hexagon:2611.00790v3

CRYSTAL STRUCTURES ON THETA BASES

JIARUI FEI

Abstract. The exchange matrix and its symmetrizer of a skew-symmetrizable seed determine two families of derivations, indexed by the frozen vertices that become sinks or sources relative to the mutable vertices. In characteristic zero they are locally nilpotent on the upper cluster algebra with polynomial frozen variables, and each family satisfies the Serre relations. We characterize the *exact pairs* whose derivations form an $\mathfrak{sl}_2$-triple with their commutator. If every chosen sink has an exact source partner, the derivations give an integrable action of a derived Kac–Moody algebra. Each Cartan component has finite or affine type, and the affine components act at level zero. Under mutable reddening, the theta basis of the regular middle algebra is normalized perfect for these derivations, and each system of exact pairs gives a normal crystal. For any finite word of sink colors, the successive raising lengths and final label parametrize this basis by the lattice points of a rational polyhedral cone. Reciprocity between opposite mutable patterns tests exactness in any seed. We illustrate this theory by various examples.

Contents

Introduction 1
1. Theta functions and coefficients 5
2. Canonical boundary derivations and degrees 8
3. Regular theta cones and tropical crystals 10
4. Tropical DT duality and exact boundary pairs 15
5. Primitive theta lifts and changes of coefficients 18
6. Exact boundary pairs and Lie actions 21
7. Perfect theta bases and crystal structures 28
8. Symmetries of theta crystals 31
9. Geometric applications 34
10. Theta F-pairings and maximal adjunction 40
11. Maximal operators and string coordinates 45
12. Pointed biperfect bases and transitions 48
Declaration of generative AI and AI-assisted technologies 53
References 53

Introduction

The least exponent of a frozen variable x_i in a Laurent polynomial gives its order along $x_i = 0$. What does the greatest exponent measure for a regular theta function? If the mutable pattern admits a reddening sequence and i is a sink relative to the mutable vertices in the chosen seed, it counts the raising steps to the top of the i-string.

For example, the two clusters in the Plücker ring of the Lagrangian Grassmannian $\mathrm{LGr}(2,4)$ satisfy $xy = b_1^2 + b_0 b_2$, with b_0, b_1, b_2 frozen. In the cluster containing x, take $R = x\partial_{b_1}$, and

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 (No. 12571038).

in the cluster containing y, take $L = y\partial_{b_1}$. Differentiating the exchange relation gives

$$R(y) = 2b_1, \qquad R(b_1) = x, \qquad L(x) = 2b_1, \qquad L(b_1) = y.$$

Thus R raises and L lowers along the string of theta functions y, b_1, x. In the x-cluster, $y = x^{-1}(b_1^2 + b_0 b_2)$ has order zero along $b_1 = 0$ and lies two raising steps below x. Here both derivations come from b_1, in different seeds. For a general seed we determine which sinks and sources raise and lower along the same strings.

We work with the theta functions of Gross, Hacking, Keel, and Kontsevich [21]. For skew-symmetric seeds, the theory of quivers with potential gives crystal operators on supported presentation weights [11]. On the coordinate ring of a maximal unipotent subgroup, Tang proved the *biperfectness* of the theta basis for the left and right simple-root derivations [42, Theorem 1.1]. Our construction uses the exchange matrix and its symmetrizing weights, including those at the frozen vertices. Under reddening it gives operators on the basis of regular theta functions even when it is unknown whether that basis spans the upper cluster algebra.

Boundary derivations and crystal strings. We work over a field $\Bbbk$ of characteristic zero. Write U and F for the mutable and frozen vertices, and $\mathcal{U}^+$ for the upper cluster algebra in which the frozen variables remain polynomial. A theta function is *finite* if its lift to principal coefficients is a Laurent polynomial. We specialize the auxiliary coefficients to 1. The specializations of the finite theta functions span the localized middle algebra. Its intersection with $\mathcal{U}^+$ is the *regular middle algebra* Mid^+. Valuative independence [7] gives a basis of this intersection consisting of the *regular* finite theta functions, those belonging to $\mathcal{U}^+$. In a fixed seed we write ϑ_g for the theta function with g-vector g and $\mathcal{B}_\theta^+$ for the labels of the regular theta functions.

Let $B = (b_{ab})$ be the exchange matrix in the convention that BD is skew-symmetric for a positive integral diagonal matrix $D = \mathrm{diag}(d_a)$. A frozen vertex i is *sink-reachable* if some seed satisfies $b_{iu} \leqslant 0$ for all $u \in U$. In such a *sink-optimizing seed* define

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

where m_i is the full exponent vector and e_i is the ith coordinate vector. Reversing the inequalities defines a *source-optimizing seed* for a source-reachable vertex k, with $M_k^\star = \prod_{u \in U} x_u^{b_{ku}}$ and $R_k^\star = M_k^\star \partial_{x_k}$. In principal coefficients, the rows of the extended exchange matrix indexed by the auxiliary variables form the *c-matrix* in the column convention. A mutation sequence is *reddening* if its final c-matrix is a negative permutation matrix [29].

Theorem A. *Suppose the mutable pattern admits a reddening sequence. In every seed $\mathcal{B}_\theta^+$ is the set of integral points of a rational polyhedral cone $\mathcal{C}^+$. For a sink-reachable vertex i, the raising operator r_i and its partial inverse l_i act on a regular label g in a sink-optimizing seed by*

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

We write 0_{cr} for an undefined value. Tropical change of coordinates identifies these partial maps in all seeds. In the optimizing seed exactly g_i raising steps are possible from g, and maximal raising extends to the linear retraction $g \mapsto g + g_i v_i$ onto the face $g_i = 0$.

The defining inequalities of $\mathcal{C}^+$ express regularity along the frozen divisors. Theta reciprocity [7] computes them from the Newton polytopes of dual theta functions. We prove that these polytopes have integral vertices in the original valuation lattice and use their inequalities to show that the entire raising string lies in the cone (Section 3).

For a set I of sink-reachable vertices, let $C = (c_{ij})_{i,j \in I}$ have diagonal entries 2 and off-diagonal entries $c_{ij} = -\deg_{x_i} M_j$, computed in a sink-optimizing seed for i. For a locally nilpotent derivation R, write $R^{(n)} = R^n/n!$ and call the largest n with $R^n f \neq 0$ the *nilpotence degree* of $f \neq 0$. A basis is *normalized perfect* for R if it is adapted to the filtration by

$\ker R^{n+1}$ and the leading term of R on a basis element of degree $n > 0$ is n times a basis element of degree $n-1$, with distinct leading elements for distinct inputs of that degree.

Theorem B. *The monomials $M_i, M_k^\star$ and the derivations $R_i, R_k^\star$ are independent of their respective optimizing seeds. The derivations are locally nilpotent on $\mathcal{U}^+$. The matrix C is a symmetrizable generalized Cartan matrix, and the R_i, $i \in I$, satisfy its Serre relations.*

Under reddening the sink and source derivations preserve Mid^+*, and its theta basis is normalized perfect for each of them. In a sink-optimizing seed for i, if $g_i > 0$, then*

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

Cao's F-invariant computes the Cartan entries [5]. To prove perfectness, we identify the terms of highest x_i-degree in ϑ_g by freezing the mutable variables that occur in M_i. Qin's freezing theorem [41] identifies their sum with a single theta function in the resulting seed. This extends Tang's highest-degree argument in full exchange rank [42, Proposition 3.2].

If the exchange relations admit integral gradings in which R_i has degree c_{ji} in the jth grading, Theorem A gives an *upper seminormal crystal* of type C. Its upper string function counts the raising steps, and its weights are given by their coroot coordinates. A lowering string can still be infinite. Exact pairs supply both the gradings and the finite lowering strings.

Exact pairs and integrable actions. Call the ordered pair (i,k) of a sink-reachable and a source-reachable vertex *exact* if, in a sink-optimizing seed for i,

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

There is a direct test on the two boundary monomials. For distinct vertices exactness is equivalent to $R_iM_k^\star = x_k$. For $k = i$ it is equivalent to $R_i^{(2)}M_i^\star = M_i$ (Proposition 6.4). In fact, the pair is exact if and only if some nonzero multiples of R_i and $R_k^\star$ form a Chevalley triple with their commutator. The two scalars are then reciprocal (Theorem 6.5). In the example of $\mathrm{LGr}(2,4)$, the two boundary monomials x and y give an exact self-pair at b_1.

We call a crystal *normal* if its restriction to any set of colors of finite Cartan type is a disjoint union of crystals of finite-dimensional irreducible highest-weight modules.

Theorem C. *Suppose each $i \in I$ has an exact source partner $\sigma(i)$. The map σ is injective and $d_i = d_{\sigma(i)}$ for every i. With $L_i = R_{\sigma(i)}^\star$ and $H_i = [R_i, L_i]$, the operators R_i, L_i, H_i define an integrable action of the derived Kac–Moody algebra $\mathfrak{g}'(C)$ on $\mathcal{U}^+$. Each connected component of C has finite or affine type, and each affine component acts at* level zero*, meaning that its canonical central element acts by zero.*

Under reddening the regular theta basis of Mid^+ *carries a normal crystal for this action, with coroot gradings given by the H_i. The same conclusion holds on $\mathcal{U}^+$ whenever it has a basis of theta functions with finite principal-coefficient lifts, even without reddening.*

To rule out indefinite Cartan type, we prove that the Lie words of length at most n in finitely many $\Bbbk$-derivations of a field of transcendence degree d over $\Bbbk$ span a space of dimension $O(n^d)$ (Lemma 6.9). The level-zero assertion follows from a rank bound on commutators applied to the Heisenberg subalgebra of each affine component (Theorem 6.10). In the skew-symmetric setting, the corresponding action and Cartan-type restriction for exact boundary representations were proved in [11, Theorem 8.6 and Corollary 8.8].

A reciprocity identity relates the original and opposite mutable patterns. Under reddening it gives a test for the exact-pair hypothesis of Theorem C in any reference seed. Write $\mathbf{a}_j = (b_{ju})_{u\in U}$ for the frozen exchange vectors. For $b \in \mathbb{Z}^U$, let F_b be the principal-coefficient F-polynomial of the theta function with mutable g-vector $-b$ and zero frozen coordinates whenever this theta function is finite. Let f_b be the componentwise greatest exponent when

it exists. With D_U and $B_{U,U}$ the mutable submatrices of D and B, write

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

Write E^{op} for the same function in the opposite mutable pattern, whose exchange matrix has the opposite sign. Under reddening f_b exists for every b. We orient the mutable tropical Donaldson–Thomas transformation T by $T^{-1}b = -b - f_b B_{U,U}$.

Theorem D. *If the principal theta functions of mutable degrees $-b$ in the original pattern and a in the opposite pattern are finite, then $\mathsf{E}(a,b) = \mathsf{E}^{\mathrm{op}}(-b,-a)$. Under reddening this gives $\mathsf{E}(a,b) = \mathsf{H}(b,Ta)$ for all $a,b \in \mathbb{Z}^U$. In this case the pair (i,k) is exact if and only if $T^{-1}\mathbf{a}_i = T\mathbf{a}_k$ and, at every frozen vertex j,*

$$\mathsf{H}(\mathbf{a}_j,\mathbf{a}_i) + \mathsf{H}(T^{-1}\mathbf{a}_i, T\mathbf{a}_j) = d_j(\delta_{ij} + \delta_{kj}),$$

where δ_{ij} is the Kronecker delta.

The equality involving T compares the mutable coordinates of the two boundary monomials. The equalities indexed by j compare their frozen coordinates. This criterion is the theta analogue of [11, Lemma 4.29]. The opposite-pattern identity also holds for formal principal theta series when both support functions are finite in the indicated directions (Theorem 4.3).

Applications. Under reddening, a family of exact pairs determines a *relative DT square* from the seed with the chosen sinks as frozen vertices to the seed with their source partners. Given an opposite cluster involution of the first seed, we follow it by the relative square. If the pairs cover the frozen vertices, the composite lifts uniquely to an opposite involution of the full seed. For a finite-type set of paired colors, at most one lift realizes Lusztig's involution. Theorem 8.2 characterizes it by opposition and its action on highest theta functions. Corollary 11.2 reduces the comparison on highest functions to the extreme rays of a rational cone determined by the highest labels and one opposite F-polynomial. For the base affine space, this identifies the relative square on mutable labels with Kashiwara's involution followed by Lusztig's involution (Proposition 9.3).

On the Plücker ring of $\mathrm{LGr}(n,2n)$, $n \geqslant 2$, the additional affine raising operator acts through the root space for $-\theta$, where θ is the highest root of $\mathfrak{sp}_{2n}$. We identify the affine action as an evaluation representation and obtain a level-zero theta crystal of type $C_n^{(1)}$ (Proposition 9.5). In each Plücker degree the finite-type Lusztig involution also reverses the arrows of color 0 (Corollary 9.6).

For a finite word of sink colors, raise a regular theta function to the top of each string in turn. The successive string lengths and the final g-vector give all lattice points of a single rational polyhedral cone (Theorem 11.1). For a reduced word for the longest element of a finite-type subsystem of paired colors, the slice at a highest label gives the usual string parameters of its crystal component. Maximal raising also sends the cluster monomials of any one seed to monomials in a single target seed (Theorem 11.3). To test adjunction for the valuation pairing on the entire regular theta basis, it suffices to test the seed variables and one squarefree cluster monomial (Theorem 10.6).

Section 5 describes how changes of coefficients affect regular theta functions and proves that opposite cluster automorphisms reverse the mutable tropical Donaldson–Thomas transformation under reddening. Section 12 classifies the changes of biperfect basis that admit compatible pointed lifts before coefficient specialization. We conclude by comparing the regular theta basis with supported generic cluster characters in the skew-symmetric case over $\mathbb{C}$, under mutable reddening and the assumption that the full Jacobian algebra is finite dimensional. The orders along the frozen divisors determine when these characters span Mid^+ (Proposition 12.9).

1. Theta functions and coefficients

We work over a field $\Bbbk$ of characteristic zero. Let $V = U \sqcup F$ be the vertex set of a seed, with mutable vertices U and frozen vertices F. We follow the matrix convention of [11]. Thus $B = (b_{ab})_{a,b\in V}$ is an integer matrix, and $D = \operatorname{diag}(d_a)$ is a positive integral diagonal matrix such that BD is skew-symmetric. The weights d_f at frozen vertices are part of the seed datum. For subsets $S, S' \subseteq V$, write $B_{S,S'}$ for the indicated submatrix and D_S for the restriction of D to S. In the usual column convention the extended exchange matrix is $\widetilde{B} = -B^{\mathsf{T}}_{U,V}$. With $[c]_+ = \max(c, 0)$, mutation at $k \in U$ gives

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

We use the mutation rules of [15, 16]. The entries $B_{F,F}$ of the square matrix do not enter the exchange relations. We write exponents as row vectors, $x^p = \prod_{a\in V} x_a^{p_a}$, and use e_a for the coordinate vector at a. Valuation directions are columns, and $\mathbb{N}$ includes zero.

For $F_0 \subseteq F$, the upper cluster algebra in which the variables x_f, $f \in F_0$, remain polynomial is

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

The intersection runs over all seeds t obtained by mutable mutations, and $x_u(t)$ denotes the corresponding seed variable. We write $\mathcal{U}^+ = \mathcal{U}_F$ and $\mathcal{U}^{\mathrm{loc}} = \mathcal{U}_\varnothing$. The localized cluster algebra $\mathcal{A}^{\mathrm{loc}}$ is generated over $\Bbbk[x_f^{\pm 1} : f \in F]$ by the mutable cluster variables.

In the principal-coefficient pattern, the *c-matrix* consists of the rows of the extended exchange matrix indexed by the principal coefficients. A mutation sequence is *reddening* if every column of its final c-matrix is nonpositive. This final matrix is then a negative permutation matrix [29]. Tropical duality identifies the existence of a reddening sequence with *injective reachability*. In this case, for a seed t there is a seed $t[1]$ whose mutable g-vectors at t are the negatives of the standard basis vectors, up to permutation. We denote by $t[-1]$ a seed with $t[-1][1] = t$ up to relabeling [38, 40].

An integral tropical point has coordinates $g \in \mathbb{Z}^V$ in the chosen seed. Adjoin independent principal coefficients y_u, $u \in U$, and lift g by taking all auxiliary coordinates to be zero. The corresponding principal theta function of [21] is *pointed at* g. Its expansion is

$$x^g \mathcal{F}_g\big((x^{-B_{u,V}} y_u)_{u\in U}\big), \qquad \mathcal{F}_g(0) = 1,$$

where $\mathcal{F}_g$ has nonnegative integral coefficients. We call the theta function *finite* when this expansion is a Laurent polynomial. Then $\mathcal{F}_g$ is its *F-polynomial*, and we write ϑ_g for the result of setting all $y_u = 1$. Under reddening every integral tropical point has a finite lift by [40, Theorem A.1.5] applied to the full-rank principal extension.

The finite specialized theta functions span the *localized middle algebra* $\mathrm{Mid}^{\mathrm{loc}}$ inside $\mathcal{U}^{\mathrm{loc}}$. We write $\mathrm{Mid}^+ = \mathrm{Mid}^{\mathrm{loc}} \cap \mathcal{U}^+$ for the *regular middle algebra*. A finite theta function is *regular* if it belongs to $\mathcal{U}^+$. We write $\mathcal{B}^+_\theta$ for the labels of the regular finite theta functions.

Lemma 1.1. *For every $F_0 \subseteq F$, the finite theta functions*

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

form a basis of $\mathrm{Mid}^{\mathrm{loc}} \cap \mathcal{U}_{F_0}$. Here ord_{x_f} is the order of vanishing along $x_f = 0$ in any seed chart. In particular, the regular theta functions form a basis of Mid^+. If $\mathrm{Mid}^{\mathrm{loc}} = \mathcal{U}^{\mathrm{loc}}$, then Θ_{F_0} is a basis of $\mathcal{U}_{F_0}$.

Proof. The valuation $\nu_f = \operatorname{ord}_{x_f}$ is independent of the seed. Indeed, at least one of the two exchange monomials is independent of x_f, and their sum remains nonzero modulo x_f. Mutation therefore induces a birational change of coordinates on $x_f = 0$ and an isomorphism of the local rings at its generic point. It follows that

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

The specialized theta functions are linearly independent by valuative independence in the zero valuation direction [7, Theorem 3.10]. For a finite linear combination with at least one nonzero coefficient, the same theorem gives

$$\nu_f\Big(\sum_g c_g\vartheta_g\Big) = \min_{c_g\neq 0}\nu_f(\vartheta_g).$$

Thus the sum belongs to $\mathcal{U}_{F_0}$ exactly when each term does. Remark 1.2 verifies the hypotheses on the scattering diagram, including the case of a zero exchange column. □

Remark 1.2. A zero column of $\widetilde{B}$ corresponds to an isolated mutable vertex. Its exchange relation is $xx' = 2$, and the theta function of index $a \in \mathbb{Z}$ is $2^{[-a]_+}x^a$. These nonzero scalar multiples of Laurent monomials satisfy valuative independence over $\Bbbk$.

For the remaining vertices, write $P = \widetilde{B}$ and let $Q^\bullet$ have columns $(d_u/\ell)e_u^*$, where ℓ is any common multiple of the d_a and e_a^* is dual to the standard coordinate vector e_a. The form $B^\bullet = (Q^\bullet)^\mathsf{T}P$ is skew-symmetric. By [7, Lemma 4.3], a wall function is a series in one monomial $x^{Pv}y^v$, where y denotes the auxiliary principal coefficients and $v \in \mathbb{N}^U\backslash\{0\}$. On each initial wall $Pv \neq 0$. A bracket of degrees v_1, v_2 with $P(v_1 + v_2) = 0$ vanishes, since its scalar factor is $B^\bullet(v_1, v_2) = B^\bullet(v_1, v_1 + v_2) = 0$. Consequently $Pv \neq 0$ on every nontrivial wall. Distinct powers of its monomial have distinct x-exponents, as required by [7, Assumption 3.9]. The isolated factors multiply the theta functions by Laurent monomials in independent variables, so valuative independence holds for the entire seed.

Proposition 1.3. *Suppose the mutable pattern admits a reddening sequence. If in some seed a linear functional is strictly positive on every nonzero column of $\widetilde{B}$, then the finite theta functions form a basis of $\mathcal{U}^{\mathrm{loc}}$, and Θ_{F_0} is a basis of $\mathcal{U}_{F_0}$ for every $F_0 \subseteq F$. In particular, this holds if $\widetilde{B}$ has full column rank.*

Proof. Split off the isolated vertices as in Remark 1.2. On the remaining seed the hypothesis gives the convexity condition of [21, Theorem 0.3(7)]. Its initial and injective seeds give opposite mutable chambers, so its cluster complex is large in the sense of [21]. Proposition 0.7 of that paper gives the localized theta basis. Tensoring with the Laurent monomial bases of the isolated factors gives the basis of $\mathcal{U}^{\mathrm{loc}}$, and Lemma 1.1 gives the basis of $\mathcal{U}_{F_0}$. □

1.1. **Opposite theta functions.** The *opposite seed* has the same variables and frozen vertices as the original seed and has exchange matrix $-B$. We identify the two rational function fields by fixing the initial variables.

Proposition 1.4. *If the mutable pattern admits a reddening sequence, the original seed and the opposite seed have the same set of finite theta functions. This holds in every extended rank when theta functions are defined by specialization from principal coefficients.*

Proof. When $\widetilde{B}$ has full column rank, this is [42, Proposition 4.2]. Its injective-reachability hypothesis follows from reddening. The proof uses the Donaldson–Thomas (DT) twist and its action on theta functions [30, Theorems 5.2 and 6.2(2)]. The identities between Laurent polynomials in that proof hold over any field of characteristic zero.

In arbitrary rank, adjoin independent principal coefficients. The extended exchange matrix now has full column rank and the same mutable principal part. The result in full rank therefore applies to this seed and its opposite. Translating a theta index in the new frozen coordinates multiplies the corresponding function by a Laurent monomial in the new variables. Setting these variables equal to 1 thus maps the two equal sets onto the original and opposite sets, respectively. □

We call a seed *sink-optimizing* for $i \in F$ if $b_{iu} \leqslant 0$ for every $u \in U$, and *source-optimizing* if $b_{iu} \geqslant 0$ for every $u \in U$. A vertex is *sink-reachable* or *source-reachable* if there is a seed of the corresponding kind.

Lemma 1.5. *Suppose the mutable pattern admits a reddening sequence. A frozen vertex i is sink-reachable if and only if it is source-reachable. If t is sink-optimizing for i, then $t[-1]$ is source-optimizing for i. If t is source-optimizing for i, then $t[1]$ is sink-optimizing for i.*

Proof. Let $a(t) = B_{i,U}(t)$. By matrix mutation, $-a(t)$ transforms as a mutable g-vector for the column exchange matrix $-B_{U,U}^{\mathsf{T}}$. This transformation depends only on the mutable part of the matrix, so we may work with principal coefficients. The mutable variables of $t[1]$ have mutable g-vectors at t equal to the negatives of the standard basis vectors, up to permutation. The same holds for the variables of t at $t[-1]$ [40, Definition 2.3.1 and Proposition 2.3.3].

If $a(t) \leqslant 0$, take the mutable monomial with exponent vector $-a(t)$ at t. Its mutable g-vector at $t[-1]$ is nonpositive and equals $-a(t[-1])$ by tropical mutation. Thus $t[-1]$ is source-optimizing. If $a(t) \geqslant 0$, choose a monomial in the variables of $t[1]$ whose mutable g-vector at t is $-a(t)$. Its exponent vector at $t[1]$ is nonnegative and equals $-a(t[1])$, so $t[1]$ is sink-optimizing. □

1.2. **Degrees and codegrees.** Under reddening, define the *codegree* of ϑ_g in a seed to be the index of the same theta function in the opposite seed. In full column rank this agrees with the codegree for the dominance order of [40, Propositions 3.3.12 and 3.4.8]. In arbitrary rank both indices are computed in the principal-coefficient extension before specialization. Writing $\mathcal{F}_g$ for the F-polynomial introduced above, we have in the original seed

$$\widehat{y}_u = x^{-B_{u,V}}, \qquad \vartheta_g = x^g \mathcal{F}_g(\widehat{y}).$$

Proposition 1.6. *Under reddening, $\mathcal{F}_g$ has a componentwise greatest exponent $f(g) \in \mathbb{N}^U$, and its coefficient is 1. Write $\eta_\theta(g)$ for the codegree of ϑ_g and $\mathcal{F}^{\mathrm{op}}_{\eta_\theta(g)}$ for its F-polynomial in the opposite seed. Then*

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

These formulas hold for arbitrary geometric coefficients. For cluster monomials they hold even without reddening, with the codegree understood as the pointed degree in the opposite seed.

Proof. Work with independent principal coefficients y_u. Under reddening, Proposition 1.4 identifies the two theta expansions. A cluster monomial belongs to the same cluster for the original and opposite exchange matrices, so in this case the two expansions exist without reddening by the separation formula [16, Corollary 6.3]. Their constant coefficients are 1 by pointedness, or for cluster monomials by sign coherence [21]. Thus in either case

$$x^g \mathcal{F}_g(\widehat{y}y) = x^h y^f G(\widehat{y}^{-1}y^{-1}), \qquad \mathcal{F}_g(0) = G(0) = 1,$$

where products and inverses in the arguments are coordinatewise. The constant term of G contributes $x^h y^f$. Comparing the coefficients of y^f gives $h = g - fB_{U,V}$ and $[y^f]\mathcal{F}_g = 1$. More generally, e occurs in $\mathcal{F}_g$ exactly when $f - e$ occurs in G, with the same coefficient. Since the exponents of G are nonnegative, f is the greatest exponent of $\mathcal{F}_g$. This proves the reversal formula. Setting the auxiliary variables equal to 1 gives the stated codegree. □

In the normalization of [30], the DT twist on the localized algebra sends the opposite theta function of index h to the original theta function of index $-h$. Its action on indices is therefore

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

On mutable g-vectors this is inverse tropical transport to the injective seed. We use this convention in Section 4.

2. Canonical boundary derivations and degrees

Fei constructs derivations from boundary representations of quivers with potential [11, Sections 4 and 8.1]. We construct the corresponding derivation from the exchange matrix. In a sink-optimizing seed for $i \in F$, let

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

where we extend ζ_i by zero on U and ∂_{x_i} fixes the other seed variables.

2.1. Canonical monomials and the degree formula.

Lemma 2.1. *The monomial M_i and the vector ζ_i are independent of the sink-optimizing seed at i. The g-vector of M_i in any seed t is*

$$g(M_i; t) = -B_{i,V}(t) + \zeta_i.$$

Proof. The formula holds in the chosen optimizing seed. By matrix mutation, $-B_{i,V}(t)$ transforms by tropical g-mutation for $-B^{\mathsf{T}}$. Adding ζ_i commutes with this transformation because its mutable part is zero. Thus the formula holds in every seed.

In a second optimizing seed the mutable coordinates of $g(M_i)$ are nonnegative. The cluster monomial property of theta functions [21] gives

$$M_i = \prod_{u \in U} x_u^{-b_{iu}} \prod_{f \in F} x_f^{\zeta_{if} - b_{if}}.$$

Each mutable cluster variable has order zero along every frozen divisor by the valuation invariance in Lemma 1.1. Taking the order along $x_f = 0$ in this identity gives $\zeta_{if} = b_{if}$, as asserted. □

For a mutable cluster monomial P, let g_P^t be the mutable part of its g-vector in a seed t, and let F_P^t be its F-polynomial with principal coefficients. For a nonzero polynomial G with exponent support $\operatorname{supp} G$ and a column vector w, write $G[w] = \max_{e \in \operatorname{supp} G} ew$. Cao's *$F$-invariant* of two mutable cluster monomials 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 symmetric, nonnegative and independent of t. It vanishes precisely when P and Q are *compatible*, meaning that they are monomials in a common cluster [5, Theorems 4.10 and 4.19, Proposition 4.30]. The principal F-polynomials and mutable g-vectors depend only on the mutable pattern.

In a sink-optimizing seed for i, write $M_i = x^m$, so $m_u = -b_{iu}$ for $u \in U$ and $m_f = 0$ for $f \in F$. Let $\rho_i(P) = \deg_{x_i} P$ in this seed.

Proposition 2.2. *For every mutable cluster monomial P, we have $d_i \rho_i(P) = \mathsf{F}_D(M_i, P)$. In particular, $\rho_i(P)$ is independent of the optimizing seed.*

Proof. Write the separation formula as $P = x^g F_P(\widehat{y})$. An exponent e of F_P contributes x_i-exponent $g_i - d_i^{-1} e D_U m_U^{\mathsf{T}}$, since $\widetilde{b}_{iu} = -b_{ui} = -(d_u/d_i) m_u$. The largest such exponent is g_i, attained at the constant term. The smallest is zero, since P has order zero along $x_i = 0$ and the coefficients are positive. Thus

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

Here $F_{M_i} = 1$ and $g_{M_i} = m_U$, so the last expression is $\mathsf{F}_D(M_i, P)/d_i$. □

Corollary 2.3. *The rational derivation R_i is independent of the sink-optimizing seed. Its kernel contains a mutable cluster monomial P if and only if P is compatible with M_i.*

Proof. By Lemma 2.1, M_i is a cluster monomial in every optimizing seed. Each mutable variable in a second such seed is therefore compatible with M_i and has ρ_i-degree zero by Proposition 2.2. Thus R_i kills these variables. On a frozen variable x_f the two derivations both take the value $\delta_{if}M_i$, where δ_{if} is the Kronecker delta. They agree on a transcendence basis of the rational function field. The description of the kernel follows from the degree formula and Cao's compatibility criterion. □

2.2. **Regularity and local nilpotence.**

Theorem 2.4. *For a sink-reachable vertex i, the derivation R_i preserves $\mathcal{U}^+$ and is locally nilpotent on it. For a source-reachable vertex j, let $M_j^\star = \prod_{u\in U} x_u^{b_{ju}}$ in a source-optimizing seed. Then $R_j^\star = M_j^\star\partial_{x_j}$ is likewise independent of the optimizing seed and locally nilpotent on $\mathcal{U}^+$. Each derivation preserves every $\mathcal{U}_{F_0}$ and is locally nilpotent there precisely when its own frozen vertex belongs to F_0.*

Proof. In a sink-optimizing seed for i, the derivation $R_i = M_i\partial_{x_i}$ preserves the seed ring with polynomial frozen variables. It also preserves the ring of each adjacent seed. If $b_{ik} = 0$, the exchange relation at k is independent of x_i, so $R_i(x_k') = 0$. In this case M_i is a monomial in the unchanged variables. If $b_{ik} < 0$, write $x_kx_k' = P = Ax_i^q + C$ with $q > 0$, where A and C are independent of x_i and x_k. Then

$$R_i(x_i) = M_i, \qquad R_i(x_k') = qAx_i^{q-1}M_i/x_k.$$

Since x_k divides M_i, substituting $x_k = P/x_k'$ gives two elements of the adjacent seed ring. The other seed variables have derivative zero, and the inverse of each mutable seed variable lies in that ring.

After all frozen variables have been inverted, the initial and adjacent seed rings are still stable under R_i. If $\widetilde{B}$ has full column rank, their intersection is $\mathcal{U}^{\mathrm{loc}}$ by [2, Corollary 1.9]. Intersecting $\mathcal{U}^{\mathrm{loc}}$ with the original optimizing seed ring recovers $\mathcal{U}^+$ by invariance of the frozen valuations. This proves stability in full column rank.

In arbitrary rank, adjoin principal frozen coefficients at the optimizing seed. The new matrix has full column rank, and the lifted derivation has multiplier M_i. For every seed t, the derivatives of its lifted seed variables lie in the extended upper algebra with polynomial frozen variables, hence in the extended seed ring at t. Setting the new coefficients equal to 1 commutes with $M_i\partial_{x_i}$ and maps these derivatives to the derivatives of the original variables. Thus R_i preserves the seed ring at t and hence their intersection $\mathcal{U}^+$.

In the optimizing seed, every element of $\mathcal{U}^+$ is polynomial in x_i and M_i is independent of x_i, which proves local nilpotence. Clearing the negative frozen orders gives $\mathcal{U}_{F_0} = \mathcal{U}^+[x_f^{-1} : f \notin F_0]$. The derivation extends to this localization and remains locally nilpotent if $i \in F_0$, since it kills all the inverted variables. If $i \notin F_0$, then $R_i^n(x_i^{-1}) = (-1)^n n! M_i^n x_i^{-n-1} \neq 0$ for every n. Applying the same argument to $-B$ gives the source assertion. □

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

$$d_ia_{ij} = d_ja_{ji} = \mathsf{F}_D(M_i, M_j).$$

It follows that $c_{ii} = 2$ and $c_{ij} = -a_{ij}$ for $i \neq j$ define a symmetrizable generalized Cartan matrix, with symmetrizer $\mathrm{diag}(d_i)$. In a sink-optimizing seed for i, the ith coordinate of $g(M_j)$ equals $\rho_i(M_j)$ by the proof of Proposition 2.2. Since $b_{ij} = \zeta_{ij}$ in that seed, Lemma 2.1 gives

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

For a source-reachable frozen vertex j and a mutable cluster monomial P, let $\rho_j^\star(P) = \deg_{x_j} P$ in a source-optimizing seed. This is the largest r with $(R_j^\star)^r P \neq 0$. In that seed all x_j-exponents in the separation formula are at least the pointed exponent. The constant

term attains the minimum, which is zero, and the maximum equals $d_j^{-1}F_P[D_U g(M_j^\star)_U^{\mathsf{T}}]$. Consequently

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

The mixed degrees thus vanish precisely when M_i and $M_j^\star$ are compatible. For a fixed mutable pattern the multipliers depend on the valued entries b_{fu}, $f \in F$, $u \in U$. The entries between frozen vertices play no part in the mutation of the extended exchange matrix.

3. Regular theta cones and tropical crystals

Fei's tropical crystal has as its elements the presentation weights of objects supported on the mutable vertices [11, Section 5]. We construct the crystal operators on the g-vectors of regular theta functions. Unless stated otherwise, the mutable pattern in this section admits a reddening sequence, so every integral tropical point has a finite principal-coefficient theta function.

3.1. **Theta reciprocity and the regular cone.** Let $M = \bigoplus_{a\in V}\mathbb{Z}e_a$ be the exponent lattice in a fixed seed, and let $N = M^*$, with dual basis $(e_a^*)_{a\in V}$. Choose a positive common multiple ℓ of the d_a. The lattices for the *chiral dual* seed datum are

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

For each $u \in U$, the original initial wall has normal $q_u = e_u^*$ and its monomial has exponent $p_u = \widetilde{B}e_u$. In the notation of [7, Sections 4.1 and 4.5], the exponent and normal of the dual initial wall are, respectively,

$$q_u^\bullet = (d_u/\ell)e_u^*, \qquad p_u^\bullet = (\ell/d_u)p_u.$$

The normal belongs to $M^\bullet$ because $-d_a b_{ua} = d_u b_{au}$. Also $e_f^* \in N^\bullet$ for every $f \in F$ by the choice of ℓ.

For a Laurent polynomial $H = \sum_a h_a x^a$ and a valuation direction n, let $\mathrm{val}_n(H) = \min_{h_a\neq 0}\langle a, n\rangle$. We use this minimum convention for both seed data, including after freezing mutable variables. Let $\vartheta_n^\vee$ denote the theta function of the dual datum indexed by $n \in N^\bullet$. When a column of $\widetilde{B}$ is zero, the dual normal vanishes. We omit that wall and take the dual of the isolated factor to be a torus. Its theta functions are Laurent monomials. By Remark 1.2, both valuations on this factor equal the pairing of the two exponents. On the remaining factors theta reciprocity applies directly [7, Corollary 3.2, Claim 4.16 and Theorem 5.19]. Adding the valuations of the factors gives

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

whenever the two values are finite.

Lemma 3.1. *Write $n = \sum_{a\in V} r_a(d_a/\ell)e_a^*$, with $r = (r_a)_{a\in V}$. Suppose the principal theta function of mutable degree $h = rB_{V,U}$ for the opposite mutable pattern is finite, and denote its F-polynomial by $\mathcal{F}_h^{\mathrm{op}}$. Then*

$$\vartheta_n^\vee = z^n\mathcal{F}^{\mathrm{op}}_{rB_{V,U}}\big((z^{q_u^\bullet})_{u\in U}\big), \tag{3.1}$$

where z^α is the Laurent monomial of exponent $\alpha \in N^\bullet$. This identity requires no reddening. Under reddening every $\vartheta_n^\vee$ is a Laurent polynomial in every seed with frozen variables inverted.

Proof. First omit the isolated zero columns. Use the basis $((d_a/\ell)e_a^*)_{a\in V}$ of $N^\bullet$ and let $E_U : \mathbb{Z}^U \hookrightarrow \mathbb{Z}^V$ be coordinate inclusion. In these coordinates the exponent and normal matrices of the dual datum are E_U and $B_{V,U}$. The lattice embedding

$$\iota : \mathbb{Z}^V \longrightarrow \mathbb{Z}^U \oplus \mathbb{Z}^V, \qquad t \longmapsto (B_{V,U}^{\mathsf{T}}t, t)$$

is a linear morphism to the seed datum with

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

Indeed, $\iota E_U = \widehat{P}$ and $\iota^{\mathsf{T}}\widehat{Q} = B_{V,U}$. The target is the principal-coefficient seed of the opposite mutable pattern together with independent frozen Laurent variables. Its mutable exchange matrix in the column convention is $B_{U,U}^{\mathsf{T}} = -\widetilde{B}_{U,U}$, and its weight matrix in the convention of [7] is ℓD_U^{-1}.

The linear-morphism theorem in its positive-chamber form [7, Theorem 4.9 and equation (21)] sends $\vartheta_n^\vee$ to the target theta function of index ιr^{T}, even when the source has no full-dimensional positive chamber. Its mutable degree is $rB_{V,U}$. In the pointed expansion each indeterminate of its F-polynomial is replaced by the monomial of exponent $\widehat{P}e_u = \iota E_U e_u$. Before specialization this identity holds coefficientwise for the formal principal series, whose uth argument after pullback is $y_u z^{q_u^\bullet}$. Here y_u is the auxiliary scattering coefficient. The principal frozen coordinates distinguish every coefficient exponent. For a finite opposite theta function, setting $y_u = 1$ gives (3.1). An isolated zero column contributes a Laurent monomial on the dual side and zero to the corresponding coordinate of $rB_{V,U}$, so the same formula includes that factor. Inverting the principal coefficient variables identifies the target with the opposite of the original principal seed, up to its independent Laurent variables. Under reddening finiteness now follows from Proposition 1.4. □

Cheung, Magee, Mandel and Muller characterize regular theta functions on a partial compactification by their boundary valuations [7, Corollary 1.2]. For each frozen vertex f these valuations are determined by one opposite F-polynomial. When $\mathcal{F}^{\mathrm{op}}_{B_{f,U}}$ is finite, define

$$A_f = \{d_f^{-1}\nu D_U : \nu \in \operatorname{supp} \mathcal{F}^{\mathrm{op}}_{B_{f,U}}\}, \qquad \mathcal{P}_f = \operatorname{conv}(A_f) \subseteq \mathbb{R}^U_{\geqslant 0}.$$

We call $\mathcal{P}_f$ the *boundary polytope* at f. Its defining set A_f contains 0.

Corollary 3.2. *Under mutable reddening, the g-vectors of the regular theta functions in any seed are precisely the lattice points of the rational polyhedral cone*

$$\mathcal{C}^+ = \{(p,q) \in M_{\mathbb{R}} : q_f + p \cdot c \geqslant 0 \text{ for all } f \in F,\ c \in \mathcal{P}_f\},$$

where $M_{\mathbb{R}} = M \otimes_{\mathbb{Z}} \mathbb{R}$ and the dot pairs mutable coordinates.

Without reddening, fix $F_0 \subseteq F$. If $\mathcal{F}^{\mathrm{op}}_{B_{f,U}}$ is finite for each $f \in F_0$, the same inequalities for $f \in F_0$ characterize those finite theta functions that belong to $\mathcal{U}_{F_0}$. This hypothesis holds if every $f \in F_0$ is source-reachable.

Proof. Write $n_f = (d_f/\ell)e_f^*$. By (3.1), the exponent support of $\vartheta_{n_f}^\vee$ is $(d_f/\ell)(e_f^* + A_f)$, where we extend each element of A_f by zero on the frozen vertices. For a finite theta function of index (p,q), reciprocity in direction n_f therefore gives

$$\operatorname{ord}_{x_f}(\vartheta_{(p,q)}) = (\ell/d_f)\operatorname{val}_{(p,q)}(\vartheta_{n_f}^\vee) = q_f + \min_{c \in \mathcal{P}_f} p \cdot c.$$

Apply the frozen valuation criterion in Lemma 1.1. Under reddening all the required theta functions are finite. More generally, for a source-reachable f, the row $B_{f,U}$ mutates as a g-vector for the opposite pattern and becomes nonnegative in a source-optimizing seed. Its opposite theta function is thus a cluster monomial. For the optimized-seed criterion with principal coefficients, see [21, Lemma 9.10(3)]. □

3.2. **Boundary polytopes and closure under raising.** The Newton polytope of $\vartheta_{e_f^*}^\vee$ is $e_f^* + \mathcal{P}_f$. Indeed, $e_f^* = (\ell/d_f)n_f$, and the tropical valuation of a theta function is homogeneous in its label [7, Lemma 2.9]. For a mutable g-vector p, the least frozen coordinate 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 $\mathbb{Z}^U_{\geqslant 0}$. The regular cone is therefore defined by the integral inequalities*

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

Proof. The support function $h_{\mathcal{P}_f}(w) = \max_{c\in\mathcal{P}_f} w \cdot c$ takes integral values on $\mathbb{Z}^U$, since reciprocity gives

$$h_{\mathcal{P}_f}(w) = -\operatorname{ord}_{x_f}(\vartheta_{(-w,0)}).$$

For a vertex c choose an integral w at which c is the unique maximizer. Such a w exists because $\mathcal{P}_f$ is rational. For every sufficiently large positive integer k, the vectors $kw + e_u$, $u \in U$, also have unique maximizer c. Hence $c_u = h_{\mathcal{P}_f}(kw + e_u) - kh_{\mathcal{P}_f}(w)$ is an integer. Each c_u is nonnegative by the description of A_f. Testing a linear inequality on the vertices suffices, which gives the formula for $\mathcal{C}^+$. □

Theorem 3.4. *In a sink-optimizing seed for i, let $m_U = -B_{i,U}$, $m_F = 0$ and $v = m - e_i$. Then*

$$g + sv \in \mathcal{C}^+ \qquad (g \in \mathcal{C}^+,\ 0 \leqslant s \leqslant g_i,\ s \in \mathbb{R}).$$

The linear map $g \mapsto g + g_i v$ retracts $\mathcal{C}^+$ onto $\mathcal{C}^+ \cap \{g_i = 0\}$. If g and s are integral, $g + sv$ indexes a regular theta function.

Proof. The opposite mutable degree defining A_i is $-m_U$. Freeze the variables indexed by $J = \{u : m_u > 0\}$ in the opposite principal-coefficient seed. The degree is now supported on frozen variables, so its theta function is a Laurent monomial. By Qin's freezing theorem [41, Theorem 5.2], setting $Y_u = 0$ for $u \in J$ leaves just the constant term of $\mathcal{F}^{\mathrm{op}}_{-m_U}$. Every nonzero $c \in A_i$ consequently has $c_u > 0$ for some $u \in J$. Thus $m_U \cdot c > 0$, and at a nonzero vertex of $\mathcal{P}_i$ Lemma 3.3 gives $m_U \cdot c \geqslant 1$.

Translation by sv changes $g_f + g_U \cdot c$ by $s(m_U \cdot c - \delta_{if})$. This change is nonnegative except at $f = i$, $c = 0$, where the new inequality is $g_i - s \geqslant 0$. This proves the first assertion. At $s = g_i$ the ith coordinate vanishes, and the map fixes every point with $g_i = 0$. The assertion about integral points follows from Corollary 3.2. □

Corollary 3.5. *In the same optimizing seed, let $\mathcal{N} = \{e^*_f + c : f \in F,\ c \in \mathrm{Vert}(\mathcal{P}_f)\}$ and $a_n = \langle v, n\rangle$. For $g \in M \cap \mathcal{C}^+$ the number of successive translations by $-v$ that stay in $\mathcal{C}^+$ is*

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

with minimum $+\infty$ if the indexing set is empty.

Proof. For $t \geqslant 0$, each normal n with $a_n \leqslant 0$ gives $\langle g - tv, n\rangle \geqslant \langle g, n\rangle \geqslant 0$. Each normal with $a_n > 0$ gives the upper bound $t \leqslant \langle g, n\rangle / a_n$. The integral values of t satisfying these bounds form the required interval. □

3.3. **Intrinsic boundary operators.** The coordinates of an integral tropical point in two adjacent seeds are related by tropical mutation. In the row convention, mutation at k sends g to g' with

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

For a mutation word w, let Φ_w be the composite of these piecewise integral-linear bijections, mutating the matrix at each step. Theta covariance identifies the regular theta indices in the two seeds. As in the introduction, 0_{cr} is the value of an undefined crystal operator, and the lattice origin indexes the function 1.

Lemma 3.6. *If a mutation word w_f reaches a source-optimizing seed for f, then $\operatorname{ord}_{x_f}(\vartheta_g) = (\Phi_{w_f} g)_f$. If such a word exists for every $f \in F$, then the regular theta indices in the starting seed t are*

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

Proof. In the source-optimizing seed, $\widetilde{b}_{fu} = (d_u/d_f)b_{fu} \geqslant 0$ for every $u \in U$. Each monomial in the pointed expansion has exponent $\Phi_{w_f} g - \nu B_{U,V}$ for some $\nu \in \mathbb{N}^U$, and the initial monomial occurs. Thus the least x_f-exponent is $(\Phi_{w_f} g)_f$. Apply the frozen valuation criterion of Lemma 1.1. □

Choose a mutation word from a reference seed to a sink-optimizing seed for i, and let Φ_i be its tropical coordinate change. In the optimizing seed write $M_i = x^{m_i}$ and $v_i = m_i - e_i$. For a regular theta index g write $a = \Phi_i g$. Define the *raising operator* and its string function by

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

Define the *lowering operator* by $l_i(g) = \Phi_i^{-1}(a - v_i)$ if $a - v_i$ is regular in the optimizing seed, and $l_i(g) = 0_{\mathrm{cr}}$ otherwise.

Theorem 3.7. *The operators r_i, l_i and the function ρ_i are independent of the chosen optimizing seed and mutation word. They commute with changes of seed coordinates. The operators are partial inverses, and*

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

In particular, ρ_i decreases by one under r_i and increases by one under l_i whenever the operator is defined.

Proof. Fix an optimizing seed and let $n = a_i$. Theorem 3.4 gives precisely n successive raising steps. Translation is injective and has the stated partial inverse.

Let $p = \Phi_i^{-1}(a + nv_i)$. Freeze the mutable variables in the support of m_i. In the pointed expansion of ϑ_g, a term with coefficient exponent ν has x_i-degree $a_i - d_i^{-1}\nu D_U(m_i)_U^{\mathsf{T}}$. Thus the freezing operator keeps precisely the part of degree a_i, which we denote by ϑ_g^0. Qin's freezing theorem identifies this part with the theta function of index a in the frozen seed [41, Theorem 5.2]. In deficient rank we may apply freezing to the principal-coefficient lift first, since all terms in each fixed x_i-degree have positive coefficients.

The theta function ϑ_p is regular, and its highest x_i-degree is zero, so freezing leaves it unchanged. Its index differs from a by nv_i, supported on variables frozen in this seed. Translation of a theta index by a frozen exponent multiplies its theta function by the corresponding monomial. It follows that

$$R_i^{(n)} \vartheta_g = x^{nv_i} \vartheta_g^0 = \vartheta_p, \qquad \deg_{R_i} \vartheta_g = n,$$

where $R_i^{(n)} = R_i^n/n!$ and $\deg_{R_i}$ is the nilpotence degree. The first equality follows by differentiating the terms of degree n in x_i.

Since R_i is canonical and the specialized theta functions are linearly independent, both n and the index p are intrinsic. For $n > 0$, any choice of optimizing seed gives a first raise of length $n - 1$ whose maximal raise is p. Among indices of length $n - 1$, maximal raising is injective, since in any one optimizing seed it is translation by $(n - 1)v_i$. This determines $r_i(g)$ independently of the choice. Its partial inverse is then independent of that choice as well. All these statements concern tropical points, and so commute with changes of seed coordinates. □

Corollary 3.8. *The* maximal raising operator $P_i(g) = r_i^{\rho_i(g)} g$ *is an idempotent map onto* $\rho_i^{-1}(0)$. *If $a = \Phi_i g$, then*

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

Proof. Iterate the translation in an optimizing seed and apply the divided-power identity in the preceding proof. □

The sink-reachable vertices for $-B$ are the source-reachable vertices for B. Proposition 1.4 identifies their regular theta functions, and Proposition 1.6 relates the indices in the two seeds. Applying the preceding construction to $-B$ therefore gives operators $r_j^\star$ and string functions $\rho_j^\star$ on $\mathcal{B}_\theta^+$. The raising length $\rho_j^\star$ is the jth pointed coordinate in a sink-optimizing seed for $-B$. In general it differs from the order of vanishing along $x_j = 0$ in Lemma 3.6. For exact boundary pairs, Section 7.3 identifies r_i and its partner $r^\star_{\sigma(i)}$ with the raising and lowering operators of a seminormal theta crystal.

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

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

where the denominator is the rational span of the mutable exchange rows. A *compatible grading* is a linear functional on $\mathbb{Q}^V$ vanishing on that span, or equivalently a functional on $L_{\mathbb{Q}}$. We call a family $(w_j)_{j\in I}$ of such functionals *adapted* if $w_j(q_i) = c_{ji}$ for all i, j. Define $\operatorname{wt}_j(g) = w_j(g)$ and $\lambda_j(g) = \rho_j(g) + \operatorname{wt}_j(g)$. The family is integral if every w_j takes integral values on $\mathbb{Z}^V$.

Corollary 3.9. *For an adapted integral family, $(r_i, l_i, \rho_i, \lambda_i, \operatorname{wt}_i)_{i\in I}$ defines an upper seminormal crystal on $\mathcal{B}_\theta^+$, with weights given in simple coroot coordinates. Adapted rational weights give the rational weight convention of* [11].

Proof. All exponents of a theta function differ by sums of exchange vectors, so every theta function is homogeneous for each compatible grading. In an optimizing seed the degree of the derivation R_i for w_j is $w_j(q_i) = c_{ji}$. Translation by v_i has the same degree. Thus $\operatorname{wt}_j(r_i g) = \operatorname{wt}_j(g) + c_{ji}$ when $r_i g$ is defined. The partial inverse and raising-length axioms follow from Theorem 3.7. By definition $\lambda_i = \rho_i + \operatorname{wt}_i$ and $\lambda_i(r_i g) = \lambda_i(g) - 1 + c_{ii} = \lambda_i(g) + 1$. □

Lemma 3.10. *Let $Q : \mathbb{Q}^I \to L_{\mathbb{Q}}$ send the ith coordinate vector to the class $[q_i]$. Adapted rational weights exist if and only if $\ker Q \subseteq \ker C$. When they exist, the space of choices is an affine space over $\operatorname{Hom}_{\mathbb{Q}}(L_{\mathbb{Q}}/\operatorname{im} Q, \mathbb{Q}^I)$.*

Proof. An adapted family is a map $W : L_{\mathbb{Q}} \to \mathbb{Q}^I$ satisfying $WQ = C$. The kernel inclusion is precisely the condition that the assignment $W(Qx) = Cx$ be well defined on $\operatorname{im} Q$. Any such map extends linearly to $L_{\mathbb{Q}}$, and the difference of two extensions factors through $L_{\mathbb{Q}}/\operatorname{im} Q$. □

Proposition 3.11. *Let $D_I = \operatorname{diag}(d_i : i \in I)$. The subspace $\ker Q$ is totally isotropic for the symmetric form with matrix $D_I C$. Adapted rational weights exist if $D_I C$ is positive semidefinite, in particular if every component of C has finite or affine type. In finite type the classes $[q_i]$ are linearly independent. Adapted rational weights also exist whenever $|I| \leqslant 2$. These assertions hold without the reddening hypothesis.*

Proof. By Section 2, $g(M_i) = -B_{i,V} + \zeta_i$, $\zeta_{ii} = 0$, and $a_{ij} = \zeta_{ji} + (d_j/d_i)\zeta_{ij}$ for $i \neq j$. For $t = (t_i)_{i\in I} \in \ker Q$, choose scalars v_u with $\sum_i t_i q_i = \sum_u v_u B_{u,V}$. Form the row vector h with coordinates v_u on U, t_i on I, and zero elsewhere. Then $hB = \sum_i t_i(\zeta_i - e_i)$. Since BD is skew-symmetric,

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

Polarization shows that the form vanishes on $\ker Q \times \ker Q$. If $D_I C$ is positive semidefinite, all isotropic vectors lie in its radical, so $\ker Q \subseteq \ker C$. Apply Lemma 3.10. In the positive definite case $\ker Q = 0$.

For two connected vertices, write $C = \left(\begin{smallmatrix} 2 & -a \\ -b & 2 \end{smallmatrix}\right)$ with $a, b > 0$. If $ab \leqslant 4$, the symmetrization is positive semidefinite. If $ab > 4$, a nonzero rational isotropic vector would force $ab(ab-4) = (ab-2)^2 - 4$ to be a square in $\mathbb{Q}$ and hence in $\mathbb{Z}$. This integer lies strictly between $(ab-3)^2$ and $(ab-2)^2$. Thus $\ker Q = 0$ in this case as well. A disconnected Cartan matrix of rank two is positive definite. □

For a subgroup H of a free abelian group G, write $\operatorname{sat}_G H = (H \otimes \mathbb{Q}) \cap G$ for its saturation. Let G be the quotient of $\mathbb{Z}^V$ by the saturation of the integral span of the mutable exchange rows, and let $L = \sum_{i \in I} \mathbb{Z}[q_i] \subseteq G$. If rational adapted weights exist, the assignment $[q_i] \mapsto Ce_i$ defines a homomorphism $L \to \mathbb{Z}^I$. It extends to G precisely when its rational extension takes integral values on $\operatorname{sat}_G L$, since the saturated subgroup is a direct summand of G. This gives a divisibility criterion in Smith normal form.

In an upper seminormal crystal, ρ_i counts the raising steps. Before an exact source partner is chosen, λ_i may be negative and may differ from the number of lowering steps. Corollary 3.5 gives that number, which may be infinite.

4. Tropical DT duality and exact boundary pairs

Exact boundary pairs are the seed-theoretic counterparts of the τ-exact pairs of [11, Section 4]. We give a criterion for exactness in terms of two principal-coefficient F-polynomials. When the mutable pattern admits a reddening sequence, theta reciprocity relates this criterion to the tropical Donaldson–Thomas transformation.

For a frozen vertex j, let $\mathbf{a}_j = B_{j,U}$. If $a = \mathbf{a}_j$ in a given seed, mutation at k gives a vector a' with coordinates

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

This is precisely the mutation rule for the negative of a mutable g-vector. It holds for every frozen vertex.

4.1. **Tropical F-polynomials and frozen coordinates.** In a fixed seed, the *projective weight* of a finite theta function with g-vector g is $b = -g_U$. Write F_b for its principal-coefficient F-polynomial. This notation is unambiguous because changing g_F multiplies the theta function by a frozen Laurent monomial. The weight b is *reachable* if it is the projective weight of a monomial in the mutable cluster variables of some seed.

For a mutable cluster monomial P of projective weight b, Proposition 1.6 gives a componentwise greatest exponent f_b of F_b, with coefficient 1, and the codegree

$$\operatorname{codeg} P = g(P) - f_b B_{U,V}.$$

These assertions hold for every reachable b, and under reddening for every integral mutable weight.

For $a \in \mathbb{R}^U$ and finite F_b, define

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

When f_b exists, also define

$$\mathsf{H}(a, b) = aD_U f_b^{\mathsf{T}} + \mathsf{E}(a, b).$$

Writing $\operatorname{Newt}(F_b)$ for the convex hull of its exponent support, these are the support functions of $-\operatorname{Newt}(F_b)$ and $f_b - \operatorname{Newt}(F_b)$, respectively, evaluated at aD_U. Both polytopes contain the origin, so their support functions are nonnegative.

Lemma 4.1. *Let $(p, q) \in \mathbb{Z}^U \oplus \mathbb{Z}^F$ index a finite theta function in the reference seed. For $j \in F$, write ν_j for the order along $x_j = 0$. Then*

$$\nu_j(\vartheta_{(p,q)}) = q_j - \frac{\mathsf{E}(\mathbf{a}_j, -p)}{d_j}. \tag{4.1}$$

In particular, if P is a mutable cluster monomial of projective weight b, then

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

Proof. A term of F_{-p} with exponent α contributes x_j-exponent $q_j - \sum_{u\in U}\alpha_u b_{uj} = q_j + d_j^{-1}\mathbf{a}_j D_U \alpha^{\mathsf{T}}$. Positivity prevents cancellation under specialization, so the minimum of these exponents gives (4.1). A mutable cluster monomial has order zero in its own seed, hence in every seed by Lemma 1.1. This gives $g(P)_j$. Proposition 1.6 then gives the codegree. □

Let i be sink-reachable and k source-reachable. Write $b_k = -g(M_k^\star)_U$ and $c_i = \operatorname{codeg}(M_i)_U$. The boundary monomial formula and its opposite give $g(M_i)_U = -\mathbf{a}_i$ and $\operatorname{codeg}(M_k^\star)_U = \mathbf{a}_k$. Consequently, in every seed,

$$c_i = -\mathbf{a}_i - f_{\mathbf{a}_i} B_{U,U}, \qquad \mathbf{a}_k = -b_k - f_{b_k} B_{U,U}.$$

The vertices (i,k) form an *exact boundary pair* if $\operatorname{codeg}(M_i) + g(M_k^\star) = e_i + e_k$. The two vertices may coincide.

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

$$\mathsf{H}(\mathbf{a}_j, \mathbf{a}_i) + \mathsf{E}(\mathbf{a}_j, b_k) = d_j(\delta_{ij} + \delta_{kj}) \qquad (j \in F).$$

These conditions can be tested in any seed and in any extended rank. Each reachable sink has at most one exact source partner, and each reachable source has at most one exact sink partner.

Proof. The negative of a g-vector and a codegree obey the same tropical mutation rule. Translation by a vector supported on frozen indices commutes with this rule. Hence exactness in one seed implies exactness in every seed. The mutable coordinates of the defining identity give $c_i = b_k$, and its frozen coordinates follow from Lemma 4.1 applied to M_i and $M_k^\star$.

For fixed i, substituting $b_k = c_i$ makes the left side independent of the choice of k. The values δ_{kj} on the right determine k. Passing to $-B$ exchanges degrees and codegrees of cluster monomials (Proposition 1.6), so the same argument gives uniqueness on the source side. □

4.2. **Mutable DT duality and the exact-pair criterion.** To define the transformation T, suppose first that the mutable pattern admits a reddening sequence. The bijection from degrees to codegrees in Proposition 1.6 commutes with translation in the frozen coordinates. It therefore induces a bijection on mutable labels. In projective coordinates, define the *mutable tropical DT transformation* T by its inverse

$$T^{-1}(b) = -b - f_b B_{U,U}.$$

For $a \geqslant 0$, the theta function of mutable degree a is a cluster monomial in the reference seed, so $F_{-a} = 1$ and $T(a) = -a$. Our T^{-1} is the action of the localized DT twist in projective coordinates. Indeed, the twist sends a pointed degree g to $-\operatorname{codeg}(\vartheta_g)$. For the boundary monomials this gives $b_k = T\mathbf{a}_k$ and $c_i = T^{-1}\mathbf{a}_i$.

Theorem 4.3. *Let $a, b \in \mathbb{Z}^U$, and write E^{op} for the support function defined from the opposite mutable pattern. If the principal theta functions of mutable degrees $-b$ in the original pattern and a in the opposite pattern are finite, then*

$$\mathsf{E}(a,b) = \mathsf{E}^{\mathrm{op}}(-b,-a).$$

For the corresponding formal principal series, the same equality holds when both directional support suprema are finite. Under reddening, for all $a, b \in \mathbb{Z}^U$ it gives $\mathsf{E}(a,b) = \mathsf{H}(b, Ta)$.

Proof. Work in the principal-coefficient seed of the mutable pattern and denote the frozen copy of u by u'. Thus $B_{u',U} = e_u$ and $d_{u'} = d_u$. Choose a common multiple ℓ of the d_u. For $g = (-b, 0)$ and $n = \ell^{-1}\sum_u d_u a_u e_{u'}^*$, the pointed expansion gives

$$\operatorname{val}_n(\vartheta_g) = \ell^{-1} \min_{\nu \in \operatorname{supp} F_b} a D_U \nu^{\mathsf{T}} = -\ell^{-1}\mathsf{E}(a,b).$$

The coordinates of n in the basis $((d_j/\ell)e_j^*)_{j\in V}$ are $(0, a)$. Its opposite mutable degree in (3.1) is therefore a. Since $\langle g, n\rangle = 0$ and $\langle g, q_u^\bullet\rangle = -b_u d_u/\ell$, that formula gives

$$\operatorname{val}_g(\vartheta_n^\vee) = -\ell^{-1} \max_{\nu \in \operatorname{supp} \mathcal{F}_a^{\mathrm{op}}} b D_U \nu^{\mathsf{T}} = -\ell^{-1}\mathsf{E}^{\mathrm{op}}(-b,-a).$$

Theta reciprocity equates the two valuations [7, Claim 4.16 and Theorem 5.19]. For formal principal series, keep the auxiliary scattering coefficients on both sides with valuation zero. The coefficientwise identity in the proof of Lemma 3.1 expresses these valuations as the negatives of the respective support suprema divided by ℓ. For integral labels the ordinary and rational broken-line valuations agree [7, Corollary 3.2], and reciprocity applies whenever both suprema are finite.

Under reddening, the theta function of degree $(-Ta, 0)$ has mutable codegree $T^{-1}(Ta) = a$. Proposition 1.6 therefore gives $\mathcal{F}_a^{\mathrm{op}}(Y) = Y^{f_{Ta}} F_{Ta}(Y^{-1})$. Its support has maximum $\mathsf{H}(b, Ta)$ in direction bD_U. □

Corollary 4.4. *Under reddening, (i, k) is an exact boundary pair if and only if $T^{-1}\mathbf{a}_i = T\mathbf{a}_k$ and*

$$\mathsf{H}(\mathbf{a}_j, \mathbf{a}_i) + \mathsf{H}(T^{-1}\mathbf{a}_i, T\mathbf{a}_j) = d_j(\delta_{ij} + \delta_{kj}) \qquad (j \in F).$$

If every frozen vertex is source-reachable, the same test holds without reddening upon substituting $c_i = \mathrm{codeg}(M_i)_U$ for $T^{-1}\mathbf{a}_i$ and $b_j = -g(M_j^\star)_U$ for $T\mathbf{a}_j$.

Proof. In Proposition 4.2, substitute $c_i = T^{-1}\mathbf{a}_i$ and $b_k = T\mathbf{a}_k$. When these two weights agree, Theorem 4.3 gives $\mathsf{E}(\mathbf{a}_j, b_k) = \mathsf{H}(T^{-1}\mathbf{a}_i, T\mathbf{a}_j)$.

For a source-reachable j, the mutable codegree of $M_j^\star$ is $\mathbf{a}_j$. Proposition 1.6 gives $\mathcal{F}_{\mathbf{a}_j}^{\mathrm{op}}(Y) = Y^{f_{b_j}} F_{b_j}(Y^{-1})$ for this cluster monomial. The first assertion of Theorem 4.3 now gives $\mathsf{E}(\mathbf{a}_j, b_k) = \mathsf{H}(b_k, b_j)$ without reddening. □

For $j \notin \{i, k\}$ both summands vanish, since each is nonnegative.

4.3. **Adding frozen vertices.** Fix the mutable pattern, its symmetrizer D_U, and a pair (i, k) with $c_i = b_k$. Write $F_i = F_{\mathbf{a}_i}$, $G_k = F_{b_k}$, and $f_i = f_{\mathbf{a}_i}$. The *admissible coefficient cone* for (i, k) is the closed rational polyhedral cone

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

Proposition 4.5. *Suppose (i, k) is exact. Extend the seed by frozen vertices, keeping the exchange entries and symmetrizing weights on the original vertices. The pair remains exact if and only if $\mathbf{a}_j = B_{j,U} \in \mathcal{C}_{i,k}$ for each new vertex j. A family of exact pairs is preserved if and only if each $\mathbf{a}_j$ lies in the intersection of their cones.*

Proof. Mutation of $\mathbf{a}_i$ or $\mathbf{a}_k$ involves only that vector and the principal matrix, so the original optimizing sequences still apply. The two boundary weights and their principal-coefficient F-polynomials are therefore unchanged, as are the exactness equations at the original vertices. At a new vertex with coefficient vector a, Proposition 4.2 requires $\mathsf{H}(a, \mathbf{a}_i) + \mathsf{E}(a, b_k) = 0$. Both terms are nonnegative. By their definitions, they vanish precisely when $a \in \mathcal{C}_{i,k}$. Applying this criterion to each pair proves the last assertion. □

Every integral $a \in \mathcal{C}_{i,k}$ occurs in such an extension. Choose $d_j = 1$, $B_{j,U} = a$ and $B_{U,j} = -D_U a^{\mathsf{T}}$, and set all new entries between frozen vertices to zero. The enlarged matrix BD is skew-symmetric.

4.4. **Comparison with boundary representations.** Consider a reference ice quiver with potential as in [11, Construction 2.5]. In particular, there are no arrows between frozen vertices in the reference seed, and we use the full quiver with potential under mutation. Write E_i and $E_j^\star$ for the boundary and dual boundary representations. For a finite-dimensional decorated representation L, let $\delta(L)$ and $\check{\delta}(L)$ be its projective and injective weights, computed from the Betti numbers in [11, (1.3)–(1.5)]. The underlying theories of mutation and presentations are due to Derksen–Weyman–Zelevinsky and Derksen–Fei [9, 8]. The relation between presentations and tropical F-polynomials is developed in [10].

Proposition 4.6. *Let i be reachable and j dually reachable in the sense of* [11, Definition 4.7]. *Then (i,j) is exact if and only if $-\check{\delta}(E_i) = \delta(E_j^\star)$. If the full Jacobian algebra is finite-dimensional, this is equivalent to $\tau^{-1}E_i \cong E_j^\star$, where τ is the Auslander–Reiten translation on its decorated representations.*

Proof. Reachability gives a seed in which E_i is simple and i is sink-optimizing. No frozen arrow enters i there, so the injective Betti numbers give $-\check{\delta}(E_i) = m_i - e_i = \operatorname{codeg}(M_i) - e_i$. Dually, in a seed where $E_j^\star$ is simple, its projective weight is $\delta(E_j^\star) = e_j - g(M_j^\star)$ [11, Corollary 4.6 and Lemma 4.19]. The boundary representations are rigid, so both signed weights obey projective-weight mutation [11, Lemmas 1.5 and 4.3]. The codegree and the negative g-vector obey the same rule. Adding a frozen coordinate vector commutes with mutation because it leaves the mutated coordinate unchanged. Thus the two formulas hold in every seed, where equality of the weights is exactly $\operatorname{codeg}(M_i) + g(M_j^\star) = e_i + e_j$.

For a finite-dimensional full Jacobian algebra, $\tau^{-1}E_i$ has projective weight $-\check{\delta}(E_i)$ [11, Section 1.1]. Both presentations are rigid, so equality of their weights is equivalent to isomorphism [8]. Compare [11, Lemma 4.29]. □

For skew-symmetrizable matrices, related formulas hold for good Jacobi-finite nondegenerate ramified quivers with potential. If mutation takes the first decorated representation to one with zero ordinary part, tropical F-polynomials compute the Hom and E-invariants of the pair [13, Proposition B.1]. For a general presentation, the corresponding formula may require a positive multiple of its weight [13, Remark B.2]. A categorical extension of Proposition 4.6 would also require a comparison of the full and mutable boundary representations.

5. Primitive theta lifts and changes of coefficients

5.1. **Primitive regular theta lifts.** Fix a seed and write a theta g-vector as $(p,q) \in \mathbb{Z}^U \oplus \mathbb{Z}^F$. For $q \in \mathbb{Z}^F$, write $x_F^q = \prod_{f\in F} x_f^{q_f}$. Frozen factorization gives $\vartheta_{(p,q)} = x_F^q \vartheta_{(p,0)}$. Thus finiteness depends only on p. Let Θ_μ be the set of p for which the principal-coefficient lift of $\vartheta_{(p,0)}$ is finite.

For $p \in \Theta_\mu$, (4.1) gives the least frozen coordinates for which $\vartheta_{(p,q)}$ is regular, namely $\beta_B(p)_f = \mathsf{E}(\mathbf{a}_f, -p)/d_f$. Each coordinate is the integer $-\nu_f(\vartheta_{(p,0)})$ and is nonnegative because F_{-p} has constant term 1. The *primitive regular theta function* of mutable label p is

$$\chi_p^B = x_F^{\beta_B(p)} \vartheta_{(p,0)}.$$

It is the unique theta function with that mutable label and with order zero at every frozen divisor. For a regular $\vartheta_{(p,q)}$, we call the vector $a = q - \beta_B(p)$ its *frozen excess*. Its fth coordinate is the order along $x_f = 0$.

Proposition 5.1. *In the fixed seed, the g-vectors of the regular finite theta functions are the vectors $(p, \beta_B(p) + a)$ with $p \in \Theta_\mu$ and $a \in \mathbb{N}^F$. The primitive regular theta functions give the decomposition*

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

The same decomposition holds for $\mathcal{U}^+(B)$ whenever the regular theta functions span that algebra. 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 particular, the g-vectors of the regular finite theta functions form a saturated additive monoid in each seed.

Proof. Equation (4.1) gives the stated parametrization. By frozen factorization, the monomials $x_F^a \chi_p^B$ for these pairs (p,a) are exactly the regular theta basis elements. This proves the direct-sum formula.

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 theta function indexed by the sum of their degrees [7, Example 2.6]. After specialization, positivity yields

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

This proves subadditivity of β_B. Homogeneity follows from $\nu_f(\vartheta_{(np,0)}) = n\nu_f(\vartheta_{(p,0)})$ [7, Lemma 2.9]. Additivity of the regular locus now follows from subadditivity. For saturation, if (np, nq) is regular and finite for some $n > 0$, then $p \in \Theta_\mu$ and $nq \geqslant \beta_B(np) = n\beta_B(p)$, so (p, q) is regular and finite. □

For a sink-reachable i and a source-reachable j, the same factorization gives

$$\deg_{R_i}(x_F^a \chi_p^B) = \deg_{R_i}(\chi_p^B) + a_i, \qquad \deg_{R_j^\star}(x_F^a \chi_p^B) = \deg_{R_j^\star}(\chi_p^B) + a_j.$$

Indeed, in the corresponding optimizing seeds these nilpotence degrees are ordinary degrees in x_i and x_j, respectively. When the theta basis is perfect for these derivations, the degrees are its upper string lengths.

Under reddening, (4.1) and Lemma 3.3 identify the two descriptions of the least frozen coordinates by the formula

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

Thus $\beta_B(p)_f$ is convex and piecewise linear with integral slopes. By Gordan's lemma, the lattice points of the rational polyhedral cone in Corollary 3.2 form a finitely generated monoid.

5.2. **Seed identifications.** A *direct seed identification* is an isomorphism Φ of function fields over $\Bbbk$ sending the variables of one seed to those of another by a permutation π. We require π to preserve the mutable and frozen index sets, their symmetrizing weights, and the extended exchange matrix $B_{U,V}$. An *opposite seed identification* has the same properties with the sign of the extended exchange matrix reversed. The matrices $B_{U,V}$ and the weights already determine the entries $B_{F,U}$ by skew-symmetrizability. Entries between two frozen indices play no part in either definition. The seeds in an identification may be any seeds in the respective mutation classes.

Proposition 5.2. *For a sink-reachable i and a source-reachable j, the boundary derivations in the target pattern satisfy*

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

under a direct identification, and

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

under an opposite identification. Both identifications give isomorphisms of the upper algebras with polynomial frozen variables. Direct identifications also carry the regular finite theta basis of one middle algebra to that of the other. For opposite identifications the same assertion holds whenever the original and opposite finite theta families agree, in particular under reddening.

Whenever the identification carries one stable perfect theta basis to the other, it intertwines the leading operators, their partial inverses, the upper string lengths, and the transported weight gradings. A direct identification sends an exact pair (i, j) to $(\pi(i), \pi(j))$, and an opposite identification sends it to $(\pi(j), \pi(i))$.

Proof. A direct identification takes a sink- or source-optimizing seed to one of the same kind, and an opposite identification exchanges the two kinds. In each case Φ takes the boundary monomial to the boundary monomial at the corresponding vertex. The identities for the derivations follow by evaluating them on the seed variables.

The isomorphism identifies corresponding Laurent rings and permutes the polynomial frozen variables, hence identifies the upper algebras. Mutation covariance of theta functions gives the basis assertion in the direct case. In the opposite case, Φ gives a direct identification

with the opposite target pattern. We may then use the equality of finite theta families, which follows under reddening from Proposition 1.4. In both cases the frozen valuations select the regular basis elements. Conjugation preserves nilpotence degrees and transports the leading term in each perfectness expansion, giving the assertions on operators and string lengths. Finally, direct identifications preserve degrees and codegrees of cluster monomials, and opposite identifications exchange them. The claims about exact pairs follow from $\operatorname{codeg}(M_i) + g(M_j^\star) = e_i + e_j$. □

Under reddening, write η_θ for the bijection determined by $\vartheta_g^B = \vartheta_{\eta_\theta(g)}^{-B}$. An opposite identification returning to the same valued ice pattern gives an *opposite cluster automorphism*. If s is its tropical map from B to $-B$, its permutation of the original theta labels is $\kappa = \eta_\theta^{-1} s$.

5.3. **Relative changes of coefficients.** Let $\gamma = (\boldsymbol{\mu}, \pi)$ consist of a mutation sequence and a permutation of the mutable vertices for which the relabeled principal exchange matrix equals the original one. Extend π by the identity on F. Give the target seed matrix $B^\gamma = \pi\boldsymbol{\mu}(B)$, symmetrizing weights $d^\gamma_{\pi(v)} = d_v$, and independent initial variables y. Denote the inverse of mutation pullback followed by relabeling by $\mathcal{T}_\gamma$. Thus the relabeled terminal cluster is sent to y, and $\mathcal{T}_\gamma(x_f) = y_f$ for each $f \in F$. We use the latter equalities to identify the two polynomial coefficient rings.

Let γ_g denote the induced change of g-vector coordinates. The mutable part of $\gamma_g(p, q)$ depends only on p and the principal pattern. Denote it by $\gamma_\mu(p)$.

Corollary 5.3. *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 boundary derivations with the same frozen index and carry primitive lifts to primitive lifts. More precisely,

$$\mathcal{T}_\gamma(x_F^a \chi_p^B) = y_F^a \chi^{B^\gamma}_{\gamma_\mu(p)}.$$

In g-vector coordinates this formula becomes

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

The isomorphisms transport the actions of exact pairs and, when the theta bases are stable and perfect, their crystals with the corresponding gradings.

Proof. Apply Proposition 5.2 to the relabeled terminal seed. The image of χ_p^B has mutable degree $\gamma_\mu(p)$ and order zero at every frozen divisor. By uniqueness it is $\chi^{B^\gamma}_{\gamma_\mu(p)}$. The formula follows by multiplication by the frozen monomial y_F^a. □

These isomorphisms compose along successive changes of coefficients. The resulting field map can depend on the mutation sequence even when the terminal exchange matrix is fixed.

A reddening sequence, with its standard mutable relabeling, returns the principal matrix to itself [30, Section 2.12]. The final c-matrix $-P$ and the matrix G of final mutable g-vectors satisfy $G^{\mathsf{T}} D_U(-P) = D_U$ [38, Theorem 1.2]. Since G is integral and unimodular, the permutation P preserves D_U. Let γ denote this transformation. After relabeling, Qin's degree–codegree formula in the principal extension identifies the mutable degree at the predecessor seed with the negative of the mutable codegree at the reference seed [40, Proposition 3.3.12]. In the convention of Section 4, this says $\gamma_\mu^{-1}(p) = -T^{-1}(-p)$, hence $\gamma_\mu(p) = -T(-p)$. The relative change of coefficients therefore sends (p, a) to $(-T(-p), a)$.

The localized theta DT twist acts by $\mathrm{DT}_\theta(g) = -\eta_\theta(g)$. Proposition 1.6 identifies its mutable action with $p \mapsto -T^{-1}(-p) = \gamma_\mu^{-1}(p)$. Since $\eta_\theta(e_f) = e_f$, this twist sends x_f to x_f^{-1} for each $f \in F$. The relative map $\mathcal{T}_\gamma$ fixes the identified polynomial coefficient ring and has target matrix B^γ.

Proposition 5.4. *Suppose the mutable pattern admits a reddening sequence. Let Φ be an opposite cluster automorphism with frozen permutation π and tropical map s from B to $-B$. Let s_μ be the mutable part of s and k the mutable part of $\kappa = \eta_\theta^{-1} s$. On regular theta functions,*

$$\Phi(x_F^a \chi_p^B) = x_F^{\pi a} \chi_{k(p)}^B, \qquad (\pi a)_{\pi(j)} = a_j.$$

Let $N(p) = -p$. For the mutable action γ_μ of normalized reddening, we have

$$k\gamma_\mu k^{-1} = \gamma_\mu^{-1}, \qquad k\gamma_\mu = N s_\mu.$$

If $k^2 = \mathrm{Id}$, then $\gamma_\mu = k(N s_\mu)$ is a product of two involutions.

Proof. Write η for the mutable part of η_θ. The opposite seed has mutable codegree map η^{-1}, so $k = \eta^{-1} s_\mu$. Negating the exchange matrix and the input to a tropical degree mutation negates its output. Starting from $-B$, the same word and relabeling thus have mutable map $N s_\mu N$. Applying Φ to $\vartheta_g^B = \vartheta_{\eta_\theta(g)}^{-B}$ and comparing the mutable degrees gives

$$\eta^{-1} s_\mu = (N s_\mu N)\eta.$$

Since $\gamma_\mu^{-1} = N\eta$, we obtain $k = \gamma_\mu(N s_\mu) = (N s_\mu)\gamma_\mu^{-1}$, or equivalently $\gamma_\mu^{-1} k = k\gamma_\mu = N s_\mu$. If $k^2 = \mathrm{Id}$, it follows that $(k\gamma_\mu)^2 = \mathrm{Id}$ and $\gamma_\mu = k(N s_\mu)$.

Finally, Φ permutes the frozen valuations by π, so it takes the unique primitive lift at p to that at $k(p)$. Multiplying by the image of x_F^a gives the formula for every regular theta function. □

Corollary 5.5. *For any seed B and $M \subseteq F$, let B_M be the seed obtained by deleting the frozen vertices outside M. Setting their variables equal to 1 sends χ_p^B to $\chi_p^{B_M}$ for each $p \in \Theta_\mu$ and induces an isomorphism*

$$\mathrm{Mid}^+(B)/(x_f - 1 : f \in F \backslash M) \xrightarrow{\sim} \mathrm{Mid}^+(B_M).$$

Write ρ_M for this specialization on $\mathrm{Mid}^+(B)$. If $i \in M$ has exact source partner j in B, it has an exact source partner in B_M precisely when $j \in M$, in which case its partner is j.

Proof. Projection away from the deleted coordinates takes each exchange vector to the corresponding one in B_M and preserves the mutable wall normals. It therefore takes each principal theta expansion to the expansion with projected label [7, Theorem 4.9]. The principal variables distinguish the terms, so the two lifts are finite for the same p. In that case their F-polynomials agree. By (4.1), $\beta_{B_M}(p)$ is the restriction of $\beta_B(p)$ to M. The free decompositions of Proposition 5.1 give the isomorphism.

For a retained vertex, both its optimizing sequences and its mutable boundary weights are unchanged. Thus Proposition 4.2 preserves the pair (i, j) if $j \in M$. Conversely, if (i, k) is exact in B_M, its mutable criterion gives $b_k = c_i = b_j$. Comparing the frozen criterion at k in the two patterns gives $\delta_{ik} + 1 = \delta_{ik} + \delta_{jk}$, hence $k = j$. □

6. Exact boundary pairs and Lie actions

For quivers with potential, the boundary representations of [11, Section 8] give Chevalley generators when the corresponding pairs are τ-exact. We derive the Chevalley–Serre relations directly from the exchange matrix. The positive generators require only sink reachability. To construct the negative generators, we characterize the source derivations that complete a given sink derivation to an $\mathfrak{sl}_2$-triple.

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 $\{x^p, x^q\} = p\Lambda q^{\mathsf{T}} x^{p+q}$. The identity $B\Lambda = D$ ensures that the bracket is log-canonical in every seed [18]. For a source boundary j, let $\zeta_j^\star = -B_{j,F}(t_j^*)$, extended by zero on U, where t_j^* is a source-optimizing seed for j. Write $Z_j^- = x^{\zeta_j^\star}$. The corresponding sink data ζ_i and Z_i^+ are as in Section 2.

Lemma 6.1. *For invertible B, 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,

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

Equations (6.1) and (6.2) also hold when B is singular.

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.

The Euler derivation $K_i = \{Z_i^+, -\}/Z_i^+$ has weight vector $\zeta_i \Lambda$. This weight is compatible because ζ_i is supported on frozen indices and $B\Lambda = D$. 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).

If B is singular, adjoin one invertible frozen variable z_a for each $a \in V$ and use the exchange matrix and symmetrizer

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

The matrix $\widehat{B}$ is invertible. The same mutation sequences optimize the original boundaries, and the lifted derivations kill the z_a. At $z_a = 1$ their multipliers and their values on seed coordinates specialize to the original ones. Positivity of the Laurent expansions shows that this specialization preserves the x_i-degree of M_j. Thus both identities descend. □

Theorem 6.2. *For distinct sink-reachable vertices i, j, the least $n \geqslant 1$ such that $(\operatorname{ad} R_i)^n R_j = 0$ is $a_{ij} + 1$. Thus the Serre exponents are determined by the boundary Cartan matrix C of Section 2. The same assertion holds for source-reachable vertices and their source Cartan matrix.*

Let i be sink-reachable and j source-reachable, and write $b = \deg_{R_i} M_j^\star$. For $i \neq j$ the least $n \geqslant 1$ such that $(\operatorname{ad} R_i)^n R_j^\star = 0$ is 1 if $b = 0$ and $b + 2$ otherwise. For $i = j$ it is 1 if $b \leqslant 1$ and $b + 1$ if $b \geqslant 2$.

Proof. We may prove vanishing after the invertible coefficient extension of Lemma 6.1. Positivity preserves the x_i-degrees of both lifted multipliers under specialization, and the lifted derivations annihilate every auxiliary variable. Their commutators therefore specialize. To prove that a commutator is nonzero, we evaluate it on a seed coordinate in the original field.

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 $(\operatorname{ad} R_i)^{a_{ij}} R_j(x_j) = R_i^{a_{ij}} M_j \neq 0$. Replacing B by $-B$ proves the assertion for sources.

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 $\mathfrak{n}_+(C) \to \operatorname{Der}_{\Bbbk}(\mathcal{U}^+)$. This extends the action of [11, Theorem 8.3] to skew-symmetrizable exchange data. The source derivations likewise give an action for their own Cartan matrix.

If the compatible gradings w_j are adapted in the sense of Lemma 3.10, the associated Euler derivations H_j commute and satisfy $[H_j, R_i] = c_{ji} R_i$. They extend the action to the positive Borel subalgebra of $\mathfrak{g}'(C)$. Exact pairs will provide these gradings together with the negative simple-root operators.

Corollary 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

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

□

6.2. **Exact pairs and their normalization.** Fix a sink-optimizing seed t_i for i and write $M_i = x^{m_i}$, $v_i = m_i - e_i$. Let j be source-reachable, with canonical monomial $M_j^\star = R_j^\star(x_j)$.

In the seed t_i, the exact-pair condition of Section 4 becomes

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

This identity can also be tested by applying R_i to the source monomial.

Proposition 6.4. *The pair (i, j) is exact if and only if $R_i M_j^\star$ is a nonzero scalar multiple of x_j for $i \neq j$, or $R_i^2 M_i^\star$ is a nonzero scalar multiple of M_i for $i = j$. In these cases the respective normalizations are*

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

Proof. In t_i, let $U_i = \{k \in U : b_{ik} < 0\}$. The principal-coefficient expansion of $M_j^\star$ is

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

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 [41, 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. Since the expansion is finite with positive coefficients on its support, specializing $y = 1$ preserves its highest x_i-degree g_i.

Every index in the support of v_i is now frozen. The factorization of theta functions by frozen monomials 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 scalar-multiple condition 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 derivative is a nonzero scalar times a single monomial. Its exponent must therefore be $g + nv_i$. Solving for g proves (6.3). □

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

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

In this notation, (6.2) reads $d_i x_i RN = d_j x_j R^\star M$.

Theorem 6.5. *Write $H = [R_i, R_j^\star]$. The pair (i, j) is exact if and only if $[H, R_i]$ is a nonzero scalar multiple of R_i. In this case H is the Euler derivation of an integral compatible grading, and $R_i, R_j^\star, H$ give an integrable $\mathfrak{sl}_2$-action on $\mathcal{U}^+$ with*

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

If $i \neq j$, then $d_i = d_j$ and $R_j^\star M_i = x_i$. In particular, nonzero multiples $sR_i, tR_j^\star$ form a Chevalley triple with their commutator if and only if (i, j) is exact and $st = 1$.

Proof. Passing from B to $-B$ exchanges degrees and codegrees of cluster monomials, even without reddening (Proposition 1.6). Thus (j, i) is exact for $-B$ whenever (i, j) is exact for B. For $i \neq j$, Proposition 6.4 applied in both patterns gives $RN = x_j$ and $R^\star M = x_i$. The mixed identity then gives $d_i = d_j$. Write d for their common value, also when $i = j$.

In the i-sink-optimizing seed, the normalization $RN = x_j$ for $i \neq j$, or $R^2N = 2M$ for $i = j$, says that the part of highest x_i-degree in N is $x_i x_j/M$. Use the invertible extension $\widehat{B}$ of Lemma 6.1, with auxiliary frozen variables z, and write $X = (x, z)$. The lifted sink multiplier is still $M = x^m$. The lifted source multiplier $\widehat{N}$ specializes to N at $z = 1$. Its Laurent coefficients are positive, so every x_i-degree occurring in $\widehat{N}$ survives specialization. Since the top term in N has coefficient one, the top part in $\widehat{N}$ is a single monomial $z^\gamma x_i x_j/M$ of coefficient one.

Write $\widehat{N} = \sum_\alpha c_\alpha X^\alpha$, let g be the exponent of $\widehat{Z}_j^-$, and let q be the exponent of the top monomial minus g. With $\widehat{\Lambda} = \widehat{B}^{-1}\widehat{D}$, Lemma 6.1 gives

$$\frac{\widehat{R}^\star(X_k)}{X_k} = \frac{1}{dx_j}\sum_\alpha c_\alpha \big((\alpha - g)\widehat{\Lambda}\big)_k X^\alpha. \tag{6.4}$$

The right side has x_i-degree at most one, with linear coefficient $z^\gamma(q\widehat{\Lambda})_k/(dM)$. For $k \neq i$, the left side is polynomial in x_i by regularity. Applying $M\partial_{x_i}$ and specializing $z = 1$ gives $H(x_k) = d^{-1}(q\widehat{\Lambda})_k x_k$ on the original coordinates $k \neq i$. For every frozen index f, including the auxiliary ones, $\widehat{R}^\star(X_f) = \delta_{jf}\widehat{N}$. Comparing linear coefficients in (6.4) gives $(q\widehat{\Lambda})_f = d\delta_{jf}$.

If $i \neq j$, then $H(x_i) = -R^\star M = -x_i$, and if $i = j$, then $H(x_i) = RN - R^\star M = 0$ by the mixed identity. Hence the vector $\widehat{w} = d^{-1}q\widehat{\Lambda} - e_i$ gives the weights of H on all original coordinates and is zero on the auxiliary ones. For $k \in U$, we have $q_k = -m_k$ and $d_k m_k = d\, b_{ki}$. Thus the kth entry of $\widehat{B}\widehat{w}^\mathsf{T} = -d^{-1}\widehat{D}q^\mathsf{T} - \widehat{B}e_i^\mathsf{T}$ vanishes, proving compatibility on the original seed.

Write $\delta = \delta_{ij}$. Since $RM = 0$, we have $HM = R(R^\star M) = (1 + \delta)M$, using $R^\star M = x_i$ for distinct vertices and $R^\star M = RN$ for $i = j$. The frozen variables x_i, x_j have weights $\delta - 1, 1 - \delta$, respectively. The source monomial N is homogeneous for the compatible grading,

and its term $x_i x_j/M$ has weight $-(1+\delta)$. In their respective optimizing seeds, $R = M\partial_{x_i}$ and $R^\star = N\partial_{x_j}$ therefore have weights 2 and -2. This proves the bracket identities.

To prove integrality, let $f \in \mathcal{U}^+$ be an eigenvector of H. By the Poincaré–Birkhoff–Witt theorem, its cyclic $\mathfrak{sl}_2$-module is spanned by $(R^\star)^b R^a f$, with $a, b \geqslant 0$. Local nilpotence bounds a and, for each such a, bounds b. The cyclic module is therefore 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 $\mathcal{U}^+$.

Conversely, suppose $[H, R] = aR$ for some nonzero scalar a. If $i = j$, the mixed identity gives $H(x_i) = 0$, so $aM = HM = R(R^\star M) = R^2 N$. Proposition 6.4 gives exactness. If $i \neq j$, write $c = (R^\star M)/x_i$. Then $H(x_i) = -x_i c$ and $H(x_j) = RN = (d_j/d_i)x_j c$. Since $Rx_j = 0$, evaluating $[H, R] = aR$ at x_j gives $Rc = 0$. Hence $HM = R(R^\star M) = Mc$, and evaluation at x_i gives $aM = 2Mc$. Thus $R^\star M = (a/2)x_i$. The same proposition applied to $-B$ makes (j, i) exact, hence (i, j) is exact for B. In either case the forward implication yields $a = 2$.

If nonzero multiples $sR, tR^\star$ form a Chevalley triple with their commutator, then $[H, R] = 2(st)^{-1}R$. It follows that the pair is exact and $st = 1$. □

The rank-one action also determines the difference of nilpotence degrees. If $Hf = mf$ and $f \neq 0$, then $\deg_{R^\star} f - \deg_R f = m$, as one sees in the simple $\mathfrak{sl}_2$-summands of the cyclic module generated by f.

Proposition 6.6. *Consider the* Langlands dual exchange datum $B^\vee = -B^{\mathsf{T}} = D^{-1}BD$, *with symmetrizer* $D^\vee = \ell D^{-1}$, *where* ℓ *is any positive common multiple of the* d_a. *Keep the same mutable and frozen index sets. 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-optimizing seed 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 6.5 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.**

Lemma 6.7. *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. Write $N = M_k^\star$ in an i-sink-optimizing seed. By exactness $g(N) = e_i + e_k - m_i$, so $g(N)_U = -(m_i)_U \leqslant 0$. The first summand of Cao's invariant $\mathsf{F}_D(M_h, N)$ vanishes because every exponent of F_{M_h} is nonnegative and its constant term is one. 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

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

The monomial N has order zero along the frozen divisor $x_h = 0$. By positivity, this order is the minimum of the indicated exponents. 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 Theorem 6.2 both give $[R_h, R_k^\star] = 0$. In the opposite seed, (k, i) is exact by Proposition 5.2. Applying the first assertion to that pair gives the source degree formula and the second commutator. □

Let I be a set of sink-reachable boundaries and assign an exact source partner $\sigma(i)$ to each $i \in I$. Write $L_i = R^\star_{\sigma(i)}$ 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 6.8. *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 the Euler derivation of an integral compatible grading, and $d_i = d_{\sigma(i)}$. The source Cartan matrix, reindexed by σ, equals C.

Proof. Uniqueness in Proposition 4.2 gives injectivity of σ. Theorem 6.5 gives the rank-one relations, the compatible integral gradings and $d_i = d_{\sigma(i)}$. Lemma 6.7 gives $[R_j, L_i] = 0$ for $j \neq i$.

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 Theorem 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. □

6.5. **Cartan type and the affine center.** In the action of Theorem 6.8, every Cartan component acts nontrivially because $R_i(x_i) = M_i \neq 0$ for each $i \in I$. The fraction field of $\mathcal{U}^+$ lies in the rational function field of a seed and hence has finite transcendence degree. Under these hypotheses, Theorem 6.10 and Corollary 6.11 show that each component has finite or affine type and that each affine component acts at level zero.

Let $C = (c_{ij})_{i,j\in I}$ be a symmetrizable generalized Cartan matrix with finite index set I. Choose a positive rational diagonal matrix D_I such that $D_I C$ is symmetric. The following lemma extends the growth estimate for Laurent polynomial derivations in [11, proof of Proposition 8.7] to arbitrary fields of finite transcendence degree.

Lemma 6.9. *Let $K/\Bbbk$ be a field extension of characteristic zero and finite transcendence degree d. Let $\mathfrak{d}$ be the Lie algebra generated by $\Bbbk$-derivations $D_1, \dots, D_r$ of K, and let $\mathfrak{d}_{\leqslant n}$ be the span of the Lie words in these generators of length at most n. Then $\dim_\Bbbk \mathfrak{d}_{\leqslant n} = O(n^d)$.*

Proof. Choose a transcendence basis $x_1, \dots, x_d$ and let $K_0 = \Bbbk(x_1, \dots, x_d)$. Adjoining the finitely many $D_a(x_j)$ to K_0 gives a finite extension L inside K. Every D_a preserves L, as one sees by differentiating a separable minimal polynomial over K_0.

Choose a finitely generated $\Bbbk$-algebra in L containing the x_j and with fraction field L. Inverting a common denominator for the derivatives of its generators gives an algebra R

preserved by all the D_a. Filter R by total degree in a finite set of algebra generators that includes the x_j. By the Leibniz rule, each D_a increases this degree by at most a fixed constant c. A Lie word of length at most n therefore sends each x_j into $R_{\leqslant 1+cn}$. Since K/K_0 is separable algebraic, a derivation of K is determined by its values on the x_j. Hence

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

by the Hilbert polynomial of the degree filtration and $\dim R = d$. □

Theorem 6.10. *Let A be a commutative domain over a field $\Bbbk$ of characteristic zero, and suppose that $\mathrm{Frac}(A)$ has finite transcendence degree over $\Bbbk$. If $\Bbbk$-derivations E_i, F_i, H_i satisfy the Chevalley–Serre relations for C and each connected component of C acts nontrivially on A, then $D_I C$ is positive semidefinite. On each affine component the canonical central element acts by zero.*

Proof. If $E_i = 0$, then $H_i = [E_i, F_i] = 0$. For every neighbor j of i, the relation $[H_i, E_j] = c_{ij} E_j$ gives $E_j = 0$. Applying the same argument along paths shows that all E_j and H_j in this component vanish. The relation $[H_j, F_j] = -2F_j$ then gives $F_j = 0$ as well. Thus $E_i \neq 0$ for every i.

Extend the derivations to $K = \mathrm{Frac}(A)$. If $D_I C$ were not positive semidefinite, it would have a negative principal minor. Factoring that minor over the connected components of its index set gives a connected principal submatrix C_J of negative determinant. In particular, C_J is nonsingular and of indefinite type. The Chevalley–Serre relations define a Lie algebra homomorphism

$$\phi : \mathfrak{g}_{\Bbbk}(C_J) \longrightarrow \mathrm{Der}_{\Bbbk}(K).$$

If $\sum_{i \in J} a_i H_i = 0$, bracketing with each $E_j \neq 0$ gives $C_J^{\mathsf{T}} a = 0$. Thus the H_i are linearly independent. Since C_J is nonsingular, its simple coroots span its Cartan subalgebra. The kernel of ϕ meets this subalgebra trivially and hence is zero [6, Section 2.1].

By Kac's finite-growth theorem, the dimensions of the principal homogeneous components of $\mathfrak{g}_{\mathbb{C}}(C_J)$ admit no polynomial bound [23, Chapter II, §6, Theorem 1]. The components of degrees m and $-m$, for $m > 0$, are spanned by brackets of m positive and m negative Chevalley generators, respectively. Their dimensions are the same over $\Bbbk$ by scalar extension from the Chevalley–Serre presentation over $\mathbb{Q}$. The injection ϕ therefore contradicts Lemma 6.9.

Now let Z be the action of the canonical central element of an affine component, and let $d = \mathrm{trdeg}_{\Bbbk} K$. Its loop realization contains, for every m, elements $u_1, \dots, u_m, v_1, \dots, v_m$ whose only nonzero brackets are $[u_j, v_j] = c$ and their negatives, where c is the canonical central element [22, Chapters 7–8]. Indeed, take opposite loop modes of a nonisotropic Cartan element and rescale the negative modes. In twisted type one uses an invariant Cartan element and degrees divisible by the order of the diagram automorphism. If a root of unity ζ is needed, we replace $\Bbbk \subset K$ by $\Bbbk(\zeta) \subset K(\zeta)$. The transcendence degree is unchanged, and the derivations extend uniquely in characteristic zero.

Write $D_1, \dots, D_{2m}$ for the resulting derivations. For $f \in K$, the matrix L with entries $L_{ab} = D_a(D_b f)$ has rank at most d over K, since the D_a span a space of dimension at most d in $\mathrm{Der}_{\Bbbk}(K)$. Hence $L - L^{\mathsf{T}} = ([D_a, D_b] f)_{a,b}$ has rank at most $2d$. If $Zf \neq 0$, the bracket relations make this alternating matrix nonsingular. Taking $m = d + 1$ gives a contradiction. □

Corollary 6.11. *Under the hypotheses of Theorem 6.10, every connected component of C has finite or affine type. If $i \neq j$, then $0 \leqslant c_{ij} c_{ji} \leqslant 4$. If $c_{ij} c_{ji} = 4$, the two vertices form an entire connected component.*

Proof. The first assertion is the classification of indecomposable symmetrizable Cartan matrices with positive semidefinite symmetrization [22, Chapter 4]. The principal minor on $\{i, j\}$ gives the upper bound. For equality, write $a = -c_{ij}$ and $b = -c_{ji}$, so $ab = 4$. The vector $v = a e_i + 2 e_j$ satisfies $v^{\mathsf{T}} D_I C v = 0$. Positive semidefiniteness implies $Cv = 0$. Thus

$ac_{ki} + 2c_{kj} = 0$ for each $k \notin \{i, j\}$. Both summands are nonpositive, so k is adjacent to neither i nor j. □

7. Perfect theta bases and crystal structures

For generic cluster characters, perfectness and normality are proved in [11, Sections 8.3–8.4]. For theta functions we first identify the terms of highest degree in a frozen variable. We then prove that the boundary derivations preserve the algebra spanned by the regular theta functions.

7.1. **The terms of highest frozen degree.** Tang proved the following perfectness statement for a full-rank extended exchange matrix when all theta functions are Laurent polynomials [42, Proposition 3.2]. See also [12, Proposition 5.11] for injective-reachable seeds of full rank. The proof below applies in arbitrary rank. It uses valuative independence to distinguish the terms of highest degree after specialization of the principal coefficients.

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 on A, then $\{\vartheta_g : g \in \mathcal{B}\}$ is a normalized perfect basis for ∂.*

For each $g \in \mathcal{B}$, the nilpotence degree $\rho_\partial(\vartheta_g)$ equals 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. Each element of A is a finite Laurent polynomial in x_a over $\Bbbk(x_v : v \neq a)$. Since $u_a = 0$, a nonzero term of negative x_a-degree survives every power of ∂ in characteristic zero. Local nilpotence therefore implies that every element is polynomial in x_a, and its nilpotence degree equals its ordinary 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 [41, 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 [7, 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.

After freezing U_a, every index in the support of v is frozen. Factorization by frozen monomials gives $x^v q_g = q_{g+v}$ in this pattern. For $g_a > 0$, the part of highest x_a-degree in $\partial \vartheta_g$ 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 $\sum_{h_a = g_a - 1} c_h q_h = g_a q_{g+v}$. The finite theta functions of the frozen pattern are linearly independent. Hence $g + v \in \mathcal{B}$, with coefficient g_a, and all other terms have $h_a \leqslant g_a - 2$. The leading map is a translation, hence injective. Applying $\partial^{g_a - 1}$ kills these other terms. Induction gives the formula for the maximal derivative. □

7.2. **Stability of the regular middle algebra.** Suppose first that the mutable pattern admits a reddening sequence. By Section 1, all principal theta functions are finite, and valuative independence after specialization gives

$$\mathrm{Mid}^{\mathrm{loc}} = \bigoplus_{g \in \mathbb{Z}^V} \Bbbk \vartheta_g, \qquad \mathrm{Mid}^{+} = \bigoplus_{g \in \mathcal{B}_\theta^+} \Bbbk \vartheta_g.$$

Theorem 7.2. *Under mutable reddening, every canonical derivation at a reachable sink or source preserves* Mid^+ *and is locally nilpotent there. Its regular theta basis is normalized perfect for each of these derivations. 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. Use the coefficient extension $\widehat{B}$ of Lemma 6.1, with the auxiliary frozen variables z inverted. Its extended exchange matrix has full column rank. The principal exchange matrix and the optimizing sequences for the original boundaries are unchanged. The lifted derivations kill z and specialize to the original ones 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 theta functions regular along the original frozen divisors form a basis of the lifted upper algebra with those variables polynomial. Theorem 2.4 gives stability, so $\widehat{R}_i \widehat{\vartheta}_{(g,0)} = \sum_{\hat{h}} c_{\hat{h}} \widehat{\vartheta}_{\hat{h}}$ is a finite expansion in regular theta functions. Setting $z = 1$ takes each summand to a regular theta function of the original pattern. The resulting finite sum 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 6.8 restrict to the stable algebra. Theta functions are homogeneous for the compatible Cartan gradings, so their integral weight decomposition restricts as well. □

More generally, let A have a theta basis perfect for each member of a family of boundary derivations D_j, $j \in J$, and let $\rho_j(g)$ be the D_j-degree of ϑ_g. Each kernel in Theorem 7.1 is spanned by a subset of this basis. Taking intersections gives, for all integers $n_j \geqslant 0$,

$$\bigcap_{j \in J} \ker(D_j^{n_j+1}|_A) = \mathrm{span}_{\Bbbk}\{\vartheta_g : \rho_j(g) \leqslant n_j \text{ for all } j \in J\}.$$

Taking all $n_j = 0$ gives a theta basis of the common invariant algebra. When the D_j are the raising operators of a finite-type exact subsystem, its basis consists of the highest-weight theta functions for that subsystem. Compare [11, Corollary 8.11] and [12, Proposition 5.5].

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. Equality of the localized cluster and upper cluster algebras gives a second criterion, valid without reddening.

Corollary 7.3. *If* $\mathcal{A}^{\mathrm{loc}} = \mathcal{U}^{\mathrm{loc}}$*, the theta functions with finite principal-coefficient lifts form a basis of* $\mathcal{U}^{\mathrm{loc}}$*, and their regular subfamily forms a basis of* $\mathcal{U}^+$*. This regular basis is normalized perfect for every reachable sink derivation.*

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 [7, 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 [36, 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?

7.3. **Normal theta crystals.** Let A be a subalgebra of $\mathcal{U}^+$ with 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, since differences of exponents in a pointed expansion belong to the span of the mutable exchange vectors, which all H_i 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) be an $\mathfrak{sl}_2$-triple acting 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 $\lambda(b) - \rho(b) = m$, where $\rho(b) = \deg_E b$ and $\lambda(b) = \deg_F b$. The two leading maps are partial inverses.*

Proof. Decompose the integrable module into finite-dimensional simple $\mathfrak{sl}_2$-modules. Filter it by the sums $V^{\leqslant N}$ of the summands of highest weight at most N. On weight m, for $N \geqslant |m|$ of the same parity, this filtration step is $\ker E^{(N-m)/2+1}$. It also equals $\ker F^{(N+m)/2+1}$ on the same weight. Thus the basis is adapted to the F-kernels and induces a basis of each successive quotient.

Let b have weight m and highest nonzero isotypic component of highest weight N. Then $\rho = (N-m)/2$ and $\lambda = (N+m)/2$, as each vector has finite support in the direct sum. If $\lambda > 0$, E from weight $m-2$ to weight m on this quotient is an isomorphism. By E-perfectness, this isomorphism sends each induced basis vector to $\rho + 1$ times an induced basis vector. It therefore determines a bijection of the two bases. On this quotient EF acts on weight m as $(\rho+1)\lambda$, so F acts by the inverse bijection with coefficient λ. The remaining terms have weight $m-2$ and highest weight at most $N-2$, hence lowering length at most $\lambda - 2$. If $\lambda = 0$, then $Fb = 0$. This gives both normalized perfectness and the inverse arrows, as in the finite-dimensional case of [12, Lemma 5.13]. □

Write r_i, l_i for the leading maps of E_i, F_i on theta labels. Let ϖ_i, $i \in I$, denote the standard basis of the coroot-coordinate lattice $\mathbb{Z}^I$. The weight of a label b is $\mathrm{wt}(b) = \sum_i \mathrm{wt}_i(b)\varpi_i$, where $H_i\vartheta_b = \mathrm{wt}_i(b)\vartheta_b$. Let $\bar{\alpha}_j = \sum_i c_{ij}\varpi_i$. 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 [11, Section 4.4]. The functions ρ_i and λ_i count the successive applications of r_i and l_i, respectively. Thus the crystal is *seminormal*. In an i-sink-optimizing seed,

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

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

Let J be an affine Cartan component and let $a_i^\vee > 0$ be the coefficients of its canonical central element in the simple coroots. By Theorem 6.10, $\sum_{i\in J} a_i^\vee H_i = 0$. If b is highest in every color of J, then $\rho_i(b) = 0$ and $\mathrm{wt}_i(b) = \lambda_i(b)$ for each $i \in J$. Thus $\sum_{i\in J} a_i^\vee \lambda_i(b) = 0$. Every lowering length $\lambda_i(b)$ is nonnegative, so each vanishes for $i \in J$. Hence b is isolated in the restriction to J.

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 theta functions and the Chevalley derivations have rational coefficients in every seed. In a finite theta expansion, linear independence therefore implies that an element with rational Laurent coefficients has rational theta coefficients. Thus the derivations preserve the $\mathbb{Q}$-span of the given basis, and we may extend this span to $\mathbb{C}$ without changing any leading arrow. For finite-type J, the restricted integrable module is a direct sum of finite-dimensional simple modules. The union of their specialized global bases is a perfect basis with the stated

component crystals [26, 34]. The perfect-basis comparison theorem identifies this crystal with that of the theta basis [3, Lemma 5.50 and Theorem 5.55]. □

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 6.8 gives the integrable action on $\mathcal{U}^+$. Theorem 7.1 gives perfectness for the raising operators. Lemma 7.4 gives perfectness for the lowering operators and identifies the two leading maps as partial inverses. Apply Corollary 7.5 and, for the final assertion, Corollary 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 $l_i = r^\star_{\sigma(i)}$. The commutators $[R_i, R^\star_{\sigma(i)}]$ act by the integral coroot weights.*

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-optimizing seed. Translation in this seed gives

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

for $0 \leqslant n \leqslant \rho_i(g)$ and $0 \leqslant n \leqslant \lambda_i(g)$, respectively. Proposition 4.2 tests whether a given source is the exact partner of i.

Recall from Section 3.4 that $[q_i]$ is the degree of R_i in the universal rational grading $L_{\mathbb{Q}}$. When can we extend the coroot grading so that the simple roots are linearly independent?

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,*

$$\operatorname{rank}\begin{pmatrix} B_{U,V} \\ (q_i)_{i\in I} \end{pmatrix} = \operatorname{rank}(B_{U,V}) + |I|.$$

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 ambient weight space has dimension $2|I| - \operatorname{rank} C$. 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. □

Adjoining frozen variables can preserve the exact pairs only under the coefficient constraints of Proposition 4.5. Such extensions need not make the simple-root degree classes independent, as the rank-two Lagrangian Grassmannian in Section 9.3 shows.

8. Symmetries of theta crystals

8.1. **Weyl group actions and Lusztig involutions.** Let $A \subseteq \mathcal{U}^+$ have a basis of finite theta functions, and suppose the exact-pair operators E_i, F_i preserve A. Write r_i, l_i for the raising and lowering maps of the normal crystal of Section 7.3. The *string reflection* in color i 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-sink-optimizing seed, $g(s_i b) = g(b) - \mathrm{wt}_i(b)(m_i - e_i)$. On each finite-type rank-two restriction, normality identifies these reflections with Kashiwara's Weyl group action on a union of highest-weight crystals [27, Theorem 11.1]. Hence they satisfy every finite braid relation and define an action of the Weyl group on theta labels. Tropical transport gives this action in every seed.

Local nilpotence also defines algebra automorphisms $\dot{s}_i = \exp(E_i)\exp(-F_i)\exp(E_i)$. For h in the coroot lattice $\bigoplus_{i\in I}\mathbb{Z}h_i$ and a vector a of weight ν, define $\Theta_h(a) = (-1)^{\nu(h)}a$. The standard Tits relations [22, 24] 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 [11, Theorem 8.16], applied to the present integrable module algebra.

Let $A^{(i)}_{\leqslant N}$ be the sum of the ith $\mathfrak{sl}_2$-summands of highest weight at most N, and set $N_i(b) = \rho_i(b) + \lambda_i(b)$. The action of $\dot{s}_i$ on the associated graded is

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

This is the perfect-basis compatibility in [11, 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.

For a set $J \subseteq I$ of finite Cartan type, the longest Weyl group element $w_{0,J}$ defines opposition on J by $w_{0,J}\alpha_i = -\alpha_{i^*}$. 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 $\xi_\lambda r_i = l_{i^*}\xi_\lambda$ [31, Section 2.4].

Write ξ_J for the involution acting by ξ_λ on each connected J-component. For the crystal of supported presentation weights, the corresponding factorization appears in [11, Proposition 6.5].

Proposition 8.1. *Suppose the exact pairs $(i,\sigma(i))$, $i \in J$, give a normal J-crystal on $\mathcal{B}^+_\theta$. Let Φ be an opposite cluster automorphism preserving this basis, with action κ and frozen permutation π. Then*

$$\kappa r_i = l_{i^*}\kappa \quad (i \in J) \qquad\Longleftrightarrow\qquad \pi(i) = \sigma(i^*) \quad (i \in J).$$

Under these equivalent conditions there are unique permutations p_λ of $\mathcal{H}_\lambda$ such that

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

Thus $\kappa = \xi_J$ precisely when κ preserves every connected J-component. Under mutable reddening, the regular basis is preserved automatically. If $\kappa = \xi_J$, its induced mutable map ξ_{mut} then satisfies

$$\xi_{\mathrm{mut}}\gamma_\mu\xi_{\mathrm{mut}} = \gamma_\mu^{-1}, \qquad \gamma_\mu = \xi_{\mathrm{mut}}\circ(Ns_\mu), \qquad (Ns_\mu)^2 = \mathrm{Id}\,.$$

Here s_μ is the mutable part of the tropical map of Φ from B to $-B$, $N(p) = -p$, and γ_μ is normalized reddening on mutable labels.

Proof. By Proposition 5.2, $\kappa r_i = r^\star_{\pi(i)}\kappa$, while $l_{i^*} = r^\star_{\sigma(i^*)}$ by exactness. Distinct source operators are distinguished by their string lengths on the frozen variables. Indeed, in a source-optimizing seed $R^\star_j$ kills x_f for $f \neq j$ and sends x_j to its canonical monomial, which it kills. Thus the source string length of x_f in color j is δ_{jf}. This proves the equivalence. Taking partial inverses now gives $\kappa l_i = r_{i^*}\kappa$. Thus $\xi_J\kappa$ commutes with all J-arrows. It preserves their string lengths and hence their differences, the restricted coroot coordinates. It therefore permutes the highest elements of each weight λ. Every element of a component is obtained from its highest element by lowering operators, so this permutation determines $\xi_J\kappa$ throughout the component. This proves the formula and the component criterion.

Under reddening, Proposition 5.2 gives basis preservation. Every mutable label has a primitive regular lift, so $\kappa = \xi_J$ implies $\xi_{\mathrm{mut}}^2 = \mathrm{Id}$ on the whole mutable lattice. Proposition 5.4 gives the asserted identities. □

In particular, $\kappa^2 = \mathrm{Id}$ if and only if each p_λ is involutive. A grading preserved by Φ forces $p_\lambda = \mathrm{Id}$ if it is constant on J-components and distinguishes their highest elements of weight λ.

8.2. **The relative DT square.** Assume mutable reddening and choose exact pairs $(i, \sigma(i))$, $i \in I \subseteq F$. Write $I' = \sigma(I)$. For $M \subseteq F$, let B_M be the seed obtained from B by deleting the frozen vertices outside M. The specialization ρ_M sets the corresponding variables equal to 1 as in Corollary 5.5.

Starting from B_I, apply the inverse of the normalized reddening transformation twice and then relabel each frozen vertex i as $\sigma(i)$. This gives a direct cluster isomorphism $\Psi : B_I \to B_{I'}$, the *relative DT square*. Indeed, the negative of the frozen exchange vector $\mathbf{a}_i = B_{i,U}$ mutates as a mutable g-vector. Since $T(p) = -\gamma_\mu(-p)$, two applications of the inverse send $\mathbf{a}_i$ to $T^{-2}\mathbf{a}_i = \mathbf{a}_{\sigma(i)}$ by Corollary 4.4. Theorem 6.8 gives $d_i = d_{\sigma(i)}$, so the relabeling also preserves the symmetrizing weights. On mutable theta labels, Ψ acts by γ_μ^{-2}.

Fix an opposite cluster involution K of B_I. Write ω for its permutation of the frozen vertices and k for its action on mutable theta labels. A *lift* of ΨK is an opposite cluster involution Φ of B satisfying $\rho_{I'}\Phi = \Psi K \rho_I$.

For $S \subseteq I$ of finite Cartan type, let $i \mapsto i^*$ denote opposition on S, and let $A_S^+ = \bigcap_{i \in S} \ker E_i$ be the algebra of S-highest vectors in Mid^+. Write n_S for the representative of the longest Weyl group element $w_{0,S}$ obtained from $n_i = \dot{s}_i^{-1}$ along a reduced word.

Theorem 8.2. *If $F = I \cup I'$, then ΨK has a unique lift.*

For any lift Φ, the induced action on regular theta labels is Lusztig's involution ξ_S if and only if $\omega(i) = i^*$ for all $i \in S$ and Φ agrees with n_S on A_S^+. If such a lift exists, it is unique and its action $\xi_{S,\mathrm{mut}}$ on mutable theta labels satisfies

$$\gamma_\mu^{-2} = \xi_{S,\mathrm{mut}} k, \qquad \xi_{S,\mathrm{mut}} = \gamma_\mu^{-1} k \gamma_\mu. \tag{8.1}$$

Proof. If a lift exists, specialization determines its frozen permutation π on I, and reversal of exact pairs determines it on I' by uniqueness of their partners. Thus

$$\pi(i) = \sigma(\omega(i)), \qquad \pi(\sigma(i)) = \omega(i) \qquad (i \in I). \tag{8.2}$$

These prescriptions agree on $I \cap I'$. Indeed, if $f = \sigma(i) \in I$, then (i, f) survives in B_I. Its opposite image under K is $(\omega(f), \omega(i))$ by Proposition 5.2, so Corollary 5.5 gives $\sigma(\omega(f)) = \omega(i)$. Thus (8.2) defines a weight-preserving involution on $I \cup I'$.

Specialize all frozen variables to 1. Now K reverses γ_μ, and ΨK acts on the common coefficient-free pattern by the involution $b = \gamma_\mu^{-2} k = \gamma_\mu^{-1} k \gamma_\mu$. Let s be the tropical degree change of ΨK to the opposite pattern and write $N(a) = -a$. By Proposition 5.4, $Ns = b\gamma_\mu$. Thus $L = sN$ is an involution, since b reverses γ_μ. For a word representing ΨK to give an opposite identification with frozen permutation π, the condition at f is $\mathbf{a}_{\pi(f)} = L(\mathbf{a}_f)$. It holds on I by construction. In particular $L(\mathbf{a}_{\omega(i)}) = \mathbf{a}_{\sigma(i)}$, giving the required condition on I' as well. These vectors and their symmetrizing weights determine the frozen entries of the extended exchange matrix. Hence every weight-preserving involution extending the prescribed permutation and satisfying this condition gives an opposite identification. Mutation commutes with setting the other frozen variables to 1, giving the required specialization. By Proposition 5.4, its action on primitive theta functions is determined by b and its action on the frozen variables by π. These data determine the field map, since the seed variables are regular theta functions, and its square is the identity. In particular coverage gives existence and uniqueness.

Now let Φ be a lift. By Proposition 8.1 and (8.2), it sends each raising arrow of color $i \in S$ to a lowering arrow of color i^* precisely when $\omega(i) = i^*$ for all $i \in S$. It then acts

by ξ_S precisely when it takes each highest theta function to the lowest one in the same component. We claim that for an S-highest label g this lowest function is $n_S\vartheta_g = \vartheta_{w_{0,S}g}$. Indeed, on an i-highest vector of weight ν, n_i acts by $F_i^{(\nu(h_i))}$. At each step of a reduced word for $w_{0,S}$, the inverse of the preceding Weyl element takes the next simple root to a positive root. The current vector is therefore i-highest for that step, and the divided-power formula of Theorem B for the source derivation takes it to the opposite end of its string. This proves the claim. Since the highest theta functions form a basis of A_S^+, the component condition is exactly agreement with n_S on A_S^+.

Two lifts acting by ξ_S agree on every seed variable, hence on the function field. For such a lift $b = \xi_{S,\mathrm{mut}}$, which gives (8.1). □

Corollary 11.2 reduces the comparison on A_S^+ to a finite set of highest labels. Once Φ and n_S agree there, it suffices to check opposition for the colors $i \in S \cap \sigma(S)$ with $\sigma(i) \in S$. In particular, if $S \cap \sigma(S) = \varnothing$, agreement on A_S^+ alone characterizes the Lusztig lift.

To see the reduction, fix $i \in S$. If $\sigma(i) \notin S$, the variable $x_{\sigma(i)}$ is S-highest, with lowering lengths δ_{ij} for $j \in S$. The lowest element of its component has raising lengths δ_{i^*j}. Since $\Phi(x_{\sigma(i)}) = x_{\omega(i)}$ is that lowest element, its raising lengths $\delta_{\omega(i)j}$ give $\omega(i) = i^*$. If $i \notin \sigma(S)$, then x_i is S-lowest. Since Φ sends the highest element of its component to x_i and $\Phi^2 = \mathrm{Id}$, its image $x_{\sigma(\omega(i))}$ is the highest element of that component. Comparing lowering lengths again gives $\omega(i) = i^*$.

When $I \cup I' \neq F$, a lift exists precisely when the involution $L = sN$ from the proof preserves the multiset of uncovered exchange vectors at each symmetrizing weight. Match the vertices over the two vectors of each two-element orbit by a bijection, and choose an involution on the vertices over each fixed vector. These choices give all lifts. When a lift exists, it is unique precisely when the uncovered vectors of each weight are distinct. For a Lusztig lift, every x_f with $f \notin S \cup \sigma(S)$ must be fixed, since all E_i, F_i for $i \in S$ kill it. Thus a lift fixing the uncovered vertices exists precisely when $L(\mathbf{a}_f) = \mathbf{a}_f$ for each of them. The opposition and highest-vector tests then decide whether it acts by ξ_S.

Consider the seed $xx' = 1 + fg$, with f, g sinks in the x-chart. The pairs (f, g) and (g, f) are exact. Their sink multipliers are x, their source multipliers are x', and $R_f x' = g$, $R_g x' = f$. Take both pairs. For $S = \{f, g\}$ the Cartan type is $A_1 \times A_1$, with trivial opposition. The common highest algebra is $\Bbbk[x]$. Its elements are Laurent polynomials in x alone, and a negative power of x introduces a denominator $1 + fg$ in the x'-chart. Here Ψ interchanges f, g and fixes x, x', and $n_S x = F_f F_g x = x'$. The opposite involution exchanging x, x' and fixing f, g agrees on the highest algebra, but (8.2) gives $\omega(f) = g$. For $S = \{f\}$ the opposite involution exchanging both x, x' and f, g has $\omega = \mathrm{Id}$. On the highest element x its value is x', whereas $n_f x = F_f x = f$.

Corollary 8.3. *If $I = F$, then Ψ is a direct automorphism of the full pattern sending color i to $\sigma(i)$. Every opposite cluster automorphism reverses Ψ on the theta basis, and $\Psi\xi_S\Psi^{-1} = \xi_{\sigma(S)}$. If ξ_S has an opposite cluster realization, then $\xi_{\sigma^j(S)}\xi_S = \Psi^{2j}$ for every $j \in \mathbb{Z}$.*

Proof. An opposite cluster automorphism with frozen permutation π reverses the exact pairs. Their uniqueness gives $\pi\sigma\pi^{-1} = \sigma^{-1}$. The same automorphism conjugates γ_μ to its inverse on mutable labels, so the formulas on primitive lifts in Corollary 5.3 and Proposition 5.4 give reversal on the full basis. Direct naturality carries each S-component and its extremal elements to a $\sigma(S)$-component. This proves the conjugation formula, whose jth iterate gives the last identity when ξ_S reverses Ψ. □

9. Geometric applications

9.1. **Maximal unipotent subgroups.** Let G be a simply connected complex semisimple group. Choose a maximal torus T, positive roots and a pinning. Write U_G for the corresponding maximal unipotent subgroup, e_i for the positive Chevalley generators, and $x_i(t) = \exp(te_i)$

for the simple root subgroups. We denote the simple roots, simple coroots and fundamental weights by α_i, $\alpha_i^\vee$ and ϖ_i. Left and right multiplication define derivations on $\mathbb{C}[U_G]$ by

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

In a fundamental representation, let v_δ be the pinned extremal vector of weight δ and v_γ^* the weight covector dual to the pinned vector of weight γ. The *generalized minor* $\Delta_{\gamma,\delta}$ is the matrix coefficient $u \mapsto v_\gamma^*(uv_\delta)$ [14, Section 1.4]. Under the conjugation action $(h \cdot f)(u) = f(h^{-1}uh)$ of T, this minor has weight $\delta - \gamma$. Both derivations raise the weight by α_i. Tang proved that the regular theta basis of $\mathbb{C}[U_G]$ is biperfect for these two families [42, Theorem 1.1]. Its bicrystal is therefore Kashiwara's $B(\infty)$ [25, Theorem 1.7].

Let w_0 be the longest Weyl group element and define i^* by $w_0\alpha_i = -\alpha_{i^*}$. On each simple factor, normalize the invariant inner product so that the shortest roots have squared length 2. Every position with letter i in a reduced-word seed has symmetrizing weight $d(\alpha_i) = (\alpha_i, \alpha_i)/2$, including the last occurrence, with frozen variable $A_i = \Delta_{\varpi_i, w_0\varpi_i}$. With the chosen pinning, D_i^R is the sink derivation at A_{i^*} and D_i^L the source derivation at A_i [42, Lemmas 5.1 and 5.2]. We index the derivations by their simple roots.

Proposition 9.1. *The Cartan matrix of each family of boundary derivations is $C = (\langle \alpha_i^\vee, \alpha_j \rangle)_{i,j}$. For each family the simple-coroot coordinates of the conjugation weight form the unique adapted compatible grading.*

Let ι be the pinned antiautomorphism of U_G fixing each $x_i(t)$. Its pullback is an opposite cluster automorphism of $\mathbb{C}[U_G]$. It permutes the regular theta basis and induces Kashiwara's involution on $B(\infty)$.

Proof. Left and right translation are faithful at the level of infinitesimal actions. Hence the least Serre exponent of D_i^L on D_j^L is $1 - c_{ij}$ for $i \neq j$, and likewise on the right. Theorem 6.2 identifies both boundary Cartan matrices with C.

The extended exchange matrix has full column rank by compatibility with the skew form of [39, Theorem 8.1.2]. There is one frozen variable per simple root, so $\dim L_{\mathbb{Q}} = \operatorname{rank} G$ in the notation of Section 3.4. For either family, the conjugation grading sends the $[q_i]$ to the independent simple roots α_i. Thus the $[q_i]$ form a basis, and the simple-coroot coordinates give the unique adapted grading.

Choose a reduced word $\mathbf{j} = (j_1, \dots, j_N)$ for w_0. Write s_j for the simple reflection at j, $v_0 = 1$ and $v_k = s_{j_1} \cdots s_{j_k}$. The *prefix minor seed* and the *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 each word the two seeds are mutation-equivalent by the construction of Goodearl and Yakimov [20, Theorem 8.2(c)], in the minor coordinates of [39, Theorems 8.1.2 and 8.2.2]. Changing the reduced word preserves this mutation class [12, Lemma 3.8]. Let $\mathbf{j}^\dagger = (j_N^*, \dots, j_1^*)$ and $\rho(k) = N + 1 - k$. The prefix of length $N - k$ in $\mathbf{j}^\dagger$ has product $v_k w_0$. The involution identity for generalized minors [14, Proposition 2.7] gives

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

Reversal interchanges the previous and next occurrences of each letter. The exchange formulas [39, Remarks 8.1.4 and A.3.1] therefore give, for every mutable k and every a,

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

Reversal also sends last occurrences to first occurrences and preserves the symmetrizing weights. Thus ι^* identifies the prefix seed with the opposite suffix seed and sends A_j to A_{j^*}.

These seeds are injective-reachable [39, Lemma 9.1.3], so Proposition 5.2 gives preservation of the regular theta basis. Since ι fixes the simple root subgroups and reverses multiplication,

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

The induced map on $B(\infty)$ fixes the element corresponding to 1 and exchanges the two crystal operators of each color. These properties characterize Kashiwara's involution [25, Theorem 1.7 and Remark 1.8]. □

If σ is the tropical map of ι^* from the reference seed to its opposite, its action on the original theta g-vectors is $\eta_\theta^{-1}\sigma$ by Proposition 5.2. For the categorical construction in simply-laced type, see [11, Proposition 10.1].

9.2. **Base affine spaces.** Let U_G^- be the opposite maximal unipotent subgroup. On the *base affine space* $X = U_G^-\backslash G$, the left action of T grades $\mathbb{C}[X]$ by the monoid P_+ of dominant integral weights. The right action of G identifies each graded summand with the simple module $L(\lambda)$ of highest weight λ. Thus

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

For a right T-weight vector $f \in R_\lambda$ we call λ its degree and denote its right weight by $\mathrm{wt}(f)$. Francone's minimal monomial lift of a reduced-word seed realizes $\mathbb{C}[X]$ as an upper cluster algebra with polynomial frozen variables [17, arXiv version, Theorem 8.3.2 and Lemma 8.3.4]. These lifted seeds belong to one mutation class [12, Corollary 3.9]. Their mutable pattern is that of U_G, so it admits a reddening sequence. The regular theta basis is homogeneous for both torus actions [12, Lemma 5.1 and Theorem 5.4].

Write E_i, F_i for the derivations of $\mathbb{C}[X]$ given by the right action of the pinned Chevalley generators. On R_λ, the *Lusztig involution* Ξ_λ exchanges the pinned highest and lowest vectors and satisfies $\Xi_\lambda E_i = F_{i^*}\Xi_\lambda$ and $\Xi_\lambda F_i = E_{i^*}\Xi_\lambda$ [31, Section 2.4]. Morier-Genoud computes this involution in parametrizations of the canonical basis [35, Corollary 2.17]. To describe its action on individual theta functions, we first identify the pinned derivations with boundary derivations. Following [11, Section 10.3], index the two families of frozen minors by

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

We keep the weights of the word positions and give p_i the weight of q_i. Since q_i has letter i^* and opposition preserves root lengths, $d_{p_i} = d_{q_i} = d(\alpha_i)$.

Proposition 9.2. *For each 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 their Cartan matrix is that of the factor. For every $\lambda \in P_+$, including on factors of type A_1, the theta basis of R_λ is normalized perfect for the pinned Chevalley derivations and its crystal is $B(\lambda)$.*

Proof. Suppose G is simple of rank at least two. Use the lifted prefix minors x_k of a reduced word $\mathbf{j} = (j_1, \dots, j_N)$ for w_0, with partial products v_k as in Section 9.1. Write $s(k)$ for the next occurrence of j_k, or $+\infty$ if there is none. An initial letter i must occur again, since otherwise $w_0\varpi_i = s_i\varpi_i$ would pair positively with a neighboring coroot, contrary to antidominance. By reversal, a final letter also occurs earlier.

Take a word ending in i^* and denote the preceding occurrence of i^* by N^-. On U_G, the pinned right derivative is $x_{N^-}\partial_{x_N}$ and $\widetilde{b}_{Nk} = -\delta_{k,N^-}$ for every mutable k [42, Lemma 5.1 and equation (15)]. The two minors have the same left degree ϖ_{i^*}, and E_i kills each highest minor x_{p_j}. Thus $E_i = x_{N^-}\partial_{x_N}$ also in the lifted seed, where $x_N = x_{q_i}$. The two positions have equal symmetrizing weight. Hence $b_{q_ik} = -\delta_{k,N^-}$ and $E_i = R_{q_i}$.

For F_i, take a word beginning with i. The root $v_k^{-1}(-\alpha_i)$ is positive for every $k \geqslant 1$, so F_i kills each extremal minor x_k. The pinned fundamental i-string gives $F_i x_{p_j} = \delta_{ij}x_1$. Homogeneity in left degree ϖ_i gives $\widetilde{b}_{p_ik} = -\sum_{r:j_r=i}\widetilde{b}_{rk}$ for each mutable k. If $j_k = i$, the predecessor and successor of k contribute 1 and -1 to the sum, with the predecessor absent only at $k = 1$. If $j_k \neq i$, there are no contributions unless i occurs in $(k, s(k))$. The prefix exchange rule then gives opposite contributions at the predecessor of the first such occurrence and at the last such occurrence. The predecessor exists because the word begins with i [39,

Remark 8.1.4]. Thus $\widetilde{b}_{p_i k} = \delta_{k1}$. The first position has letter i and hence the same weight as p_i. Thus $b_{p_i k} = \delta_{k1}$ and $F_i = R^{\star}_{p_i}$.

The pinned Chevalley relations give exactness by Theorem 6.5. Since $\mathrm{Lie}(G)$ acts faithfully, its Serre exponents determine the boundary Cartan matrix by Theorem 6.2. Theorem 7.1 and Corollary 7.5 give normalized perfectness and normality. The Chevalley derivations preserve each R_λ, so its theta basis has both properties. The module is irreducible of highest weight λ, hence its crystal is $B(\lambda)$. See also [12, Corollaries 5.12 and 5.15].

In type A_1 there are no mutable vertices. The degree-m monomials in the highest and lowest coordinates of the standard representation form the usual normalized perfect theta basis of $L(m\varpi_1)$ [11, Remark 10.6]. Tensoring over the simple factors gives the result for G. □

Write $Q = \{q_i\}$ and $P = \{p_i\}$. On each simple factor of rank at least two let $\Psi : B_Q \to B_P$ be the relative DT square for (q_i, p_i). On a factor of type A_1 use the relabeling $q_1 \mapsto p_1$ between the two one-variable patterns. Use the product of these maps for G. The specialization ρ_Q recovers the unipotent prefix seed. Indeed, on the open cell hu, a lifted prefix minor is its restriction to U_G multiplied by its fundamental weight evaluated at h [17, arXiv version, Lemma 8.3.4].

Proposition 9.3. *The pinned involution $\Xi = \bigoplus_{\lambda \in P_+} \Xi_\lambda$ is a degree-preserving opposite cluster automorphism. It exchanges x_{p_i} with $x_{q_i *}$ and acts on the theta indices of R_λ by ξ_λ. For the pinned antiautomorphism ι of U_G in Proposition 9.1, we have*

$$\rho_P \Xi = \Psi\, \iota^* \rho_Q.$$

Let k_∞ be the mutable action of ι^*. In the common mutable coordinates the ξ_λ induce a single involution ξ_{mut} satisfying

$$\gamma_\mu^{-2} = \xi_{\mathrm{mut}} k_\infty, \qquad \xi_{\mathrm{mut}} = \gamma_\mu^{-1} k_\infty \gamma_\mu.$$

Proof. On a simple factor of rank at least two, x_{q_i} restricts to A_{i*} on U_G. Thus ι^* takes q_i to q_{i*} by Proposition 9.1. The pairs (q_i, p_i) of Proposition 9.2 cover the frozen vertices. Theorem 8.2 therefore gives a unique lift of $\Psi\iota^*$. By (8.2), it takes x_{p_i} to $x_{q_i *}$. Let S be the full set of simple colors. In degree $\lambda = \sum_i m_i \varpi_i$, the highest-weight line is spanned by $\prod_i x_{p_i}^{m_i}$, so the minors x_{p_i} generate A_S^+. On each generator the lift agrees with n_S, since the pinning gives $n_S x_{p_i} = \Delta_{\varpi_i, w_0 \varpi_i} = x_{q_i *}$ [14, Section 1.4]. Since ι^* induces opposition on Q, Theorem 8.2 identifies the action of the lift on theta labels with Lusztig's involution. By Proposition 5.2, the lift intertwines F_i with E_{i*}. It therefore agrees with Ξ_λ on all of R_λ, which is generated from the highest vector by the F_i. The specialization follows from the lifting statement, and the two mutable identities are (8.1) with $k = k_\infty$.

In type A_1 the mutable lattice is zero and ι^* is the identity on $\mathbb{C}[U_G]$. Exchanging $p = x_{p_1}$ and $q = x_{q_1}$ sends $p^{m-j} q^j$ to $p^j q^{m-j}$ in degree m. Taking the product over the simple factors proves the proposition. □

In particular $\xi_{\mathrm{mut}} \gamma_\mu \xi_{\mathrm{mut}} = \gamma_\mu^{-1}$, and $\gamma_\mu = \xi_{\mathrm{mut}} \circ (-s_\mu)$ for the mutable tropical coordinate change s_μ from the lifted seed to its opposite, by Proposition 5.4.

Let $\mathcal{B}^+_{\theta,\lambda}$ be the set of g-vectors of regular theta functions of degree λ in the reference seed. It consists of the integral points in a bounded degree slice of the regular theta cone [12, Theorem 5.4]. In simply-laced type, Genz, Koshevoy and Schumann identify the polyhedral parametrizations given by the cluster potential and the decoration function with the classical string and Lusztig parametrizations [19]. For $\lambda, \mu, \nu \in P_+$, the tensor product multiplicity formula [25, Proposition 1.2] gives

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

Here $\rho_i(g)$ is the E_i-string length of ϑ_g. Each upper bound on a string length gives finitely many linear inequalities in the reference seed [12, Proposition 5.5]. With the alternative

realization G/U_G and the left action of G, the summand of right T-degree λ is $L(\lambda)^*$. Its crystal involution is $\xi_{-w_0\lambda}$, which equals ξ_λ in types B_n and C_n.

9.3. **Lagrangian Grassmannians.** Let $X = \mathrm{LGr}(n, 2n) = G/P_n$ for $G = \mathrm{Sp}_{2n}(\mathbb{C})$, $n \geqslant 2$, where P_n stabilizes a Lagrangian subspace. The Plücker ring of X is $\mathcal{R} = \bigoplus_{m\geqslant 0} \mathcal{R}_m$, with $\mathcal{R}_m = L(m\varpi_n)^*$. Let b_0 be the highest Plücker coordinate, whose zero divisor is the Schubert divisor with multiplicity one. For $1 \leqslant i \leqslant n$, let b_i be the minimal lift of the minor at the last occurrence of i in a reduced word for the open cell. On this cell the terminal n-minor is the lowest Plücker coordinate divided by b_0. If the numerator vanished along the Schubert divisor, the quotient would be regular on projective X, hence constant. The two coordinates have distinct weights, so b_n is the lowest coordinate. The lifted word seed gives $\mathcal{R}$ the structure of an upper cluster algebra with polynomial frozen variables. Its theta basis is homogeneous for both Plücker degree and torus weight [12, Theorems 3.4 and 5.4].

In the standard realization of C_n with orthonormal basis $e_1, \ldots, e_n$, the highest root is $\theta = 2e_1$. Choose root vectors $X_\theta, X_{-\theta}$ with $[X_\theta, X_{-\theta}] = H_\theta$, the coroot of θ. We use the same symbols for their derivations on $\mathcal{R}$. Write $\mathcal{E}$ for the Euler derivation, which acts on $\mathcal{R}_m$ by multiplication by m.

Lemma 9.4. *The Plücker ring has no nonzero homogeneous derivations of negative degree. For every integer $q \geqslant 0$ there is a G-module isomorphism*

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

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

Proof. Write the symplectic space as $E \oplus E^*$, with E of torus weights $e_1, \ldots, e_n$, and let $z = b_n$. Over $z \neq 0$ a point of the cone is determined by its scale z and a Lagrangian transverse to E^*. Such a Lagrangian is the graph of a symmetric map $S : E \to E^*$. The opposite unipotent radical of P_n fixes z and acts by all translations $S \mapsto S + C$. A derivation commuting with these translations has coefficients independent of S in the frame consisting of $\mathcal{E} = z\partial_z$ and the constant translation fields. If its degree is q, these coefficients are scalar multiples of z^q.

The translations form the irreducible $\mathrm{GL}(E)$-module $\mathrm{Sym}^2(E^*)$. Its lowest weight line is the root space of $-\theta = -2e_1$. Thus a derivation fixed by the negative maximal unipotent subgroup restricts to $z^q(a\mathcal{E} + bX_{-\theta})$. If $q < 0$, evaluating at b_0 gives $ab_0 + bX_{-\theta}b_0 \in z\mathcal{R} \cap \mathcal{R}_1 = \mathbb{C}z$. The root string through b_0 has length one because $\langle \varpi_n, \theta^\vee \rangle = 1$. The three nonzero vectors have distinct weights $\varpi_n, \varpi_n - \theta, -\varpi_n$ for $n \geqslant 2$, so $a = b = 0$.

For $q \geqslant 0$ both fields extend to $\mathcal{R}$ and give the stated lowest weight lines. Since $\mathcal{R}_1$ generates $\mathcal{R}$, its degree-q derivations embed in $\mathrm{Hom}(\mathcal{R}_1, \mathcal{R}_{q+1})$, with $\mathcal{R}_j = 0$ for $j < 0$. Complete reducibility gives the decomposition, and the longest Weyl group element gives the two highest weight lines. □

Proposition 9.5. *Every frozen vertex is both sink-reachable and source-reachable. The exact source partner of the sink i is $\sigma(i) = n - i$ for $0 \leqslant i \leqslant n$. The boundary derivations give a degree-preserving action of the derived affine Lie algebra of type $C_n^{(1)}$ at level zero. After a choice of finite pinning, this action is the pullback of the geometric $\mathfrak{sp}_{2n}$-action by evaluation of the loop parameter at a nonzero scalar. The homogeneous theta basis is normalized perfect for its Chevalley generators. In degree m, the crystal on colors $1, \ldots, n$ is $B(m\varpi_n)$. If $m > 0$, no element is highest for all the affine colors.*

Proof. Let $J = \{1, \ldots, n-1\}$ and let w_J be the longest element of the parabolic Weyl subgroup generated by the s_j for $j \in J$. The word

$$(n, n-1, \ldots, 1)(n, n-1, \ldots, 2) \cdots (n)$$

sends e_a to $-e_{n+1-a}$, so it represents $w = w_0 w_J$. It is reduced because its length $n(n+1)/2$ counts the positive roots sent to negative roots, namely $2e_a$ and $e_a + e_b$ for $a < b$. We use

symmetrizing weights $d_i = 1$ for letters $i < n$ and $d_n = d_0 = 2$, and orient the exchange matrix so that the positive simple-root derivations act at sinks.

For $1 \leqslant i \leqslant n$, let X_i, Y_i be the positive and negative Chevalley derivations of the geometric G-action. In a seed adapted to X_i, the only moving frozen variable is b_i [12, Proposition 4.6]. For $i < n$ one has $w^{-1}\alpha_i = \alpha_{n-i}$, so Y_i acts at b_{n-i}. The index n is a left descent of w, and Y_n acts at b_0 [12, Corollary 4.9 and Proposition 4.11]. In these seeds the active vertices are sinks for X_i and sources for Y_i. Since the mutable pattern admits a reddening sequence [12, Lemma 5.1], Lemma 1.5 shows that every frozen vertex is both sink-reachable and source-reachable.

In the adapted monomial formulas of [12, Theorem 4.1], the factor at a mutable index k is d_i/d_k. Since $d_i = d_{n-i}$, including $d_n = d_0$, the mutable exponents agree with those of R_i and $R^\star_{n-i}$, respectively. Each geometric derivation is therefore a scalar times a frozen monomial times its boundary derivation. All frozen exponents are nonnegative, since the geometric derivative of the active variable is regular in the adapted seed.

Each b_j has positive Plücker degree. A nonconstant frozen factor b of degree d would force the image of $\mathcal{R}_1$ into $b\mathcal{R}_{1-d}$, since the boundary derivation preserves $\mathcal{R}$. This space has dimension at most one, whereas every simple-root operator has rank at least two on $\mathcal{R}_1$. For the long root, the extremal weights $(\pm 1, \dots, \pm 1)$ give 2^{n-1} disjoint strings differing in the last sign. For a short root $\alpha_i = e_i - e_{i+1}$, an extremal weight with signs $(+,-)$ in positions $i, i+1$ pairs to two with $\alpha_i^\vee$ and gives a string of length two. Thus

$$X_i = c_i R_i, \qquad Y_i = c_i' R^\star_{n-i}$$

for nonzero scalars c_i, c_i'.

Since $X_i, Y_i, [X_i, Y_i]$ form a Chevalley triple, Theorem 6.5 gives exactness of $(i, n-i)$ and $c_i c_i' = 1$, including when $i = n - i$.

It remains to pair the sink 0 and the source n. Lemma 6.7, applied to each finite pair, gives

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

Both boundary derivations are homogeneous in Plücker degree, as their monomial formulas show. By Lemma 9.4 their degrees q, q' are nonnegative, and the commutators place them in the lowest and highest weight spaces, respectively. Since R_0 kills b_n and $R_n^\star$ kills b_0, their Euler terms vanish. Thus for nonzero scalars c, c' we have

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

In a sink-optimizing seed for 0, the mutable monomial $R_0 b_0$ is a unit in the Laurent ring with polynomial frozen variables. Since $X_{-\theta} b_0$ belongs to this ring and b_n is a nonunit there, $R_0 b_0 = c b_n^q X_{-\theta} b_0$ forces $q = 0$. The same argument in a source-optimizing seed for n gives $q' = 0$. The opposite root vectors $X_{-\theta}, X_\theta$ form a Chevalley triple with $-H_\theta$. Theorem 6.5 gives exactness of $(0, n)$ and $cc' = 1$.

Use the finite pinning in which R_i and $R^\star_{n-i}$ are the positive and negative simple-root operators for $1 \leqslant i \leqslant n$. In the loop realization of the derived affine algebra, the additional Chevalley generators may be taken to be $X_{-\theta} \otimes t$ and $X_\theta \otimes t^{-1}$ [22, Chapter 7]. Evaluation at $t = c$ gives exactly $R_0 = c X_{-\theta}$ and $R_n^\star = c^{-1} X_\theta$. Their commutator is $H_0 = -H_\theta$. The eigenvalues of the finite simple coroots and of $-H_\theta$ on the simple and affine root vectors give the Cartan matrix of type $C_n^{(1)}$. The central element $h_0 + \theta^\vee$ acts by zero. All generators preserve degree. Theorem 7.2 gives normalized perfectness, and Corollary 7.5 identifies the finite-color crystal of $\mathcal{R}_m = L(m\varpi_n)^*$ with $B(m\varpi_n)$, since $w_0 = -1$. The level-zero argument in Section 7.3 shows that a vertex highest in all affine colors would be isolated. For $m > 0$ this is impossible, since $B(m\varpi_n)$ is connected on the finite colors and has more than one vertex. □

Corollary 9.6. *For $m \geqslant 0$, the Lusztig involution of $B(m\varpi_n)$ is the unique involution ξ_m of the degree-m theta crystal satisfying*

$$\xi_m r_i = l_i \xi_m, \qquad \xi_m l_i = r_i \xi_m \qquad (0 \leqslant i \leqslant n).$$

Proof. Let $J_+ = \{1, \dots, n\}$ and $J_- = \{0, \dots, n-1\}$. Under the evaluation action of Proposition 9.5, the roots for J_- are obtained from the usual simple roots for J_+ by replacing $2e_n$ with $-2e_1$. They form another simple system for $\mathfrak{sp}_{2n}$. In each system the longest Weyl group element acts by -1 on the full torus weight lattice. Thus the finite-type Lusztig involutions ξ_{J_+} and ξ_{J_-} both negate the torus weight and reverse the arrows of every common color $1, \dots, n-1$.

Their composite preserves the full torus weight and commutes with the crystal operators of the common colors. The simple roots of these colors, together with the full torus, define the Levi subgroup GL_n. The restriction of $\mathcal{R}_m = L(m\varpi_n)^*$ to this subgroup is multiplicity-free [1, Theorem 5.4 and Table 2, case 1.3]. Hence each highest label for the common colors is determined by its full GL_n-weight. The composite fixes these labels and therefore every label obtained from them by lowering. It follows that $\xi_{J_+} = \xi_{J_-}$, so ξ_{J_+} reverses every affine color. Uniqueness follows because $B(m\varpi_n)$ is connected and its highest vertex must map to its unique lowest vertex. □

For the analogous involution on projected level-zero path crystals, see [32, Section 4.5]. By Corollary 8.3, the relative DT square Ψ is a direct cluster automorphism of the Plücker ring acting on colors by $i \mapsto n - i$. It exchanges the two finite color sets in the preceding proof, so on theta labels it commutes with the involution given in each degree by ξ_m.

In rank two there is one mutable vertex, so the normalized reddening map on mutable degrees is $p \mapsto -p$. The map Ψ interchanges b_0 and b_2 and fixes x and y, the primitive regular lifts of the labels 1 and -1 in the x-chart. Their Plücker exchange relation is $xy = b_1^2 + b_0 b_2$. Write $E_i = R_i$ and $F_i = R^\star_{2-i}$, taking each partial derivative in the cluster containing its monomial coefficient.

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

In the order $0, 1, 2$, the Cartan matrix is $\left(\begin{smallmatrix} 2 & -1 & 0 \\ -2 & 2 & -2 \\ 0 & -1 & 2 \end{smallmatrix}\right)$.

To test the full weight-lattice criterion after adjoining frozen variables $z_1, \dots, z_s$, write the exchange relation as $xy = P(z)b_1^2 + Q(z)b_0 b_2$ with coefficient monomials P, Q. The same two seeds still optimize the original frozen vertices, so their six boundary derivations have the formulas above. The exchange relation now gives $E_0 x = Q b_2$ and $E_1^2 y = 2Px$. Exactness of $(0, 2)$ forces $Q = 1$, and exactness of $(1, 1)$ forces $P = 1$. Thus in the grading quotient the original relation $\deg x + \deg y = 2\deg b_1 = \deg b_0 + \deg b_2$ still holds. The three raising derivations have degree shifts $\deg y - \deg b_0$, $\deg x - \deg b_1$ and $\deg y - \deg b_2$, respectively. Since $[q_0] + 2[q_1] + [q_2] = 0$, Proposition 7.8 rules out independent simple-root classes in every frozen extension preserving the exact pairs.

10. Theta F-pairings and maximal adjunction

In an optimizing seed, maximal raising keeps the terms of highest x_i-degree and multiplies their sum by a monomial. On the F-polynomial it selects a face of the Newton polytope. We first use this face to compare Cao's symmetric F-invariant before and after maximal raising. We then determine the full domain on which maximal raising admits an adjoint for the valuation pairing $\mathcal{K}$. For the corresponding categorical operators, see [11, Appendix A].

10.1. **Principal coefficients and the symmetric F-invariant.** Fix a seed t. Adjoin invertible frozen variables z_k, $k \in U$, so that the extended exchange matrix in the column

convention is $\widehat{\widetilde{B}}_t = \begin{pmatrix} \widetilde{B}_t \\ I_U \end{pmatrix}$. In the square-matrix convention this extension has exchange matrix and symmetrizer

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

where E_U includes the mutable coordinate lattice in the full coordinate lattice. Let g be a tropical point with finite principal-coefficient lift, and let g_t denote its coordinates at t. Its *principal-coefficient F-polynomial* $\mathcal{F}_{t,g}$ is determined by

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

We use Y_k, $k \in U$, as the indeterminates of this polynomial. The exponents of z distinguish its terms. Theta positivity and pointedness give nonnegative integral coefficients and constant term one [21]. For a cluster variable this is the F-polynomial of Fomin and Zelevinsky [16, Definition 3.3]. In the projective-weight convention of Section 4, $\mathcal{F}_{t,g} = F_{-(g_t)_U}$.

Lemma 10.1. *Fix a principal extension at one seed. If a finite theta function has degree (g_t, a_t) at another seed t, its canonical F-polynomial there is $\mathcal{F}_{t,g}$, where g is the projected tropical point. 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, 38]. 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, denote the new auxiliary variables by u_k and apply 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. □

The affine injection $\alpha \mapsto (g_t + \alpha\widetilde{B}_t^{\mathsf{T}}, \alpha)$ sends the Newton polytope of $\mathcal{F}_{t,g}$ onto that of the lifted theta function. The coefficient at each vertex is therefore one [7, Theorem 3.15].

Write $N_{t,g} = \operatorname{Newt}(\mathcal{F}_{t,g})$, $h_{t,g}(w) = \max_{\alpha\in N_{t,g}} \alpha w$, and $w_{t,g} = D_U(g_t)_U^{\mathsf{T}}$. We write F for the extension of the invariant F_D of Section 2 to these theta functions, suppressing the fixed symmetrizer D.

Proposition 10.2. *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 [5, Theorem 4.10]. *For cluster monomials with trivial coefficients this also recovers* [5, Proposition 4.30].

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 [5, Theorem 4.10]. Lemma 10.1 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, distinct coefficient degrees give distinct Laurent

exponents. Thus the principal-coefficient polynomial is also Cao's canonical polynomial in the original seed. □

Two theta functions are *F-compatible* when their F-invariant vanishes.

10.2. **Maximal raising and Newton faces.** Assume mutable reddening for the rest of this section. Fix a seed in which i is a sink, and use its coordinates throughout. Let m be the exponent vector of M_i, so $m_U = -B_{i,U}$ and $m_F = 0$. Write

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

For $g \in \mathcal{B}_\theta^+$, the element $P_i g = r_i^{\rho_i(g)} g$ is the highest element of its i-string. We omit t from the notation for F-polynomials, Newton polytopes and support functions and from $w_{t,g}$. Since N_g lies in the nonnegative orthant and contains zero, $N_g^0 = N_g \cap q^\perp$ is a face. Denote its support function by h_g^0.

Proposition 10.3. *In the chosen seed,*

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

The polynomial identity holds before coefficient specialization.

Proof. Take the principal extension at this seed. 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 3.8, applied in this extension, gives $R_i^{(g_i)} \widehat{\vartheta}_{(g,0)} = \widehat{\vartheta}_{(g+g_i v,0)}$. 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. □

We write $\mathsf{F}(g,h)$ for $\mathsf{F}(\vartheta_g, \vartheta_h)$. For regular labels g, h, 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 10.4. *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 functional q vanishes on both faces N_g^0 and N_h^0. This gives the first two identities. Since a face has no larger support value than the whole polytope, both A_i terms are nonnegative. Either term vanishes precisely when the indicated maximizing face meets the corresponding face perpendicular to q. Nonnegativity of the coefficient exponents and subadditivity of support functions 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 decrease 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 $t > (\alpha w - h_g^0(w))/(\alpha q)$, and there are only finitely many exponents.

For an isolated zero exchange column at k, the 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.

For the symmetric F-invariant, a map Q on regular labels satisfying $\mathsf{F}(P_i g, h) = \mathsf{F}(g, Qh)$ for all regular g, h exists if and only if $M_i = 1$. Indeed, if $m_k > 0$, take $g = e_i$ and let $\vartheta_h = x'_k$ be the adjacent cluster variable. Its F-polynomial is $1 + Y_k$, so the left side is $d_k m_k > 0$, while every pairing with x_i is zero. If $M_i = 1$, maximal raising removes an independent frozen factor and $Q = \mathrm{Id}$ works.

10.3. **Adjunction for the valuation pairing.** For a real direction $n \in N^\bullet_{\mathbb{R}} = N^\bullet \otimes_{\mathbb{Z}} \mathbb{R}$, let $\mathcal{K}(g, n) = \mathrm{val}_n(\vartheta_g)$ be the minimum of n on the Laurent support of ϑ_g. In the chosen seed, $P_i g = g + g_i v$. The transpose of this linear map is

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

It fixes every coordinate except the ith, which becomes $\langle m, n\rangle$. Both P_i and Q_i are idempotent. Moreover, Q_i preserves the chiral-dual lattice $N^\bullet$. Indeed, if $n_a = (d_a/\ell)z_a$, the new integral coordinate at i is $\sum_{k\in U} b_{ki} z_k$, since $d_k m_k = d_i b_{ki}$.

Denote theta functions in the seed with J frozen by a superscript 0. Qin's freezing theorem selects the terms of coefficient degree $\alpha_J = 0$ in the principal extension [41, Theorem 5.2]. After specialization, Proposition 10.3 gives $\vartheta_{P_i g} = x^{g_i v}\vartheta_g^0$. Every exponent a of ϑ_g^0 has $a_i = g_i$, and hence $\langle a, Q_i n\rangle = \langle a + g_i v, n\rangle$. Taking minima gives the first equality below for every real n. For $n \in N^\bullet$, freezing preserves injective reachability by [41, Theorem 5.7] (see [37, Theorem 3.1.3] for quivers), so the dual theta function is finite by Lemma 3.1. Theta reciprocity in the frozen seed gives the second equality.

$$\mathcal{K}(P_i g, n) = \mathrm{val}_{Q_i n}(\vartheta_g^0) = \mathrm{val}_g(\vartheta^{\vee,0}_{Q_i n}). \tag{10.1}$$

Let $\mathcal{S}_i$ be the support of ϑ_{-m} in the chosen seed and define

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

The pointing exponent $-m$ belongs to $\mathcal{S}_i$, and Q_i fixes the pairing with m. Thus

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

Lemma 10.5. *For $n \in N^\bullet$, the equality $W_i(n) = 0$ holds if and only if $\vartheta^{\vee,0}_{Q_i n} = \vartheta^\vee_{Q_i n}$.*

Proof. Write $q = Q_i n$ and let $A_q \subseteq \mathbb{N}^U$ be the set of coefficient degrees occurring in the dual theta expansion. Writing Z^η for the dual torus monomial of exponent $\eta \in N^\bullet$, we have

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

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 $\alpha_J = 0$ for every contributing degree. By (3.1), the sum is given by the F-polynomial in the opposite principal-coefficient pattern. Freezing J there sets $Y_k = 0$ for $k \in J$ by Qin's freezing theorem [41, Theorem 5.2]. The remaining mutable degree is the restriction of the original one, so (3.1) in the frozen seed gives $\vartheta_q^{\vee,0}$ from exactly these terms. The dual monomials for distinct degrees are distinct. □

Theorem 10.6. *For a real direction $n \in N^\bullet_{\mathbb{R}}$, there exists a real direction q such that*

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

if and only if n belongs to the rational polyhedral cone

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

The adjoint direction is unique and equals $Q_i n$. For $n \in N^\bullet$ it also belongs to $N^\bullet$. The map Q_i retracts Ω_i onto $\Omega_i \cap v^\perp$.

Let I_k be the variable of the injective seed $t[1]$ with mutable degree $-e_k$ in the chosen seed t. It suffices to test the identity on the variables of t and on the squarefree cluster monomial $\prod_{k\in J} I_k$, with the empty product understood as 1. The cone Ω_i is the whole valuation space exactly when $M_i = 1$.

Proof. Testing the seed variables determines q. Maximal raising fixes x_a for $a \neq i$ and sends x_i to x^m, so $q = Q_i n$. For $k \in J$, write $a_k = -e_k + c_k$ for the degree of I_k at t, where c_k is supported on F. Its pointing term gives $\delta_k(q) = \langle a_k, q\rangle - \mathrm{val}_q(I_k) \geqslant 0$. Frozen translation gives $\vartheta_{-m} = \prod_{k\in J}(x^{-c_k} I_k)^{m_k}$, hence $W_i(n) = \sum_{k\in J} m_k \delta_k(q)$. Thus $n \in \Omega_i$ precisely when q attains the minimum on each $\mathrm{Newt}(I_k)$ at a_k.

The squarefree product has degree $a = \sum_{k\in J} a_k$. After freezing J this degree is supported on frozen variables, so $\vartheta^0_a = x^a$. By (10.1) the difference $\mathcal{K}(P_i a, n) - \mathcal{K}(a, q)$ is $\sum_{k\in J} \delta_k(q)$. Its vanishing is therefore equivalent to $n \in \Omega_i$.

For $n \in \Omega_i \cap N^\bullet$, Lemma 10.5 gives $\vartheta^{\vee,0}_{Q_i n} = \vartheta^\vee_{Q_i n}$. Equation (10.1) and reciprocity give the identity for every regular g. Positive homogeneity extends it to the rational points of Ω_i. The cone is rational polyhedral by the formula for W_i, and both sides are continuous minima of finitely many linear forms, so it holds throughout Ω_i. Since $Q_i^2 = Q_i$ and $\langle m, Q_i n\rangle = \langle m, n\rangle$, we have $W_i(Q_i n) = W_i(n)$. The image of Q_i is $v^\perp$, giving the retraction.

If $M_i = 1$, then $m = 0$ and W_i vanishes identically. Otherwise choose $k \in J$ and take $n = e^*_k$. The adjacent variable x'_k has two Laurent terms whose exponents differ from its pointing exponent by 0 and $p_k = -B^{\mathsf{T}}_{k,V}$, respectively. Freezing k keeps the pointing term. Since $Q_i n = e^*_k + m_k e^*_i$ and $\langle p_k, Q_i n\rangle = -m_k b_{ki} < 0$, the other term has strictly smaller valuation. Thus adjunction fails on x'_k in this direction.

□

For example, let $xx' = 1 + f$, with x mutable and f a frozen sink. Then $R_f = x\partial_f$, and maximal raising sends f to x and x' to 1. In the coordinates (x, f), the support of $\vartheta_{-e_x} = x'/f = x^{-1}(1 + f^{-1})$ gives

$$Q_f(u, w) = (u, u), \qquad \Omega_f = \{(u, w) : u \leqslant 0\}.$$

For any real n and any $g \in \mathcal{B}^+_\theta$, one has

$$\Delta_i(g; n) = \mathcal{K}(P_i g, n) - \mathcal{K}(g, Q_i n) \geqslant 0.$$

Indeed, (10.1) compares the minimum of $Q_i n$ on the face of $\mathrm{Newt}(\vartheta_g)$ of highest x_i-degree with its minimum on the whole polytope. Equality holds exactly when the minimizing face meets that highest-degree face. For integral n, reciprocity expresses the two minima as valuations of the frozen and original dual theta functions at g, extending Δ_i to a homogeneous piecewise-linear function of g.

Given a word $P_{i_r} \cdots P_{i_1}$ and a direction $n_r = n$, define $n_{k-1} = Q_{i_k} n_k$ in the corresponding optimizing seed, successively for $k = r, \ldots, 1$. If $n_k \in \Omega_{i_k}$ at each step, iteration gives

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

The adjunction characterizes its domain and the restriction of Q_i to that domain independently of the optimizing seed. Indeed, valuations on the seed variables distinguish dual tropical points.

Suppose the original exponent lattice admits a rational skew form Γ with $\widetilde{B}^{\mathsf{T}}\Gamma = (D_U \mid 0_F)$. In our convention, Cao's *tropical invariant* [5, Definition 4.6 and Theorem 4.10] 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 10.6 therefore gives, for an integral target h', the criterion

$$\begin{aligned}\mathsf{A}_\Gamma(P_i g,h) = \mathsf{A}_\Gamma(g,h') \quad &\text{for every regular } g\\ \iff \quad -\Gamma h^{\mathsf{T}} \in \Omega_i \quad &\text{and} \quad \Gamma(h')^{\mathsf{T}} = Q_i \Gamma h^{\mathsf{T}}.\end{aligned}$$

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.

11. Maximal operators and string coordinates

11.1. **String cones and highest theta functions.** Assume mutable reddening. For a word $\mathbf{i} = (i_1,\dots,i_r)$ in sink-reachable colors, start at $g^{(0)} = g \in \mathcal{B}_\theta^+$ and set

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

The *string datum* is $\mathbf{s_i}(g) = (n_1,\dots,n_r)$. From $\Psi_{\mathbf{i}}(g) = (\mathbf{s_i}(g), g^{(r)})$ we recover g by applying $l_{i_r}^{n_r},\dots,l_{i_1}^{n_1}$ to the final label in that order.

Theorem 11.1. *In each reference seed, $\Psi_{\mathbf{i}}$ extends to a homogeneous piecewise integral-linear embedding of $\mathbb{R}^V$ with a piecewise integral-linear inverse on its image. The image of $\mathcal{C}^+$ is a convex rational polyhedral cone and*

$$\Psi_{\mathbf{i}}(\mathcal{B}_\theta^+) = \Psi_{\mathbf{i}}(\mathcal{C}^+) \cap (\mathbb{Z}^r \times \mathbb{Z}^V).$$

These pairs form a finitely generated saturated monoid. The string data alone form a finitely generated monoid.

Proof. Let Φ_i be tropical transport to a sink-optimizing seed. For $a = \Phi_i(g)$ a single step gives $(n,p) = (a_i, \Phi_i^{-1}(a + a_i v_i))$ and has inverse $(n,p) \mapsto \Phi_i^{-1}(\Phi_i(p) - nv_i)$. Composing the steps gives the asserted embedding and inverse, each with finitely many rational cones of linearity. In particular the inverse takes lattice points to lattice points.

For indeterminates $z_1,\dots,z_r$, local nilpotence gives a polynomial-valued algebra homomorphism $\mathcal{E}_{\mathbf{i}}(f) = \exp(z_r R_{i_r}) \cdots \exp(z_1 R_{i_1}) f$. In lexicographic order, comparing the first coordinate first, its leading term at ϑ_g is $z^{\mathbf{s_i}(g)}\vartheta_{g^{(r)}}$ by Corollary 3.8. For equal strings, distinct labels have distinct final labels, so the leading coefficients at the largest string in a theta expansion cannot cancel. The negative of the leading exponent defines a valuation on Mid^+. For the corresponding flag variety valuation see [28, Theorem 4.1].

Write $\Psi_{\mathbf{i}}(g) = (s,p)$ and $\Psi_{\mathbf{i}}(h) = (t,q)$. The leading term for $\vartheta_g\vartheta_h$ is $z^{s+t}\vartheta_p\vartheta_q$. In principal coefficients the product at p,q contains the lift at $p+q$ with coefficient one, and positivity preserves this contribution on specialization [7, Corollary 2.5 and Example 2.6]. Comparing leading coefficients in the theta expansion therefore gives a label k with $\Psi_{\mathbf{i}}(k) = (s+t, p+q)$. Thus the lattice points of the image are closed under addition.

By Corollary 3.2, $\mathcal{C}^+$ is rational polyhedral, so its image is a finite union of rational polyhedral cones. Clearing denominators gives closure under addition for the rational points of this union, and density and closedness give it for all real points. The image is therefore the cone generated by the finitely many rational generators of its pieces. Gordan's lemma and projection to the string coordinates give the monoid assertions. □

For a reduced word for the longest element in finite Cartan type, exact partners give a normal crystal. The final label is the highest label of its component by Corollary 7.7 and [33, Section 1]. Thus the fiber at a highest label p has as its lattice points the strings of its

component. It is a polytope, since it contains $(0,p)$ and its finitely many lattice points rule out a nonzero rational recession direction. For the classical string cone see [4, Proposition 3.5].

Corollary 11.2. *Assume mutable reddening and let S be a set of sink-reachable colors. The highest labels of $A_S^+ = \bigcap_{i\in S} \ker R_i$ in Mid^+ are the lattice points of a rational polyhedral cone H_S in each seed. Write $\mathbf{1} = (1)_{u\in U}$ and $\delta(g) = \mathsf{E}^{\mathrm{op}}(g_U, \mathbf{1})$, the degree of $\mathcal{F}_g$ for $\deg Y_u = d_u$. The cone*

$$\widehat{H}_S = \{(g,m) : g \in H_S,\ m \geqslant \delta(g)\}$$

is pointed and rational polyhedral. The functions ϑ_g for (g,m) in its Hilbert basis*, the minimal additive generating set of its lattice points, generate A_S^+.*

Now choose exact pairs $(i,\sigma(i))$, $i \in I$, as in Theorem 8.2, and suppose $S \subseteq I$ has finite Cartan type. For a lift Φ with label action κ, the equality $\Phi = n_S$ on A_S^+ holds if and only if $\kappa(g) = w_{0,S}g$ for each primitive generator (g,m) of a nonvertical extreme ray.

Proof. For a word listing S once, the face on which all string coordinates vanish has the form $\{0\} \times H_S$ by Theorem 11.1. Its lattice points label the common kernel by perfectness. Theorem 4.3 identifies the weighted degree $\deg_D \mathcal{F}_g$ with $\mathsf{E}^{\mathrm{op}}(g_U, \mathbf{1})$. This is the support function of the Newton polytope of the opposite F-polynomial of mutable degree $-\mathbf{1}$, evaluated at $-g_U D_U$. Thus δ is nonnegative, subadditive and positively homogeneous, and its epigraph over H_S is rational polyhedral. If this cone contained a line, it would contain opposite nonzero lattice points $\pm(g,m)$ with $m = \delta(g) = \delta(-g) = 0$. Both principal lifts would be monomials $x^{\pm g}$. Laurentness after mutation at u forces $\pm g_u \geqslant 0$, since a negative exponent leaves a nonunit exchange binomial in the denominator. The two monomials are then frozen, and their regularity forces $g = 0$. This proves pointedness.

Let $g = \sum_j g_j$ with all g_j highest and $\delta(g) = \sum_j \delta(g_j)$. Specialize the auxiliary variable at each u to t^{d_u}. The principal lift at h then has t-degree $\delta(h)$. The product has a finite theta expansion by Proposition 1.3. Its unshifted term is the lift at g, with coefficient one. Every other summand is a positive multiple of t^k times a lift at h, with $k > 0$, because its auxiliary exponent is nonzero and nonnegative [7, Corollary 2.5 and Example 2.6]. Positivity gives $k + \delta(h) \leqslant \sum_j \delta(g_j) = \delta(g)$. At $t = 1$ the product lies in A_S^+, whose highest theta basis is part of the localized theta basis. By positivity, every label that occurs is therefore highest, and

$$\prod_j \vartheta_{g_j} = \vartheta_g + \sum_{\delta(h)<\delta(g)} c_h \vartheta_h, \qquad c_h \in \mathbb{Z}_{\geqslant 0}.$$

In any decomposition $(g,\delta(g)) = \sum_j (g_j, m_j)$ into lattice points of the epigraph, subadditivity gives $\delta(g) \leqslant \sum_j \delta(g_j) \leqslant \sum_j m_j = \delta(g)$. Taking a Hilbert basis decomposition and using the product formula, induction on $\delta(g)$ proves the generation assertion. In degree zero the sum of lower terms is empty. For a related finite-generation theorem see [21, Theorem 8.19].

Let $\tau = w_{0,S}^{-1}\kappa$. This injective label map commutes with positive integer scaling by Proposition 5.4 and the string reflections of Section 8.1. On highest labels h we have $n_S\vartheta_h = \vartheta_{w_{0,S}h}$ by the proof of Theorem 8.2. If the ray tests hold, the two algebra maps agree on every product of the corresponding theta functions. Thus τ permutes the highest labels in its support. For each nonzero highest g, some positive integral multiple $a(g,\delta(g))$ is a sum of primitive ray generators. The preceding degree inequality forces equality throughout and excludes vertical summands. In this product ag is the unique label of largest degree. By homogeneity $\tau(g)$ is highest, and if $\tau(g) \neq g$ then $\delta(\tau(g)) < \delta(g)$. A moved highest label of least degree would therefore map to a fixed label, contradicting injectivity. Necessity follows by evaluation on the ray labels. □

The ray functions need not generate the highest algebra. For example, take the B_2 seed $(x_1, x_2; f, h)$ with weights $(2,1,1,1)$ and $\widetilde{B}^{\mathsf{T}} = \left(\begin{smallmatrix} 0 & -2 & 2 & 0 \\ 1 & 0 & -1 & 1 \end{smallmatrix}\right)$. Number the cluster variables cyclically. In the (x_3,x_4)-chart, $R_f = x_4\partial_f$, $x_2x_4 = x_3 + fh$ and $x_3x_5 = x_4^2 + h^2$. Since $M_h^\star = x_2$, the identity $R_f x_2 = h$ makes (f,h) a covering exact pair. The six mutable variables have order zero at both frozen divisors and f-degrees $2,1,0,0,0,1$ in this chart. Their cluster

cones cover the mutable plane, so the highest monomials for $S = \{f\}$ are powers of h times monomials in (x_3, x_4) or (x_4, x_5).

In the initial seed the triples $(g_U, \delta(g))$ for x_3, x_4, x_5 are $(-1, 2, 2), (-1, 1, 3), (-1, 0, 4)$, with zero frozen coordinates. The middle triple is the midpoint of the others and degrees add in each cluster. The nonvertical rays therefore give x_3, x_5, h. Changing the sign of x_4 in the (x_3, x_4)-chart fixes all three, so their algebra does not contain x_4.

11.2. **Cluster monomials.** For cluster monomials we can follow an entire cluster through a maximal raising step. Here an *ordinary cluster monomial* is a monomial with nonnegative exponents in the mutable and frozen variables of one seed. In the generic basis setting, [11, Corollary A.3] treats generalized cluster monomials.

Theorem 11.3. *Let R_i be a canonical sink derivation and choose a sink-optimizing seed for i. For every seed t there is a seed s such that $R_i^{(d)} z$ is an ordinary monomial of s for every ordinary cluster monomial z of t, where $d = \deg_{R_i} z$. One can reach s from the chosen optimizing seed without mutating any vertex adjacent to i in that seed. On the monomials of t, the map $z \mapsto R_i^{(\deg_{R_i} z)} z$ is multiplicative and acts on exponents by a matrix of nonnegative integers. The same conclusions hold for source derivations. Both statements hold without reddening.*

Proof. In the chosen sink-optimizing seed write $M_i = x^{m_i}$ and let $J = \{u \in U : b_{iu} < 0\}$. Let q_z be the sum of the Laurent terms of z having maximal x_i-exponent. Then $R_i^{(d)} z = x^{d(m_i - e_i)} q_z$. Lift the monomial z, in its own seed, to the principal extension based at the optimizing seed. Terms of coefficient degree ν have x_i-exponent equal to the pointing exponent minus $\sum_{u \in J} \nu_u b_{ui}$. Hence q_z is the specialization of the part with $\nu_J = 0$. Qin's freezing theorem identifies this part with a localized cluster monomial of the pattern with J frozen [41, Theorem 5.11]. The chamber of t lies in a reachable chamber of that pattern [41, Lemmas 5.3 and 5.6]. All monomials of t have their degrees in this same chamber, so its seed s works for all of them. Specializing the auxiliary variables to one gives the same assertion in the original coefficients. The factor $x^{d(m_i - e_i)}$ uses only variables frozen in this pattern. Consequently $R_i^{(d)} z = x(s)^c$, with $c_u \geqslant 0$ for every $u \in U \backslash J$.

To prove $c_u \geqslant 0$ for $u \in J$, note that $b_{iv} = 0$ for every $v \in U \backslash J$ in the optimizing seed. The matrix mutation rule shows inductively that mutations at these vertices leave $B_{i,U}$ unchanged. In particular $b_{iu}(s) < 0$. 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 x_i-degrees of products add and their highest-degree parts multiply, so maximal raising is multiplicative. The images of the variables of t supply the nonnegative exponent vectors of this monomial map. Applying the argument to the opposite exchange matrix proves the assertion for a source. □

For a fixed word and initial seed t, induction gives a single seed at each step containing the images of all monomials of t. The exponent matrices multiply. Since degrees add, each string coordinate is an integral linear function of the exponents in t.

Remark 11.4. Let $A \subseteq \mathcal{U}^+$ be a subalgebra stable under the derivations of $\mathbf{i}$, with a basis of regular theta functions having finite principal-coefficient lifts. Theorem 7.1 then gives the maximal divided-power formula used above, even without reddening. For a $\Bbbk$-algebra homomorphism $\epsilon : A \to \Bbbk$, let $\Phi_{\mathbf{i}}^*(f)$ be the polynomial obtained by applying ϵ to the coefficients of $\mathcal{E}_{\mathbf{i}}(f)$. 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 be the matrix with vth row $\mathbf{d}_v$. Suppose $-B_{u,V}(t)\mathsf{M} \leqslant_{\text{lex}} 0$ 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, the assumption on ϵ preserves the leading term of $\mathcal{E}_{\mathbf{i}}(\vartheta_g)$. The images of the seed variables are nonzero, so we may substitute them in the Laurent expansion at t. Each exchange monomial then has nonpositive lexicographic exponent. The F-polynomial has constant term one and nonnegative integer coefficients. Since the field has characteristic zero and the seed variables map to monomials with coefficient one, terms of equal exponent cannot cancel. The leading exponent is therefore $g(t)\mathsf{M}$. If M has full row rank, any rational right inverse C recovers $g(t) = \mathbf{s}_{\mathbf{i}}(g)C$. For the unipotent factorization charts, see [11, Appendix A.3].

In an i-sink-optimizing seed, $c_{ij} = -g(M_j)_i$ for distinct sink colors i, j, and $g(M_k^\star)$ determines whether k is the exact partner of i (Section 2.3 and Proposition 4.2). For an exact pair (i, k), transport e_i and e_k from their respective optimizing seeds to the reference seed by the max-plus tropical exchange relation. By positivity, the resulting vectors give the degrees of the seed variables under R_i and $R_k^\star$, respectively. Subtracting the sink degrees from the source degrees gives their coroot weights (Lemma 7.4).

12. Pointed biperfect bases and transitions

Assume mutable reddening. For an injective-reachable seed of full column rank, Qin parametrizes the compatibly pointed bases of the localized upper cluster algebra over $\mathbb{Z}$ by changes of basis triangular for the dominance order [40, Theorem 5.1.3]. Here we work with the regular theta basis and allow arbitrary extended rank. We determine which lower terms survive coefficient specialization and show that they are independent of the full-rank frozen extension. Perfectness forces their coefficients to agree along crystal strings. We describe the remaining parameters using the coefficient graph of [11, Section 9 and Appendix B], then determine which of them preserve linearity over the polynomial ring in the frozen variables. For quivers with potential, we also determine when supported generic cluster characters span the regular middle algebra.

Write $A = \mathrm{Mid}^+$ and $\mathcal{B} = \mathcal{B}_\theta^+$. Fix the weight gradings and the boundary derivations D_a, $a \in \mathcal{I}$, for which the theta basis is normalized perfect. Write r_a for the raising map, l_a for its partial inverse and ρ_a for its string length.

12.1. **Dominance after coefficient specialization.** Let $L = \mathbb{Z}^V$ be the exponent lattice of a reference seed. Choose a coefficient extension of full column rank obtained by adjoining invertible frozen variables z and keeping the original exchange data. Write K for the lattice of auxiliary exponents, so that $\widehat{L} = L \oplus K$. The principal extension has $K = \mathbb{Z}^U$. If the original extended exchange matrix has full column rank, the identity extension $K = 0$ is also available. Let sp denote specialization of the auxiliary variables to one. Frozen translation and the separation formula give

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

The full-rank extension has a theta basis of its localized upper algebra by Proposition 1.3. Its distinct specializations form the theta basis of the localized middle algebra.

For tropical points $\widehat{g}, \widehat{h}$ in the extension, write $\widehat{h} \prec_t \widehat{g}$ if their coordinates at t satisfy

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

This is the strict dominance order at t. Write $\widehat{h} \prec \widehat{g}$ if this inequality holds at every seed. Let $\widehat{D}_g$ consist of the tropical points $\widehat{h} \prec (g, 0)$, using reference coordinates on the right. These points are lower than $(g, 0)$ at a seed t and at its predecessor $t[-1]$. Every label in this full-rank extension indexes a theta function compatibly pointed in all seeds. It therefore has a support dimension in the sense of [40, Proposition 3.4.8], so the common lower set is finite by [40, Proposition 3.4.6(2)].

Definition 12.1. For $g, h \in L$, we write $h \lhd g$ if $(h, c) \prec (g, 0)$ for some $c \in K$. We call this relation the *projected dominance order* and write $E_g = \{h \in \mathcal{B} : h \lhd g\}$. A family $(b_g)_{g \in \mathcal{B}}$

admits compatible pointed lifts if each b_g is the specialization of an element of the localized upper algebra of the extension that is pointed at the tropical representative of $(g,0)$ in every seed.

Lemma 12.2. *The relation $\lhd$ is a strict partial order with finite lower sets. It is independent of the reference seed and of the chosen full-rank frozen extension. In a fixed seed, the points of $\widehat{D}_g$ over each h are indexed by the same set of exchange exponents in every such extension.*

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.

To compare extensions, fix $h \in L$ and a vector $n \in \mathbb{N}^U \backslash \{0\}$ such that $h - g = -nB_{U,V}$. In any extension with vertex set $\widehat{V}$, the equation $\widehat{h} - \widehat{g} = -n\widehat{B}_{U,\widehat{V}}$ fixes the auxiliary coordinates of $\widehat{h}$ once $\widehat{g} = (g,0)$ is fixed. Conversely, full column rank recovers n from the two lifted labels.

Under mutation at k, the same difference in the new seed has signed exchange exponents

$$n'_j = n_j \quad (j \neq k), \qquad n'_k = -n_k + \sum_{j \in U} n_j[-b_{jk}]_+ + [g_k]_+ - [h_k]_+.$$

Indeed, write the tropical mutation of a as $L_k(a) - [a_k]_+ \widehat{B}_{k,\widehat{V}}$, where L_k is its linear branch for $a_k \leqslant 0$. For $j \neq k$, matrix mutation gives $L_k(\widehat{B}_{j,\widehat{V}}) = \widehat{B}'_{j,\widehat{V}} + [-b_{jk}]_+ \widehat{B}'_{k,\widehat{V}}$, and $L_k(\widehat{B}_{k,\widehat{V}}) = -\widehat{B}'_{k,\widehat{V}}$. Subtracting the two tropical mutations gives the formula. It involves only original mutable data, even when iterated. The lifted point belongs to $\widehat{D}_g$ precisely when these signed exponents are nonnegative in every seed. They cannot all vanish, since tropical mutation preserves distinctness of the lifts. Thus the same initial vectors n index the fiber over h in every extension. □

Theorem 12.3. *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. The transition matrix and its inverse are unitriangular for $\lhd$, and each basis element has a finite expansion in the other basis. The class of these bases is independent of the full-rank frozen extension.

Proof. Expand the lift in the localized theta basis of the extension. In any seed the theta terms have distinct leading degrees. In a finite theta expansion of a pointed element, the theta function of its leading degree has coefficient one, and all other degrees are strictly smaller [40, Lemma 3.4.9]. Applying this at each seed 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 12.2. 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 $c(g,h)$ such that $(h, c(g,h)) \prec (g,0)$. 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 linear map $\vartheta_g \mapsto b_g$ as $1 + N$. The labels occurring in $N^k\vartheta_g$ lie at the ends of strictly descending chains of length k from g. Such a chain has length at most $|E_g|$. Hence N is locally nilpotent, and the geometric series $\sum_{k\geqslant 0}(-N)^k$ gives the inverse on each basis vector. □

Let $P_g \subset L$ be the projection of $\widehat{D}_g$ onto its original coordinates.

Corollary 12.4. *For every $f \in \vartheta_g + \bigoplus_{h\in P_g} \Bbbk\vartheta_h$, the lifts of f compatibly pointed at $(g, 0)$ form an affine space of dimension $|\widehat{D}_g| - |P_g|$. A finite set of original seeds, depending only on g, suffices to test the prescribed pointing of an element of the localized upper algebra in every such extension.*

If the original extended matrix has full column rank, every element of its localized upper algebra compatibly pointed at g has a unique lift compatibly pointed at $(g, 0)$ in each such extension.

Proof. Specialization adds the coefficients over each fiber of $\widehat{D}_g \to P_g$. Thus prescribing f imposes one independent, solvable linear equation per fiber. By Lemma 12.2 both cardinalities in the dimension formula are independent of the extension.

Start with t and $t[-1]$. For each point in their finite common strict lower set outside $\widehat{D}_g$, choose one further seed where the dominance inequality fails. The exchange exponents in Lemma 12.2 show that the same finite set of seeds works for every extension. If an element is pointed at the prescribed degree in all these seeds, its theta expansion has leading term $\widehat{\vartheta}_{(g,0)}$ and every remaining label belongs to $\widehat{D}_g$, by [40, Lemma 3.4.9]. Such an expansion is pointed in every seed.

When the original extended matrix has full column rank, take the identity extension. A compatibly pointed element has coefficient one at ϑ_g and all its other labels are strictly lower in every seed [40, Lemma 3.4.9]. Every fiber for the identity extension is a singleton. Lemma 12.2 gives the same fiber sizes for any other extension, so the dimension formula gives a unique lift. □

If ϑ_g is a localized cluster monomial in A, then $\widehat{D}_g$ is empty. Lift the monomial in a seed containing it and multiply by an auxiliary Laurent monomial to give it degree $(g, 0)$ in the reference seed. In the seed containing the monomial, compatibility with its predecessor forces support dimension zero, so the lift is the unique compatibly pointed element of that degree [40, Lemma 3.4.12]. Any $(h, c) \in \widehat{D}_g$ would give a second one by adding $\widehat{\vartheta}_{(h,c)}$. Thus every family with compatible pointed lifts has $b_g = \vartheta_g$ at these labels. In particular, $b_0 = \vartheta_0 = 1$.

12.2. **Perfectness and transition parameters.** For $n \geqslant 0$ let $K_{a,n} = \ker(D_a^{n+1}|_A)$, and set $K_{a,n} = 0$ for $n < 0$. Perfectness identifies $K_{a,n}$ with the span of the ϑ_g of length $\rho_a(g) \leqslant n$. The derivation D_a induces a map $K_{a,n}/K_{a,n-1} \to K_{a,n-1}/K_{a,n-2}$.

For $g \neq h$, call (g, h) *admissible* if $h \in E_g$, the prescribed weights agree, and $\rho_a(h) \leqslant \rho_a(g)$ for every $a \in \mathcal{I}$. By convention $T_{g,g} = 1$ for $g \in \mathcal{B}$, and $T_{0_{\mathrm{cr}},h} = T_{g,0_{\mathrm{cr}}} = 0$. The label $0 \in \mathcal{B}$ denotes the unit, and 0_{cr} denotes an undefined crystal operator value.

Proposition 12.5. *A basis (b_g) as in Theorem 12.3 preserves the prescribed weights and is normalized perfect with the prescribed crystal maps if and only if its off-diagonal coefficients vanish on inadmissible pairs and satisfy*

$$T_{r_a g,h} = T_{g,l_a h} \qquad \text{if } \rho_a(g) > 0, \quad \rho_a(h) = \rho_a(g) - 1. \tag{12.1}$$

Proof. Perfectness with the prescribed lengths requires $b_g \in K_{a,\rho_a(g)}$, so its coefficients vanish whenever $\rho_a(h) > \rho_a(g)$. Homogeneity forces vanishing when the weights differ. Conversely, the admissibility conditions preserve the weights and all the filtrations under the change of basis. The finite inverse series shows that the inverse preserves them as well. For $\rho_a(g) = n > 0$, the induced map of D_a sends the class of ϑ_g to $n[\vartheta_{r_a g}]$. Commutation with the change of basis on these quotients is therefore equivalent to (12.1), after canceling n. This is the criterion

for normalized perfectness with the prescribed crystal operators [11, Proposition 9.1 and Appendix B.1]. □

Following [11, Theorem 9.12], form the *coefficient graph* on ordered pairs $(g,h) \in \mathcal{B}^2$ with $g \neq h$. Its edges are

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

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 set of connected components containing no marked vertex.

Corollary 12.6. *The bases with compatible pointed lifts that preserve the chosen gradings and are normalized perfect for all the prescribed crystal operators 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 diagonal instances of (12.1) read $1 = 1$. For the off-diagonal entries, when $l_a h$ is defined the equation equates the entries at the ends of an edge. Otherwise it sets the entry at $(r_a g, h)$ equal to zero, as required by the second mark. Thus a coefficient is constant on its component and vanishes there if the component contains a mark. Each E_g is finite, so every assignment of the remaining coefficients gives a basis. □

For a finite set of indices G, only the unmarked components meeting a pair (g,h) with $g \in G$ contribute to $(b_g)_{g\in G}$. A group acting on the theta basis and preserving the order, the chosen gradings, and the crystal data up to permutation of the colors preserves edges and marks. For a $\Bbbk$-linear action, the equivariant changes of basis have one parameter per orbit of unmarked components, as in [11, Corollary 9.14]. An admissible pair with $\rho_a(h) < \rho_a(g)$ for every a is isolated and gives a free parameter.

Linearity over the frozen variables. Let Ψ be the map $\vartheta_g \mapsto b_g$. By frozen translation, Ψ is linear over $\Bbbk[x_F] = \Bbbk[x_j : j \in F]$ if and only if $b_{g+e_j} = x_j b_g$ for $g \in \mathcal{B}, j \in F$. This is the polynomial version of the coefficient-factorization condition in [40, Remark 5.1.4].

Corollary 12.7. *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 there is some* $j \in F$ *for which* $g - e_j \in \mathcal{B}$ *and* $h - e_j \notin \mathcal{B}$. *Among the transitions of Corollary 12.6, those that are* $\Bbbk[x_F]$*-linear have one parameter per unmarked 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. □

In particular, a coefficient-linear change of basis is determined by the images of the primitive regular theta functions of Proposition 5.1. Each regular theta function is a frozen monomial times exactly one of them, and Ψ commutes with that multiplication.

12.3. **Supported presentation weights and regular theta functions.** We ask which primitive regular theta functions occur as leading terms of supported generic cluster characters in the projected dominance order of Definition 12.1.

Let $(\Delta, \mathcal{S})$ be a skew-symmetric reference ice quiver with potential as in [11, Construction 2.5]. Assume that its full Jacobian algebra is finite dimensional and that its mutable pattern admits a reddening sequence. Work over $\mathbb{C}$ and take $D = I$. A projective presentation weight δ is *μ-supported* if the cokernel of a general presentation of weight δ is supported on the mutable vertices. Write $\mathcal{T} = -\operatorname{trop}(\Delta, \mathcal{S})$ for the negatives of these weights, $C_{\mathrm{gen}}(\delta)$ for the generic cluster character, and V_{gen} for the span of these characters. Let $\check{\delta}$ be the injective presentation weight of the same general decorated representation.

Lemma 12.8. *If $g = -\delta \in \mathcal{T}$, then ϑ_g is regular, $\eta_\theta(g) = -\check{\delta}$, and there is a finite expansion*

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

Proof. Adjoin an auxiliary frozen source $z_u \to u$ for each mutable u, as in [11, proof of Proposition 8.1]. The extended exchange matrix now has full rank. The radical of the projective at z_u is the old projective at u, so the enlarged Jacobian algebra is still finite dimensional. A presentation of weight $(\delta, 0)$ has the old cokernel extended by zero and hence the same quotient Grassmannians. Its generic character is regular, is pointed at $(g, 0)$ [11, Lemma 7.6], and is compatibly pointed under mutation [40, proof of Theorem 5.2.2].

Qin's comparison of pointed bases in full rank [40, Theorems 5.1.3 and 5.2.2] gives a unitriangular theta expansion of this character in the localization. Proposition 1.3 and regularity imply that each theta term is regular. Specializing all z_u to 1 preserves regularity. The leading term becomes ϑ_g, and Lemma 12.2 shows that all remaining terms have labels strictly below g. Finally, the support-dimension formula [40, Propositions 3.4.8 and 5.2.4] identifies the greatest F-exponent with the mutable dimension vector of the general cokernel. The codegree formula of Proposition 1.6 and the relation between projective and injective weights [11, (1.6)] give $\eta_\theta(g) = -\check{\delta}$. □

Write J_μ for the Jacobian algebra of the mutable quiver with potential. For each $j \in F$, let $\mathcal{E}_j^\mu$ be the mutable decorated representation extracted at j in [11, Section 6.2]. Given a mutable presentation weight b, let $q_j(b) = \mathrm{e}_{J_\mu}(\mathcal{E}_j^\mu, b)$ be the generic dimension of the Derksen–Fei E-space for $\mathcal{E}_j^\mu$ and a general presentation of weight b. Write $q(b) = (q_j(b))_{j \in F}$. Extension by zero gives a supported weight $\delta^0 = (b, -q(b))$. Since $D = I$, the least frozen coordinates of a regular theta function with mutable label $-b$ are $\beta_B(-b)_j = \mathsf{E}(\mathbf{a}_j, b)$ by (4.1).

Proposition 12.9. *For every frozen vertex j, we have*

$$\mathrm{ord}_{x_j=0}(\vartheta_{-\delta^0}) = q_j(b) - \mathsf{E}(\mathbf{a}_j, b) \geqslant 0.$$

This order is zero if j admits a sink- or source-optimizing seed. It is also zero at every frozen vertex if b is reachable.

Among labels with mutable part $-b$, the frozen coordinates form $q(b) + \mathbb{N}^F$ for $\mathcal{T}$ and $\beta_B(-b) + \mathbb{N}^F$ for $\mathcal{B}_\theta^+$. Consequently, $V_{\mathrm{gen}} = \mathrm{Mid}^+$ exactly when $q_j(b) = \mathsf{E}(\mathbf{a}_j, b)$ for every b and j. In particular, the supported generic characters form a basis of Mid^+ if every frozen vertex admits an optimizing seed.

Proof. By [11, Definition 6.7 and Lemma 6.9], the supported weights with mutable part b are $\delta^0 - (0, a)$, $a \in \mathbb{N}^F$. Their negatives have frozen coordinates $q(b) + a$. Since $D = I$, (4.1) gives $\nu_j(\vartheta_{(-b,q)}) = q_j - \mathsf{E}(\mathbf{a}_j, b)$. Nonnegativity of these orders gives the stated frozen coordinates for regular theta functions. For $q = q(b)$ the function is regular by Lemma 12.8.

If j admits either kind of optimizing seed, it admits a source-optimizing seed by Lemma 1.5. Mutate the extension by zero to that seed. Mutation preserves the vanishing of its ordinary and decoration spaces at all frozen vertices, and its presentation weight transforms by [11, Theorem 2.2(4)]. No nonzero representation space maps into j at this seed. Hence both projective Betti numbers at j vanish by [11, (1.3)–(1.5)]. The jth degree coordinate is therefore zero and equals the frozen order by Lemma 3.6. If b is reachable, choose a seed in which its general mutable representation is purely negative. Its extension by zero is then purely negative as well, and the corresponding theta function is a mutable cluster monomial. All its frozen orders vanish.

By Lemma 12.8, replacing ϑ_g by $C_{\mathrm{gen}}(-g)$ for $g \in \mathcal{T}$ is a unitriangular change in the regular theta basis. Each lower set is finite by Lemma 12.2, so this change is invertible. It gives the decomposition

$$\mathrm{Mid}^+ = V_{\mathrm{gen}} \oplus \bigoplus_{h \in \mathcal{B}_\theta^+ \setminus \mathcal{T}} \mathbb{C}\vartheta_h.$$

Thus the characters span exactly when the frozen coordinates agree for every b. □

Corollary 12.10. *On μ-supported weights, $\delta \mapsto -\delta$ intertwines the raising operators of* [11] *with the theta raising operators and preserves their upper string lengths, for both reachable and dually reachable colors. If $\mathcal{T} = \mathcal{B}_\theta^+$, the regular theta basis of* Mid^+ *is biperfect in the sense of Berenstein–Kazhdan for the two presentation crystals with their compatible gradings. If also* $\mathrm{Mid}^+ = \mathcal{U}^+$*, this proves the assertion about theta functions in* [11, Conjecture 8.14]. *Both equalities hold whenever the supported generic characters span $\mathcal{U}^+$.*

Proof. In a sink-optimizing seed, the raising derivation is $M_i\partial_{x_i}$, the upper string length is $-\delta_i$, and raising changes δ by $e_i - m_i$ [11, Lemmas 5.3–5.4 and 6.2]. These are exactly the theta formulas for $g = -\delta$, and both sets of formulas commute with mutation. For a dually reachable color the same argument uses $\eta_\theta(-\delta) = -\check{\delta}$. Theorem 7.2 gives normalized perfectness. Every theta function is homogeneous for each compatible grading because differences between exponents in its pointed expansion are integral combinations of mutable exchange rows. Finally, if $V_{\mathrm{gen}} = \mathcal{U}^+$, the inclusions $V_{\mathrm{gen}} \subseteq \mathrm{Mid}^+ \subseteq \mathcal{U}^+$ give equality, and Proposition 12.9 gives $\mathcal{T} = \mathcal{B}_\theta^+$. □

Declaration of generative AI and AI-assisted technologies

The author used ChatGPT to explore proof arguments and revise the exposition. The author takes full responsibility for the mathematical content.

References

[1] R. Avdeev and A. Petukhov, *Branching rules related to spherical actions on flag varieties*, Algebr. Represent. Theory **23** (2020), 541–581.

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

[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] A. Berenstein and A. Zelevinsky, *Tensor product multiplicities, canonical bases and totally positive varieties*, Invent. Math. **143** (2001), 77–128.

[5] P. Cao, *F-invariant in cluster algebras*, arXiv:2306.11438v4 (2025).

[6] 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.

[7] M.-W. Cheung, T. Magee, T. Mandel, and G. Muller, *Valuative independence and cluster theta reciprocity*, arXiv:2505.09585v2 (2026).

[8] H. Derksen and J. Fei, *General presentations of algebras*, Adv. Math. **278** (2015), 210–237.

[9] H. Derksen, J. Weyman, and A. Zelevinsky, *Quivers with potentials and their representations I: Mutations*, Selecta Math. (N.S.) **14** (2008), 59–119.

[10] J. Fei, *Tropical F-polynomials and general presentations*, J. Lond. Math. Soc. (2) **107** (2023), 2079–2120.

[11] J. Fei, *Crystal structure of upper cluster algebras*, arXiv:2309.08326v4 (2026).

[12] J. Fei, *Schubert Cox rings and root operators*, Hexagon 2611.00668v1 (2026).

[13] J. Fei, *Ramified quivers with potentials, monodromy characters, and cluster bases I*, arXiv:2610.09868v2 (2026).

[14] S. Fomin and A. Zelevinsky, *Double Bruhat cells and total positivity*, J. Amer. Math. Soc. **12** (1999), 335–380.

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

[16] S. Fomin and A. Zelevinsky, *Cluster algebras IV: Coefficients*, Compos. Math. **143** (2007), 112–164.

[17] L. Francone, *The minimal monomial lifting of cluster algebras I: Branching problems*, Adv. Math. **500** (2026), 111102. arXiv:2310.11808.

[18] M. Gekhtman, M. Shapiro, and A. Vainshtein, *Cluster algebras and Poisson geometry*, Mosc. Math. J. **3** (2003), 899–934.

[19] V. Genz, G. Koshevoy, and B. Schumann, *Polyhedral parametrizations of canonical bases & cluster duality*, Adv. Math. **369** (2020), 107178.

[20] K. R. Goodearl and M. T. Yakimov, *Quantum cluster algebra structures on quantum nilpotent algebras*, Mem. Amer. Math. Soc. **247** (2017), no. 1169.

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

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

[23] V. G. Kac, *Simple irreducible graded Lie algebras of finite growth*, Math. USSR-Izv. **2** (1968), 1271–1311.
[24] V. G. Kac and D. H. Peterson, *Defining relations of certain infinite dimensional groups*, in *Élie Cartan et les mathématiques d'aujourd'hui*, Astérisque, no. S131 (1985), 165–208.
[25] J. Kamnitzer, *Perfect bases in representation theory: three mountains and their springs*, in *Proceedings of the International Congress of Mathematicians 2022*, vol. 4, EMS Press, 2976–2996.
[26] M. Kashiwara, *On crystal bases of the q-analogue of universal enveloping algebras*, Duke Math. J. **63** (1991), 465–516.
[27] M. Kashiwara, *On crystal bases*, in *Representations of Groups* (Banff, 1994), CMS Conf. Proc. **16** (1995), 155–197.
[28] K. Kaveh, *Crystal bases and Newton–Okounkov bodies*, Duke Math. J. **164** (2015), 2461–2506.
[29] 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.
[30] Y. Kimura, F. Qin, and Q. Wei, *Twist automorphisms and Poisson structures*, SIGMA **19** (2023), 105, 39 pp.
[31] C. Lenart, *On the combinatorics of crystal graphs, I. Lusztig's involution*, Adv. Math. **211** (2007), 204–243.
[32] C. Lenart, S. Naito, D. Sagaki, A. Schilling, and M. Shimozono, *A uniform model for Kirillov–Reshetikhin crystals II. Alcove model, path model, and $P = X$*, Int. Math. Res. Not. IMRN (2017), no. 14, 4259–4319.
[33] P. Littelmann, *Cones, crystals, and patterns*, Transform. Groups **3** (1998), 145–179.
[34] G. Lusztig, *Canonical bases arising from quantized enveloping algebras*, J. Amer. Math. Soc. **3** (1990), 447–498.
[35] S. Morier-Genoud, *Geometric lifting of the canonical basis and semitoric degenerations of Richardson varieties*, Trans. Amer. Math. Soc. **360** (2008), 215–235.
[36] G. Muller, *$\mathcal{A} = \mathcal{U}$ for locally acyclic cluster algebras*, SIGMA **10** (2014), 094, 8 pp.
[37] G. Muller, *The existence of a maximal green sequence is not invariant under quiver mutation*, Electron. J. Combin. **23** (2016), no. 2, Paper 2.47.
[38] T. Nakanishi and A. Zelevinsky, *On tropical dualities in cluster algebras*, in *Algebraic Groups and Quantum Groups*, Contemp. Math. **565** (2012), 217–226.
[39] F. Qin, *Dual canonical bases and quantum cluster algebras*, arXiv:2003.13674v3 (2023).
[40] F. Qin, *Bases for upper cluster algebras and tropical points*, J. Eur. Math. Soc. **26** (2024), 1255–1312.
[41] F. Qin, *Applications of the freezing operators on cluster algebras*, J. Algebra Appl. **25** (2026), 2550357, 30 pp.
[42] Q. Tang, *Biperfectness of the theta basis of a maximal unipotent subgroup*, Hexagon 2610.00603v1 (2026).

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