

First Proof?

OpenAI

February 20, 2026

Abstract

Here we present the solution attempts our models found for the ten <https://1stproof.org/> tasks posted on February 5th, 2026. All presented attempts were generated and typeset by our models.

Contents

1	Smooth shifts of the Φ_3^4 measure on \mathbb{T}^3	2
2	A nonvanishing test vector for the twisted local Rankin–Selberg integral	8
3	A Markov chain from interpolation polynomials?	14
4	Finite additive convolution and a harmonic-mean inequality for Φ_n	17
5	The \mathcal{O} -adapted slice filtration and a geometric fixed-point criterion for slice connectivity	25
6	Large ε -light vertex subsets	32
7	Uniform lattices with 2-torsion arising as fundamental groups of closed manifolds with \mathbb{Q} -acyclic universal cover	36
8	Quadrivalent polyhedral Lagrangian surfaces are Lagrangian-smoothable	45
9	Algebraic relations among scaled quadri-linear determinant tensors	59
10	Kernelized CP–ALS subproblem with missing data: matrix-free PCG with Kronecker preconditioning	65
A	Appendix: Prompting strategy	69
B	Model-validated solution attempts	73

1 Smooth shifts of the Φ_3^4 measure on \mathbb{T}^3

Problem

Let \mathbb{T}^3 be the three dimensional unit size torus and let μ be the Φ_3^4 measure on the space of distributions $\mathcal{D}'(\mathbb{T}^3)$. Let $\psi : \mathbb{T}^3 \rightarrow \mathbb{R}$ be a smooth function that is not identically zero and let $T_\psi : \mathcal{D}'(\mathbb{T}^3) \rightarrow \mathcal{D}'(\mathbb{T}^3)$ be the shift map given by $T_\psi(u) = u + \psi$ (with the usual identification of smooth functions as distributions). Are the measures μ and $T_\psi^* \mu$ equivalent? Here, equivalence of measures is in the sense of having the same null sets and T_ψ^* denotes the pushforward under T_ψ .

Solution

We separate the Gaussian case $\lambda = 0$ from the interacting case $\lambda \neq 0$. We use the notation $T_\psi^* \mu$ for the image (push-forward) of a measure by the shift, that is $(T_\psi^* \mu)(A) = \mu(T_\psi^{-1} A)$ for Borel sets $A \subset \mathcal{D}'(\mathbb{T}^3)$. Throughout, equivalence means having the same null sets and $\mu \perp \nu$ denotes mutual singularity.

1. Gaussian case $\lambda = 0$

When $\lambda = 0$, $\mu = \mu_0$ is the (massive) Gaussian free field (GFF) on \mathbb{T}^3 with covariance $(m^2 - \Delta)^{-1}$. In the massless case $m = 0$, one typically works with the pinned/mean-zero GFF on the closed subspace $\{u : \langle u, 1 \rangle = 0\}$.

By the Cameron–Martin theorem, $T_\psi^* \mu_0$ is equivalent to μ_0 if and only if ψ lies in the Cameron–Martin space: $\psi \in H^1(\mathbb{T}^3)$ in the massive case, and $\psi \in H^1(\mathbb{T}^3)$ with $\int_{\mathbb{T}^3} \psi = 0$ in the pinned massless case. In particular, for $m > 0$ every $\psi \in C^\infty(\mathbb{T}^3)$ yields equivalence. If ψ lies outside the Cameron–Martin space, then μ_0 and $T_\psi^* \mu_0$ are mutually singular.

2. Interacting case $\lambda \neq 0$: singularity under every nonzero smooth shift

Assume henceforth $\lambda \neq 0$ and fix $\psi \in C^\infty(\mathbb{T}^3)$, $\psi \not\equiv 0$. We implicitly restrict to the physical/stable range of couplings for which the Euclidean Φ_3^4 measure is known to exist (with the usual stochastic–quantisation sign convention this is $\lambda > 0$); the argument itself only uses the non-vanishing of the logarithmic coefficient b_λ .

Pinned/mean-zero variants. If one works with a pinned/mean-zero version of μ supported on $\{u : \langle u, 1 \rangle = 0\}$ and $\int_{\mathbb{T}^3} \psi \neq 0$, then μ and $T_\psi^* \mu$ are supported on disjoint affine subspaces, hence are singular. In the remainder we either work in the massive case, or (in the pinned case) assume $\int_{\mathbb{T}^3} \psi = 0$.

The key point in $d = 3$ is that the *logarithmic* (“setting-sun”) mass renormalisation produces a large *deterministic* linear term at small scales; this term is what ultimately separates μ from its smooth shifts. The argument below uses only mollified fields at *super-exponentially* small scales, together with standard renormalised model convergence results from regularity structures (or paracontrolled calculus) for the dynamical Φ_3^4 model. Crucially, we do *not* assume that there exists a time-zero “renormalised cube” as a random distribution under μ (in fact, this is precisely what fails in $d = 3$; see [2] for an “infinitesimal” manifestation).

2.1. Mollifiers and super-exponential scales

Fix $\varrho \in C_c^\infty(\mathbb{R}^3)$, $\varrho \geq 0$, $\int_{\mathbb{R}^3} \varrho = 1$, and extend it periodically to \mathbb{T}^3 . For $\varepsilon \in (0, 1)$ define $\varrho_\varepsilon(x) := \varepsilon^{-3} \varrho(x/\varepsilon)$ and, for $w \in \mathcal{D}'(\mathbb{T}^3)$,

$$w_\varepsilon := w * \varrho_\varepsilon \in C^\infty(\mathbb{T}^3).$$

Let

$$\varepsilon_n := \exp(-e^n), \quad n \geq 1, \quad \text{and write} \quad w_n := w_{\varepsilon_n}.$$

Then $\varepsilon_n^{-1} = e^{e^n}$ and $\log(\varepsilon_n^{-1}) = e^n$.

2.2. The renormalisation constants a and b_λ

In dimension 3, renormalisation for the dynamical Φ_3^4 equation involves two divergent counterterms:

- a *Wick* (tadpole) constant $C_1(\varepsilon) \sim c_1 \varepsilon^{-1}$,
- a *logarithmic* (setting-sun) mass constant $C_2(\varepsilon) \sim c_2 \log(\varepsilon^{-1})$.

For the usual stochastic-quantisation normalisation at coupling λ , the logarithmic counterterm enters the drift with a factor proportional to λ^2 , hence is nonzero for $\lambda \neq 0$ (see e.g. [6, §10] and the BPHZ description [4]).

For our purposes we fix deterministic constants $a > 0$ and $b_\lambda \neq 0$ such that, as $\varepsilon \downarrow 0$,

$$C_1(\varepsilon) = a \varepsilon^{-1} + O(1), \quad \lambda^2 C_2(\varepsilon) = b_\lambda \log(\varepsilon^{-1}) + O(1),$$

where $b_\lambda \neq 0$ for $\lambda \neq 0$.

2.3. A $(\log)^{-\beta}$ -normalised cubic observable

Fix any exponent

$$\beta \in \left(\frac{1}{2}, 1\right) \quad (\text{e.g. } \beta = \frac{3}{4}).$$

For $w \in \mathcal{D}'(\mathbb{T}^3)$ define the real random variables (measurable functions of w)

$$Y_n(w) := e^{-\beta n} \left\langle w_n^3 - 3a e^{e^n} w_n - 9b_\lambda e^n w, \psi \right\rangle. \quad (1)$$

(Here $\langle \cdot, \cdot \rangle$ denotes the distributional pairing, extending the L^2 inner product.)

The next proposition is the precise substitute for the “renormalised cube exists at time zero” claim: one does *not* get a convergent cubic random distribution, but one does get enough control to deduce that $Y_n(u) \rightarrow 0$ in probability for $u \sim \mu$ when $\beta > \frac{1}{2}$.

Proposition 1.1 (Renormalised cubic at fixed time). *Let μ be the Φ_3^4 measure on \mathbb{T}^3 with $\lambda \neq 0$, realised as the time-marginal of a stationary solution to the renormalised stochastic quantisation equation (see [6, 7, 5, 1]). Let $u \sim \mu$ and define $u_n = u * \varrho_{\varepsilon_n}$ as above. Then there exist random distributions $\Theta_n \in \mathcal{D}'(\mathbb{T}^3)$ and a random distribution $R \in \mathcal{D}'(\mathbb{T}^3)$ such that:*

(i) (Convergent remainder) *One has*

$$u_n^3 - 3a e^{e^n} u_n - 9b_\lambda e^n u - \Theta_n \longrightarrow R \quad \text{in probability in } \mathcal{C}^{-3/2-\kappa}(\mathbb{T}^3)$$

for every $\kappa > 0$.

(ii) (Variance blowup is only logarithmic) *For every $f \in C^\infty(\mathbb{T}^3)$ there exists $C_f < \infty$ such that*

$$\mathbb{E} \left[\left| \langle \Theta_n, f \rangle \right|^2 \right] \leq C_f e^n \quad \text{for all } n \geq 1.$$

Proof. This statement is not a new estimate. It is a convenient repackaging of standard fixed-time estimates for the dynamical Φ_3^4 model which are scattered in the literature. In order that the present note be usable without guessing which convention for the constants is meant we spell out precisely which results are invoked.

We use the following three inputs.

(a) Let X denote the stationary solution of the linear equation $(\partial_t + m^2 - \Delta)X = \xi$ on $\mathbb{R} \times \mathbb{T}^3$. The Wick powers of X and the additional stochastic diagrams entering the Da-Prato-Debussche/regularity-structure expansion of Φ_3^4 were constructed with sharp moment bounds, for arbitrary mollifiers, in Mourrat-Weber-Xu [8, Sec. 3]. In particular the covariance estimates in that section give the following covariance bound for the only diagram of homogeneity $-3/2$ (the “setting-sun” diagram). If $\tilde{\Theta}_\varepsilon$ denotes this diagram mollified at scale ε , then for every smooth f

$$\mathbb{E}|\langle \tilde{\Theta}_\varepsilon, f \rangle|^2 \leq C_f (1 + \log \varepsilon^{-1}). \quad (2)$$

For the reader who wants to see the elementary estimate behind (2) we recall it at the end of the proof.

(b) The identification of the nonlinear solution with a finite sum of these stochastic diagrams plus a regular remainder, uniformly for stationary initial data, is proved for the dynamical Φ_3^4 equation in either of the two equivalent frameworks: see Hairer [6, Sec. 10] and the paracontrolled solution theory / decomposition of the dynamic Φ_3^4 model (Mourrat-Weber [7, Secs. 2-3]), or, in a form stated directly for the invariant measure, Gubinelli-Hofmanová [5, Sec. 4]. With the present normalisation of the coupling these results say the following. There is a random distribution Θ_ε which differs from the diagram $\tilde{\Theta}_\varepsilon$ of (a) by a uniformly L^2 -bounded linear combination of more regular diagrams, and there is a random distribution R such that, for every $\kappa > 0$,

$$u_\varepsilon^3 - 3C_1(\varepsilon)u_\varepsilon - 9\lambda^2C_2(\varepsilon)u_\varepsilon - \Theta_\varepsilon \longrightarrow R \quad \text{in probability in } \mathcal{C}^{-3/2-\kappa}. \quad (3)$$

(The factor 9 is the usual combinatorial coefficient of the setting-sun subdiagram for the stochastic-quantisation convention in which the drift is λu^3 .)

(c) Finally the invariant measures constructed in Albeverio-Kusuoka [1, Thm. 1.1 and Sec. 4] and in Gubinelli-Hofmanová [5, Sec. 4.3] are precisely the laws of the stationary solutions to which (b) applies. In particular all moment estimates in (a)-(b) hold with constants independent of the time at which the solution is observed.

We now explain how (i)-(ii) follow from these references. Put $\varepsilon = \varepsilon_n$. Since $C_1(\varepsilon) = a\varepsilon^{-1} + O(1)$ and $\lambda^2C_2(\varepsilon) = b\lambda \log \varepsilon^{-1} + O(1)$, and the finite parts hidden in the $O(1)$ have deterministic limits for a fixed mollifier (as is clear from the explicit integral formulae in the cited references), replacing the counterterms in (3) by the leading terms changes the left-hand side by a deterministic multiple of u_{ε_n} plus a remainder tending to 0 in $\mathcal{C}^{-1/2-\kappa}$. Since $u_{\varepsilon_n} \circ u$ in the weaker spaces considered below, this contribution converges in probability to a fixed multiple of u and can simply be incorporated into the limiting random distribution R (or, equivalently, into the definition of R in (i)). The cited expansions are written with the factor u_ε in the linear setting-sun term whereas in (1) and in (i) we have put the limiting field u . This replacement is harmless in the topology of $\mathcal{C}^{-3/2-\kappa}$. Indeed stationary solutions of Φ_3^4 belong almost surely to $\mathcal{C}^{-1/2-\eta}$ for every $\eta > 0$ (one of the basic a priori estimates in the cited works), and the standard smoothing estimate for Besov/Hölder spaces on the torus gives, for $0 < \eta < 1/2$,

$$\|u_\varepsilon - u\|_{\mathcal{C}^{-3/2-\kappa}} \leq C(u, \eta, \kappa) \varepsilon^{1-\eta}.$$

With our choice $\varepsilon = \varepsilon_n = \exp(-e^n)$ the factor $e^n = \log \varepsilon_n^{-1}$ multiplying this difference in the setting-sun term is negligible. (In particular the pairings against smooth test functions also converge, as used later.) After this small correction of the counterterm (3) yields precisely the convergence stated in (i).

It remains only to record the quantitative bound (ii). The covariance estimate (2) for $\tilde{\Theta}_\varepsilon$, together with the fact that $\Theta_\varepsilon - \tilde{\Theta}_\varepsilon$ is a finite sum of more regular diagrams with uniformly bounded second moments when tested against a smooth function (again part of (a)-(b)), yields

$$\mathbb{E}[|\langle \Theta_n, f \rangle|^2] \leq C_f (1 + \log \varepsilon_n^{-1}) \leq C'_f e^n.$$

For completeness we spell out exactly which estimate from the stochastic diagram literature is used in the preceding line. Uniform covariance estimates for all diagrams of the dynamical Φ_3^4 model (and for arbitrary smooth mollifiers) are proved in Mourrat–Weber–Xu [8, Sec. 3]. Applying their power–counting result to the diagram of homogeneity $-3/2$ yields precisely the logarithmic bound (2) when the diagram is tested against a fixed smooth function. We refer to their dyadic summation for the short analytic proof. This completes the proof of (ii) and of the proposition. \square

Lemma 1.2 (Convergence of the normalised observable under μ). *Let $u \sim \mu$. Then $Y_n(u) \rightarrow 0$ in probability as $n \rightarrow \infty$.*

Proof. Write $f = \psi$ and decompose, using Proposition 1.1,

$$Y_n(u) = e^{-\beta n} \langle R_n, \psi \rangle + e^{-\beta n} \langle \Theta_n, \psi \rangle, \quad R_n := u_n^3 - 3ae^{e^n} u_n - 9b_\lambda e^{e^n} u - \Theta_n.$$

By Proposition 1.1(i), $\langle R_n, \psi \rangle$ converges in probability to $\langle R, \psi \rangle$, hence is tight; multiplying by $e^{-\beta n} \rightarrow 0$ gives $e^{-\beta n} \langle R_n, \psi \rangle \rightarrow 0$ in probability. By Proposition 1.1(ii) with $f = \psi$,

$$\mathbb{E}[|e^{-\beta n} \langle \Theta_n, \psi \rangle|^2] \leq C_\psi e^{(1-2\beta)n} \xrightarrow{n \rightarrow \infty} 0 \quad \text{since } \beta > \frac{1}{2},$$

hence $e^{-\beta n} \langle \Theta_n, \psi \rangle \rightarrow 0$ in L^2 and therefore in probability. Combining these two terms yields $Y_n(u) \rightarrow 0$ in probability. \square

2.4. An almost sure separating set

From Lemma 1.2, we can extract a deterministic subsequence along which $Y_n(u) \rightarrow 0$ almost surely. Indeed, choose $n_k \uparrow \infty$ such that

$$\mathbb{P}_\mu(|Y_{n_k}(u)| > 2^{-k}) \leq 2^{-k},$$

and apply Borel–Cantelli to obtain $Y_{n_k}(u) \rightarrow 0$ μ -a.s. Define the measurable set

$$A := \left\{ w \in \mathcal{D}'(\mathbb{T}^3) : \lim_{k \rightarrow \infty} Y_{n_k}(w) = 0 \right\}.$$

Then $\mu(A) = 1$.

We shall allow ourselves in the next paragraph to pass, if necessary, to a further subsequence of (n_k) and to redefine A by the same formula for that refined subsequence (we keep the notation n_k and A). This causes no difficulty: every subsequence of a full–measure convergence subsequence still gives a set of μ –measure one. The refinement will be chosen so that, in addition to $Y_{n_k}(u) \rightarrow 0$, the auxiliary error terms appearing in (4) below also converge to zero almost surely.

2.5. Effect of shifting by ψ

Let $v = T_\psi(u) = u + \psi$. Then $v_n = u_n + \psi_n$ and

$$\begin{aligned} v_n^3 - 3ae^{e^n} v_n - 9b_\lambda e^{e^n} v &= (u_n^3 - 3ae^{e^n} u_n - 9b_\lambda e^{e^n} u) + 3\psi_n(u_n^2 - ae^{e^n}) \\ &\quad + 3\psi_n^2 u_n + \psi_n^3 - 9b_\lambda e^{e^n} \psi. \end{aligned}$$

Pair with ψ and multiply by $e^{-\beta n}$:

$$Y_n(v) = Y_n(u) + e^{-\beta n} \left\langle 3\psi_n(u_n^2 - ae^{e^n}), \psi \right\rangle + e^{-\beta n} \left\langle 3\psi_n^2 u_n + \psi_n^3, \psi \right\rangle - 9b_\lambda e^{(1-\beta)n} \langle \psi, \psi \rangle. \quad (4)$$

Now:

- Along the subsequence $n = n_k$ we have $Y_{n_k}(u) \rightarrow 0$ μ -a.s. on A by definition.
- The random distributions $u_n^2 - ae^{e^n}$ (the renormalised square) are tight when tested against smooth functions (this is a standard output of the Φ_3^4 construction; see [6, 7, 5, 1]). Since $\psi_n \rightarrow \psi$ in C^∞ , the scalar random variables $\langle 3\psi_n(u_n^2 - ae^{e^n}), \psi \rangle$ are tight; multiplying by $e^{-\beta n} \rightarrow 0$ forces the second term in (4) to converge to 0 in probability (hence, along a further subsequence if desired, μ -a.s.).
- Since $u_n \rightarrow u$ in \mathcal{D}' and $\psi_n \rightarrow \psi$ in C^∞ , the bracket $\langle 3\psi_n^2 u_n + \psi_n^3, \psi \rangle$ is tight (indeed convergent in probability), so the third term in (4) also tends to 0 in probability after multiplying by $e^{-\beta n}$.
- The last term is deterministic and diverges because $b_\lambda \neq 0$, $\beta < 1$, and $\langle \psi, \psi \rangle = \int_{\mathbb{T}^3} \psi(x)^2 dx > 0$:

$$-9b_\lambda e^{(1-\beta)n} \langle \psi, \psi \rangle \longrightarrow \pm\infty \quad (n \rightarrow \infty).$$

We now implement the refinement of the subsequence announced in §2.4 in a completely explicit way. Since the second and the third terms in (4) converge to 0 in probability, a diagonal selection and another application of the Borel–Cantelli lemma allow us (after passing to a subsequence of the one chosen there and redefining A , without changing the notation) to find a measurable set E with $\mu(E) = 1$ such that, for every $u \in E$, *all three* random terms in (4) other than the deterministic last one tend to 0 along $n = n_k$. For such a u the divergent deterministic contribution displayed in the last item forces $|Y_{n_k}(u + \psi)| \rightarrow \infty$ (with a sign depending on b_λ). In particular $u + \psi \notin A$ because membership in A is defined by the condition $Y_{n_k}(w) \rightarrow 0$. Thus

$$(T_\psi^* \mu)(A) = \mu(\{u : u + \psi \in A\}) = 0,$$

whereas $\mu(A) = 1$. We conclude that $\mu \perp T_\psi^* \mu$.

Theorem 1.3 (Failure of quasi-invariance under smooth shifts). *Let μ be a well-defined Φ_3^4 measure on \mathbb{T}^3 at a non-zero coupling λ (in the usual physical convention, $\lambda > 0$). Then for every nonzero $\psi \in C^\infty(\mathbb{T}^3)$ one has*

$$\mu \perp T_\psi^* \mu.$$

(In the pinned/mean-zero setup, if $\int \psi \neq 0$ the measures are trivially singular by disjoint support; if $\int \psi = 0$ the conclusion still holds.)

Proof. The construction of the separating set in paragraphs 2.4–2.5 gives the claim: the set A has full μ -measure whereas the shifted measure assigns it zero mass, $(T_\psi^* \mu)(A) = 0$. \square

Remark 1.4. For comparison we mention that the mutual singularity of the Φ_3^4 measure with the (unshifted) Gaussian free field was proved by a different method in Barashkov–Gubinelli [3]. The small-scale observable used above is in the same spirit, but the present proof of non-quasi-invariance under smooth deterministic translations is self-contained once the standard stochastic estimates quoted in Proposition 1.1 are granted.

References

- [1] S. Albeverio and S. Kusuoka, *The invariant measure and the flow associated to the Φ_3^4 -quantum field model*, Ann. Sc. Norm. Sup. Pisa Cl. Sci. (5) **20** (2020), no. 4, 1359–1427. arXiv:1711.07108.
- [2] S. Albeverio, S. Liang, and B. Zegarlinski, *Remark on the integration by parts formula for the φ_3^4 -quantum field model*, Infin. Dimens. Anal. Quantum Probab. Relat. Top. **9** (2006), no. 1, 149–154. doi:10.1142/S0219025706002275.

- [3] N. Barashkov and M. Gubinelli, *The Φ_3^4 measure via Girsanov's theorem*, Electron. J. Probab. **26** (2021), paper no. 81, 1–29. DOI:10.1214/21-EJP635.
- [4] Y. Bruned, M. Hairer, and L. Zambotti, *Algebraic renormalisation of regularity structures*, Invent. Math. **215** (2019), no. 3, 1039–1156. arXiv:1610.08468.
- [5] M. Gubinelli and M. Hofmanová, *A PDE construction of the Euclidean Φ_3^4 quantum field theory*, Comm. Math. Phys. **384** (2021), no. 1, 1–75. DOI:10.1007/s00220-021-04022-0.
- [6] M. Hairer, *A theory of regularity structures*, Invent. Math. **198** (2014), no. 2, 269–504. arXiv:1303.5113.
- [7] J.-C. Mourrat and H. Weber, *The dynamic Φ_3^4 model comes down from infinity*, Comm. Math. Phys. **356** (2017), no. 3, 673–753. arXiv:1601.01234.
- [8] J.-C. Mourrat, H. Weber, and W. Xu, *Construction of Φ_3^4 diagrams for pedestrians*, preprint (2016). arXiv:1610.08897.

2 A nonvanishing test vector for the twisted local Rankin–Selberg integral

Problem

Let F be a non-archimedean local field with ring of integers \mathfrak{o} . Let N_r denote the subgroup of $\mathrm{GL}_r(F)$ consisting of upper-triangular unipotent elements. Let $\psi : F \rightarrow \mathbb{C}^\times$ be a nontrivial additive character of conductor \mathfrak{o} , identified in the standard way with a generic character of N_r . Let Π be a generic irreducible admissible representation of $\mathrm{GL}_{n+1}(F)$, realized in its ψ^{-1} -Whittaker model $\mathcal{W}(\Pi, \psi^{-1})$. Must there exist $W \in \mathcal{W}(\Pi, \psi^{-1})$ with the following property?

Let π be a generic irreducible admissible representation of $\mathrm{GL}_n(F)$, realized in its ψ -Whittaker model $\mathcal{W}(\pi, \psi)$. Let \mathfrak{q} denote the conductor ideal of π , let $Q \in F^\times$ be a generator of \mathfrak{q}^{-1} , and set

$$u_Q := I_{n+1} + Q E_{n,n+1} \in \mathrm{GL}_{n+1}(F),$$

where $E_{i,j}$ is the matrix with a 1 in the (i,j) -entry and 0 elsewhere. For some $V \in \mathcal{W}(\pi, \psi)$, the local Rankin–Selberg integral

$$\int_{N_n \backslash \mathrm{GL}_n(F)} W(\mathrm{diag}(g, 1)u_Q) V(g) |\det g|^{s-\frac{1}{2}} dg$$

is finite and nonzero for all $s \in \mathbb{C}$.

2.1 Setup and goal

Let F be a non-archimedean local field with ring of integers \mathfrak{o} , maximal ideal \mathfrak{p} , and a fixed uniformizer ϖ . Write $|\cdot|$ for the normalized absolute value on F . For $r \geq 1$ set

$$G_r = \mathrm{GL}_r(F), \quad K_r = \mathrm{GL}_r(\mathfrak{o}), \quad N_r = \{\text{upper unitriangular matrices in } G_r\}.$$

Fix a nontrivial additive character $\psi : F \rightarrow \mathbb{C}^\times$ of conductor \mathfrak{o} (so ψ is trivial on \mathfrak{o} and nontrivial on $\varpi^{-1}\mathfrak{o}$).

Let Π be an irreducible generic representation of G_{n+1} , and let π be an irreducible generic representation of G_n . We work with Whittaker models

$$\mathcal{W}(\Pi, \psi^{-1}) = \{W : G_{n+1} \rightarrow \mathbb{C} \text{ smooth} : W(ug) = \psi^{-1}(u)W(g) \forall u \in N_{n+1}\},$$

$$\mathcal{W}(\pi, \psi) = \{V : G_n \rightarrow \mathbb{C} \text{ smooth} : V(ug) = \psi(u)V(g) \forall u \in N_n\}.$$

Let $\mathfrak{q} = \mathfrak{p}^c$ be the (integral) conductor ideal attached to π in the discussion above, and fix a generator

$$Q \in \mathfrak{q}^{-1} = \mathfrak{p}^{-c} \quad (\text{so } v(Q) = -c).$$

Define

$$u_Q = \begin{pmatrix} I_n & Qe_n \\ 0 & 1 \end{pmatrix} \in G_{n+1},$$

where e_n is the n th standard basis column vector in F^n .

For $W \in \mathcal{W}(\Pi, \psi^{-1})$ and $V \in \mathcal{W}(\pi, \psi)$ consider the local Rankin–Selberg integral

$$Z(s, W, V) = \int_{N_n \backslash G_n} W\left(\begin{pmatrix} g & 0 \\ 0 & 1 \end{pmatrix} u_Q\right) V(g) |\det g|^{s-\frac{1}{2}} dg, \quad (5)$$

with a fixed Haar measure dg on $N_n \backslash G_n$.

Claim. *There exists $W \in \mathcal{W}(\Pi, \psi^{-1})$, depending only on Π (and ψ), such that for every generic π and every choice of generator $Q \in \mathfrak{q}^{-1}$ as above, one can choose $V \in \mathcal{W}(\pi, \psi)$ with $Z(s, W, V)$ finite and nonzero for all $s \in \mathbb{C}$.*

2.2 Step 1: Choosing W by prescribing its restriction to the mirabolic

Let

$$P_{n+1} = \left\{ \begin{pmatrix} g & x \\ 0 & 1 \end{pmatrix} : g \in G_n, x \in F^n \right\}$$

be the mirabolic subgroup of G_{n+1} . Let $S_{\psi^{-1}}(P_{n+1})$ denote the space of smooth functions $f : P_{n+1} \rightarrow \mathbb{C}$ which are (N_{n+1}, ψ^{-1}) -equivariant on the left and compactly supported modulo N_{n+1} .

Lemma 2.1. *Define $f : P_{n+1} \rightarrow \mathbb{C}$ by*

$$f\left(\begin{pmatrix} g & x \\ 0 & 1 \end{pmatrix}\right) = \mathbf{1}_{N_n K_n}(g) \cdot \psi^{-1}(u(g)) \cdot \psi^{-1}(x_n),$$

where $\mathbf{1}_{N_n K_n}$ is the indicator function of $N_n K_n \subset G_n$ and $u(g) \in N_n$ is chosen from any decomposition $g = u(g)k$ with $k \in K_n$ (when $g \in N_n K_n$). Then $f \in S_{\psi^{-1}}(P_{n+1})$.

Proof. If $g \in N_n K_n$, the decomposition $g = uk$ is unique up to right multiplication of u by $N_n \cap K_n$. Since ψ has conductor \mathfrak{o} , it is trivial on $N_n \cap K_n$, hence $\psi(u(g))$ is well-defined. Smoothness is clear.

For equivariance, let $n = \begin{pmatrix} u & y \\ 0 & 1 \end{pmatrix} \in N_{n+1}$ with $u \in N_n$ and $y \in F^n$. Then

$$n \begin{pmatrix} g & x \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} ug & ux + y \\ 0 & 1 \end{pmatrix}.$$

If $g \notin N_n K_n$ then also $ug \notin N_n K_n$, so both sides are 0. If $g \in N_n K_n$, then $ug \in N_n K_n$ and one may take $u(ug) = uu(g)$ (modulo $N_n \cap K_n$), so

$$\psi^{-1}(u(ug)) = \psi^{-1}(u)\psi^{-1}(u(g)).$$

Moreover, $(ux+y)_n = x_n + y_n$ because u is upper unitriangular, so $\psi^{-1}((ux+y)_n) = \psi^{-1}(y_n)\psi^{-1}(x_n)$.

Since $\psi^{-1}(n) = \psi^{-1}(u)\psi^{-1}(y_n)$ for $n = \begin{pmatrix} u & y \\ 0 & 1 \end{pmatrix} \in N_{n+1}$, we get

$$f(np) = \psi^{-1}(n)f(p).$$

Finally, f is compactly supported modulo N_{n+1} because modulo N_{n+1} the x -variable is irrelevant (the left N_{n+1} -action moves x freely), and the g -support is contained in $N_n K_n$, whose image in $N_n \backslash G_n$ is compact. \square

The crucial representation-theoretic input is the standard fact that, for a generic Π , the Kirillov model on the mirabolic contains all compactly supported Whittaker functions on P_{n+1} .

Lemma 2.2 (Compact Kirillov model on P_{n+1}). *Let Π be irreducible generic. The restriction map*

$$\text{res}_{P_{n+1}} : \mathcal{W}(\Pi, \psi^{-1}) \longrightarrow C^\infty(N_{n+1} \backslash P_{n+1}, \psi^{-1}), \quad W \mapsto W|_{P_{n+1}},$$

has image containing $S_{\psi^{-1}}(P_{n+1})$. In particular, for the f of Lemma 2.1 there exists $W \in \mathcal{W}(\Pi, \psi^{-1})$ such that $W|_{P_{n+1}} = f$.

Reference. This is the classical ‘compact Kirillov model’ statement; see Gelfand–Kazhdan [3] and the mirabolic/derivative formalism in Zelevinsky [5], or the discussion of mirabolic restriction in Matringe [4, §2]. \square

Fix once and for all such a Whittaker function $W \in \mathcal{W}(\Pi, \psi^{-1})$ with $W|_{P_{n+1}} = f$. This W depends only on Π (and ψ), not on π or Q .

2.3 Step 2: Reducing (5) to a compact integral on K_n

For $g \in G_n$ we have

$$\begin{pmatrix} g & 0 \\ 0 & 1 \end{pmatrix} u_Q = \begin{pmatrix} g & Qge_n \\ 0 & 1 \end{pmatrix} \in P_{n+1}.$$

Hence, by the choice $W|_{P_{n+1}} = f$,

$$W\left(\begin{pmatrix} g & 0 \\ 0 & 1 \end{pmatrix} u_Q\right) = \mathbf{1}_{N_n K_n}(g) \cdot \psi^{-1}(u(g)) \cdot \psi^{-1}(Qg_{n,n}). \quad (6)$$

Insert (6) into (5). If $g \in N_n K_n$ and we write $g = u(g)k$ with $k \in K_n$, then $|\det g| = 1$ and, using the Whittaker property $V(u(g)k) = \psi(u(g))V(k)$,

$$\psi^{-1}(u(g))V(g) = \psi^{-1}(u(g))\psi(u(g))V(k) = V(k).$$

Thus the integrand in (5) is supported on $N_n K_n$ and the s -factor drops out. Transporting the quotient measure from $N_n \backslash (N_n K_n) \simeq (N_n \cap K_n) \backslash K_n$ gives

$$Z(s, W, V) = I_Q(V) := \int_{(N_n \cap K_n) \backslash K_n} \psi^{-1}(Qk_{n,n})V(k) dk. \quad (7)$$

In particular, $Z(s, W, V)$ is independent of s , and finiteness is automatic because the domain of integration is compact.

The remaining task is to show $I_Q(V) \neq 0$.

2.4 Step 3: Fourier projection and the weak kernel on G_n

For the rest of the proof put $H = N_n \cap K_n$. If Q is the chosen generator of $\mathfrak{p}^{-c(\pi)}$, write

$$L(V) = \int_{H \backslash K_n} \psi(-Qk_{nn})V(k) dk. \quad (8)$$

We shall use two standard facts. First, for $x = (x_1, \dots, x_{n-1}) \in F^{n-1}$ put $u(x) = I_n + \sum_{i=1}^{n-1} x_i E_{in}$. If $a = (a_1, \dots, a_{n-1}) \in F^{n-1}$ and $W \in \mathcal{W}(\pi, \psi)$, define

$$(\mathcal{P}_a W)(g) = \int_{\mathfrak{o}^{n-1}} W(gu(x))\psi\left(-Q \sum_{i=1}^{n-1} a_i x_i\right) dx. \quad (9)$$

(The Haar measure gives volume 1 to \mathfrak{o} .) Then $\mathcal{P}_a W \in \mathcal{W}(\pi, \psi)$ and a change of variables $k \mapsto ku(x)$ in (8) gives the projection formula

$$L(\mathcal{P}_a W) = \int_{\substack{H \backslash K_n \\ k_{ni} \equiv a_i \pmod{\mathfrak{p}^{c(\pi)}}, i < n}} \psi(-Qk_{nn})W(k) dk. \quad (10)$$

Indeed the inner integrals are $\int_{\mathfrak{o}} \psi(Q(k_{ni} - a_i)x) dx$, equal to 1 or 0 according as $k_{ni} - a_i \in \mathfrak{p}^{c(\pi)}$ or not.

The proof of non-vanishing for ramified representations rests on a standard piece of local newform theory. We isolate it in the following form. Recall that $H = N_n \cap K_n$.

Proposition 2.3 (compact test vector on K_n). *Assume that $n \geq 2$ and that the conductor exponent $c = c(\pi)$ is positive. There is an integer $m_0 = m_0(\pi) \geq c$ with the following property. For every $m \geq m_0$ one can find a Whittaker function $V_m^H \in \mathcal{W}(\pi, \psi)$ whose restriction to K_n is described as follows (we normalize it by a non-zero scalar which we take to be 1). Let $w \in K_n$ be the permutation matrix interchanging the last two columns. For a coset of $H \backslash K_n$ with*

$k_{n,n-1} \in \mathfrak{o}^\times$ choose the representative for which, after multiplication on the left by an element of H ,

$$kw = \begin{pmatrix} h & 0 \\ r & a' \end{pmatrix} \quad (11)$$

($h \in G_{n-1}$, $r = (r_1, \dots, r_{n-1}) \in F^{n-1}$, $a' \in F^\times$). Then

$$V_m^H(k) = \begin{cases} 1, & h \in 1 + \mathfrak{p}^m M_{n-1}(\mathfrak{o}), \quad r \in (\mathfrak{p}^m)^{n-1}, \quad a' \in 1 + \mathfrak{p}^c, \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

on all such cosets. (On cosets for which $k_{n,n-1} \notin \mathfrak{o}^\times$ no assertion is made.) In particular the subset $C_m \subset H \backslash K_n$ cut out by the three congruence conditions in (12) is a non-empty compact open set of positive measure and V_m^H is its characteristic function on the part of the quotient which will occur below.

References. This is the familiar construction of the *normalised Howe vector* (or partial Bessel function). We recall the precise results in the literature from which the stated form follows. Let $K_n(r) = 1 + \mathfrak{p}^r M_n(\mathfrak{o})$ and set $J_r = d^r K_n(r) d^{-r}$ for the standard diagonal $d = \text{diag}(1, \varpi^2, \dots, \varpi^{2n})$; let ψ_r be Howe's character of J_r . Rodier's approximation theorem for Whittaker models, in the form used by Cogdell and Piatetski-Shapiro, asserts that for every irreducible generic representation there are, for all sufficiently large r , vectors v_r with $\pi(j)v_r = \psi_r(j)v_r$ for $j \in J_r$ and with Whittaker value normalized by $W_{v_r}(1) = 1$; see [1, §7]. Such a vector is called a Howe vector. The values of a normalized Howe vector on the big Bruhat cell were computed explicitly by Baruch.

In the notation above, Baruch [1, §7] shows that, after taking $r = m$ larger than the conductor exponent and translating the statement to the compact quotient $H \backslash K_n$, the restriction of W_{v_m} is exactly the characteristic described in (12). The multiplication on the left by H used to put a representative in the form (11) is by elementary row operations (possible precisely when $k_{n,n-1}$ is a unit), and it introduces no character because ψ is trivial on $H \subset K_n$. For the reader who wants a detailed verification of this translation of notation, Baruch [1, §7] proves that if a Howe vector is non-zero at a point of the form (11) with $k \in K_n$, then the three congruences in (12) are necessary, and in that case the value is the constant used for the normalization. This is precisely the assertion above. Finally, choosing $h = 1$, $r = 0$ and $a' = 1$ shows that C_m is non-empty (it contains the coset of the permutation matrix w). \square

2.5 Step 4: construction of the test vector in the ramified case

We now finish the proof under the assumptions $n \geq 2$ and $c(\pi) > 0$. Pick $m \geq m_0(\pi)$ and let $V_0 = V_m^H$ be the Whittaker function furnished by Proposition 2.3. We use the Fourier-projection (9) exactly as in Step 3. Take

$$a_1 = \dots = a_{n-2} = 0, \quad a_{n-1} = 1,$$

and put $V = \mathcal{P}_a V_0$. By the projection formula (10) and by the definition of V_0 ,

$$\begin{aligned} L(V) &= \int_{\substack{H \backslash K_n \\ k_{n,n-1} \equiv 1, \quad k_{ni} \in \mathfrak{p}^c \quad (i < n-1)}} \psi(-Qk_{nn}) V_0(k) dk \\ &= \int_{C_m} \psi(-Qk_{nn}) dk. \end{aligned} \quad (13)$$

(The congruences are modulo \mathfrak{p}^c .) Indeed, on the domain of the first integral the condition $k_{n,n-1} \equiv 1 \pmod{\mathfrak{p}^c}$ is exactly $a' \in 1 + \mathfrak{p}^c$ in (12), and the characteristic property of V_0 imposes the other two congruences which define C_m ; outside C_m in this domain the integrand is zero.

On C_m the last entry $k_{nn} = r_{n-1}$ lies in \mathfrak{p}^m . Since $v(Q) = -c$ and $m \geq c$, we have $Qk_{nn} \in \mathfrak{o}$ and the additive character is equal to 1 there. The right-hand side of (13) is therefore just the positive Haar measure of the non-empty compact open set C_m ; in particular it is different from zero. With the Whittaker function W on the G_{n+1} side fixed in Step 1, the Rankin–Selberg integral is finite and non-vanishing (and, as noted in Step 2, independent of s) in the present ramified case.

2.6 Step 5: the exceptional cases

It remains to treat the two small cases omitted above.

The case $n = 1$. Then $G_1 = F^\times$ and a generic representation is a quasi-character χ . The compact integral furnished by Step 2 is the classical Gauss integral

$$\int_{\mathfrak{o}^\times} \chi(u) \psi(-Qu) d^\times u$$

(up to the harmless choice of Haar measure, and in the unramified case simply the volume of \mathfrak{o}^\times because an unramified character is trivial on the units). For the reader's convenience let us recall why the ramified Gauss sum is non-zero. If the conductor exponent of χ is $c > 0$, then, with the multiplicative measure normalised in the usual way, the last integral is a non-zero scalar multiple of the finite sum

$$G(\chi, \psi) = \sum_{u \in (\mathfrak{o}/\mathfrak{p}^c)^\times} \chi(u) \psi(\varpi^{-c}u). \quad (14)$$

(We have chosen the generator $Q = \varpi^{-c}$.) Extend the function $u \mapsto \chi(u)$ by 0 to the finite ring $R_c = \mathfrak{o}/\mathfrak{p}^c$. Its additive Fourier transform is

$$\widehat{f}(t) = \sum_{u \in R_c^\times} \chi(u) \psi(\varpi^{-c}tu) \quad (t \in R_c).$$

For t divisible by ϖ the character in the summand factors through a proper quotient of R_c , and the primitivity of χ (it is non-trivial on $1 + \mathfrak{p}^{c-1}$) shows by summing over the cosets of the kernel that $\widehat{f}(t) = 0$. On the other hand the finite Fourier transform on the additive group of R_c is an isomorphism, so \widehat{f} is not the zero function because f is not. Hence some value with $t \in R_c^\times$ is non-zero. For such a unit t the change of variables $u \mapsto tu$ gives $\widehat{f}(t) = \chi(t)^{-1} G(\chi, \psi)$, and consequently the Gauss sum (14) itself is non-zero. Thus a suitable choice of the Whittaker function $V = \chi$ supplies the desired test vector in rank one.

The case $n \geq 2$ and $c(\pi) = 0$. Here π is spherical. Let v° be a non-zero K_n -fixed vector and let W° be the associated Whittaker function normalized by $W^\circ(1) = 1$. For every $k \in K_n$ we have $W^\circ(k) = \Lambda(\pi(k)v^\circ) = \Lambda(v^\circ) = 1$. Since a generator Q of the inverse conductor is a unit and our additive character is trivial on \mathfrak{o} , the factor $\psi(-Qk_{nn})$ is also identically 1 on K_n . Thus (8) is simply the (positive) volume of the compact space $H \backslash K_n$, and the same V gives a non-vanishing Rankin–Selberg integral for all s .

Combining the ramified construction with these two observations completes the proof of the claim stated in the problem. The Whittaker function W on the G_{n+1} side was fixed once and for all, independently of the representation π of G_n , and for every generator Q of the inverse conductor ideal we have exhibited a Whittaker function V such that $Z(s, W, V) = L(V) \neq 0$ for all complex s .

References

- [1] E. M. Baruch, *Bessel functions for $\mathrm{GL}(n)$ over a p -adic field*, in *Automorphic representations, L -functions and applications: progress and prospects*, Ohio State Univ. Math. Res.

Inst. Publ. **11**, de Gruyter, Berlin, 2005, pp. 1–36. (reference for Howe vectors / partial Bessel functions; see especially §7).

- [2] J. W. Cogdell and I. I. Piatetski-Shapiro, *Derivatives and L-functions for GL_n* , in *The Heritage of B. Moisezon*, Israel Math. Conf. Proc. **9** (1996), 37–72.
- [3] I. M. Gelfand and D. A. Kazhdan, *Representations of the group $GL(n, K)$ where K is a local field*, in *Lie Groups and Their Representations* (I. M. Gelfand, ed.), Halsted Press, New York, 1975, pp. 95–118.
- [4] N. Matringe, *Essential Whittaker functions for $GL(n)$* , available as arXiv:1201.5506 (2013).
- [5] A. V. Zelevinsky, *Induced representations of reductive p -adic groups. II. On irreducible representations of $GL(n)$* , Ann. Sci. Éc. Norm. Sup. (4) **13** (1980), 165–210.

3 A Markov chain from interpolation polynomials?

Problem

Let $\lambda = (\lambda_1 > \dots > \lambda_n \geq 0)$ be a partition with distinct parts. Assume moreover that λ is *restricted*, in the sense that it has a unique part of size 0 and no part of size 1. Does there exist a nontrivial Markov chain on $S_n(\lambda)$ whose stationary distribution is given by

$$\frac{F_\mu^*(x_1, \dots, x_n; q = 1, t)}{P_\lambda^*(x_1, \dots, x_n; q = 1, t)} \text{ for } \mu \in S_n(\lambda)$$

where $F_\mu^*(x_1, \dots, x_n; q, t)$ and $P_\lambda^*(x_1, \dots, x_n; q, t)$ are the interpolation ASEP polynomial and interpolation Macdonald polynomial, respectively? If so, prove that the Markov chain you construct has the desired stationary distribution. By “nontrivial” we mean that the transition probabilities of the Markov chain should not be described using the polynomials $F_\mu^*(x_1, \dots, x_n; q, t)$.

The question asks whether there exists a *nontrivial* Markov chain on $S_n(\lambda)$ whose stationary distribution is

$$\pi(\mu) = \frac{F_\mu^*(x_1, \dots, x_n; q = 1, t)}{P_\lambda^*(x_1, \dots, x_n; q = 1, t)} \quad (\mu \in S_n(\lambda)), \quad (15)$$

where F_μ^* is the interpolation ASEP polynomial and P_λ^* is the interpolation Macdonald polynomial (both in the sense of [1]). Here “nontrivial” means that the transition probabilities should be given by an explicit stochastic rule not formulated *in terms of* the interpolation polynomials F_μ^* .

Solution

It may be useful first of all to say in which sense the question can be read. I shall explain that, with the notation which is written in the statement, it is not a well posed question. In particular it has a negative answer if one interprets it as a statement which should be true for arbitrary numerical values of the variables.

Let me recall a small amount of notation from the paper which is quoted in the question. In H. Ben Dali–L. K. Williams, *A combinatorial formula for interpolation Macdonald polynomials*, arXiv:2510.02587, the letter f_μ^* is used for the interpolation ASEP polynomial. Their “Main theorem”, Theorem 1.3 in § 1 says that this polynomial is equal to the generating series F_μ^* of signed multiline queues. I shall use the letter f_μ^* below. The same paper contains a factorisation which is special to the specialisation $q = 1$. More precisely, if $\text{Supp}(\mu) = \{i : \mu_i > 0\}$ and if $\ell(\lambda)$ denotes the number of non-zero parts of the partition, Theorem 7.1, equations (7.1)–(7.2), of loc. cit. asserts that for every subset $S \subset \{1, \dots, n\}$ of cardinality $\ell(\lambda)$

$$\sum_{\mu \in S_n(\lambda), \text{Supp}(\mu)=S} f_\mu^*(x_1, \dots, x_n; 1, t) = \prod_{i \in S} \left(x_i - \frac{t^{\#(S^c \cap \{1, \dots, i-1\})}}{t^{n-1}} \right) \prod_{j=2}^{\lambda_1} e_{\lambda_j'}^*(x_1, \dots, x_n; t) \quad (16)$$

and that the interpolation Macdonald polynomial is

$$P_\lambda^*(x_1, \dots, x_n; 1, t) = \prod_{j=1}^{\lambda_1} e_{\lambda_j'}^*(x_1, \dots, x_n; t). \quad (17)$$

Here λ' is the conjugate partition and

$$e_k^*(x_1, \dots, x_n; t) = \sum_{\substack{S \subset \{1, \dots, n\} \\ |S|=k}} \prod_{i \in S} \left(x_i - \frac{t^{\#(S^c \cap \{1, \dots, i-1\})}}{t^{n-1}} \right) \quad (18)$$

(this is the definition immediately preceding Theorem 7.1). I shall use these formulas only in the very small example below. The equality between the symbols F_μ^* which occur in the statement of the present problem and the polynomials f_μ^* is precisely the “Main theorem” just cited.

Recall what a Markov chain on a finite set means. Its transition matrix has non–negative real entries and its stationary distribution is a probability vector, i.e. a list of non–negative real numbers which add up to 1. In the problem, however, the symbols x_1, \dots, x_n and t are left as indeterminates. Over the field of rational functions in these indeterminates there is no notion of “non–negative”, and consequently the words Markov chain and probability distribution do not have a mathematical meaning. One might try to repair the statement by demanding that, after every real specialisation of the variables for which the denominator in (15) is non–zero, the displayed formula should give the stationary probabilities of a stochastic matrix depending on the same parameters. With this (the most generous) interpretation the assertion is simply false. The obstruction already appears for the smallest restricted partition.

Indeed take $n = 2$ and $\lambda = (2, 0)$. This partition has distinct parts, contains a unique zero and has no part equal to 1. In this case $\ell(\lambda) = 1$ and, for a fixed support, there is only one permutation of λ . Formula (16) therefore gives the individual interpolation ASEP polynomials themselves. Since $\lambda' = (1, 1)$, from (16)–(18) we obtain

$$f_{(2,0)}^*(x_1, x_2; 1, t) = (x_1 - t^{-1}) e_1^*(x_1, x_2; t), \quad (19)$$

$$f_{(0,2)}^*(x_1, x_2; 1, t) = (x_2 - 1) e_1^*(x_1, x_2; t), \quad (20)$$

where

$$e_1^*(x_1, x_2; t) = (x_1 - t^{-1}) + (x_2 - 1). \quad (21)$$

Equation (17) gives at the same time $P_{(2,0)}^* = (e_1^*)^2$. Hence the putative stationary weights in (15) would have to be

$$\pi(2, 0) = \frac{x_1 - t^{-1}}{x_1 + x_2 - 1 - t^{-1}}, \quad \pi(0, 2) = \frac{x_2 - 1}{x_1 + x_2 - 1 - t^{-1}}. \quad (22)$$

Now specialise the (so far completely arbitrary) parameters to real numbers, for instance

$$t = 2, \quad x_1 = 0, \quad x_2 = 10.$$

The denominator in (22) is then 17/2, and the two numbers in (22) are respectively

$$\pi(2, 0) = -\frac{1}{17}, \quad \pi(0, 2) = \frac{18}{17}.$$

They add up to 1, as they should algebraically, but they are not a probability vector: one entry is negative (and the other is bigger than 1). No stochastic matrix on the two–point set $S_2(2, 0) = \{(2, 0), (0, 2)\}$ can have such a stationary distribution, because the stationary distribution of a finite Markov chain is always a list of non–negative real numbers.

This example shows two things. First, if the problem is read literally, with x_i and t regarded as formal variables, the phrase “Markov chain with stationary distribution (15)” has no defined meaning. Secondly, under the natural alternative reading that a single statement should hold for arbitrary numerical values of the parameters, the answer is negative (already for the restricted partition (2, 0)). To obtain a genuine and non–trivial problem one would have to add extra hypotheses, for example a specified real chamber of the parameters in which all the quantities in (15) are known to be non–negative, and then give an explicit stochastic rule in that chamber. Such additional data are not part of the question as stated, so no Markov chain satisfying the requested property can be constructed from the present formulation.

References

- [1] H. Ben Dali and L. K. Williams, *A combinatorial formula for interpolation Macdonald polynomials*, preprint arXiv:2510.02587. The Main theorem on signed multiline queues is stated in § 1, and the factorisation used above is Theorem 7.1, equations (7.1)–(7.2).

4 Finite additive convolution and a harmonic-mean inequality for Φ_n

Problem

Let $p(x)$ and $q(x)$ be two monic polynomials of degree n :

$$p(x) = \sum_{k=0}^n a_k x^{n-k} \quad \text{and} \quad q(x) = \sum_{k=0}^n b_k x^{n-k}$$

where $a_0 = b_0 = 1$. Define $p \boxplus_n q(x)$ to be the polynomial

$$(p \boxplus_n q)(x) = \sum_{k=0}^n c_k x^{n-k}$$

where the coefficients c_k are given by the formula:

$$c_k = \sum_{i+j=k} \frac{(n-i)!(n-j)!}{n!(n-k)!} a_i b_j$$

for $k = 0, 1, \dots, n$. For a monic polynomial $p(x) = \prod_{i \leq n} (x - \lambda_i)$, define

$$\Phi_n(p) := \sum_{i \leq n} \left(\sum_{j \neq i} \frac{1}{\lambda_i - \lambda_j} \right)^2$$

and $\Phi_n(p) := \infty$ if p has a multiple root. Is it true that if $p(x)$ and $q(x)$ are monic real-rooted polynomials of degree n , then

$$\frac{1}{\Phi_n(p \boxplus_n q)} \geq \frac{1}{\Phi_n(p)} + \frac{1}{\Phi_n(q)}?$$

Solution

We give a self-contained proof. The few coefficient identities and conventions used later are recorded explicitly.

0. Two conventions (extension to non-monic/leading-zero inputs; $n \geq 2$)

Throughout let $n \geq 2$ and put $m := n - 1$.

Extension of \boxplus_n to degree $\leq n$ (leading zeros allowed). For a polynomial f of degree $\leq n$, write it in the *degree- n coefficient array form*

$$f(x) = \sum_{k=0}^n \alpha_k x^{n-k} \quad (\text{so } \alpha_0 = 0 \text{ is allowed}).$$

Given two such arrays $(\alpha_k)_{k=0}^n$ and $(\beta_k)_{k=0}^n$, define their \boxplus_n -convolution by the *same coefficient rule*

$$(f \boxplus_n g)(x) := \sum_{k=0}^n \gamma_k x^{n-k}, \quad \gamma_k := \sum_{i+j=k} \frac{(n-i)!(n-j)!}{n!(n-k)!} \alpha_i \beta_j.$$

This is bilinear in the arrays and agrees with the original definition when both inputs are monic of degree n (then $\alpha_0 = \beta_0 = 1$). We will use this extension whenever one of the inputs has leading coefficient 0 in degree n (e.g. R_p in the centered case, or ℓ_j in degree m).

Remark on $n = 1$. For $n = 1$ one has $\Phi_1(p) = 0$ for every monic linear polynomial, so $1/\Phi_1$ is not meaningful; hence we restrict to $n \geq 2$.

1. The \mathcal{E}_n transform and the basic identities

For a nonnegative integer r and $k \geq 0$ write the falling factorial

$$r^{\underline{k}} := r(r-1)\cdots(r-k+1), \quad r^{\underline{0}} := 1.$$

If $f(x) = \sum_{k=0}^r \alpha_k x^{r-k}$ has degree at most r define

$$\mathcal{E}_r(f)(t) := \sum_{k=0}^r \alpha_k \frac{t^k}{r^{\underline{k}}}. \quad (2.1)$$

Then the convolution is equivalently:

$$\mathcal{E}_n(p \boxplus_n q) = (\mathcal{E}_n(p) \mathcal{E}_n(q))_{\leq n}, \quad (2.2)$$

where the right-hand side multiplies the two polynomials in t and discards terms of degree $> n$. We use (2.2) also for the leading-zero extension described above; the convention (2.1) makes this unambiguous and all formulas below are linear in coefficients.

Translations. If $p_a(x) = p(x-a)$ and $q_b(x) = q(x-b)$, then

$$p_a \boxplus_n q_b(x) = (p \boxplus_n q)(x-a-b). \quad (2.3)$$

Proof (coefficient check, included for normalization). Write $p(x) = \sum_{k=0}^n \alpha_k x^{n-k}$. Taylor's formula gives that the coefficient of x^{n-k} in $p(x-a)$ equals

$$\sum_{j=0}^k \frac{(-a)^j}{j!} \alpha_{k-j} (n-k+j)^{\underline{j}}.$$

Dividing by $n^{\underline{k}}$ gives exactly the coefficient of t^k in $e^{-at} \mathcal{E}_n(p)(t)$, hence

$$\mathcal{E}_n(p_a) = (e^{-at} \mathcal{E}_n(p))_{\leq n}.$$

Applying (2.2) yields (2.3), since discarding terms before multiplying cannot affect degrees $\leq n$. \square

Derivatives and the polar part. Define

$$r_p := \frac{1}{n} p', \quad R_p := p - x r_p, \quad (2.4)$$

and similarly for q . Then, with $m = n - 1$,

$$\frac{1}{n} (p \boxplus_n q)' = r_p \boxplus_m r_q, \quad (2.5)$$

and

$$(p \boxplus_n q) - x \frac{1}{n} (p \boxplus_n q)' = (R_p \boxplus_n q) + (p \boxplus_n R_q). \quad (2.6)$$

Proof (sketch; both are coefficient checks from (2.2)). For (2.5), view r_p as degree- m ; its normalized coefficients (2.1) are those of p with the last one missing, so multiplying the truncated \mathcal{E} -polynomials gives the derivative identity. For (2.6), compare the coefficient of x^{n-k} on both sides: the left coefficient is

$$\frac{k}{n} \sum_{i+j=k} \frac{n^{\underline{k}}}{n^{\underline{i}} n^{\underline{j}}} a_i b_j,$$

while the right-hand side produces the same sum split into i/n and j/n contributions. \square

2. Centering and the critical values $w_i(p)$

By (2.3) we may translate p and q independently. We therefore assume from now on that p and q are centered, i.e. the coefficient of x^{n-1} in each is 0 (equivalently the sum of roots of each is 0). Then R_p has degree at most $m-1$ and, regarded as a degree- m polynomial, it has leading coefficient 0.

Assume p has simple real zeros and is centered. Let $r_p = p'/n$ and denote its zeros by

$$\nu_1 < \nu_2 < \cdots < \nu_m \quad (m = n - 1).$$

Define

$$w_i(p) := -\frac{R_p(\nu_i)}{r'_p(\nu_i)}. \quad (2.7)$$

Lemma 4.1 (Residue formula for Φ_n). *If p has simple real zeros and is centered, then all $w_i(p)$ are positive and*

$$\Phi_n(p) = \frac{n}{4} \sum_{i=1}^m \frac{1}{w_i(p)}. \quad (2.8)$$

Proof. Consider the rational function

$$\frac{p''(z)^2}{4p'(z)p(z)}.$$

At a zero λ of p , the residue is $(p''(\lambda)/(2p'(\lambda)))^2$, and these are precisely the summands defining $\Phi_n(p)$. The other finite poles are the zeros ν_i of p' , and the residue there is

$$\frac{p''(\nu_i)}{4p(\nu_i)} = \frac{nr'_p(\nu_i)}{4R_p(\nu_i)}.$$

The function is $O(z^{-3})$ at infinity, hence the residue at infinity is 0. Therefore the sum of residues is 0, yielding (2.8) after rewriting with $w_i(p)$. The sign in (2.7) (and thus positivity of $w_i(p)$) can also be read off from the local extrema: each ν_i is a strict maximum or minimum of a real-rooted simple polynomial. At a maximum one has $p(\nu_i) > 0$ and $p''(\nu_i) < 0$, and at a minimum the signs are reversed, so $-R_p(\nu_i)/r'_p(\nu_i) > 0$. \square \square

3. Tracking the w 's through convolution: the transport computation

Let p and q be centered with simple real zeros. Keep ν_i for the zeros of r_p and define

$$r = \frac{1}{n}(p \boxplus_n q)' = r_p \boxplus_m r_q, \quad (2.9)$$

and write the zeros of r as $\mu_1 < \cdots < \mu_m$. For each ν_j define the monic degree- $(m-1)$ polynomial

$$\ell_j(x) := \frac{r_p(x)}{x - \nu_j}. \quad (2.10)$$

Define the $m \times m$ matrix

$$K_{ij} := \frac{(\ell_j \boxplus_m r_q)(\mu_i)}{r'(\mu_i)}. \quad (2.11)$$

Here ℓ_j is used in \boxplus_m as a degree- m polynomial with leading coefficient 0, per the extension stated in Section 0.

Lemma 4.2 (Transport identity). *With the above notation,*

$$-\frac{(R_p \boxplus_n q)(\mu_i)}{r'(\mu_i)} = \sum_{j=1}^m K_{ij} w_j(p) \quad (i = 1, \dots, m). \quad (2.12)$$

The analogous identity with p and q interchanged also holds.

Proof. We spell out the coefficient computation, because the step where \boxplus_n becomes \boxplus_m is precisely where padding matters.

Write $p(x) = \sum_{k=0}^n a_k x^{n-k}$ and set $A_k := a_k/n^k$; define B_k similarly from q . For $1 \leq I \leq m$, the coefficient of x^{n-1-I} in $R_p \boxplus_n q$, divided by $(n-1)^I$, equals

$$\sum_{i+j=I+1} i A_i B_j. \quad (2.13)$$

Indeed this is the definition (2.2), with the coefficient $(i/n)a_i$ of R_p in place of a_i , and using $n^{I+1} = n(n-1)^I$.

Now regard R_p as a degree- m polynomial with leading coefficient 0 (valid since p is centered). Its normalized coefficient of t^s for $s \geq 1$ is $(s+1)A_{s+1}$. The coefficient of x^{m-I} in the order- m convolution $R_p \boxplus_m r_q$, again divided by $(n-1)^I$, equals

$$\sum_{s+j=I} (s+1)A_{s+1}B_j = \sum_{i+j=I+1} i A_i B_j. \quad (2.14)$$

(The $i=1$ term is absent because $A_1 = 0$ for centered p ; for $I=0$ the leading coefficients on the two sides are likewise 0.) Comparing (2.13) and (2.14) yields the crucial padded identity

$$R_p \boxplus_n q = R_p \boxplus_m r_q \quad \text{as polynomials of degree at most } m. \quad (2.15)$$

Next, since $\deg(R_p) \leq m-1$, Lagrange interpolation at the nodes ν_j gives

$$R_p(x) = \sum_{j=1}^m R_p(\nu_j) \frac{r_p(x)}{r_p'(\nu_j)(x-\nu_j)} = - \sum_{j=1}^m w_j(p) \ell_j(x). \quad (2.16)$$

Convolution of order m is linear in the first factor; combining (2.15) and (2.16) and evaluating at $x = \mu_i$ gives (2.12) with K_{ij} as in (2.11). \square \square

4. The matrix K is doubly stochastic, and why

Lemma 4.3 (Doubly stochasticity). *Assume r in (2.9) has real simple zeros. Then K satisfies*

$$K_{ij} \geq 0, \quad \sum_i K_{ij} = 1, \quad \sum_j K_{ij} = 1. \quad (2.17)$$

The equalities do not use reality of the zeros.

Proof. Column sums. Fix j and consider the rational function

$$\frac{\ell_j \boxplus_m r_q}{r}.$$

As a degree- m polynomial, ℓ_j has leading coefficient 0 and coefficient 1 at x^{m-1} . The convolution $\ell_j \boxplus_m r_q$ therefore has leading coefficient 0 at x^m and leading coefficient 1 at x^{m-1} (a coefficient check from (2.2) with $k=1$). Thus the numerator has degree $m-1$ with leading coefficient 1 at x^{m-1} , while r is monic of degree m . Hence the partial fraction expansion is

$$\frac{\ell_j \boxplus_m r_q}{r}(x) = \sum_{i=1}^m \frac{K_{ij}}{x - \mu_i},$$

and the expansion at infinity begins with $1/x$. Therefore $\sum_i K_{ij} = 1$.

Row sums. Use the identity $\sum_{j=1}^m \ell_j = r'_p$ (obtained by writing $r_p(x) = \prod_j (x - \nu_j)$ and differentiating), and claim

$$\sum_{j=1}^m (\ell_j \boxplus_m r_q) = r'. \quad (2.18)$$

This is again a coefficient check using \mathcal{E}_m . Let $\mathcal{E}_m(r_p)(t) = A(t)$ and $\mathcal{E}_m(r_q)(t) = B(t)$. Then the normalized coefficient polynomial of the leading-zero derivative r'_p in degree m is $(tA(t))_{\leq m}$. After convolution with r_q it becomes $(tA(t)B(t))_{\leq m}$. On the other hand $\mathcal{E}_m(r) = (A(t)B(t))_{\leq m}$ by (2.2), and the leading-zero derivative of r has normalized polynomial $(t(A(t)B(t)))_{\leq m}$. These coincide because any terms of $A(t)B(t)$ of degree $> m$ disappear after multiplying by t and truncating to degree m . Thus (2.18) holds. Evaluating (2.18) at $x = \mu_i$ gives

$$\sum_{j=1}^m (\ell_j \boxplus_m r_q)(\mu_i) = r'(\mu_i),$$

and dividing by $r'(\mu_i)$ yields $\sum_j K_{ij} = 1$.

Nonnegativity. Assume (for the moment) that the interlacing–preservation theorem of the next subsection is known for simple polynomials of the same (actual) degree. There is a small point of interpretation here, because in our application ℓ_j has degree $m - 1$ (it is being padded by a zero leading coefficient in the order- m convolution). We record explicitly the standard limiting device which reduces this case to the theorem just quoted.

For $\varepsilon > 0$ put

$$\ell_j^{(\varepsilon)}(x) := \ell_j(x) + \varepsilon r_p(x).$$

Then $\ell_j^{(\varepsilon)}$ has real zeros – namely ν_k for $k \neq j$ and the additional zero $\nu_j - 1/\varepsilon$ – and is in non-strict proper position with r_p . If one wants the hypotheses of the interlacing theorem literally, move the common zeros by arbitrarily small alternating perturbations (and divide by the positive leading coefficient) to obtain simple degree- m polynomials which interlace r_p ; after applying the theorem to those polynomials and to r_q , let the perturbations tend to zero. By continuity of the zeros (Hurwitz’s theorem, or elementary continuity of the roots as functions of the coefficients) it follows that $\ell_j^{(\varepsilon)} \boxplus_m r_q$ interlaces $r = r_p \boxplus_m r_q$. Finally we let $\varepsilon \downarrow 0$ and use bilinearity of the convolution to get the desired interlacing of $\ell_j \boxplus_m r_q$ with r (allowing coincidences).

Since the coefficient of $\ell_j \boxplus_m r_q$ at x^{m-1} is positive, such an interlacing (even with coincidences) implies that $(\ell_j \boxplus_m r_q)(\mu_i)$ has the same sign as $r'(\mu_i)$, or is zero, and hence $K_{ij} \geq 0$. \square

5. Real-rootedness and interlacing preservation for \boxplus_n (non-circular)

We now prove the following key theorem. The proof proceeds by a self-contained induction and, at the same time, supplies the deferred nonnegativity of K in all degrees.

Obreschkoff (Hermite–Kakeya–Obreschkoff) theorem. Let f, g be real polynomials of degree n without common zeros and with leading coefficients of the same sign. Then the zeros of f and g interlace if and only if every nontrivial linear combination $af + bg$ has only real zeros. (Proof: consider $R = f/g$; interlacing $\Leftrightarrow R$ strictly monotone between poles; strict monotonicity \Leftrightarrow every horizontal line meets the graph n times.)

Theorem 4.4 (Real-rootedness and interlacing preservation). *If p, q are monic real-rooted polynomials of degree n , then so does $p \boxplus_n q$. Moreover, if p_1, p_2 are two such polynomials and their zeros interlace, then $p_1 \boxplus_n q$ and $p_2 \boxplus_n q$ interlace.*

Proof. We first treat monic polynomials of exact degree n with simple zeros. Multiplying an input by a non-zero constant does not change its zeros (and merely scales the convolution in that factor), and polynomials with multiple zeros will be recovered at the end by approximation.

For the real-rootedness assertion translations allow us to assume that p and q are centered; once it is proved in that situation formula (2.3) removes the centering. We prove, by induction on n , the slightly stronger statement that simple inputs give simple outputs. Simultaneously we use that in all smaller degrees the matrices of Lemma 4.3 are non-negative; this non-negativity was reduced in Lemma 4.3 to the interlacing part of the theorem in the smaller degree and is therefore part of the induction hypothesis.

Base $n = 1$ is trivial. Assume these assertions known up to degree $n - 1$. Let p, q be centered of degree n with simple real zeros, and define $r_p = p'/n$, $r_q = q'/n$. By the induction hypothesis,

$$r = r_p \boxplus_{n-1} r_q$$

has real simple zeros $\mu_1 < \dots < \mu_{n-1}$. Also, by the induction hypothesis again, the matrices K, \tilde{K} arising from the degree $n - 1$ derivative convolution are nonnegative doubly stochastic, so $(Kw^p)_i + (\tilde{K}w^q)_i > 0$.

Lemma 4.2 applied to both $(R_p \boxplus_n q)$ and $(p \boxplus_n R_q)$ together with the split (2.6) yields, at the points μ_i ,

$$(p \boxplus_n q)(\mu_i) = -r'(\mu_i) \left((Kw^p)_i + (\tilde{K}w^q)_i \right). \quad (2.19)$$

The parenthesized quantity is strictly positive. For a monic degree- $(n - 1)$ polynomial r the sign of $r'(\mu_i)$ is $(-1)^{(n-1)-i}$; hence the values $(p \boxplus_n q)(\mu_i)$ alternate in sign.

A monic polynomial of degree n whose derivative has real simple zeros and whose values at these critical points alternate in sign has n real simple zeros (one in each interval $(-\infty, \mu_1), (\mu_1, \mu_2), \dots, (\mu_{n-1}, \infty)$). Indeed the sign just computed gives $(p \boxplus_n q)(\mu_{n-1}) < 0$ while the polynomial is positive for large positive x , and $(p \boxplus_n q)(\mu_1)$ has sign $(-1)^{n-1}$, opposite to the sign $(-1)^n$ at large negative x ; between two consecutive critical points the derivative has a fixed sign, so the alternation yields exactly one crossing in each interval and none at a critical point. Thus $p \boxplus_n q$ is real-rooted in the centered case, hence in general by (2.3).

This completes the induction for real-rootedness.

To prove interlacing preservation, let p_1, p_2 be two interlacing degree- n polynomials. By Obreschkoff, every linear combination $h = ap_1 + bp_2$ is real-rooted. If the leading coefficient of h is non-zero we divide by it and apply the real-rootedness part just proved (multiplying the convolution afterwards by the same constant); if the leading coefficient vanishes we approximate, say, by $ap_1 + (b + \varepsilon)p_2$ and pass to the limit. Consequently

$$(ap_1 + bp_2) \boxplus_n q$$

is real-rooted for all a, b . By linearity of \boxplus_n in the first factor,

$$(ap_1 + bp_2) \boxplus_n q = a(p_1 \boxplus_n q) + b(p_2 \boxplus_n q).$$

Hence every linear combination of $p_1 \boxplus_n q$ and $p_2 \boxplus_n q$ is real-rooted. Applying the Obreschkoff theorem in this limiting form (equivalently, factoring out any common zeros first) yields that $p_1 \boxplus_n q$ and $p_2 \boxplus_n q$ interlace.

It remains only to spell out the harmless limiting convention that was used in Lemma 4.3. The passage from simple to multiple roots in the real-rootedness assertion itself is obtained by the same coefficientwise approximation, since the set of real-rooted polynomials is closed. The interlacing theorem just proved for simple polynomials with non-zero leading coefficient also covers the padded leading-zero inputs occurring there: the approximation $\ell_j^{(\varepsilon)} = \ell_j + \varepsilon r_p$ (together with an arbitrarily small perturbation to remove common zeros) reduces that situation to the strict case, and bilinearity plus continuity of the roots allow one to pass to the

limit. Multiple common roots in the statement of the theorem itself are dealt with by the same perturbation.

Finally, the nonnegativity part of Lemma 4.3 now holds in every degree, because it was reduced to interlacing preservation in the relevant lower degree, which we have proved by induction. \square

6. The key decomposition of the convolved critical values

Keep p, q centered and simple, and keep the notation above. For the convolved polynomial define its w -numbers by

$$w_i(p \boxplus_n q) := -\frac{(p \boxplus_n q - x r)(\mu_i)}{r'(\mu_i)}. \quad (2.20)$$

Then by (2.6) and Lemma 4.2 (and its $p \leftrightarrow q$ analogue),

$$w_i(p \boxplus_n q) = (K w^p)_i + (\tilde{K} w^q)_i, \quad (2.21)$$

where K, \tilde{K} are nonnegative doubly stochastic matrices and w^p, w^q are the vectors of w -numbers of p and q .

7. The one-line estimate and conclusion

Define

$$A_p := \sum_{i=1}^m \frac{1}{w_i(p)}, \quad A_q := \sum_{i=1}^m \frac{1}{w_i(q)},$$

and similarly for $p \boxplus_n q$. If $S_i := (K w^p)_i$, then by Jensen's inequality for the convex function $x \mapsto 1/x$ (using the row sums of K to form convex combinations and the column sums afterwards),

$$\sum_{i=1}^m \frac{1}{S_i} \leq \sum_{i=1}^m \frac{1}{w_i(p)} = A_p, \quad (2.22)$$

and similarly for $T_i := (\tilde{K} w^q)_i$ one has $\sum_i 1/T_i \leq A_q$.

For positive S, T and every real α ,

$$\frac{1}{S+T} \leq \frac{\alpha^2}{S} + \frac{(1-\alpha)^2}{T}, \quad (2.23)$$

since the difference is $(\alpha T - (1-\alpha)S)^2 / (ST(S+T))$. Summing (2.23) with $S = S_i, T = T_i$ and using (2.21) and (2.22) yields

$$\sum_{i=1}^m \frac{1}{w_i(p \boxplus_n q)} \leq \alpha^2 A_p + (1-\alpha)^2 A_q.$$

Choosing $\alpha = A_q / (A_p + A_q)$ gives

$$\sum_{i=1}^m \frac{1}{w_i(p \boxplus_n q)} \leq \frac{A_p A_q}{A_p + A_q}. \quad (2.24)$$

By Lemma 4.1,

$$\Phi_n(p) = \frac{n}{4} A_p, \quad \Phi_n(q) = \frac{n}{4} A_q, \quad \Phi_n(p \boxplus_n q) = \frac{n}{4} \sum_{i=1}^m \frac{1}{w_i(p \boxplus_n q)}.$$

Multiplying (2.24) by $n/4$ and inverting gives

$$\frac{1}{\Phi_n(p \boxplus_n q)} \geq \frac{1}{\Phi_n(p)} + \frac{1}{\Phi_n(q)}.$$

This proves the inequality for centered simple-rooted polynomials.

8. Removing centering and multiple roots

If the roots are not centered, translate p and q to make them centered, apply the proved inequality, and translate back using (2.3). The value of Φ_n is unchanged by a common translation of the roots.

If a polynomial has a multiple root we interpret $1/\Phi_n$ as 0. Choose sequences $p^{(s)}, q^{(s)}$ of monic real-rooted polynomials with simple roots converging coefficientwise to the given ones. The map $(p, q) \mapsto p \boxplus_n q$ is polynomial in the coefficients, hence continuous, and for simple polynomials the quantities in (2.8) are continuous functions of the roots. Moreover (2.8) shows that as two roots of a real-rooted polynomial coalesce the sum defining Φ_n tends to $+\infty$, so its reciprocal tends to 0. Applying the proved inequality to $p^{(s)}, q^{(s)}$ and taking \liminf therefore gives the desired inequality in the limit. In particular, if the right-hand side has a positive limit the convolved polynomials cannot acquire a multiple root in the limit (otherwise the left-hand side would tend to 0), and if the right-hand side tends to 0 the estimate is immediate. Thus the convention on multiple roots is consistent and the inequality survives passage to the limit.

Conclusion

For all $n \geq 2$ and all monic real-rooted degree- n polynomials p, q (with $\Phi_n = \infty$ on multiple roots),

$$\boxed{\frac{1}{\Phi_n(p \boxplus_n q)} \geq \frac{1}{\Phi_n(p)} + \frac{1}{\Phi_n(q)}}.$$

5 The \mathcal{O} -adapted slice filtration and a geometric fixed-point criterion for slice connectivity

Problem

Fix a finite group G . Let \mathcal{O} denote an incomplete transfer system associated to an N_∞ operad. Define the slice filtration on the G -equivariant stable category adapted to \mathcal{O} and state and prove a characterization of the \mathcal{O} -slice connectivity of a connective G -spectrum in terms of the geometric fixed points.

Solution

Conventions. By a *localizing subcategory* of a stable category we mean a full subcategory closed under equivalences, cofibres (hence suspensions), extensions, and arbitrary coproducts (not necessarily under inverse suspension). For a collection of objects \mathcal{C} we write $\text{Loc}\langle\mathcal{C}\rangle$ for the smallest localizing subcategory containing \mathcal{C} .

We write $\Phi^H: \text{Sp}^G \rightarrow \text{Sp}$ for geometric H -fixed points. When we speak of the *connectivity* of $\Phi^H(X)$ we mean the connectivity of the underlying nonequivariant spectrum. Let $\text{Sp}_{\geq k} \subset \text{Sp}$ denote the usual full subcategory of k -connective spectra.

A G -spectrum X will be called *connective* if $\Phi^H(X) \in \text{Sp}_{\geq 0}$ for every $H \leq G$.

Lemma 5.1 (Connective spectra are generated by spheres). *For every integer $c \geq 0$ one has*

$$\text{Loc}_{\text{Sp}}\langle S^m \mid m \geq c \rangle = \text{Sp}_{\geq c}.$$

Proof. Since $\text{Sp}_{\geq c}$ is closed under coproducts, cofibres, and extensions, and contains S^m for all $m \geq c$, it contains $\text{Loc}_{\text{Sp}}\langle S^m \mid m \geq c \rangle$. Thus

$$\text{Loc}_{\text{Sp}}\langle S^m \mid m \geq c \rangle \subseteq \text{Sp}_{\geq c}.$$

Conversely, let $X \in \text{Sp}_{\geq c}$. By the standard cellular approximation theorem for spectra (see [1, § III.2]) one can build a CW-approximation of X using only cells in degrees $\geq c$: start with $X_{-1} = 0$ and inductively construct a sequence

$$X_{-1} \longrightarrow X_0 \longrightarrow X_1 \longrightarrow \cdots \longrightarrow X$$

so that $X_r \rightarrow X$ is $(c+r)$ -connective and X_r is obtained from X_{r-1} by attaching a wedge of spheres $\bigvee \Sigma^m S^0$ with $m \geq c+r$ (e.g. by killing the kernel of $\pi_{c+r}(X_{r-1}) \rightarrow \pi_{c+r}(X)$ and then surjecting onto $\pi_{c+r+1}(X)$). Taking the sequential homotopy colimit (mapping telescope) yields an equivalence

$$\text{hocolim}_r X_r \simeq X.$$

Each X_r lies in $\text{Loc}_{\text{Sp}}\langle S^m \mid m \geq c \rangle$, and this subcategory is closed under sequential homotopy colimits because in a stable category $\text{hocolim}_r X_r \simeq \text{cofib}(\bigoplus_r X_r \xrightarrow{1\text{-shift}} \bigoplus_r X_r)$ is built from coproducts and a cofiber. Hence $X \in \text{Loc}_{\text{Sp}}\langle S^m \mid m \geq c \rangle$. \square

Transfer systems from N_∞ -operads. Write $K \rightarrow H$ for the relation in the transfer system. For each $H \leq G$ set

$$\mathcal{O}_H := \{ K \leq H \mid K \rightarrow H \}.$$

If \mathcal{O} comes from an N_∞ -operad (equivalently an indexing system), then for each fixed H the family \mathcal{O}_H is closed under H -conjugation and finite intersections (this reflects closure of admissible H -sets under products and subobjects; see [2, §3]).

Consequently, \mathcal{O}_H has a smallest subgroup:

$$H_{\mathcal{O}} := \bigcap_{\substack{K \leq H \\ K \rightarrow H}} K. \quad (23)$$

This subgroup is normal in H and still belongs to \mathcal{O}_H .

Define the index and the real permutation representation

$$d_H := [H : H_{\mathcal{O}}] \in \mathbb{Z}_{\geq 1}, \quad \rho_H^{\mathcal{O}} := \mathbb{R}[H/H_{\mathcal{O}}]$$

(the real vector space on the H -set of cosets, with the induced H -action).

Lemma 5.2 (Conjugation and restriction invariance). *Let \mathcal{O} be the transfer system under discussion (coming from an N_{∞} -operad, so that the subgroups $H_{\mathcal{O}}$ of (23) are defined).*

1. (Conjugation invariance.) *For any $g \in G$ and $L \leq G$, conjugation induces a bijection between subgroups of L which transfer to L and subgroups of $g^{-1}Lg$ which transfer to $g^{-1}Lg$. Consequently*

$$(g^{-1}Lg)_{\mathcal{O}} = g^{-1}L_{\mathcal{O}}g,$$

and in particular $d_{g^{-1}Lg} = d_L$.

2. (Restriction to subgroups.) *Let $H \leq G$ and let $\mathcal{O}|_H$ denote the restricted transfer system on H (i.e. $K \rightarrow_{\mathcal{O}|_H} L$ iff $K \rightarrow_{\mathcal{O}} L$ for $K \leq L \leq H$). Then for every $L \leq H$ one has*

$$L_{\mathcal{O}|_H} = L_{\mathcal{O}} \quad \text{and hence} \quad d_L^{\mathcal{O}|_H} = d_L^{\mathcal{O}}.$$

Proof. (1) If $K \leq L$ and $K \rightarrow L$, closure of the transfer relation under conjugation implies $g^{-1}Kg \rightarrow g^{-1}Lg$. This gives a bijection between the indexing sets in the defining intersections, so

$$(g^{-1}Lg)_{\mathcal{O}} = \bigcap_{\substack{K' \leq g^{-1}Lg \\ K' \rightarrow g^{-1}Lg}} K' = \bigcap_{\substack{K \leq L \\ K \rightarrow L}} g^{-1}Kg = g^{-1} \left(\bigcap_{\substack{K \leq L \\ K \rightarrow L}} K \right) g = g^{-1}L_{\mathcal{O}}g.$$

Taking indices yields $d_{g^{-1}Lg} = d_L$.

(2) By definition of restriction, the collection $\{K \leq L \mid K \rightarrow_{\mathcal{O}|_H} L\}$ coincides with $\{K \leq L \mid K \rightarrow_{\mathcal{O}} L\}$ whenever $L \leq H$, so the defining intersections agree. \square

Lemma 5.3 (Monotonicity and an orbit-count estimate). *For the above transfer system let $K \leq H \leq G$.*

1. *One has $K_{\mathcal{O}} \leq K \cap H_{\mathcal{O}}$ and hence $d_K \geq [K : K \cap H_{\mathcal{O}}]$.*
2. *The number of K -orbits in $H/H_{\mathcal{O}}$ satisfies*

$$|K \backslash H/H_{\mathcal{O}}| = \frac{|H : H_{\mathcal{O}}|}{|K : K \cap H_{\mathcal{O}}|} \geq \frac{d_H}{d_K}.$$

Proof. (1) By the restriction axiom for transfer systems, from $H_{\mathcal{O}} \rightarrow H$ we obtain $K \cap H_{\mathcal{O}} \rightarrow K$. Since $K_{\mathcal{O}}$ is the smallest subgroup in \mathcal{O}_K (by the same construction (23) applied to K), we have $K_{\mathcal{O}} \leq K \cap H_{\mathcal{O}}$. Taking indices gives $d_K = [K : K_{\mathcal{O}}] \geq [K : K \cap H_{\mathcal{O}}]$.

(2) Because $H_{\mathcal{O}} \trianglelefteq H$, the stabilizer in K of any coset in $H/H_{\mathcal{O}}$ is exactly $K \cap H_{\mathcal{O}}$. Hence every K -orbit has cardinality $[K : K \cap H_{\mathcal{O}}]$, so

$$|K \backslash H/H_{\mathcal{O}}| = \frac{|H : H_{\mathcal{O}}|}{|K : K \cap H_{\mathcal{O}}|}.$$

The inequality follows from (1): $|K : K \cap H_{\mathcal{O}}| \leq d_K$ and $|H : H_{\mathcal{O}}| = d_H$. \square

The \mathcal{O} -adapted (regular) slice filtration. For each integer n define $\tau_{\geq n}^{\mathcal{O}} \subset \mathrm{Sp}^G$ to be the localizing subcategory generated by the \mathcal{O} -slice cells

$$G_+ \wedge_H S^{m\rho_H^{\mathcal{O}}} \quad (H \leq G, m \in \mathbb{Z}, md_H \geq n). \quad (24)$$

Here S^V denotes the representation sphere of a (virtual) real H -representation V ; for $m < 0$ the notation means the corresponding desuspension by the virtual representation $m\rho_H^{\mathcal{O}}$. The full subcategories $\tau_{\geq n}^{\mathcal{O}}$ form a decreasing filtration of Sp^G ; in what follows we only use these subcategories themselves.

We will use the following (well-known) result.

Lemma 5.4 (Geometric fixed points detect equivalences). *A map $f: X \rightarrow Y$ in Sp^G is an equivalence if and only if $\Phi^H(f)$ is an equivalence of spectra for every subgroup $H \leq G$.*

Proof. Only the “if” direction needs proof. Let $F := \mathrm{fib}(f)$. Since each Φ^H is exact, $\Phi^H(F) \simeq \mathrm{fib}(\Phi^H(f)) \simeq 0$ for all $H \leq G$. By [3, Prop. 2.52], this implies $F \simeq 0$, hence f is an equivalence. \square

The fixed-point characterization of \mathcal{O} -slice connectivity.

Theorem 5.5. *Let X be a connective G -spectrum and let $n \geq 0$. Then*

$$X \in \tau_{\geq n}^{\mathcal{O}} \iff \Phi^H(X) \in \mathrm{Sp}_{\geq [n/d_H]} \text{ for every } H \leq G. \quad (25)$$

Proof. Set

$$\mathcal{P}_n(G) := \{ Y \in \mathrm{Sp}^G \mid \Phi^H(Y) \in \mathrm{Sp}_{\geq [n/d_H]} \text{ for all } H \leq G \}.$$

Because each Φ^H is exact and preserves coproducts, and each $\mathrm{Sp}_{\geq k}$ is localizing in Sp , the class $\mathcal{P}_n(G)$ is a localizing subcategory of Sp^G .

Step 1: $\tau_{\geq n}^{\mathcal{O}} \subseteq \mathcal{P}_n(G)$. Since $\mathcal{P}_n(G)$ is localizing, it suffices to check the generators (24). Fix a generator

$$Y = G_+ \wedge_H S^{m\rho_H^{\mathcal{O}}} \quad \text{with } md_H \geq n.$$

Because $n \geq 0$ and $d_H \geq 1$, this forces $m \geq 0$.

Fix $K \leq G$. A standard double-coset formula for geometric fixed points of induced spectra (see e.g. [3, Prop. 2.46]) gives an equivalence

$$\Phi^K(G_+ \wedge_H Z) \simeq \bigvee_{\substack{[g] \in K \backslash G/H \\ K^g \leq H}} \Sigma^\infty(W_G K / W_H K^g)_+ \wedge \Phi^{K^g}(Z), \quad (26)$$

natural in the H -spectrum Z , where $K^g := g^{-1}Kg$ and $W_L M = N_L(M)/M$ denotes the Weyl group (the subgroup $W_H K^g$ is viewed inside $W_G K$ by conjugation with g). Applying this with $Z = S^{m\rho_H^{\mathcal{O}}}$ and using that geometric fixed points of a representation sphere are the sphere of the fixed subrepresentation,

$$\Phi^{K^g}(S^{m\rho_H^{\mathcal{O}}}) \simeq S^{m(\rho_H^{\mathcal{O}})^{K^g}},$$

we obtain

$$\Phi^K(Y) \simeq \bigvee_{\substack{[g] \in K \backslash G/H \\ K^g \leq H}} \Sigma^\infty(W_G K / W_H K^g)_+ \wedge S^{m(\rho_H^{\mathcal{O}})^{K^g}}.$$

If the indexing set is empty then this wedge is the zero spectrum (and hence is connective in every degree), so there is nothing to prove. Otherwise smashing with Σ^∞ of a finite set merely produces a finite wedge of the same sphere, and it remains to understand the dimension of $(\rho_H^{\mathcal{O}})^{K^g}$. Since

$\rho_H^\mathcal{O} = \mathbb{R}[H/H_\mathcal{O}]$ is a permutation representation, the fixed subspace has dimension equal to the number of orbits:

$$\dim((\rho_H^\mathcal{O})^{K^g}) = |K^g \backslash H/H_\mathcal{O}|.$$

Because $K^g \leq H$, Lemma 5.3(2) applied to $K^g \leq H$ gives

$$|K^g \backslash H/H_\mathcal{O}| \geq \frac{d_H}{d_{K^g}}.$$

By Lemma 5.2(1) we have $d_{K^g} = d_K$, hence

$$|K^g \backslash H/H_\mathcal{O}| \geq \frac{d_H}{d_K}.$$

Therefore every sphere summand has dimension at least $m(d_H/d_K) \geq n/d_K$, hence at least $\lceil n/d_K \rceil$ since it is integral. Thus $\Phi^K(Y) \in \mathrm{Sp}_{\geq \lceil n/d_K \rceil}$ for all K , so $Y \in \mathcal{P}_n(G)$, and $\tau_{\geq n}^\mathcal{O} \subseteq \mathcal{P}_n(G)$ follows.

Step 2: $\mathcal{P}_n(G) \subseteq \tau_{\geq n}^\mathcal{O}$. We prove this by induction on $|G|$. When $G = e$, we have $d_e = 1$ and

$$\tau_{\geq n}^\mathcal{O} = \mathrm{Loc}_{\mathrm{Sp}} \langle S^m \mid m \geq n \rangle = \mathrm{Sp}_{\geq n}$$

by Lemma 5.1, so the claim is the ordinary Postnikov connectivity statement.

Assume now $G \neq e$ and that the statement holds for all proper subgroups of G (with the restricted transfer system). Let X be connective and assume $X \in \mathcal{P}_n(G)$, i.e. $\Phi^H(X) \in \mathrm{Sp}_{\geq \lceil n/d_H \rceil}$ for all $H \leq G$.

Let \mathcal{F} be the family of proper subgroups of G and consider isotropy separation:

$$E\mathcal{F}_+ \wedge X \longrightarrow X \longrightarrow \tilde{E}\mathcal{F} \wedge X. \quad (27)$$

Step 2a: $E\mathcal{F}_+ \wedge X \in \tau_{\geq n}^\mathcal{O}$. For each proper $H < G$, the restricted H -spectrum $\mathrm{res}_H^G X$ is connective and satisfies the same geometric fixed-point bounds for all $L \leq H$:

$$\Phi^L(\mathrm{res}_H^G X) \simeq \Phi^L(X) \in \mathrm{Sp}_{\geq \lceil n/d_L \rceil}.$$

Here $\tau_{\geq n}^\mathcal{O}(\mathrm{Sp}^H)$ denotes the filtration for the restricted transfer system $\mathcal{O}|_H$; by Lemma 5.2(2) the subgroups $L_{\mathcal{O}|_H}$ (and hence the integers d_L) agree with those computed in G for all $L \leq H$. Therefore the inductive hypothesis applies and gives

$$\mathrm{res}_H^G X \in \tau_{\geq n}^\mathcal{O}(\mathrm{Sp}^H).$$

A G -CW filtration of $E\mathcal{F}$ has cells $G/H \times D^r$ with $H \in \mathcal{F}$ and $r \geq 0$. Smashing with X , the successive cofibres in the skeletal filtration are wedges of spectra

$$G/H_+ \wedge S^r \wedge X \simeq G_+ \wedge_H \Sigma^r(\mathrm{res}_H^G X).$$

Induction $G_+ \wedge_H (-)$ is exact and preserves coproducts, hence sends localizing subcategories to localizing subcategories; moreover it sends each $\mathcal{O}|_H$ -slice cell $H_+ \wedge_L S^{m\rho_L^{\mathcal{O}|_H}}$ (with $L \leq H$) to the corresponding \mathcal{O} -slice cell $G_+ \wedge_L S^{m\rho_L^\mathcal{O}}$ (using Lemma 5.2(2) to identify $\rho_L^{\mathcal{O}|_H} = \rho_L^\mathcal{O}$). Thus

$$G_+ \wedge_H (\tau_{\geq n}^\mathcal{O}(\mathrm{Sp}^H)) \subseteq \tau_{\geq n}^\mathcal{O}(\mathrm{Sp}^G).$$

Since suspension preserves $\tau_{\geq n}^\mathcal{O}$, each skeletal stage of $E\mathcal{F}_+ \wedge X$ lies in $\tau_{\geq n}^\mathcal{O}$. Finally, $\tau_{\geq n}^\mathcal{O}$ is closed under sequential homotopy colimits (mapping telescopes) because it is closed under coproducts and cofibres; hence the colimit $E\mathcal{F}_+ \wedge X$ lies in $\tau_{\geq n}^\mathcal{O}(\mathrm{Sp}^G)$.

Step 2b: $\tilde{E}\mathcal{F} \wedge X \in \tau_{\geq n}^{\mathcal{O}}$. Set $Z := \tilde{E}\mathcal{F} \wedge X$. Then Z is concentrated over G :

$$\Phi^H(Z) \simeq 0 \quad (H < G), \quad \Phi^G(Z) \simeq \Phi^G(X).$$

Let $c := \lceil n/d_G \rceil \in \mathbb{Z}_{\geq 0}$. By assumption $\Phi^G(X) \in \mathrm{Sp}_{\geq c}$, hence $\Phi^G(Z) \in \mathrm{Sp}_{\geq c}$. By Lemma 5.1, this implies

$$\Phi^G(Z) \in \mathrm{Loc}_{\mathrm{Sp}}\langle S^m \mid m \geq c \rangle. \quad (28)$$

Consider the exact functor

$$L: \mathrm{Sp} \rightarrow \mathrm{Sp}^G, \quad L(W) := \tilde{E}\mathcal{F} \wedge \epsilon^*W,$$

where ϵ^* denotes the trivial G -action. The spectrum $L(W)$ is concentrated over G and satisfies $\Phi^G(L(W)) \simeq W$.

Lemma 5.6 (Spectra concentrated over G). *Let $\mathcal{C} \subset \mathrm{Sp}^G$ be the full subcategory of G -spectra T such that $E\mathcal{F}_+ \wedge T \simeq 0$ (equivalently $\Phi^H(T) \simeq 0$ for every proper subgroup $H < G$). Then*

$$\Phi^G: \mathcal{C} \longrightarrow \mathrm{Sp}$$

*is an equivalence and a quasi-inverse is $L(W) = \tilde{E}\mathcal{F} \wedge \epsilon^*W$. In particular, for every $Z \in \mathcal{C}$ there is a natural map*

$$\tilde{E}\mathcal{F} \wedge \epsilon^*(\Phi^G Z) \longrightarrow Z$$

which is an equivalence.

Proof. We shall use two standard pieces of equivariant stable homotopy theory. First, smashing with $\tilde{E}\mathcal{F}$ is a smashing Bousfield localization of Sp^G : its local objects are precisely the spectra T with $E\mathcal{F}_+ \wedge T \simeq 0$, and the functor $L_{\mathcal{F}}(Y) = \tilde{E}\mathcal{F} \wedge Y$ is left adjoint to the inclusion of this full subcategory (for example [4, §3.3]). Secondly, for a finite group the geometric fixed point functors are exact, preserve arbitrary homotopy colimits and are strong symmetric monoidal, and for a suspension spectrum one has $\Phi^H(\Sigma^\infty A) \simeq \Sigma^\infty(A^H)$; we refer to [5, §V.4] (see also [3, §2]).

Let us first check that the two descriptions of \mathcal{C} agree. If $T \simeq \tilde{E}\mathcal{F} \wedge T$, then for every proper $H < G$ the preceding properties give $\Phi^H(T) \simeq \Phi^H(\tilde{E}\mathcal{F}) \wedge \Phi^H(T) \simeq 0$. Conversely, if all these proper geometric fixed points vanish, the map $T \rightarrow \tilde{E}\mathcal{F} \wedge T$ in the isotropy–separation cofibre sequence is an equivalence after applying every Φ^H (for $H = G$ because $\Phi^G(E\mathcal{F}_+) \simeq 0$), and hence is an equivalence by Lemma 5.4. Thus \mathcal{C} is exactly the local subcategory for this smashing localization.

Put $L(W) = \tilde{E}\mathcal{F} \wedge \epsilon^*W$. If $X \in \mathcal{C}$, the localization adjunction and the ordinary adjunction $\epsilon^* \dashv (-)_{\mathrm{cat}}^G$ for categorical fixed points give natural isomorphisms of morphism sets

$$[L(W), X]_{\mathrm{Sp}^G} \cong [\epsilon^*W, X]_{\mathrm{Sp}^G} \cong [W, X_{\mathrm{cat}}^G]_{\mathrm{Sp}} = [W, \Phi^G X]_{\mathrm{Sp}}, \quad (29)$$

where in the last step we used $X \simeq \tilde{E}\mathcal{F} \wedge X$. Hence $L: \mathrm{Sp} \rightarrow \mathcal{C}$ is left adjoint to the restricted functor Φ^G .

The unit of this adjunction is an isomorphism. Indeed the monoidality just recalled and the calculation of the geometric fixed points of the universal spaces give natural equivalences

$$\Phi^G L(W) \simeq \Phi^G(\tilde{E}\mathcal{F}) \wedge \Phi^G(\epsilon^*W) \simeq S^0 \wedge W \simeq W.$$

For the second equivalence note that the exact coproduct–preserving functor $\Phi^G \epsilon^*$ agrees with the identity on all suspension spheres, and hence on all spectra by the universal property of the stable homotopy category. Under these identifications the adjunction unit is the inverse equivalence: for $W = S^0$ it is induced by the collapse map $S^0 \circ \tilde{E}\mathcal{F}$ (after applying Φ^G), and since both functors are exact and preserve arbitrary coproducts while S^0 is a compact generator

of Sp , the unit is an isomorphism for every W . By the triangular identity for the adjunction (29), the counit

$$\epsilon_X : \tilde{E}\mathcal{F} \wedge \epsilon^*(\Phi^G X) \longrightarrow X, \quad X \in \mathcal{C},$$

therefore becomes an equivalence after applying Φ^G . For a proper subgroup both source and target have trivial geometric fixed points, so Lemma 5.4 shows that ϵ_X itself is an equivalence. This counit is the displayed natural map and the two functors are quasi-inverse equivalences. \square

Applying L to (28) and using exactness and coproduct preservation, we obtain

$$Z \simeq L(\Phi^G Z) \in \mathrm{Loc}_{\mathrm{Sp}^G} \langle L(S^m) \mid m \geq c \rangle. \quad (30)$$

Thus it suffices to show $L(S^m) \in \tau_{\geq n}^{\mathcal{O}}$ for every $m \geq c$.

Claim 5.7. If $m \geq 0$ and $md_G \geq n$, then $L(S^m) \in \tau_{\geq n}^{\mathcal{O}}(\mathrm{Sp}^G)$.

Proof. The generator $S^{m\rho_G^{\mathcal{O}}}$ belongs to $\tau_{\geq n}^{\mathcal{O}}$ by definition since $md_G \geq n$.

We first show $E\mathcal{F}_+ \wedge S^{m\rho_G^{\mathcal{O}}} \in \tau_{\geq n}^{\mathcal{O}}$. For a proper subgroup $H < G$, the restricted H -spectrum $\mathrm{res}_H^G S^{m\rho_G^{\mathcal{O}}}$ is connective and, for $L \leq H$,

$$\Phi^L(\mathrm{res}_H^G S^{m\rho_G^{\mathcal{O}}}) \simeq S^{m(\rho_G^{\mathcal{O}})^L}.$$

Since $\rho_G^{\mathcal{O}} = \mathbb{R}[G/G_{\mathcal{O}}]$ is a permutation representation, $\dim((\rho_G^{\mathcal{O}})^L) = |L \backslash G/G_{\mathcal{O}}|$, and Lemma 5.3(2) (with $H = G$) yields $|L \backslash G/G_{\mathcal{O}}| \geq d_G/d_L$. Hence

$$\Phi^L(\mathrm{res}_H^G S^{m\rho_G^{\mathcal{O}}}) \in \mathrm{Sp}_{\geq \lceil n/d_L \rceil}.$$

By the inductive hypothesis applied to the group H (with the restricted transfer system, and using Lemma 5.2(2) to identify the same d_L), we conclude $\mathrm{res}_H^G S^{m\rho_G^{\mathcal{O}}} \in \tau_{\geq n}^{\mathcal{O}}(\mathrm{Sp}^H)$ for all proper $H < G$. The same G -CW cellular argument as in Step 2a (using exactness/coproduct preservation of induction and closure under sequential homotopy colimits) then shows $E\mathcal{F}_+ \wedge S^{m\rho_G^{\mathcal{O}}} \in \tau_{\geq n}^{\mathcal{O}}(\mathrm{Sp}^G)$.

Now in the cofiber sequence

$$E\mathcal{F}_+ \wedge S^{m\rho_G^{\mathcal{O}}} \longrightarrow S^{m\rho_G^{\mathcal{O}}} \longrightarrow \tilde{E}\mathcal{F} \wedge S^{m\rho_G^{\mathcal{O}}},$$

the first two terms lie in $\tau_{\geq n}^{\mathcal{O}}$, so the third term does as well:

$$\tilde{E}\mathcal{F} \wedge S^{m\rho_G^{\mathcal{O}}} \in \tau_{\geq n}^{\mathcal{O}}. \quad (31)$$

Set $Z_0 := \tilde{E}\mathcal{F} \wedge S^{m\rho_G^{\mathcal{O}}}$. Then Z_0 is concentrated over G , so $\Phi^H(Z_0) \simeq 0$ for every proper subgroup $H < G$, while

$$\Phi^G(Z_0) \simeq S^{m(\rho_G^{\mathcal{O}})^G} \simeq S^m,$$

since $(\rho_G^{\mathcal{O}})^G \cong \mathbb{R}$ (constant functions on $G/G_{\mathcal{O}}$). Therefore Lemma 5.6 (applied to $Z_0 \in \mathcal{C}$) supplies a natural equivalence

$$L(S^m) = L(\Phi^G Z_0) \xrightarrow{\simeq} Z_0 = \tilde{E}\mathcal{F} \wedge S^{m\rho_G^{\mathcal{O}}}.$$

Since the right-hand side lies in $\tau_{\geq n}^{\mathcal{O}}$ by (31), and $\tau_{\geq n}^{\mathcal{O}}$ is closed under equivalences, it follows that $L(S^m) \in \tau_{\geq n}^{\mathcal{O}}$. \square

Now if $m \geq c = \lceil n/d_G \rceil$, then $m \geq 0$ and $md_G \geq n$, so Claim 5.7 applies. Therefore all generators in (30) lie in $\tau_{\geq n}^{\mathcal{O}}$, hence $Z \in \tau_{\geq n}^{\mathcal{O}}$.

Step 2c: conclude. In (27), both end terms lie in the localizing subcategory $\tau_{\geq n}^{\mathcal{O}}$, which is closed under extensions, hence $X \in \tau_{\geq n}^{\mathcal{O}}$.

Combining Steps 1 and 2 proves (79). \square

Remark 5.8 (Two extremes). If \mathcal{O} allows all transfers, then $H_{\mathcal{O}} = e$ and $d_H = |H|$, and Theorem 5.5 recovers the usual regular slice-connectivity criterion. If \mathcal{O} only allows trivial transfers, then $H_{\mathcal{O}} = H$ and $d_H = 1$ for all H , and the filtration $\tau_{\geq n}^{\mathcal{O}}$ reduces to the ordinary Postnikov filtration.

References

- [1] J. F. Adams, *Stable Homotopy and Generalised Homology*, University of Chicago Press, Chicago Lectures in Mathematics, 1974.
- [2] A. J. Blumberg and M. A. Hill, *Operadic multiplications in equivariant spectra, norms, and transfers*, *Adv. Math.* **285** (2015), 658–708.
- [3] M. A. Hill, M. J. Hopkins, and D. C. Ravenel, *On the nonexistence of elements of Kervaire invariant one*, *Ann. of Math. (2)* **184** (2016), no. 1, 1–262.
- [4] M. Hovey, J. H. Palmieri, and N. P. Strickland, *Axiomatic stable homotopy theory*, *Mem. Amer. Math. Soc.* **128** (1997), no. 610.
- [5] M. A. Mandell and J. P. May, *Equivariant orthogonal spectra and S-modules*, *Mem. Amer. Math. Soc.* **159** (2002), no. 755.

6 Large ε -light vertex subsets

Problem

For a graph $G = (V, E)$, let $G_S = (V, E(S, S))$ denote the graph with the same vertex set, but only the edges between vertices in S . Let L be the Laplacian matrix of G and let L_S be the Laplacian of G_S . I say that a set of vertices S is ε -light if the matrix $\varepsilon L - L_S$ is positive semidefinite. Does there exist a constant $c > 0$ so that for every graph G and every ε between 0 and 1, V contains an ε -light subset S of size at least $c\varepsilon|V|$?

Solution

We prove the claim with an explicit constant $c = 1/256$.

Throughout we write $n = |V|$. Some degenerate cases are immediate and will be set aside. If $n = 0$ the assertion is vacuous. If $\varepsilon = 0$ the empty set is 0-light and has the required size. If the graph has no edges (so that its Laplacian is the zero matrix), then $L_S = 0$ for every S and we may simply take all vertices. Thus in the main part of the proof we assume $n \geq 1$, $\varepsilon \in (0, 1]$, and that the Laplacian has positive rank. All Loewner inequalities and traces below are taken on the subspace $\text{range}(L) = (\ker L)^\perp$, and I denotes the identity on that space. The exceptional cases are revisited in Step 6.

Step 1: Normalization on the Laplacian range. Let $\ker L$ be the space of vectors that are constant on each connected component of G . Let L^\dagger be the Moore–Penrose pseudoinverse, and define

$$L^{-1/2} := (L^\dagger)^{1/2}.$$

Then $L^{-1/2}$ acts as the inverse square root on $\text{range}(L) = (\ker L)^\perp$ and as 0 on $\ker L$.

For an edge $e = \{u, v\}$ define the rank-one edge Laplacian

$$L_e := (e_u - e_v)(e_u - e_v)^\top,$$

so that $L = \sum_{e \in E} L_e$. All sums over edges below are taken with multiplicity if the graph has parallel edges. Define

$$A_e := L^{-1/2} L_e L^{-1/2}.$$

Each A_e is positive semidefinite on $\text{range}(L)$, whose dimension is $d := \text{rank}(L) \leq n$. Moreover, on $\text{range}(L)$ we have

$$\sum_{e \in E} A_e = L^{-1/2} \left(\sum_{e \in E} L_e \right) L^{-1/2} = L^{-1/2} L L^{-1/2} = I. \quad (32)$$

Also, for any $S \subseteq V$,

$$L^{-1/2} L_S L^{-1/2} = \sum_{e \in E(S, S)} A_e \quad \text{on } \text{range}(L). \quad (33)$$

Therefore, it suffices to find S such that on $\text{range}(L)$,

$$\sum_{e \in E(S, S)} A_e \preceq \varepsilon I, \quad (34)$$

because then (33) implies $L^{-1/2} L_S L^{-1/2} \preceq \varepsilon I$, i.e.

$$x^\top L_S x \leq \varepsilon x^\top L x \quad \text{for all } x \perp \ker L.$$

If $P = L^{-1/2}L^{1/2} = L^{1/2}L^{-1/2}$ denotes the orthogonal projection onto $\text{range}(L)$, the displayed Loewner inequality is equivalent to $x^T L_S x \leq \varepsilon x^T L x$ for every $x \in \text{range}(L)$ by taking $z = L^{1/2}x$ (where $L^{1/2}$ denotes the positive square root of L , acting as zero on $\ker L$) in the quadratic form. Vectors in $\ker L$ are constant on each connected component and are therefore also annihilated by L_S ; by symmetry no mixed terms occur between $\text{range}(L)$ and $\ker L$. Hence the inequality holds for all $x \in \mathbb{R}^V$, which is exactly $\varepsilon L - L_S \succeq 0$.

Step 2: A one-sided BSS barrier lemma. The following lemma is a one-sided variant of the barrier method introduced by Batson, Spielman and Srivastava [1]; we give a complete proof for the reader's convenience. For a PSD matrix $M \succeq 0$ on a d -dimensional space and a scalar $u > \lambda_{\max}(M)$, define the potential

$$\Phi_u(M) := \text{tr}(uI - M)^{-1}.$$

Lemma 6.1 (One-sided barrier). *Assume $M \prec uI$, let $u' > u$, and put $U := (u'I - M)^{-1}$. If $B \succeq 0$ satisfies*

$$\text{tr}(BU) + \frac{\text{tr}(BU^2)}{\Phi_u(M) - \Phi_{u'}(M)} \leq 1, \quad (35)$$

then $M + B \prec u'I$ and $\Phi_{u'}(M + B) \leq \Phi_u(M)$.

Proof. Let $K := B^{1/2}UB^{1/2} \succeq 0$. The hypothesis (35) implies $\text{tr}(K) < 1$: the second summand there is non-negative, and if it were zero then the positive semidefinite matrix $B^{1/2}U^2B^{1/2}$ would have trace zero and hence vanish; since U is invertible on our space this forces $B = 0$. Consequently every eigenvalue of K is < 1 , so in particular $\|K\| < 1$ and $(I - K)$ is invertible. By the Sherman–Morrison–Woodbury identity (which can be verified by multiplying the two sides),

$$(u'I - M - B)^{-1} = U + UB^{1/2}(I - K)^{-1}B^{1/2}U,$$

so $u'I - M - B \succ 0$, i.e. $M + B \prec u'I$.

Taking traces, using cyclicity of the trace and the elementary fact that $\text{tr}(XC) \leq \text{tr}(YC)$ whenever $0 \preceq X \preceq Y$ and $C \succeq 0$, together with $(I - K)^{-1} \preceq (1 - \text{tr}K)^{-1}I$ (valid for PSD K with $\text{tr}K < 1$), we obtain

$$\Phi_{u'}(M + B) \leq \Phi_{u'}(M) + \frac{\text{tr}(BU^2)}{1 - \text{tr}(BU)}.$$

A short rearrangement shows that (35) is equivalent to the bound that the right-hand side is at most $\Phi_u(M)$. This yields $\Phi_{u'}(M + B) \leq \Phi_u(M)$. \square

We will also use the following inequality: if $u' = u + \delta$, then

$$\Phi_u(M) - \Phi_{u'}(M) \geq \delta \text{tr}(u'I - M)^{-2} = \delta \text{tr}(U^2). \quad (36)$$

Indeed, diagonalizing M with eigenvalues λ_j gives

$$\Phi_u(M) - \Phi_{u'}(M) = \sum_j \frac{\delta}{(u - \lambda_j)(u' - \lambda_j)} \geq \sum_j \frac{\delta}{(u' - \lambda_j)^2}.$$

Step 3: A partial coloring process. Fix $\varepsilon \in (0, 1]$ and set

$$r := \left\lceil \frac{16}{\varepsilon} \right\rceil, \quad u_0 := \frac{\varepsilon}{2}, \quad \delta := \frac{\varepsilon}{n}, \quad k := \left\lfloor \frac{n}{4} \right\rfloor. \quad (37)$$

We will color k vertices, one at a time, using r colors.

At time t ($0 \leq t \leq k$), let $T \subseteq V$ be the set of colored vertices, $|T| = t$, and $\text{col} : T \rightarrow \{1, \dots, r\}$ the coloring. Define the PSD matrix (on $\text{range}(L)$)

$$M_t := \sum_{\substack{uv \in E \\ u, v \in T \\ \text{col}(u) = \text{col}(v)}} A_{uv}. \quad (38)$$

Thus M_t contains the contributions from edges whose endpoints are already colored and share the same color.

Let $R := V \setminus T$ be the uncolored vertices, $m := |R| = n - t$. For $v \in R$ and $\gamma \in \{1, \dots, r\}$ define the prospective increment obtained by coloring v with γ :

$$B_v^\gamma := \sum_{\substack{u \in T \\ \text{col}(u) = \gamma \\ uv \in E}} A_{uv}. \quad (39)$$

Then if we color v with color γ , we have $M_{t+1} = M_t + B_v^\gamma$.

Step 4: Inductive barrier invariant. Let $u_t := u_0 + t\delta$. We maintain the invariant

$$M_t \prec u_t I \quad \text{and} \quad \Phi_{u_t}(M_t) \leq \Phi_{u_0}(0) = \frac{d}{u_0}. \quad (40)$$

This holds at $t = 0$ since $M_0 = 0$.

Assume it holds for some $t < k$. Set $u = u_t$, $u' = u_{t+1} = u_t + \delta$, and

$$U := (u' I - M_t)^{-1}.$$

We claim there exists a choice of $(v, \gamma) \in R \times \{1, \dots, r\}$ for which the barrier condition (35) holds with $M = M_t$ and $B = B_v^\gamma$.

Consider the average over a uniformly random pair (v, γ) :

$$\frac{1}{mr} \sum_{v \in R} \sum_{\gamma=1}^r \left[\text{tr}(B_v^\gamma U) + \frac{\text{tr}(B_v^\gamma U^2)}{\Phi_u(M_t) - \Phi_{u'}(M_t)} \right]. \quad (41)$$

Observe that

$$\sum_{v \in R} \sum_{\gamma=1}^r B_v^\gamma = \sum_{\substack{uv \in E \\ u \in T, v \in R}} A_{uv} \preceq \sum_{e \in E} A_e = I \quad \text{on } \text{range}(L),$$

because the left-hand side is a sub-sum of the PSD matrices $\{A_e\}$ in (128). If $X \preceq Y$ and $C \succeq 0$, then $\text{tr}(XC) \leq \text{tr}(YC)$ because $\text{tr}(C^{1/2}(Y - X)C^{1/2}) \geq 0$. Applying this observation with $C = U$ and with $C = U^2$ (both positive semidefinite) we get

$$\sum_{v, \gamma} \text{tr}(B_v^\gamma U) \leq \text{tr}(U), \quad \sum_{v, \gamma} \text{tr}(B_v^\gamma U^2) \leq \text{tr}(U^2).$$

Therefore (41) is at most

$$\frac{\text{tr}(U)}{mr} + \frac{\text{tr}(U^2)}{mr(\Phi_u(M_t) - \Phi_{u'}(M_t))}. \quad (42)$$

By the inductive hypothesis, $\text{tr}(U) = \Phi_{u'}(M_t) \leq \Phi_u(M_t) \leq d/u_0$; the middle inequality uses that, for fixed M_t , the function $s \mapsto \Phi_s(M_t)$ decreases as the barrier level s increases. By (36),

$$\Phi_u(M_t) - \Phi_{u'}(M_t) \geq \delta \text{tr}(U^2),$$

so (in the non-trivial case $d > 0$, where $\text{tr}(U^2) > 0$) the second term in (42) is at most $1/(\delta mr)$. Hence the average (41) is at most

$$\frac{d/u_0}{mr} + \frac{1}{\delta mr}. \quad (43)$$

As long as $t < k = \lfloor n/4 \rfloor$, we have $m = n - t \geq 3n/4$ and $d \leq n$. Using the choices (37), we bound

$$\frac{d/u_0}{mr} \leq \frac{n/(\varepsilon/2)}{(3n/4) \cdot (16/\varepsilon)} = \frac{1}{6}, \quad \frac{1}{\delta mr} \leq \frac{1}{(\varepsilon/n) \cdot (3n/4) \cdot (16/\varepsilon)} = \frac{1}{12},$$

so the average (43) is < 1 . Therefore there exists at least one pair (v, γ) for which (35) holds. Applying Lemma 6.1 yields

$$M_{t+1} \prec u_{t+1}I \quad \text{and} \quad \Phi_{u_{t+1}}(M_{t+1}) \leq \Phi_{u_t}(M_t) \leq \frac{d}{u_0}.$$

Thus the invariant (162) propagates to $t + 1$, completing the induction for $t = 0, 1, \dots, k$.

Step 5: Extracting a large ε -light set. After k steps, the colored set T (with $|T| = k$) is partitioned into r color classes S_1, \dots, S_r . By definition of M_k ,

$$M_k = \sum_{a=1}^r \sum_{uv \in E(S_a, S_a)} A_{uv} = \sum_{a=1}^r L^{-1/2} L_{S_a} L^{-1/2} \quad \text{on } \text{range}(L).$$

From the invariant, $M_k \preceq u_k I$ with $u_k = u_0 + k\delta \leq \varepsilon/2 + \varepsilon/4 = 3\varepsilon/4$. Since each summand is PSD, each is dominated by the sum. Let S be the largest color class. Then

$$L^{-1/2} L_S L^{-1/2} \preceq M_k \preceq \frac{3\varepsilon}{4} I \preceq \varepsilon I \quad \text{on } \text{range}(L).$$

As explained in Step 1, this implies $L_S \preceq \varepsilon L$, i.e. $\varepsilon L - L_S \succeq 0$, so S is ε -light.

Step 6: Size lower bound. Among the k colored vertices, the largest color class has size at least k/r . If $n \geq 4$, then $k = \lfloor n/4 \rfloor \geq n/8$. Also,

$$r = \left\lceil \frac{16}{\varepsilon} \right\rceil \leq \frac{16}{\varepsilon} + 1 \leq \frac{32}{\varepsilon}.$$

Hence

$$|S| \geq \frac{k}{r} \geq \frac{n/8}{32/\varepsilon} = \frac{\varepsilon n}{256}.$$

The construction above was used only under the standing assumptions made at the beginning (in particular that the graph has at least one edge). If the graph is edgeless, taking $S = V$ is trivially ε -light. It remains to look at small values of n : for $1 \leq n \leq 3$ any single vertex set $S = \{v\}$ has $L_S = 0$ and hence is ε -light, and it satisfies $|S| = 1 \geq \varepsilon n/256$ because $\varepsilon \leq 1$. The cases $n = 0$ or $\varepsilon = 0$ were disposed of at the start. Thus in all cases there exists an ε -light set S with

$$|S| \geq \frac{\varepsilon}{256} |V|.$$

This proves the statement with the universal constant $c = 1/256$.

References

- [1] J. Batson, D. A. Spielman and N. Srivastava, *Twice-Ramanujan sparsifiers*, SIAM Journal on Computing **41** (2012), no. 6, 1704–1721.

7 Uniform lattices with 2–torsion arising as fundamental groups of closed manifolds with \mathbb{Q} –acyclic universal cover

Problem

Suppose that Γ is a uniform lattice in a real semisimple group, and that Γ contains some 2–torsion. Is it possible for Γ to be the fundamental group of a compact manifold without boundary whose universal cover is acyclic over the rational numbers \mathbb{Q} ?

Solution

We construct a uniform lattice Γ in a real semisimple group, containing a central element of order 2, and a closed manifold M with $\pi_1(M) \cong \Gamma$ such that \widetilde{M} is \mathbb{Q} –acyclic.

7.0.1 Step 0: A torsion–free uniform lattice in $SO^+(n, 1)$

Fix an odd integer $n = 2m + 1 \geq 5$. It is classical that there exist closed hyperbolic n –manifolds in every dimension. For example, standard arithmetic constructions using anisotropic quadratic forms over totally real fields give compact hyperbolic orbifolds in all dimensions (see [1] and the discussion in [4, Ch. II]). Passing to the identity component and using Selberg’s lemma [6] we obtain a torsion–free finite index subgroup. Thus we may choose a torsion–free cocompact lattice

$$L < SO^+(n, 1)$$

such that

$$N = L \backslash \mathbb{H}^n \tag{44}$$

is a closed orientable hyperbolic n –manifold (hence an aspherical $K(L, 1)$).

7.0.2 Step 1: The spin–lift lattice with central 2–torsion

Consider the spin covering

$$1 \longrightarrow \{\pm 1\} \longrightarrow \text{Spin}(n, 1) \xrightarrow{p} SO^+(n, 1) \longrightarrow 1.$$

Define

$$\Gamma := p^{-1}(L) < \text{Spin}(n, 1).$$

Then Γ is discrete and cocompact in $\text{Spin}(n, 1)$ (because p is a finite covering and L is discrete and cocompact in $SO^+(n, 1)$). Thus Γ is a uniform lattice in the real semisimple group $\text{Spin}(n, 1)$. We shall use below that Γ is finitely presented; this follows for lattices in Lie groups (for instance from the theorem of Borel–Serre on arithmetic groups, or from the fact that a compact fundamental domain gives a finite presentation, cf. [4, Ch. II, Thm. 4.7]).

Let $z := -1 \in \text{Spin}(n, 1)$ be the nontrivial element of the kernel. Then $z \in \Gamma$ is central and has order 2, hence Γ contains 2–torsion.

Remark 7.1 (Only 2–torsion). Since L is torsion–free, any finite order element $\gamma \in \Gamma$ satisfies $p(\gamma) \in L$ finite order, hence $p(\gamma) = 1$ and $\gamma \in \ker(p) = \{\pm 1\}$. Therefore the only nontrivial finite order element in Γ is the central involution z .

7.0.3 Step 2: A projective $\mathbb{Q}\Gamma$ –Poincaré complex by extension of scalars

Let $\mathbb{Q}\Gamma$ be the rational group ring, and define the central idempotent

$$e := \frac{1+z}{2} \in \mathbb{Q}\Gamma.$$

We now relate the corner algebra $e(\mathbb{Q}\Gamma)e$ to $\mathbb{Q}L$.

Lemma 7.2 (The corner algebra). *There is a canonical ring isomorphism*

$$e(\mathbb{Q}\Gamma)e \cong \mathbb{Q}L,$$

and $e\mathbb{Q}\Gamma$ is free as a left module over $e(\mathbb{Q}\Gamma)e$.

Proof. Choose a set-theoretic section $s : L \rightarrow \Gamma$ of the projection $p|_{\Gamma} : \Gamma \rightarrow L$ (i.e. $p(s(\lambda)) = \lambda$). Since $\ker(p) = \{\pm 1\} = \langle z \rangle$ is central, there is a 2-cocycle $\varepsilon : L \times L \rightarrow \{0, 1\}$ such that

$$s(\lambda)s(\mu) = s(\lambda\mu)z^{\varepsilon(\lambda,\mu)}.$$

Multiplying by e kills the ambiguity $z^{\varepsilon(\lambda,\mu)}$ because $ez = e$. Hence in $\mathbb{Q}\Gamma$ we have

$$(es(\lambda))(es(\mu)) = es(\lambda)s(\mu) = es(\lambda\mu)z^{\varepsilon(\lambda,\mu)} = es(\lambda\mu).$$

Thus the rule $\lambda \mapsto es(\lambda)$ defines a multiplicative map $\mathbb{Q}L \rightarrow e(\mathbb{Q}\Gamma)e$ sending the group basis of L to elements of the corner. If $s'(\lambda) = s(\lambda)z^{\delta(\lambda)}$ is another section, then $es'(\lambda) = es(\lambda)$, so the resulting map is independent of the choice of section. It is an isomorphism for the following elementary reason. The elements $\{s(\lambda), zs(\lambda) \mid \lambda \in L\}$ are pairwise distinct and form a subset of the group basis of $\mathbb{Q}\Gamma$; hence a relation $\sum a_{\lambda}es(\lambda) = 0$ would give $\sum a_{\lambda}s(\lambda) + \sum a_{\lambda}zs(\lambda) = 0$ and all a_{λ} are zero. Thus the displayed map identifies the bases $\{\lambda\}$ and $\{es(\lambda)\}$ of the two \mathbb{Q} -vector spaces, and it is onto because every element of the corner is a \mathbb{Q} -linear combination of the $es(\lambda)$.

For the module statement we use that the idempotent e is central. Consequently $e\mathbb{Q}\Gamma = e\mathbb{Q}\Gamma e$ (for $ex = exe$). Thus, viewed as a left module over the corner algebra $e(\mathbb{Q}\Gamma)e$, the module $e\mathbb{Q}\Gamma$ is just the regular module of this ring, in particular it is free of rank one. This proves the lemma. \square

Now fix a finite CW structure on N and let c_i be the number of i -cells. The universal cover $\tilde{N} \simeq \mathbb{H}^n$ is contractible, so the cellular chain complex $C_*(\tilde{N}; \mathbb{Q})$ is a finite free chain complex of right $\mathbb{Q}L$ -modules which resolves the trivial right $\mathbb{Q}L$ -module \mathbb{Q} . We extend scalars along $\mathbb{Q}L \cong e(\mathbb{Q}\Gamma)e$ by tensoring with the $(\mathbb{Q}L, \mathbb{Q}\Gamma)$ -bimodule $e\mathbb{Q}\Gamma$:

$$P_* := C_*(\tilde{N}; \mathbb{Q}) \otimes_{\mathbb{Q}L} e\mathbb{Q}\Gamma.$$

Then each chain group is

$$P_i \cong (e\mathbb{Q}\Gamma)^{c_i} \quad (0 \leq i \leq n), \quad (45)$$

hence P_* is a finite chain complex of finitely generated projective right $\mathbb{Q}\Gamma$ -modules.

Lemma 7.3 (Homology of P_*). *The complex P_* is acyclic in positive degrees and has $H_0(P_*) \cong \mathbb{Q}$ with the trivial Γ -action. More precisely,*

$$H_i(P_*) = 0 \quad (i > 0), \quad H_0(P_*) \cong \mathbb{Q}.$$

Proof. Because $e\mathbb{Q}\Gamma$ is free as a left $\mathbb{Q}L$ -module (Lemma 7.2), it is flat, so tensoring the exact augmented cellular complex of \tilde{N} over $\mathbb{Q}L$ preserves exactness in positive degrees. Thus $H_i(P_*) = 0$ for $i > 0$ and

$$H_0(P_*) \cong \mathbb{Q} \otimes_{\mathbb{Q}L} e\mathbb{Q}\Gamma \cong \mathbb{Q} \otimes_{\mathbb{Q}L} \mathbb{Q}L \cong \mathbb{Q}$$

as a vector space (using the corner isomorphism of Lemma 7.2 in the middle).

It remains to identify the residual right action of Γ . Let $[1 \otimes e]$ be the generator displayed above. If $\gamma = s(\mu)z^{\varepsilon} \in \Gamma$ then, in the tensor product over $\mathbb{Q}L$,

$$[1 \otimes e]\gamma = [1 \otimes es(\mu)z^{\varepsilon}] = [1 \otimes es(\mu)] = [1 \cdot \mu \otimes e] = [1 \otimes e],$$

because the augmentation makes $1 \cdot \mu = 1$ and $ez = e$. Thus the right action is trivial and $H_0(P_*) \cong \mathbb{Q}$ as a trivial right Γ -module. \square

For later use we make explicit that induction by the central idempotent is compatible with the duality which enters the definition of a symmetric Poincaré complex. Recall that our chain modules are right modules and that the dual of a right module is again regarded as a right module by means of the standard involution on the group ring.

Lemma 7.4 (Induction and duality). *Let $A = \mathbb{Q}\Gamma$ and put $A_+ = eAe (= eA)$. Via Lemma 7.2 we identify A_+ with $\mathbb{Q}L$. The functor*

$$- \otimes_{\mathbb{Q}L} eA : \text{Proj}(\mathbb{Q}L) \longrightarrow \text{Proj}(A)$$

identifies the category of finitely generated projective right $\mathbb{Q}L$ -modules with the full subcategory of finitely generated projective right A -modules P satisfying $Pe = P$. Under this identification the duality $P \mapsto \text{Hom}_A(P, A)$ (with the involution convention just mentioned) corresponds to the usual duality over $\mathbb{Q}L$. Consequently a symmetric Poincaré chain equivalence on a finite projective $\mathbb{Q}L$ -complex tensors to a symmetric Poincaré chain equivalence on the induced A -complex.

Proof. Because e is a central idempotent the ring A is the direct product $eA \times (1 - e)A$, and a right A -module is the same thing as a pair of modules over the two factors. The induction functor sends a $\mathbb{Q}L (= A_+)$ -module C to the pair $(C, 0)$; the inverse on the indicated full subcategory is $P \mapsto Pe = P$. Projectivity and exactness are preserved by this equivalence.

It remains only to note that the duals match. If $P = Pe$ and $f : P \rightarrow A$ is A -linear, then $f(P) = f(Pe) = f(P)e \subset Ae = eA$; hence

$$\text{Hom}_A(P, A) = \text{Hom}_A(P, eA) = \text{Hom}_{eAe}(Pe, eAe).$$

Via the isomorphism $eAe \cong \mathbb{Q}L$ this is precisely the ordinary dual of the corresponding $\mathbb{Q}L$ -module, and the involutions agree because $e^* = e$. Applying the equivalence degreewise to a chain complex shows that a chain equivalence and its adjoint remain such after tensoring, which is the Poincaré assertion. \square

The manifold N determines an n -dimensional (symmetric) Poincaré chain complex structure on $C_*(\tilde{N}; \mathbb{Q})$ over $\mathbb{Q}L$ (coming from the fundamental class and a cellular diagonal approximation). By Lemma 7.4 and the exactness of $- \otimes_{\mathbb{Q}L} e\mathbb{Q}\Gamma$ this Poincaré structure extends to P_* by extension of scalars. Thus:

Proposition 7.5 (A projective $\mathbb{Q}\Gamma$ -Poincaré complex). *The chain complex P_* is a finite n -dimensional projective Poincaré chain complex over $\mathbb{Q}\Gamma$ with $H_0(P_*) \cong \mathbb{Q}$ and $H_i(P_*) = 0$ for $i > 0$.*

It is useful to keep track of orientations for the later surgery step. The orientation module of this Poincaré complex (and hence of all subsequent \mathbb{Q} -Poincaré spaces we construct) is trivial: the manifold N is orientable and the homomorphism $\Gamma \rightarrow SO^+(n, 1)$ factors through the orientation-preserving group, while the central element z acts trivially on the fundamental class.

7.0.4 Step 3: Vanishing of Wall's finiteness obstruction and realization by a finite CW complex

We now compute the Wall finiteness obstruction and then appeal to the standard realization theorem over subrings of \mathbb{Q} .

The Wall finiteness obstruction of a finitely dominated CW complex (or, equivalently, of a finite projective chain complex modeling its universal cover) lies in the reduced group $\tilde{K}_0(\mathbb{Q}\Gamma)$. In our setting it is represented by the alternating sum of chain modules:

$$o(P_*) = \sum_{i=0}^n (-1)^i [P_i] = \sum_{i=0}^n (-1)^i c_i [e\mathbb{Q}\Gamma] = \chi(N) [e\mathbb{Q}\Gamma] \in K_0(\mathbb{Q}\Gamma). \quad (46)$$

Since N is a closed orientable manifold of odd dimension n , we have $\chi(N) = 0$. Therefore $o(P_*) = 0$ already in $K_0(\mathbb{Q}\Gamma)$, hence also in $\tilde{K}_0(\mathbb{Q}\Gamma)$.

We shall use the following precise form of Wall's realization theorem. Recall that Γ is finitely presented (Step 1).

Theorem 7.6 (Wall–Ranicki realization over localizations). *Let $R \subset \mathbb{Q}$ be a subring and let π be a finitely presented group. Let C_* be a finite chain complex of finitely generated projective right $R\pi$ -modules, concentrated in degrees $0, \dots, n$, with $H_i(C_*) = 0$ for $i > 0$ and $H_0(C_*) \cong R$ (trivial action). Its Wall obstruction is the class $\sum (-1)^i [C_i] \in \tilde{K}_0(R\pi)$. If this class is zero, then there exists a finite connected CW complex X with $\pi_1(X) \cong \pi$ such that*

$$C_*(\tilde{X}; R) \simeq C_*$$

as chain complexes of $R\pi$ -modules. If in addition C_ carries an n -dimensional symmetric Poincaré chain-complex structure, X may be chosen to be a finite R -Poincaré complex of formal dimension n (elementary algebraic expansions do not change the Poincaré type).*

For the reader's convenience let us also indicate references for this statement. Wall constructs, from an algebraic $R\pi$ -complex (for R a localization of \mathbb{Z} , in particular a subring of \mathbb{Q}), a finitely dominated CW complex with the prescribed cellular $R\pi$ -chains and proves that the obstruction above is the only one to making it finite; see Wall [8, Thm. F and §§ 2–5]. The algebraic reformulation for projective complexes, and the fact that elementary expansions do not change symmetric structures, is spelled out in Ranicki [5]. Let us also make explicit the small piece of algebra which allows one to pass from rational matrices to honest cell attachments. If the obstruction vanishes, adding finitely many elementary contractible projective complexes replaces a finite projective resolution by a *based free* resolution. Choose a finite presentation of π and choose the bases in degrees 0 and 1 (and a relator summand in degree 2) so that the corresponding part of the differential is the cellular differential of the presentation 2-complex. By elementary changes of basis over $R\pi$ one may further suppose that any additional basis vectors in degree 2 have zero boundary (the relator columns already generate $\ker \partial_1$). In the present paper we only need the case $R = \mathbb{Q}$. The remaining boundary matrices have entries in $\mathbb{Q}\pi$, and then there is a simple simultaneous clearing of denominators: writing these matrices in the chosen bases, pick diagonal change-of-basis matrices T_i (with non-zero integer entries) inductively for $i \geq 2$, starting with $T_0 = T_1 = 1$, so that $T_{i-1}^{-1} \partial_i T_i$ has coefficients in $\mathbb{Z}\pi$ for every i (for a fixed T_{i-1} one merely takes the entries of the diagonal of T_i large enough to clear the denominators in each column). The conjugated complex is isomorphic to the original over $R\pi$. Wall's cellular realization then starts from the presentation 2-complex for π and attaches cells in dimensions ≥ 3 to realize this integral free complex up to R -chain equivalence; the relative Hurewicz theorem identifies the necessary homotopy and homology classes.

Proposition 7.7 (Realization by a finite CW complex). *There exists a finite CW complex X with $\pi_1(X) \cong \Gamma$ such that, as chain complexes of right $\mathbb{Q}\Gamma$ -modules,*

$$C_*(\tilde{X}; \mathbb{Q}) \simeq P_*. \tag{47}$$

Moreover X is a finite \mathbb{Q} -Poincaré complex of formal dimension n , and \tilde{X} is \mathbb{Q} -acyclic in positive degrees.

Proof. The computation (46) gives a vanishing obstruction. Applying Theorem 7.6 with $R = \mathbb{Q}$, $\pi = \Gamma$ and $C_* = P_*$ yields a finite CW complex whose cellular rational chain complex is chain-homotopy equivalent to P_* . The last assertions follow from the Poincaré structure on P_* (Proposition 7.5) and from Lemma 7.3. \square

For later use let us spell out explicitly why the finite complex just obtained is a genuine Poincaré space (and not merely a complex with the right algebraic chains). The symmetric

Poincaré chain equivalence on the finite projective resolution P_* is equivalent to the following cohomological statement: there is a class

$$\alpha \in H_n(\Gamma; \mathbb{Q})$$

whose cap product induces isomorphisms $H^i(\Gamma; M) \xrightarrow{\cong} H_{n-i}(\Gamma; M)$ (with the usual orientation twist) for every $\mathbb{Q}\Gamma$ -module M . This is simply another way of expressing the algebraic Poincaré condition; see for example Brown's discussion of duality over coefficient rings in [2, Ch. VIII, §§ 9–10] or Ranicki [5].

Since $C_*(\tilde{X}; \mathbb{Q})$ is a free resolution of the trivial module, the classifying map induces isomorphisms $H_*(X; M) \cong H_*(\Gamma; M)$ and $H^*(X; M) \cong H^*(\Gamma; M)$ for every local system of \mathbb{Q} -vector spaces M . Let $[X] \in H_n(X; \mathbb{Q})$ be the image of α under this identification (equivalently, using the rational homology equivalence $g : X \rightarrow N$ constructed in Step 4, the inverse image of the hyperbolic fundamental class). By the naturality of the cap product, capping with $[X]$ yields the Poincaré duality isomorphisms on X with arbitrary local coefficients. Thus X is a finite \mathbb{Q} -Poincaré complex of formal dimension n , and the duality agrees with the algebraic symmetric structure transported from P_* .

7.0.5 Step 4: construction of a *geometric rational normal invariant*

The point at which one has to be a little careful is the passage from the algebraic Poincaré complex of Step 2 to a datum to which the (geometric) surgery machine applies. We give the details here. We shall exhibit an honest stable real vector bundle over the finite complex X which is a *rational reduction* of the Spivak normal fibration of X , and we shall check that the algebraic normal complex associated to it is exactly the one obtained by tensoring the normal complex of the hyperbolic manifold N with the idempotent module.

We begin by fixing a map to the hyperbolic manifold. Let $c_X : X \rightarrow B\Gamma$ be the classifying map of the universal covering of the finite complex furnished by Proposition 7.7. Since $C_*(\tilde{X}; \mathbb{Q})$ is a free resolution of the trivial $\mathbb{Q}\Gamma$ -module, c_X induces isomorphisms on homology with trivial rational coefficients. The extension $1 \rightarrow \{\pm 1\} \rightarrow \Gamma \xrightarrow{p} L \rightarrow 1$ gives a map of classifying spaces $Bp : B\Gamma \rightarrow BL$ and, because the order of the kernel is invertible in \mathbb{Q} , the transfer (or the Lyndon–Hochschild–Serre spectral sequence) shows that Bp is a \mathbb{Q} -homology equivalence [2, Ch. VII, § 6]. Finally BL is (canonically up to homotopy) the aspherical manifold $N = L \backslash \mathbb{H}^n$. We choose once and for all a cellular representative of the composite

$$g : X \xrightarrow{c_X} B\Gamma \xrightarrow{Bp} BL \simeq N. \quad (48)$$

It is a rational homology equivalence. Orient X in such a way that $g_*[X] = [N]$ in $H_n(-; \mathbb{Q})$.

Let ν_N be the stable normal bundle of the smooth manifold N and put

$$\xi := g^* \nu_N \quad (49)$$

(up to adding a trivial summand in order to have an honest vector bundle of large rank over the finite skeleton of BO). We shall use the following elementary fact, which is often left implicit in accounts of rational surgery.

Lemma 7.8 (Rational reductions of the Spivak fibration). *Let Y^n be a finite \mathbb{Q} -Poincaré complex and let η be an oriented stable real vector bundle over Y . Denote by $U \in H^r(T(\eta); \mathbb{Q})$ the rational Thom class ($r = \text{rank } \eta$). Then there is, after adding a trivial summand to η if necessary, a stable map of spectra*

$$\rho : S^{n+r} \longrightarrow T(\eta) \quad (50)$$

such that the Hurewicz image of ρ is the Thom image of the fundamental class of Y (over \mathbb{Q}). Consequently (η, ρ) is a \mathbb{Q} -normal invariant of Y : the sphere fibration of η is fibre homotopy equivalent to the Spivak normal fibration after rationalization and the cap product with the Thom class realizes the prescribed Poincaré duality.

Proof. We spell out the standard facts which enter. First of all a finite R -Poincaré complex (with R a localization of \mathbb{Z}) has an R -Spivak normal fibration: this is a stable oriented spherical fibration ν_Y^R over Y together with an R -local Thom–Pontryagin collapse $c_\nu : S^{n+k} \rightarrow T(\nu_Y)$ whose cap product realizes the Poincaré fundamental class. The construction and uniqueness are due to Spivak and Wall; for localizations see Sullivan’s notes on localization and, in the form used in surgery with fundamental group, Hausmann [3, § 1, Prop. 1.1 and Cor. 1.4] or Ranicki [5, Ch. 9, esp. §§9.2–9.3]. We shall use this result with $R = \mathbb{Q}$.

We next recall why such a fibration is automatically reducible after rationalization. Let BSG denote the classifying space of stable *oriented* spherical fibrations (the identity component of the stable group of homotopy self-equivalences of spheres). For $i > 0$ one has $\pi_i(BSG) = \pi_{i-1}^S$, the stable homotopy groups of spheres. By Serre’s theorem these groups are finite. Consequently the rationalization $BSG_{(0)}$ is contractible (the spaces involved are nilpotent). If η is an oriented stable real vector bundle over Y , its underlying stable sphere fibration is classified by the composite $Y \rightarrow BO \rightarrow BSG$; after rationalization this map is null. The same is true for the Spivak fibration, and there is therefore a fibre-homotopy equivalence of stable spherical fibrations over Y

$$J(\eta)_{(0)} \simeq \nu_{Y(0)}. \quad (51)$$

(The choice is unique up to homotopy.) Stabilizing by adding a trivial bundle we may suppose that the two representatives have the same fibre dimension, say r .

The equivalence (51) induces a rational equivalence of Thom spectra and carries the rational Thom class of the Spivak fibration to a Thom class $U \in H^r(T(\eta); \mathbb{Q})$. Composing the Spivak collapse c_ν with the inverse Thom-space equivalence gives a stable map in the rational stable category $S_{(0)}^{n+r} \rightarrow T(\eta)_{(0)}$. Since Thom spaces of bundles over finite complexes are finite spectra, such a rational map can be represented by an honest stable map after multiplying by a non-zero integer: the stable Hurewicz theorem (or, equivalently, the Atiyah–Hirzebruch spectral sequence together with the finiteness of the stable stems; see [7, Ch. IX, Thm. 3.1]) shows that $\pi_*^s(E) \otimes \mathbb{Q} \cong H_*(E; \mathbb{Q})$ for finite spectra. We choose a representative and, if necessary, rescale the Thom class by the inverse integer (over the coefficient field \mathbb{Q}). This gives the map (50). By construction its Hurewicz image is precisely the Thom image of $[Y]$, and capping with the Thom class (equivalently, transporting the Thom class of the Spivak fibration) yields the Poincaré duality isomorphisms with arbitrary local systems of \mathbb{Q} -vector spaces. Thus (η, ρ) is a rational reduction of the Spivak fibration, i.e. a \mathbb{Q} -normal invariant in the sense used in rational surgery. \square

We apply the lemma to the bundle $\xi = g^*\nu_N$. It remains to check that this *geometric* invariant is compatible with the algebraic calculations of Step 5. We make this explicit because the idempotent construction is slightly unusual.

Proposition 7.9 (Identification with the idempotent normal complex). *Let the chain equivalence (47) be chosen as follows. Lift the map (48) to a cellular p -equivariant map of universal covers $\tilde{g} : \tilde{X} \rightarrow \tilde{N} = \mathbb{H}^n$; it induces a chain map of right $\mathbb{Q}\Gamma$ -complexes*

$$C_*(\tilde{X}; \mathbb{Q}) \longrightarrow C_*(\tilde{N}; \mathbb{Q}) \otimes_{\mathbb{Q}L} e\mathbb{Q}\Gamma = P_* \quad (52)$$

(the module $e\mathbb{Q}\Gamma$ is just the regular $\mathbb{Q}L$ -module with the right action inflated along p , so the formula is the evident one). Since both complexes are projective resolutions of the trivial module, the comparison theorem shows that (52) is a chain-homotopy equivalence; we take this as the equivalence in (47). With this choice the algebraic normal structure on $C_*(\tilde{X}; \mathbb{Q})$ associated, in Ranicki’s sense, to the rational normal invariant (ξ, ρ) of Lemma 7.8 is exactly the structure obtained from the normal complex of the manifold N by tensoring over $\mathbb{Q}L$ with the bimodule $e\mathbb{Q}\Gamma$.

Proof. Only naturality has to be checked. The cellular symmetric Poincaré structure on $C_*(\tilde{N}; \mathbb{Q})$ is obtained from a diagonal approximation and from the fundamental class of N ; the algebraic normal structure associated to ν_N is represented by a Thom cocycle for this bundle. Pulling the Thom cocycle back by g gives the Thom cocycle of ξ , and our orientation of X was chosen so that fundamental classes correspond. The diagonal approximation on X may be taken cellularly natural with respect to g (two choices give chain-homotopic symmetric structures). Ranicki’s “symmetric construction” and his definition of an algebraic normal complex are functorial for exact functors of projective module categories compatible with involution and for maps of spaces: see [5, Chs. 1–3]. Applying this functoriality to the exact functor $-\otimes_{\mathbb{Q}L} e\mathbb{Q}\Gamma$ and to the chain equivalence (52) gives the asserted identification. Concretely, on the level of modules the equality follows from the identities $es(\lambda\mu) = es(\lambda)s(\mu)$ (Lemma 7.2), and on homology from the naturality of cap product. \square

In particular the finite Poincaré complex X constructed in Step 3 is equipped with an explicit stable vector bundle (indeed the pull-back of a bundle from a manifold) representing the algebraic normal datum whose surgery obstruction we compute next. This fills the gap between the idempotent algebra and the geometric input required by the rational surgery theorem.

7.0.6 Step 5: Vanishing of the rational surgery obstruction

Let $n = 2m + 1 \geq 5$. Given an n -dimensional \mathbb{Q} -Poincaré complex X with fundamental group Γ and a chosen normal invariant, there is a surgery obstruction in the (quadratic/symmetric) L -group $L_n^h(\mathbb{Q}\Gamma)$ for the standard involution (we are in the orientable, trivial-orientation case); if it vanishes then X is normally cobordant to a closed manifold, and in fact can be realized by a closed manifold mapping to X inducing π_1 -isomorphism and $\mathbb{Q}\Gamma$ -homology equivalence. This is the rational surgery theorem in the presence of fundamental group, due to Hausmann and Ranicki (building on Sullivan–Wall), see [3, 5].

The key point in this construction is that the surgery obstruction of this normal datum is zero for a very concrete reason. Because the idempotent e is central (and is fixed by the involution) the rational group ring splits as a product of rings with involution

$$\mathbb{Q}\Gamma \cong e\mathbb{Q}\Gamma e \times (1 - e)\mathbb{Q}\Gamma(1 - e),$$

and algebraic L -theory is additive for such products. The normal complex chosen in Step 4 was obtained on the projective complex P_* . The finite CW complex X of Step 3 has cellular chains free over $\mathbb{Q}\Gamma$, but the chain equivalence (47) is obtained from P_* by adding elementary (contractible) summands. In the cobordism group $L_n^h(\mathbb{Q}\Gamma)$ such summands represent the zero element. Consequently the surgery obstruction of X is represented by the induced normal complex on P_* , and all of its chain modules are of the form $(e\mathbb{Q}\Gamma)^{c_i}$; the class therefore lies entirely in the first factor of the displayed product. Under the ring isomorphism $e\mathbb{Q}\Gamma e \cong \mathbb{Q}L$ this first component is precisely the result of tensoring the identity normal complex of the manifold N with the bimodule $e\mathbb{Q}\Gamma$.

Now the identity normal map $\text{id}_N : N \rightarrow N$ has surgery obstruction $0 \in L_n^h(\mathbb{Q}L)$ (indeed it is already a manifold). Algebraic surgery is functorial for exact functors of projective module categories compatible with duality: applying $-\otimes_{\mathbb{Q}L} e\mathbb{Q}\Gamma$ to a null-cobordism of the identity normal complex gives a null-cobordism of the induced normal complex (see Ranicki [5]). Thus the component of the obstruction in $L_n^h(e\mathbb{Q}\Gamma e)$ is zero, and hence the obstruction of X in $L_n^h(\mathbb{Q}\Gamma)$ is zero as well.

7.0.7 Step 6: Rational surgery produces a closed manifold with \mathbb{Q} -acyclic universal cover

We now apply the rational surgery theorem.

Theorem 7.10 (Rational surgery realization). *Let $R \subset \mathbb{Q}$ be a subring (equivalently a localization of \mathbb{Z}) and let Y^q be a finite connected R -Poincaré complex, $q \geq 5$, with fundamental group π . Suppose that Y is equipped with an R -normal invariant: by this we mean, in the topological language, a stable real vector bundle (or stable spherical fibration) over Y together with an R -Thom class whose cap product realizes the Poincaré fundamental class, and in the algebraic language equivalently an algebraic normal structure on the symmetric chain complex $C_*(\tilde{Y}; R)$ in Ranicki's sense. Associated to this datum is a surgery obstruction*

$$\sigma(Y) \in L_q^h(R\pi)$$

(for the standard orientation character). *If $\sigma(Y) = 0$, then there exists a closed smooth q -manifold M and a normal map $f : M \rightarrow Y$ representing the chosen invariant such that*

$$\pi_1(M) \xrightarrow{\cong} \pi_1(Y) = \pi, \quad f_* : H_*(M; R\pi) \xrightarrow{\cong} H_*(Y; R\pi).$$

In dimensions $q \not\equiv 0 \pmod{4}$ no further numerical conditions occur; in dimension $4k$ the equality of the signature with the Hirzebruch L -class determined by the chosen Pontryagin data is the additional requirement built into the choice of normal invariant.

For the reader who wants precise sources we spell out where this statement is proved. Hausmann constructs, for localizations of \mathbb{Z} , normal maps which are homology equivalences modulo the chosen localization and develops the corresponding obstruction theory in [3, §§ 1–3]. The translation between such (localized) normal invariants and algebraic normal complexes, and the functoriality of the obstruction under exact functors of module categories, is part of Ranicki's algebraic surgery framework; see [5]. Because in our application the normal invariant is represented by an honest stable vector bundle (Lemma 7.8 and Proposition 7.9) the output of the theorem may be taken in the smooth category.

We apply Theorem 7.10 with $R = \mathbb{Q}$, $q = n = 2m + 1$, $\pi = \Gamma$, and $Y = X$ equipped with the normal invariant chosen in Step 4. The obstruction vanishes by Step 5. Hence we obtain:

Proposition 7.11. *There exists a closed smooth n -manifold M and a map $f : M \rightarrow X$ inducing an isomorphism on π_1 and an isomorphism on $\mathbb{Q}\Gamma$ -homology:*

$$\pi_1(M) \xrightarrow{\cong} \pi_1(X) = \Gamma, \quad H_*(M; \mathbb{Q}\Gamma) \xrightarrow{\cong} H_*(X; \mathbb{Q}\Gamma).$$

7.0.8 Step 7: The universal cover of M is \mathbb{Q} -acyclic

Let $\tilde{f} : \tilde{M} \rightarrow \tilde{X}$ be the lift of f to universal covers. Recall that with our convention of cellular *right* Γ -modules the chain complex computing homology of a connected complex with coefficients in the left regular module $\mathbb{Q}\Gamma$ is $C_*(\tilde{Y}; \mathbb{Z}) \otimes_{\mathbb{Z}\Gamma} \mathbb{Q}\Gamma$ for such a complex Y , which identifies (by choosing one lift of each cell) with the ordinary cellular complex of the universal cover with rational coefficients. Thus the $\mathbb{Q}\Gamma$ -homology equivalence in Proposition 7.11 is equivalently the statement that \tilde{f} induces an isomorphism

$$H_*(\tilde{M}; \mathbb{Q}) \xrightarrow{\cong} H_*(\tilde{X}; \mathbb{Q}).$$

Using Proposition 7.7 and Lemma 7.3 we have

$$H_i(\tilde{X}; \mathbb{Q}) \cong H_i(P_*) = \begin{cases} \mathbb{Q}, & i = 0, \\ 0, & i > 0. \end{cases}$$

Therefore

$$H_i(\tilde{M}; \mathbb{Q}) = \begin{cases} \mathbb{Q}, & i = 0, \\ 0, & i > 0, \end{cases}$$

so the universal cover \tilde{M} is \mathbb{Q} -acyclic in positive degrees.

7.0.9 Conclusion

We have produced, for each odd $n = 2m + 1 \geq 5$, a uniform lattice $\Gamma < \text{Spin}(n, 1)$ containing the central involution $z = -1$ and a closed smooth n -manifold M with $\pi_1(M) \cong \Gamma$ such that \widetilde{M} is acyclic over \mathbb{Q} .

Theorem 7.12. *Yes: there exist uniform lattices Γ in real semisimple groups containing 2-torsion which occur as the fundamental groups of closed manifolds whose universal covers are \mathbb{Q} -acyclic.*

7.0.10 Compatibility with complete Euler characteristic obstructions

For context, we briefly explain why this example does not contradict familiar torsion/Euler characteristic obstructions. Brown defines the *complete Euler characteristic* $\tilde{\chi}(\Gamma)$ of a group of finite type, whose coefficients at conjugacy classes of finite order elements can be expressed, for cocompact lattices, in terms of Euler characteristics of centralizers (see [2, Ch. IX, §7]). In many settings, existence of a finite \mathbb{Q} -acyclic universal cover forces these coefficients to vanish away from the identity class.

In our example, by Remark 7.1, the only nontrivial finite order element is the central involution z , and its centralizer is all of Γ . The (rational) Euler characteristic of a group with a finite normal subgroup is multiplicative with the factor $1/|F|$ (Brown [2, Ch. IX, §7]); consequently

$$\chi_{\mathbb{Q}}(\Gamma) = \chi(L)/2 = \chi(N)/2.$$

Equivalently, applying the Hattori–Stallings trace to the Wall element (46) shows that the complete Euler characteristic has coefficient $\chi(N)/2$ both at the identity class and at the class of z (and no other finite classes). For odd n we have $\chi(N) = 0$, so the coefficient at z indeed vanishes. In contrast, in even dimensions the Gauss–Bonnet formula for a compact hyperbolic manifold gives $\chi(N) \neq 0$, and already the finiteness obstruction (46) is detected by this trace and is nonzero. Thus the odd-dimensional spin-lift construction gives a genuine positive answer precisely in the case compatible with these obstructions.

References

- [1] A. Borel and Harish–Chandra, *Arithmetic subgroups of algebraic groups*, Ann. of Math. (2) **75** (1962), 485–535.
- [2] K. S. Brown, *Cohomology of Groups*, Graduate Texts in Mathematics **87**, Springer–Verlag, New York, 1982.
- [3] J.-C. Hausmann, *Manifolds with a given homology and fundamental group*, Comment. Math. Helv. **53** (1978), 113–134.
- [4] M. S. Raghunathan, *Discrete Subgroups of Lie Groups*, Ergebnisse der Mathematik und ihrer Grenzgebiete **68**, Springer–Verlag, 1972.
- [5] A. Ranicki, *Algebraic and Geometric Surgery*, Oxford Mathematical Monographs, Oxford University Press, 2002.
- [6] A. Selberg, *On discontinuous groups in higher-dimensional symmetric spaces*, in *Contributions to Function Theory (Internat. Colloq. Function Theory, Bombay, 1960)*, Tata Institute of Fundamental Research, Bombay, 1960, pp. 147–164.
- [7] R. M. Switzer, *Algebraic Topology—Homotopy and Homology*, Classics in Mathematics, Springer–Verlag, Berlin, 2002.
- [8] C. T. C. Wall, *Finiteness conditions for CW complexes*, Ann. of Math. (2) **81** (1965), 56–69.

8 Quadrivalent polyhedral Lagrangian surfaces are Lagrangian-smoothable

Problem

A polyhedral Lagrangian surface K in \mathbb{R}^4 is a finite polyhedral complex all of whose faces are Lagrangians, and which is a topological submanifold of \mathbb{R}^4 . A Lagrangian smoothing of K is a Hamiltonian isotopy K_t of smooth Lagrangian submanifolds, parameterised by $(0, 1]$, extending to a topological isotopy, parametrised by $[0, 1]$, with endpoint $K_0 = K$.

Let K be a polyhedral Lagrangian surface with the property that exactly 4 faces meet at every vertex. Does K necessarily have a Lagrangian smoothing?

Solution

Yes. We construct explicit local smoothing models at each vertex and each edge, glue these models into K inside pairwise disjoint *affine symplectic* (hence Darboux) neighborhoods chosen compatibly from the outset, and then verify that the resulting Lagrangian isotopy has vanishing flux and hence is Hamiltonian.

0. Conventions and standing hypotheses.

Symplectic conventions. Work in (\mathbb{R}^4, ω) with ω exact. Fix once and for all a global primitive

$$\lambda := p_1 dq_1 + p_2 dq_2$$

in some fixed global Darboux coordinates (q_1, q_2, p_1, p_2) , so $d\lambda = \omega$. In these coordinates

$$\omega = dp_1 \wedge dq_1 + dp_2 \wedge dq_2, \quad \omega(\partial_{p_i}, \partial_{q_j}) = \delta_{ij}, \quad \omega(\partial_{q_i}, \partial_{p_j}) = -\delta_{ij}.$$

We use the Hamiltonian convention

$$\iota_{X_H} \omega = -dH. \tag{53}$$

(With $\omega = \sum dp_i \wedge dq_i$, this is the convention giving Hamilton's equations $\dot{q}_i = \partial H / \partial p_i$ and $\dot{p}_i = -\partial H / \partial q_i$.)

If λ_{std} denotes the standard primitive in some other Darboux chart, then $\lambda_{\text{std}} - \lambda$ is a closed 1-form on \mathbb{R}^4 and hence exact. Thus, whenever we work in a local Darboux chart, we will freely replace the local primitive by adding an exact correction term so that all computations ultimately refer to the same fixed global λ .

Polyhedral hypotheses. We make explicit the local cell-structure properties used below.

- $K \subset \mathbb{R}^4$ is a compact (or, more generally, properly embedded) topological 2-manifold *without boundary*.
- K is given a *face-to-face* finite polyhedral structure: each 2-cell is a compact convex polygon contained in an affine Lagrangian plane, and the intersection of any two cells is a (possibly empty) common face of each.
- Every vertex is *quadrivalent*: exactly four 2-faces meet at each vertex (equivalently, the link of each vertex in K is a 4-cycle).

Under these hypotheses, each point of the interior of an edge has a neighborhood in K homeomorphic to a disk and modeled by exactly two faces meeting along that edge (no “branching” along edges). Likewise, at a vertex, the four incident faces appear as four planar sectors meeting cyclically.

1. Linear normal form at a quadrivalent vertex, including all sign patterns.

Let x be a vertex of K , and let l_1, l_2, l_3, l_4 be the four *oriented* edge rays emanating from x , numbered in cyclic order in the link. Let

$$P_i := \langle l_i, l_{i+1} \rangle \quad (i \bmod 4)$$

denote the affine Lagrangian planes of the four incident faces (after translating x to the origin we identify these planes with linear Lagrangian subspaces). Each incident face is a planar sector in P_i bounded by l_i and l_{i+1} .

Lemma 8.1 (Eliminating coplanar adjacencies). *Let $K \subset (\mathbb{R}^4, \omega)$ be a polyhedral Lagrangian surface, and let x be a vertex with emanating rays l_1, l_2, l_3, l_4 in cyclic order, and $P_i = \langle l_i, l_{i+1} \rangle$ as above.*

Assume that two adjacent faces are coplanar, i.e. $P_i = P_{i+1}$ for some i . Then x is not a genuine quadrivalent singularity of the underlying embedded surface K : after a cyclic relabeling (so $P_1 = P_2$), exactly one of the following holds.

- (a) $P_1 = P_2 = P_3 = P_4$ (so x is a smooth point of K).
- (b) l_1 and l_3 are collinear (opposite rays on a line ℓ) and $P_3 = P_4$ (so x lies in the interior of a geometric edge of K).

Consequently, since the existence of a Lagrangian smoothing depends only on the subset K and not on the chosen polyhedral structure, one may, without loss of generality, discard such inessential vertices and assume in the vertex analysis that no two adjacent faces are coplanar.

Proof. By cyclic relabeling assume $P_1 = P_2 =: L$. Choose nonzero vectors $v_i \in l_i$. Then $v_1, v_2, v_3 \in L$.

First suppose that v_1 and v_3 are linearly independent. Then $L = \langle v_1, v_3 \rangle$. Since $P_4 = \langle v_4, v_1 \rangle$ is Lagrangian, $\omega(v_4, v_1) = 0$, and since $P_3 = \langle v_3, v_4 \rangle$ is Lagrangian, $\omega(v_4, v_3) = 0$. Hence v_4 is ω -orthogonal to $\langle v_1, v_3 \rangle = L$, i.e. $v_4 \in L^\omega$. Because L is Lagrangian, $L^\omega = L$, so $v_4 \in L$. It follows that $P_3 = \langle v_3, v_4 \rangle \subset L$ and $P_4 = \langle v_4, v_1 \rangle \subset L$, hence $P_3 = P_4 = L$, proving (a). Since K is a topological 2-manifold and lies in the affine plane L near x , x is a smooth point.

Now suppose that v_1 and v_3 are linearly dependent. Then l_1 and l_3 are collinear, hence opposite rays on a common line ℓ . Therefore

$$P_3 = \langle l_3, l_4 \rangle = \langle \ell, l_4 \rangle = \langle l_1, l_4 \rangle = P_4,$$

proving (b). The union of the two coplanar faces in $P_1 = P_2$ is a polyhedral half-plane in P_1 bounded by ℓ , and similarly the union of the two coplanar faces in $P_3 = P_4$ is a half-plane in P_3 bounded by ℓ . Thus the germ at x is an “edge wedge” and x is an interior point of that geometric edge.

In either case x is inessential as a vertex and may be removed from the polyhedral structure without changing the subset K . \square

Henceforth assume *no two adjacent faces are coplanar*. Fix nonzero vectors $v_i \in l_i$. Then $v_3 \notin \langle v_1, v_2 \rangle$, so

$$\dim \langle v_1, v_2, v_3, v_4 \rangle \geq 3. \quad (54)$$

Because each P_i is Lagrangian, we have

$$\omega(v_1, v_2) = \omega(v_2, v_3) = \omega(v_3, v_4) = \omega(v_4, v_1) = 0.$$

Thus the only potentially nonzero pairings among the v_i are $\omega(v_1, v_3)$ and $\omega(v_2, v_4)$. If both were zero, then ω would vanish on $\langle v_1, v_2, v_3, v_4 \rangle$, making it isotropic of dimension ≥ 3 , impossible in a symplectic 4-space. Hence at least one of $\omega(v_1, v_3)$ or $\omega(v_2, v_4)$ is nonzero; after a cyclic relabeling we may assume

$$\omega(v_1, v_3) \neq 0. \quad (55)$$

Transversality of opposite planes. Under the standing hypothesis that no adjacent faces are coplanar, the opposite planes $P_1 = \langle v_1, v_2 \rangle$ and $P_3 = \langle v_3, v_4 \rangle$ are transverse. Indeed, if $P_1 \cap P_3 \neq \{0\}$ then there exists $0 \neq w = av_3 + bv_4 \in P_1$. Pairing with v_1 gives

$$0 = \omega(v_1, w) = a\omega(v_1, v_3) + b\omega(v_1, v_4) = a\omega(v_1, v_3),$$

since $\omega(v_1, v_4) = 0$ (adjacent in the cyclic order). Thus $a = 0$ and $w = bv_4 \in P_1$, so $v_4 \in P_1$ and hence $P_4 = \langle v_4, v_1 \rangle \subset P_1$, i.e. P_4 and P_1 are coplanar, contradiction. Therefore $P_1 \cap P_3 = \{0\}$.

Since P_1 and P_3 are transverse Lagrangian planes, ω induces a nondegenerate pairing

$$\omega|_{P_3 \times P_1} : P_3 \times P_1 \rightarrow \mathbb{R}, \quad (w, u) \mapsto \omega(w, u),$$

because if $w \in P_3$ pairs trivially with all $u \in P_1$, then $w \in P_1^\omega = P_1$, hence $w = 0$.

Vertex Darboux coordinates and sign pattern. Choose a basis $e_1 = v_1, e_2 = v_2$ of P_1 . Let f^1, f^2 be the ω -dual basis of P_3 with respect to the pairing $\omega|_{P_3 \times P_1}$, i.e.

$$\omega(f^j, e_i) = \delta_i^j. \quad (56)$$

Write

$$v_3 = a_1 f^1 + a_2 f^2, \quad v_4 = b_1 f^1 + b_2 f^2.$$

Because $P_2 = \langle v_2, v_3 \rangle$ is Lagrangian, $\omega(v_2, v_3) = 0$. Using $\omega(e_2, f^1) = -\omega(f^1, e_2) = 0$ and $\omega(e_2, f^2) = -\omega(f^2, e_2) = -1$, we obtain

$$0 = \omega(v_2, v_3) = a_1 \omega(e_2, f^1) + a_2 \omega(e_2, f^2) = -a_2,$$

so $a_2 = 0$ and hence $v_3 = a_1 f^1$. Similarly, since $P_4 = \langle v_4, v_1 \rangle$ is Lagrangian we have $\omega(v_1, v_4) = 0$, and using $\omega(e_1, f^1) = -1$ and $\omega(e_1, f^2) = 0$ we get $b_1 = 0$, hence $v_4 = b_2 f^2$.

Thus v_3 is a nonzero scalar multiple of f^1 and v_4 is a nonzero scalar multiple of f^2 . Define the *vertex sign pattern*

$$\sigma_1 := \text{sign}(a_1) \in \{\pm 1\}, \quad \sigma_2 := \text{sign}(b_2) \in \{\pm 1\}. \quad (57)$$

These signs are determined by the oriented rays l_3 and l_4 and *cannot* be changed by rescaling rays, since only *positive* rescalings of half-lines are allowed.

Rescale v_3 and v_4 by positive factors so that $v_3 = \sigma_1 f^1$ and $v_4 = \sigma_2 f^2$. Let (q_1, q_2, p_1, p_2) be the linear symplectic coordinate system determined by declaring

$$e_i = \partial_{q_i}, \quad f^i = \partial_{p_i},$$

so $\omega = dp_1 \wedge dq_1 + dp_2 \wedge dq_2$ and $\lambda_{\text{std}} = p_1 dq_1 + p_2 dq_2$ in these coordinates. Then the four rays take the *signed* standard form

$$l_1 = \mathbb{R}_+ \partial_{q_1}, \quad l_2 = \mathbb{R}_+ \partial_{q_2}, \quad l_3 = \mathbb{R}_+ (\sigma_1 \partial_{p_1}), \quad l_4 = \mathbb{R}_+ (\sigma_2 \partial_{p_2}). \quad (58)$$

In particular, there are *four* possible sign patterns $(\sigma_1, \sigma_2) \in \{\pm 1\}^2$ at a vertex.

Since each face is a planar sector bounded by the corresponding rays, after shrinking to a small ball around x (and identifying it with its image under this affine symplectomorphism) the polyhedral *star* of x is the Lagrangian cone

$$(u, v) \mapsto (u^+, v^+, \sigma_1(-u)^+, \sigma_2(-v)^+), \quad s^+ := \max\{s, 0\}. \quad (59)$$

This is the only place where the ‘‘four faces at each vertex’’ hypothesis is used.

2. A one-dimensional rounded corner of zero action (with a genuine construction).

We work first in the symplectic plane $(\mathbb{R}_{(q,p)}^2, dp \wedge dq)$, compatible with the global convention, and we use the Liouville form $\lambda_{(q,p)} := p dq$ (so $d\lambda_{(q,p)} = dp \wedge dq$).

Lemma 8.2 (Embedded rounding with prescribed jets and zero Liouville action). *There exists a smooth embedded curve*

$$\bar{\gamma} : \mathbb{R} \rightarrow \mathbb{R}_{(q,p)}^2, \quad r \mapsto (\bar{q}(r), \bar{p}(r))$$

such that:

1. $\bar{\gamma}(r) = (0, -r)$ for all $r \leq -2$, and $\bar{\gamma}(r) = (r, 0)$ for all $r \geq 2$;
2. $\bar{\gamma}$ agrees identically with these half-axes on open neighborhoods of $(-\infty, -2]$ and $[2, \infty)$ (hence all jets match at $r = \pm 2$);
3. $\bar{\gamma}$ is embedded;
4. its Liouville action is zero:

$$\int_{-\infty}^{\infty} \bar{p}(r) \bar{q}'(r) dr = 0. \quad (60)$$

Proof. Step 1: choose a base rounding with good geometry. Choose once and for all a small number, say $\eta = 1/4$. We first construct a smooth non-decreasing function $\bar{q} : \mathbb{R} \rightarrow [0, \infty)$ such that

$$\bar{q}(r) = 0 \text{ for } r \leq -2 + \eta, \quad \bar{q}(r) = r \text{ for } r \geq 2 - \eta, \quad \bar{q}'(r) > 0 \text{ for } -2 + \eta < r < 2 - \eta.$$

(The function may be chosen flat to infinite order at the two points where it starts and where it joins the line r .) In particular the identities required in (1) hold, and the stronger equalities on the slightly larger intervals give the ‘‘open neighbourhood’’ property (2).

Next choose a smooth function $\bar{p}_0 : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$\bar{p}_0(r) = -r \text{ for } r \leq -2 + \eta, \quad \bar{p}_0(r) = 0 \text{ for } r \geq 2 - \eta,$$

and $\bar{p}_0(r) > 0$ on the intermediate interval $(-2 + \eta, 2 - \eta)$. Then the curve $\bar{\gamma}_0(r) := (\bar{q}(r), \bar{p}_0(r))$ is embedded: on the incoming axial neighbourhood it lies on the positive p -axis and on the outgoing neighbourhood on the positive q -axis, while on the middle interval it is a graph over the strictly increasing q -coordinate and hence has no self-intersections (nor does it meet the axial pieces except at the transition endpoints).

Let

$$A_0 := \int_{-\infty}^{\infty} \bar{p}_0(r) \bar{q}'(r) dr$$

(the integral converges because the product $\bar{p}_0 \bar{q}'$ is supported in the compact interval $[-2 + \eta, 2 - \eta]$.)

Step 2: adjust the action by a compactly supported vertical perturbation. Choose a nonzero smooth bump function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ supported in $(-1, 1)$ (so ϕ vanishes identically near ± 2), and set

$$c := \int_{-\infty}^{\infty} \phi(r) \bar{q}'(r) dr.$$

Since $\bar{q}'(r) > 0$ throughout the support interval $(-1, 1)$ (by our choice of $\eta = 1/4$) and ϕ is not identically zero, we can choose ϕ so that $c \neq 0$ (e.g. take $\phi \geq 0$, not identically zero).

For $s \in \mathbb{R}$ define

$$\bar{p}_s(r) := \bar{p}_0(r) + s \phi(r), \quad \bar{\gamma}_s(r) := (\bar{q}(r), \bar{p}_s(r)).$$

Then $\bar{\gamma}_s$ agrees with $\bar{\gamma}_0$ (hence with the prescribed half-axes) for $|r| \geq 1$, so the jets at $r = \pm 2$ are unchanged. Moreover, the perturbation is supported in $(-1, 1) \subset (-2 + \eta, 2 - \eta)$ where the q -coordinate is strictly increasing, so for every s the curve remains a graph over q on the only region that changes and hence stays embedded.

Finally, the action depends affinely on s :

$$\int \bar{p}_s \bar{q}' dr = \int \bar{p}_0 \bar{q}' dr + s \int \phi \bar{q}' dr = A_0 + s c.$$

Choose $s_* := -A_0/c$. Then $\bar{\gamma} := \bar{\gamma}_{s_*}$ satisfies (60) and has all required properties. \square

Fix such a curve $\bar{\gamma}(r) = (\bar{q}(r), \bar{p}(r))$ once and for all. For $\varepsilon > 0$ define the rescaled curve $\gamma_\varepsilon(r) = (q_\varepsilon(r), p_\varepsilon(r))$ by

$$q_\varepsilon(r) = \varepsilon \bar{q}(r/\varepsilon), \quad p_\varepsilon(r) = \varepsilon \bar{p}(r/\varepsilon), \quad (61)$$

and define the $\varepsilon = 0$ limit to be the union of half-axes

$$q_0(r) = r^+, \quad p_0(r) = (-r)^+.$$

Then for each $\varepsilon > 0$, γ_ε is smooth and embedded, equals the coordinate half-axes for $|r| \geq 2\varepsilon$, and all derivatives match at the transition points $r = \pm 2\varepsilon$ (because $\bar{\gamma}$ was chosen to agree identically with axes near ± 2).

Define

$$F_\varepsilon(r) := \int_{-\infty}^r p_\varepsilon(s) q'_\varepsilon(s) ds. \quad (62)$$

By construction and (60), choosing the integration constant so that F_ε vanishes along the incoming axis, we have

$$F_\varepsilon(r) = 0 \quad \text{whenever } \gamma_\varepsilon(r) \text{ lies on either coordinate axis.} \quad (63)$$

3. Vertex smoothing model (with the vertex sign pattern).

Fix a vertex x and its Darboux coordinates (58) with sign pattern (σ_1, σ_2) . For $\varepsilon > 0$ define the product map

$$\Phi_\varepsilon^{(\sigma_1, \sigma_2)} : \mathbb{R}_{(u,v)}^2 \rightarrow \mathbb{R}^4, \quad \Phi_\varepsilon^{(\sigma_1, \sigma_2)}(u, v) = (q_\varepsilon(u), q_\varepsilon(v), \sigma_1 p_\varepsilon(u), \sigma_2 p_\varepsilon(v)). \quad (64)$$

This is an embedding (product of embeddings). With $\lambda_{\text{std}} = p_1 dq_1 + p_2 dq_2$ in these vertex coordinates,

$$\begin{aligned} (\Phi_\varepsilon^{(\sigma_1, \sigma_2)})^* \lambda_{\text{std}} &= \sigma_1 p_\varepsilon(u) q'_\varepsilon(u) du + \sigma_2 p_\varepsilon(v) q'_\varepsilon(v) dv \\ &= d(\sigma_1 F_\varepsilon(u) + \sigma_2 F_\varepsilon(v)). \end{aligned} \quad (65)$$

Hence $\Phi_\varepsilon^{(\sigma_1, \sigma_2)}$ parametrizes an *exact* Lagrangian surface.

Moreover, for $|u| \geq 2\varepsilon$ the u -factor is already in the axial regime so $F_\varepsilon(u) = 0$ and (65) restricts to $\sigma_2 dF_\varepsilon(v)$, which will match the primitive on the corresponding edge strip. If both $|u| \geq 2\varepsilon$ and $|v| \geq 2\varepsilon$ then the primitive vanishes identically. As $\varepsilon \rightarrow 0$, $\Phi_\varepsilon^{(\sigma_1, \sigma_2)}$ converges uniformly on compact sets to the signed cone (59).

To use this as a local replacement near x we shall not fix the parameter domain once for all at this point. After the edge charts have been chosen (Step 4) we take a small polyhedral ball U_x in the above affine coordinates and we let D_x be the inverse image, in the *cone* parametrisation (59), of $K \cap U_x$. The boundary of this set is then rounded slightly in the parameter plane. What will be important for the gluing is the following elementary property, which we impose by shrinking U_x and by choosing the rounding in collars. There is a collar of ∂D_x which, for the fixed scale ε_0 , lies in the union of the axial regions $\{|u| \geq 2\varepsilon_0\}$ and $\{|v| \geq 2\varepsilon_0\}$; moreover the portions of the boundary that do not belong to the planned overlaps with edge strips lie where *both* inequalities hold. On such collars the parametrization is literally planar (in fact affine linear) and all jets match the adjoining pieces. In particular only a small neighbourhood of the corner, not the whole product surface, is inserted into K . Notice that D_x need not be a

rectangle: on the collars which will overlap an edge its boundary may be a graph in the other variable. This harmless flexibility will be used below in order that the vertex and edge images coincide as subsets of \mathbb{R}^4 .

4. Edge normal form and an edge smoothing model compatible with vertices, including edge signs.

Linear normal form along an edge, with sign. Let $P, P' \subset \mathbb{R}^4$ be Lagrangian planes meeting along a line $\ell = P \cap P'$. Choose $0 \neq e \in \ell$. Choose $a \in P \setminus \ell$ and $b \in P' \setminus \ell$ pointing to the prescribed inward sides of the corresponding faces (these choices are well-defined up to positive scaling). Then $\omega(b, a) \neq 0$, for otherwise $\langle e, a, b \rangle$ would be a 3-dimensional isotropic subspace, impossible in (\mathbb{R}^4, ω) . Define

$$\varepsilon(P, P'; a, b) := \text{sign}(\omega(b, a)) \in \{\pm 1\}.$$

By positive rescaling of a and/or b , we may and do assume

$$\omega(b, a) = \varepsilon(P, P'; a, b) \in \{\pm 1\}. \quad (66)$$

Choose $f \in \mathbb{R}^4$ such that

$$\omega(f, e) = 1, \quad \omega(f, a) = 0, \quad \omega(f, b) = 0. \quad (67)$$

The three vectors e, a, b are linearly independent. Indeed, if (say) b belonged to the span of e and a , then the two incident faces would lie in the same affine Lagrangian plane and the edge would be inessential; alternatively this would contradict the preceding non-isotropy argument. The non-degeneracy of ω identifies \mathbb{R}^4 with its dual. Prescribing the values in (67) on the independent subspace $\langle e, a, b \rangle$ and extending them arbitrarily to a covector on all of \mathbb{R}^4 therefore gives a solution f ; any two such solutions differ by a multiple of the common line $\langle e \rangle = (\langle e, a, b \rangle)^\omega$, and all of them will serve our purpose. Since a vector in the span of e, a, b pairs trivially with e , no solution lies in that span, so (e, a, f, b) is automatically a basis. With this choice the only non-zero pairings among the ordered quadruple are

$$\omega(f, e) = 1, \quad \omega(\varepsilon^{-1}b, a) = 1$$

(up to skew-symmetry). This is precisely the symplectic normal form stated below.

In the associated affine symplectic coordinates

$$(x, y, p_x, p_y) \quad (68)$$

defined by

$$e = \partial_x, \quad a = \partial_y, \quad f = \partial_{p_x}, \quad b = \varepsilon \partial_{p_y} \quad (\varepsilon := \varepsilon(P, P'; a, b)),$$

we have $\omega = dp_x \wedge dx + dp_y \wedge dy$ and $\lambda_{\text{std}} = p_x dx + p_y dy$, and the two incident half-planes take the standard *signed* form

$$\{p_x = p_y = 0, y \geq 0\} \quad \text{and} \quad \{p_x = 0, y = 0, \varepsilon p_y \geq 0\}. \quad (69)$$

(When $\varepsilon = +1$ the second half-plane is $p_y \geq 0$; when $\varepsilon = -1$ it is $p_y \leq 0$.)

Edge data and scale functions. Fix a geometric edge e of K with endpoints x_- and x_+ . Let P_e and P'_e be the two Lagrangian face planes containing e , and let $\ell_e = P_e \cap P'_e$ be the edge line.

Choose a nonzero tangent vector $e_e \in \ell_e$ pointing along the edge (any choice is fine; it only affects the parametrization along the edge, not the subset we construct). Choose inward-pointing transverse vectors $a_e \in P_e \setminus \ell_e$ and $b_e \in P'_e \setminus \ell_e$, and normalize them by (66); denote the resulting sign by

$$\varepsilon_e := \text{sign}(\omega(b_e, a_e)) \in \{\pm 1\}, \quad \omega(b_e, a_e) = \varepsilon_e. \quad (70)$$

Choose f_e satisfying (67) for $(e, a, b) = (e_e, a_e, b_e)$ and define the associated affine symplectic coordinates (68) on a tubular neighborhood U_e of the interior of e (and, after shrinking U_e , also on endpoint collars inside U_e). In these coordinates $K \cap U_e$ is exactly the wedge (69).

Recall that in Step 1 the non-zero vectors chosen on the vertex rays were arbitrary up to positive scale. We now fix those choices once and for all. Thus, at each endpoint x_{\pm} , when the face P_e is the one spanned by the edge under consideration and its adjacent ray, we take as representative of that adjacent ray exactly the vector used in the vertex Darboux coordinates. Since it points into the same inward half-plane as a_e it can be written

$$v_{\text{adj},\pm} = \kappa_{\pm} a_e + \alpha_{\pm} e_e, \quad (71)$$

with $\kappa_{\pm} > 0$. In the vertex normal form at x_{\pm} we take the basis vector e_2 to be precisely this representative. The construction of the symplectically dual basis in Step 1 then fixes a representative of the adjacent ray in the other face P'_e : it is the unique positive multiple for which the pairing with $v_{\text{adj},\pm}$ has absolute value 1. Using the normalization $\omega(b_e, a_e) = \varepsilon_e = \pm 1$ (and the fact that the pairings with the tangent direction e_e vanish), this representative has the expression

$$v'_{\text{adj},\pm} = \kappa_{\pm}^{-1} b_e + \beta_{\pm} e_e, \quad (72)$$

for some $\beta_{\pm} \in \mathbb{R}$. Thus no extra sign is hidden in the coefficients: all κ_{\pm} are *positive* and the only sign information of the wedge is carried by ε_e . (The equality of the sign ε_e with the vertex sign used below is checked explicitly in the compatibility argument.)

Choose a smooth function $\kappa_e(s) > 0$ of the edge coordinate $s = x$ such that $\kappa_e(s) = \kappa_-$ on a collar near x_- and $\kappa_e(s) = \kappa_+$ on a collar near x_+ .

Edge smoothing strip (including the edge sign). Define

$$\Psi_{\varepsilon,e}(s, r) = \left(s, \kappa_e(s) q_{\varepsilon}(r), -\varepsilon_e \frac{\kappa'_e(s)}{\kappa_e(s)} p_{\varepsilon}(r) q_{\varepsilon}(r), \varepsilon_e \kappa_e(s)^{-1} p_{\varepsilon}(r) \right), \quad (73)$$

defined on a finite strip $(s, r) \in I_e \times [-R_e, R_e]$ inside the chosen tube U_e (with I_e an interval parameterizing the portion of the edge we modify, and R_e chosen so that the axial regime already holds on the boundary $|r| = R_e$ for all $\varepsilon \leq \varepsilon_0$).

For all sufficiently small ε (and after possibly shrinking the strip once and for all) this parametrization is an embedding. Indeed the vector $\partial_s \Psi_{\varepsilon,e}$ has x -component 1 whereas $\partial_r \Psi_{\varepsilon,e}$ has x -component 0, so the immersion property reduces to the fact that the rounded curve γ_{ε} is a regular embedding in the (q, p) -plane. If two image points have different s -coordinates they are obviously distinct; for fixed s the pair $(q_{\varepsilon}(r), p_{\varepsilon}(r))$ (and hence also $(\kappa(s)q_{\varepsilon}(r), \varepsilon_e \kappa(s)^{-1} p_{\varepsilon}(r))$) determines r because γ_{ε} is embedded, so no self-intersections occur. Finally the corrective p_x -coordinate in (73) is of size $O(\varepsilon^2)$ on the support of the rounding; by taking ε_0 small we ensure that the whole image lies in the prescribed tubular neighbourhood U_e .

With $\lambda_{\text{std}} = p_x dx + p_y dy$ in these edge coordinates, a direct computation gives

$$\begin{aligned} \Psi_{\varepsilon,e}^* \lambda_{\text{std}} &= p_x ds + p_y dy \\ &= \left(-\varepsilon_e \frac{\kappa'_e}{\kappa_e} p_{\varepsilon} q_{\varepsilon} \right) ds + (\varepsilon_e \kappa_e^{-1} p_{\varepsilon}) (\kappa q'_e dr + \kappa' q_{\varepsilon} ds) \\ &= \varepsilon_e p_{\varepsilon}(r) q'_e(r) dr = d(\varepsilon_e F_{\varepsilon}(r)), \end{aligned} \quad (74)$$

so $\Psi_{\varepsilon,e}$ parametrizes an *exact* Lagrangian strip. The third coordinate in (73) is precisely the correction term ensuring (74) when κ_e varies with s . By (63), the primitive $\varepsilon_e F_{\varepsilon}(r)$ vanishes on collars where the strip has already entered the axial regime. When κ_e is constant (near the endpoints), the third coordinate vanishes and the strip is simply a rescaled copy of the one-dimensional rounding (with the correct sign ε_e in the p_y -direction).

Compatibility with the vertex model on an open collar. We now make explicit choices ensuring that the vertex and edge charts match on overlaps as *subsets* (and with matching primitives), not merely at the level of tangent cones.

Fix an endpoint x of the edge e . Choose the vertex Darboux coordinates at x as in Step 1 so that the edge ray corresponding to e is

$$l_1 = \mathbb{R}_+ \partial_{q_1}$$

(in particular, the two incident face planes along e at x are $P_1 = \langle \partial_{q_1}, \partial_{q_2} \rangle$ and $P_4 = \langle \partial_{q_1}, \sigma_2 \partial_{p_2} \rangle$ for the appropriate vertex sign σ_2). In the edge coordinates (68) chosen above, e is the x -axis and the same two faces are in the standard position (69) with edge sign $\varepsilon_e = \text{sign}(\omega(b_e, a_e))$.

At the same time we recall the harmless freedom, noted in Step 1, to rescale a chosen non-zero vector on any ray by a positive constant (and similarly to rescale the tangent vector used to parametrize the edge). One may use this to arrange that the coordinate along the common edge in the vertex chart has the same absolute scale as the edge coordinate—in which case the scale factor ρ_x below equals 1 at the endpoint. We keep a factor in the formulas to cover the case where no such simultaneous normalization is imposed.

Because both the vertex coordinates and the edge coordinates were chosen to be *affine symplectic* coordinates on \mathbb{R}^4 , the transition map between them on the overlap $U_x \cap U_e$ is the restriction of a *linear* symplectic map (after translating x to the origin). The requirement that it preserve the two incident Lagrangian planes and the common edge line forces the transition to have a concrete finite-dimensional form. Concretely:

Lemma 8.3 (Vertex–edge transition normal form on the wedge). *Let (q_1, q_2, p_1, p_2) and (x, y, p_x, p_y) be linear symplectic coordinates on (\mathbb{R}^4, ω) with*

$$\omega = dp_1 \wedge dq_1 + dp_2 \wedge dq_2 = dp_x \wedge dx + dp_y \wedge dy.$$

Assume that the common edge line is the q_1 -axis in the first coordinates and the x -axis in the second, and that the two incident Lagrangian planes are

$$P_1 = \{p_1 = p_2 = 0\}, \quad P_4 = \{p_1 = q_2 = 0\}$$

in (q, p) -coordinates and

$$\{p_x = p_y = 0\}, \quad \{p_x = y = 0\}$$

in (x, y, p) -coordinates. After replacing the symplectic pair (x, p_x) by $(c^{-1}x, cp_x)$ for some positive constant c (which leaves the two half-planes in the same normal form), there are numbers $\delta \in \{\pm 1\}$, $\kappa > 0$ and $\alpha, \beta, \gamma \in \mathbb{R}$ such that the coordinate change has the form

$$\begin{aligned} x &= \delta q_1 + \alpha q_2 + \beta p_2 + \gamma p_1, \\ y &= \kappa q_2 + \delta \kappa \beta p_1, \\ p_x &= \delta p_1, \\ p_y &= \kappa^{-1} p_2 - \delta \alpha \kappa^{-1} p_1. \end{aligned} \tag{75}$$

Here δ records the relative orientation of the two coordinates along the edge (if the positive q_1 -ray is sent to the positive x -ray then $\delta = +1$). In particular, restricting to the union $P_1 \cup P_4$ (where $p_1 = 0$) we obtain the simpler formulas

$$x = \delta q_1 + \alpha q_2 + \beta p_2, \quad y = \kappa q_2, \quad p_y = \kappa^{-1} p_2. \tag{76}$$

Proof. We spell out the elementary linear algebra since we shall use the signs below. Let T be the linear change of coordinates. The conditions that $T(P_1) = \{p_x = p_y = 0\}$ imply that p_x and p_y contain no q -terms, and the condition that $T(P_4) = \{p_x = y = 0\}$ implies in addition that

p_x is a non-zero multiple of p_1 and that y contains no q_1 - or p_2 -term. Thus, before imposing the symplectic equations, we may write

$$\begin{aligned} x &= Gq_1 + Hq_2 + Ip_1 + Jp_2, & y &= Eq_2 + Fp_1, \\ p_x &= Ap_1, & p_y &= Cp_1 + Dp_2, \end{aligned}$$

with $A, G, E \neq 0$. (The coefficient G has the sign of the orientation with which the edge is parametrised in the two charts.) The equality $T^*(dp_x \wedge dx + dp_y \wedge dy) = dp_1 \wedge dq_1 + dp_2 \wedge dq_2$ is obtained by a direct comparison of coefficients:

$$\begin{aligned} dp_x \wedge dx + dp_y \wedge dy &= AG dp_1 \wedge dq_1 + (AH + CE) dp_1 \wedge dq_2 \\ &\quad + (AJ - DF) dp_1 \wedge dp_2 + DE dp_2 \wedge dq_2. \end{aligned}$$

Hence

$$AG = 1, \quad AH + CE = 0, \quad AJ - DF = 0, \quad DE = 1. \quad (77)$$

We are still free to rescale the *edge* symplectic coordinates by a positive constant, $x \mapsto x/|G|$, $p_x \mapsto |G|p_x$ (leaving y, p_y fixed). This preserves the normal form of the two half-planes and replaces G by $\delta := \text{sign}(G)$. By the first equation in (77) the same operation then gives $A = \delta$. The sign of E records whether the inward half-line in the vertex face is sent to the inward half-line in the edge chart. With our choices it is positive, and we put $\kappa := E$. Set $\alpha := H$, $\beta := J$, $\gamma := I$, and use the remaining equations in (77) to solve for C and F . We obtain precisely (75). Formula (76) is the restriction to $p_1 = 0$. \square

Return to the geometric situation at the endpoint x of the edge e . Let $s_x \in \mathbb{R}$, $\delta_x \in \{\pm 1\}$, a positive scale $\rho_x > 0$, and $\kappa_x > 0$, $\alpha_x, \beta_x \in \mathbb{R}$ be the parameters from Lemma 8.3: here s_x is the value of the edge coordinate $s = x$ at the vertex (the affine origin in the edge chart need not be the vertex), δ_x records whether the oriented edge ray in the vertex chart points in the positive or negative s -direction, and ρ_x accounts for a possible residual difference of scale if one has not normalized the representative of the ray. Then, restricting the vertex smoothing model $\Phi_\varepsilon^{(\sigma_1, \sigma_2)}$ to the collar where the edge-direction parameter u is already in the axial regime ($u \geq 2\varepsilon$ so $q_\varepsilon(u) = u$, $p_\varepsilon(u) = 0$), and using (76) (undoing the optional rescaling of the edge coordinate recorded by ρ_x), we find that points of the vertex model near this edge have edge-chart coordinates

$$(s_x + \delta_x \rho_x u + \alpha_x q_\varepsilon(v) + \beta_x (\sigma_2 p_\varepsilon(v)), \kappa_x q_\varepsilon(v), 0, \kappa_x^{-1} (\sigma_2 p_\varepsilon(v))). \quad (78)$$

On the other hand, on the edge collar where $\kappa_e \equiv \kappa_x$ (so $\kappa'_e = 0$) the edge strip (73) has image

$$(s, \kappa_x q_\varepsilon(r), 0, \varepsilon_e \kappa_x^{-1} p_\varepsilon(r)).$$

It remains to identify the edge sign with the vertex sign. Write the chosen inward transverse vectors in the vertex coordinates at x in the form

$$a_e = A \partial_{q_2} + \alpha' \partial_{q_1}, \quad b_e = B (\sigma_2 \partial_{p_2}) + \beta' \partial_{q_1},$$

with $A, B > 0$ (adding tangent components along the edge does not change the inward half-planes). The pairings with ∂_{q_1} vanish because both incident planes are Lagrangian and contain that vector, hence

$$\omega(b_e, a_e) = AB \omega(\sigma_2 \partial_{p_2}, \partial_{q_2}) = AB \sigma_2.$$

After the positive rescaling used in the normalization (66) the sign is unchanged, so $\varepsilon_e = \sigma_2$ and the last coordinate agrees. Thus the *images coincide* on the overlap, with the smooth change of parameters

$$(s, r) = (s_x + \delta_x \rho_x u + \alpha_x q_\varepsilon(v) + \beta_x \sigma_2 p_\varepsilon(v), v).$$

In particular, the vertex patch and the edge strip agree on an *open collar* in $U_x \cap U_e$ as subsets of \mathbb{R}^4 . More precisely the primitives computed with the two local Liouville forms have the same differential on this collar, and on the smaller axial subcollar both primitives themselves vanish by (63); this is the fact that will be used in the global exactness argument after we have corrected the local primitives by exact terms coming from the fixed global form λ . The same argument applies at the other endpoint x_+ .

It may be useful to spell out that we have not lost any generality by treating only the case in which, at the vertex under consideration, the chosen edge is the ray l_1 . The Darboux coordinates of Step 1 are fixed once for all at a vertex, and the other three incident edges are the rays l_2, l_3, l_4 . On the collar of the edge l_2 the vertex primitive (65) reduces to $\sigma_1 dF_\varepsilon(u)$ (because the parameter v is axial), and a direct computation exactly like the one above gives

$$\text{sign } \omega(\sigma_1 \partial_{p_1}, \partial_{q_1}) = \sigma_1.$$

Thus the sign ε_e of that edge is σ_1 . For the edge l_3 the parameter u is axial on the negative side and the primitive reduces to $\sigma_2 dF_\varepsilon(v)$; taking the inward vectors in the faces P_2 and P_3 to be respectively ∂_{q_2} and $\sigma_2 \partial_{p_2}$ gives $\text{sign } \omega(\sigma_2 \partial_{p_2}, \partial_{q_2}) = \sigma_2$. Finally the edge l_4 has sign σ_1 and primitive $\sigma_1 dF_\varepsilon(u)$. Equivalently one can cyclically relabel the four rays and repeat the proof. Hence every incident edge of the fixed vertex coordinates is compatible with the corresponding edge strip, the coefficient of F_ε being precisely the edge sign.

5. Choosing neighborhoods and gluing the global smoothing (with finite strips/disks).

Since K has finitely many vertices and edges, choose:

- for each vertex x , a small closed ball U_x centered at x , such that the U_x are pairwise disjoint and such that in an affine symplectic coordinate chart on U_x the germ $K \cap U_x$ is exactly the planar star (59) restricted to a small disk in the (u, v) -domain;
- for each edge e , a symplectic tubular neighborhood U_e of the *interior* of e (excluding tiny endpoint subsegments), such that different U_e are disjoint, and such that U_e meets only the two endpoint balls U_{x_-} and U_{x_+} (and only in collar overlaps).

Shrink U_x and U_e further so that on each overlap $U_x \cap U_e$ we are in the collar where:

1. the corresponding κ_e has already become constant (equal to the endpoint value κ_x);
2. the vertex parameter in the edge direction is already in the axial regime for the *fixed* scale ε_0 ;
3. the edge transverse parameter r satisfies $|r| \geq 2\varepsilon_0$ on a nonempty open subset of the overlap (so the smoothing is stationary there for *all* $t \in (0, 1]$ with $\varepsilon(t) \leq \varepsilon_0$).

Such choices are possible because all charts are local and we may shrink neighborhoods independently.

Fix $\varepsilon_0 > 0$ so small that, in each chart, the vertex and edge modifications supported in $|u|, |v|, |r| < 2\varepsilon_0$ lie entirely inside the chosen neighborhoods. Define

$$\varepsilon(t) := \varepsilon_0 t, \quad t \in (0, 1].$$

Definition of the smoothed surfaces. For each $t \in (0, 1]$, define K_t by replacing:

- in each U_x , the polyhedral star $K \cap U_x$ by the image of the chosen small domain D_x under the vertex smoothing map $\Phi_{\varepsilon(t)}^{(\sigma_1(x), \sigma_2(x))}$ composed with the inverse of the chosen affine symplectomorphism $U_x \rightarrow \mathbb{R}^4$;

- in each U_e , the wedge $K \cap U_e$ by the image of a finite strip $I_e \times [-R_e, R_e]$ under the edge smoothing map $\Psi_{\varepsilon(t),e}$ composed with the inverse of the chosen affine symplectomorphism $U_e \rightarrow \mathbb{R}^4$;
- outside $\bigcup_x U_x \cup \bigcup_e U_e$, leave K unchanged.

By the compatibility established in Step 4 (explicit overlap equality on open collars), on each nonempty overlap $U_x \cap U_e$ the vertex patch and the edge strip define *the same subset* of \mathbb{R}^4 and match smoothly (their parameterizations differ only by a smooth diffeomorphism of the collar in the parameter domain). Moreover, on the outer boundaries of U_x and U_e the smoothed pieces are *literally* the original planar pieces because $q_{\varepsilon(t)}$ and $p_{\varepsilon(t)}$ are exactly axial there (and likewise $F_{\varepsilon(t)}$ vanishes there).

Therefore, for each $t > 0$ the set K_t is a smooth embedded Lagrangian surface.

Topological isotopy to K at $t = 0$. The last point which needs a little care is the choice of a single parameter surface for all values of the smoothing parameter. In Lemma 8.3 the change of variables on a vertex–edge collar contains the functions q_ε and p_ε , so if one glued the vertex and edge charts by that formula the gluing would depend on t . We freeze the gluing at the reference scale $t = 1$ and compensate by a harmless reparametrisation of the edge strips.

Let Σ be the smooth surface obtained by gluing the vertex domains D_x to the edge strips $I_e \times [-R_e, R_e]$ by the transition maps, on the endpoint collars,

$$(s, r) = T_{x,1}(u, r) := s_x + \delta_x \rho_x u + \alpha_x q_{\varepsilon_0}(r) + \beta_x \varepsilon_e p_{\varepsilon_0}(r),$$

where $\varepsilon_0 = \varepsilon(1)$, and by rounding these maps in the stationary axial part; outside the chosen neighbourhoods Σ is just the given polyhedral surface. (For the other three edges at a vertex the notation is the analogous one from the compatibility discussion of Step 4.)

For $t \in (0, 1]$ we now define an embedding $f_t : \Sigma \rightarrow \mathbb{R}^4$. On a vertex chart we use the map $\Phi_{\varepsilon(t)}$ as before. On an edge chart we compose $\Psi_{\varepsilon(t),e}$ with a small diffeomorphism of the parameter strip which only changes the edge coordinate. Choose cut–off functions χ_{x-}, χ_{x+} of the coordinate s , supported in the endpoint collars of I_e and equal to 1 on the actual overlaps with the vertex balls (the collars have been shrunk so that the tapering occurs outside the overlaps). Put

$$\begin{aligned} \Theta_{t,e}(s, r) := & (s + \chi_{x-}(s)(\alpha_{x-}[q_{\varepsilon(t)} - q_{\varepsilon_0}](r) + \beta_{x-}\varepsilon_e[p_{\varepsilon(t)} - p_{\varepsilon_0}](r)) \\ & + \chi_{x+}(s)(\alpha_{x+}[q_{\varepsilon(t)} - q_{\varepsilon_0}](r) + \beta_{x+}\varepsilon_e[p_{\varepsilon(t)} - p_{\varepsilon_0}](r)), r). \end{aligned}$$

For ε_0 sufficiently small this is a diffeomorphism of the strip onto itself (the shifts are $O(\varepsilon_0)$ and vanish in the axial region $|r| \geq 2\varepsilon_0$). We set f_t on the edge chart equal to the affine inverse of $\Psi_{\varepsilon(t),e} \circ \Theta_{t,e}$. On a collar where $\chi_x = 1$ the reference transition $T_{x,1}$ is thereby changed exactly into the transition $T_{x,t}$ of Step 4, so the vertex and edge formulas agree on the fixed gluing of Σ ; away from the collars $\Theta_{t,e}$ is the identity. Since $\Psi_{\varepsilon,e}^* \lambda_{\text{std}} = d(\varepsilon_e F_\varepsilon(r))$ contains no ds -term, this reparametrisation will not affect the primitive used later.

Together with the identity outside the chosen neighbourhoods these local definitions give a smooth family of embeddings $f_t : \Sigma \rightarrow \mathbb{R}^4$ whose images are precisely the sets K_t . The derivative $\partial_t f_t$ is supported in the union of the vertex balls and edge strips.

As $t \rightarrow 0^+$, the maps f_t converge uniformly on compact subsets to the obvious PL parametrization of K (because each local model converges uniformly to the corresponding cone/wedge and the support shrinks like $O(\varepsilon(t))$). Thus f_t extends continuously at $t = 0$ to an embedding f_0 with image K . In the sense of a family of embeddings this already gives an isotopy of the surface. If one insists on an *ambient* topological isotopy of R^4 , the standard isotopy–extension theorem for locally flat embedded compact submanifolds (applied to the track $(t, x) \mapsto f_t(x)$) promotes this family to such an ambient isotopy. In either interpretation $\{K_t\}_{t \in [0,1]}$ is a topological isotopy with endpoint $K_0 = K$.

6. From exactness of the deformation to a single Hamiltonian isotopy on all of $(0, 1]$.

We use the standard exactness/flux criterion in an exact symplectic manifold, written with the Hamiltonian convention (53) and with explicit attention to smooth dependence on t .

Lemma 8.4 (Hamiltonian criterion in the exact case). *Let $(M, \omega = d\lambda)$ be an exact symplectic manifold and let $f_t : \Sigma \rightarrow M$ be a smooth family of embeddings for $t \in (0, 1]$ with images $L_t = f_t(\Sigma)$ smooth Lagrangian submanifolds.*

Assume:

1. *there is a fixed compact set $C \subset \Sigma$ such that $\partial_t f_t$ is supported in C for all $t \in (0, 1]$ (equivalently, the family is stationary on $\Sigma \setminus C$), and the trace $\bigcup_{t \in (0, 1]} f_t(C)$ is contained in a fixed compact subset of M ;*
2. *there is a fixed closed 1-form β on Σ and a smooth function $S : (0, 1] \times \Sigma \rightarrow \mathbb{R}$ such that for each t*

$$f_t^* \lambda = \beta + dS_t, \quad S_t := S(t, \cdot), \quad (79)$$

and S_t is independent of t on $\Sigma \setminus C$ (in particular, $\partial_t S_t$ is supported in C).

Then there exists a smooth function $H : (0, 1] \times M \rightarrow \mathbb{R}$ with compact support contained in a fixed compact subset of M such that the (time-dependent) Hamiltonian flow ϕ_H^{t, t_0} (defined by (53)) satisfies

$$\phi_H^{t, t_0}(L_{t_0}) = L_t$$

for all $t, t_0 \in (0, 1]$ (in particular, $\{L_t\}_{t \in (0, 1]}$ is a Hamiltonian isotopy, generated by a single time-dependent Hamiltonian on all of $(0, 1]$).

Proof. Let V_t be the deformation vector field along L_t , defined by $V_t \circ f_t = \partial_t f_t$. Differentiate (79) in t . Since β is fixed and closed, we obtain

$$\frac{d}{dt} f_t^* \lambda = d(\partial_t S_t).$$

On the other hand, by Cartan's formula,

$$\frac{d}{dt} f_t^* \lambda = f_t^*(\mathcal{L}_{V_t} \lambda) = f_t^*(d(\iota_{V_t} \lambda) + \iota_{V_t} d\lambda) = d(\lambda(V_t) \circ f_t) + f_t^*(\iota_{V_t} \omega).$$

Comparing gives

$$f_t^*(\iota_{V_t} \omega) = d(\partial_t S_t - \lambda(V_t) \circ f_t). \quad (80)$$

Thus the 1-form $\iota_{V_t} \omega$ restricted to L_t is exact, with primitive

$$h_t := (\partial_t S_t - \lambda(V_t) \circ f_t) \circ f_t^{-1} \in C^\infty(L_t).$$

By hypothesis, $\partial_t S_t$ and V_t are supported in C , hence h_t has compact support in L_t (and depends smoothly on t because S and f do).

For the next step we need extensions of these functions off the moving submanifolds. We recall explicitly the elementary fact that this can be done with smooth dependence on the parameter. Equip M with a fixed Riemannian metric. For each compact subinterval $[\tau, 1] \subset (0, 1]$ the map $(t, x) \mapsto f_t(x)$ is a proper embedding of $[\tau, 1] \times C$, hence the normal exponential map gives a tubular neighbourhood with projection depending smoothly on t . On the countable cover $\{[1/(k+1), 1/k]\}_k$ these neighbourhoods can be chosen compatible on overlaps and patched by a partition of unity in the parameter (or, equivalently, by using the exponential map in the smoothly varying normal bundle of the embedded track). Thus we obtain open sets U_t and projections $\pi_t : U_t \rightarrow L_t$ depending smoothly for all $t \in (0, 1]$; their radii may shrink as $t \rightarrow 0$, which is harmless. The union of all supports will still lie in a fixed compact subset of M by

hypothesis (1). Extend h_t to U_t by $\tilde{h}_t := h_t \circ \pi_t$ and multiply by a bump function χ_t supported in U_t and equal to 1 near L_t . The bump functions too can be chosen smooth in t by the same local-in-time construction, with supports contained in that fixed compact subset.

Define the Hamiltonian

$$H_t := -\chi_t \tilde{h}_t, \quad H(t, \cdot) := H_t.$$

Then H is smooth on $(0, 1] \times M$ and compactly supported in M .

Let X_{H_t} be the Hamiltonian vector field defined by (53). Along L_t , since $\chi_t \equiv 1$ and $\tilde{h}_t|_{L_t} = h_t$, we have $H_t|_{L_t} = -h_t$, hence

$$\iota_{X_{H_t}} \omega|_{L_t} = -dH_t|_{L_t} = dh_t = \iota_{V_t} \omega|_{L_t}$$

by (80). Thus $V_t - X_{H_t}$ is tangent to L_t . Tangential components only reparametrize L_t , so the time-dependent Hamiltonian flow (with initial time t_0) carries L_{t_0} to L_t for all $t, t_0 \in (0, 1]$. \square

Application to K_t . Let Σ and $f_t : \Sigma \rightarrow \mathbb{R}^4$ be as in Step 5. By construction $\partial_t f_t$ is supported in a fixed compact subset $C \subset \Sigma$ and the family is stationary on $\Sigma \setminus C$; moreover all points of $f_t(C)$ lie in the fixed finite union of vertex and edge neighborhoods chosen in Step 5, hence in a compact subset of \mathbb{R}^4 .

Fix a reference time $t_0 = 1$ and set

$$\beta := f_{t_0}^* \lambda.$$

Since L_{t_0} is Lagrangian, β is closed.

We claim that there exists a smooth function $S : (0, 1] \times \Sigma \rightarrow \mathbb{R}$, supported in C and identically 0 on a collar of ∂C , such that $f_t^* \lambda = \beta + dS_t$ for all $t \in (0, 1]$.

Indeed, cover C by finitely many vertex charts and edge charts from Step 5, and include also the stationary open set $\Sigma \setminus C$. In each vertex chart U_x , let λ_{std} denote the standard Liouville form in the chosen vertex Darboux coordinates, and choose a smooth function G_x on U_x with

$$\lambda_{\text{std}} = \lambda + dG_x$$

(which exists because $H^1(\mathbb{R}^4) = 0$). Then (65) gives

$$f_t^* \lambda = d(\sigma_1(x)F_{\varepsilon(t)}(u) + \sigma_2(x)F_{\varepsilon(t)}(v)) - d(G_x \circ f_t) = dP_{t,x} \quad \text{on } U_x,$$

where

$$P_{t,x} := \sigma_1(x)F_{\varepsilon(t)}(u) + \sigma_2(x)F_{\varepsilon(t)}(v) - G_x \circ f_t.$$

Similarly, in each edge chart U_e choose G_e with $\lambda_{\text{std}} = \lambda + dG_e$ and use (74) to write

$$f_t^* \lambda = d(\varepsilon_e F_{\varepsilon(t)}(r)) - d(G_e \circ f_t) = dP_{t,e} \quad \text{on } U_e,$$

where

$$P_{t,e} := \varepsilon_e F_{\varepsilon(t)}(r) - G_e \circ f_t.$$

On the stationary region $\Sigma \setminus C$ we set $P_{t,\text{out}} := 0$.

Define local difference functions

$$S_{t,x} := P_{t,x} - P_{t_0,x}, \quad S_{t,e} := P_{t,e} - P_{t_0,e}, \quad S_{t,\text{out}} := 0.$$

Then $dS_{t,*} = f_t^* \lambda - f_{t_0}^* \lambda = f_t^* \lambda - \beta$ on each chart.

On any connected overlap of two charts, the difference of two such local functions has zero differential, hence is constant. By the way the overlaps were chosen in Step 5, each overlap

contains a nonempty open subset where the smoothing is already in the axial regime for the fixed scale ε_0 ; hence for every $t \in (0, 1]$ we have

$$F_{\varepsilon(t)} = F_{\varepsilon(t_0)} = 0$$

on that open subset (by (63)), and also $f_t = f_{t_0}$ there (the family is stationary on the axial collar), so $G \circ f_t = G \circ f_{t_0}$. Therefore each local difference function $S_{t,*}$ vanishes on that open subset of the overlap, forcing the overlap-constant to be 0. Thus the $S_{t,*}$ agree on overlaps and glue to a globally defined smooth function $S_t : \Sigma \rightarrow \mathbb{R}$, supported in C and identically 0 near ∂C .

Moreover, since each local formula depends smoothly on t and the local functions agree *identically* on overlaps (not just up to a t -dependent constant), the glued map $(t, x) \mapsto S(t, x)$ is smooth on $(0, 1] \times \Sigma$.

We have shown $f_t^* \lambda = \beta + dS_t$ with β fixed and S_t stationary on $\Sigma \setminus C$. Applying Lemma 8.4, the isotopy $\{K_t\}_{t \in (0, 1]}$ is Hamiltonian, generated by a single smooth compactly supported Hamiltonian $H(t, \cdot)$ on all of $(0, 1]$.

Conclusion. We have constructed a family $\{K_t : t \in (0, 1]\}$ of smooth embedded Lagrangian surfaces such that:

- $K_t \rightarrow K$ in C^0 as $t \rightarrow 0^+$ and $\{K_t\}_{t \in [0, 1]}$ extends to a topological isotopy with $K_0 = K$;
- for $t > 0$ the family is a Hamiltonian isotopy (indeed generated by one smooth time-dependent Hamiltonian on $(0, 1]$).

It is worth pointing out that, in the compact case, the statement is consistent with the familiar topological restrictions on Lagrangian submanifolds of the standard symplectic \mathbb{R}^4 . Suppose that K is compact and orientable. For any compatible almost-complex structure J the normal bundle of the smoothed surface K_t is $J(TK_t)$, hence is isomorphic (as an oriented real plane bundle) to the tangent bundle. The Euler number of the normal bundle of an oriented compact embedded surface in \mathbb{R}^4 is its self-intersection number, which is zero because $H_2(\mathbb{R}^4) = 0$. Consequently $\chi(K) = 0$; in particular a quadrivalent polyhedral Lagrangian sphere cannot occur.

Therefore every quadrivalent polyhedral Lagrangian surface $K \subset \mathbb{R}^4$ admits a Lagrangian smoothing.

References

- [1] A. Weinstein, *Symplectic manifolds and their Lagrangian submanifolds*, Adv. Math. **6** (1971), 329–346. DOI: 10.1016/0001-8708(71)90020-X.
- [2] D. McDuff and D. Salamon, *Introduction to Symplectic Topology*, 3rd ed., Oxford University Press, Oxford, 2017.
- [3] M. W. Hirsch, *Differential Topology*, Graduate Texts in Mathematics, Vol. 33, Springer, New York, 1976. DOI: 10.1007/978-1-4684-9449-5.
- [4] R. D. Edwards and R. C. Kirby, *Deformations of spaces of imbeddings*, Ann. of Math. (2) **93** (1971), no. 1, 63–88. DOI: 10.2307/1970753.

9 Algebraic relations among scaled quadri-linear determinant tensors

Problem

Let $n \geq 5$. Let $A^{(1)}, \dots, A^{(n)} \in \mathbb{R}^{3 \times 4}$ be Zariski-generic. For $\alpha, \beta, \gamma, \delta \in [n]$, construct $Q^{(\alpha\beta\gamma\delta)} \in \mathbb{R}^{3 \times 3 \times 3 \times 3}$ so that its (i, j, k, ℓ) entry for $1 \leq i, j, k, \ell \leq 3$ is given by

$$Q_{ijkl}^{(\alpha\beta\gamma\delta)} = \det[A^{(\alpha)}(i, :); A^{(\beta)}(j, :); A^{(\gamma)}(k, :); A^{(\delta)}(\ell, :)].$$

Here $A(i, :)$ denotes the i th row of a matrix A , and semicolon denotes vertical concatenation. We are interested in algebraic relations on the set of tensors $\{Q^{(\alpha\beta\gamma\delta)} : \alpha, \beta, \gamma, \delta \in [n]\}$.

More precisely, does there exist a polynomial map $\mathbf{F} : \mathbb{R}^{81n^4} \rightarrow \mathbb{R}^N$ that satisfies the following three properties?

- The map \mathbf{F} does not depend on $A^{(1)}, \dots, A^{(n)}$.
- The degrees of the coordinate functions of \mathbf{F} do not depend on n .
- Let $\lambda \in \mathbb{R}^{n \times n \times n \times n}$ satisfy $\lambda_{\alpha\beta\gamma\delta} \neq 0$ for precisely $\alpha, \beta, \gamma, \delta \in [n]$ that are not identical. Then $\mathbf{F}(\lambda_{\alpha\beta\gamma\delta} Q^{(\alpha\beta\gamma\delta)} : \alpha, \beta, \gamma, \delta \in [n]) = 0$ holds if and only if there exist $u, v, w, x \in (\mathbb{R}^*)^n$ such that $\lambda_{\alpha\beta\gamma\delta} = u_\alpha v_\beta w_\gamma x_\delta$ for all $\alpha, \beta, \gamma, \delta \in [n]$ that are not identical.

Solution

We give an explicit construction of such a map \mathbf{F} (in fact, with uniform degree 5). In the proof it is a little cleaner to work first over the algebraic closure \mathbb{C} . At the end of the argument I explain why, for real data, the factors which are obtained over \mathbb{C} may in fact be chosen real. All polynomials which occur have real (indeed integral) coefficients.

Step 0: packaging the tensors. Let

$$R := [n] \times \{1, 2, 3\}, \quad s = (\alpha, i) \in R,$$

and write $c(s) = \alpha$ (camera index) and $r(s) = i$ (row index). Given an array $x = (x_{ijkl}^{(\alpha\beta\gamma\delta)}) \in \mathbb{C}^{81n^4}$, define a single tensor $\mathcal{X} \in (\mathbb{C}^{3n})^{\otimes 4}$ by

$$\mathcal{X}_{stuv} := x_{r(s)r(t)r(u)r(v)}^{(c(s)c(t)c(u)c(v))}, \quad s, t, u, v \in R. \quad (81)$$

This identifies \mathbb{C}^{81n^4} with $\mathbb{C}^{(3n)^4}$.

For $p \in \{1, 2, 3, 4\}$, let $\text{Flat}_p(\mathcal{X})$ be the mode- p flattening: it is the matrix obtained by grouping the p th index as a row index and the other three indices as a column index. Thus

$$\text{Flat}_p(\mathcal{X}) \in \mathbb{C}^{3n \times (3n)^3}.$$

Step 1: the polynomial map. Define \mathbf{F} by

$$\mathbf{F}(x) := \left(\text{all } 5 \times 5 \text{ minors of } \text{Flat}_p(\mathcal{X}) \text{ for } p = 1, 2, 3, 4 \right), \quad (82)$$

where \mathcal{X} is obtained from x via (81). Each coordinate of \mathbf{F} is a determinant of a 5×5 submatrix of a flattening, hence is a polynomial of degree 5 in the entries of x . The definition uses only the input tensor x (and n through the index ranges), and does not involve the matrices $A^{(\alpha)}$.

Step 2: genericity hypotheses on the cameras. Stack all camera rows into a single matrix

$$M \in \mathbb{C}^{3n \times 4}, \quad \text{whose rows are } a_s \in \mathbb{C}^{1 \times 4} \text{ (} s \in R \text{),}$$

so that the 3×4 block of rows indexed by $\{(\alpha, 1), (\alpha, 2), (\alpha, 3)\}$ equals $A^{(\alpha)}$. After removing a proper algebraic subset of $(\mathbb{C}^{3 \times 4})^n$, we may assume:

(G1) each $A^{(\alpha)}$ has rank 3 (so its row space $U_\alpha \subset \mathbb{C}^{1 \times 4}$ is a hyperplane);

(G2) M has rank 4 (equivalently, $\text{im}(M) \subset \mathbb{C}^{3n}$ is 4-dimensional);

(G3) for every ordered triple $(\beta, \gamma, \delta) \in [n]^3$ that is not constant,

$$\text{span}\{u \wedge v \wedge w : u \in U_\beta, v \in U_\gamma, w \in U_\delta\} = \bigwedge^3 \mathbb{C}^4.$$

Let us spell out why this open set is non-empty. All three requirements are conditions given by the non-vanishing of polynomial functions of the entries of the matrices. Thus it is enough to know that none of these polynomials is identically equal to zero. For (G1) this is clear (one 3×3 minor of $A^{(\alpha)}$ has to be non-zero), and for (G2) it suffices to note, for instance, that one may take a first camera whose three rows are the standard row vectors e_1, e_2, e_3 and a second camera with rows e_1, e_2, e_4 (along with arbitrary further blocks); the stacked matrix then has rank 4.

For the reader's convenience I also record an explicit verification for (G3). Put $V = \mathbb{C}^4$ and fix an ordered triple of indices which is not constant. The span which occurs in (G3) is the image of the multilinear map $U_\beta \otimes U_\gamma \otimes U_\delta \rightarrow \bigwedge^3 V$, $(u, v, w) \mapsto u \wedge v \wedge w$; after bases have been chosen in the three hyperplanes, the condition that this image have dimension four is the non-vanishing of some 4×4 minor and hence is polynomial. This polynomial is not identically zero. Indeed, if two of the hyperplanes are equal we may (after a change of coordinates) take $U_\beta = U_\gamma = \langle e_1, e_2, e_3 \rangle$ and $U_\delta = \langle e_1, e_2, e_4 \rangle$; the wedges

$$e_2 \wedge e_3 \wedge e_1, \quad e_2 \wedge e_4 \wedge e_1, \quad e_3 \wedge e_4 \wedge e_1, \quad e_2 \wedge e_3 \wedge e_4$$

(which are obtained from suitable choices of u, v, w) already form a basis of $\bigwedge^3 V$. If the three hyperplanes are pairwise distinct we may take $U_\beta = \langle e_2, e_3, e_4 \rangle$, $U_\gamma = \langle e_1, e_3, e_4 \rangle$ and $U_\delta = \langle e_1, e_2, e_4 \rangle$, from which the four basis wedges of $\bigwedge^3 V$ are obtained just as easily (for example $e_2 \wedge e_3 \wedge e_1, e_2 \wedge e_4 \wedge e_1, e_3 \wedge e_4 \wedge e_1, e_2 \wedge e_3 \wedge e_4$). Permuting the roles of the indices gives the remaining cases. Consequently, for every fixed triple the failure of (G3) is a proper algebraic subset of the parameter space.

The ambient space $(\mathbb{C}^{3 \times 4})^n$ is irreducible, and a finite intersection of non-empty Zariski open subsets of an irreducible variety is non-empty. Thus there are cameras satisfying (G1)–(G3), and in fact the set of such cameras is a Zariski open dense subset. All the defining polynomials have real coefficients, so this open subset contains real points (equivalently its real points are Euclidean dense). Hence a Zariski-generic real choice of $A^{(1)}, \dots, A^{(n)}$ satisfies (G1)–(G3).

Step 3: linear-algebra preliminaries. Let $W := \text{im}(M) \subset \mathbb{C}^{3n}$, so $\dim W = 4$ by (G2).

Lemma 9.1 (Cofactor vector). *For any $t, u, v \in R$ there exists a (unique) vector $\omega_{tuv} \in \mathbb{C}^4$ such that*

$$\det(x, a_t, a_u, a_v) = x \omega_{tuv}^T \quad \forall x \in \mathbb{C}^{1 \times 4}.$$

Moreover, the column of $\text{Flat}_1(\mathcal{Y})$ indexed by (t, u, v) , where $\mathcal{Y}_{stuv} = \det(a_s, a_t, a_u, a_v)$, equals $M \omega_{tuv}$.

Proof. The map $x \mapsto \det(x, a_t, a_u, a_v)$ is linear in $x \in \mathbb{C}^{1 \times 4}$, hence is given by $x \mapsto x \omega^T$ for a unique $\omega \in \mathbb{C}^4$. The second claim is immediate since $(\det(a_s, a_t, a_u, a_v))_{s \in R} = (a_s \omega_{tuv}^T)_{s \in R} = M \omega_{tuv}$. \square

Lemma 9.2 (Diagonal stabilizer is scalar). *Assume (G1)–(G2) and let $D = \text{diag}(d_\alpha I_3)_{\alpha=1}^n \in \mathbb{C}^{3n \times 3n}$ with each $d_\alpha \in \mathbb{C}$. If $DW \subseteq W$, then all d_α are equal.*

Proof. Since $\text{rank}(M) = 4$, the columns of M form a basis of W . The inclusion $DW \subseteq W$ therefore defines a (unique) linear endomorphism $H \in M_4(\mathbb{C})$ by

$$DM = MH.$$

(The matrix H need not be invertible if some of the d_α vanish.) Restricting to the block of rows belonging to camera α gives

$$d_\alpha A^{(\alpha)} = A^{(\alpha)} H.$$

Thus every row of $A^{(\alpha)}$, and hence every vector in the row space U_α , is a left eigenvector of H with eigenvalue d_α . Any two hyperplanes in \mathbb{C}^4 have non-zero intersection (in fact of dimension at least two), so if $y \in U_\alpha \cap U_{\alpha'}$ is non-zero we have $yH = d_\alpha y = d_{\alpha'} y$, which forces $d_\alpha = d_{\alpha'}$. Hence all d_α are equal. \square

Step 4: the tensor slice and the “easy” direction. Given scalars $\lambda \in \mathbb{C}^{n \times n \times n \times n}$, consider the tensor $\mathcal{T} \in (\mathbb{C}^{3n})^{\otimes 4}$ defined by

$$\mathcal{T}_{stuv} := \lambda_{c(s)c(t)c(u)c(v)} \det(a_s, a_t, a_u, a_v). \quad (83)$$

This \mathcal{T} is exactly the packaging (81) of the family $(\lambda_{\alpha\beta\gamma\delta} Q^{(\alpha\beta\gamma\delta)})_{\alpha,\beta,\gamma,\delta}$.

Assume first that $\lambda_{\alpha\beta\gamma\delta} = u_\alpha v_\beta w_\gamma x_\delta$ for all non-identical quadruples and some $u, v, w, x \in (\mathbb{C}^*)^n$. Fix a triple (β, γ, δ) that is not constant. For the mode-1 flattening, the column indexed by (t, u, v) with $c(t) = \beta, c(u) = \gamma, c(v) = \delta$ equals

$$(\mathcal{T}_{stuv})_{s \in R} = \text{diag}(u_\alpha I_3)_{\alpha=1}^n \cdot (v_\beta w_\gamma x_\delta \cdot (M\omega_{tuv})) \in \text{diag}(u_\alpha I_3)W,$$

by Lemma 9.1. If the suffix happens to be constant, say (β, β, β) , the very same formula is still valid. Indeed ω_{tuv} is then the covector associated with $a_t \wedge a_u \wedge a_v \in \bigwedge^3 U_\beta$, so the entries of $M\omega_{tuv}$ in the block of rows belonging to camera β are all zero (the determinant vanishes on the whole hyperplane U_β); multiplying by $v_\beta w_\beta x_\beta \text{diag}(u_\alpha I_3)$ therefore reproduces the column of \mathcal{T} as well – for the row block $\alpha = \beta$ both sides are simply zero, independently of the value of $\lambda_{\beta\beta\beta\beta}$. Hence *all* columns of $\text{Flat}_1(\mathcal{T})$, for arbitrary suffixes, lie in the fixed 4-plane $\text{diag}(u_\alpha I_3)W$, so $\text{rank}(\text{Flat}_1(\mathcal{T})) \leq 4$. The same argument applies to the other three flattenings. Therefore all 5×5 minors in (82) vanish, i.e. $\mathbf{F}(\mathcal{T}) = 0$.

Step 5: vanishing of minors forces one-mode factorization. Assume now that $\mathbf{F}(\mathcal{T}) = 0$, i.e. every 5×5 minor of every $\text{Flat}_p(\mathcal{T})$ vanishes. Equivalently,

$$\text{rank}(\text{Flat}_p(\mathcal{T})) \leq 4 \quad \text{for } p = 1, 2, 3, 4. \quad (84)$$

We use the standing assumption on λ :

$$\lambda_{\alpha\beta\gamma\delta} \neq 0 \quad \text{iff } (\alpha, \beta, \gamma, \delta) \text{ is not all identical.} \quad (85)$$

In particular, if a triple (β, γ, δ) is not constant, then $\lambda_{\alpha\beta\gamma\delta} \neq 0$ for *all* α , hence the diagonal matrix

$$D_{\beta\gamma\delta} := \text{diag}(\lambda_{\alpha\beta\gamma\delta} I_3)_{\alpha=1}^n \quad (86)$$

is invertible.

Fix a triple (β, γ, δ) that is not constant. Consider the subcollection of columns of $\text{Flat}_1(\mathcal{T})$ with fixed camera suffix (β, γ, δ) and varying row choices within those cameras. By Lemma 9.1 these columns are precisely

$$D_{\beta\gamma\delta} \cdot M\omega_{tuv} \quad \text{with } c(t) = \beta, c(u) = \gamma, c(v) = \delta.$$

Condition (G3) implies that the vectors ω_{tuv} (with this suffix) span \mathbb{C}^4 : by (G1) the three rows in each camera form a basis of the corresponding hyperplane, and the cofactor construction is precisely the standard isomorphism $\bigwedge^3 \mathbb{C}^4 \simeq (\mathbb{C}^4)^*$ (we have identified a covector with a column vector by means of the chosen coordinates). Hence these columns span $D_{\beta\gamma\delta}W$. Therefore:

$$\text{span}\{\text{columns of } \text{Flat}_1(\mathcal{T}) \text{ with suffix } (\beta, \gamma, \delta)\} = D_{\beta\gamma\delta}W. \quad (87)$$

Now choose three distinct cameras b, c, d (possible since $n \geq 5$) and let U_1 denote the full column space of $\text{Flat}_1(\mathcal{T})$. By (87) applied to (b, c, d) we have $D_{bcd}W \subseteq U_1$, and since D_{bcd} is invertible and $\dim W = 4$, we get $\dim(D_{bcd}W) = 4$. Combined with (84) for $p = 1$, this forces

$$U_1 = D_{bcd}W.$$

Applying (87) to any non-constant triple (β, γ, δ) yields $D_{\beta\gamma\delta}W \subseteq U_1 = D_{bcd}W$, equivalently

$$H_{\beta\gamma\delta}W = W, \quad H_{\beta\gamma\delta} := D_{bcd}^{-1}D_{\beta\gamma\delta}.$$

But $H_{\beta\gamma\delta}$ is diagonal of the form $\text{diag}(h_\alpha I_3)_{\alpha=1}^n$. By Lemma 9.2, $H_{\beta\gamma\delta}$ must be a scalar multiple of the identity, so there exists $f_{\beta\gamma\delta} \in \mathbb{C}^*$ such that

$$\frac{\lambda_{\alpha\beta\gamma\delta}}{\lambda_{abcd}} = f_{\beta\gamma\delta} \quad \forall \alpha.$$

Setting $u_\alpha := \lambda_{abcd}$ gives the *mode-1 factorization*

$$\lambda_{\alpha\beta\gamma\delta} = u_\alpha f_{\beta\gamma\delta} \quad \text{whenever } (\beta, \gamma, \delta) \text{ is not constant.} \quad (88)$$

Repeating the same argument for the other flattenings (choosing, for the reference triples, any three distinct cameras such as the b, c, d above with the roles of the modes permuted) gives vectors $v, w \in (\mathbb{C}^*)^n$ —one may take for instance $v_\beta = \lambda_{b\beta cd}$ and $w_\gamma = \lambda_{bc\gamma d}$ —and functions g, h such that

$$\lambda_{\alpha\beta\gamma\delta} = v_\beta g_{\alpha\gamma\delta} \quad \text{whenever } (\alpha, \gamma, \delta) \text{ is not constant,} \quad (89)$$

$$\lambda_{\alpha\beta\gamma\delta} = w_\gamma h_{\alpha\beta\delta} \quad \text{whenever } (\alpha, \beta, \delta) \text{ is not constant.} \quad (90)$$

Step 6: gluing the one-mode factorizations. Let

$$E_1 := \{(\beta, \gamma, \delta) \in [n]^3 : (\beta, \gamma, \delta) \text{ is not constant}\}.$$

From (88)–(89) we produce a function of two indices. Fix $(\gamma, \delta) \in [n]^2$. Choose indices α_0, β_0 such that $(\beta_0, \gamma, \delta) \in E_1$ and $(\alpha_0, \gamma, \delta)$ is not constant (e.g. if $\gamma = \delta$, take $\alpha_0 \neq \gamma$ and $\beta_0 \neq \gamma$; otherwise any choice works). Define

$$r_{\gamma\delta} := \frac{\lambda_{\alpha_0\beta_0\gamma\delta}}{u_{\alpha_0}v_{\beta_0}}. \quad (91)$$

This is well-defined (independent of the choice): indeed, whenever both (88) and (89) apply we have

$$\frac{\lambda_{\alpha\beta\gamma\delta}}{u_\alpha v_\beta} = \frac{f_{\beta\gamma\delta}}{v_\beta} = \frac{g_{\alpha\gamma\delta}}{u_\alpha},$$

so the quotient depends only on (γ, δ) . In particular, for any fixed (γ, δ) and any β with $(\beta, \gamma, \delta) \in E_1$ we may choose an index α' (different from $\gamma = \delta$ if these two are equal) for which $(\alpha', \gamma, \delta)$ is not constant; applying the displayed equality to (α', β) shows that $f_{\beta\gamma\delta}/v_\beta = r_{\gamma\delta}$. Substituting this value in (88) gives, for every α (even if (α, γ, δ) should be constant),

$$\lambda_{\alpha\beta\gamma\delta} = u_\alpha v_\beta r_{\gamma\delta} \quad \text{for all } (\beta, \gamma, \delta) \in E_1 \text{ and all } \alpha. \quad (92)$$

Now use (90) to split $r_{\gamma\delta}$ multiplicatively. Fix $\delta \in [n]$ and choose α_0, β_0 with $\alpha_0 \neq \delta$ and $\beta_0 \neq \delta$. Then $(\alpha_0, \beta_0, \delta)$ is not constant, so (90) applies, and moreover $(\beta_0, \gamma, \delta) \in E_1$ for every γ (since $\beta_0 \neq \delta$). Thus, for all γ ,

$$u_{\alpha_0} v_{\beta_0} r_{\gamma\delta} = \lambda_{\alpha_0\beta_0\gamma\delta} = w_\gamma h_{\alpha_0\beta_0\delta}.$$

Hence $r_{\gamma\delta}/w_\gamma$ is independent of γ ; define $x_\delta \in \mathbb{C}^*$ by

$$x_\delta := \frac{h_{\alpha_0\beta_0\delta}}{u_{\alpha_0}v_{\beta_0}}.$$

Then $r_{\gamma\delta} = w_\gamma x_\delta$, and substituting into (92) yields

$$\lambda_{\alpha\beta\gamma\delta} = u_\alpha v_\beta w_\gamma x_\delta \quad \text{whenever } (\beta, \gamma, \delta) \in E_1. \quad (93)$$

It remains to treat the triples *not* in E_1 , i.e. $(\beta, \gamma, \delta) = (\beta, \beta, \beta)$. For $\alpha \neq \beta$, the quadruple $(\alpha, \beta, \beta, \beta)$ is not all identical, so $\lambda_{\alpha\beta\beta\beta} \neq 0$ by (85). Fix such $\alpha \neq \beta$ and choose $\eta \in [n]$ with $\eta \neq \beta$. Then (90) (with $\delta = \beta$ and (α, β, β) not constant) gives

$$\lambda_{\alpha\beta\gamma\beta} = w_\gamma h_{\alpha\beta\beta} \quad \forall \gamma. \quad (94)$$

Taking $\gamma = \eta$ and using (93) for the quadruple $(\alpha, \beta, \eta, \beta)$ (which has $(\beta, \eta, \beta) \in E_1$ since $\eta \neq \beta$) yields

$$w_\eta h_{\alpha\beta\beta} = \lambda_{\alpha\beta\eta\beta} = u_\alpha v_\beta w_\eta x_\beta,$$

so $h_{\alpha\beta\beta} = u_\alpha v_\beta x_\beta$. Plugging $\gamma = \beta$ into (94) gives

$$\lambda_{\alpha\beta\beta\beta} = w_\beta h_{\alpha\beta\beta} = u_\alpha v_\beta w_\beta x_\beta.$$

Thus the factorization holds for all quadruples that are not all identical:

$$\lambda_{\alpha\beta\gamma\delta} = u_\alpha v_\beta w_\gamma x_\delta \quad \text{for every } (\alpha, \beta, \gamma, \delta) \text{ not all identical.}$$

Step 7: returning to real scalars. It remains only to justify the passage from the complex argument to the real statement formulated in the problem. We shall use the following elementary observation. Suppose that all the numbers $\lambda_{\alpha\beta\gamma\delta}$ are real and satisfy (85), and suppose that for some complex vectors $u, v, w, x \in (\mathbb{C}^*)^n$ the equality

$$\lambda_{\alpha\beta\gamma\delta} = u_\alpha v_\beta w_\gamma x_\delta \quad (95)$$

holds for every non-identical quadruple. Then the four vectors may be chosen with real (non-zero) entries.

For completeness I give the short proof. Choose three distinct indices $p, q, r \in [n]$ (this is the only place where fewer than five indices would in fact have sufficed). From (95) we get, for example,

$$\frac{u_\alpha}{u_p} = \frac{\lambda_{\alpha p q r}}{\lambda_{p p q r}} \in \mathbb{R}^* \quad (\alpha \in [n]),$$

and in the same way $v_\beta/v_p = \lambda_{p\beta q r}/\lambda_{p p q r}$, $w_\gamma/w_p = \lambda_{p q \gamma r}/\lambda_{p q p r}$ and $x_\delta/x_p = \lambda_{p q r \delta}/\lambda_{p q r p}$ are real and non-zero. (Every displayed denominator is legitimate because the corresponding quadruple is not all identical.) Thus all entries of (say) u have the same complex phase up to a sign, and the same is true for v, w, x . Write $u_p = |u_p|e^{i\theta}$, $v_p = |v_p|e^{i\phi}$, $w_p = |w_p|e^{i\psi}$ and $x_p = |x_p|e^{i\chi}$. Since, for example, the product corresponding to the non-identical quadruple (p, p, q, r) , $u_p v_p w_q x_r = \lambda_{p p q r}$, is a non-zero real number, the sum $\theta + \phi + \psi + \chi$ is congruent to 0 modulo π (the possible signs of the real ratios such as w_q/w_p and x_r/x_p only add integer multiples of π).

Define $\tilde{u}_\alpha = e^{-i\theta} u_\alpha$, $\tilde{v}_\beta = e^{-i\phi} v_\beta$, $\tilde{w}_\gamma = e^{-i\psi} w_\gamma$ and $\tilde{x}_\delta = e^{-i\chi} x_\delta$. These numbers are all real and non-zero, and for every non-identical quadruple their product equals $e^{-i(\theta+\phi+\psi+\chi)} \lambda_{\alpha\beta\gamma\delta}$. If this common factor is -1 rather than $+1$, we simply change the sign of one of the four real vectors. In either case (95) holds with real factors.

Conclusion. The map \mathbf{F} defined by (82) is independent of the cameras, each coordinate has degree 5 (independent of n), and for Zariski-generic $A^{(1)}, \dots, A^{(n)}$ it satisfies the desired characterization. More explicitly, for every real array λ obeying (85),

$$\mathbf{F}((\lambda_{\alpha\beta\gamma\delta} Q^{(\alpha\beta\gamma\delta)})_{\alpha,\beta,\gamma,\delta}) = 0 \iff \exists u, v, w, x \in (\mathbb{R}^*)^n \text{ such that } \lambda_{\alpha\beta\gamma\delta} = u_\alpha v_\beta w_\gamma x_\delta$$

for all non-identical quadruples. Conversely, any such real factorization (and a fortiori any complex one) makes all the minors in (82) vanish, as was shown in Step 4. This is precisely the polynomial map required in the statement of the problem. \square

10 Kernelized CP–ALS subproblem with missing data: matrix-free PCG with Kronecker preconditioning

Problem

Given a d -way tensor $\mathcal{T} \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ such that the data is unaligned (meaning the tensor \mathcal{T} has missing entries), we consider the problem of computing a CP decomposition of rank r where some modes are infinite-dimensional and constrained to be in a Reproducing Kernel Hilbert Space (RKHS). We want to solve this using an alternating optimization approach, and our question is focused on the mode- k subproblem for an infinite-dimensional mode. For the subproblem, then CP factor matrices $A_1, \dots, A_{k-1}, A_{k+1}, \dots, A_d$ are fixed, and we are solving for A_k .

Our notation is as follows. Let $N = \prod_i n_i$ denote the product of all sizes. Let $n \equiv n_k$ be the size of mode k , let $M = \prod_{i \neq k} n_i$ be the product of all dimensions except k , and assume $n \ll M$. Since the data are unaligned, this means only a subset of \mathcal{T} 's entries are observed, and we let $q \ll N$ denote the number of observed entries. We let $T \in \mathbb{R}^{n \times M}$ denote the mode- k unfolding of the tensor \mathcal{T} with all missing entries set to zero. The vec operation creates a vector from a matrix by stacking its columns, and we let $S \in \mathbb{R}^{N \times q}$ denote the selection matrix (a subset of the $N \times N$ identity matrix) such that $S^T \text{vec}(T)$ selects the q known entries of the tensor \mathcal{T} from the vectorization of its mode- k unfolding. We let $Z = A_d \odot \dots \odot A_{k+1} \odot A_{k-1} \odot \dots \odot A_1 \in \mathbb{R}^{M \times r}$ be the Khatri-Rao product of the factor matrices corresponding to all modes except mode k . We let $B = TZ$ denote the MTTKRP of the tensor \mathcal{T} and Khatri-Rao product Z .

We assume $A_k = KW$ where $K \in \mathbb{R}^{n \times n}$ denotes the psd RKHS kernel matrix for mode k . The matrix W of size $n \times r$ is the unknown for which we must solve. The system to be solved is

$$[(Z \otimes K)^T S S^T (Z \otimes K) + \lambda(I_r \otimes K)] \text{vec}(W) = (I_r \otimes K) \text{vec}(B). \quad (96)$$

Here, I_r denotes the $r \times r$ identity matrix. This is a system of size $nr \times nr$. Using a standard linear solver costs $O(n^3 r^3)$, and explicitly forming the matrix is an additional expense.

Explain how an iterative preconditioned conjugate gradient linear solver can be used to solve this problem more efficiently. Explain the method and choice of preconditioner. Explain in detail how the matrix-vector products are computed and why this works. Provide complexity analysis. We assume $n, r < q \ll N$. Avoid any computation of order N .

Solution (matrix-free PCG with Kronecker preconditioning)

We show how to solve (96) efficiently using a matrix-free preconditioned conjugate gradient (PCG) method. The central idea is (i) to express all masked contractions using only the observed indices, and (ii) to choose a preconditioner that is spectrally close to the true operator and admits a fast inverse via Kronecker eigenstructure.

Observed-index notation (eliminating S and avoiding Z). Let the set of observed entries in the mode- k unfolding be

$$\Omega = \{(i_\ell, j_\ell)\}_{\ell=1}^q, \quad i_\ell \in \{1, \dots, n\}, \quad j_\ell \in \{1, \dots, M\}.$$

(Here j_ℓ encodes the multi-index over modes $\neq k$.) For each observation ℓ , define the corresponding row of the Khatri–Rao product by

$$z_\ell^T := e_{j_\ell}^T Z \in \mathbb{R}^{1 \times r}, \quad \ell = 1, \dots, q,$$

and collect these rows into the matrix $Z_\Omega \in \mathbb{R}^{q \times r}$. Crucially, Z_Ω can be formed without ever constructing $Z \in \mathbb{R}^{M \times r}$: for $s = 1, \dots, r$,

$$(z_\ell)_s = \prod_{m \neq k} A_m(i_m^{(\ell)}, s),$$

where $(i_m^{(\ell)})_{m \neq k}$ is the multi-index corresponding to j_ℓ . Thus Z_Ω is only $q \times r$.

Also define the row-selection matrix $R_\Omega \in \mathbb{R}^{q \times n}$ by

$$(R_\Omega)_{\ell, i_\ell} = 1, \quad (R_\Omega)_{\ell, i} = 0 \text{ for } i \neq i_\ell.$$

Then for any $U \in \mathbb{R}^{n \times r}$, $(R_\Omega U)_{\ell, :} = U_{i_\ell, :}$.

Matrix-free operator application. It is convenient to work with matrices rather than the vectorized unknown. For any $X \in \mathbb{R}^{n \times r}$, define $x = \text{vec}(X)$. The Kronecker identity

$$(Z \otimes K) \text{vec}(X) = \text{vec}(K X Z^T) \quad (97)$$

implies that the masked product $SS^T(Z \otimes K) \text{vec}(X)$ extracts exactly the q scalars

$$s_\ell(X) := e_{i_\ell}^T K X z_\ell = (K X)_{i_\ell, :} \cdot z_\ell, \quad \ell = 1, \dots, q.$$

Equivalently, with $\mathbf{1}_r$ the r -vector of ones,

$$s(X) = (R_\Omega K X \odot Z_\Omega) \mathbf{1}_r \in \mathbb{R}^q. \quad (98)$$

A direct expansion shows

$$\begin{aligned} & (Z \otimes K)^T S S^T (Z \otimes K) \text{vec}(X) \\ &= \sum_{\ell=1}^q s_\ell(X) (z_\ell \otimes K e_{i_\ell}) = \text{vec} \left(K \sum_{\ell=1}^q e_{i_\ell} s_\ell(X) z_\ell^T \right) = \text{vec}(K R_\Omega^T \text{Diag}(s(X)) Z_\Omega). \end{aligned} \quad (99)$$

Therefore the full coefficient operator in (96) can be applied without forming any N - or M -dimensional objects:

$$A \text{vec}(X) = \text{vec}(K R_\Omega^T \text{Diag}(s(X)) Z_\Omega + \lambda K X), \quad \text{where } s(X) \text{ is given by (98)}. \quad (100)$$

Right-hand side without forming $B = TZ$. Let $t \in \mathbb{R}^q$ collect the observed tensor values in the unfolding: $t_\ell := T_{i_\ell, j_\ell}$. Then the sparse MTTKRP satisfies

$$B = TZ = R_\Omega^T \text{Diag}(t) Z_\Omega \in \mathbb{R}^{n \times r}, \quad (101)$$

and the right-hand side becomes

$$b = (I_r \otimes K) \text{vec}(B) = \text{vec}(KB). \quad (102)$$

Again, no $M \times r$ matrix Z is formed, and no computation scales with N .

PCG formulation. Assume K is positive definite, or replace it by $K + \varepsilon I$ with a small nugget $\varepsilon > 0$ (standard in kernel methods). Then the operator in (96) is symmetric positive definite (SPD), and CG applies. PCG iterates on the linear system

$$A w = b, \quad w = \text{vec}(W),$$

using only: (i) matrix-free applications of A via (100), and (ii) applications of a preconditioner inverse P^{-1} described next.

A Kronecker preconditioner from mean masking. Let $\rho := q/N$ be the observation density. If the observed set is approximately uniform, a common and effective approximation is

$$SS^T \approx \rho I.$$

Under this replacement,

$$(Z \otimes K)^T(\rho I)(Z \otimes K) = \rho(Z^T Z) \otimes (K^T K) = \rho \Gamma \otimes K^2, \quad \Gamma := Z^T Z \in \mathbb{R}^{r \times r}.$$

This motivates the SPD preconditioner

$$P := \rho \Gamma \otimes K^2 + \lambda I_r \otimes K. \quad (103)$$

The matrix Γ is cheap to form without Z by the standard CP identity

$$\Gamma = Z^T Z = (A_d^T A_d) * \cdots * (A_{k+1}^T A_{k+1}) * (A_{k-1}^T A_{k-1}) * \cdots * (A_1^T A_1), \quad (104)$$

where $*$ denotes the Hadamard product. This requires only the $r \times r$ Gram matrices of the other modes.

Fast application of P^{-1} via eigendecompositions. Compute once the eigendecompositions

$$K = U \text{Diag}(\sigma_1, \dots, \sigma_n) U^T, \quad \Gamma = V \text{Diag}(\gamma_1, \dots, \gamma_r) V^T,$$

with $U \in \mathbb{R}^{n \times n}$, $V \in \mathbb{R}^{r \times r}$ orthogonal and $\sigma_i, \gamma_a \geq 0$. Then (103) becomes

$$P = (V \otimes U) \text{Diag}\left(\rho \gamma_a \sigma_i^2 + \lambda \sigma_i\right)_{i=1, \dots, n; a=1, \dots, r} (V \otimes U)^T.$$

Hence for any residual $r = \text{vec}(R)$ with $R \in \mathbb{R}^{n \times r}$,

$$\hat{R} := U^T R V, \quad \hat{Y}_{i,a} := \frac{\hat{R}_{i,a}}{\rho \gamma_a \sigma_i^2 + \lambda \sigma_i}, \quad P^{-1} r = \text{vec}(U \hat{Y} V^T). \quad (105)$$

Importantly, K^2 is never formed explicitly; only the eigenvalues σ_i^2 are used. If some $\sigma_i = 0$, either add a nugget $K \leftarrow K + \varepsilon I$ or restrict to $\text{range}(K)$; in either case, P remains SPD.

Matrix-free matvec algorithm (what PCG actually computes). Given $x = \text{vec}(X)$ with $X \in \mathbb{R}^{n \times r}$, compute $y = Ax$ as:

1. $U \leftarrow KX$.
2. For $\ell = 1, \dots, q$, compute $s_\ell \leftarrow U_{i_\ell, :} \cdot z_\ell$ (rowwise dot product).
3. Accumulate $G \in \mathbb{R}^{n \times r}$ by scatter-add:

$$G_{i_\ell, :} += s_\ell z_\ell^T \quad (\ell = 1, \dots, q), \quad \text{equivalently } G = R_\Omega^T \text{Diag}(s) Z_\Omega.$$

4. Return $y = \text{vec}(KG + \lambda U)$.

The right-hand side is computed once using (101)–(102): form $H = R_\Omega^T \text{Diag}(t) Z_\Omega$ by the same scatter pattern as in step 3 (with s replaced by t), then set $b = \text{vec}(KH)$.

Why PCG converges quickly. Let e_m denote the PCG error after m iterations. Since A and P are SPD, standard PCG theory yields

$$\frac{\|e_m\|_A}{\|e_0\|_A} \leq 2 \left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^m, \quad \kappa := \kappa(P^{-1/2} A P^{-1/2}).$$

When the mask is approximately uniform random, SS^T is spectrally close to ρI and thus A is close (in a spectral sense) to the Kronecker-structured approximation (103), which typically clusters the eigenvalues of $P^{-1}A$ and yields small iteration counts.

Complexity (avoiding any $\mathcal{O}(N)$ work). All costs below are expressed in terms of (n, r, q) (and small mode sizes for forming Gram matrices), and no object of size N or M is ever formed.

Precomputation.

- Build $Z_\Omega \in \mathbb{R}^{q \times r}$ from observed multi-indices: $\mathcal{O}(qr(d-1))$ multiplications (or $\mathcal{O}(qr)$ if factor rows are accessed efficiently).
- Form $\Gamma = Z^T Z$ via (104): $\mathcal{O}(\sum_{m \neq k} n_m r^2)$ to form each $A_m^T A_m$ plus $\mathcal{O}((d-1)r^2)$ Hadamard products.
- Eigendecompositions: $\mathcal{O}(n^3 + r^3)$.
- Right-hand side $b = \text{vec}(KH)$ with $H = R_\Omega^T \text{Diag}(t) Z_\Omega$: $\mathcal{O}(qr + n^2 r)$.

Per PCG iteration.

- Two dense kernel multiplies (KX and KG): $\mathcal{O}(n^2 r)$ each, i.e. $\mathcal{O}(n^2 r)$ up to constants.
- Two observation-driven contractions (compute s_ℓ and scatter-add G): $\mathcal{O}(qr)$.
- Apply P^{-1} via (105) (two basis changes and diagonal scaling): $\mathcal{O}(n^2 r + nr^2)$.
- Vector inner products and saxpys on \mathbb{R}^{nr} : $\mathcal{O}(nr)$ (lower order).

Thus one PCG iteration costs

$$\mathcal{O}(qr + n^2 r + nr^2),$$

and T_{cg} iterations cost

$$\mathcal{O}\left(T_{\text{cg}}(qr + n^2 r + nr^2)\right).$$

This is substantially smaller than the $\mathcal{O}((nr)^3)$ cost of a dense solve, and the method never performs any computation of order N .

A Appendix: Prompting strategy

For the solutions presented in the main body, we orchestrated the model work manually. Here we describe a simple, programmatic, and problem-independent setup that we tried after the original deadline and we expect it produces results of roughly equivalent quality.

A.1 Solution strategy

The framework is simple:

- Generate a small number of seed ideas.
- Prompt the model to solve the given problem using each of the seed ideas.
- Repeat up to 3 times:
 - Verify (i) correctness of the proof and (ii) validity of any cited material and bibliographic references.
 - If gaps are found, prompt the model to revise the drafted solution.
- If the verifications pass, typeset the resulting solution.

In Section B, we share the solutions to Problem #6 for one such run. Here we generated four solutions for every seed idea (see Section A.4.1). For the idea of using BSS, our internal model produced 3/4 solutions that were then judged (by the same model) to be correct. For two of the verified runs, there were zero rounds of revisions; for the third run, there was one round of revisions.

A.2 Prompt templates

A.2.1 Generate ideas template

You are a problem idea generator! Your goal is to generate FIVE powerful high-level approaches to solving the given problem. You can use your own knowledge or a literature search. Your ideas will be passed along to a team of technical problem solvers who will work on each idea in parallel.

Problem

{{problem}}

Instructions

You should output a JSON dictionary inside a markdown block in the following format:

```
““json
{
  "idea_1": "brief description of the first promising idea",
  "idea_2": "brief description of the second promising idea",
  "idea_3": "brief description of the third promising idea",
  "idea_4": "brief description of the fourth promising idea",
  "idea_5": "brief description of the fifth promising idea",
}
““
```

A.3 Solve template

```
{{problem}}
{% if idea %}
You should pursue the following approach:
{{idea}}
Keep working hard until you have a complete and rigorous solution.
{% endif %}
```

A.3.1 Verify template

I need you to check the following solution very carefully and let me know if you find any gaps that cannot easily be fixed.

```
{% if check_refs %}You should also carefully check that all the bibliographic references are valid.{% endif %}
```

Problem

```
{{problem}}
```

Solution?

```
{{solution}}
```

Instructions

Something is considered a "gap" if it is a critical error that cannot easily be fixed. If the proposed solution has no gaps, you should set the "correct" key to true below and you provide an empty list "[]" for the "gaps" key. If there are gaps, you should set the "correct" key to false below and you should provide a list of gaps as specified.

Your output should be in the format:

```
<CORRECT>one of 'true' or 'false'</CORRECT>
<GAPS>
  [required section if correct is false]
  - Brief explanation of gap_1
  - Brief explanation of gap_2
  - ...
</GAPS>
```

A.3.2 Solution refinement template

You previously submitted a solution to the given problem, which you will find below. Unfortunately, the referee found some issues with your solution that you will need to address carefully and rigorously. Keep working until you have revised the argument, or if you find it to be impossible to fix, you should work on a new approach.

Problem

{{problem}}

Solution

{{solution}}

Critique

{{gaps}}

A.3.3 Typesetting template

Problem

{{problem}}

Solution

{{solution}}

Please give me a full, rigorous, and complete latex draft of the provided argument. The draft should be written in a professional mathematical style.

Clearly demarcate the problem and the solution.

Use a single `\section{...}` (but as many subsection as you wish).

If you need citations to known literature, you may include a bibliography (use the "`\bibitem`" format for bibliography entries).

A.4 Solution attempts for Problem #6

A.4.1 Generated ideas

- Normalize by the Laplacian: write $L = \sum_e L_e$, set $A_e := L^{-1/2} L_e L^{-1/2}$ so $\sum_e A_e = I$ on $\text{range}(L)$. Then build an r -coloring with $r = \Theta(1/\varepsilon)$ via a deterministic Batson–Spielman–Srivastava-style one-sided barrier potential $\Phi_u(M) = \text{tr}((uI - M)^{-1})$, maintaining that the monochromatic-edge sum $M = \sum_{(u,v) \text{ same color}} A_{uv}$ stays $\preceq \varepsilon I$. Finally take the largest color class S , which inherits $L_S \preceq \varepsilon L$ and has $|S| = \Omega(\varepsilon n)$.
- SDP + rounding: formulate a relaxation maximizing $\sum_v x_v$ with $x_v \in [0, 1]$ and a lifted constraint ensuring $\sum_{uv \in E} X_{uv} A_{uv} \preceq \varepsilon I$ where $X \approx xx^T$ and $\text{diag}(X) = x$. Show a uniform fractional solution with total mass $\Theta(\varepsilon n)$ is feasible (after choosing r /parameters). Then apply pipage/swap rounding (or randomized rounding with a pessimistic estimator) together with matrix Chernoff-style bounds to round to an integral S while preserving the PSD constraint up to constants.
- Interlacing-polynomials/MSS approach: view M as a random sum of rank-1 matrices $w_e w_e^T$ (with $w_e := L^{-1/2}(e_u - e_v)$) gated by random vertex coloring. Compute/bound the expected characteristic (or mixed characteristic) polynomial of M and use real-rootedness + interlacing to show there exists a coloring whose $\lambda_{\max}(M)$ is at most the largest root of the expectation (about $O(1/r)$). This bypasses tail bounds and can potentially give near-optimal constants.
- Matrix-discrepancy/paving/online-load-balancing viewpoint: interpret vertex coloring into r bins as distributing the edge matrices A_e into monochromatic loads (each monochro-

matic edge contributes to exactly one bin). Use a discrepancy-style theorem (e.g., Banaszczyk/Spencer analogs, or Kadison–Singer-type paving) or an online mirror-descent potential to choose colors so that the PSD load in every bin (or at least in the total monochromatic load) is $O(1/r) \cdot I$. Then take a largest bin to get an ε -light S of size $\Omega(\varepsilon n)$.

- Probabilistic method with matrix concentration: color vertices i.i.d. into $r \approx \Theta(1/\varepsilon)$ colors and study $M := L^{-1/2} (\sum \text{monochromatic edges } L_e) L^{-1/2}$. One has $\mathbb{E}[M] = (1/r)I$. Reveal colors sequentially and apply a matrix Azuma/Freedman inequality; crucially, one vertex exposure only changes M by a PSD sum over incident edges, which is $\preceq I$ because it is a sub-sum of $\{A_e\}$. Conclude $\Pr(M \preceq (C/r)I) > 0$, hence some coloring works; pick the largest color class for $|S| \geq \Omega(n/r) = \Omega(\varepsilon n)$.

B Model-validated solution attempts

B.1 Large ε -light vertex subsets

The answer is yes. In fact one can take $c = 1/204$.

Theorem B.1. *There exists an absolute constant $c > 0$ such that for every finite undirected graph $G = (V, E)$ with Laplacian L and every $\varepsilon \in [0, 1]$, there exists a set $S \subseteq V$ with $|S| \geq c\varepsilon|V|$ such that*

$$L_S \preceq \varepsilon L.$$

One may take $c = 1/204$.

Conventions and Laplacian decompositions Let $V = \{1, \dots, n\}$, and let e_1, \dots, e_n be the standard basis of $\mathbb{R}^V \cong \mathbb{R}^n$. Suppose each edge $e = \{i, j\}$ carries a weight $w_e \geq 0$ (the unweighted case corresponds to $w_e \equiv 1$). Define the *edge Laplacian*

$$L_e := w_e(e_i - e_j)(e_i - e_j)^\top,$$

so that the graph Laplacian is the sum

$$L = \sum_{e \in E} L_e.$$

For any $F \subseteq E$, write

$$L_F := \sum_{e \in F} L_e;$$

this equals the Laplacian of the weighted graph (V, F) with inherited weights. In particular, for $S \subseteq V$ and $F = E(S, S)$, we have $L_S = L_F$.

For symmetric matrices A, B , write $A \preceq B$ if $B - A$ is positive semidefinite (PSD), i.e. $x^\top Ax \leq x^\top Bx$ for all x .

Reduction to the non-constant subspace Let $\ker(L)$ be the nullspace of L ; it consists of vectors that are constant on each connected component of G . Let

$$H := \ker(L)^\perp$$

be its orthogonal complement. On H , the restriction of L is positive definite. Let L^\dagger denote the Moore–Penrose pseudoinverse of L , and write $L^{\dagger/2} := (L^\dagger)^{1/2}$ for its (unique) PSD square root. Equivalently, $L^{\dagger/2}$ acts as $L^{-1/2}$ on H and as 0 on $\ker(L)$.

For each edge e , define the normalized PSD matrix (acting on H)

$$A_e := L^{\dagger/2} L_e L^{\dagger/2}. \tag{106}$$

Then each A_e is PSD on H , and we have the identity

$$\sum_{e \in E} A_e = L^{\dagger/2} L L^{\dagger/2} = P_H, \tag{107}$$

where P_H is the orthogonal projection onto H . Thus, viewing all operators as acting on H , we may regard (128) as

$$\sum_{e \in E} A_e = I_H.$$

Finally, note that if $F \subseteq E$ then $\ker(L) \subseteq \ker(L_F)$: indeed any vector constant on each component of G is constant on each component of (V, F) since (V, F) is a subgraph of G . Consequently L and L_F both vanish on $\ker(L)$, and to prove $L_F \preceq \varepsilon L$ on all of \mathbb{R}^V it suffices to prove it on H .

B.1.1 A partial coloring lemma

Lemma B.2. *Let $G = (V, E)$ be as above, and let A_e be as in (106) acting on H , so that $\sum_{e \in E} A_e = I_H$. Fix an integer $r \geq 1$. Then one can color at least $\lfloor n/2 \rfloor$ vertices using colors $\{1, \dots, r\}$ such that, if M denotes the sum of A_e over all edges whose endpoints are both colored and have received the same color, then*

$$M \preceq \frac{44}{r} I_H.$$

Proof of Lemma B.2 All matrices and traces in this proof are taken on the space H . Let $m := \dim(H)$, so $m \leq n$.

Potential function. For $u > 0$ and a symmetric matrix X with $X \prec uI_H$, define

$$\Phi^u(X) := \text{tr}((uI_H - X)^{-1}).$$

This is finite exactly when $X \prec uI_H$.

Setup. We build a partial r -coloring iteratively. At any stage, for each color $i \in \{1, \dots, r\}$ let C_i be the set of vertices currently colored with i , and let U be the set of uncolored vertices.

Let M be the sum of A_e over all edges whose endpoints are both colored and have the same color:

$$M := \sum_{\substack{\{x,y\} \in E \\ x,y \notin U \\ \text{col}(x) = \text{col}(y)}} A_{\{x,y\}}. \quad (108)$$

If we color a vertex $v \in U$ with color i , then the only new monochromatic edges created are those between v and already-colored vertices in C_i . Thus the increment to M is the PSD matrix

$$X_{v,i} := \sum_{\substack{x \in C_i \\ \{x,v\} \in E}} A_{\{x,v\}}. \quad (109)$$

The update is $M \leftarrow M + X_{v,i}$.

A resolvent estimate.

Claim B.3. Let A be symmetric with $A \prec uI_H$, and set $R := (uI_H - A)^{-1}$. Let $X \succeq 0$ and define $Y := R^{1/2} X R^{1/2}$. If $\|Y\| < 1$, then

$$\text{tr}((uI_H - A - X)^{-1}) - \text{tr}(R) \leq \frac{\text{tr}(R X R)}{1 - \|Y\|}.$$

Proof. We have

$$(uI_H - A - X)^{-1} = R^{1/2} (I_H - Y)^{-1} R^{1/2}.$$

Since $\|Y\| < 1$ and $Y \succeq 0$, the Neumann series converges and

$$(I_H - Y)^{-1} = \sum_{k \geq 0} Y^k.$$

Moreover, for $k \geq 1$, $Y^k \preceq \|Y\|^{k-1} Y$, hence

$$(I_H - Y)^{-1} - I_H = \sum_{k \geq 1} Y^k \preceq \sum_{k \geq 1} \|Y\|^{k-1} Y = \frac{1}{1 - \|Y\|} Y.$$

Taking traces and using cyclicity gives the desired inequality. \square

Choosing a good vertex-color pair. Fix a stage with current sets $(C_i)_{i=1}^r$, U , and a matrix M satisfying $M \prec uI_H$ for some $u > 0$. Define the slightly larger barrier

$$u' := u + \frac{4}{r|U|}, \quad (110)$$

and let

$$R := (u'I_H - M)^{-1}.$$

Define the *boundary matrix* between colored and uncolored vertices:

$$B := \sum_{\substack{\{x,y\} \in E \\ x \notin U, y \in U}} A_{\{x,y\}}.$$

Since B is a sub-sum of $\sum_{e \in E} A_e = I_H$ (and each $A_e \succeq 0$), we have

$$0 \preceq B \preceq I_H. \quad (111)$$

Observe that every boundary edge $\{x, v\}$ with $v \in U$ and $x \notin U$ appears in exactly one $X_{v,i}$, namely with $i = \text{col}(x)$. Hence

$$\sum_{v \in U} \sum_{i=1}^r X_{v,i} = B. \quad (112)$$

Taking traces and using (111) yields

$$\sum_{v \in U} \sum_{i=1}^r \text{tr}(RX_{v,i}R) = \text{tr}(RBR) \leq \text{tr}(R^2), \quad (113)$$

and similarly

$$\sum_{v \in U} \sum_{i=1}^r \text{tr}(RX_{v,i}) = \text{tr}(RB) \leq \text{tr}(R). \quad (114)$$

Claim B.4. There exists a pair $(v, i) \in U \times \{1, \dots, r\}$ such that simultaneously

$$\text{tr}(RX_{v,i}R) \leq \frac{2 \text{tr}(R^2)}{r|U|}, \quad \text{tr}(RX_{v,i}) \leq \frac{2 \text{tr}(R)}{r|U|}. \quad (115)$$

Proof. Let $N := r|U|$ and index the N pairs (v, i) as $p = 1, \dots, N$. Write $a_p := \text{tr}(RX_{v,i}R)$ and $b_p := \text{tr}(RX_{v,i})$. By (113), $\sum_{p=1}^N a_p \leq \text{tr}(R^2)$. If at least $N/2$ indices had $a_p > 2 \text{tr}(R^2)/N$, then $\sum_p a_p > (N/2) \cdot 2 \text{tr}(R^2)/N = \text{tr}(R^2)$, a contradiction. Hence fewer than $N/2$ indices violate the first inequality. The same argument using (114) shows fewer than $N/2$ indices violate the second inequality. Therefore fewer than N indices violate at least one inequality, so at least one index satisfies both. \square

The iterative construction. We initialize

$$M_0 := 0, \quad u_0 := \frac{40}{r}, \quad (116)$$

and $U_0 := V$ (all vertices uncolored). Note that

$$\Phi^{u_0}(M_0) = \text{tr}((u_0 I_H)^{-1}) = \frac{m}{u_0} \leq \frac{n}{u_0} = \frac{nr}{40}. \quad (117)$$

Assume inductively that after coloring t vertices we have (M_t, u_t, U_t) and

$$M_t \prec u_t I_H, \quad \Phi^{u_t}(M_t) \leq \Phi^{u_0}(M_0). \quad (118)$$

We will show how to color one more vertex provided $|U_t| \geq n/2$.

Set $U := U_t$, $M := M_t$, $u := u_t$, and define u' by (110). Let $R := (u'I_H - M)^{-1}$. Note that increasing the barrier decreases the potential:

$$\Phi^{u'}(M) \leq \Phi^u(M) \leq \Phi^{u_0}(M_0),$$

so in particular

$$\text{tr}(R) = \Phi^{u'}(M) \leq \Phi^{u_0}(M_0) \leq \frac{nr}{40}. \quad (119)$$

Choose a pair (v, i) satisfying (115) (guaranteed by Claim B.4), and set $X := X_{v,i}$ and $M^+ := M + X$. From (115) and (119), when $|U| \geq n/2$ we have

$$\text{tr}(RX) \leq \frac{2 \text{tr}(R)}{r|U|} \leq \frac{2(nr/40)}{r(n/2)} = \frac{1}{10}. \quad (120)$$

Let $Y := R^{1/2}XR^{1/2}$. Since $X \succeq 0$, we have $Y \succeq 0$ and $\|Y\| \leq \text{tr}(Y) = \text{tr}(RX) \leq 1/10$, so $\|Y\| < 1$. Claim B.3 applies and gives

$$\Phi^{u'}(M^+) - \Phi^{u'}(M) \leq \frac{\text{tr}(RXR)}{1 - \|Y\|} \leq \frac{\frac{2 \text{tr}(R^2)}{r|U|}}{1 - 1/10} = \frac{20}{9} \cdot \frac{\text{tr}(R^2)}{r|U|}, \quad (121)$$

where we used (115) in the second inequality.

Next we estimate the drop from moving the barrier from u to u' . Let $\lambda_1, \dots, \lambda_m$ be the eigenvalues of M on H , and set $s_j := u - \lambda_j > 0$. Then $u' - \lambda_j = s_j + \Delta$ with $\Delta := u' - u = 4/(r|U|)$. Hence

$$\begin{aligned} \Phi^u(M) - \Phi^{u'}(M) &= \sum_{j=1}^m \left(\frac{1}{s_j} - \frac{1}{s_j + \Delta} \right) = \sum_{j=1}^m \frac{\Delta}{s_j(s_j + \Delta)} \\ &= \sum_{j=1}^m \frac{\Delta}{(s_j + \Delta)^2} \cdot \frac{s_j + \Delta}{s_j} \geq \sum_{j=1}^m \frac{\Delta}{(s_j + \Delta)^2} = \Delta \text{tr}(R^2) \\ &= \frac{4}{r|U|} \text{tr}(R^2). \end{aligned} \quad (122)$$

Combining (121) and (122) yields

$$\Phi^{u'}(M^+) \leq \Phi^u(M) - \left(4 - \frac{20}{9}\right) \frac{\text{tr}(R^2)}{r|U|} \leq \Phi^u(M).$$

Thus $\Phi^{u'}(M^+) \leq \Phi^u(M) \leq \Phi^{u_0}(M_0)$, proving the potential part of the invariant at the next step.

Moreover, since $\|Y\| < 1$ we have $I_H - Y \succ 0$, hence

$$u'I_H - M^+ = u'I_H - M - X = R^{-1/2}(I_H - Y)R^{-1/2} \succ 0,$$

so $M^+ \prec u'I_H$. Therefore the full invariant (162) is preserved with the updates $M_{t+1} := M^+$ and $u_{t+1} := u'$.

Termination and bound on M . We repeat the above step until $T := \lfloor n/2 \rfloor$ vertices have been colored. During this process, for all $t < T$ we have $|U_t| = n - t \geq n/2$. Hence, summing the barrier increments (110) gives

$$u_T = u_0 + \sum_{t=0}^{T-1} \frac{4}{r|U_t|} = \frac{40}{r} + \frac{4}{r} \sum_{t=0}^{T-1} \frac{1}{n-t} \leq \frac{40}{r} + \frac{4}{r} \cdot \frac{T}{n/2} \leq \frac{44}{r}. \quad (123)$$

Since $M_T \prec u_T I_H$, we conclude

$$M_T \preceq \frac{44}{r} I_H,$$

which is the desired inequality. This completes the proof of Lemma B.2. \square

B.1.2 Derivation of Theorem B.9 from the lemma

Fix $\varepsilon \in (0, 1]$ and set

$$r := \left\lceil \frac{50}{\varepsilon} \right\rceil. \quad (124)$$

Apply Lemma B.2 with this r . We obtain a partial r -coloring of $T := \lfloor n/2 \rfloor$ vertices, and a matrix M (sum of A_e over monochromatic edges among colored vertices) satisfying

$$M \preceq \frac{44}{r} I_H.$$

Among the r color classes, one contains at least T/r colored vertices; call that class S . Let $F := E(S, S)$ be the edge set inside S . Every edge in F is monochromatic (with that chosen color), hence

$$\sum_{e \in F} A_e \preceq M \preceq \frac{44}{r} I_H.$$

Since $r \geq 50/\varepsilon$, we have $(44/r) \leq (44/50)\varepsilon < \varepsilon$, and therefore

$$\sum_{e \in F} A_e \preceq \varepsilon I_H. \quad (125)$$

Conjugation back to Laplacians. By definition $A_e = L^{\dagger/2} L_e L^{\dagger/2}$, so summing (125) over $e \in F$ gives

$$L^{\dagger/2} L_F L^{\dagger/2} \preceq \varepsilon I_H. \quad (126)$$

Let $h \in H$ and set $z := L^{1/2}h$ (where $L^{1/2}$ denotes the PSD square root of L on H). Then $h = L^{\dagger/2}z$ and

$$h^\top L_F h = z^\top (L^{\dagger/2} L_F L^{\dagger/2}) z \leq \varepsilon z^\top z = \varepsilon h^\top L h.$$

Thus $L_F \preceq \varepsilon L$ on H . Since $\ker(L) \subseteq \ker(L_F)$, the same inequality holds on all of \mathbb{R}^V . Because $L_S = L_F$, we conclude

$$L_S \preceq \varepsilon L,$$

so S is ε -light.

Cardinality bound. If $n = 1$, take $S = V$. If $n \geq 2$, then $T = \lfloor n/2 \rfloor \geq n/4$, hence

$$|S| \geq \frac{T}{r} \geq \frac{n}{4r}.$$

From (124), $r \leq 50/\varepsilon + 1 \leq 51/\varepsilon$, so

$$|S| \geq \frac{n}{4} \cdot \frac{\varepsilon}{51} = \frac{\varepsilon n}{204}.$$

Therefore Theorem B.9 holds with $c = 1/204$.

Finally, for $\varepsilon = 0$ the conclusion is trivial: any S satisfies $|S| \geq 0$, and choosing (for instance) a single vertex gives $L_S = 0$ and hence $0 \cdot L - L_S \succeq 0$.

B.2 Large ε -light induced subgraphs

B.2.1 Problem

Let $G = (V, E)$ be a finite undirected graph (possibly with positive edge weights). For a subset $S \subseteq V$, define the induced-edge set

$$E(S, S) = \{ \{u, v\} \in E : u \in S, v \in S \},$$

and define the graph

$$G_S = (V, E(S, S)),$$

i.e. the graph on the same vertex set but retaining only edges with both endpoints in S .

Let L denote the (weighted) Laplacian matrix of G and L_S the Laplacian of G_S (so L_S is the Laplacian formed from the edges in $E(S, S)$, still viewed as a $|V| \times |V|$ matrix indexed by V). For $\varepsilon \in (0, 1]$, call S ε -light if the matrix

$$\varepsilon L - L_S$$

is positive semidefinite.

Question. Does there exist a universal constant $c > 0$ such that for every graph $G = (V, E)$ and every $\varepsilon \in (0, 1]$, there exists an ε -light subset $S \subseteq V$ with

$$|S| \geq c\varepsilon|V|?$$

B.2.2 Solution

We answer the question affirmatively, with an absolute constant $c > 0$.

Preliminaries and equivalent formulations Throughout, all matrices are real symmetric, and for symmetric matrices X, Y we write $X \preceq Y$ to mean $Y - X$ is positive semidefinite (Loewner order).

Recall that for any Laplacian L and any vector $x \in \mathbb{R}^V$,

$$x^\top Lx = \sum_{\{i,j\} \in E} w_{ij} (x_i - x_j)^2,$$

where $w_{ij} > 0$ are the (possibly unit) edge weights. In particular, $L \succeq 0$ and $L\mathbf{1} = 0$, where $\mathbf{1}$ denotes the all-ones vector.

The condition that S is ε -light is equivalent to each of the following:

$$\begin{aligned} \varepsilon L - L_S &\succeq 0 \\ \iff L_S &\preceq \varepsilon L \\ \iff x^\top L_S x &\leq \varepsilon x^\top Lx \quad \forall x \in \mathbb{R}^V. \end{aligned}$$

Indeed, the last equivalence is the standard characterization of Loewner order by quadratic forms.

Reduction to connected components We first prove the claim when G is connected and $|V| \geq 2$. The general case then follows by applying the construction in each connected component and taking unions.

Indeed, if G has connected components $G^{(1)}, \dots, G^{(k)}$ on vertex sets $V^{(1)}, \dots, V^{(k)}$, then (after a vertex ordering) L is block-diagonal with blocks $L^{(t)}$. If for each component we find $S^{(t)} \subseteq V^{(t)}$ such that

$$L_{S^{(t)}} \preceq \varepsilon L^{(t)},$$

then setting $S := \bigcup_t S^{(t)}$ yields L_S block-diagonal with blocks $L_{S^{(t)}}$, and thus

$$L_S \preceq \varepsilon L.$$

Moreover, the sizes add: $|S| = \sum_t |S^{(t)}|$.

Thus it suffices to work componentwise. From now on, assume G is connected with $n := |V| \geq 2$.

Normalized edge matrices on the mean-zero subspace Let

$$H = \left\{ x \in \mathbb{R}^V : \sum_{i \in V} x_i = 0 \right\}.$$

Since G is connected, L is positive definite on H , hence invertible there. We write L^{-1} and $L^{-1/2}$ for the inverse and inverse square root of L as operators on H .

For each edge $e = \{i, j\} \in E$ define the rank-one Laplacian contribution

$$L_e = w_{ij}(e_i - e_j)(e_i - e_j)^\top,$$

so that $L = \sum_{e \in E} L_e$. Define the normalized edge matrices on H by

$$A_e = L^{-1/2} L_e L^{-1/2} \quad (e \in E), \quad (127)$$

viewed as operators on H .

Then each $A_e \succeq 0$ and has rank 1 on H , and moreover

$$\sum_{e \in E} A_e = L^{-1/2} \left(\sum_{e \in E} L_e \right) L^{-1/2} = L^{-1/2} L L^{-1/2} = I_H, \quad (128)$$

where I_H denotes the identity operator on H .

For each vertex $v \in V$, define

$$B_v = \sum_{e \ni v} A_e, \quad (129)$$

the sum of normalized edge matrices incident to v . Since every edge has two endpoints,

$$\sum_{v \in V} B_v = 2 \sum_{e \in E} A_e = 2I_H. \quad (130)$$

Let $d := \dim H = n - 1$.

Taking traces in (130) yields $\sum_v \operatorname{tr}(B_v) = 2d$. Hence the set

$$T = \{v \in V : \operatorname{tr}(B_v) \leq 4\} \quad (131)$$

satisfies $|T| \geq n/2$, since otherwise $\sum_v \operatorname{tr}(B_v) > 4 \cdot (n/2) = 2n$, contradicting $2d = 2n - 2 < 2n$.

In particular,

$$m := |T| \geq \frac{n}{2} \geq \frac{d}{2}. \quad (132)$$

A partial coloring lemma via a moving barrier We now prove the key lemma controlling the sum of A_e over monochromatic edges of a partial coloring.

Lemma B.5 (Partial coloring with spectral control). *Let H be a real inner-product space of dimension $d \geq 1$. Let $\{A_e\}_{e \in \mathcal{E}}$ be a finite family of positive semidefinite operators on H satisfying*

$$\sum_{e \in \mathcal{E}} A_e = I_H. \quad (133)$$

Let T be a finite set of “vertices” of size $m := |T|$ and suppose that to each $v \in T$ there is associated an operator

$$B_v = \sum_{e \in \mathcal{E}(v)} A_e,$$

where $\mathcal{E}(v) \subseteq \mathcal{E}$ are some (possibly overlapping) subfamilies, such that

$$\sum_{v \in T} B_v \preceq 2I_H, \quad (134)$$

and for all $v \in T$,

$$\text{tr}(B_v) \leq \tau \quad (135)$$

for some $\tau \geq 1$. Assume moreover that $m \geq d/2$.

Then there exists an absolute constant $C_0 > 0$ with the following property: for every integer $r \geq 1$, one can choose a subset $T' \subseteq T$ with $|T'| \geq m/2$ and a map $\text{col} : T' \rightarrow \{1, \dots, r\}$ (a partial r -coloring) such that the operator

$$M = \sum_{\substack{e \text{ joins } u, v \in T' \\ \text{col}(u) = \text{col}(v)}} A_e \quad (136)$$

satisfies

$$M \preceq \frac{C_0 \tau}{r} I_H. \quad (137)$$

Proof. Fix $r \geq 1$. We construct the partial coloring iteratively, coloring one vertex at a time and ensuring that a suitable spectral barrier is respected.

Notation for increments. Suppose at some stage we have already colored a set $U \subseteq T$ (with colors in $\{1, \dots, r\}$), and let $R := T \setminus U$ denote the set of uncolored vertices. Let M denote the current operator sum (136) determined by already colored vertices.

For a vertex $v \in R$ and color $j \in \{1, \dots, r\}$, define the prospective increment

$$X_j(v) = \sum_{\substack{u \in U \\ \text{col}(u) = j}} A_{uv}, \quad (138)$$

where A_{uv} denotes A_e if there is an “edge” e joining u and v in the underlying structure, and $A_{uv} := 0$ if there is no such e . Coloring v with color j adds exactly $X_j(v)$ to M .

By construction, for fixed v we have

$$\sum_{j=1}^r X_j(v) \preceq B_v, \quad (139)$$

since the left-hand side sums A_{uv} over those edges from v to already colored vertices, a subfamily of $\mathcal{E}(v)$, and all operators are positive semidefinite.

Barrier potential. Given $u > 0$ and an operator $M \prec uI_H$, define

$$\Phi_u(M) = \text{tr}((uI_H - M)^{-1}). \quad (140)$$

We maintain an upper barrier u and ensure that as we update M and increment u , the potential does not increase.

Set the initial barrier to be

$$u_0 = \frac{1024 \tau}{r}, \quad (141)$$

and start with $M_0 = 0$ and $U = \emptyset$.

At any stage with uncolored set R , define the barrier increment

$$\delta(R) = \frac{80}{r|R|}. \quad (142)$$

A resolvent estimate. Let $M \prec uI_H$ and set $u' = u + \delta$ with $\delta > 0$, and put

$$A = u'I_H - M.$$

Let $X \succeq 0$ satisfy $\|A^{-1/2}XA^{-1/2}\| < 1$. Then

$$\mathrm{tr}(A - X)^{-1} \leq \mathrm{tr}(A^{-1}) + \frac{\mathrm{tr}(A^{-2}X)}{1 - \|A^{-1/2}XA^{-1/2}\|}. \quad (143)$$

Moreover, since for each eigenvalue λ of M we have

$$\frac{1}{u - \lambda} - \frac{1}{u' - \lambda} = \frac{\delta}{(u - \lambda)(u' - \lambda)} \geq \frac{\delta}{(u' - \lambda)^2},$$

we also have

$$\Phi_u(M) - \mathrm{tr}(A^{-1}) \geq \delta \mathrm{tr}(A^{-2}). \quad (144)$$

Estimate (143) follows by writing $Y = A^{-1/2}XA^{-1/2}$ and observing

$$(A - X)^{-1} - A^{-1} = A^{-1/2}((I - Y)^{-1}Y)A^{-1/2},$$

and using $(I - Y)^{-1}Y \preceq (1 - \|Y\|)^{-1}Y$ for $0 \preceq Y \prec I$ and taking traces.

Choosing a good vertex and color. Suppose at some stage the current barrier is u and current operator is $M \prec uI_H$. Let R be the uncolored set and define $A = (u + \delta(R))I_H - M$. Let

$$S_1 := \mathrm{tr}(A^{-1}), \quad S_2 := \mathrm{tr}(A^{-2}).$$

We claim there exists a vertex $v \in R$ such that both

$$\mathrm{tr}(A^{-2}B_v) \leq \frac{8}{|R|}S_2, \quad \mathrm{tr}(A^{-1}B_v) \leq \frac{8}{|R|}S_1 \quad (145)$$

hold. Indeed, by (134) applied to the current $R \subseteq T$ (note that $B_v \succeq 0$),

$$\sum_{w \in R} B_w \preceq \sum_{w \in T} B_w \preceq 2I_H,$$

and thus

$$\frac{1}{|R|} \sum_{w \in R} \mathrm{tr}(A^{-k}B_w) = \frac{1}{|R|} \mathrm{tr} \left(A^{-k} \sum_{w \in R} B_w \right) \leq \frac{2}{|R|} \mathrm{tr}(A^{-k}) = \frac{2}{|R|} S_k, \quad k \in \{1, 2\}.$$

By Markov's inequality, fewer than $|R|/4$ vertices violate each inequality in (145), so some $v \in R$ satisfies both.

Fix such a $v \in R$. Using (139), we have for $k \in \{1, 2\}$

$$\sum_{j=1}^r \mathrm{tr}(A^{-k}X_j(v)) = \mathrm{tr} \left(A^{-k} \sum_{j=1}^r X_j(v) \right) \leq \mathrm{tr}(A^{-k}B_v).$$

Hence the average over j is at most $\frac{1}{r}\text{tr}(A^{-k}B_v)$. By (145), that average is at most $\frac{8}{r|R|}S_k$. Again by Markov, at least half of the colors $j \in \{1, \dots, r\}$ satisfy simultaneously

$$\text{tr}(A^{-2}X_j(v)) \leq \frac{32}{r|R|}S_2, \quad \text{tr}(A^{-1}X_j(v)) \leq \frac{32}{r|R|}S_1. \quad (146)$$

Choose such a color j and set $X := X_j(v)$.

Verifying the norm condition. We next ensure that the condition $\|A^{-1/2}XA^{-1/2}\| < 1$ holds automatically with room to spare. We will maintain inductively that the potential never increases:

$$\Phi_u(M) \leq \Phi_{u_0}(0) = \text{tr}(u_0^{-1}I_H) = \frac{d}{u_0}.$$

In particular, at the current step we have

$$S_1 = \text{tr}(A^{-1}) = \Phi_{u+\delta(R)}(M) \leq \Phi_u(M) \leq \frac{d}{u_0}.$$

Using (146) with $k = 1$,

$$\text{tr}(A^{-1}X) \leq \frac{32}{r|R|}S_1 \leq \frac{32}{r|R|} \cdot \frac{d}{u_0}.$$

Using the choice (141) of u_0 , this becomes

$$\text{tr}(A^{-1}X) \leq \frac{32}{r|R|} \cdot \frac{d}{1024\tau/r} = \frac{32d}{1024\tau|R|}. \quad (147)$$

We will stop the procedure once $|R| = m/2$; during the construction, therefore, $|R| \geq m/2$. Since $m \geq d/2$ by assumption, it follows that $d/|R| \leq 4$ during all steps up to this stopping time. With $\tau \geq 1$, (147) gives

$$\text{tr}(A^{-1}X) \leq \frac{32 \cdot 4}{1024} \leq \frac{1}{8}.$$

Since $A^{-1/2}XA^{-1/2} \succeq 0$, we have

$$\|A^{-1/2}XA^{-1/2}\| \leq \text{tr}(A^{-1/2}XA^{-1/2}) = \text{tr}(A^{-1}X) \leq \frac{1}{8},$$

and thus $\|A^{-1/2}XA^{-1/2}\| < 1$ as required.

Potential non-increase. Apply (143) with this A and X , using the bound $\|A^{-1/2}XA^{-1/2}\| \leq 1/8$:

$$\Phi_{u+\delta(R)}(M+X) = \text{tr}(A-X)^{-1} \leq \text{tr}(A^{-1}) + \frac{\text{tr}(A^{-2}X)}{1-1/8} \leq \text{tr}(A^{-1}) + \frac{8}{7}\text{tr}(A^{-2}X).$$

Using (146) with $k = 2$, we obtain

$$\Phi_{u+\delta(R)}(M+X) \leq \text{tr}(A^{-1}) + \frac{8}{7} \cdot \frac{32}{r|R|}S_2 \leq \text{tr}(A^{-1}) + \frac{64}{r|R|}S_2.$$

On the other hand, (144) yields

$$\Phi_u(M) \geq \text{tr}(A^{-1}) + \delta(R)S_2 = \text{tr}(A^{-1}) + \frac{80}{r|R|}S_2.$$

Since $80 > 64$, combining the last two displays gives

$$\Phi_{u+\delta(R)}(M+X) \leq \Phi_u(M). \quad (148)$$

In particular, $M + X \prec (u + \delta(R))I_H$, so the new barrier remains valid and the iterative step can continue.

Stopping time and final barrier. We repeat the step until $m/2$ vertices have been colored, i.e. until $|R| = m/2$. Starting at barrier u_0 and adding $\delta(R) = 80/(r|R|)$ at each step, the final barrier u_{final} satisfies

$$\begin{aligned} u_{\text{final}} &\leq u_0 + \sum_{s=m/2+1}^m \frac{80}{rs} \\ &\leq \frac{1024\tau}{r} + \frac{80}{r} \sum_{s=m/2+1}^m \frac{1}{s} \\ &\leq \frac{1024\tau}{r} + \frac{80}{r} \log 2 \\ &\leq \frac{1024\tau + 140}{r} \leq \frac{1200\tau}{r}. \end{aligned}$$

By construction, at the end we have an operator M (as in (136)) satisfying $M \prec u_{\text{final}}I_H$. Thus (137) holds with $C_0 = 1200$. \square

Extracting an ε -light subset We now prove the main theorem.

Theorem B.6. *There exists an absolute constant $c > 0$ such that for every finite weighted graph $G = (V, E)$ and every $\varepsilon \in (0, 1]$, there exists $S \subseteq V$ with*

$$|S| \geq c\varepsilon|V|$$

and

$$L_S \preceq \varepsilon L,$$

i.e. S is ε -light.

Proof. By the connected-components reduction above, it suffices to assume G is connected with $n = |V| \geq 2$.

Let H be the mean-zero subspace and define A_e and B_v by (127)–(129). Let T be the set (131). Then $m := |T| \geq d/2$ by (132), and for $v \in T$ we have $\text{tr}(B_v) \leq 4$. Moreover, from (130),

$$\sum_{v \in T} B_v \preceq \sum_{v \in V} B_v = 2I_H.$$

Thus the hypotheses of Lemma B.5 hold with $\tau = 4$.

Choose an integer

$$r = \left\lceil \frac{C}{\varepsilon} \right\rceil,$$

where $C \geq 1$ is a sufficiently large absolute constant, chosen below.

Apply Lemma B.5 to obtain a subset $T' \subseteq T$ with $|T'| \geq m/2$ and a partial coloring $\text{col} : T' \rightarrow \{1, \dots, r\}$ such that the corresponding operator M defined by (136) satisfies

$$M \preceq \frac{C_0\tau}{r} I_H = \frac{4C_0}{r} I_H. \quad (149)$$

For each color $p \in \{1, \dots, r\}$, let $S_p := \{v \in T' : \text{col}(v) = p\}$. Choose p maximizing $|S_p|$, and set $S := S_p$. Since the colored set has size $|T'| \geq m/2 \geq n/4$, we have

$$|S| \geq \frac{|T'|}{r} \geq \frac{n}{4r}. \quad (150)$$

Let M_p denote the portion of M coming from monochromatic edges within color class p :

$$M_p = \sum_{\{u,v\} \in E(S,S)} A_{uv}.$$

Then $M_p \succeq 0$ and $M_p \preceq M$, so by (149),

$$M_p \preceq \frac{4C_0}{r} I_H. \quad (151)$$

On the other hand, by definition of A_e and the fact that

$$L_S = \sum_{e \in E(S,S)} L_e,$$

we have, as operators on H ,

$$L^{-1/2} L_S L^{-1/2} = \sum_{e \in E(S,S)} L^{-1/2} L_e L^{-1/2} = \sum_{e \in E(S,S)} A_e = M_p. \quad (152)$$

Combining (152) and (151) gives

$$L^{-1/2} L_S L^{-1/2} \preceq \frac{4C_0}{r} I_H.$$

Congruencing by $L^{1/2}$ (valid on H) yields

$$L_S \preceq \frac{4C_0}{r} L \quad \text{on } H. \quad (153)$$

Equivalently, for all $x \in H$,

$$x^\top L_S x \leq \frac{4C_0}{r} x^\top L x.$$

To extend from H to all of \mathbb{R}^V , write any $x \in \mathbb{R}^V$ uniquely as $x = y + \alpha \mathbf{1}$ with $y \in H$ and $\alpha \in \mathbb{R}$. Since both L and L_S annihilate constants, $L\mathbf{1} = 0$ and $L_S\mathbf{1} = 0$, we have $x^\top L x = y^\top L y$ and $x^\top L_S x = y^\top L_S y$. Thus (153) implies the same inequality for all $x \in \mathbb{R}^V$:

$$x^\top L_S x \leq \frac{4C_0}{r} x^\top L x \quad \forall x \in \mathbb{R}^V. \quad (154)$$

Therefore $L_S \preceq (4C_0/r)L$ as matrices on \mathbb{R}^V .

Choose C so that $C \geq 4C_0$. Then $r = \lceil C/\varepsilon \rceil$ satisfies

$$\frac{4C_0}{r} \leq \frac{4C_0}{C/\varepsilon} = \varepsilon,$$

and hence (154) implies $L_S \preceq \varepsilon L$, i.e. S is ε -light.

Finally, from (150) and $r = \lceil C/\varepsilon \rceil \leq C/\varepsilon + 1$,

$$|S| \geq \frac{n}{4r} \geq \frac{n}{4(C/\varepsilon + 1)} = \frac{\varepsilon}{4(C + \varepsilon)} n \geq \frac{\varepsilon}{4(C + 1)} n,$$

since $\varepsilon \leq 1$. Thus the theorem holds with

$$c = \frac{1}{4(C + 1)},$$

where C is any absolute constant at least $4C_0$ (for instance, taking $C = 4800$ suffices when $C_0 = 1200$ from Lemma B.5).

This completes the connected case and hence, by the componentwise reduction, the general case. \square

Remarks

- The proof gives an explicit (though very small) absolute constant $c > 0$; the constants were not optimized.
- The barrier method used in Lemma [B.5](#) is closely related to the approach introduced by Batson–Spielman–Srivastava in their work on sparsification; the argument here is a vertex-coloring variant.

References

- [1] J. Batson, D. Spielman, and N. Srivastava, *Twice-Ramanujan sparsifiers*, SIAM Journal on Computing **41** (2012), no. 6, 1704–1721.

B.3 Existence of large ε -light vertex sets

B.3.1 Solution overview

We prove a stronger quantitative statement: for every $\varepsilon \in (0, 1]$ and every graph $G = (V, E)$ with $n := |V|$, there exists an ε -light set S with

$$|S| \geq \frac{\varepsilon n}{400}.$$

The proof is deterministic and uses a one-sided matrix barrier argument in the spirit of Batson–Spielman–Srivastava [1].

B.3.2 Preliminaries: edge Laplacians and a normalization

Fix an arbitrary orientation of each edge $e \in E$. Let $b_e \in \mathbb{R}^V$ be the corresponding incidence vector: if $e = uv$ is oriented from u to v , then $(b_e)_u = 1$, $(b_e)_v = -1$, and all other coordinates are 0. Define the rank-one *edge Laplacian* by

$$L_e := b_e b_e^\top.$$

Then the Laplacian of G satisfies

$$L = \sum_{e \in E} L_e.$$

Let L^+ denote the Moore–Penrose pseudoinverse of L . Let

$$H := \text{range}(L) \subseteq \mathbb{R}^V,$$

equivalently, H is the subspace of vectors whose coordinate-sum is 0 on each connected component of G . Note that $b_e \in H$ for every $e \in E$. Let $d := \dim(H)$; then $d \leq n$.

Define, for each $e \in E$,

$$A_e := (L^+)^{1/2} L_e (L^+)^{1/2},$$

viewed as a self-adjoint operator on H . Since $(L^+)^{1/2} L (L^+)^{1/2}$ is the orthogonal projector onto $\text{range}(L) = H$, we have the identity on H :

$$\sum_{e \in E} A_e = I_H. \tag{155}$$

Conversely, for any set of edges $F \subseteq E$, define

$$M_F := \sum_{e \in F} A_e \quad (\text{an operator on } H).$$

Then

$$\sum_{e \in F} L_e = L^{1/2} M_F L^{1/2}. \tag{156}$$

Indeed, since $L^{1/2}(L^+)^{1/2}$ is the orthogonal projector P_H onto H , we have

$$L^{1/2} M_F L^{1/2} = L^{1/2} (L^+)^{1/2} \left(\sum_{e \in F} L_e \right) (L^+)^{1/2} L^{1/2} = P_H \left(\sum_{e \in F} L_e \right) P_H.$$

Because each $L_e = b_e b_e^\top$ has range contained in H (as $b_e \in H$), we get $P_H L_e P_H = L_e$, whence (156).

In particular, for a vertex set $S \subseteq V$,

$$L_S = \sum_{e \in E(S,S)} L_e = L^{1/2} M_{E(S,S)} L^{1/2}.$$

Therefore, to prove that S is ε -light, it suffices to show

$$M_{E(S,S)} \preceq \varepsilon I_H, \quad (157)$$

because then for all $x \in \mathbb{R}^V$,

$$x^\top L_S x = (L^{1/2}x)^\top M_{E(S,S)}(L^{1/2}x) \leq \varepsilon \|L^{1/2}x\|^2 = \varepsilon x^\top Lx.$$

B.3.3 A one-sided matrix barrier lemma

For a self-adjoint operator M on H and a scalar u with $M \prec uI_H$, define the potential

$$\Phi^u(M) := \operatorname{tr}((uI_H - M)^{-1}).$$

Lemma B.7 (One-sided barrier). *Let M be self-adjoint on H with $M \prec uI_H$, and let $\Delta \succeq 0$. Let $u' = u + \delta$ with $\delta > 0$. Assume that*

$$\operatorname{tr}((u'I_H - M)^{-2}\Delta) \leq \frac{1}{4}(\Phi^u(M) - \Phi^{u'}(M)) \quad (158)$$

and

$$\operatorname{tr}((u'I_H - M)^{-1}\Delta) \leq \frac{1}{4}. \quad (159)$$

Then $M + \Delta \prec u'I_H$ and

$$\Phi^{u'}(M + \Delta) \leq \Phi^u(M).$$

Proof. Let $X := u'I_H - M$, so $X \succ 0$, and set $R := X^{-1}$. Define $Y := R^{1/2}\Delta R^{1/2} \succeq 0$. By (159),

$$\operatorname{tr}(Y) = \operatorname{tr}(R\Delta) \leq \frac{1}{4}.$$

Since $Y \succeq 0$, its largest eigenvalue satisfies $\lambda_{\max}(Y) \leq \operatorname{tr}(Y) \leq 1/4$, hence $I_H - Y \succ 0$ and therefore $X - \Delta \succ 0$. Thus $u'I_H - (M + \Delta) = X - \Delta \succ 0$, i.e. $M + \Delta \prec u'I_H$.

Moreover,

$$(X - \Delta)^{-1} = R^{1/2}(I_H - Y)^{-1}R^{1/2}.$$

Hence

$$\begin{aligned} \Phi^{u'}(M + \Delta) - \Phi^{u'}(M) &= \operatorname{tr}((X - \Delta)^{-1} - X^{-1}) \\ &= \operatorname{tr}\left(R^{1/2}((I_H - Y)^{-1} - I_H)R^{1/2}\right) \\ &= \operatorname{tr}(R^{1/2}(I_H - Y)^{-1}YR^{1/2}). \end{aligned}$$

For $t \in [0, 1/4]$ we have $\frac{t}{1-t} \leq \frac{4}{3}t$, and by functional calculus this implies

$$(I_H - Y)^{-1}Y \preceq \frac{4}{3}Y.$$

Therefore,

$$\Phi^{u'}(M + \Delta) - \Phi^{u'}(M) \leq \frac{4}{3} \operatorname{tr}(R^{1/2}YR^{1/2}) = \frac{4}{3} \operatorname{tr}(R\Delta R) = \frac{4}{3} \operatorname{tr}(R^2\Delta).$$

On the other hand, diagonalizing M shows the resolvent identity

$$\Phi^u(M) - \Phi^{u'}(M) = \delta \operatorname{tr}((uI_H - M)^{-1}(u'I_H - M)^{-1}). \quad (160)$$

Since $(uI_H - M)^{-1} \succeq (u'I_H - M)^{-1} = R$ and these operators commute (both are functions of M), we have

$$\operatorname{tr}((uI_H - M)^{-1}(u'I_H - M)^{-1}) \geq \operatorname{tr}(R^2),$$

and consequently

$$\Phi^u(M) - \Phi^{u'}(M) \geq \delta \operatorname{tr}(R^2).$$

Combining this with (158) yields

$$\operatorname{tr}(R^2 \Delta) \leq \frac{1}{4}(\Phi^u(M) - \Phi^{u'}(M)).$$

Thus,

$$\Phi^{u'}(M + \Delta) - \Phi^{u'}(M) \leq \frac{4}{3} \cdot \frac{1}{4}(\Phi^u(M) - \Phi^{u'}(M)) \leq \Phi^u(M) - \Phi^{u'}(M),$$

which implies $\Phi^{u'}(M + \Delta) \leq \Phi^u(M)$. □

B.3.4 A deterministic colouring lemma

We use the trace inner product $\langle A, B \rangle := \operatorname{tr}(AB)$ on operators on H .

Lemma B.8 (Colouring lemma). *Assume that the matrices $\{A_e\}_{e \in E}$ satisfy (155). Let $n := |V|$ and $d := \dim(H) \leq n$. Then for every integer $r \geq 1$, one can colour at least $\lfloor n/2 \rfloor$ vertices with r colours so that, if*

$$M := \sum_{\substack{uv \in E \\ u, v \text{ coloured} \\ \text{colour}(u) = \text{colour}(v)}} A_{uv},$$

then

$$M \preceq \frac{32(1 + \log 2)}{r} I_H. \quad (161)$$

Proof. If $E = \emptyset$, the statement is trivial. Assume $E \neq \emptyset$.

We build the colouring iteratively. At some stage, let C be the set of coloured vertices, $U := V \setminus C$ the uncoloured vertices, and $s := |U|$. Let C_1, \dots, C_r be the colour classes in C . For $v \in U$ and $i \in \{1, \dots, r\}$ define

$$\Delta_i(v) := \sum_{u \in C_i} A_{uv},$$

where the summand is 0 if $uv \notin E$. If the current monochromatic-edge sum is M_k , then colouring v with colour i updates it to $M_k + \Delta_i(v)$.

We maintain inductively a barrier $u_k > 0$ such that

$$M_k \prec u_k I_H \quad \text{and} \quad \Phi^{u_k}(M_k) \leq \frac{d}{u_k}, \quad (162)$$

where we start with

$$M_0 = 0, \quad u_0 := \frac{32}{r}.$$

Clearly $M_0 \prec u_0 I_H$ and $\Phi^{u_0}(M_0) = \operatorname{tr}(u_0^{-1} I_H) = d/u_0$.

We will colour vertices until $m := \lfloor n/2 \rfloor$ vertices have been coloured. During this process we always have $s \geq n/2$.

Inductive step. Assume (162) holds at step k , with current uncoloured set size $s \geq n/2$. Set

$$\delta := \frac{32}{rs}, \quad u' := u_k + \delta, \quad R := (u' I_H - M_k)^{-1}.$$

Define the crossing-edge operator

$$A_{\text{cr}} := \sum_{v \in U} \sum_{i=1}^r \Delta_i(v).$$

This is the sum of A_e over edges e with one endpoint in C and one endpoint in U . In particular, $A_{\text{cr}} \succeq 0$ and, since it is a partial sum of the A_e and $\sum_e A_e = I_H$, we have

$$A_{\text{cr}} \preceq I_H. \quad (163)$$

Consider all rs pairs (v, i) with $v \in U$ and $1 \leq i \leq r$. We bound the averages of the two quantities appearing in Lemma B.7.

First, using (163),

$$\frac{1}{rs} \sum_{v \in U} \sum_{i=1}^r \langle R^2, \Delta_i(v) \rangle = \frac{1}{rs} \langle R^2, A_{\text{cr}} \rangle \leq \frac{1}{rs} \text{tr}(R^2).$$

Moreover, diagonalizing M_k yields the identity

$$\Phi^{u_k}(M_k) - \Phi^{u'}(M_k) = \delta \text{tr}((u_k I_H - M_k)^{-1} (u' I_H - M_k)^{-1}),$$

and since $(u_k I_H - M_k)^{-1} \succeq (u' I_H - M_k)^{-1} = R$ and these commute, we have

$$\Phi^{u_k}(M_k) - \Phi^{u'}(M_k) \geq \delta \text{tr}(R^2).$$

Therefore,

$$\frac{1}{rs} \sum_{v,i} \frac{\langle R^2, \Delta_i(v) \rangle}{\Phi^{u_k}(M_k) - \Phi^{u'}(M_k)} \leq \frac{1}{rs} \cdot \frac{\text{tr}(R^2)}{\delta \text{tr}(R^2)} = \frac{1}{rs\delta} = \frac{1}{32}. \quad (164)$$

Second, similarly,

$$\frac{1}{rs} \sum_{v \in U} \sum_{i=1}^r \langle R, \Delta_i(v) \rangle = \frac{1}{rs} \langle R, A_{\text{cr}} \rangle \leq \frac{1}{rs} \text{tr}(R).$$

Since $\text{tr}(R) = \Phi^{u'}(M_k) \leq \Phi^{u_k}(M_k) \leq d/u_0$ by the inductive hypothesis, we get

$$\frac{1}{rs} \sum_{v,i} \langle R, \Delta_i(v) \rangle \leq \frac{1}{rs} \cdot \frac{d}{u_0} = \frac{d}{32s}. \quad (165)$$

Because $s \geq n/2$ and $d \leq n$, the right-hand side is at most $1/16$.

Combining (164) and (165), the average over all pairs (v, i) of

$$\frac{\langle R^2, \Delta_i(v) \rangle}{\Phi^{u_k}(M_k) - \Phi^{u'}(M_k)} + \langle R, \Delta_i(v) \rangle$$

is at most $1/32 + 1/16 = 3/32$. Hence there exists at least one pair (v, i) for which this sum is $\leq 3/32$. Since both summands are nonnegative, each is $\leq 3/32 < 1/4$. Thus, for this pair (v, i) , both barrier conditions (158)–(159) of Lemma B.7 hold (with $M := M_k$, $u := u_k$, $\Delta := \Delta_i(v)$). Applying Lemma B.7 gives

$$M_k + \Delta_i(v) \prec u' I_H \quad \text{and} \quad \Phi^{u'}(M_k + \Delta_i(v)) \leq \Phi^{u_k}(M_k).$$

We colour v with colour i , define $M_{k+1} := M_k + \Delta_i(v)$ and $u_{k+1} := u'$, and the invariant (162) holds at step $k+1$.

Summing the barrier increments. After $m = \lfloor n/2 \rfloor$ steps, we have

$$u_m = u_0 + \sum_{k=0}^{m-1} \frac{32}{rs_k},$$

where s_k is the number of uncoloured vertices at step k . Since s_k decreases by 1 each step and starts at n , we have $s_k \geq n - k$ and thus

$$u_m \leq \frac{32}{r} + \frac{32}{r} \sum_{j=n-m+1}^n \frac{1}{j} \leq \frac{32}{r} + \frac{32}{r} \log \left(\frac{n}{n-m} \right) \leq \frac{32(1 + \log 2)}{r},$$

because $n - m \geq n/2$. Finally, $M_m \prec u_m I_H$ implies (161) (with $M := M_m$), completing the proof. \square

B.3.5 Proof of the main theorem

Theorem B.9. *There exists a universal constant $c > 0$ such that for every graph $G = (V, E)$ with $n := |V|$ and every $\varepsilon \in (0, 1]$, there is an ε -light set $S \subseteq V$ with*

$$|S| \geq c\varepsilon n.$$

In particular, one may take $c = 1/400$.

Proof. If $E = \emptyset$, then $L = 0$ and $L_S = 0$ for every S , so $S = V$ is ε -light and $|S| = n$, which satisfies the claim. Assume henceforth that $E \neq \emptyset$ and form the normalized matrices A_e on H as in (155).

Set

$$r := \left\lceil \frac{100}{\varepsilon} \right\rceil.$$

Apply Lemma B.8 to obtain a partial colouring of at least $\lfloor n/2 \rfloor$ vertices using r colours, and let $T \subseteq V$ be the set of coloured vertices. Let T_1, \dots, T_r be the colour classes. For each j , define

$$M_j := \sum_{uv \in E(T_j, T_j)} A_{uv}.$$

The matrix M from Lemma B.8 is $M = \sum_{j=1}^r M_j$. Since $M_j \succeq 0$ for each j , we have $M_j \preceq M$, hence by (161),

$$M_j \preceq \frac{32(1 + \log 2)}{r} I_H \preceq \varepsilon I_H,$$

where the final inequality holds by the choice of r (the constant 100 comfortably dominates $32(1 + \log 2)$).

Let S be a largest colour class among the T_j 's. Then

$$|S| \geq \frac{|T|}{r} \geq \frac{\lfloor n/2 \rfloor}{r}.$$

If $n \geq 2$, then $\lfloor n/2 \rfloor \geq n/3$, so

$$|S| \geq \frac{n}{3r} \geq \frac{\varepsilon n}{3(100 + \varepsilon)} \geq \frac{\varepsilon n}{400}.$$

If $n = 1$, the unique vertex forms a set of size $1 \geq (\varepsilon/400)n$.

It remains to check that S is ε -light. Let $F := E(S, S)$. By (156) and the definition of M_F ,

$$L_S = \sum_{e \in F} L_e = L^{1/2} M_F L^{1/2}, \quad \text{where} \quad M_F = \sum_{e \in F} A_e = M_j$$

for the corresponding colour class j . Since $M_F \preceq \varepsilon I_H$, for every $x \in \mathbb{R}^V$ we have

$$x^\top L_S x = (L^{1/2} x)^\top M_F (L^{1/2} x) \leq \varepsilon (L^{1/2} x)^\top (L^{1/2} x) = \varepsilon x^\top L x.$$

Equivalently, $\varepsilon L - L_S \succeq 0$, so S is ε -light. □

References

- [1] J. Batson, D. A. Spielman, and N. Srivastava, *Twice-Ramanujan Sparsifiers*, SIAM Journal on Computing **41** (2012), no. 6, 1704–1721.