the cluster monomial to the extension and normalize its auxiliary degree to zero. Qin's uniqueness lemma [34, Lemma 3.4.12] identifies every compatibly pointed lift of that degree with this monomial. Specialize to obtain the claim. □

In particular, $b_g = \vartheta_g$ in every such basis if and only if no unmarked component meets a pair with first entry g. Finite support for each first index does not force the graph components to be finite. A path to a marked pair proves that a coefficient vanishes; proving that a parameter is free may require control of an entire component. An admissible pair with $\rho_a(h) < \rho_a(g)$ for every a is an isolated unmarked vertex and gives an independent parameter.

11.3. **Polynomial coefficient linearity.** We may also require the transition map Ψ to be $\Bbbk[x_F]$-linear. Frozen factorization makes this equivalent to $b_{g+e_j} = x_j b_g$ for $g \in \mathcal{B}, j \in F$. This is the polynomial version of the coefficient-factorization condition in [34, Remark 5.1.4].

Corollary 11.8. *Add to the preceding graph the edges $(g,h) \longleftrightarrow (g+e_j, h+e_j)$ for $j \in F$. Also mark (g,h) whenever $g - e_j \in \mathcal{B}$ but $h - e_j \notin \mathcal{B}$ for some $j \in F$. Then the $\Bbbk[x_F]$-linear normalized biperfect transitions with compatible pointed lifts have one parameter per unmarked connected component of the enlarged graph.*

Proof. Expanding $b_{g+e_j} = x_j b_g$ in the theta basis and comparing the coefficients of ϑ_{h+e_j} gives $T_{g+e_j,h+e_j} = T_{g,h}$, the equality imposed by the new edge. For $k \in \mathcal{B}\backslash(\mathcal{B}+e_j)$, the coefficient of ϑ_k in $x_j b_g$ is zero, so $T_{g+e_j,k} = 0$ at the marked pair. Conversely, these conditions give equality of every theta coefficient on the two sides. Thus Ψ commutes with multiplication by each x_j and is $\Bbbk[x_F]$-linear. The finite support condition remains $h \in E_g$. □

The string-profile restrictions on these transitions, including lexicographic triangularity and the partial order obtained when every crystal element can be raised to the unit, follow directly from [12, Appendix B.2]. They require no pointedness or exchange-rank hypothesis.

Declaration of generative AI and AI-assisted technologies

During the development of the paper, the author used ChatGPT 6.0 to assist in exploring and formulating proof arguments. The author subsequently checked, revised, and independently verified all arguments and takes full responsibility for the mathematical content.

References

[1] M. F. Atiyah, *Complex analytic connections in fibre bundles*, Trans. Amer. Math. Soc. **85** (1957), 181–207.

[2] A. Berenstein, S. Fomin, and A. Zelevinsky, *Cluster algebras III: Upper bounds and double Bruhat cells*, Duke Math. J. **126** (2005), 1–52. https://arxiv.org/abs/math/0305434.

[3] A. Berenstein and D. Kazhdan, *Geometric and unipotent crystals II: From unipotent bicrystals to crystal bases*, in Quantum Groups, Contemp. Math. **433** (2007), 13–88.

[4] R. Bott, *Homogeneous vector bundles*, Ann. of Math. (2) **66** (1957), 203–248.

[5] D. Bump and A. Schilling, *Crystal Bases: Representations and Combinatorics*, World Scientific, 2017.

[6] P. Cao, *F-invariant in cluster algebras*, arXiv:2306.11438v4 (2025). https://arxiv.org/abs/2306.11438v4.

[7] L. Carbone, T. Coelho, S. H. Murray, F. Thurman, and S. Zhu, *Growth of root multiplicities along imaginary root strings in Kac–Moody algebras*, Comm. Algebra **54** (2026), 909–924. https://arxiv.org/abs/2403.01687.

[8] M.-W. Cheung, T. Magee, T. Mandel, and G. Muller, *Valuative independence and cluster theta reciprocity*, arXiv:2505.09585v2 (2026). https://arxiv.org/abs/2505.09585v2.

[9] H. Derksen and J. Fei, *General presentations of algebras*, Adv. Math. **278** (2015), 210–237. https://doi.org/10.1016/j.aim.2015.03.012.

[10] H. Derksen, J. Weyman, and A. Zelevinsky, *Quivers with potentials and their representations I: Mutations*, Selecta Math. (N.S.) **14** (2008), 59–119. https://arxiv.org/abs/0704.0649.

[11] J. Fei, *Tropical F-polynomials and general presentations*, J. Lond. Math. Soc. (2) **107** (2023), 2079–2120. https://arxiv.org/abs/1911.10513.

[12] J. Fei, *Crystal Structure of Upper Cluster Algebras*, arXiv:2309.08326v4 (2026). https://arxiv.org/abs/2309.08326v4.

[13] J. Fei, *Schubert Cox rings and root operators*, Hexagon 2611.00668v1 (2026). https://hexagonmath.org/2611.00668v1.

[14] J. Fei, *Cluster Realization of KR Cartan algebras*, manuscript, 2026.

[15] J. Fei, *Ramified Quivers with Potentials, Monodromy Characters, and Cluster Bases I*, arXiv:2610.09868v2 (2026). https://arxiv.org/abs/2610.09868v2.

[16] S. Fomin and A. Zelevinsky, *Double Bruhat cells and total positivity*, J. Amer. Math. Soc. **12** (1999), 335–380. https://arxiv.org/abs/math/9802056.

[17] S. Fomin and A. Zelevinsky, *Cluster algebras I: Foundations*, J. Amer. Math. Soc. **15** (2002), 497–529. https://arxiv.org/abs/math/0104151.

[18] S. Fomin and A. Zelevinsky, *Cluster algebras IV: Coefficients*, Compos. Math. **143** (2007), 112–164. https://arxiv.org/abs/math/0602259.

[19] L. Francone, *The minimal monomial lifting of cluster algebras I: Branching problems*, Adv. Math. **500** (2026), 111102. https://doi.org/10.1016/j.aim.2026.111102.

[20] V. Genz, G. Koshevoy, and B. Schumann, *Polyhedral parametrizations of canonical bases & cluster duality*, Adv. Math. **369** (2020), 107178. https://doi.org/10.1016/j.aim.2020.107178.

[21] M. Gross, P. Hacking, S. Keel, and M. Kontsevich, *Canonical bases for cluster algebras*, J. Amer. Math. Soc. **31** (2018), 497–608. https://doi.org/10.1090/jams/890.

[22] V. G. Kac, *Infinite-Dimensional Lie Algebras*, 3rd ed., Cambridge University Press, 1990.

[23] V. G. Kac and D. H. Peterson, *Defining relations of certain infinite dimensional groups*, Astérisque, no. S131 (1985), 165–208. https://numdam.org/item/AST_1985__S131__165_0/.

[24] J. Kamnitzer, *Perfect bases in representation theory: Three mountains and their springs*, Proc. Int. Cong. Math. 2022, vol. 4, 2976–2996. https://doi.org/10.4171/ICM2022/132.

[25] M. Kashiwara, *On crystal bases of the q-analogue of universal enveloping algebras*, Duke Math. J. **63** (1991), 465–516.

[26] B. Keller, with an appendix by L. Demonet, *A survey on maximal green sequences*, in Representation Theory and Beyond, Contemp. Math. **758** (2020), 267–286. https://arxiv.org/abs/1904.09247.

[27] Y. Kimura, F. Qin, and Q. Wei, *Twist automorphisms and Poisson structures*, SIGMA **19** (2023), 105, 39 pp. https://sigma-journal.com/2023/105/sigma23-105.pdf.

[28] C. Lenart, *On the combinatorics of crystal graphs, I. Lusztig's involution*, Adv. Math. **211** (2007), 204–243. https://doi.org/10.1016/j.aim.2006.08.002.

[29] G. Lusztig, *Canonical bases arising from quantized enveloping algebras*, J. Amer. Math. Soc. **3** (1990), 447–498.

[30] S. Morier-Genoud, *Geometric lifting of the canonical basis and semitoric degenerations of Richardson varieties*, Trans. Amer. Math. Soc. **360** (2008), no. 1, 215–235. https://arxiv.org/abs/math/0504538.

[31] G. Muller, $\mathcal{A} = \mathcal{U}$ *for locally acyclic cluster algebras*, SIGMA **10** (2014), 094, 8 pp. https://sigma-journal.com/2014/094/sigma14-094.pdf.

[32] T. Nakanishi and A. Zelevinsky, *On tropical dualities in cluster algebras*, Contemp. Math. **565** (2012), 217–226. https://arxiv.org/abs/1101.3736.

[33] F. Qin, *Dual canonical bases and quantum cluster algebras*, arXiv:2003.13674v3 (2023). https://arxiv.org/abs/2003.13674v3.
[34] F. Qin, *Bases for upper cluster algebras and tropical points*, J. Eur. Math. Soc. **26** (2024), 1255–1312. https://doi.org/10.4171/JEMS/1308.
[35] F. Qin, *Applications of the freezing operators on cluster algebras*, J. Algebra Appl. **25** (2026), 2550357, 30 pp. https://arxiv.org/abs/2407.03186.
[36] Q. Tang, *Biperfectness of the theta basis of a maximal unipotent subgroup*, Hexagon 2610.00603v1 (2026). https://hexagonmath.org/2610.00603v1.

School of Mathematical Sciences, Shanghai Jiao Tong University, China
Email address: jiarui@sjtu.edu.cn