hexagon:2611.00790v2

CRYSTAL STRUCTURES ON THETA BASES

JIARUI FEI

Abstract. Suppose the mutable part of a skew-symmetrizable seed admits a reddening sequence. We prove that the regular theta basis of the middle algebra is normalized perfect for the derivations attached to frozen sinks and sources reachable by mutation. Its g-vectors are the lattice points of a rational polyhedral cone in every seed. When i is a sink, the theta function with index g can be raised exactly g_i times, and its highest nonzero divided derivative is again a theta function.

Local nilpotence and the Serre relations for each family of derivations hold on the upper cluster algebra with polynomial frozen variables even without reddening. We characterize exact sink–source pairs by a tropical criterion. Their Chevalley generators give integrable actions whose connected Cartan components are of finite or affine type. Under reddening the resulting theta crystals are normal. We apply these results to Lusztig's involution on the theta basis of the base affine space and to a level-zero crystal of type $C_n^{(1)}$ on the Plücker ring of the Lagrangian Grassmannian.

Contents

Introduction 1
1. Theta functions and coefficients 4
2. Canonical boundary derivations and degrees 7
3. Regular theta cones and tropical crystals 9
4. Tropical DT duality and exact boundary pairs 14
5. Primitive theta lifts and changes of coefficients 18
6. Exact boundary pairs and Lie actions 21
7. Perfect theta bases and crystal structures 27
8. Cartan type and geometric applications 32
9. Theta F-pairings and maximal adjunction 39
10. Maximal operators and string coordinates 43
11. Pointed biperfect bases and transitions 45
Declaration of generative AI and AI-assisted technologies 48
References 48

Introduction

The two clusters of the Plücker ring of the Lagrangian Grassmannian $\mathrm{LGr}(2,4)$ are related by the exchange $xy = b_1^2 + b_0 b_2$, where b_0, b_1, b_2 are frozen. In the cluster containing x, consider the derivation $R = x\partial_{b_1}$, which fixes x, b_0, b_2. Differentiating the exchange relation gives

$$R(y) = 2b_1, \qquad R(b_1) = x, \qquad R(x) = 0.$$

The three theta functions y, b_1, x thus form a raising string. Its length is the degree of y in b_1 on the x-chart, and its top element is $R^2(y)/2 = x$. Along the frozen divisor $b_1 = 0$, the function y has order zero.

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).

We prove that the same distinction governs theta functions for skew-symmetrizable seeds. A frozen vertex that becomes a sink defines a derivation of the upper cluster algebra. When the mutable pattern admits a reddening sequence, the degree filtration of this derivation determines strings in the regular theta basis. We characterize the source derivations that act as lowering operators on the same strings. When each chosen sink has such a partner, the raising and lowering derivations generate an integrable action of a derived Kac–Moody algebra. The possible connected Cartan types are finite and affine.

The theta functions are those of Gross, Hacking, Keel, and Kontsevich [20]. For skew-symmetric seeds, crystal operators were constructed on supported presentation weights in [11], where quivers with potential relate them to generic cluster characters. Tang proved the *biperfectness* of the theta basis for the coordinate ring of a maximal unipotent subgroup, establishing perfectness for both the left and the right simple-root derivations [37, Theorem 1.1]. Here the exchange matrix and its symmetrizing weights determine the operators. The construction applies to regular theta functions even when it is unknown whether they span 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 $R_k^\star = M_k^\star\partial_{x_k}$. A *reddening sequence* is a sequence of mutable mutations for which the final c-matrix in principal coefficients is a negative permutation matrix [27].

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 have the following formulas in a sink-optimizing seed.*

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

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 [36] identifies their sum with a single theta function in the resulting seed. This extends Tang's highest-degree argument in full exchange rank [37, 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). The opening example is an exact self-pair at b_1, with boundary monomials x and y.

One can also test exactness in any reference seed from the mutable pattern, the symmetrizing weights, and the frozen exchange vectors $\mathbf{a}_j = (b_{ju})_{u\in U}$. 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. Under reddening its support has a componentwise greatest element f_b. Write D_U and $B_{U,U}$ for the mutable submatrices of D and B. For $a, b \in \mathbb{Z}^U$ define

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

We orient the mutable tropical Donaldson–Thomas transformation T by $T^{-1}b = -b - f_bB_{U,U}$.

Theorem C. *Under reddening, $\mathsf{E}(a,b) = \mathsf{H}(b,Ta)$ for all $a, b \in \mathbb{Z}^U$. 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 checks the mutable coordinates of the defining identity. The remaining equalities check its frozen coordinates. This criterion is the theta analogue of [11, Lemma 4.29].

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 D. *Suppose each $i \in I$ has an exact source partner $\sigma(i)$. These partners are distinct and $d_i = d_{\sigma(i)}$. 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 is of finite or affine type.*

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 8.1). 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].

Applications. On the base affine space, the pinned Lusztig involution permutes the theta functions in each highest-weight summand and induces the corresponding crystal involution (Theorem 8.6). On the Plücker ring of $\mathrm{LGr}(n, 2n)$, $n \geqslant 2$, the affine raising operator is a scalar multiple of the negative highest-root derivation. The affine action factors through evaluation of the loop parameter at a nonzero scalar, and its theta crystal has type $C_n^{(1)}$ and level zero (Proposition 8.8).

Two further results concern maximal raising. It takes all monomials of a given cluster to monomials in one common cluster (Theorem 10.2). We also determine the real valuation directions for which maximal raising has an adjoint. They form a rational polyhedral cone that can be tested on the seed variables and one further cluster monomial (Theorem 9.6).

In the skew-symmetric case with mutable reddening and finite-dimensional Jacobian algebra, we compute the orders of theta functions indexed by supported presentation weights along the frozen divisors. These orders determine when the supported generic characters span Mid^+ (Proposition 4.8). Section 5 describes the effect of changes of coefficients on the regular theta functions and proves that opposite cluster automorphisms reverse the mutable tropical Donaldson–Thomas transformation under reddening. Section 11 classifies the changes of biperfect basis that admit compatible pointed lifts before coefficient specialization.

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 $D = \mathrm{diag}(d_a)$ 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 [27]. 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 [33, 35].

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 [20] 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 [35, Theorem A.1.5].

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. We say that *upper theta spanning* holds when $\mathrm{Mid}^{\mathrm{loc}} = \mathcal{U}^{\mathrm{loc}}$.

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

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

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^+. Under upper theta spanning, Θ_{F_0} is a basis of $\mathcal{U}_{F_0}$.

Proof. The valuation $\nu_f = \mathrm{ord}_{x_f}$ 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 $v \in \mathbb{N}^U \setminus \{0\}$. 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 $\ker \widetilde{B} \cap \mathbb{R}^U_{\geqslant 0} = \{0\}$ in some seed, 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. By Gordan's theorem of the alternative, some linear functional is strictly positive on every column of $\widetilde{B}$. This is the convexity condition of [20, Theorem 0.3(7)]. The initial and

injective seeds give opposite mutable chambers, so the cluster complex is large in the sense of [20]. Proposition 0.7 of that paper gives the basis of $\mathcal{U}^{\mathrm{loc}}$. Lemma 1.1 gives the remaining assertions. □

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 [37, Proposition 4.2]. Its injective-reachability hypothesis follows from reddening. The proof uses the Donaldson–Thomas (DT) twist and its action on theta functions [28, 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 $t[-1]$ [35, 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 [35, 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 [20]. 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 [28], 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 [20] 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 $\mathcal{U}^+$.*

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 $x_k x_k' = P = A x_i^q + C$ with $q > 0$, where A and C are independent of x_i and x_k. Then

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

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. This proves local nilpotence. 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

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

It follows that $c_{ii} = 2$ and $c_{ij} = -a_{ij}$ for $i \neq j$ define a symmetrizable generalized Cartan matrix, with symmetrizer $\operatorname{diag}(d_i)$. 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, define $\rho_j^\star(P) = \deg_{x_j} P$ in a source-optimizing seed. This equals the nilpotence degree of P for $R_j^\star$, that 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. Throughout this section the mutable pattern 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 $\operatorname{val}_n(H) = \min_{h_a \neq 0} \langle a, n \rangle$. 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 $n \in N^\bullet$. Theta reciprocity [7, Corollary 3.2, Claim 4.16 and Theorem 5.19] gives the pairing

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

whenever the two values are finite.

Lemma 3.1. *For every $n \in N^\bullet$, the dual theta function $\vartheta_n^\vee$ is a Laurent polynomial in each seed with the frozen variables inverted.*

Proof. An isolated mutable vertex with zero exchange column contributes the rank-one factor of Remark 1.2. Its dual normal is zero, so the dual factor is a Laurent monomial. We may omit these factors.

Use the basis $((d_a/\ell)e_a^*)_{a\in V}$ to identify $N^\bullet$ with $\mathbb{Z}^V$, and write $E_U : \mathbb{Z}^U \hookrightarrow \mathbb{Z}^V$ for coordinate inclusion. In this basis and its dual, the exponent and normal matrices of the dual datum are E_U and $B_{V,U}$. Consider the injective lattice map

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

It is a linear morphism of seed data to

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

since $\iota E_U = \widehat{P}$ and $\iota^{\mathsf{T}}\widehat{Q} = B_{V,U}$. The target is a principal-coefficient $\mathcal{A}$-seed with mutable matrix $B_{U,U}^{\mathsf{T}} = -\widetilde{B}_{U,U}$, together with independent frozen Laurent variables. Its weight matrix in the convention of [7] is ℓD_U^{-1}, since $D_U B_{U,U}^{\mathsf{T}}$ is skew-symmetric. Inverting each auxiliary coefficient as a torus coordinate identifies this seed with the opposite of the original mutable principal-coefficient seed, up to the independent Laurent variables. By reddening and Proposition 1.4, all its integral theta functions are Laurent polynomials.

The linear-morphism theorem and its positive-chamber formulation [7, Theorem 4.9 and Section 4.4] send $\vartheta_z^\vee$ to the target theta function of index ιz. This formulation applies even if the source datum has no full-dimensional positive chamber. Since ι is injective, distinct exponents in the source remain distinct in the target. The target has finite support, hence so does $\vartheta_z^\vee$. □

Cheung, Magee, Mandel and Muller describe the theta basis of a partial compactification by its boundary valuations [7, Corollary 1.2]. The preceding finiteness lemma gives a finite set of inequalities in every seed.

Corollary 3.2. *In each seed the g-vectors of regular theta functions are the lattice points of a rational polyhedral cone. More precisely, let $\mathcal{S}_f^\vee$ be the exponent support of $\vartheta_{e_f^*}^\vee$ in that seed. Then*

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

where $M_{\mathbb{R}} = M \otimes_{\mathbb{Z}} \mathbb{R}$.

Proof. By Lemma 3.1, $\mathcal{S}_f^\vee$ is a finite subset of the rational lattice $N^\bullet$. Reciprocity gives $\operatorname{ord}_{x_f}(\vartheta_g) = \min_{a\in\mathcal{S}_f^\vee}\langle g, a\rangle$. The frozen valuation criterion in the proof of Lemma 1.1 now gives the assertion. □

3.2. **Boundary polytopes and closure under raising.** The initial monomial of $\vartheta_{e_f^*}^\vee$ has exponent e_f^*, and every bend adds a nonnegative integral combination of the dual wall exponents $q_u^\bullet$. Thus we may write

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

where z^α denotes the dual Laurent monomial of exponent α. The set A_f contains 0. We identify its elements with their mutable coordinates and call $\mathcal{P}_f = \operatorname{conv}(A_f)$ the *boundary polytope* at f. For a mutable g-vector p, regularity at f requires that its frozen coordinate be at least

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

where the dot denotes the pairing of mutable coordinates.

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

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

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

Proof. Let $0\neq c\in A_i$. A broken line contributing to $z^{e_i^*+c}$ has a first nontrivial bend. Its wall normal pairs nontrivially with the incoming exponent e_i^* and is proportional to $\sum_u h_u p_u$ for some $h\in\mathbb{N}^U\backslash\{0\}$. Indeed, the rescaled wall normals of the dual datum are the original p_u [7, Lemma 4.3 and Section 4.4]. Since

$$\left\langle e_i^*, \textstyle\sum_u h_u p_u\right\rangle = -d_i^{-1}\sum_u d_u m_u h_u,$$

some u has $m_u h_u > 0$. At this bend the exponent increases by a positive multiple of $\sum_u h_u(d_u/\ell)e_u^*$, and every later bend adds a nonnegative vector. Hence $m_U\cdot c>0$. By Lemma 3.3, $m_U\cdot c\geqslant 1$ at every nonzero vertex of $\mathcal{P}_i$.

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 $\mathcal{N} = \{e_f^* + c : f\in F,\ c\in\operatorname{Vert}(\mathcal{P}_f)\}$ 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 [36, 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. □

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 $\mathrm{wt}_j(g) = w_j(g)$ and $\lambda_j(g) = \rho_j(g) + \mathrm{wt}_j(g)$. The family is integral if every w_j takes integral values on $\mathbb{Z}^V$.

Corollary 3.9. *For an adapted integral family, $(r_i, l_i, \rho_i, \lambda_i, \mathrm{wt}_i)_{i\in I}$ 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 $\mathrm{wt}_j(r_i g) = \mathrm{wt}_j(g) + c_{ji}$ when $r_i g$ is defined. The partial inverse and raising-length axioms follow from Theorem 3.7. By definition $\lambda_i = \rho_i + \mathrm{wt}_i$ 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.

3.5. **The opposite crystal.** 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.

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 be a monomial in the mutable variables of a cluster and let $b = -g(P)_U$ in the reference seed. For every $j \in F$,*

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

Proof. The order of P along $x_j = 0$ is zero in any seed in which P is a mutable monomial, and frozen valuations are invariant under mutation. In the reference seed, the term of the separation formula indexed by ν has x_j-exponent

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

Positivity prevents cancellation when specialization identifies monomials. The least exponent is therefore zero, which gives the first equality. The codegree formula gives the second. □

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, for any rank of the extended exchange matrix.

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 $M_k^\star$. □

4.2. **The mutable DT transformation.** Suppose 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. *Under reddening, for all $a, b \in \mathbb{Z}^U$, we have that*

$$\mathsf{E}(a,b) = \mathsf{H}(b, Ta).$$

Proof. The polynomials F_b and the transformation T depend only on the mutable pattern. We may therefore work with principal coefficients, writing V and B for the vertex set and exchange matrix of that seed. Fix a, b, and let $c = Ta$. Choose finite theta functions L, P with full degrees $-\widetilde{b}, -\widetilde{c}$ whose mutable parts are $-b, -c$, respectively. The vector $\widetilde{a} = \operatorname{codeg} P$ has mutable part $T^{-1}c = a$, and $\widetilde{a} = -\widetilde{c} - f_c B_{U,V}$.

The principal extension admits an integral skew-symmetric matrix Γ satisfying

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

An explicit choice is given in the proof of Proposition 9.2. It follows that

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

We apply theta reciprocity in the form of [7, Claim 5.7 and Theorem 5.19]. In the notation of that paper, take $P_0 = -B_{U,V}^{\mathsf{T}}$, let Q_0 be the inclusion of the mutable coordinate lattice in $\mathbb{Z}^V$, and take the diagonal matrix D_U^{-1}. Then $Q_0^\bullet = Q_0 D_U$. The identities $\Lambda P_0 = Q_0^\bullet$ and $\Lambda^{\mathsf{T}} P_0 = -Q_0^\bullet$ for $\Lambda = -\Gamma$ give the required Λ-structure.

Take $m_1 = -\widetilde{a}^{\mathsf{T}}$ and $m_2 = -\widetilde{b}^{\mathsf{T}}$ in the reciprocity identity. The positive-chamber theta function at m_2 is L. By [7, Lemma 4.14], the negative-chamber theta function at m_1 is the positive-chamber theta function for $(P_0, -Q_0)$ with the same label. Negation sends this datum to $(-P_0, Q_0)$ and the label to $\widetilde{a} = \operatorname{codeg} P$. Proposition 1.6 thus identifies the negative-chamber polynomial with $P(x^{-1})$. All these functions are finite, so their valuations

can be read from their Laurent polynomials [7, Corollary 3.2]. Replacing x by x^{-1} negates the valuation direction. Reciprocity now reads

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

By the pointed expansions, the two sides are $-\widetilde{b}\Gamma\widetilde{a}^{\mathsf{T}} - \mathsf{E}(a,b)$ and $-\widetilde{c}\Gamma\widetilde{b}^{\mathsf{T}} - \mathsf{E}(b,c)$. Substituting the identity for $\widetilde{b}\Gamma\widetilde{a}^{\mathsf{T}}$ gives $\mathsf{E}(a,b) = bD_U f_c^{\mathsf{T}} + \mathsf{E}(b,c) = \mathsf{H}(b,c)$. □

4.3. **The boundary criterion.**

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*

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

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 $j \notin \{i,k\}$ both summands vanish, since each is nonnegative. For fixed i, the left side depends only on j. Subtracting $d_j\delta_{ij}$ leaves $d_j\delta_{kj}$, so an exact partner k, if it exists, is unique.

4.4. **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.

4.5. **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. We use the boundary representations E_i, the dual boundary representations $E_j^\star$, and the Auslander–Reiten translation τ of [11, Sections 1 and 4]. Their projective presentation weights correspond to the negatives of g-vectors. 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, Section 4]. *Then (i,j) is an exact boundary pair if and only if $\tau^{-1}E_i \cong E_j^\star$.*

Proof. In a sink-optimizing seed for i, the projective weight of $\tau^{-1}E_i$ is $m_i - e_i$. Dually, the projective weight of $E_j^\star$ is $e_j - g(M_j^\star)$ in a source-optimizing seed for j [11, Section 4]. Both expressions for the latter weight obey projective-weight mutation, so they also agree in the chosen seed for i. There $\mathrm{codeg}(M_i) = m_i$. Equality of the two presentation weights is

thus precisely the defining equation for exactness. Both presentations are rigid, and a rigid presentation is determined up to isomorphism by its weight [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.

4.6. **Supported presentation weights and regular theta functions.** 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.

We use the order $h \lhd g$ obtained by projecting strict dominance in the principal extension. Thus some lift of h has strictly smaller degree than the lift $(g, 0)$ in every seed (Definition 11.1).

Lemma 4.7. *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 [35, proof of Theorem 5.2.2].

Qin's comparison of pointed bases in full rank [35, 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 11.2 shows that all remaining terms have labels strictly below g. Finally, the support-dimension formula [35, 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 $\delta^0 = (b, -q(b))$. For comparison, set $\beta_B(-b)_j = \mathsf{E}(\mathbf{a}_j, b)$.

Proposition 4.8. *For every frozen vertex j, we have*

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

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 $q(b) + a$. For a theta function with g-vector $(-b, q)$, the minimum x_j-exponent in the separation formula is $q_j - \mathsf{E}(\mathbf{a}_j, b)$ since $D = I$. This gives the stated frozen coordinates for regular theta functions. Taking $q = q(b)$ gives the formula for the order, whose nonnegativity follows from Lemma 4.7.

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 4.7, 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 11.2, 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 4.9. *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 4.8 gives $\mathcal{T} = \mathcal{B}_\theta^+$. □

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.

Let ν_f denote order along $x_f = 0$. Frozen valuations are invariant under mutation, and an element of $\mathcal{U}^{\mathrm{loc}}$ belongs to $\mathcal{U}^+$ precisely when all these orders are nonnegative. In the separation formula for $\vartheta_{(p,q)}$, a monomial of F_{-p} with exponent α contributes $\mathbf{a}_f D_U \alpha^{\mathsf{T}}/d_f$ to the x_f-exponent. Positivity ensures that terms of least exponent cannot cancel upon specialization. Taking their minimum gives

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

Define $\beta_B(p)_f = \mathsf{E}(\mathbf{a}_f, -p)/d_f$. This is an integer by (5.1) 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*

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

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 (5.1) 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 [20, 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, (5.1) 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 a field isomorphism Φ 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 [28, Section 2.12]. 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 [35, 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,*

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

Write $G = \gamma_\mu$ for the mutable action of normalized reddening and $N(p) = -p$. Then

$$kGk^{-1} = G^{-1}, \qquad kG = Ns_\mu.$$

If $k^2 = \mathrm{Id}$, then $G = k(Ns_\mu)$ 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 $Ns_\mu N$. Applying Φ to $\vartheta_g^B = \vartheta^{-B}_{\eta_\theta(g)}$ and comparing the mutable degrees gives

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

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

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. □

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 $\mathfrak{n}_+(C)$ [21], 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*

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

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 [36, 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 $x^{r v_i} q_g = q_{g + r v_i}$. If $n = g_i$, we therefore have $R_i^n M_j^\star = n! q_{g + n v_i}$. Under

(6.3), $n = 1 + \delta_{ij}$, and $g + nv_i$ equals e_j when $i \neq j$ and m_i when $i = j$. The corresponding functions are x_j and M_i, proving the normalization.

Conversely, either normalization implies $g_i = 1 + \delta_{ij}$. The pointed term x^g contributes $n!x^{g+nv_i}$ to the maximal derivative. All contributions have nonnegative coefficients, while the normalized derivative is a single monomial. Its exponent must therefore be $g + nv_i$. Solving for g proves (6.3). □

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

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

These derivations define an integrable $\mathfrak{sl}_2$ action on $\mathcal{U}^+$. If $i \neq j$, then $d_i = d_j$ and $R_j^\star M_i = x_i$.

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

Proof. Assume first that B is invertible and use $\Lambda = B^{-1}D$. Write $G = Z_j^- = x^g$ for the frozen Laurent monomial of Lemma 6.1. The source formula gives

$$R^\star = \frac{1}{d_j x_j}\left(\{N, -\} - N\frac{\{G, -\}}{G}\right). \tag{6.4}$$

On a seed coordinate it reads

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

Suppose $i \neq j$. The normalization $RN = x_j$ gives $N = x_i x_j/M + C$, with C independent of x_i. Write $a = d_i/d_j$ and $q = e_i + e_j - m - g$. Differentiating (6.5) gives $[R, R^\star](x_k) = d_j^{-1}(q\Lambda)_k x_k$ for $k \neq i$. Equation (6.2) gives $R^\star M = ax_i$, so $[R, R^\star](x_i) = -ax_i$. Comparing the coefficients of x_i on the two sides of (6.5) for each frozen index $k = f$, using $R^\star(x_f) = \delta_{jf}N$ and $N = x_i x_j/M + C$, gives $d_j^{-1}(q\Lambda)_i = 0$, $d_j^{-1}(q\Lambda)_j = 1$, and all remaining frozen entries vanish. Thus H is diagonal with weights $w = d_j^{-1}q\Lambda - ae_i$.

This grading is compatible. Indeed, $Bw^\mathsf{T} = -d_j^{-1}Dq^\mathsf{T} - aBe_i^\mathsf{T}$. For $k \in U$, we have $q_k = -m_k$ and $m_k = -b_{ik} = (d_i/d_k)b_{ki}$, so the kth entry vanishes. The identities $RM = R^\star N = 0$, $RN = x_j$, and $R^\star M = ax_i$ imply $HM = aM$ and $HN = -N$. Hence $[H, R] = 2aR$. Every diagonal derivation preserves a log-canonical bracket and commutes with the Euler derivation $\{G, -\}/G$. Applying H to (6.4) therefore gives $[H, R^\star] = -2R^\star$. Jacobi now yields $0 = [H, [R, R^\star]] = 2(a - 1)H$. Since $H(x_j) = x_j$, we obtain $a = 1$, proving the distinct case.

For $i = j$, write $x = x_i$ and $d = d_i$. The normalization $R^2N = 2M$ gives $N = x^2/M + Ax + C$, where A, C are independent of x. Let $q = 2e_i - m - g$. For $k \in U$, (6.5) takes the form

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

with A_k, C_k independent of x. Polynomiality in the original frozen variable x forces $C_k = 0$. It follows that $[R, R^\star](x_k) = d^{-1}(q\Lambda)_k x_k$. The bracket kills every frozen variable, including x, because $RN = R^\star M$. Comparing the coefficients of x^2/M in $R^\star(x_f)$ gives $d^{-1}(q\Lambda)_i = 1$ and zero at the other frozen entries. Thus H is diagonal with weights $w = d^{-1}q\Lambda - e_i$. As in the case $i \neq j$, one has $q_k = -m_k$ for $k \in U$, so $B_{U,V}w^\mathsf{T} = 0$.

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

For general B, use the invertible coefficient extension of Lemma 6.1, keeping the original frozen variables polynomial. Denote the lifted operators by $\widehat{R}$ and $\widehat{R^\star}$ and their multipliers by $\widehat{M}$ and $\widehat{N}$. In the lifted optimizing seed for i, Laurent positivity gives nonnegative integral coefficients for $\widehat{N}$. Thus every nonzero term of $\widehat{R}\widehat{N}$ has a positive integral coefficient, and every nonzero term of $\widehat{R}^2\widehat{N}$ has a positive even coefficient. For $i \neq j$ the first derivative specializes to x_j. For $i = j$ the second specializes to $2M$. Consequently, for some integral γ,

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

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

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

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

The diagonal bracket in the extension kills each auxiliary variable. Its compatible grading therefore descends to a compatible grading for B. To prove integrality, let $f \in \mathcal{U}^+$ 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, for $i \neq j$, evaluation of $[R, R^\star] = H$ on x_j gives $RN = x_j$. For $i = j$, apply $[H, R] = 2R$ to x_i. Since $H(x_i) = 0$, we find $HM = 2M$. The mixed identity gives $RN = R^\star M$, whence $R^2 N = R(R^\star M) = [R, R^\star]M = 2M$. The preceding proposition proves exactness. □

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 $j \neq k$. Each sink has at most one exact source partner.

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 $[R_h, R_k^\star] = 0$. To obtain the second degree formula and commutator, apply the same argument to (k,i) for $-B$. This reversed pair is exact by Proposition 6.4. Indeed, for $k \neq i$, Theorem 6.5 gives $R_k^\star M_i = x_i$. For $k = i$, the mixed identity gives $R_i^\star M_i = R_i M_i^\star$. Applying $R_i^\star$ and using $R_i^\star M_i^\star = 0$ and $H M_i^\star = -2M_i^\star$, where $H = [R_i, R_i^\star]$, yields $(R_i^\star)^2 M_i = 2M_i^\star$. For $j \neq k$ the vanishing of $[R_i, R_j^\star]$ excludes exactness of (i,j) by Theorem 6.5. □

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

Theorem 6.8. *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. Theorem 6.5 gives the rank-one relations, the compatible integral gradings, and the equality $d_i = d_{\sigma(i)}$. Lemma 6.7 gives $[R_j, L_i] = 0$ for $j \neq i$. If $\sigma(i) = \sigma(j)$ with $i \neq j$, then $H_i = [R_i, L_j] = 0$, contrary to $[H_i, R_i] = 2R_i \neq 0$.

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

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

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

By the sharp sink Serre relation, the first vanishing power is $n = 1 - c_{ij}$. Substitution gives $u_{ij} = c_{ij}$ and therefore $v_{ij} = -c_{ij}$. Dually, L_j is killed by $\operatorname{ad} R_i$ and has weight $-c_{ij}$. At the first vanishing power of $\operatorname{ad} L_i$, the dual rank-one identity gives that power as $1 - c_{ij}$. Such a power exists by the source Serre relation of 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. □

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 [37, 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 [36, 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 [20, 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 [32, 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 $\ker E^{(N-m)/2+1}$. Thus E-perfectness gives a basis on 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 $\bar{\varpi}_i$, $i \in I$, denote the standard basis of the coroot-coordinate lattice $\mathbb{Z}^I$. The weight of a label b is $\mathrm{wt}(b) = \sum_i \mathrm{wt}_i(b)\bar{\varpi}_i$, where $H_i\vartheta_b = \mathrm{wt}_i(b)\vartheta_b$. Let $\bar{\alpha}_j = \sum_i c_{ij}\bar{\varpi}_i$. The Cartan brackets and Lemma 7.4 give the crystal axioms with these coroot-coordinate weights. On affine components these are classical weights, with the null root mapped to zero, as in [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.

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 [25, 30]. 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. The boundary criterion in Section 4.3 determines when the source maps provide these inverses.

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

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

The class of q_i records how R_i changes the degree in the universal rational grading. Each compatible rational grading is a linear form on $L_{\mathbb{Q}}$.

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

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

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 8.4 shows.

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

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

It is an involution and sends weight ν to $\nu - \nu(h_i)\bar{\alpha}_i$. In an i-optimizing seed, $g(s_i b) = g(b) - \operatorname{wt}_i(b)(m_i - e_i)$. On each finite-type rank-two restriction, normality identifies these reflections with Kashiwara's Weyl group action on a union of highest-weight crystals [26, 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 [21, 23] give

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

together with the braid relations. This is the Tits extension action of [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

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

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.

7.4.2. *Opposite symmetries and Lusztig involutions.* Let $J \subseteq I$ have finite type and let $w_{0,J}$ be the longest element of its Weyl group. Define i^* 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 $\xi_\lambda r_i = l_{i^*}\xi_\lambda$ [29, 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 7.9. *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*

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

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

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

Here s_μ is the mutable part of the tropical map of Φ from B to $-B$, $N(p) = -p$, and G is the mutable map of the normalized reddening transformation of Proposition 5.4.

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 $\xi^2_{\mathrm{mut}} = \mathrm{Id}$ 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.

Corollary 7.10. *Assume mutable reddening and suppose every $i \in F$ has an exact source partner $\sigma(i)$. Take F as the crystal index set. Then σ is a permutation and γ^{-2}, followed by the frozen relabeling $i \mapsto \sigma(i)$, is a direct cluster automorphism Ψ. On the regular theta basis it permutes the colors by σ and acts by*

$$\Psi(p, a) = (G^{-2}(p), \sigma a).$$

Every opposite cluster automorphism with action κ and frozen permutation π satisfies

$$\pi\sigma\pi^{-1} = \sigma^{-1}, \qquad \kappa\Psi\kappa^{-1} = \Psi^{-1}$$

on frozen indices and on the regular theta basis, respectively.

Proof. Lemma 6.7 gives uniqueness of each partner, and Theorem 6.8 gives injectivity and $d_{\sigma(i)} = d_i$. Thus σ permutes the finite set F. Corollary 4.4 gives $\mathbf{a}_{\sigma(i)} = T^{-2}\mathbf{a}_i$. The negative frozen row $-\mathbf{a}_i$ transforms as a mutable pointed degree, as in the proof of Lemma 1.5. After γ^{-2} the row at i is $T^{-2}\mathbf{a}_i = \mathbf{a}_{\sigma(i)}$. Relabeling i by $\sigma(i)$ therefore restores every coefficient entry b_{iu} with $u \in U$. The principal matrix has returned to its initial value. Since $d_{\sigma(i)} = d_i$, skew-symmetrizability also restores each b_{ui}. These are precisely the data of a direct identification of the valued ice exchange pattern. Corollary 5.3 and Proposition 5.2 give the basis formula and the action on colors.

An opposite identification takes $(i, \sigma(i))$ to $(\pi\sigma(i), \pi(i))$. Uniqueness of the partner gives $\sigma\pi\sigma = \pi$. In particular $\kappa r_i = l_{\pi\sigma(i)}\kappa$ and $\kappa l_i = r_{\pi\sigma(i)}\kappa$. Conjugation by π sends σ to σ^{-1}, and conjugation by the mutable map k of κ sends G to G^{-1} by Proposition 5.4. The two coordinate formulas give the last identity. □

8. Cartan type and geometric applications

8.1. **Growth of Lie algebras of derivations.** Let $C = (c_{ij})_{i,j\in I}$ be a symmetrizable generalized Cartan matrix and let D_I be a positive rational diagonal matrix such that $D_I C$ 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 8.1. *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 write $K_0 = \Bbbk(x_1, \dots, x_d)$. The elements $D_a(x_j)$ generate a finite extension L of K_0 inside K. Each D_a maps K_0 into L. Differentiating the separable minimal polynomial of any $y \in L$ over K_0 then gives $D_a(y) \in L$. Restriction to L is injective on $\mathfrak{d}$, since K/L is algebraic and separable.

Write $L = \Bbbk(y_1, \dots, y_s)$ and choose a common denominator $q \in \Bbbk[y_1, \dots, y_s]$ for the $D_a(y_b)$. The derivations preserve $R = \Bbbk[y_1, \dots, y_s, q^{-1}]$. Filter R by total degree in these $s+1$ algebra generators. Since there are finitely many generators and derivations, the Leibniz rule gives a constant c such that $D_a(R_{\leqslant m}) \subseteq R_{\leqslant m+c}$ for every a, m. A Lie word of length at most n thus maps every generator into $R_{\leqslant 1+cn}$. Its values on the generators determine it on L, so

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

The last estimate follows from the Hilbert polynomial, since R has Krull dimension d. □

Theorem 8.2. *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.*

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 [22, 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 8.1. □

Corollary 8.3. *Under the hypotheses of Theorem 8.2, 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 [21, 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 $v = ae_i + 2e_j$ 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. □

8.2. **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

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

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 [37, Theorem 1.1]. Its bicrystal is therefore Kashiwara's $B(\infty)$ [24, Theorem 1.7].

Let w_0 be the longest element of the Weyl group and define i^* by $w_0\alpha_i = -\alpha_{i^*}$. The frozen variables are $A_j = \Delta_{\varpi_j, w_0\varpi_j}$. With the chosen pinning, D_i^R is the sink derivation at A_{i^*} and D_i^L is the source derivation at A_i [37, Lemmas 5.1 and 5.2]. We index both families by their simple roots.

Proposition 8.4. *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.

A reduced-word minor seed has one variable for each position in the word, and $\operatorname{rank} G$ of these variables are frozen [34, Theorem 8.1.2]. Its extended exchange matrix has full rank by the triangular minor obtained by taking, for each nonterminal occurrence of a letter, the column at its next occurrence [12, proof of Proposition 3.6]. Thus the grading quotient $L_{\mathbb{Q}}$ of Section 3 has dimension $\operatorname{rank} G$. For either family, the conjugation grading sends the degree shifts $[q_i]$ to α_i. Its coroot coordinates take integral values and give the columns of the nonsingular matrix C on these shifts. The $[q_i]$ therefore form a basis of $L_{\mathbb{Q}}$, which proves uniqueness.

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

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

Goodearl and Yakimov proved that these seeds are mutation-equivalent for each word, as stated in [34, 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 [34, Remarks 8.1.4 and A.3.1] therefore give, for every mutable k and every a,

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

Here the superscripts indicate the suffix and prefix seeds, respectively. Reversal sends last occurrences to first occurrences and preserves the symmetrizing weights. Thus ι^* identifies the prefix seed with the opposite of the suffix seed and sends A_j to A_{j^*}.

These seeds are injective-reachable [34, 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 [24, 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].

8.3. **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 $\operatorname{wt}(f)$. 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 and their mutable pattern admits a reddening sequence. The regular theta basis is homogeneous for both torus actions [12, Corollary 3.9, 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 $\Xi_\lambda F_i = E_{i^*}\Xi_\lambda$ [29, Section 2.4]. Morier-Genoud computes this involution in parametrizations of the canonical basis [31, 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

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

Proposition 8.5. *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. First suppose that G is simple of rank at least two. Fix a reduced word $\mathbf{j} = (j_1, \dots, j_N)$ for w_0, with partial products $v_k = s_{j_1} \cdots s_{j_k}$. Its lifted prefix variables are $x_k = \Delta_{\varpi_{j_k},v_k\varpi_{j_k}}$. Let $s(k)$ denote the next occurrence of j_k, or $+\infty$ if there is none. If a word begins with i, it contains another i. Otherwise $w_0\varpi_i = s_i\varpi_i$, which pairs positively with a neighboring coroot, whereas $w_0\varpi_i$ is antidominant. Reversing the word gives the same conclusion for its last letter.

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 [37, 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 $F_i x_{p_j} = \delta_{ij}x_1$. Let a_j be the first occurrence of j. Homogeneity of the lifted exchanges gives $\widetilde{b}_{p_jk} = -\sum_{r:j_r=j}\widetilde{b}_{rk}$ for each mutable k. The prefix exchange rule then yields

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

Indeed, for $j = j_k$ the contributions from consecutive occurrences cancel except when $k = a_j$. For $j \neq j_k$, intervals of letter j crossing $(k, s(k))$ contribute with opposite signs in pairs. The sole unpaired crossing occurs when $k < a_j < s(k)$. Since $a_i = 1$, we have $\widetilde{b}_{p_ik} = \delta_{k1}$. The equality $d_{p_i} = d_1$ gives $b_{p_ik} = \delta_{k1}$, hence $F_i = R^\star_{p_i}$.

The equalities $E_i x_{p_i} = 0$ and $[E_i, F_i]x_{p_i} = x_{p_i}$ imply $E_iF_ix_{p_i} = x_{p_i}$. Proposition 6.4 gives exactness. Since $\mathfrak{g} = \mathrm{Lie}(G)$ acts faithfully on $\mathbb{C}[X]$, its Serre exponents determine the boundary Cartan matrix by Theorem 6.2. Theorem 7.1 and Corollary 7.5 give a normal crystal on the theta basis of the coordinate ring. The Chevalley derivations preserve the degree, so perfectness restricts to the theta basis of each R_λ. Since R_λ is irreducible of highest weight λ, this basis has crystal $B(\lambda)$. See also [12, Corollaries 5.12 and 5.15]. For a factor of type A_1, the seed has no mutable vertices. Its theta functions are the monomials in the highest and lowest coordinates of the standard representation. In degree m these monomials form the usual normalized perfect basis of $L(m\varpi_1)$ [11, Remark 10.6]. Tensoring over the simple factors gives the assertion for G. □

Theorem 8.6. *The direct sum $\Xi = \bigoplus_{\lambda\in P_+}\Xi_\lambda$ is a degree-preserving opposite cluster automorphism. It exchanges x_{p_i} with $x_{q_{i^*}}$. If ξ_λ denotes Lusztig's involution on the crystal $B(\lambda)$ of theta indices, then*

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

On mutable g-vectors, the maps ξ_λ induce a single involution ξ_{mut} of $\mathbb{Z}^U$. Let s_μ be the mutable part of the tropical map of Ξ from the lifted seed to its opposite, and let γ_μ be the

normalized mutable reddening map. Then

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

Proof. For each letter i of a reduced word for w_0, adjoin $-\infty$ and $+\infty$ to its ordered list of occurrences. The intervals between successive members of this list index the seed variables of degree ϖ_i. The first interval indexes x_{p_i} and the last indexes x_{q_i*}. Reflection of positions gives a bijection from these intervals to the intervals for the reversed word. For two consecutive intervals of the same letter, reflection reverses their order. For interlacing intervals of different letters, it reverses the order of interlacing and preserves the Cartan factor. The exchange rule thus changes the sign of every entry, including those involving the initial frozen variables by (8.1) and the exchange formula of [2, Definitions 2.2–2.3]. Since the two lifted seeds are mutation-equivalent, this bijection defines an opposite cluster automorphism Φ. It preserves degree because it preserves the letter of each interval.

Reflection exchanges x_{p_i} with x_{q_i*}. Propositions 5.2 and 8.5 give $\Phi F_i = E_{i*}\Phi$ and $\Phi E_i = F_{i*}\Phi$. On an A_1 factor these identities follow by interchanging the two coordinates. For $\lambda = \sum_i m_i\varpi_i$, reflection sends the pinned highest vector $\prod_i x_{p_i}^{m_i}$ to the pinned lowest vector $\prod_i x_{q_i*}^{m_i}$. The highest vector generates R_λ under the F_i, so $\Phi|_{R_\lambda} = \Xi_\lambda$.

By Proposition 5.2, Φ permutes the regular theta basis. On the connected crystal $B(\lambda)$ it exchanges r_i with l_{i*} and sends the highest element to the lowest. Its action is therefore ξ_λ. Proposition 5.4 gives the common mutable map and both identities. □

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 [24, Proposition 1.2] gives

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

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.

8.4. **Lagrangian Grassmannians.** Let $G = \mathrm{Sp}_{2n}(\mathbb{C})$, with $n \geqslant 2$. The Lagrangian Grassmannian $X = \mathrm{LGr}(n, 2n)$ is G/P_n, where P_n is the maximal parabolic subgroup stabilizing a Lagrangian subspace. Its Plücker ring is $\mathcal{R} = \bigoplus_{m\geqslant 0} \mathcal{R}_m$, with $\mathcal{R}_m = L(m\varpi_n)^*$. Denote the section defining the Schubert divisor by b_0 and the minimal lift of the terminal minor of letter i by b_i for $1 \leqslant i \leqslant n$. Here the terminal minor is the minor at the last occurrence of the letter in a reduced word for the open cell. The two sections b_0, b_n are the highest and lowest Plücker coordinates. Indeed, b_0 cuts out the complement of the open cell with multiplicity one. On that cell the terminal n-minor is the opposite Plücker coordinate divided by b_0. The numerator has order zero along the Schubert divisor, so its lift is minimal. The lifted word seed gives $\mathcal{R}$ the structure of an upper cluster algebra with polynomial frozen variables and a homogeneous theta basis [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 8.7. *For every integer $q \geqslant 0$, the G-module of homogeneous derivations of Plücker degree q decomposes as*

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

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. The punctured affine cone is the complement of the zero section in $\mathcal{O}_X(-1)$. Its homogeneous vector fields of degree q are the global sections of the middle term of the twisted Atiyah sequence [1]

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

Such a vector field gives a derivation of $\mathcal{R}$, since its values on the degree-one generators extend across the vertex by normality and the vertex has codimension at least two. Conversely, each homogeneous derivation restricts to such a vector field.

For the tautological Lagrangian subbundle $\mathcal{S}$, the tangent bundle is $T_X \simeq \mathrm{Sym}^2(\mathcal{S}^*)$. Its fiber at the standard flag is an irreducible Levi module of lowest weight $-\theta$. Borel–Weil–Bott gives $H^0(X, T_X(q)) \simeq L(\theta + q\varpi_n)^*$ and $H^1(X, \mathcal{O}_X(q)) = 0$ [4]. Complete reducibility splits the sequence of global sections as G-modules. The first summand consists of multiples of $\mathcal{E}$ by degree-q sections. The vector fields $b_0^q X_\theta$ and $b_n^q X_{-\theta}$ project nontrivially to $T_X(q)$ and have weights $\pm(\theta + q\varpi_n)$, respectively. This identifies the four lines. □

Proposition 8.8. *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 $1, \dots, n$ is $B(m\varpi_n)$. If $m > 0$, no element is highest for all the affine colors.*

Proof. Let $J = \{1, \dots, n-1\}$ and let w_J be the longest element of the parabolic Weyl subgroup generated by the s_j for $j \in J$. The minimal representative $w = w_0 w_J$ sends e_a to $-e_{n+1-a}$ and has reduced word

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

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 and let $H_i^{\mathrm{geom}} = [X_i, Y_i]$. 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.

The mutable exponents in the adapted monomial formulas of [12, Theorem 4.1] are those of R_i and $R_{n-i}^\star$, respectively. Indeed the factor at a mutable index k in those formulas is d_i/d_k, with $d_i = d_{n-i}$, including $d_n = d_0$. It follows that each geometric derivation is its corresponding boundary derivation times a scalar and a frozen monomial.

We show that the frozen monomial is constant. Each b_j has positive Plücker degree. If a frozen monomial b of degree $d > 0$ divided a geometric derivation, its image on $\mathcal{R}_1$ would lie in $b\mathcal{R}_{1-d}$, which has dimension at most one because the boundary derivation preserves $\mathcal{R}$. Yet the long simple-root operators have rank at least 2^{n-1} on $\mathcal{R}_1$. Indeed, the extremal weights $(\pm 1, \dots, \pm 1)$ give 2^{n-1} distinct pairs differing by $2e_n$. For a short simple root $\alpha_i = e_i - e_{i+1}$ choose an extremal weight with signs $(+, -)$ in positions $i, i+1$. Its pairing with $\alpha_i^\vee$ is two, so both operators have rank at least two on that string alone. We have proved

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

for nonzero scalars c_i, c_i'.

Fix i, let $k = n - i$, and write $z = b_k$. In the adapted seed for Y_i, the vector z has lowering length one. Its raising length for X_i is zero when $k \neq i$ and one when $k = i$. In a

finite-dimensional $\mathfrak{sl}_2$-module, the lowering length minus the raising length of a weight vector is its weight. Indeed, an irreducible component of largest highest weight attains both lengths. Thus $H_i^{\mathrm{geom}} z = z$ for $k \neq i$ and $H_i^{\mathrm{geom}} z = 0$ for $k = i$. The rank-one relations give

$$X_i Y_i z = z \quad (k \neq i), \qquad \tfrac{1}{2} X_i^2 Y_i z = X_i z \quad (k = i).$$

Consequently the maximal divided R_i-derivative of the cluster monomial $R_k^\star z$ equals $(c_i c_i')^{-1} z$ or $(c_i c_i')^{-1} R_i z$, respectively. By Theorem 7.1, that derivative is a theta function. Both z and, in the second case, $R_i z$ are cluster monomials, hence themselves theta functions. Linear independence gives $c_i c_i' = 1$. Proposition 6.4 now proves exactness of all finite pairs.

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

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

The exchange entries between 0 and the mutable vertices are not all zero. Indeed, in the seed adapted to Y_n, the nonzero cluster monomial $R_0^\star b_0$ has degree one. The same argument with X_n and b_n applies at n. Since mutation of these entries is invertible and fixes the zero vector, the multipliers of R_0 and $R_n^\star$ are nonconstant mutable monomials in their respective optimizing seeds. Both derivations are homogeneous. As b_0, b_n have degree one and nonconstant cluster variables have positive degree, their degrees q, q' are nonnegative.

The commutators and Lemma 8.7 place R_0 in the span of the two lowest weight vectors in degree q. Evaluation at b_n eliminates the multiple of $\mathcal{E}$, since $R_0 b_n = X_{-\theta} b_n = 0$. Evaluation at b_0 similarly eliminates the Euler term in $R_n^\star$. Hence 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 equality $\langle \varpi_n, \theta^\vee \rangle = 1$ implies $X_{-\theta} X_\theta b_n = b_n$. Hence $R_0(R_n^\star b_n) = cc' b_n$ is the maximal divided derivative of a cluster monomial. As above, Theorem 7.1 gives $cc' = 1$, and $(0, n)$ is exact.

Use the finite pinning in which R_i and $R_{n-i}^\star$ 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 $X_\theta \otimes t^{-1}$ [21, Chapter 7]. Evaluation at $t = c$ gives exactly $R_0 = c X_{-\theta}$ and $R_n^\star = c^{-1} X_\theta$. Their commutator is $H_0 = -H_\theta$, so the canonical central element acts by zero. All generators preserve degree. Theorem 7.1 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 $w_0 = -1$. Its only highest element corresponds to b_0^m, and $R_0(b_0^m) \neq 0$ for $m > 0$. □

By Corollary 7.10 the affine cone over the Lagrangian Grassmannian has a direct cluster automorphism taking color i to $n - i$. In rank two there is one mutable vertex, so the normalized reddening map on mutable degrees is $p \mapsto -p$. The automorphism interchanges b_0 and b_2 and fixes both mutable variables x, y. Their Plücker exchange relation is $xy = b_1^2 + b_0 b_2$. Write $E_i = R_i$ and $F_i = R_{2-i}^\star$. In the following formulas each partial derivative is taken 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	$x\partial_{b_2}$	$y\partial_{b_1}$	$x\partial_{b_0}$

For example, E_0 fixes y, so differentiating the relation gives $E_0 x = b_2$ and hence $E_0(F_0 b_2) = b_2$. Similarly $E_1 y = 2b_1$ gives $E_1^2(F_1 b_1) = 2x$. In the order $0, 1, 2$, the Cartan matrix is $\left(\begin{smallmatrix} 2 & -1 & 0 \\ -2 & 2 & -2 \\ 0 & -1 & 2 \end{smallmatrix}\right)$.

In this example the full weight-lattice criterion of Proposition 7.8 fails after every frozen extension preserving the exact pairs. After adjoining frozen variables $z_1, \ldots, z_s$, the exchange

relation has the form $xy = P(z)b_1^2 + Q(z)b_0b_2$ 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 $E_0x = Qb_2$ and $E_1^2y = 2Px$. 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. Their classes satisfy $[q_0] + 2[q_1] + [q_2] = 0$, so every frozen extension preserving these exact pairs leaves the three classes dependent.

9. 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].

9.1. **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 [20]. 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 9.1. *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 [20, 33]. 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 $1+x^{-B_{k,V}(t)}z^{e_kC_t^{\mathsf{T}}}$ 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 [20] 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 9.2. *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 9.1 identifies his canonical polynomials in every seed with the principal-coefficient polynomials defined above. This proves seed independence. Since $0 \in N_{t,g}$, both support values are nonnegative. Integrality and invariance under frozen factors follow from the definition. When $\widetilde{B}$ has full column rank, 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.

9.2. **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 9.3. *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 9.4. *Maximal raising satisfies*

$$\begin{aligned} \mathsf{F}(P_i g, P_i h) &= \mathsf{F}(g, h) - A_i(g; h) - A_i(h; g), \\ \mathsf{F}(P_i g, h) &= \mathsf{F}(g, h) - A_i(g; h) + B_i(g; h), \\ 0 \leqslant B_i(g; h) &\leqslant g_i \mathsf{F}(M_i, \vartheta_h), \end{aligned}$$

$$g_i = \operatorname{ord}_{x_i}(\vartheta_g) + \frac{1}{d_i}\mathsf{F}(M_i, \vartheta_g).$$

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.

9.3. **Adjunction for the valuation pairing.** 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 [36, Theorem 5.2]. After specialization, Proposition 9.3 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 and then applying theta reciprocity in the seed with J frozen gives

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

For any real direction $n \in N^\bullet_{\mathbb{R}} = N^\bullet \otimes_{\mathbb{Z}} \mathbb{R}$, we use $\mathcal{K}(g,n) = \operatorname{val}_n(\vartheta_g)$, the minimum of the linear form n on the Laurent support of ϑ_g. 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 9.5. *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

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

Reciprocity at $(-m, q)$ gives $W_i(n) = \max_{\alpha \in A_q} \sum_{k \in J} \frac{d_k m_k}{\ell} \alpha_k$. Each summand is nonnegative, and $m_k > 0$ exactly for $k \in J$. The maximum vanishes exactly when every contributing degree has $\alpha_J = 0$. Freezing deletes precisely the degrees involving an index in J. To see this for the chiral dual, apply Qin's freezing theorem in the principal realization of Lemma 3.1. Its injective map sends each dual bend exponent to the corresponding principal bend exponent. The same map respects the two data after the directions in J are removed. Thus a degree survives in the realization exactly when it survives in the dual expansion. Distinct dual degrees give distinct exponents, which proves the equivalence. □

Theorem 9.6. *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 $\Omega_i \cap v^\perp$.

It suffices to test the identity on the variables of the chosen seed and one additional cluster monomial, all of whose exponents are nonnegative. This monomial is independent of n.

Proof. First take $n \in \Omega_i \cap N^\bullet$. By Lemma 9.5, $\vartheta^{\vee,0}_{Q_i n} = \vartheta^\vee_{Q_i n}$. Substituting in (9.1) and applying theta reciprocity gives $\mathcal{K}(P_i g, n) = \mathcal{K}(g, Q_i n)$ for every regular g. The formula for W_i identifies Ω_i with the rational polyhedral cone in the statement. By positive homogeneity the identity holds at every rational point of this cone. For each g the two sides are continuous functions of n, since the theta functions have finite Laurent support. Density of the rational points gives the identity on all of Ω_i.

Conversely, let n, q be real directions and test the identity on the seed variables. Maximal raising fixes x_a for $a \neq i$ and sends x_i to x^m. Thus $q_a = n_a$ for $a \neq i$ and $q_i = \langle m, n \rangle$, which forces $q = Q_i n$.

Choose a normalized reddening sequence from the optimizing seed. After relabeling, the mutable cluster variables I_k of the terminal seed have degrees $-e_k + c_k$ in the optimizing seed, with c_k supported on F. Since I_k is regular at each original frozen divisor, its pointing exponent has nonnegative frozen coordinates. Thus $c_k \geqslant 0$. Let $c = \sum_{k \in J} m_k c_k$, $f = c + e_i$, and $\gamma = -m + f$. Frozen translation gives

$$\vartheta_\gamma = x^f \vartheta_{-m} = x_i \prod_{k \in J} I_k^{m_k}.$$

In particular, ϑ_γ is a cluster monomial with nonnegative exponents. In the seed with J frozen, γ is supported on frozen indices, so $\vartheta^0_\gamma = x^\gamma$. The highest nonzero divided power therefore gives $\vartheta_{P_i \gamma} = x^{\gamma + \gamma_i v}$. Its valuation at n is $\langle \gamma, Q_i n \rangle$. Frozen translation therefore gives

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

Equivalently, for $q = Q_i n$ the difference is the sum of the nonnegative quantities $m_k(\langle -e_k + c_k, q \rangle - \mathrm{val}_q(I_k))$, $k \in J$. Thus this one monomial detects the failure of adjunction at every real $n \notin \Omega_i$. Finally $Q_i^2 = Q_i$ and $\langle m, Q_i n \rangle = \langle m, n \rangle$, so $W_i(Q_i n) = W_i(n)$. Since Q_i has image $v^\perp$, it restricts to the claimed retraction.

□

Already in rank one the cone can be proper. 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, for integral n this is the minimum over the frozen dual support minus the minimum over the full dual support, both evaluated at g. For such n the same formula extends Δ_i to a homogeneous piecewise-linear function of g. Positive homogeneity and continuity in n give the inequality for every real n. The test monomial in the proof realizes $\Delta_i(\gamma; n) = W_i(n)$, so a single $g = \gamma$ detects the common zero set Ω_i.

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 9.6 therefore gives, for an integral target h', the criterion

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

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.

10. Maximal operators and string coordinates

10.1. **String data and reconstruction.** Assume mutable reddening and fix a word $\mathbf{i} = (i_1, \ldots, i_r)$ in sink-reachable colors. Starting at $g^{(0)} = g \in \mathcal{B}_\theta^+$, successively raise to the top of each string,

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

We call $\mathbf{s_i}(g) = (n_1, \ldots, n_r)$ the *string datum* of g for this word. The pair $\Psi_{\mathbf{i}}(g) = (\mathbf{s_i}(g), g^{(r)})$ determines g. Applying $l_{i_r}^{n_r}, \ldots, l_{i_1}^{n_1}$ to $g^{(r)}$, in that order, recovers g. Corollary 3.8 gives the corresponding identity for theta functions,

$$R_{i_r}^{(n_r)} \cdots R_{i_1}^{(n_1)} \vartheta_g = \vartheta_{g^{(r)}}.$$

Each string coordinate is also a valuation. In a sink-optimizing seed for i, the direction $-e_i^*$ defines the valuation $-\deg_{R_i}$ on nonzero elements of $\mathcal{U}^+$, since $R_i = M_i \partial_{x_i}$. It extends to the function field by taking differences on quotients and hence defines a dual tropical point n_i^- independent of the optimizing seed. Thus $n_k = -\mathcal{K}(g^{(k-1)}, n_{i_k}^-)$. Whenever the successive adjoints of Theorem 9.6 are defined, they express n_k as the negative of a single pairing with the original g. For a categorical counterpart, see [11, Appendix A.2].

Proposition 10.1. *Fix a mutation sequence to an optimizing seed for each color in* $\mathbf{i}$*. The resulting map* $\Psi_{\mathbf{i}}$ *extends to a homogeneous piecewise integral-linear embedding* $\mathbb{R}^V \longrightarrow \mathbb{R}^r \times \mathbb{R}^V$ *with a piecewise integral-linear inverse on its image. Under mutable reddening, the image of the real regular cone* $\mathcal{C}^+$ *is a finite union of rational polyhedral cones. For each rational polyhedral cone* $K \subseteq \mathcal{C}^+$ *on which this extension is linear, it maps the lattice points of* K *bijectively onto those of* $\Psi_{\mathbf{i}}(K)$*. They form a finitely generated saturated monoid in the image.*

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

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

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

The regular cone is rational polyhedral by Corollary 3.2. Intersecting it with the finitely many cones of linearity gives the asserted finite union. For any rational polyhedral cone K as in the statement, $\Psi_{\mathbf{i}}(K)$ is rational polyhedral, and the inverse takes its lattice points back to lattice points of K. The lattice points of the image form a saturated monoid, finitely generated by Gordan's lemma. □

10.2. **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 10.2. *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 seed pattern with J frozen. Moreover, the chamber for t is contained in a single reachable chamber of that pattern [36, Lemmas 5.3 and 5.6, proof of Proposition 5.5, and Theorem 5.11]. The corresponding seed s works for every monomial of t. 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 10.3. Let A be a characteristic-zero algebra with regular theta basis, stable under the boundary derivations for a word $\mathbf{i} = (i_1, \dots, i_r)$. For a $\Bbbk$-algebra homomorphism $\epsilon : A \to \Bbbk$ and indeterminates $z_1, \dots, z_r$, define

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

Suppose the variables of a seed t belong to A and satisfy $\Phi^*_{\mathbf{i}}(x_v) = z^{\mathbf{d}_v}$ with $\mathbf{d}_v \in \mathbb{N}^r$. Let M be the matrix with vth row $\mathbf{d}_v$. Suppose $-B_{u,V}(t)\mathsf{M} \leqslant_{\mathrm{lex}} 0$ for every mutable u, comparing the first coordinate first. For every $\vartheta_g \in A$ with $\epsilon(\vartheta_{g^{(r)}}) \neq 0$, one then has $\mathbf{s}_{\mathbf{i}}(g) = g(t)\mathsf{M}$. Indeed, successive maximal divided derivatives give the leading exponent $\mathbf{s}_{\mathbf{i}}(g)$ with coefficient $\epsilon(\vartheta_{g^{(r)}})$. In the pointed expansion, on the other hand, each exchange monomial has nonpositive lexicographic exponent after substitution. 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 $g(t)\mathsf{M}$. Any integral right inverse of M thus recovers $g(t)$ from the string, in particular when M is square and unimodular. For the unipotent factorization charts, see [11, Appendix A.3].

The same tropical maps determine the Cartan entries and the exact partners. Write Φ_i and $\Phi^\star_k$ for tropical transport to a sink-optimizing seed for i and a source-optimizing seed for k, respectively. Let m_i and n_k be the exponents of the corresponding boundary monomials. For distinct sink colors i, j, the maximal x_i-exponent of M_j gives $c_{ij} = -(\Phi_i \Phi_j^{-1} m_j)_i$. By Proposition 4.2, (i, k) is exact if and only if $\Phi_i(\Phi^\star_k)^{-1} n_k = e_i + e_k - m_i$.

For an exact pair (i, k), the degrees of the seed variables under R_i and $R^\star_k$ are e_i and e_k in their respective optimizing seeds. To compute their degrees in the reference seed, transport each vector using the tropical exchange relation

$$q'_u = \max\left\{\sum_a [b_{ua}]_+ q_a,\ \sum_a [-b_{ua}]_+ q_a\right\} - q_u, \qquad q'_a = q_a \quad (a \neq u).$$

Indeed, positivity of the exchange relation and of the cluster variables in their optimizing seed prevents cancellation of the terms of highest degree. Let α_i and β_i denote the resulting vectors for the sink and the source derivation, respectively. Their difference $\beta_i - \alpha_i$ gives the coroot weights of the reference seed variables. Indeed, if a vector has weight w for $\mathfrak{sl}_2$ and p is the largest highest weight of a simple summand in which its component is nonzero, then its raising and lowering lengths are $(p - w)/2$ and $(p + w)/2$, respectively.

11. Pointed biperfect bases and transitions

Assume mutable reddening. Which changes of the regular theta basis preserve its crystal operators? For the bases considered here, each basis vector has a lift to principal coefficients with a prescribed leading term in every seed. The dominance order controls which lower terms can survive specialization. Perfectness then forces their coefficients to agree along crystal strings. We describe the remaining parameters by the coefficient graph of [11, Section 9 and Appendix B], and determine which parameters also preserve linearity over the polynomial ring in the frozen variables.

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.

11.1. **Dominance after coefficient specialization.** Let $L = \mathbb{Z}^V$ be the exponent lattice of a reference seed. Choose its principal extension with auxiliary frozen variables z_k, $k \in U$, and auxiliary exponent lattice $K = \mathbb{Z}^U$. If the original extended exchange matrix has full column rank, we may take $K = 0$ and work in the original seed. In either case the exponent lattice is $\widehat{L} = L \oplus K$. Let sp denote specialization of every auxiliary variable to one. Frozen translation gives

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

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 [35, Proposition 3.4.8], so the common lower set is finite by [35, Proposition 3.4.6(2)].

Definition 11.1. 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 11.2. *The relation $\lhd$ is a strict partial order with finite lower sets. For a fixed global coefficient extension, it is independent of the choice of initial seed.*

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

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

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

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

Every such family is a basis. The transition matrix and its inverse are unitriangular for $\lhd$, and each basis element has a finite expansion in the other basis.

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

Proof. Expand 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 [35, 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 11.2. After specialization, collecting the coefficients of each ϑ_h gives the required formula. Since $b_g \in A$, independence of the localized theta functions makes all coefficients outside $\mathcal{B}$ vanish.

Conversely, choose for each $h \in E_g$ a witness $c(g,h)$ 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 11.4. *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 $|\widehat{D}_g| - |P_g|$. Moreover, there is a finite set of seeds depending only on g such that an element of the localized upper algebra of the extension is compatibly pointed at $(g,0)$ if it is pointed at the corresponding degrees in these seeds.*

Proof. Specialization adds the coefficients over each fiber of $\widehat{D}_g \to P_g$. Prescribing f therefore imposes $|P_g|$ independent linear equations on $|\widehat{D}_g|$ coefficients, with a solution for every f in the indicated affine space.

For the second assertion 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. If the 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 [35, Lemma 3.4.9]. Such an expansion is pointed in every seed. □

11.2. **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 $g \neq 0$, $h \in E_g$, the prescribed weights agree, and $\rho_a(h) \leqslant \rho_a(g)$ for every $a \in \mathcal{I}$. 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 $0 \in \mathcal{B}$ corresponds to the unit.

Proposition 11.5. *Let (b_g) be a basis as in Theorem 11.3. It satisfies $b_0 = 1$, each b_g has the prescribed weights of ϑ_g, and the basis is normalized perfect with the prescribed crystal maps if and only if its off-diagonal coefficients vanish on inadmissible pairs and satisfy*

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

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 and $b_0 = 1$ force the remaining vanishing conditions. Conversely, these conditions preserve the unit, the weights, and all the filtrations under the change of basis. The finite inverse series shows that the inverse also preserves them. 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 (11.1), 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 11.6. *The bases with compatible pointed lifts that fix the unit, 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 (11.1) 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.

Every localized cluster monomial in A belongs to every basis with compatible pointed lifts. Indeed, lift the monomial in a seed where it is a monomial, then translate its auxiliary degree to zero in the reference seed using an auxiliary Laurent monomial. By [35, Lemma 3.4.12], this is the unique compatibly pointed lift of that degree.

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 [35, Remark 5.1.4].

Corollary 11.7. *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 11.6, 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.

Declaration of generative AI and AI-assisted technologies

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

References

[1] M. F. Atiyah, *Complex analytic connections in fibre bundles*, Trans. Amer. Math. Soc. **85** (1957), 181–207.
[2] A. Berenstein, S. Fomin, and A. Zelevinsky, *Cluster algebras III: Upper bounds and double Bruhat cells*, Duke Math. J. **126** (2005), 1–52.
[3] A. Berenstein and D. Kazhdan, *Geometric and unipotent crystals II: From unipotent bicrystals to crystal bases*, in *Quantum Groups*, Contemp. Math. **433** (2007), 13–88.
[4] R. Bott, *Homogeneous vector bundles*, Ann. of Math. (2) **66** (1957), 203–248.
[5] 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] M. Gross, P. Hacking, S. Keel, and M. Kontsevich, *Canonical bases for cluster algebras*, J. Amer. Math. Soc. **31** (2018), 497–608.
[21] V. G. Kac, *Infinite-Dimensional Lie Algebras*, 3rd ed., Cambridge University Press, 1990.
[22] V. G. Kac, *Simple irreducible graded Lie algebras of finite growth*, Math. USSR-Izv. **2** (1968), 1271–1311.
[23] 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.
[24] 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.
[25] M. Kashiwara, *On crystal bases of the q-analogue of universal enveloping algebras*, Duke Math. J. **63** (1991), 465–516.
[26] M. Kashiwara, *On crystal bases*, in *Representations of Groups* (Banff, 1994), CMS Conf. Proc. **16** (1995), 155–197.
[27] 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.
[28] Y. Kimura, F. Qin, and Q. Wei, *Twist automorphisms and Poisson structures*, SIGMA **19** (2023), 105, 39 pp.
[29] C. Lenart, *On the combinatorics of crystal graphs, I. Lusztig's involution*, Adv. Math. **211** (2007), 204–243.
[30] G. Lusztig, *Canonical bases arising from quantized enveloping algebras*, J. Amer. Math. Soc. **3** (1990), 447–498.
[31] S. Morier-Genoud, *Geometric lifting of the canonical basis and semitoric degenerations of Richardson varieties*, Trans. Amer. Math. Soc. **360** (2008), 215–235.
[32] G. Muller, $\mathcal{A} = \mathcal{U}$ *for locally acyclic cluster algebras*, SIGMA **10** (2014), 094, 8 pp.
[33] T. Nakanishi and A. Zelevinsky, *On tropical dualities in cluster algebras*, in *Algebraic Groups and Quantum Groups*, Contemp. Math. **565** (2012), 217–226.
[34] F. Qin, *Dual canonical bases and quantum cluster algebras*, arXiv:2003.13674v3 (2023).
[35] F. Qin, *Bases for upper cluster algebras and tropical points*, J. Eur. Math. Soc. **26** (2024), 1255–1312.
[36] F. Qin, *Applications of the freezing operators on cluster algebras*, J. Algebra Appl. **25** (2026), 2550357, 30 pp.
[37] 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