THE CURVATURE OF A MOVE SET: DIRICHLET–WILSON BOUNDS AND CERTIFIED TRIVIALIZATION COSTS FOR NON-COMMUTATIVE TRANSPORT ON GRAPHS

Anonymous authors
Paper under double-blind review

ABSTRACT

Local search improves a candidate solution by applying moves from a fixed set, but its theory usually describes the search space through the objective. Moves need not commute: the two orders of a pair of moves may end at different candidates. We ask whether the resulting loop defect can be bounded from both sides, computed cheaply and certified against true optima. We model the moves as non-commutative transport on a base graph that records which move sequences are assumed equivalent. Loop holonomy measures the mismatch, and we study the minimum cost of making every loop close. A fractional packing of cycle defects and a representation-spectrum bound certify this cost from below; a spanning-tree gauge and a cycle-balanced section bound it from above. The bracket is exact on a single cycle; with the discrete metric, the cost equals a weighted frustration index on finite groups and vanishes on the state graph of any move set. Thus curvature depends on the relations the base graph assumes. We also connect the cost to three other readings: the point-level non-commutativity defect equals plaquette holonomy, the connection Laplacian separates into cut and holonomy branches with a corrected Cheeger inequality, and an ordered cohomological certificate bounds the distance a local algorithm must travel without exceeding graph distance. On 450 rotation instances a dual certifies the optima to within 3×10^{-8}, and interval arithmetic proves exactness for 490 instances. The cheap lower certificates recover on average 96 to 100 percent of the certified optimum. On the commutation lattice of two rotations, the certified cost per plaquette reaches eight times the single-plaquette value on a 32×32 box. These results give a computable, certifiable curvature relative to the relations a method assumes.

1 INTRODUCTION

Local search with a fixed set of moves is a common approach to hard combinatorial problems, from 2-opt reversals for routing to flip-and-repair rules for satisfiability. Its search space is studied through objective autocorrelation along random walks (Weinberger, 1990; Stadler, 1996), networks of local optima and their basins (Ochoa et al., 2008), and models of tunable ruggedness (Kauffman & Levin, 1987). The neighbourhood graph's spectral gap is another structural descriptor (Chung, 1997).

None of these descriptors records how the moves compose. Two moves applied in different orders can lead to different candidates, so carrying a candidate around a closed loop need not return it to its starting point. The mismatch after the loop is the holonomy of the move set (Figure 1a). It depends only on the operators and the base graph recording which sequences are treated as interchangeable; the objective-based descriptors above do not record it. Related areas study particular forms of this inconsistency. Group synchronization asks how well noisy relative rotations on a graph can be explained by absolute rotations (Singer, 2011; Rosen et al., 2019); signed and gain graphs study when edge labels can be made consistent (Harary, 1953; Zaslavsky, 1989); connection and sheaf Laplacians turn the consistency of local transports into a spectral quantity (Bandeira et al., 2013; Hansen & Ghrist, 2019; Bodnar et al., 2022). These results do not by themselves give a certified

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

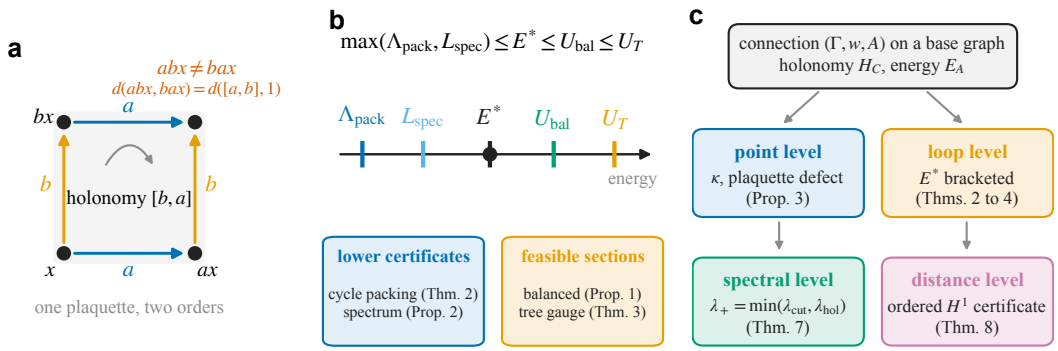

Figure 1: One connection, four conditional readings. (a) On the commutation lattice the two orders of a and b bound one plaquette; its holonomy is the commutator $[a, b] = b^{-1}a^{-1}ba$, and in the group model $d(abx, bax) = d([a, b], 1)$ (Proposition 3). (b) The bracket of Section 3. (c) The four levels and the results linking them, each under its stated hypotheses.

two-sided bound on the cost of making a general non-commutative transport consistent, nor connect that cost to point-level non-commutativity, spectral behaviour and travel distance.

The natural quantity to study is the minimal trivialization cost. Given a transport A_e in a compact group on each edge $e : u \to v$, it is the least value of $\sum_e w_e\, d(A_e x_u, x_v)^2$ over all assignments x_v of group elements to the vertices. Informally, it measures how far the transports are from being explained by one consistent choice of states, and it vanishes exactly when every loop has trivial holonomy (Theorem 1). Unfortunately, two obstacles stand between this definition and a usable theory. The cost is a non-convex minimization over a product of groups, so an optimizer returns only an upper bound; the quality of any cheaper certificate then cannot be judged without the optimum. Worse, the elementary certificates are of unknown quality. The holonomy defect of a single cycle bounds the cost from below, and a spanning tree that pushes the whole defect onto its co-tree edges bounds it from above; on a unit-weight cycle of length n the two differ by the factor n (Proposition 1).

We bracket the minimal trivialization cost by provable certificates that are cheap to compute, and we measure their gap against optima that are themselves certified. The main result is a Dirichlet–Wilson bracket:

$$\Lambda_{\mathrm{pack}}(A) \;\le\; E^*(A) \;\le\; \min_T U_T(A),$$

Here Λ_{pack} is the fractional cycle-packing lower bound, and U_T is the energy of a spanning-tree section (Section 3). A representation-spectrum bound strengthens the lower side; a cycle-balanced section improves the upper side. Intuitively, curvature forces a cost, while a tree shows one way to pay it. The bracket is exact on a single cycle of a connected compact Lie group with the geodesic distance, and on finite groups with the discrete metric the cost is the weighted frustration index of the gain graph. The ingredients are classical: a triangle inequality, LP duality, a tree gauge, the Ky Fan principle and a published Cheeger inequality. Our contributions are the base-graph lemma, capacity-corrected packing, exactness and tree-ratio results, partial-versus-full separation, certified-optimum layer and tightness map.

Around the bracket we prove three further readings of the same transport (Figure 1c). It turns out that on the state graph of any move set every loop closes, so curvature is always relative to a declared set of assumed relations (Lemma 2). Given that choice, the point-level non-commutativity defect equals the plaquette holonomy (Section 4), the connection Laplacian splits into a cut branch and a holonomy branch with a branch-aware Cheeger inequality (Section 5), and an ordered cohomological certificate bounds the distance a local algorithm must travel (Section 6).

We evaluate the bounds under a protocol fixed in advance, using optima certified by a tangent dual (450 instances) and by directed-rounding interval arithmetic (490 instances). The cheap lower certificates recover on average 96 to 100 percent of the certified optimum. On the icosahedron the spectral bound stays active while cycle packing reaches only 58 to 72 percent, even with all 12,878 simple cycles. On the commutation lattice of two rotations the certified cost per plaquette rises to 8.1 times the single-plaquette value on a 32×32 box. Three tempting strengthenings fail: a

full-frustration square lower bound, a distance certificate beyond graph distance, and a point-level defect that detects commutative holonomy. The four readings are therefore linked only under their stated conditions.

2 CONNECTIONS, HOLONOMY AND THE TRIVIALIZATION ENERGY

This section fixes the objects used throughout. Let $\Gamma = (V, E)$ be a finite connected graph without loops, $n = |V| \geq 2$, with edge weights $w_e > 0$. Each undirected edge is stored with one orientation $e : u \to v$ and reverse transport $A_{\bar{e}} = A_e^{-1}$, and energies count each undirected edge once. Let K be a compact group with a continuous bi-invariant distance d, so $d(agb, ahb) = d(g, h)$ for all $a, b, g, h \in K$. A finite group may carry the discrete metric $d(g, h) = c\,\mathbf{1}[g \neq h]$, and a compact Lie group the geodesic distance of a bi-invariant Riemannian metric. For a section $x = (x_v) \in K^V$ we define the energy and its minimum:

$$E_A(x) = \sum_{e:u\to v} w_e\, d(A_e x_u, x_v)^2, \qquad E^*(A) = \min_{x\in K^V} E_A(x), \tag{1}$$

where the minimum exists by compactness. Edge by edge, $E_A(x)$ measures how far the transported state $A_e x_u$ misses the assigned state x_v. For a closed walk $C = (e_1, \ldots, e_m)$ with edges listed in traversal order, the holonomy is $H_C(A) = A_{e_m} \cdots A_{e_1}$. Changing the base point conjugates H_C and reversing the walk inverts it, so by bi-invariance $d(H_C, 1)$ depends on neither choice.

Theorem 1 (zero energy is trivializability)**.** *$E_A \geq 0$, and $E_A(x) = 0$ if and only if $A_e x_u = x_v$ on every edge. The following are also equivalent: (a) $E^*(A) = 0$; (b) $H_C(A) = 1$ for every closed walk C; (c) for one spanning tree, every fundamental cycle pulled back to the root has trivial holonomy.*

Put another way, a connection costs nothing exactly when it is a change of coordinates in disguise, so that one choice of states makes every transport consistent. The proof in Appendix A propagates a root value along the tree. The energy is invariant under the local gauge $x'_v = h_v x_v$, $A'_{u\to v} = h_v A_{u\to v} h_u^{-1}$ and under the global right action $x_v \mapsto x_v q$, but not under a global left action. Vanishing plaquette holonomy is flatness, which is weaker than condition (b): flatness alone need not give zero energy, whereas zero energy is exactly trivializability.

Lemma 1 (gauge-fixed form)**.** *With $A^x_e = x_v^{-1} A_e x_u$, $E_A(x) = \sum_e w_e\, d(A^x_e, 1)^2$.*

Thus E_A is the squared size of the connection in the gauge x, and E^* is the cost of the best gauge. By contrast, the Wilson and Manton actions of lattice gauge theory depend on plaquette holonomies alone (Wilson, 1974; Manton, 1980); the name of the bracket records this difference.

3 THE DIRICHLET–WILSON BRACKET

We now bound E^* from both sides: holonomy around cycles forces a cost from below, and explicit sections pay a known cost from above. For a simple cycle C write $s_C = \sum_{e\in C} w_e^{-1}$ and $h_C(A) = d(H_C(A), 1)^2 / s_C$. A function $\mu \geq 0$ on a finite family $\mathcal{C}$ of simple cycles is a *fractional packing* if $\sum_{C\ni e} \mu_C \leq 1$ for every edge; think of each edge as having one unit of capacity to share among the cycles through it.

Theorem 2 (cycle lower bound and fractional packing)**.** *For every x and every simple cycle C, $h_C(A) \leq \sum_{e\in C} w_e\, d(A_e x_u, x_v)^2 \leq E_A(x)$. Consequently, every fractional packing μ on $\mathcal{C}$ gives the bounds:*

$$E_A(x) \;\geq\; \sum_{C\in\mathcal{C}} \mu_C\, h_C(A), \qquad E^*(A) \;\geq\; \Lambda_{\mathrm{pack}}(A;\mathcal{C}) := \max_{\mu \text{ packing}} \sum_{C\in\mathcal{C}} \mu_C\, h_C(A). \tag{2}$$

Concretely, a cycle with non-trivial holonomy must pay its defect somewhere along its edges, and the packing prevents an edge from being charged twice. The single-cycle bound follows from a triangle inequality and weighted Cauchy–Schwarz; the packing bound multiplies it by μ_C and exchanges the sums (Appendix B). Any feasible μ is a certificate, and a floating-point LP solution can be repaired by directed capacity scaling. Cycles that repeat an edge need the capacity $\sum_C \mu_C n_C(e) \leq 1$ with multiplicities $n_C(e)$; without it the bound is false (Section 7).

Theorem 3 (tree-gauge upper bound)**.** *Fix a root r and a spanning tree T, let P_v be the tree transport from r to v, and for a non-tree edge $e : u \to v$ let $B_e = P_v^{-1} A_e P_u$. Then $E^*(A) \le U_T(A) := \sum_{e \notin T} w_e \, d(B_e, 1)^2$.*

Put differently, a tree gauge makes every tree edge consistent and charges each co-tree edge the full holonomy of its fundamental cycle: the section $x_v = P_v$ has zero residual on tree edges and residual $d(B_e, 1)$ on co-tree edges, where B_e is that holonomy. Combining the two theorems gives the bracket:

$$\Lambda_{\mathrm{pack}}(A; \mathcal{C}) \le E^*(A) \le \min_{T \in \mathcal{T}} U_T(A). \tag{3}$$

The bracket holds for any cycle family $\mathcal{C}$ and tree family $\mathcal{T}$ examined.

Proposition 1 (exactness on a cycle; unicycles)**.** *Let K be a connected compact Lie group with the geodesic distance of a bi-invariant Riemannian metric. If Γ is a single cycle C then $E^*(A) = d(H_C, 1)^2 / s_C$, attained by placing increments of a minimal geodesic from 1 to H_C on the edges in proportion to w_e^{-1}. If Γ is a connected unicyclic graph with cycle C, the same value is attained by a cycle-balanced section that propagates the balanced cycle frame along the pendant trees at zero residual; the tree gauge that deletes a cycle edge e_* has energy $w_{e_*} d(H_C, 1)^2$, so its best ratio is $U_T / E^* = w_{\min} s_C$. With one designated cycle edge of weight ϵ and unit weights elsewhere the ratio is exactly $1 + (|C| - 1)\epsilon$. For $\mathrm{SO}(3)$ with $d(H_C, 1) < \pi$ the minimizer is unique modulo the global right action.*

In words, on a single cycle the optimal section spreads the holonomy evenly, in proportion to the edge resistances w_e^{-1}, whereas a tree gauge places all of it on one edge; the ratio $w_{\min} s_C$ is the price of that choice. The proposition needs the geodesic distance: its proof splits the holonomy along a minimal geodesic, and chordal distances and finite groups admit no such fractional roots.

Proposition 2 (representation-spectrum lower bound)**.** *Let $\rho : K \to \mathrm{O}(q)$ be a representation with $\|\rho(g) - \rho(h)\|_F^2 \le C_\rho \, d(g, h)^2$, and let Δ_ρ be the normalized connection Laplacian of the edge maps $\rho(A_e)$, that is, the degree-normalized matrix of the form $z \mapsto \sum_{e : u \to v} w_e \|\rho(A_e) z_u - z_v\|^2$ (studied in Section 5). Write $d_v = \sum_{e \ni v} w_e$ and $\mathrm{vol}(V) = \sum_v d_v$. Then the smallest eigenvalues bound the cost from below:*

$$E^*(A) \;\ge\; L_{\mathrm{spec}}(A) := \frac{\mathrm{vol}(V)}{C_\rho} \sum_{j=1}^{q} \lambda_j(\Delta_\rho). \tag{4}$$

For the standard representations of $\mathrm{SO}(2)$ and $\mathrm{SO}(3)$ with the angular distance one may take $C_\rho = 2$.

The sum of the q smallest eigenvalues is the Ky Fan relaxation of the corresponding chordal energy, as in orthogonal synchronization (Bandeira et al., 2013). The proposition transfers this relaxation to the geodesic cost through C_ρ. The proof stacks the blocks $\sqrt{d_v}\rho(x_v)/\sqrt{\mathrm{vol}(V)}$ into an isometry and applies the Ky Fan principle. Neither lower certificate dominates the other: the loop bound wins on cycles and on 39 of 300 theta records, the spectral bound on every icosahedron record (Section 7).

Theorem 4 (finite groups)**.** *Let K be finite, $\delta_K = \min_{g \ne 1} d(g, 1)$, $D_K = \max_g d(g, 1)$, and let $\ell_w(A)$ be the least total weight of an edge set whose removal leaves a trivializable connection on every component. Then $\delta_K^2 \, \ell_w(A) \le E^*(A) \le D_K^2 \, \ell_w(A)$; under the discrete metric $d = c\,\mathbf{1}[g \ne h]$, $E^*(A) = c^2 \ell_w(A)$.*

The theorem places the finite-group cost between two multiples of the minimum edge-removal weight; the discrete metric makes the relation exact. For $K = \mathbb{Z}/2$, this is signed-graph frustration and balance (Harary, 1953; Atay & Liu, 2020); for general finite K, it is gain-graph balance (Zaslavsky, 1989). Magnetic Laplacians use a related index (Lange et al., 2015), while our exact identity follows from Theorem 4. On finite groups the cycle certificate of Theorem 2 is loose by a factor of order $|C|$ (Figure 5), and the right certificate is the hitting/packing dual of ℓ_w (Appendix B).

4 Point level and local level

Sections 4 to 6 read the same connection at three further levels: the point and spectral levels are tied to E^* by stated inequalities, and the distance level reads a sheaf on the same base graph (Figure 1c). First, however, we must say which graph carries a move set.

Base graph of a move set. A move set is a family of operators $\sigma_1, \dots, \sigma_p$ on a metric space (X, d_X); in the group model $X = K$ and $\sigma_i(x) = a_i x$. A *base graph* for the move set has move-labelled edges, with $A_e = a_i$ on edges labelled i; a section assigns a state per vertex, and its energy is the mismatch between the labelled moves and those states. The state graph, with edges $x \to \sigma_i x$, is a base graph without curvature: the tautological section $x_v = v$ has zero energy for every move set, free or not. Curvature therefore lives on base graphs that identify move sequences the state graph keeps apart, and the following lemma says which.

Lemma 2 (base graphs and assumed relations)**.** *Let F_p be the free group on the move symbols, $N \trianglelefteq F_p$, $Q = F_p/N$, and let Γ_Q be the Cayley graph of Q with generator-labelled edges and $A_e = a_i$. Then (i) the closed walks of Γ_Q are the words in N, and the holonomy of a closed walk is the value of its word in K; (ii) $E^* = 0$ on every finite connected subgraph of Γ_Q if and only if $i \mapsto a_i$ factors through Q, that is, every relation assumed in Q holds in K; (iii) for $Q = \langle a_1, \dots, a_p \rangle \le K$, the state graph, $E^* = 0$ with $x_q = q$; for $Q = \mathbb{Z}^p$, the commutation lattice, every unit square is a plaquette with holonomy $c_{ij} = a_j^{-1} a_i^{-1} a_j a_i$, and $E^* = 0$ on all boxes if and only if the a_i commute.*

The lemma says that a base graph encodes which move sequences a method treats as equivalent, and the holonomy measures how wrong that assumption is: the commutation lattice is the base graph of a method that applies moves in any order, and the right-angled Artin group (the free group modulo only the declared commutation relations) of an independence relation is that of a method reordering only declared-independent moves. Bit flips commute, so the flip move set of satisfiability is flat on its commutation lattice, where no curvature enters its difficulty. Bit flips F_i and adjacent swaps S_j on $\{0,1\}^N$, by contrast, generate an isometric hyperoctahedral action under the Hamming metric. The commutator of F_i with a swap that moves bit i is a double flip, of displacement 2 at every point, so their commutation lattice is curved, and Theorem 5 gives explicit plaquette bounds there.

Square lemma. The point-level non-commutativity defect of a move set is $\kappa(x) = \binom{p}{2}^{-1} \sum_{i<j} d_X(\sigma_i\sigma_j x, \sigma_j\sigma_i x)^2$, the mean squared mismatch between the two orders of each pair of moves. It needs no objective values, and in the group model it is a plaquette quantity of the commutation lattice.

Proposition 3 (group model)**.** *Let $a_1, \dots, a_p \in K$ and $c_{ij} = a_j^{-1} a_i^{-1} a_j a_i$. Then (i) $\kappa(x) = \binom{p}{2}^{-1} \sum_{i<j} d(c_{ij}, 1)^2$ for every x; (ii) the plaquette with forward labels $a_i, a_j, a_i^{-1}, a_j^{-1}$ has holonomy c_{ij}, so with unit weights $E^* \ge d(c_{ij}, 1)^2/4$ for every plaquette present in Γ, and pairwise edge-disjoint plaquettes add up to $E^* \ge \binom{p}{2}\kappa/4$, while plaquettes sharing edges combine only through the capacities of Theorem 2; (iii) if the a_i commute then $\kappa \equiv 0$ although H_C and E^* can be arbitrary (for instance $\mathrm{SO}(2)$ on a cycle with non-zero total angle).*

In effect, in the group model κ does not depend on the point at all: it is the average plaquette holonomy, and each plaquette present in the base graph forces an energy of at least a quarter of its squared commutator. Part (iii) explains why κ, E^*, the spectral gap λ_+ (the smallest positive eigenvalue of the connection Laplacian, defined in Section 5), and the distance certificate are four conditionally linked readings rather than four readings of one curvature: κ reads only the plaquette, or commutator, part of the holonomy and misses flat but non-trivializable transport. On plaquette loops this identity is exact; on general loops only the one-directional bound (ii) holds.

Theorem 5 (isometric square lemma)**.** *Let each σ_i be an isometric bijection of (X, d_X), $\Phi_{ij} = \sigma_j^{-1}\sigma_i^{-1}\sigma_j\sigma_i$, and $a_{ij}(x) = d_X(x, \Phi_{ij}x)$. Then $d_X(\sigma_i\sigma_j x, \sigma_j\sigma_i x) = a_{ij}(x)$, $|a_{ij}(x) - a_{ij}(y)| \le 2d_X(x,y)$, and for any 4-cycle whose edge maps are isometric bijections composing to Φ, any section y with $d_X(y_{v_0}, x_0) \le \rho$ has energy $E \ge \frac{1}{4}\big(d_X(x_0, \Phi x_0) - 2\rho\big)_+^2$. For $N = \binom{p}{2}$ pairwise edge-disjoint plaquettes with the same ball constraint, $E_{\mathrm{total}} \ge \frac{N}{4}\big(\kappa(x_0) - 4\rho\sqrt{\kappa(x_0)}\big)_+$.*

The localized bound trades off (i) the commutator displacement $d_X(x_0, \Phi x_0)$, which depends only on the moves and the reference state, against (ii) the radius ρ within which the section is pinned; it is informative while 2ρ stays below the displacement. The theorem replaces the group model's three hypotheses with one, that each move is an isometric bijection; we make no claim for non-injective repair operators.

Theorem 6 (local spectrum at a trivializable connection)**.** *Let K be a compact Lie group with a bi-invariant Riemannian metric, let x^* satisfy $E_A(x^*) = 0$, and write $x_v = x_v^* \exp(\xi_v)$ in a normal*

neighbourhood. Then $E_A(x) = \sum_{e:u\to v} w_e\|\xi_v - \xi_u\|^2 + O(\|\xi\|^3)$ *and* $\mathrm{Hess}_{x^*} E_A = 2L_\Gamma \otimes I_{\mathfrak{k}}$ *with* L_Γ *the unnormalized graph Laplacian. Modulo the global right action, the linearized gradient flow contracts as* $e^{-2\lambda_2(L_\Gamma)t}$ *in norm and* $e^{-4\lambda_2 t}$ *in energy, and the nonlinear flow contracts at every rate below* $2\lambda_2$ *in a sufficiently small neighbourhood (Appendix C).*

Near a zero-energy connection, then, the energy is a graph Dirichlet form in disguise, so gradient flow relaxes like heat diffusion on the base graph, at a rate set by its Fiedler value. The statement is local: a discrete step of size τ has modal rate $-\log|1 - 2\tau\lambda_j|/\tau$, and convergence from arbitrary initializations is not claimed.

5 The spectral branch

We now turn to the spectrum. After a choice of gauge, the connection Laplacian separates into a part governed by the graph and a part governed by holonomy. Let every edge carry $Q_e \in \mathrm{O}(q)$, either $\rho(A_e)$ or an independent orthogonal map, and let $z_v \in \mathbb{R}^q$. With $d_v = \sum_{e\ni v} w_e$ and $D = \mathrm{diag}(d_v)$, the connection Laplacian L_Q has quadratic form $\sum_{e:u\to v} w_e\|Q_e z_u - z_v\|^2$. Define $\Delta_Q = (D^{-1/2}\otimes I)L_Q(D^{-1/2}\otimes I)$ and $\Delta_\Gamma = D^{-1/2}L_\Gamma D^{-1/2}$, where L_Γ is the unnormalized scalar graph Laplacian. The Rayleigh quotient is $\mathcal{R}(z) = z^\top L_Q z / \sum_v d_v\|z_v\|^2$. A unit physical section has $\|z_v\| = 1$, or $\|\phi_v\| = \sqrt{d_v}$ in normalized coordinates $\phi_v = \sqrt{d_v}z_v$; the distinction matters on non-regular graphs.

Proposition 4 (common root and branch decomposition)**.** *Fix a root* r *and a spanning tree with tree transports* $P_v \in \mathrm{O}(q)$*. Let* $B_e = P_v^\top Q_e P_u$ *for co-tree edges,* $\mathcal{K}_r = \bigcap_{e\notin T}\ker(B_e - I)$*,* $k = \dim\mathcal{K}_r$*, and* $\mathcal{K}_v = P_v\mathcal{K}_r$*. Then* $Q_e\mathcal{K}_u = \mathcal{K}_v$*, the orthogonal complements are preserved,* $\ker\Delta_Q$ *has dimension* k*, and* $\Delta_Q \simeq (\Delta_\Gamma\otimes I_k)\oplus\Delta_{\mathrm{hol}}$ *with* Δ_{hol} *positive definite when* $q > k$*. Hence* $\lambda_+(\Delta_Q) = \min\{\lambda_{\mathrm{cut}}, \lambda_{\mathrm{hol}}\}$*, where* $\lambda_{\mathrm{cut}} = \lambda_2(\Delta_\Gamma)$ *if* $k > 0$ *and* $\lambda_{\mathrm{hol}} = \lambda_1(\Delta_{\mathrm{hol}})$ *if* $q > k$*, with the missing branch set to* $+\infty$*.*

Directions fixed by every holonomy behave like ordinary graph signals (the cut branch). The remaining directions are pinned by holonomy (the holonomy branch), and the smaller branch eigenvalue sets the smallest positive eigenvalue. Fixed spaces at different base points lie in different fibres; intersecting them in one coordinate system gives incorrect dimensions after a local gauge. On the fibrewise complements $\mathcal{K}_v^\perp$ we define the partial and full frustrations:

$$\eta_{\mathrm{p}} = \min\{\mathcal{R}(z) : z\neq 0,\ z_v\in\mathcal{K}_v^\perp,\ \|z_v\|\in\{0,1\}\},\quad \eta_{\mathrm{f}} = \min\{\mathcal{R}(z) : z_v\in\mathcal{K}_v^\perp,\ \|z_v\| = 1\ \forall v\}. \tag{5}$$

The partial frustration lets a test section switch off at some vertices, while the full frustration forces a unit vector everywhere; both use the original graph's degrees in the denominator. Bandeira et al. (2013) prove $\eta_{\mathrm{p}}^2/10 \le \lambda_{\mathrm{hol}} \le \eta_{\mathrm{p}}$ in their Theorem 2.2 and $\lambda_2(\Delta_\Gamma)\eta_{\mathrm{f}}/44 \le \lambda_{\mathrm{hol}} \le \eta_{\mathrm{f}}$ in their Theorem 2.4; the full version carries the spectral gap of the base graph, which matters below.

Theorem 7 (branch-aware Cheeger inequality)**.** *Let* h_Γ *be the normalized Cheeger constant of* Γ*,* $h_{\mathrm{cut}} = h_\Gamma$ *if* $k > 0$ *and* $+\infty$ *otherwise,* $h_{\mathrm{hol}} = \eta_{\mathrm{p}}/\sqrt{5}$ *if* $q > k$ *and* $+\infty$ *otherwise, and* $h_{\mathrm{p}}^\star = \min\{h_{\mathrm{cut}}, h_{\mathrm{hol}}\}$*. Then:*

$$\frac{(h_{\mathrm{p}}^\star)^2}{2} \;\le\; \lambda_+(\Delta_Q) \;\le\; \min\{2h_{\mathrm{cut}},\ \sqrt{5}\,h_{\mathrm{hol}}\} \;\le\; \sqrt{5}\,h_{\mathrm{p}}^\star. \tag{6}$$

The constant $h_{\mathrm{p}}^\star$ collects two independent obstructions: (i) the cut branch, present only if some direction is fixed by every holonomy ($k > 0$), whose constant depends on the graph alone through h_Γ; and (ii) the holonomy branch, present only if some direction is not ($q > k$), whose constant depends on the transports through η_{p}. Each branch is the corresponding classical inequality (Appendix D). Since each branch is controlled only within constant factors, the branch attaining $h_{\mathrm{p}}^\star$ need not be the branch attaining λ_+, and the verification records both. Naturally, one might hope to use the simpler full frustration η_{f} instead. Unfortunately, the square lower bound then fails:

Proposition 5 (no full-frustration square bound)**.** *Take two unit-weight triangles joined by one edge of weight* $\epsilon > 0$*, with sign transports* $+1$ *on one triangle and* -1 *on the other. Then the common fixed-space dimension is* $k = 0$*,* $\eta_{\mathrm{f}} = 2/(6+\epsilon)$ *and* $\lambda_+ \le \epsilon/(6+\epsilon)$*. For* $\epsilon = 10^{-3}$*,* $\eta_{\mathrm{f}}^2/10 > \lambda_+$*.*

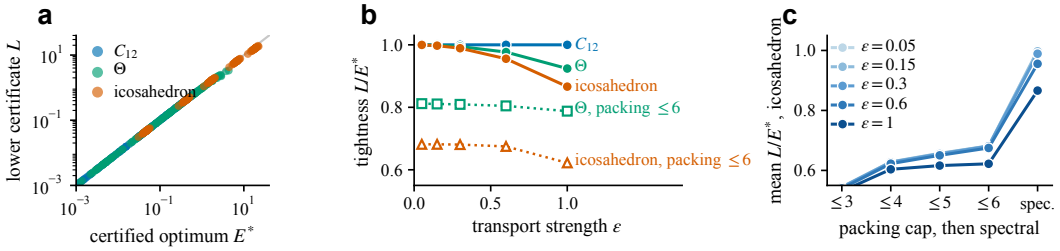

Figure 2: Dirichlet–Wilson bracket on certified optima (V2, R5). 450 $\mathrm{SO}(3)$ instances, $\{C_{12}, \Theta, \text{icosahedron}\} \times \varepsilon \in \{0.05, 0.15, 0.3, 0.6, 1\} \times 30$ underlying random inputs, each transport a rotation by ε about a random axis; optima certified to width 2.84×10^{-8} (Appendix F). (a) Lower certificate $L = \max(\Lambda_{\text{pack}}, L_{\text{spec}})$ against E^*. (b) Mean tightness L/E^* by strength, with the packing branch (cycles of length at most 6) shown separately. (c) Certificate ladder on the icosahedron.

The surprising part is that dropping the common fixed space (by using the full frustration) does not restore the square bound: a small eigenvalue may localize its eigenvector on the balanced triangle, while a full unit section must pay for the frustrated one. Figure 3a shows the same failure on a two-cycle family with a variable bridge, where the partial sandwich holds throughout.

6 The Ordered Cohomological Distance Certificate

The last reading asks how far a local algorithm must travel between two states, and whether cohomology can certify more than graph distance. Let $\mathcal{F}$ be a real cellular sheaf on Γ (vector spaces on vertices and edges with linear restriction maps), with coboundary $\delta : C^0 \to C^1$ and $H^1 = \operatorname{coker} \delta$, and fix vertices $s \neq t$ and $\ell \geq 1$. An *ordered ℓ-local certificate with m levels* is a map $\rho : V \to \{0, \dots, m\}$ with $\rho(s) = 0$, $\rho(t) = m$ and $|\rho(u) - \rho(v)| \leq 1$ whenever $d_\Gamma(u, v) \leq \ell$, together with cochains $\omega_1, \dots, \omega_{m-1}$ with linearly independent classes in H^1, each supported on an s–t edge separator crossed at its level; $\mathfrak{b}_\ell$ is the largest feasible m. Geometrically, a certificate slices the graph into m ordered layers that a step of length at most ℓ cannot skip, and each crossing is witnessed by an independent cohomology class.

Theorem 8 (distance certificate)**.** *Every certificate satisfies $m \leq \lceil d_\Gamma(s,t)/\ell \rceil$, and every trajectory from s to t whose steps move distance at most ℓ takes $T \geq m$ steps; hence $\mathfrak{b}_\ell \leq \lceil d_\Gamma(s,t)/\ell \rceil$. The bound is attained on the path family with k zero-restriction edges ($H^1 \cong \mathbb{R}^k$, $\mathfrak{b}_1 = k + 1 = d_\Gamma(s,t)$) and not in general: k classes stacked on one bridge give $\mathfrak{b}_1 = 1$, and a cycle with a constant sheaf gives at most two levels.*

Put simply, the certificate lower-bounds the number of steps but can never beat breadth-first distance. Its content is this ceiling, which refutes the tempting reading "$\dim H^1 > 0$ implies $T \geq \dim H^1 + 1$": a hole is not a wall. Its one advantage is that classes and separators are local data, in principle assembled without expanding the state graph; we give no efficient construction (Appendix E).

7 Verification

So far every statement has been a theorem; we now test each one, and the tightness of the bracket, numerically. Each module has a criterion fixed in advance and an evidence level (A analytic, B verified interval, C numerical, D feasible, X exploratory), and a development batch precedes a single confirmation batch (Table 2, Appendix H). Fourteen consistency checks (V0), among them the fixed-space pitfall of Proposition 4, run first and all pass.

Analytic family (V1). On the $\mathrm{SO}(3)$ commutator plaquette, $\psi = d([a,b],1)$, $\kappa = \psi^2$, $E^* = \psi^2/4$ and the full rank-3 spectrum are closed-form (Appendix D.1); they agree with independent computation to 3.2×10^{-14} and with an optimizer to 0.06% (Figure 7 in Appendix G).

Bracket study (V2). The panel crosses $\mathrm{SO}(2)$ and $\mathrm{SO}(3)$ with C_{12}, the theta graph Θ and the icosahedron at five strengths, with 30 underlying random inputs per cell (900 records). Let $L = \max(\Lambda_{\text{pack}}, L_{\text{spec}})$ and let $U = U_{\text{best}}$ be the lowest energy among the computed feasible sections;

Table 1: Certificate cost and tightness: median milliseconds per instance (one core) and mean ratio to the certified optimum on 30 icosahedron and 30 theta $\mathrm{SO}(3)$ records, and on one ring-plus-matching graph per size (degree 3, $\varepsilon = 0.3$; ratios use the best feasible value from 8 restarts (2 for $n = 768$)). Packing caps 6 and 4; the tangent dual needs the optimizer's candidate.

	icosahedron		theta		$n = 192$		$n = 768$	
certificate	ms	ratio	ms	ratio	ms	ratio	ms	ratio
spectral, Prop. 2	0.37	0.96	0.11	0.99	32	0.97	1064	0.97
packing, Thm. 2	7.9	0.68	0.83	0.84	1.4	0.02	3.2	0.01
tree gauge, Thm. 3	0.19	3.80	0.04	3.45	1.6	12.4	7.8	16.3
L-BFGS-B, all restarts	967	1	154	1	5338	1	5189	1
tangent dual, App. F	1.0	1.00	0.33	1.00	27	1.00	709	1.00

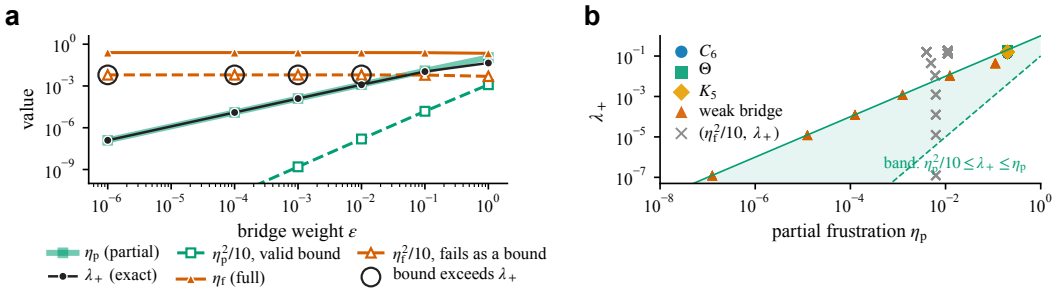

Figure 3: Spectral branch (V3). (a) Two unit-weight 4-cycles joined by a bridge of weight ε, one frustrated (common fixed-space dimension $k = 0$), all quantities exact: $\eta_{\mathrm{p}}^2/10 \le \lambda_+ \le \eta_{\mathrm{p}}$ holds throughout, while $\eta_{\mathrm{f}}^2/10$ exceeds λ_+ once $\varepsilon \le 10^{-2}$ (circled). (b) All 120 rank-one instances lie in the band $[\eta_{\mathrm{p}}^2/10, \eta_{\mathrm{p}}]$; the crosses ($\eta_{\mathrm{f}}^2/10$) leave it in 18 cases.

every record satisfies $L \le U \le U_T$. A floating-point global certificate (R5) bounds the optimum on all 450 $\mathrm{SO}(3)$ records: the tangent chordal dual handles 448 records, and a strengthened $\mathrm{conv}\,\mathrm{SO}(3)$ dual handles the other 2. The largest gap to the feasible value is 2.84×10^{-8}, so these values are within 3×10^{-8} of the optima (Figure 2). Mean tightness L/E^* is 1.000 on C_{12}, 0.979 on Θ and 0.962 on the icosahedron; at $\varepsilon = 1$ it falls to 0.924 and 0.866 on the latter two graphs.

The representation-spectrum bound is active on all 150 icosahedron records and 144 of 150 theta records; a single cycle gives more on the other 6. It reduces the mean $\mathrm{SO}(3)$ gap $U - L$ from 0.687 to 0.197. On the icosahedron, however, the packing LP reaches only 0.62 to 0.68 of the optimum as the cap grows (Figure 2c). Even with all 12,878 simple cycles, it reaches 0.58 to 0.72 on a held-out batch of 40 instances (Appendix G). For $\mathrm{SO}(2)$, which has no global certificate, we report only L/U (Appendix G); enumeration on the finite panel (240 records, S_3 and Q_8) confirms $E^* = \ell_w$.

What each certificate costs. The candidate-free certificates take 0.1 to 10 milliseconds per 12-vertex instance, the 32-restart search that supplies the feasible value about one second, and the tangent dual, once a candidate exists, as little as the spectral bound (Table 1). On sparse random graphs with 768 vertices the spectral bound still reaches 0.97 of the best feasible value in one second, while short-cycle packing collapses for want of short cycles. SE-Sync-type certifiers (Rosen et al., 2019) bound the chordal objective, not E^*.

Exact-real certificates (R6 to R11). Floating-point certificates leave one question open: does the real instance encoded by the stored binary64 data have its global minimizer where the candidate is? For the 450 $\mathrm{SO}(3)$ records and 40 additional icosahedron instances, directed-rounding interval arithmetic proves that every global minimizer lies in a fixed-root ball of radius at most 1.64×10^{-4} on which the chart Hessian is at least $0.0186\, I$ (Appendix F). All 490 pass with a global gap of at most 7.22×10^{-11}; the statement concerns these instances only.

Spectral branch (V3). We build connections with prescribed fixed dimension k for all nine legal pairs (q, k) with $q \leq 3$ on C_6, Θ, K_5 and a weak-bridge graph, under a Haar-random $\mathrm{O}(q)$ gauge at every vertex (1080 records). On all 1080 the common-root dimension is the requested k and the spectrum splits into the two sub-spectra to 2×10^{-8}. For $q = 1$, with frustrations exact by enumeration, equation 6 holds on all 240 rank-one records and the full-frustration bound fails on the 18 weak-bridge records with $\varepsilon \leq 10^{-2}$ (Figure 3). For $q > 1$ they are upper estimates (Appendix G).

Curvature of a two-move set. What does the curvature of an actual move set look like? We compute the bracket and the tangent dual on $m \times m$ boxes ($m \leq 32$) of the commutation lattice of $a = R_z(\theta)$ and $b = R_x(\varphi)$, with commutator angle ψ (Lemma 2). At $\psi = 0.21$ every optimum is certified (relative gap below 2×10^{-8}). The energy per plaquette rises from $\psi^2/4$ at $m = 2$ to $8.1\,\psi^2/4$ at $m = 32$, with values $7.4, 7.9, 8.1$ at $m = 16, 24, 32$. On these boxes the curvature behaves like a density set by the commutator rather than the single-plaquette value. Along the same series the spectral certificate falls from 1.00 to 0.72 of the optimum, and unit-square packing to 0.06; only the candidate-derived dual closes the bracket on these lattices. At $\psi = 0.47, 0.94, 1.80$, the dual certifies only boxes with $m \leq 4, 6, 3$, respectively. Longer runs change the feasible values by at most 5 percent, but the larger boxes remain uncertified: their spectral lower bound is 0.39 to 0.90 of the feasible value (Figure 11).

Distance certificate (V4), local rates (V5), square lemma (V6). The path family attains the ceiling of Theorem 8 with $\lceil (k+1)/\ell \rceil$ levels, matched by an executed ℓ-local walk. Even cycles with antipodal endpoints fall short: they have two levels against $\lceil n/(2\ell) \rceil$ steps (Figure 8). Gradient flows match the rates of Theorem 6 to four digits (Figure 9), and 3000 hyperoctahedral elements confirm Theorem 5. The TSP association (V7) is reported in Appendix G.

8 Related work

We distinguish five related lines of work. *Synchronization.* Chordal E^* reduces to angular synchronization for $\mathrm{SO}(2)$ (Singer, 2011; Wang & Singer, 2013; Bandeira et al., 2017) and rotation averaging for $\mathrm{SO}(3)$ (Hartley et al., 2013). Strong duality can certify low-noise rotation averages (Eriksson et al., 2018); SE-Sync certifies poses (Rosen et al., 2019), while Burer–Monteiro theory analyzes low-rank semidefinite programs (Boumal et al., 2016). Vision uses loop constraints (Zach et al., 2010; Huang & Guibas, 2013); our packing LP turns them into a lower bound, and our tangent dual uses the spectrahedral description of $\mathrm{conv}\,\mathrm{SO}(3)$ (Saunderson et al., 2015) and interval methods (Rump, 2010). *Gain graphs.* Signed and gain graphs study balance (Harary, 1953; Zaslavsky, 1989; Atay & Liu, 2020); magnetic Laplacians use related frustration (Lange et al., 2015). *Sheaf and connection Laplacians.* Sheaf theory (Hansen & Ghrist, 2019; 2021), neural sheaf diffusion (Bodnar et al., 2022) and connection Cheeger theory (Bandeira et al., 2013) share a Laplacian framework. The connection Laplacian is its orthogonal-bundle case; our decomposition separates consensus from holonomy for orthogonal sheaves without fixing a gauge. *Topological obstructions.* Section 6 follows the path-to-obstruction scheme of Herlihy & Shavit (1999). Cellular-sheaf cohomology supplies its language (Curry et al., 2016); Abramsky & Brandenburger (2011) study a distinct local–global obstruction. Stronger lower bounds may use distributed-locality or overlap-gap arguments (Linial, 1992; Gamarnik, 2021); our potential ρ gives only graph distance. *Curvature on graphs.* Ricci-type graph curvature (Ollivier, 2009; Forman, 2003) has been linked to over-squashing (Topping et al., 2022). Those ignore transport labels; ours vanishes when loops close, even on a positively curved graph, yet can be non-zero on one cycle.

9 Limitations and conclusion

The base graph's assumed relations are a modelling choice. Except for Theorem 6, our theorems use no objective or acceptance rule and establish neither search convergence nor complexity. Exactness holds on one cycle; exact-real certificates cover the listed 12-vertex instances, and the $\mathrm{SO}(2)$ panel has no exact certificate. We report no learning experiment, and non-orthogonal learned sheaves lie outside Section 5. The lattice experiment does not link curvature to search difficulty. The result is a computable, certifiable curvature relative to assumed relations: cheap certificates average within 4 percent of the $\mathrm{SO}(3)$ panel's optima and lie within 28 percent of the optimum on a 1024-vertex box.

AI USE STATEMENT

In this work, we used generative AI tools for three tasks: to assist with the experiments, including writing the experiment, certificate-timing and figure-generation scripts; to write the first draft of the manuscript; and for the early literature survey and the bibliography check. The research question, the theory and the scope of the claims were set by the authors. AI-assisted code was verified by rerunning it on the released records and comparing every reported number with the regenerated output, and the authors revised the draft and checked the cited literature against the original sources. We have reviewed all AI-assisted work and take responsibility for the final content of this work, including text, claims and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

The work is mathematical and computational, uses no human subjects or personal data, and involves synthetic instances only. We are aware of no ethical concern beyond the accurate scoping of claims, which the text addresses explicitly.

REPRODUCIBILITY STATEMENT

Complete proofs are in Appendices A to E. The certification pipeline, including the exact-real reification, the strictly complementary dual, the fixed-root radius conversion and the interval Hessian test, is specified in Appendix F. The protocol in Appendix H specifies seeds, identities, evidence levels, four-state checks, budgets and runtimes. The supplementary code release provides all JSON records, scripts and one-command verification.

REFERENCES

Samson Abramsky and Adam Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13(11):113036, 2011. doi: 10.1088/1367-2630/13/11/113036.

Fatihcan M. Atay and Shiping Liu. Cheeger constants, structural balance, and spectral clustering analysis for signed graphs. *Discrete Mathematics*, 343(1):111616, 2020. doi: 10.1016/j.disc.2019.111616.

Afonso S. Bandeira, Amit Singer, and Daniel A. Spielman. A Cheeger inequality for the graph connection Laplacian. *SIAM Journal on Matrix Analysis and Applications*, 34(4):1611–1630, 2013. doi: 10.1137/120875338.

Afonso S. Bandeira, Nicolas Boumal, and Amit Singer. Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *Mathematical Programming*, 163:145–167, 2017. doi: 10.1007/s10107-016-1059-6.

Cristian Bodnar, Francesco Di Giovanni, Benjamin Paul Chamberlain, Pietro Liò, and Michael M. Bronstein. Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs. In *Advances in Neural Information Processing Systems*, volume 35, pp. 18527–18541, 2022. doi: 10.52202/068431-1346.

Nicolas Boumal, Vladislav Voroninski, and Afonso S. Bandeira. The non-convex Burer–Monteiro approach works on smooth semidefinite programs. In *Advances in Neural Information Processing Systems*, volume 29, pp. 2757–2765, 2016.

Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1997. doi: 10.1090/cbms/092.

Justin Curry, Robert Ghrist, and Vidit Nanda. Discrete Morse theory for computing cellular sheaf cohomology. *Foundations of Computational Mathematics*, 16(4):875–897, 2016. doi: 10.1007/s10208-015-9266-8.

Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

Anders Eriksson, Carl Olsson, Fredrik Kahl, and Tat-Jun Chin. Rotation averaging and strong duality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 127–135, 2018. doi: 10.1109/CVPR.2018.00021.

Robin Forman. Bochner's method for cell complexes and combinatorial Ricci curvature. *Discrete & Computational Geometry*, 29(3):323–374, 2003. doi: 10.1007/s00454-002-0743-x.

David Gamarnik. The overlap gap property: A topological barrier to optimizing over random structures. *Proceedings of the National Academy of Sciences*, 118(41):e2108492118, 2021. doi: 10.1073/pnas.2108492118.

Paul J. Goulart and Yuwen Chen. Clarabel: An interior-point solver for conic programs with quadratic objectives. *Mathematical Programming Computation*, 2026. doi: 10.1007/s12532-026-00320-7. arXiv:2405.12762.

Jakob Hansen and Robert Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3(4):315–358, 2019. doi: 10.1007/s41468-019-00038-7.

Jakob Hansen and Robert Ghrist. Opinion dynamics on discourse sheaves. *SIAM Journal on Applied Mathematics*, 81(5):2033–2060, 2021. doi: 10.1137/20M1341088.

Frank Harary. On the notion of balance of a signed graph. *Michigan Mathematical Journal*, 2(2): 143–146, 1953. doi: 10.1307/mmj/1028989917.

Richard Hartley, Jochen Trumpf, Yuchao Dai, and Hongdong Li. Rotation averaging. *International Journal of Computer Vision*, 103(3):267–305, 2013. doi: 10.1007/s11263-012-0601-0.

Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *Journal of the ACM*, 46(6):858–923, 1999. doi: 10.1145/331524.331529.

Qi-Xing Huang and Leonidas Guibas. Consistent shape maps via semidefinite programming. *Computer Graphics Forum*, 32(5):177–186, 2013. doi: 10.1111/cgf.12184.

Qi Huangfu and J. A. Julian Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1):119–142, 2018. doi: 10.1007/s12532-017-0130-5.

Stuart Kauffman and Simon Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45, 1987. doi: 10.1016/S0022-5193(87)80029-2.

Carsten Lange, Shiping Liu, Norbert Peyerimhoff, and Olaf Post. Frustration index and Cheeger inequalities for discrete and continuous magnetic Laplacians. *Calculus of Variations and Partial Differential Equations*, 54(4):4165–4196, 2015. doi: 10.1007/s00526-015-0935-x.

Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992. doi: 10.1137/0221015.

N. S. Manton. An alternative action for lattice gauge theories. *Physics Letters B*, 96(3–4):328–330, 1980. doi: 10.1016/0370-2693(80)90778-9.

Gabriela Ochoa, Marco Tomassini, Sébastien Vérel, and Christian Darabos. A study of NK landscapes' basins and local optima networks. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pp. 555–562, 2008. doi: 10.1145/1389095.1389204.

Yann Ollivier. Ricci curvature of Markov chains on metric spaces. *Journal of Functional Analysis*, 256(3):810–864, 2009. doi: 10.1016/j.jfa.2008.11.001.

David M. Rosen, Luca Carlone, Afonso S. Bandeira, and John J. Leonard. SE-Sync: A certifiably correct algorithm for synchronization over the special Euclidean group. *The International Journal of Robotics Research*, 38(2–3):95–125, 2019. doi: 10.1177/0278364918784361.

Siegfried M. Rump. Verification methods: Rigorous results using floating-point arithmetic. *Acta Numerica*, 19:287–449, 2010. doi: 10.1017/S096249291000005X.

James Saunderson, Pablo A. Parrilo, and Alan S. Willsky. Semidefinite descriptions of the convex hull of rotation matrices. *SIAM Journal on Optimization*, 25(3):1314–1343, 2015. doi: 10.1137/14096339X.

Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011. doi: 10.1016/j.acha.2010.02.001.

Peter F. Stadler. Landscapes and their correlation functions. *Journal of Mathematical Chemistry*, 20 (1):1–45, 1996. doi: 10.1007/BF01165154.

Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. In *International Conference on Learning Representations (ICLR)*, 2022.

Lanhui Wang and Amit Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference: A Journal of the IMA*, 2(2):145–193, 2013. doi: 10.1093/imaiai/iat005.

Edward Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63(5):325–336, 1990. doi: 10.1007/BF00202749.

Kenneth G. Wilson. Confinement of quarks. *Physical Review D*, 10(8):2445–2459, 1974. doi: 10.1103/PhysRevD.10.2445.

Christopher Zach, Manfred Klopschitz, and Marc Pollefeys. Disambiguating visual relations using loop constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1426–1433, 2010. doi: 10.1109/CVPR.2010.5539801.

Thomas Zaslavsky. Biased graphs. I. Bias, balance, and gains. *Journal of Combinatorial Theory, Series B*, 47(1):32–52, 1989. doi: 10.1016/0095-8956(89)90063-4.