A candidate proof of the fixed-Cartan root property for c-vectors

Jason Turner

18 November 2026

Abstract

We present a candidate proof that every c-vector of a skew-symmetrizable cluster pattern with principal coefficients is a root of the Kac–Moody root system associated with the Cartan companion of the initial exchange matrix. The argument constructs finite directed-girth covers and compares their classical scattering data with those of auxiliary seeds carrying marked vertices. This is a working draft; its technical steps and cited identifications require independent specialist verification. No priority claim is made.

hexagon:2611.00606v1

1. Introduction and conventions

Let $B = (b_{ij})$ be an integral skew-symmetrizable $r \times r$ matrix: positive integers d_i exist for which $D = \operatorname{diag}(d_i)$ satisfies $DB = -B^T D$. Its Cartan companion $A = A(B)$ has diagonal entries 2 and off-diagonal entries $a_{ij} = -|b_{ij}|$. Let $\Delta(A)$ be the ordinary root system of the associated symmetrizable Kac–Moody algebra, identified with a subset of $\mathbb{Z}^r$ using its simple roots. Both real and imaginary roots are included. The initial c-matrix is the identity; later c-matrices belong to the principal-coefficient cluster pattern obtained by mutation from B.

Candidate assertion. For every such B and every vertex of its cluster pattern, each column c of its c-matrix belongs to $\Delta(A(B))$. The Cartan matrix is fixed at the initial seed; it is not the Cartan companion of a mutated seed or a transposed Cartan matrix.

Nakanishi and Stella formulated this statement as Conjecture 2.3 in their 2014 paper, explicitly attributing the conjecture to joint work with Andrei Zelevinsky [NS, §2.3, p. 8]. Their finite-type result is Theorem 1.3; Section 2.5 also presents previously known acyclic skew-symmetrizable and general skew-symmetric cases, Theorems 2.15 and 2.16, respectively. The analogous statement for d-vectors fails in general [NS, §2.3]. Searches for this draft have not located a proof or disproof of the identical general assertion; this is not a determination of the present status or a claim of novelty.

The positive cone $\mathbb{Z}^r_{\geq 0}$ consists of nonnegative integer vectors, $h(n)$ denotes their coordinate sum, and $\operatorname{prim}(n)$ is the primitive lattice vector on $\mathbb{R}_{>0}n$. Put $\Omega = DB$. An arrow $i \to j$ has $b_{ij} < 0$. For θ in the real dual of the root lattice, $\theta(n)$ denotes evaluation. The root assertion is not an assertion that the vector is a Schur root.

For completeness, we specify the c-matrix convention. Write $[x]_+ = \max(x, 0)$. From an initial pair (B, I_r), mutate at k by

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

$$c'_{ij} = \begin{cases} -c_{ij}, & j = k, \\ c_{ij} + c_{ik}[b_{kj}]_+ + [-c_{ik}]_+ b_{kj}, & j \neq k. \end{cases}$$

Here both right sides use the pre-mutation matrices. The second formula is the componentwise tropical coefficient mutation; no prior choice of a sign for a column is needed to define it. Sign coherence, known for skew-symmetrizable cluster patterns, says each column has one weak sign. It is used only in the endpoint that relates a c-column to a positive wall normal. Unimodularity follows by induction from the equivalent sign-coherent column mutation: at each step one column is negated and the others are changed by integral multiples of it.

For the root arguments, a generalized Cartan matrix A has diagonal entries 2, off-diagonal entries nonpositive integers, and symmetric zero pattern. It is symmetrizable if some positive diagonal matrix makes it symmetric. Its simple reflection is $s_i(x) = x - (Ax)_i e_i$. A positive root has nonnegative simple-root coordinates and connected support. The fundamental imaginary region consists of nonzero nonnegative integral vectors with connected support and all Cartan row evaluations nonpositive. Every vector in that region is a positive imaginary root, and every positive imaginary root is Weyl-conjugate to one in that region; see [Kac80, §1.3] or [Kac, Chapter 5]. In particular, an integral primitive point on an imaginary-root ray is again a root: the same integral Weyl word sends it to the region. A positive real root is a Weyl translate of a simple root and is primitive because the simple reflections are integral unimodular.

External inputs. The argument uses the following results. For a symmetrizable generalized Cartan matrix, roots are stable under its Weyl group; positive roots have connected support; and positive imaginary roots are the positive Weyl translates of the nonzero nonnegative integral vectors with connected support and nonpositive Cartan row evaluations [Kac80, §1.3, Propositions 1.1–1.2]. For reducible Cartan matrices, apply these statements to the connected component containing the support. Real roots are Weyl translates of simple roots. For an ordinary finite acyclic quiver over a finite field, the dimension vector of an indecomposable representation is a positive root of its underlying symmetric Cartan matrix [Kac80, §2.4, Theorem 1(a)]; absolute indecomposability is not required. Over characteristic zero, the positive part of a symmetrizable Kac–Moody algebra has its Chevalley–Serre presentation [GK, Theorem 2]. For skew-symmetrizable cluster patterns, column sign coherence holds, the principal C- and G-matrices are integral unimodular, and $G_t^T D C_t = D$ [NII, Theorem 3.4 and Proposition 2.3(a)]. The cluster scattering diagram, its incoming walls and cluster facets, and the faithful actions and local-factor identifications used below are the cited results of [NII, Theorems 5.27 and 6.13(a)] and [NIII, Propositions 1.8, 1.9 and 2.19, Lemma 2.11, Definition 3.11 and Theorems 3.10 and 3.12].

Invariant covering projections have a close antecedent in Mandel–Qin [MQ, Appendix A, especially Lemma A.5]. Their Theorem A.7 gives a positive containment under its stated covering, symmetry and lattice hypotheses; it is not an equality of the projected and downstairs diagrams. The construction below proves the required invariant Lie identity directly and does not invoke that containment or identify its lattice conventions with ours. Bridgeland [Br, §11.7] corrects an earlier higher-rank argument; we do not invoke the withdrawn proof of his former Lemma 11.4.

The proposed strategy constructs at each finite height a reference element by projecting invariant scattering data from an ordinary quiver whose short induced subquivers are acyclic. A finite-field Hall argument relates local coefficients to semistable objects. A second comparison uses auxiliary vertices indexed by strongly constrained positive degrees. Uniform induction on height transfers the root-ray property from the reference to the cluster element. The endpoint reads a c-facet coefficient before truncating the action. The projections are distinct: the cover projection is a Lie homomorphism only on invariants; the marked map is defined on the full positive Lie algebra.

2. Finite-dimensional scattering algebra and the input coordinates

Fix integral $D = \mathrm{diag}(d_i) > 0$ with $\Omega = DB$ alternating and $N = \mathbb{Z}^r$. Positive degrees are $n \in \mathbb{Z}^r_{\geq 0} \setminus \{0\}$; $h(n) = \sum_i n_i$. "Generic on $n^\perp$ through H" means $\theta(n) = 0$ and $\theta(a) \neq 0$ for every positive a with $h(a) \leq H$ not parallel to n. At a fixed height H use the finite nilpotent graded Lie algebra with basis X_n up to H, bracket $[X_a, X_b] = \Omega(a,b) X_{a+b}$ (zero beyond H), and its BCH group. All classical Lie algebras here are over $\mathbb{Q}$. We write $\Omega(a,b) = a^T D B b$, use BCH for the Baker–Campbell–Hausdorff group, and let $[X_n]F$ denote the coefficient of X_n in a Lie element F. A degree is mixed if at least two coordinates are nonzero. For any $\theta \in N^*_{\mathbb{R}}$, $g = g_+(\theta) g_0(\theta) g_-(\theta)$ has unique sign factors supported in $\theta > 0$, $\theta = 0$, and $\theta < 0$: determine each next homogeneous degree from the previous ones; at the top unknown degree the BCH expression adds the three coordinates, only one of which can occur. Define $p(n) = \Omega(-, n)$ and $I_n(g) = [X_n] \log g_0(p(n))$. The map $g \mapsto (I_n(g))_n$ is triangular with the identity on every associated homogeneous layer, hence polynomially invertible at finite height.

Within the $p(n)$-neutral Lie algebra projection to the n-ray is a Lie homomorphism to an abelian Lie algebra: if $a + b = kn$ and $\Omega(a, n) = 0$, then $\Omega(a, b) = 0$. Thus for $\theta \in n^\perp$ sufficiently close to $p(n)$ avoiding other walls to this cutoff and keeping all strict earlier signs, $[X_n] \log g_0(\theta) = I_n(g)$. Namely factor the old neutral group again by this small perturbation; along n the positive and negative refined factors project trivially. This is valid at central as well as nonsingular directions. Genericity on a given target ray only excludes the finitely many nonparallel integral degrees under consideration.

Define C_B by the input $H_i = \sum_{k\geq 1} \frac{(-1)^{k-1}}{d_i k^2} X_{ke_i}$ on the i-th simple ray, 0 at all mixed grades. These are the ordinary cluster scattering inputs: Nakanishi III, (1.4), (1.7)–(1.8), Lemma 3.8, Theorem 3.12. In his convention $b_{ij} = \{\delta_i e_i, e_j\}$. Taking M a common multiple of the d_i, $\delta_i = M/d_i$ and $\{\ ,\ \} = \Omega/M$, the Lie map $X_n^{\mathrm{Nak}} \mapsto X_n/M$ sends the standard simple-ray logarithm multiplied by δ_i to the displayed input. If $m = (m_i)$ is the usual g-coordinate, then $\theta_i = m_i/\delta_i = d_i m_i/M$. One never identifies a wall degree with its Hamiltonian monomial after losing radical vectors.

We shall need, for each B and H, a group element R at this cutoff having: (i) every generic $[X_n] \log R_0(\theta) \neq 0$ has $\mathrm{prim}(n)$ an $A(B)$-root; (ii) its inputs on every simple ray equal the complete simple-ray inputs of C_B; (iii) a nonzero mixed $I_n(R)$ requires connected supp n and $n_i \leq \min(P_i, N_i)$ for all i, with $P_i = \sum_j [b_{ij}]_+ n_j$ and $N_i = \sum_j [-b_{ij}]_+ n_j$. We use only (i)–(iii); equality of R and C_B is not assumed. Section 3 constructs it independently of the sought theorem.

3. A directed-girth reference and an acyclic finite-field check

Put $M = \mathrm{lcm}_i\, d_i$, $a_i = M/d_i$. Use a first fiber over i indexed by t mod a_i. For each $i \to j$ (by convention $b_{ij} < 0$), let $g = \gcd(a_i, a_j)$. For $s \in \mathbb{Z}/a_i$ and $t \in \mathbb{Z}/a_j$ satisfying $s \equiv t \pmod g$, put $m = |b_{ij}|g/a_i$ arrows $(i,s) \to (j,t)$, and otherwise none. Integrality of m is the coprime consequence of $a_j b_{ij} = -a_i b_{ji}$. For each fixed target in j-fiber the i-column-margin is $|b_{ij}|$, and the reverse margin is $|b_{ji}|$. Take $L > \max(H, 2)$ layers mod L and replace each directed arrow $u \to v$ by $(u,l) \to (v, l+1)$ with its multiplicity for every l. This preserves those one-sided margins, has no internal fiber arrows and has no directed cycles of length $\leq H$. Translation of all first coordinates and independent common layer translation generate a finite symmetry group transitive on each fiber E_i of size K/d_i, $K = LM$. The quantifiers are $\forall H\, \exists Q_H$. Nothing requires this static cover to commute with any cluster mutation or to be independent of H; it need not be a global cluster unfolding. Denote the ordinary skew matrix upstairs by Q, entries positive for arrows $j \to i$, and let π sum coordinates in each fiber. Ordinary height is preserved exactly.

Let C_Q be the ordinary classical canonical group element. By uniqueness of its inputs it is invariant under the finite symmetry. On unnormalized sums over degree orbits in the **invariant** Lie algebra define $\Phi(\sum_{v\in O} X_v) = |O| X_{\pi v}/K$. For two independent uniform orbit members v,w the expected coordinates over color i are $(\pi v)_i/|E_i|$, and the signed total ij-incidence is $|E_j| b_{ij}$. Thus the expected $Q(v,w)$ equals $\Omega(\pi v, \pi w)/K$. Every summand in their double bracket projects to the same grade, proving $\Phi[U,V] = [\Phi U, \Phi V]$. This assertion neither applies to individual noninvariant monomials nor carries over to the quantum commutator. For $\widetilde{\theta} = \pi^* \theta$ all sign factors are invariant by uniqueness, and Φ preserves their signs or sends terms to zero. Hence for $R = \Phi(C_Q)$,

$$[X_n] \log R_0(\theta) = \frac{1}{K} \sum_{\pi v = n} [X_v] \log (C_Q)_0(\widetilde{\theta}).$$

A nonzero sum only forces one contributing exact upstairs term, not an individual generic upstairs wall or the upstairs root status of its grade. At any contributing v put $J = \mathrm{supp}\, v$. Quotient by degrees involving a vertex outside J. This quotient sends C_Q to C_{Q_J}: the weight $p_Q(w)$ restricts to $p_{Q_J}(w)$ for each surviving w, so its inputs and signed factors are the correct ones. The same naturality preserves the desired v coefficient at the **given** $\widetilde{\theta}|_J$; this weight can be singular. J has at most H vertices so the induced quiver is acyclic.

Write m_{ij} for the multiplicity of the arrow $i \to j$. Here is an elementary check that a nonzero classical coefficient in this acyclic case gives an actual semistable object of dimension v at the very same weight, even if the skew matrix is singular. In the positive associative quantum torus over $\mathbb{Q}(t)$ let $X_a * X_b = t^{Q(a,b)} X_{a+b}$ and $q = t^2$. Give each simple the logarithmic input $L_i = \sum_{k\geq 1} \frac{(-1)^{k-1} t^k}{k(t^{2k}-1)} X_{ke_i}$. Sign-factorization and I-coordinate triangular uniqueness work identically with the commutator. The all-objects twisted mass of finite-dimensional ordinary F_q representations of the given quiver is

$$S = \sum_z \frac{t^{\sum_i z_i + \sum_{i\to j} m_{ij} z_i z_j}}{\prod_i \prod_{1\leq k\leq z_i} (q^k - 1)} X_z.$$

For each prime q it counts objects by $t^{e(z,z)}/|\operatorname{Aut}|$ with e the arrow Euler form. To see the formula, write $E(z) = \sum_{i\to j} m_{ij} z_i z_j$. There are $q^{E(z)}$ choices of arrow matrices, while $|\operatorname{GL}_{z_i}(q)| = q^{z_i(z_i-1)/2} \prod_{k=1}^{z_i} (q^k - 1)$. Multiplication by $t^{e(z,z)}$ leaves numerator exponent $\sum_i z_i + E(z)$, as displayed. Write $f_i = \sum_{k\geq 0} \frac{t^k X_{ke_i}}{\prod_{j=1}^{k} (q^j - 1)}$; its ordinary logarithm is L_i by $f_i(qX_i) = (1 + tX_i) f_i(X_i)$. Because Q_J is acyclic, S is the product of f_i in any reverse

topological order (sinks first): each $i \to j$ cross term acquires $t^{Q(e_j,e_i)z_iz_j} = t^{m_{ij}z_iz_j}$. Its input is zero in every mixed degree w. To see this quotient to supp w. If edgeless, the logs add and have only pure terms. Otherwise pick a nonisolated source i in a component with arrows, arrange its f_i at the right end. Writing $s(\alpha), t(\alpha)$ for the tail and head of an arrow and counting parallel arrows separately, one has

$$p(w)_i = \sum_{\alpha:\, t(\alpha)=i} w_{s(\alpha)} - \sum_{\alpha:\, s(\alpha)=i} w_{t(\alpha)} < 0.$$

Indeed the first sum is zero at this source of the induced support, whereas the second is positive. Its entire factor is in $G_-(p(w))$ and disappears from the neutral factor; all remaining terms have no i coordinate. Its pure inputs are the simple L_i. By triangular uniqueness S is exactly the acyclic canonical quantum element.

For classical specialization set $t = \exp(\hbar/2)$, $Y_n = X_n/\hbar$. The bracket of Y_a,Y_b has coefficient $2\sinh(\hbar Q(a,b)/2)/\hbar$, regular and reducing to $Q(a,b)$, and L_i in the Y-basis reduces to the ordinary dilogarithm H_i with $d_i = 1$. Regular BCH and the triangular recursion imply the local canonical quantum and classical *log* coefficients match under this specialization in the Y-basis. Thus a nonzero classical v-term implies a nonzero rational quantum X_v term; it is nonzero at $t = \sqrt{q}$ for all sufficiently large primes q.

Compare [Re, Proposition 2.5 and Lemma 4.7] for integral weights. The middle factor of S at any θ is the twisted mass of the zero-slope semistable representations. Indeed the finite HN filtration at slope $\theta(w)/h(w)$ gives a unique subobject flag of strictly decreasing semistable slopes. The finite Hall multiplication with the subobject class on the left counts it. The extension groupoid with fixed endpoint isomorphism classes has mass $|\operatorname{Ext}^1(B,A)|/(|\operatorname{Hom}(B,A)|\,|\operatorname{Aut} A|\,|\operatorname{Aut} B|)$. Multiplying by the integration weights relative to the two endpoint masses gives $t^{e(a,b)+e(b,a)}q^{-e(b,a)}$; this equals $t^{Q(a,b)}$ because the groupoid mass of extensions of the quotient B by the subobject A contributes $|\operatorname{Ext}^1(B,A)|/|\operatorname{Hom}(B,A)| = q^{-e(b,a)}$ and $e(a,b) - e(b,a) = Q(a,b)$. Across all slopes this gives $S_+S_0S_-$, hence the three factors by uniqueness. A nonzero degree-v term in log S_0 implies some semistable object of exactly dimension v: any nonzero product term is a list of semistable objects with total v, whose direct sum remains semistable: for any subobject, its intersection with the first summand and its image in the remaining summands have nonpositive weight, so the subobject also has nonpositive weight. This is existence only, not positivity of logs.

(i) At a generic target θ on $n^\perp$ take any indecomposable direct summand of this object, with dimension w. It and its complement are subobjects of a zero-slope semistable, so $\theta(\pi w) = 0$. Thus πw is positive on the n-ray. Ordinary Kac 1980, §2.4 Theorem 1(a) gives that w is an ordinary positive root already over the finite field. Embed it into the full ordinary Q Cartan by zero padding. The blockwise commuting simple reflections $\widehat{S}_i$ satisfy $\pi\widehat{S}_i = s_i\pi$ by the signed-column margins and no intrafiber edges. Whenever a positive projected vector has more than one supported color, reflect at any positive A-row; the upstairs root retains a positive coordinate outside that fiber, hence remains positive, while the projected height decreases. If supported on one fiber, upstairs connected support and the absence of intrafiber edges force a single simple. Otherwise the terminal vector is positive, connected (image of a connected root support with no fictitious interblock edge) and has A-row evaluations ≤ 0; it is in the usual fundamental imaginary region. Undo the reflections. This proves πw is a root; $\operatorname{prim}(\pi w)$ is a root too by the same Weyl reduction for integral vectors on an imaginary ray, and real roots are primitive. No claim is made that the original contributing singular term v itself is a root.

(ii) Quotient the upstairs canonical algebra to E_i. That full subquiver is edgeless, so its canonical log is just the sum of unit simple dilogs. Φ multiplies by $|E_i|/K = 1/d_i$ along the i-axis, giving the correct complete simple input (or compute directly at the attractor after the quotient).

(iii) If $I_n(R) \neq 0$ at a mixed n, choose generic $\theta \in n^\perp$ sufficiently near $p(n)$ preserving all strict signs. The preceding exact local construction provides an actual semistable ordinary representation of dimension v and $\pi v = n$. If $\theta_i > 0$ there can be no subobject supported at a single i-sheet, so the outgoing aggregate at that sheet is injective. Summing the dimension inequalities over its fiber (set omitted coordinates to zero) gives $n_i \le N_i$. If $\theta_i < 0$ no quotient at such a sheet is possible, so the incoming aggregate surjects, giving $n_i \le P_i$. The identity $p(n)_i = d_i(P_i - N_i)$ gives the sign of $P_i - N_i$, hence in each case this gives the smaller bound. At equality choose θ_i nonzero, possible since n is mixed. Root-ray property (i) guarantees connected support. This proves all three claimed reference properties. In fact a disconnected support can also be discarded by common sign-zero summand reasoning.

4. One scaled marked vertex per mixed strong integer degree

Call a mixed positive integral degree u *strong* if its support is connected and $u_i \le \min(P_i(u), N_i(u))$ for every old vertex i. This predicate is independent of the height cutoff. Fix the finite list $U = \{u : u \text{ is strong and } h(u) \le H\}$. Introduce one vertex per u, $Te_{a_u} = u$, and keep $Te_i = e_i$ on old labels. D' is D on originals and 1 at marks; $\Omega' = T^*\Omega$, $B' = (D')^{-1}\Omega'$ integral. Let C' be the canonical **classical** element of this genuine auxiliary skew-symmetrizable seed. For $\lambda_u \in \mathbb{Q}$ define the positive-degree Lie homomorphism

$$\Phi_\lambda(X_z) = \left(\prod_{u\in U} \lambda_u^{z_{a_u}}\right) X_{Tz},$$

allowing zero λ with $0^0 = 1$. For every z, $h(Tz) \ge h(z)$; source degrees $> H$ are killed after mapping to target height H, while source degrees $\le H$ with $h(Tz) > H$ are sent to zero. Thus the map descends to these common finite-height quotients; no infinite fibers or negative dimensions occur. Exact pullback gives the Lie identity; it will in general be noninjective. Sign-factor uniqueness gives $(\Phi_\lambda C')_0(\theta) = \Phi_\lambda(C'_0(T^*\theta))$. For $Tz = n$, the source attractor $p'(z) = \Omega'(-, z)$ equals $T^*p(n)$. Put $\chi_\lambda(z) = \prod_{u\in U} \lambda_u^{z_{a_u}}$ and write I'_z for the input coordinate in the enlarged seed. Applying the Lie map to the logarithm of the middle factor gives

$$I_n(\Phi_\lambda C') = \sum_{\substack{z>0\\ Tz=n}} \chi_\lambda(z) I'_z(C').$$

Here $z > 0$ denotes a nonzero degree with nonnegative coordinates. The prescribed local inputs $I'_z(C')$ vanish off all coordinate axes. The surviving original-axis terms give $I_n(C_B)$; a marked axis contributes only at $z = ke_{a_u}$, hence only when $ku = n$. Consequently

$$I_n(\Phi_\lambda C') = I_n(C_B) + \sum_{u\in U,\, k\in\mathbb{Z}_{>0},\, ku=n} \ell_k \lambda_u^k, \qquad \ell_k = \frac{(-1)^{k-1}}{k^2}.$$

Set λ_n at each $n \in U$ in order of height by subtracting the $k \ge 2$ previously fixed terms from $I_n(R) - I_n(C_B)$ and dividing by ℓ_1. The row bounds are homogeneous and positive multiples preserve support: if $u \in U$, $k \ge 1$, and $h(ku) \le H$, then $ku \in U$. Thus outside U no marked axis contributes in any retained degree; Section 3 excludes a mixed reference input there. The prescribed inputs match in every degree through H, including those with zero defect. Thus $\Phi_\lambda C' = R$ and $\Phi_0 C' = C_B$. These signed scalars are weights of formal monomials, not multiplicities of representations or negative vertex dimensions. This uses no polynomial interpolation or large-rank grid.

For any actual positive root z of $A' = A(B')$ meeting a mark, Tz is an **imaginary** $A(B)$ root. To verify, form the integral GCM A^* by replacing each old-to-mark entry $-|(Bu)_i|$ with $(Au)_i$ and its reverse by $d_i(Au)_i$, leaving all other entries unchanged. The row min bound gives $(Au)_i = 2u_i - P_i - N_i \le -|P_i - N_i|$, hence $A^* \le A'$ off diagonal and D' symmetrizes it. The positive Chevalley–Serre algebra for A^* surjects grade-preservingly onto that for A' (the larger target Cartan entry imposes the stronger iterated-adjoint relation). In particular z is an A^*-root. Repeatedly reflect at an old row with strictly positive A^*-evaluation. The marked coordinate persists, so each root stays positive, and the sum of old coordinates falls by a positive integer. Along each such reflection T intertwines with the corresponding ordinary old $A(B)$ reflection by the exact choice of the A^* row. At the terminal root z' the vector Tz' has all A-evaluations nonpositive. Its support is connected: each occupied mark is replaced by a connected true old support, an old-to-mark edge $(Au)_i < 0$ attaches to that support, and a mark-to-mark edge $\Omega(u, v) \ne 0$ requires a real old edge between the two supports. Its positive vector lies in the fundamental imaginary region, and conjugating back proves the assertion. The marks may be linearly dependent; no injectivity is used. Exact graded details appear separately in Section 6.5.

5. Uniform height induction and the genuine c-facet

Induct over H simultaneously for **all** finite ranks and matrices: every generic nonzero degree-n local log term of C_B with $h(n) \le H$ has prim(n) a positive $A(B)$-root. Height 1 is immediate. If $h(n) = H$ and prim(n) is not a root, property (i) gives $[X_n] \log R_0(\theta) = 0$. Suppose the corresponding C_B term were nonzero. Subtracting $\Phi_\lambda - \Phi_0$ on the local logs therefore forces some exact $Tz = n$ with a positive mark to have $[X_z] \log C'_0(T^*\theta) \ne 0$. Although $T^*\theta$ can be singular, in the source box below z every neutral positive charge maps onto the n-ray; their mutual brackets vanish because $\Omega' = T^*\Omega$. The sign-refinement argument now preserves this **exact** z term after a small

generic perturbation in $z^\perp$. Terms outside the z-box cannot affect its coefficient, so the finite induction statement at ordinary height $h(z)$ applies. Each marked image has height at least 2; hence $h(z) < H$. Its primitive part is therefore a source root; by the A^* construction its image lies on an ambient imaginary root ray. Since that image is parallel to n, prim(n) must be a root after all: its integral primitive vector can follow the same Weyl word into the fundamental region because the Weyl matrices are integral unimodular and all positivity/row signs and support are unchanged by positive rescaling. This is the contradiction. This is an induction on total ordinary height, not on rank or on the number of times one mutates. It never asserts a nonprimitive log harmonic must itself be a root.

Endpoint: read the primitive coefficient before truncating the action. Assume first that B is nonsingular. We spell out the two monomial conventions. Let $N^\circ = \bigoplus_i \mathbb{Z}\delta_i e_i$, with $\delta_i = M/d_i$ as in Section 2, and $M^\circ = \mathrm{Hom}(N^\circ, \mathbb{Z})$. Nakanishi's map $p^* : N \to M^\circ$ is determined by $p^*(v)(w) = \{w, v\} = \Omega(w, v)/M$ for $w \in N^\circ$; as a real functional it is our $p(v)/M$. In the basis e_i^*/δ_i of M°, its coordinate column is Bv, since $p^*(v)(\delta_i e_i) = (Bv)_i$. Write y^v for the formal monomial of degree v in the completed positive monoid algebra $\mathbb{Q}[[N_{\geq 0}]]$; x^m denotes the exponent-m monomial in the dual-lattice wall-function algebra, and $\widehat{y}^v = x^{p^*(v)}$. Thus the hat changes the exponent lattice, not the vector v.

For a seed t, let G_t be its principal g-matrix, let $\sigma(G_t)$ be the cone spanned by its columns, and let $\sigma_j(G_t)$ be the face spanned by all but the jth column. A dual-lattice point with usual coordinates $m = (m_i)$ corresponds to $\theta(e_i) = m_i/\delta_i$.

We record the comparison of the two wall languages in the present nonsingular case. Let n be primitive and let a normalized wall function satisfy

$$\log f = \sum_{k \geq 1} b_k x^{kp^*(n)}.$$

The corresponding element in Nakanishi III's completed group is

$$g_{n,f} = \exp\left(\delta(n) \sum_{k \geq 1} \frac{b_k}{k} X_{kn}^{\mathrm{Nak}}\right).$$

Indeed the x-representation sends X_{kn}^{Nak} to the derivation taking x^m to $k\langle n, m\rangle x^{m+kp^*(n)}$. Since $\langle n, p^*(n)\rangle = 0$, the derivation annihilates each $x^{kp^*(n)}$; exponentiation sends x^m to $x^m f^{\delta(n)\langle n,m\rangle}$, the normalized wall action of [NII, Definition 5.12]. Reversing the crossing replaces the group element by its inverse. Nonsingularity makes p^* injective; it identifies the positive degree monoid with $P = p^*(N_{\geq 0})$ and preserves the corresponding completion filtration. The completed x-representation is faithful [NIII, Proposition 1.9]. Thus equal path-ordered wall actions, including consistency around loops, lift to equal group elements. For the initial wall $1 + x^{p^*(e_i)}$, the lifted element is $\Psi[e_i]^{\delta_i}$. Lift the particular consistent cluster diagram of [NII, Theorem 5.27]. Its other walls are outgoing by part (b), and the lifting preserves their supports and incoming or outgoing status. Thus the lifted diagram has exactly the initial incoming elements above. The cited uniqueness of cluster scattering diagrams [NIII, Definition 3.11, Lemma 3.8, Theorems 3.10 and 3.12] identifies its equivalence class with that of the canonical element defined in Section 2.

For a seed t, the facet $\sigma_j(G_t)$ of its g-cone opposite the jth generator has primitive positive normal $n = |c_j|$ (the principal C_t is unimodular) and ordinary wall function $1 + \widehat{y}^n$ by Nakanishi II, Example 5.16(b). By Definition 6.12 and Theorem 6.13(a) this is an **essential wall**: in a representative equivalent to the full consistent diagram it is present and every other wall meets this facet in dimension less than $r - 1$. The relative interior of the facet is open in $n^\perp$. Every lower-dimensional intersection just described is contained in a proper linear subspace of $n^\perp$, since the wall supports are convex polyhedral cones. Also exclude $n^\perp \cap a^\perp$ for every lattice vector $a \in N \setminus \mathbb{Q}n$, including vectors of mixed sign. Each intersection is a proper subspace of $n^\perp$. There are only countably many walls and lattice vectors, and a countable union of proper subspaces cannot cover the nonempty relative interior: each has zero relative Lebesgue measure in positive dimension, and there is no proper nonempty subspace in dimension zero. Choose θ outside this union. Since n is primitive, $\ker(\theta) \cap N = \mathbb{Z}n$. Thus $n^\perp$ is the unique rational hyperplane containing θ, as required by Nakanishi III, Definition 2.8. This full generality implies the finite-height condition of Section 2; no neighborhood common to all cutoffs is required. Thus the intrinsic local element is the actual simple binomial wall automorphism, not an arbitrary member of a collection whose factors might cancel. With X_w^{Nak} denoting Nakanishi's Lie generator, set

$$\Psi[n] = \exp\left(\sum_{k \geq 1} \frac{(-1)^{k-1}}{k^2} X_{kn}^{\mathrm{Nak}}\right).$$

This is Nakanishi III, Definition 1.10, (1.50); the normalization exponent is not included in this definition. The group element for the wall with its positive normal (before choosing a crossing orientation) is $\Psi[n]^{\delta(n)}$, from II Example 5.13(b), equation (5.20), and Definition 5.12 and III Proposition 1.12, §2.3–2.4; at a general point it agrees with the neutral factor by III Lemma 2.11 and Proposition 2.19, equation (2.43). The exponent is the least positive rational δ with $\delta n \in N^\circ = \bigoplus_i \mathbb{Z}(M/d_i)e_i$, hence $\delta(n) = M/\gcd_i(d_i n_i) = M/d_j$ by $G^T DC = D$ and unimodularity of the principal g-matrix G; compare [NII, Proposition 2.3(a)]. Rescaling the generators as in Section 2, namely $X_w^{\mathrm{Nak}} \mapsto X_w/M$, gives

$$[X_n] \log(C_B)_0(\theta) = \delta(n)/M = 1/d_j \neq 0,$$

where $(C_B)_0(\theta)$ is the middle sign factor of the canonical element defined in Section 2. Reversing an oriented crossing inverts it.

For clarity, the normalized x-action of $\Psi[n]^{\delta(n)}$ multiplies x^m by $(1+x^{p^*(n)})^{\langle \delta(n)n,m\rangle}$, whereas its y-action multiplies y^v by $(1+y^n)^{\delta(n)\{n,v\}}$. Here the first exponent is evaluation of the dual vector m at $\delta(n)n \in N^\circ$. These formulas agree on $m = p^*(v)$ for $v \in N_{\geq 0}$, that is, on the monoid $P = p^*(N_{\geq 0})$. We compare the full completed actions first: Nakanishi III, Propositions 1.8 and 1.9 give faithfulness of the y- and x-representations, respectively, under the present nondegeneracy assumption. Equivalently choose $v \geq 0$ with $\Omega(n,v) \neq 0$ (some simple v works by nonsingularity) and read the first nontrivial action on y^v through order $h(n)+h(v)$, or higher. The y^{n+v} term is killed in a module truncated already at $h(n)$; never infer nonvanishing from that insufficient truncation. **Only after** recovering the abstract log coordinate do we reduce the group element to degree $H = h(n)$. The coordinate at n is retained there, as required by the preceding induction. If an identification uses nonsingularity, apply it to the mutable double $\begin{pmatrix} B & I \\ -I & 0 \end{pmatrix}$, symmetrizer $\mathrm{diag}(D,D)$, and mutate only old labels. Its old columns are $(c,0)$ by the principal mutation recurrence; the double is invertible. A root supported on old labels is a root of the fixed original Cartan by the ordinary parabolic embedding/restriction (the positive Lie algebra in such degrees uses only old Chevalley generators and old Serre relations): in the free positive Lie algebra, the ideal generated by a relation containing a new generator has positive new multidegree, so its intersection with the old-only subalgebra is zero. The remaining defining relations are exactly the old Serre relations. Negatives follow by symmetry. This is the entire candidate deduction.

6. Details of the finite-height arguments

This section expands the finite combinatorial parts of the proposed proof. All arguments are at a fixed finite cutoff H. Only the truncations of a fixed canonical element C_B are passed to an inverse limit. The height-dependent reference elements, covers, and marked sets need not be compatible as H varies.

6.1. Triangular factorization and local refinement

For any weight θ, partition the homogeneous Lie algebra as a vector space into the positive, zero and negative θ-subspaces. Each is a Lie subalgebra. The ordered multiplication map from their three BCH groups to the full group is bijective. Indeed, at height k, once terms of height less than k have been chosen, the unknown height-k terms enter the logarithm of the product only as their sum. Each degree belongs to exactly one sign class, so its unknown coefficient is determined uniquely. The same argument is natural for any Lie homomorphism that maps each source homogeneous degree to a single target degree (possibly sending the homogeneous component to zero), preserves the three sign classes under the chosen weights, and descends to the finite quotients. In particular, it applies after regrading by T; equality of ordinary source and target heights is not required.

The same height-by-height calculation shows that the coordinate $I_n(g)$ is the degree-n coefficient of $\log g$ plus a polynomial in coefficients of height less than $h(n)$. The map sending g to its family of I-coordinates is therefore a triangular polynomial bijection. This remains true when Ω is degenerate: degrees are never identified modulo its radical.

For the refinement statement, let $\theta_0 = p(n)$. In the θ_0-neutral subalgebra, projection onto degrees on the n-ray is a Lie homomorphism to an abelian algebra. For if $a+b = kn$ and $\theta_0(a) = 0$, then

$$\Omega(a,b) = k\Omega(a,n) - \Omega(a,a) = 0.$$

Choose a small perturbation $\theta \in n^\perp$ which does not reverse any of the finitely many strict signs at the chosen cutoff. Refactor the old neutral group into positive, neutral and negative θ-parts and absorb the positive and negative

pieces into the corresponding old sign factors. Their projections to the n-ray vanish. Thus the degree-n coefficient of the neutral logarithm is unchanged. Choosing θ outside the finitely many nonparallel hyperplanes gives the required generic weight. Nothing in this argument requires $p(n)$ to be nonzero.

6.2. The directed-girth cover and invariant projection

Put $a_i = M/d_i$. For $i \to j$, let $g = \gcd(a_i, a_j)$, $a_i = gp$ and $a_j = gq$, with p and q coprime. Skew-symmetrizability gives $a_j|b_{ij}| = a_i|b_{ji}|$, so p divides $|b_{ij}|$. The prescribed multiplicity is $|b_{ij}|/p$. Each target sheet has p compatible source labels and hence exactly $|b_{ij}|$ incoming arrows from the i-fiber. Each source sheet has q compatible target labels and hence exactly $|b_{ji}|$ outgoing arrows to the j-fiber. Advancing the layer coordinate by one on every arrow implies that the length of every directed cycle is a positive multiple of L. Therefore every induced subquiver on at most H vertices is acyclic. Simultaneous translation of the first coordinates and simultaneous translation of the layers preserve the arrows and act transitively on each fiber. No claim about compatibility of this cover with mutation is used.

Let $f_i = K/d_i$ be the cardinality of the i-fiber. Let U and V be orbit sums, with uniformly and independently chosen representatives v and w. Their expected coordinates at a sheet of color i are $(\pi v)_i/f_i$ and $(\pi w)_i/f_i$. The signed incidence in the block (i, j) is $f_j b_{ij}$, and hence

$$\mathbb{E}\, Q(v, w) = \sum_{i,j} \frac{(\pi v)_i}{f_i} f_j b_{ij} \frac{(\pi w)_j}{f_j} = \frac{\Omega(\pi v, \pi w)}{K}.$$

The bracket of the two orbit sums counts all ordered pairs. Consequently applying the map with scale $1/K$ to this bracket gives the bracket of the two projected orbit sums. This proves the Lie identity on the invariant subalgebra, even though the two coordinates of an individual representative need not be independent. The scale is essential. Sign factorization is preserved because $\widetilde{\theta}(v) = \theta(\pi v)$, and its uniqueness implies that all three factors of an invariant element are invariant.

For an exact contributing upstairs degree v, quotient the positive algebra by all degrees involving a vertex outside $J = \operatorname{supp}(v)$. This is a Lie quotient since nonnegative degrees cannot lose an outside coordinate under addition. The input at any surviving degree w is computed at $p_Q(w)$ restricted to J, exactly $p_{Q_J}(w)$. By triangular uniqueness the image of the upstairs canonical element is the canonical element of the induced subquiver Q_J. Sign-factorization at the prescribed weight also commutes with the quotient. Here $|J| \le h(v) \le H$, and Q_J is acyclic. In particular the coefficient is retained at the original, possibly singular, weight; it is not replaced by a generic upstairs coefficient.

6.3. The finite-field semistability implication

Write

$$e(a, b) = \sum_i a_i b_i - \sum_{i \to j} m_{ij} a_i b_j.$$

Then $e(a, b) - e(b, a) = Q(a, b)$. Counting all representations of a fixed dimension z over a field with $q = t^2$ elements gives the mass displayed in Section 3: the number of choices of arrow maps is $q^{\sum_{i \to j} m_{ij} z_i z_j}$, and division by the product of the general linear group orders counts isomorphism classes with automorphism weights. A reverse topological product of the simple generating functions gives the same expression; for every arrow $i \to j$ the commutation factor is $t^{m_{ij} z_i z_j}$.

For a mixed w, restrict to its support. If that induced quiver has no arrows, the simple factors commute and its logarithm has only simple-ray terms. Otherwise choose a nonisolated source i and place its simple factor last in a reverse topological order. Since $p(w)_i < 0$, that factor belongs to the negative subgroup, and all remaining factors have zero i-coordinate. Refactoring the latter and absorbing the last factor into the negative subgroup shows that the neutral factor has no coefficient in degree w. Thus the all-object mass has the prescribed mixed-zero and simple-ray inputs. Triangular uniqueness identifies it with the acyclic quantum canonical element.

At any real weight θ, set $\mu(a) = \theta(a)/h(a)$, and call a nonzero representation semistable if every nonzero subrepresentation has slope at most its own. Here the slope of an object means the slope of its dimension vector. For real weights, choose among subobjects of maximal slope one U of maximal dimension. If V has the same maximal slope, the intersection has no larger slope; additivity for intersection and sum shows that $U + V$ also has maximal slope, so maximal dimension forces $V \subseteq U$. A subobject of the quotient with slope at least $\mu(U)$ would pull

back to a subobject contradicting maximal slope or maximal dimension. Recursion on dimension gives a filtration by semistable quotients with strictly decreasing slopes. For uniqueness, let $0 = M_0 \subset M_1 \subset \cdots \subset M_s = M$ be any filtration with semistable quotients of strictly decreasing slopes $\mu_1 > \cdots > \mu_s$. For any nonzero subobject W, the successive quotients of $W \cap M_i$ inject into M_i/M_{i-1}; each nonzero quotient has slope at most μ_i. Their dimension-weighted average is therefore at most μ_1, with equality only if $W \subseteq M_1$. As M_1 itself has slope μ_1, it is exactly the unique maximal subobject of maximal slope, namely U. Induction on M/U proves uniqueness of the entire filtration. Its finite Hall multiplication identity factors the all-object mass in that order. With subobject A of dimension a on the left and quotient B of dimension b on the right, the exponent is seen directly from the two-term complex

$$\bigoplus_x \mathrm{Hom}_{\mathbb{F}_q}(B_x, A_x) \xrightarrow{\partial} \bigoplus_{\alpha: x \to y} \mathrm{Hom}_{\mathbb{F}_q}(B_x, A_y), \qquad (\partial f)_\alpha = A_\alpha f_x - f_y B_\alpha.$$

The kernel consists exactly of representation maps $B \to A$. After splitting an extension at each vertex as vector spaces, its arrow matrices are upper triangular blocks with off-diagonal entries in the right-hand space. An arbitrary tuple of such entries gives an extension because the quiver has no relations; changing the splittings changes that tuple by a coboundary. Hence the cokernel is $\mathrm{Ext}^1(B, A)$. Rank–nullity gives $\dim \mathrm{Hom}(B, A) - \dim \mathrm{Ext}^1(B, A) = e(b, a)$. The extension groupoid with endpoints fixed up to isomorphism has mass $q^{-e(b,a)}/(|\operatorname{Aut} A|\,|\operatorname{Aut} B|)$: with endpoints identified, classes are indexed by $\mathrm{Ext}^1(B, A)$ and each has endpoint-fixing automorphism group the additive group of $\mathrm{Hom}(B, A)$; then divide by both endpoint automorphism orders. Combining this with the integration twist produces the quantum torus factor $t^{e(a,b)-e(b,a)}$. Hence the zero-slope semistable mass is precisely the middle sign factor. This argument applies at singular θ as well.

A nonzero coefficient of $\log S_0$ in degree v implies that one summand in its formal logarithmic expansion is nonzero. The degrees in that product are dimensions of zero-slope semistable objects and add to v; their direct sum is again semistable of dimension v. This conclusion does not require v itself to be the dimension of an indecomposable representation or any positivity of logarithmic coefficients.

Finally, with $t = \exp(\hbar/2)$ and $Y_n = X_n/\hbar$, the commutator of Y_a and Y_b has coefficient

$$\frac{2\sinh(\hbar Q(a, b)/2)}{\hbar},$$

which is regular at $\hbar = 0$ and specializes to $Q(a, b)$. The logarithm of each simple quantum factor, expressed in the Y-basis, has constant term equal to the corresponding classical dilogarithm input. The finite triangular recursions and BCH polynomials commute with specialization, including after neutral factorization. A nonzero classical coefficient therefore comes from a nonzero rational quantum coefficient. For each fixed degree and weight it survives evaluation at $t = \sqrt{q}$ for all sufficiently large primes q, proving existence of the needed semistable finite-field object.

6.4. Projection of ordinary roots

For a related cyclic-automorphism folding statement, see Hubery [Hu, Proposition 15]. The argument here is given directly for the finite cover and the fiber projections used above.

Let w be the dimension vector of an indecomposable summand of the semistable object at a target-generic weight. The summand and its complement are both subobjects, so each has weight zero. Target genericity implies that πw lies on the n-ray. The finite-field direction of Kac's theorem says that w is a positive root of the ordinary Cartan matrix upstairs. It remains a root after zero padding to the full cover by the parabolic root-space embedding.

All simple reflections within an i-fiber commute. Let their product be $\widehat{S}_i$. The column-margin identities give $\pi \widehat{S}_i = s_i \pi$, where s_i is the simple reflection for $A(B)$. If a positive projected root occupies more than one color and its i-th Cartan row evaluation is positive, apply $\widehat{S}_i$. At every intermediate reflection a positive coordinate outside the i-fiber remains, so the reflected upstairs root cannot be negative. The projected height drops by that positive row evaluation. The process terminates. If only one color remains, connectedness of root support and absence of edges inside a fiber force a single simple root. Otherwise the terminal image is positive with connected support and all Cartan row evaluations nonpositive, hence lies in the fundamental imaginary region. The original projected vector is a root by applying the inverse Weyl word.

If the projected root is a nonprimitive multiple, it cannot be real. Apply the same integral Weyl word to the primitive point on its ray. Positive rescaling does not change the sign of the coordinates or row evaluations or the

support, so that point also reaches the fundamental imaginary region. It too is a root. These statements use the fundamental-region description in Kac [Kac80, §1.3] and do not assert that the original singular coefficient degree v is a root.

A finite-dimensional representation decomposes as a finite direct sum of indecomposables by repeatedly splitting decomposable summands; each split decreases their dimensions. At a single vertex the positive free Lie algebra on one Chevalley generator is one-dimensional, since its self-bracket is zero; its only positive root degree is therefore the simple degree.

6.5. Marked comparison and the induction box

For a marked vector u, the two row inequalities give

$$2u_i - P_i - N_i \leq -|P_i - N_i|,$$

that is, $(Au)_i \leq -|(Bu)_i|$. Thus the replacement A^* has off-diagonal entries no larger than A'. In the positive Chevalley–Serre presentations [GK, Theorem 2], the longer Serre relation for A^* is a consequence of the shorter relation for A': apply the same adjoint operator further times. Hence the defining ideal for A^* is contained in that for A', and there is a degree-preserving surjection from the positive algebra for A^* onto that for A'. Nonvanishing of the A' root space in degree z forces nonvanishing of the A^* root space there. The direction of this surjection is important.

For an old label i, the replacement was designed so that $(A^*z)_i = (ATz)_i$. Thus T intertwines the old simple reflections of A^* with those of A. If z meets a marked vertex, old reflections leave at least one marked coordinate positive, and therefore each reflected root stays positive. Whenever an old row evaluation is strictly positive, reflection reduces the sum of old coordinates by a positive integer. At termination all rows of $A(Tz)$ are nonpositive. To check connectedness, replace each occupied marked vertex by its connected old support. An old–mark edge attaches to this support: if the old vertex lies outside it, $(Au)_i < 0$ forces a true old edge to the support. A mark–mark edge has $\Omega(u, v) \neq 0$, so there is a true old edge between the two supports. This proves connectedness of the resulting support. The terminal image belongs to the fundamental imaginary region, and undoing the old reflections proves the root-replacement claim.

For the induction step retain only source degrees componentwise at most the particular z. The complementary degrees form an ideal. If a and b in this box are neutral for $T^*\theta$, target genericity implies that Ta and Tb are both parallel to n: their positive images have height at most H. Consequently $\Omega'(a, b) = \Omega(Ta, Tb) = 0$. The neutral Lie algebra in this finite box is therefore abelian. Choose $\eta \in z^\perp$ arbitrarily near the pulled-back weight, preserving every strict sign in the box and avoiding every hyperplane in the full source lattice arising from a positive degree of height at most $h(z)$ not parallel to z. Such an η exists because the excluded finite collection consists of proper hyperplanes of $z^\perp$. Refine the old middle factor by η: its image in the box is abelian, so the degree-z coefficient remains in the new neutral factor unchanged. The quotient ensures that degrees outside the box cannot contribute to this coefficient. The canonical element at the lower cutoff is the projection of the higher-cutoff element, so the coefficient is available to the induction at height $h(z)$. This conclusion concerns the marked map and must not be transferred to the cover projection. Since each mark has image of height at least two, $h(z) < h(Tz) = H$. Uniform induction over all finite ranks is thus applicable to $\mathrm{prim}(z)$, whose image is a root on the n-ray. The primitive vector $\mathrm{prim}(n)$ is a root by the fundamental-region ray argument above.

References and source guide

[Hu] Andrew Hubery, *Quiver representations respecting a quiver automorphism: a generalisation of a theorem of Kac*, arXiv:math/0203195, Lemma 14(2) and Proposition 15.

[MQ] Travis Mandel and Fan Qin, *Bracelets bases are theta bases*, arXiv:2301.11101, Appendix A, Lemma A.5 and Theorem A.7.

[NS] Tomoki Nakanishi and Salvatore Stella, *Diagrammatic description of c-vectors and d-vectors of cluster algebras of finite type*, Electronic Journal of Combinatorics 21 (2014), P1.3. arXiv:1210.6299; §2.3, p. 8; §2.5.

[NII] Tomoki Nakanishi, *Cluster algebras and scattering diagrams, Part II: Cluster patterns and scattering diagrams*, arXiv:2103.16309, Definitions 5.12 and 6.12, Examples 5.13(b), 5.16(b), Theorem 6.13(a).

[NIII] Tomoki Nakanishi, *Cluster algebras and scattering diagrams, Part III: Cluster scattering diagrams*, arXiv:2111.00800, (1.8), Proposition 1.12, Lemma 3.8 and Theorem 3.12.

[Kac80] Victor G. Kac, *Infinite root systems, representations of graphs and invariant theory*, Inventiones Mathematicae 56 (1980), 57–92, DOI:10.1007/BF01403155, §1.3 and §2.4, Theorem 1(a).

[Kac] Victor G. Kac, *Infinite-Dimensional Lie Algebras*, 3rd ed., Cambridge University Press, 1990, Chapter 5, DOI:10.1017/CBO9780511626234.

[GK] Ofer Gabber and Victor G. Kac, *On defining relations of certain infinite-dimensional Lie algebras*, Bulletin of the American Mathematical Society 5 (1981), 185–189, DOI:10.1090/S0273-0979-1981-14940-5, Theorem 2 (over characteristic zero).

[Re] Markus Reineke, *The Harder–Narasimhan system in quantum groups and cohomology of quiver moduli*, arXiv:math/0204059, Proposition 2.5 and Lemma 4.7.

[Br] Tom Bridgeland, *Scattering diagrams, Hall algebras and stability conditions*, arXiv:1603.00416, version 4, §11.7 (correction).