Coarse information and finite rank obstructions for flat lattice Hodge forms

Lluis Eriksson

6 November 2026 — consolidated revision candidate

Abstract

On a finite periodic lattice with unit spacing, a flat Hodge form cannot have a positive coercivity constant uniform under refinement when supplemented only by linear observations in a fixed finite-dimensional real space and linear fine-space repairs of uniformly bounded rank. We construct a real physical Fourier band, prove a necessary information budget and a separate bounded-rank sector obstruction, and explain why a classical singular-Schur condition prevents reversing that budget. Constructed derivative repairs and differentiable observations with locally upper-quadratic penalties admit tangent obstructions, including an exact regular nonlinear-sector argument. For nonnegative weights, a subquadratic directional spatial mean forces output-rank divergence. For signed weights, the spatial mean of the pointwise positive part gives a sufficient condition; ordinary signed means do not, as an actual fixed-output counterexample with every mean zero shows. An attributed classical square cancellation delimits the sufficient criterion; a finite-sum extension covering every eligible pair is displayed analytically. A full volume coexact packet also gives a quantitative rank budget under spatial means of the pointwise positive total coefficient on its two occupied field components. A derived integer cutoff yields positive output rank density on every sufficiently large lattice, while exact trace averaging improves a sufficient low frequency constant. Formal endpoints are separately pinned, and additional analytic arguments are identified. These are finite flat statements, with no interacting or continuum Yang–Mills mass-gap claim or new general spectral-principle claim. The integrated continuation also proves a critical inverse-moment sufficient condition for global spectral observations, together with limits on spatial localization. Rectangular means have a general upper bound with exponent $d/4$ and the target exponent $d/2$ for an explicitly restricted scaled-cube class. Two constructive counterexamples exclude a single-global-threshold algorithm and a universal exact cube decomposition, while preserving successful rectangular partitions for those same profiles. The general critical rectangular question remains open. Self-contained local and supplied-cube proofs are included, along with an actual physical fixed-map obstruction and its separately scoped exact finite penalty threshold.

1 Purpose and relation to the existing record

This text consolidates a continuation of the exact-kernel obstruction in [1]. That paper exhibits fields in the kernel of a particular line-integral block map. Changing the map can remove a particular witness. Here the obstruction depends instead on the amount of coarse information: a whole low-energy subspace is too large to be seen by a fixed-dimensional measurement together with a bounded-rank repair. The map need not factor through the original block variable, be local, or obey a norm bound; it is real-linear in the unrestricted-penalty results. The penalty

applied to its output may be nonlinear. Later sections treat differentiable observations with a locally upper-quadratic penalty and exact nonlinear regular sectors. The restricted-sector result is proved separately and does not follow merely from failure of an estimate on the full space.

Fourier diagonalization, rank-nullity and finite-dimensional Schur complements are classical. For the singular Schur complement and pseudoinverse background, see Albert and Carlson–Haynsworth–Markham [3, 4]. Finite-difference Hodge theory has a long history; Dodziuk's work [2] provides historical context, rather than this paper's specific lattice normalization. We make no claim to a new general spectral principle or an exhaustive priority search. The contribution is a concrete, reproducibly checked constraint on proposed repair routes in this programme. This consolidation is intended for the existing paper dossier, rather than an additional deposit for each lemma.

Gauge-fixing and block averaging in renormalization-group constructions require their own domain and scaling analysis. Dimock [5] develops covariant axial gauge for an abelian model and discusses its relation to Balaban's programme. His introduction distinguishes a finite unit lattice from a torus with small spacing, equivalently a large unit-spacing volume. Our unrestricted, unit-spacing flat form is not identified with every constrained fluctuation form in those constructions. This negative diagnostic does not refute their gauge-fixed estimates or establish an interacting bound.

2 Spaces and the actual flat form

Fix integers $d \geq 1$, $n \geq 1$ and $N_c \geq 2$. Set $X_n = (\mathbb{Z}/n\mathbb{Z})^d$ and $m = N_c^2 - 1$. The real one-cochain space is

$$E_n = \{A : X_n \times \{1, \dots, d\} \longrightarrow \mathbb{R}^m\}, \qquad \|A\|^2 = \sum_{x \in X_n} \sum_{i=1}^{d} \|A_i(x)\|^2.$$

The internal coordinates are those of the repository's Euclidean Lie-coordinate model. No identification with a matrix Lie algebra is required here. For positive unordered plaquettes $i < k$, define

$$\begin{aligned} (\mathrm{d}A)_{ik}(x) &= A_k(x + e_i) - A_k(x) - A_i(x + e_k) + A_i(x), \\ (\delta A)(x) &= \sum_i (A_i(x) - A_i(x - e_i)). \end{aligned}$$

The flat operator $K_{0,n}$ is the actual repository Hodge operator, with

$$\langle A, K_{0,n} A \rangle = \|\mathrm{d}A\|^2 + \|\delta A\|^2. \tag{1}$$

The plaquette norm sums each pair once. All differences have unit spacing; there is no hidden n^2 or block-side squared factor.

Choose a direction j and two distinct internal unit coordinates u, v. They exist since $m \geq 3$. For $f : X_n \to \mathbb{C}$, put

$$(E_j f)_i(x) = \begin{cases} (\Re f(x))u + (\Im f(x))v, & i = j, \\ 0, & i \neq j. \end{cases}$$

This is a real-linear injection and $\|E_j f\|^2 = \sum_x |f(x)|^2$. Writing $D_i f(x) = f(x + e_i) - f(x)$, the actual curl and divergence give

$$\langle E_j f, K_{0,n} E_j f \rangle = \sum_{i \neq j} \sum_x |D_i f(x)|^2 + \sum_x |f(x) - f(x - e_j)|^2 = \sum_i \sum_x |D_i f(x)|^2. \tag{2}$$

The last equality is finite translation invariance. The exclusive two plaquette incidences and the divergence term account for every direction exactly once, with no extra factor two.

3 The original block kernel and its repair budget

We retain the principal results of the historical paper [1] before explaining the new information-budget obstruction. The original Appendix A (an even-block square-wave witness with quotient $8/\ell$) remains in the preserved v1; it is not reproduced here. Use positive integers ℓ, N', set $n = \ell N'$, and for a coarse site y put $B_y = \{\ell y + r : 0 \le r_i < \ell\}$. The original line-integral map is

$$(Q_\ell A)_i(y) = \ell^{-d} \sum_{x \in B_y} \sum_{k=0}^{\ell-1} A_i(x + ke_i).$$

For a constant field $A_i(x) = v_i$ and $S = \sum_i \|v_i\|^2 > 0$,

$$\|A\|^2 = \ell^d (N')^d S, \qquad \|\ell^\sigma Q_\ell A\|^2 = \ell^{2\sigma+2} (N')^d S.$$

A uniform Poincare estimate must therefore have $\sigma \ge (d-2)/2$ if it is to survive this sector. Equality alone makes the two powers identical. Larger exponents pass this necessary test too; neither equality nor the inequality is a sufficient full-space condition. This follows by cancelling $(N')^d S$ to obtain $\ell^{d-2-2\sigma} \le C_P$. The real-exponent comparison is analytic here.

For $d \ge 2$ choose distinct directions p, q and the real isometry $\iota(z) = (\Re z)u + (\Im z)v$ from Section 2. For $\ell \ge 2$ define

$$A_i^{(\ell)}(x) = \mathbf{1}_{i=p}\, \iota\big(e^{2\pi \mathrm{i} x_q/\ell}\big), \qquad \lambda_\ell = |e^{2\pi \mathrm{i}/\ell} - 1|^2 = 4\sin^2(\pi/\ell).$$

The character is well defined on $X_{\ell N'}$. The sum over the q-offset in every block is a nontrivial root-of-unity sum and vanishes. Shifts in the line direction p do not change the profile. Hence

$$Q_\ell A^{(\ell)} = 0, \quad \delta A^{(\ell)} = 0, \quad \|A^{(\ell)}\|^2 = (\ell N')^d, \quad \langle A^{(\ell)}, K_{0,n} A^{(\ell)} \rangle = (\ell N')^d \lambda_\ell. \tag{3}$$

Only the plaquette pair containing p, q contributes to the curl energy; its difference has constant squared norm λ_ℓ at every site. This also explains the exact energy in (3), including periodic boundary jumps.

Consequently any scale-wise estimate

$$\|A\|^2 \le C_P(\ell)\big(\langle A, K_{0,n} A\rangle + \Psi_\ell(Q_\ell A)\big), \qquad \Psi_\ell(0) = 0,$$

forces $C_P(\ell) \ge \lambda_\ell^{-1} \ge \ell^2/(4\pi^2)$. No fixed $C_P > 0$ before all block sides can work. This includes $\Psi_\ell(y) = \|T_\ell y\|^2$ for arbitrary linear postconditioners, and all critical or faster scalar rescalings. There is no regularity or sign assumption on the scalar function Ψ_ℓ.

The original spectral strengthening is retained analytically. For $1 \le r \le R \le \lfloor (\ell-1)/2 \rfloor$, replace the profile by $\iota(ce^{2\pi \mathrm{i} r x_q/\ell})$, $c \in \mathbb{C}$. These are mutually orthogonal real two-dimensional spaces in $\ker Q_\ell$. Derivative orthogonality must be checked too: every forward difference multiplies a character by its frequency multiplier, so the same root-of-unity Gram sum annihilates cross terms between distinct frequencies. Thus for $H_\ell = K_{0,n} + (T_\ell Q_\ell)^*(T_\ell Q_\ell)$ the sum of these spaces has dimension $2R$ and maximum Rayleigh quotient at most $4\sin^2(\pi R/\ell) \le 4\pi^2 R^2/\ell^2$. Finite-dimensional min–max gives $\mu_{2R}(H_\ell) \le 4\pi^2 R^2/\ell^2$. Taking $R = \lfloor \ell^\alpha \rfloor$, $0 < \alpha < 1$, for

sufficiently large ℓ, yields a growing cluster below $4\pi^2\ell^{-2(1-\alpha)}$. Eigenvalue language is restricted to this linear self-adjoint operator, not applied to a general nonlinear penalty.

Finally, for any linear changed measurement $\widetilde{Q}_\ell$, linear postconditioner T_ℓ and positive self-adjoint fine term V_ℓ, put

$$\varepsilon_\ell = \frac{\|T_\ell \widetilde{Q}_\ell A^{(\ell)}\|}{\|A^{(\ell)}\|}, \qquad \nu_\ell = \frac{\langle A^{(\ell)}, V_\ell A^{(\ell)}\rangle}{\|A^{(\ell)}\|^2}.$$

Set $\widetilde{H}_\ell = K_{0,n} + (T_\ell\widetilde{Q}_\ell)^*(T_\ell\widetilde{Q}_\ell) + V_\ell$. Expansion gives the exact repaired quotient

$$\frac{\langle A^{(\ell)}, \widetilde{H}_\ell A^{(\ell)}\rangle}{\|A^{(\ell)}\|^2} = \lambda_\ell + \varepsilon_\ell^2 + \nu_\ell.$$

A full-space constant $c > 0$ therefore requires $\varepsilon_\ell^2 + \nu_\ell \geq c - \lambda_\ell$ at every scale, hence $\liminf_{\ell\to\infty}(\varepsilon_\ell^2 + \nu_\ell) \geq c$. On the single witness line, uniform coercivity is equivalent to $\inf_{\ell\geq 2}(\lambda_\ell + \varepsilon_\ell^2 + \nu_\ell) > 0$; this is not full-space sufficiency. In particular if $\widetilde{Q}_\ell = Q_\ell + R_\ell$, $T_\ell = \ell^{(d-2)/2} I$ and $\eta_\ell = \|R_\ell A^{(\ell)}\|/\|A^{(\ell)}\|$, then $\varepsilon_\ell^2 = \ell^{d-2}\eta_\ell^2$. In dimension four and with $V_\ell = 0$, leakage $o(\ell^{-1})$ cannot even repair this line. The liminf, real-exponent, spectral multiplicity and witness-line equivalence are analytic statements. The older exact kernel, norm-form energy, finite budget and dimension-four factorized no-go have their own preserved formal source and reproductions, separately anchored below.

4 A constructed low band

For a nonnegative integer R with $2R < n$, let $\Lambda_R = \{-R, \ldots, R\}^d$ and $B = |\Lambda_R| = (2R+1)^d$. Define

$$\chi_k(x) = \exp(2\pi \mathrm{i}\, k \cdot x/n), \qquad f_a = \sum_{k\in\Lambda_R} a_k \chi_k, \qquad S_R = \{E_j f_a : a \in \mathbb{C}^{\Lambda_R}\} \subset E_n.$$

Lemma 4.1. *The real subspace S_R has dimension $2B$, and for every $A \in S_R$,*

$$\langle A, K_{0,n} A\rangle \leq b_R(n)\|A\|^2, \qquad b_R(n) = \frac{4\pi^2 d R^2}{n^2}. \tag{4}$$

Proof. For $k, \ell \in \Lambda_R$, the product root-of-unity sum is $\sum_x \overline{\chi_k(x)}\chi_\ell(x) = n^d \mathbf{1}_{k=\ell}$. The no-alias condition ensures that coordinate differences, of absolute value at most $2R < n$, cannot be nonzero multiples of n. Thus synthesis is injective and its complex image has dimension B. As a real space that image has dimension $2B$; applying the real-linear injection E_j preserves dimension. The factor two represents independent real and imaginary coefficients stored in two real internal coordinates, not two copies of a real scalar band.

For completeness, energy orthogonality is needed as well as mass orthogonality. For each i,

$$D_i f_a = \sum_k a_k (e^{2\pi \mathrm{i} k_i/n} - 1)\chi_k.$$

Applying the same Gram sum to these derivative amplitudes gives

$$\|E_j f_a\|^2 = n^d \sum_k |a_k|^2,$$
$$\langle E_j f_a, K_{0,n} E_j f_a\rangle = n^d \sum_k |a_k|^2 \sum_i |e^{2\pi \mathrm{i} k_i/n} - 1|^2.$$

Since $|e^{\mathrm{i}t} - 1| \leq |t|$ and $|k_i| \leq R$, each sum of multipliers is at most $4\pi^2 d R^2/n^2$. This proves the bound for every member of the constructed image, not only its individual generators. □

5 The finite information budget

Let F be a finite-dimensional real vector space, $C : E_n \to F$ a linear measurement and $V : E_n \to E_n$ a linear operator. No symmetry or positivity of V is needed. Let $\Phi : F \to \mathbb{R}$ be any function with $\Phi(0) = 0$.

Theorem 5.1. *If $c > b_R(n)$ and*

$$c\|A\|^2 \le \langle A, K_{0,n}A\rangle + \Phi(CA) + \langle A, VA\rangle \quad \textit{for all } A \in E_n, \tag{5}$$

then $2(2R+1)^d \le \dim F + \operatorname{rank} V$.

Proof. If the budget is smaller than $\dim S_R$, restrict the map (C, V) to S_R. Its image has dimension at most $\dim F + \operatorname{rank} V$. Rank-nullity yields a nonzero $A \in S_R$ with $CA = VA = 0$. Equation (5) and Lemma 4.1 imply $c\|A\|^2 \le b_R(n)\|A\|^2$, a contradiction. □

The codomain dimension is a conservative measurement budget. We do not label this formal endpoint as a sharper theorem with $\operatorname{rank} C$ in its place. No regularity, sign, polynomial form or growth restriction on Φ enters.

Proposition 5.2. *If $P : E_n \to E_n$ is linear, $c > b_R(n)$ and $\dim F + \operatorname{rank} P < 2(2R+1)^d$, the inequality*

$$c\|A\|^2 \le \langle A, K_{0,n}A\rangle + \Phi(CA) \qquad (A \in \ker P)$$

is false.

Proof. Apply the same rank-nullity argument to $(C, P)|_{S_R}$. Its nonzero kernel vector belongs to the stipulated sector as well as the unobserved band. □

For example, $d = 4, n = 16, R = 1$ gives dimension 162 and $b_R(n) = \pi^2/16 < 1$. Twelve measurement coordinates and a repair of rank at most 149 cannot yield (5) with $c = 1$. This illustrates a necessary budget, not the programme's chosen target constant.

6 Uniform refinement and sector exclusion

Fix a positive integer N' and put $n_L = (L+1)N'$ for nonnegative integers L. The finite-dimensional space F is fixed. Allow arbitrary families of linear maps $C_L : E_{n_L} \to F$, arbitrary functions $\Phi_L : F \to \mathbb{R}$ with $\Phi_L(0) = 0$, and linear maps $V_L : E_{n_L} \to E_{n_L}$ with $\operatorname{rank} V_L \le r$ for one fixed nonnegative integer r.

Theorem 6.1. *There is no $c > 0$ for which (5) holds for every L and every $A \in E_{n_L}$ with these families. Likewise, for linear P_L of rank at most r, there is no positive uniform c for the sector estimates of Proposition 5.2 on all $\ker P_L$.*

Proof. Set $M = \dim F$ and $R = M + r + 1$. Since $d \ge 1$, $2(2R+1)^d > M + r$. Choose L so large that $2R < n_L$ and $b_R(n_L) < c$. Here is a fully finite choice argument, without an assumed limit theorem: with $a = 4\pi^2 dR^2$, take an integer $L > \max\{2R, a/c, 1\}$. Since $N' \ge 1$, $n_L \ge L + 1 > L > 1$, and $a < cL \le cn_L^2$. Theorem 5.1 now contradicts the rank bound. Use Proposition 5.2 for the separate sector conclusion. □

The constants c, r precede the quantifier over L. Arbitrary rescalings of C_L preserve its fixed codomain and do not escape the theorem. Removing harmonic or constant directions is covered only when a concrete linear constraint describing the proposed sector has uniformly bounded rank. A quotient or its norm is not silently identified with such a sector. Linearity is essential: if nonlinear observations were allowed, the single real observation $C_L(A) = \|A\|^2$, with $\Phi_L(y) = y$ and $V_L = 0$, would give the uniform lower bound $c = 1$ since $K_{0,n_L} \succeq 0$. This counterexample explains why finite-dimensional output alone is not the hypothesis.

7 Constructed derivative repairs and tangent observations

The bounded-rank obstruction does not exhaust the available repairs. For the actual componentwise forward gradient $\mathcal{G}A = (D_k A_i)_{x,i,k}$, finite periodic summation and cancellation of the mixed differences give

$$\|\mathcal{G}A\|^2 = \langle A, K_{0,n} A\rangle.$$

This is an identity on the whole cochain space, with unordered plaquettes counted once; it is not an assumed comparison on a supplied band. Let W be the diagonal multiplication by real weights $w_{x,i,k}$, and construct $V_w = \mathcal{G}^* W \mathcal{G}$. If $w_{x,i,k} \le M$ and $M \ge 0$, then

$$\langle A, V_w A\rangle = \sum_{x,i,k} w_{x,i,k} \|(D_k A_i)(x)\|^2 \le M \langle A, K_{0,n} A\rangle. \tag{6}$$

A lower bound on the weights is unnecessary for this upper estimate. Consequently the constructed Fourier band satisfies

$$q_w(A) := \langle A, (K_{0,n} + V_w)A\rangle \le (1 + M) b_R(n) \|A\|^2 \qquad (A \in S_R).$$

These operator identities and inequalities are formally proved in the separately pinned source reported below.

Theorem 7.1 (Finite-dimensional tangent obstruction)**.** *Let E be a finite-dimensional real inner-product space, F a finite-dimensional real normed space, and $S \subseteq E$ a linear subspace with $\dim S > \dim F$. Let $q : E \to \mathbb{R}$ satisfy $q(x) \le b\|x\|^2$ on S. Suppose $C : E \to F$ is Fréchet differentiable at zero, $C(0) = 0$, and for some $\delta > 0$ and finite $K \in \mathbb{R}$,*

$$\Phi(y) \le K\|y\|^2 \qquad (\|y\| < \delta). \tag{7}$$

If $c > b$, the full-space estimate $c\|x\|^2 \le q(x) + \Phi(C(x))$ is impossible.

Proof. Put $L = DC(0)$. Rank-nullity applied to the actual restriction $L|_S$ produces $v \in S \cap \ker L$, $v \neq 0$. Differentiability gives $\|C(tv)\|/\|tv\| \to 0$ as nonzero real $t \to 0$. For small such t, (7) applies and the proposed bound implies

$$c \le b + K \frac{\|C(tv)\|^2}{\|tv\|^2}.$$

Taking the limit gives $c \le b$, a contradiction. The kernel vector and the remainder limit are derived, not introduced as hypotheses of the theorem. □

The three focal Lean endpoints include this theorem, its remainder lemma and the actual physical specialization with threshold $(1 + M)b_R(n)$. The proof does not require an explicit assumption $\Phi(0) = 0$. A local-ball version is also valid by the same shrinking-ray argument, while the original formal endpoint above is stated on the full space. The later local kernel and average-weight endpoint in Section 9 has its own separately reproduced proof.

This result trades hypotheses. The global linearity of C is replaced by differentiability at zero, while an arbitrary zero-vanishing penalty is replaced by genuine quadratic upper control. Merely asking C and Φ to be C^1 with zero derivatives at zero is insufficient: for $1 < \alpha < 2$,

$$C(x) = \|x\|^\alpha, \qquad \Phi(y) = |y|^{2/\alpha} \quad \Longrightarrow \quad \Phi(C(x)) = \|x\|^2.$$

Both derivatives at zero vanish. Off zero, $\|DC(x)\| = \alpha\|x\|^{\alpha-1}$ and $|\Phi'(y)| = (2/\alpha)|y|^{2/\alpha-1}$ tend to zero, so both maps are indeed C^1 there and at zero. But $\Phi(y)/|y|^2 = |y|^{2/\alpha-2}$ is unbounded above near zero. On the actual lattice, setting $w = 0$ gives a uniform lower bound with $c = 1$ because $K_{0,n} \succeq 0$. This is an exact counterexample, not a numerical test.

8 Exact nonlinear regular sectors

The next theorem is an analytic consequence of rank-nullity and the classical implicit function theorem. The actual curve construction and local-energy obstruction now have separately pinned formal proofs, identified in the evidence section. The pinned formal library supplies an implicit parametrization with derivative the inclusion of the derivative kernel [6]; we construct and use that parametrization explicitly below rather than assuming the curve needed for the conclusion.

Theorem 8.1 (Regular-sector tangent obstruction)**.** *Let E be a finite-dimensional real inner-product space and F, G finite-dimensional real normed spaces. Let $T : E \to E$ be linear and set $q(x) = \langle x, Tx\rangle$. Suppose a linear subspace S satisfies $q(v) \le b\|v\|^2$ on S and $\dim S > \dim F + \dim G$. Let C be differentiable at zero with $C(0) = 0$ and let Φ satisfy (7). Let $P : E \to G$ be C^1 on a neighborhood of zero with $P(0) = 0$ and $A = DP(0)$ surjective. For every $c > b$ and every $r > 0$, there is no estimate*

$$c\|x\|^2 \le q(x) + \Phi(C(x)) \qquad (P(x) = 0,\ \|x\| < r). \tag{8}$$

Proof. Put $L = DC(0)$. Apply rank-nullity to $(L, A)|_S : S \to F \times G$ to obtain $v \in S$, $v \neq 0$, with $Lv = Av = 0$. Choose a linear complement W of $H = \ker A$. Surjectivity implies $A|_W : W \to G$ is bijective. Apply the implicit function theorem to $(u, w) \mapsto P(u + w)$ at $(0, 0) \in H \times W$. It produces a C^1 map ψ near zero in H with

$$\psi(0) = 0, \quad P(u + \psi(u)) = 0, \quad D\psi(0) = -(A|_W)^{-1}A|_H = 0.$$

Thus $\gamma(t) = tv + \psi(tv)$ lies in the exact sector for small real t, $\gamma(0) = 0$, and $\gamma(t) = tv + e(t)$ with $\|e(t)\| = o(|t|)$. In particular $\|\gamma(t)\|/|t| \to \|v\| > 0$, so the nonzero points of this curve eventually lie in the stipulated ball.

The curve need not remain in S. This is why the quadratic form assumption is used: direct expansion gives

$$q(\gamma(t)) = t^2 q(v) + t\langle v, Te(t)\rangle + t\langle e(t), Tv\rangle + \langle e(t), Te(t)\rangle.$$

Since T is bounded in finite dimension, division by t^2 gives $q(\gamma(t))/t^2 \to q(v)$. Differentiability of C and $Lv = 0$ give $C(\gamma(t)) = Le(t) + o(|t|) = o(|t|)$. Apply the assumed bound along this

exact curve and use (7). Division by t^2 and the limit yield $c\|v\|^2 \le q(v) \le b\|v\|^2$, contrary to $c > b$. □

Unlike the preceding full-space result, an arbitrary energy with an upper bound only on S would not justify this proof: the constructed exact-sector curve generally leaves S. The displayed expansion supplies the missing comparison for the actual quadratic energy.

A concrete curved example also shows why an exact vector in S is not the right object. Write $E = \mathbb{R}^4$ with coordinates (u_1, u_2, y, z), put $P(u, y, z) = z - \|u\|^2$, $C(u, y, z) = y$, $\Phi(y) = y^2$, $S = \{z = 0\}$ and $q(u, y, z) = z^2$. Then $DP(0)$ is surjective and $\dim S = 3 > 1 + 1$. But $S \cap \ker C \cap P^{-1}(0) = \{0\}$: the linear band contains no nonzero exact-sector unseen vector. Nevertheless $\gamma(t) = (t, 0, 0, t^2)$ lies in the exact sector and has repaired Rayleigh quotient $t^2/(1 + t^2) \to 0$. The implicit-curve argument is genuinely needed.

Nor can quadratic q in Theorem 8.1 be replaced merely by a C^1 energy stationary at zero with the same upper bound on S. With the same E, C, Φ, S, instead put

$$P(u, y, z) = z - \|u\|^{4/3}, \qquad e(u, y, z) = |z|^{3/2}.$$

Both maps are C^1, $DP(0)(u, y, z) = z$ is surjective, $De(0) = 0$, and $e = 0$ on S, so $b = 0$. On the exact sector,

$$e(x) + \Phi(C(x)) = \|u\|^2 + y^2, \qquad \|x\|^2 = \|u\|^2 + y^2 + \|u\|^{8/3}.$$

For $\|x\| < 1$, $\|u\|^{8/3} \le \|u\|^2$, so the local sector bound with $c = 1/2 > b$ holds. Differentiability of the norm powers at zero follows from their exponents being greater than one; their gradients off zero tend to zero. This exact counterexample refutes a proposed C^1-energy weakening, while the actual quadratic Hodge/repair energy retains the required expansion.

The regularity condition cannot simply be dropped. Take $E = \mathbb{R}^3$, $F = G = \mathbb{R}$, $C = 0$, $\Phi = 0$, $T = 0$, $S = E$, and $P(x) = \|x\|^2$. Then P is smooth, $DP(0) = 0$, and the strict dimension budget holds, but $P^{-1}(0) = \{0\}$ and the estimate with $c = 1$ is vacuous on the sector. This refutes the proposal to replace submersion by bounded derivative rank. It does not refute Theorem 8.1, because $DP(0)$ is not surjective. Critical constraint maps require a separate analysis.

Proposition 8.2 (Uniform physical consequence)**.** *Let $n_L \to \infty$ with fixed d, N_c, and let F, G have fixed finite dimensions. For each L, allow its own C_L, P_L, Φ_L satisfying the hypotheses of Theorem 8.1. Let q_L be the actual flat Hodge form plus the constructed weighted derivative repair with weights $w_L \le M_L$, $M_L \ge 0$, and $M_L/n_L^2 \to 0$. Even on the exact sectors $P_L = 0$ and arbitrarily small balls of positive radius r_L, no positive uniform coercivity constant exists.*

Proof. Choose a fixed integer R with $2(2R+1)^d > \dim F + \dim G$. For all large L, $2R < n_L$ and the constructed band obeys $q_L \le (1 + M_L) b_R(n_L) \|\cdot\|^2$. Its coefficient tends to zero. For any proposed $c > 0$, choose one such L with this coefficient below c and apply Theorem 8.1 at that fixed lattice. The curve is shrunk only after choosing L: no uniform implicit-function radius, inverse-derivative norm, penalty constant or r_L is required. □

These are finite unit-spacing obstructions for regular nonlinear sectors, not identifications of a physical gauge-fixed domain or a quotient norm.

9 Directional mean weights and local nonlinear observations

Pointwise bounds on a derivative repair are stronger than the following actual spatial-mean condition. Fix $d \geq 2$, $N_c \geq 2$, and distinct directions i, j. For nonnegative componentwise weights $w_n(k,(x,l))$, put

$$Q_n(A) = \langle A, K_{0,n}A\rangle + \sum_{x,k,l} w_n(k,(x,l))\|A_l(x+e_k) - A_l(x)\|^2, \qquad b_{ij}(n) = n^{-d}\sum_x w_n(i,(x,j)).$$

The repair is the actual G^*W_nG of the componentwise forward gradient. In particular $Q_n \geq 0$; signed weights are not included in this section. Let $D_n = \ker\delta$ be the actual flat divergence-free space.

Theorem 9.1 (Local nonlinear average-weight obstruction)**.** *For each $n \geq 1$, let F_n be a finite-dimensional real normed space, $C_n : E_n \to F_n$ be Fréchet differentiable at zero with $C_n(0) = 0$, and let $\Phi_n : F_n \to \mathbb{R}$ satisfy*

$$\Phi_n(y) \leq K_n\|y\|^2 \quad (\|y\| < \delta_n), \qquad \delta_n > 0.$$

Suppose one fixed $c > 0$ and positive radii r_n satisfy

$$c\|A\|^2 \leq Q_n(A) + \Phi_n(C_n(A)) \quad (A \in D_n,\ \|A\| < r_n). \tag{9}$$

For every integer $0 \leq R < n$,

$$c(2R - \dim F_n) \leq 2(1 + b_{ij}(n))\sum_{s=1}^{R}|e^{2\pi\mathrm{i}s/n} - 1|^2 \leq \frac{8\pi^2(1+b_{ij}(n))R^3}{n^2}. \tag{10}$$

Consequently, if $(1 + b_{ij}(n))/n^2 \to 0$, then $\dim F_n \to \infty$. No uniform bound on K_n, δ_n^{-1} or r_n^{-1} is required.

Proof. Write $L_n = DC_n(0)$. If $v \in D_n \cap \ker L_n$, differentiability gives $C_n(tv) = o(|t|)$. For all sufficiently small nonzero t, both the local coercivity and penalty bound apply. Divide by t^2 and let $t \to 0$ to get

$$c\|v\|^2 \leq Q_n(v). \tag{11}$$

This argument is performed separately at each n; it does not impose uniform neighborhoods or derivative remainder rates.

Use the two internal coordinates u, v of Section 2 to regard a complex number as a real internal vector. For $1 \leq s \leq R$ and $a \in \{1, \mathrm{i}\}$, define $\psi_{s,a}$ by taking only component j nonzero and setting it to $n^{-d/2}$ times the internal vector corresponding to $ae^{2\pi\mathrm{i}sx_i/n}$. The finite-character orthogonality identity and the real orthogonality of the two quadratures give an orthonormal family of $2R$ fields. This also covers even n; the two quadratures use two internal coordinates, not two allegedly independent scalar sine and cosine functions. Since $i \neq j$, these fields are independent of x_j, so $\delta\psi_{s,a} = 0$. Their span $S_{R,n}$ is thus an actual $2R$-dimensional subspace of D_n.

Every nonzero forward difference of $\psi_{s,a}$ is in direction i and has the same squared norm at every site, namely $n^{-d}|e^{2\pi\mathrm{i}s/n} - 1|^2$. The exact Hodge and repair energies are therefore

$$Q_n(\psi_{s,a}) = (1 + b_{ij}(n))|e^{2\pi\mathrm{i}s/n} - 1|^2.$$

In particular the trace of the compression of Q_n to $S_{R,n}$ is the middle expression in (10); no invariance of $S_{R,n}$ under the weighted operator is asserted or needed.

For completeness, the positive trace/rank argument is as follows. Set $U = S_{R,n} \cap \ker L_n$. Rank-nullity gives $\dim U \geq 2R - \dim F_n$. Extend an orthonormal basis of U to one of $S_{R,n}$. Equation (11) bounds the terms on U below by c; positivity of Q_n makes every remaining diagonal term nonnegative. Hence $c(2R - \dim F_n) \leq \operatorname{tr}_{S_{R,n}} Q_n$. Finally $|e^{2\pi \mathrm{i} s/n} - 1| \leq 2\pi s/n \leq 2\pi R/n$, which proves the last inequality in (10).

Fix R. The right side tends to zero under the stated mean condition. Eventually it is less than c. If $\dim F_n < 2R$, integrality would make the left side at least c. Thus $\dim F_n \geq 2R$ eventually. Arbitrarily large fixed R prove the claimed divergence. □

The condition concerns one directional spatial average, not the maximum of all weights. For example, a weight of size n^{d+1} at one site in component (i, j) and zero elsewhere has $b_{ij}(n) = n$, while its maximum is not subquadratic. Theorem 9.1 still excludes bounded output dimension under its other hypotheses. This example is a direct calculation, not an additional formal endpoint. The nonnegative-weight condition is essential to the trace argument as used here: the complement of the derivative kernel must contribute nonnegative energy. This is not a universal extension of older statements allowing signed weights with a pointwise absolute cap.

9.1 A constructed quadratic-weight boundary example

The exponent 2 cannot be replaced by an unrestricted allowance of quadratic-size mean weights. Let $\mu_j(A) = n^{-d} \sum_x A_j(x)$, with output in the same fixed space $F = (\mathbb{R}^{N_c^2 - 1})^d$ for every n. A finite path comparison on the actual periodic lattice gives

$$\|A - \operatorname{lift} \mu(A)\|^2 \leq dn^2 \|GA\|^2, \qquad \|\operatorname{lift} \mu(A)\|^2 = n^d \|\mu(A)\|^2.$$

One way to obtain the first inequality is to write each difference of site values as a path of at most dn forward edges, apply finite Cauchy–Schwarz, and average the pairs of sites; the displayed constant is deliberately nonoptimal. The exact modular-torus reindexing and its transport to the physical cochain norm are also proved in the companion formal construction. Using $\|x + y\|^2 \leq 2\|x\|^2 + 2\|y\|^2$, choose actual weights $w_n(k, (x, l)) = 2dn^2$ and the quadratic penalty $\Phi_n(y) = 2n^d \|y\|^2$. Then

$$\|A\|^2 \leq \langle A, G^* W_n G A \rangle + \Phi_n(\mu(A)) \leq Q_n(A) + \Phi_n(\mu(A))$$

on the entire space, hence also on D_n. The mean observation is an actual continuous linear map with finite-dimensional output type independent of n, not a postulated observation certificate. Here $b_{ij}(n) = 2dn^2$ and the subquadratic hypothesis fails. The example allows $K_n = 2n^d$ to vary, exactly as the obstruction does. It proves neither an optimal boundary constant nor positive coercivity for an interacting gauge operator. The path/mean mechanism is classical; see the primary exposition in [8].

9.2 Why the penalty assumption remains

The C^1 counterexample after Theorem 7.1 also applies to the actual nonnegative Q_n and to D_n. A second useful counterexample retains a polynomial observation and a globally Lipschitz penalty:

$$C(A) = \|A\|^2 \in \mathbb{R}, \qquad \Phi(y) = c|y|, \qquad DC(0) = 0, \qquad \Phi(C(A)) = c\|A\|^2.$$

This example already gives the bound for the actual admissible choice $w = 0$, with $Q_n = K_{0,n} \succeq 0$. The zero-form interpretation is only an abstract illustration, not an instance of this nonnegative-weight lattice family. No finite local quadratic upper constant for Φ exists, since $c|y|/y^2 = c/|y|$. Likewise, $C(A) = \|A\|^{3/2}$ and $\Phi(y) = c|y|^{4/3}$ are both C^1 with zero derivatives at zero, yet their composition is exactly quadratic. These are complete analytic counterexamples, not new Lean claims.

There is a distinct analytic regularity tradeoff, not a hypothesis removal for merely differentiable observations. If, on each fixed derivative-kernel direction, $C(tv) = t^2 H_v + o(t^2)$, and $\Phi(0) = 0$, $D\Phi(0) = 0$, then $\Phi(C(tv)) = o(t^2)$. Indeed $\|C(tv)\|/t^2$ is bounded and $|\Phi(y)| \le \epsilon\|y\|$ near zero, including at zero; multiply the two estimates and let $\epsilon \downarrow 0$. A C^2 observation and a nonnegative differentiable penalty vanishing at zero suffice for this step, because a differentiable local minimum has zero derivative. No actual observation's C^2 regularity is proved here, and no formal composition of this alternate route is claimed.

The strengthened mean-weight result is an output-rank divergence theorem. Its fixed-R argument does not assert a positive density of observations per lattice volume. Fixed $c > 0$, nonnegative weights and the stated penalty condition remain explicit. In particular no interacting or continuum Yang–Mills hypothesis, ultraviolet activity estimate, or physical spectral gap is supplied by this section.

10 Signed weights and the limits of mean criteria

Write $R_w = G^* W G$ and $Q_w(A) = \langle A, K_{0,n} A\rangle + \langle A, R_w A\rangle$. For signed weights this form need not be nonnegative. Its spatial means alone no longer supply the positive trace comparison used above. The following sufficient condition is nevertheless available without a lower bound on any original weight. Define the *pointwise* positive part and its selected directional mean by

$$w^+(k,(x,l)) = \max\{w(k,(x,l)),0\}, \qquad b^+_{ij}(n) = n^{-d}\sum_x w^+(i,(x,j)).$$

In general b^+_{ij} is not the positive part of the signed mean.

Theorem 10.1 (A sufficient signed-weight condition)**.** *Keep the observation, penalty, local divergence-free domain and fixed $c > 0$ of Theorem 9.1, but use arbitrary signed actual weights in the local coercivity inequality. For every $0 \le R < n$,*

$$c(2R - \dim F_n) \le \frac{8\pi^2(1 + b^+_{ij}(n))R^3}{n^2}.$$

If $(1 + b^+_{ij}(n))/n^2 \to 0$, then $\dim F_n \to \infty$. The original $C_n, L_n, \Phi_n, c, K_n, \delta_n, r_n$ are retained.

Proof. The exact weighted-gradient energy gives, for every actual cochain,

$$\langle A, R_w A\rangle = \sum_{x,k,l} w(k,(x,l))\|D_k A_l(x)\|^2 \le \sum_{x,k,l} w^+(k,(x,l))\|D_k A_l(x)\|^2 = \langle A, R_{w^+} A\rangle.$$

The inequality uses only nonnegativity of the squared differences. Thus the original local coercivity inequality transfers upward to the constructed nonnegative repair R_{w^+}, and Theorem 9.1 applies. Nonnegativity of w^+ and the comparison are proved from their actual definitions, not supplied as fields of a comparison structure. □

This changes the explicit mean hypothesis for signed weights. It neither proves the original weights nonnegative nor removes the local regularity, quadratic penalty or fixed positive coercivity requirements.

10.1 Zero signed means with fixed output

There is no counterpart using only ordinary signed directional means, even when *every* such mean is zero. Fix a site $o \in X_n$, set $V = n^d$, $a = 2dn^2$, and construct actual weights

$$w_n(k,(x,l)) = a(1 - V\mathbf{1}_{x=o}).$$

Since $|X_n| = V$, their ordinary spatial means are all exactly zero. Let $\mu(A)$ be the directional mean above, and let $J_n(A) = (D_k A_l(o))_{k,l}$ collect all actual forward differences at o. Use the same finite output type for every side:

$$F = (\mathbb{R}^m)^d \times (\mathbb{R}^m)^{d\times d}, \quad C_n(A) = (\mu(A), J_n(A)), \quad \|(y,z)\|_F = \max\{\|y\|_2, \|z\|_2\}.$$

Its dimension is $(d + d^2)m$, independent of n. The companion code proves the fixed finite output type; this displayed numerical dimension is an analytic calculation, not an additional named oracle endpoint.

Proposition 10.2. *The actual continuous linear observation C_n and the nonnegative quadratic penalty*

$$\Phi_n(y,z) = 2V\|y\|_2^2 + aV\|z\|_2^2$$

satisfy $C_n(0) = 0$, $DC_n(0) = C_n$, $\Phi_n(y,z) \le V(2+a)\|(y,z)\|_F^2$, and

$$\|A\|^2 \le Q_{w_n}(A) + \Phi_n(C_n(A)) \qquad (A \in E_n).$$

Consequently all the original local observation and penalty hypotheses can hold with fixed output and $c = 1$, although all signed means vanish.

Proof. Linearity follows directly from the finite sums and differences. The penalty bound follows from each projection norm being at most the specified product max norm. The exact energy expansion is

$$\langle A, R_{w_n} A\rangle = a\|GA\|^2 - aV\|J_n(A)\|_2^2.$$

Adding the penalty cancels the negative term and leaves

$$Q_{w_n}(A) + \Phi_n(C_n(A)) = \langle A, K_{0,n}A\rangle + a\|GA\|^2 + 2V\|\mu(A)\|_2^2 \ge \|A\|^2.$$

The last step is the actual mean/path estimate of the preceding section, with $a = 2dn^2$, together with $K_{0,n} \succeq 0$. □

The coefficient $K_n = V(2+a)$ may grow with the lattice side, exactly as allowed by the proposed ordinary-mean extension. This is not a transfer of the desired coercivity into an assumed observation certificate. The weights, observation, penalty and cancellation are constructed. For $n > 1$ their positive-part mean is $a(1 - V^{-1})$, of quadratic rather than subquadratic order, so they do not contradict Theorem 10.1.

As an analytic diagnostic, take $n \ge 3$, $i \ne j$, a unit internal vector u, and $A_j(x) = u\mathbf{1}_{x_i=o_i}$, with other components zero. Then $\delta A = 0$, $\|GA\|^2 = 2n^{d-1}$ and $\|J_n(A)\|^2 = 1$, giving

$$Q_{w_n}(A) = n^{d-1}(2 - a(n-2)) < 0.$$

This explains the failure of the original positive trace argument. The diagnostic field calculation is analytic; the formal fixed-output family theorem does not claim to check this additional field explicitly.

10.2 An attributed square cancellation

Fix distinct directions i, k, a component $j \neq i$ (for instance $j = k$), and $o \in X_n$. For an amplitude $A_0 \geq 0$, define the repair weights additively by

$$\begin{aligned} w(p) = {} & A_0 \mathbf{1}_{p=(i,(o,j))} - 3A_0 \mathbf{1}_{p=(k,(o,j))} \\ & - 3A_0 \mathbf{1}_{p=(i,(o+e_k,j))} - 3A_0 \mathbf{1}_{p=(k,(o+e_i,j))}. \end{aligned}$$

Writing $D_s A_j(x) = A_j(x + e_s) - A_j(x)$, commuting shifts give

$$D_i A_j(o) = D_k A_j(o) + D_i A_j(o + e_k) - D_k A_j(o + e_i).$$

The squared norm of this sum is at most three times the sum of the three squared norms. The actual finite weighted-energy sum therefore proves $\langle A, R_w A \rangle \leq 0$ for every cochain. For $n \geq 2$ the positive edge is separate from all negative contributions, so $w^+ = A_0 \mathbf{1}_{p=(i,(o,j))}$ and $b^+_{ij} = A_0/n^d$. Choosing $A_0 = n^d n^2$ gives $b^+_{ij} = n^2$.

If the original local coercivity hypotheses hold for this actual repair, they imply the same hypotheses with zero repair, since $Q_w(A) \leq \langle A, K_{0,n} A \rangle$. Applying Theorem 9.1 with zero weights again forces $\dim F_n \to \infty$. The selected normalized positive-part mean is $1 + n^{-2}$ and does not tend to zero. This last limit interpretation is an elementary analytic consequence of the formal exact mean identity, not a separate negated-limit oracle endpoint. This single sparse cycle does not make every directional criterion fail: other eligible means are zero. Its formal endpoints establish the selected mean and derived order, not an all-directions counterexample.

There is, however, a finite sum that covers every eligible pair. This extension is analytic here. For every j and every $i \neq j$, use the preceding cycle with $k = j$, the same site o and amplitude A_0, and add all its weight functions to obtain $\overline{w}$. Linearity of the exact weighted-energy sum gives $\langle A, R_{\overline{w}} A \rangle \leq 0$. For every off-diagonal pair $i \neq j$, direct inspection of the four indicators gives

$$\overline{w}(i, (x, j)) = A_0 \mathbf{1}_{x=o} - 3A_0 \mathbf{1}_{x=o+e_j}.$$

Every diagonal component has only nonpositive contributions. For $n \geq 2$ the two displayed sites are distinct, so each eligible pair has $\overline{b}^+_{ij} = A_0/n^d$. With $A_0 = n^d n^2$, *all* off-diagonal positive-part means equal n^2, while the derived domination by zero still gives the original rank obstruction. Thus no eligible mean condition in Theorem 10.1 applies to this sum. The conclusion follows from its actual energy order, not from an assumed comparison field. The finite sum and its all-pairs mean calculation are not silently included in the single-cycle formal oracle reports.

The cancellation mechanism is classical, rather than a new general spectral principle. Restricted to the four active vertices, the negative of this scalar repair has three positive path conductances $3A_0$ and a direct negative conductance $-A_0$. For $A_0 > 0$ the path resistance is $1/A_0$, so this is the critical case of Zelazo and Bürger's Theorem III.3 [9]. Isolated vertices contribute zero blocks; the internal vector coordinates sum scalar copies. This comparison attributes the mechanism; it does not establish originality of the whole consolidated dossier.

The conditional obstruction is nonvacuous. An analytic check takes $F_n = E_n$, $C_n = I$ and $\Phi_n(y) = \|y\|^2 - \langle y, R_w y \rangle \geq 0$ (or replace w by $\overline{w}$). The operator norm gives the upper quadratic

bound $(1 + \|R_w\|)\|y\|^2$, while $Q_w(A) + \Phi_n(A) = \langle A, K_{0,n}A\rangle + \|A\|^2 \geq \|A\|^2$. This example has growing output, as the theorem requires; it is not an additional formal endpoint or a compressed observation construction. None of these signed-weight statements proves a volume-density bound, identifies an interacting Hessian, or supplies a physical mass gap.

11 Full volume rank from signed total coefficient means

The directional strip in Section 9 detects divergence of the observation dimension. A full dimensional coexact packet yields a positive volume density under a stronger, explicitly stated mean hypothesis. Original weights may have either sign. The relevant quantity is the spatial mean of $\max\{1 + w, 0\}$, not an ordinary signed mean or a positive part taken after averaging. The extra 1 is the actual flat Hodge coefficient, already absorbed into the envelope; it must not be added a second time.

In this section write $D \geq 2$ for the number of lattice directions, $V_n = n^D$, and E_n for the actual real cochain space on $(\mathbb{Z}/n\mathbb{Z})^D$, with fixed $N_c \geq 2$. Use field components $0, 1$ and the internal real plane representing a complex scalar from Section 2. Put $D_i A_l(x) = A_l(x+e_i) - A_l(x)$ and retain the original form

$$Q_{w_n}(A) = \langle A, K_{0,n}A\rangle + \sum_{x,i,l} w_n(i,(x,l))\|D_i A_l(x)\|^2.$$

The actual all-cochain flat Hodge identity gives $Q_{w_n}(A) = \sum_{x,i,l}(1 + w_n(i,(x,l)))\|D_i A_l(x)\|^2$.

Theorem 11.1 (Finite full volume mean budget)**.** *Let $1 \leq R$ and $2R < n$. Assume $m \geq 0$ and*

$$\sum_x \max\{1 + w_n(i,(x,l)), 0\} \leq V_n m \qquad (0 \leq i < D,\ l = 0, 1). \tag{12}$$

No mean bound on the other field components, original weight lower bound or prescribed uniform pointwise cap is assumed. Let $C_n : E_n \to F_n$ be Fréchet differentiable at zero, $C_n(0) = 0$, with F_n finite dimensional. Suppose $\Phi_n(y) \leq K_n\|y\|^2$ for $\|y\| < \delta_n$, $\delta_n > 0$, and the original small divergence-free coercivity inequality holds with $c \geq 0$ and some radius $r_n > 0$. Set

$$H = 2R(2R+1)^{D-1}, \qquad \Lambda(k) = 4\sum_{i=0}^{D-1} \sin^2(\pi k_i/n), \qquad \mathcal{K} = \{1, \ldots, R\} \times \{-R, \ldots, R\}^{D-1}.$$

Then

$$c(H - \dim_{\mathbb{R}} F_n) \leq m \sum_{(k,q)\in\mathcal{K}\times\{0,1\}} \Lambda(k) \leq mH\,\frac{4\pi^2 DR^2}{n^2}. \tag{13}$$

Proof. For $k \in \mathcal{K}$ use the actual scalar character χ_k and the coexact field with components given by its backward differences in directions $1, 0$, with opposite signs. Commuting differences give actual divergence zero. Since $k_0 \neq 0 \pmod n$, the normalization denominator is positive. Counting Fourier orthogonality and both internal real quadratures produce an orthonormal family of H actual real fields, spanning S_n. The condition $2R < n$ prevents aliasing of their scalar frequencies; no scalar real-valued conjugacy constraint is imposed on the internal plane. Every vector of S_n is supported on field components $0, 1$.

Define actual coefficients $v_n(i,(x,l)) = \max\{1 + w_n(i,(x,l)), 0\}$ for $l = 0, 1$, and zero otherwise, and the positive form $P_n = G^* V_n^{\mathrm{wt}} G$ with those coefficients. The notation V_n^{wt} here denotes

the weight multiplication operator, not lattice volume. Then $P_n \succeq 0$ on all cochains and $Q_{w_n}(A) \le \langle A, P_n A\rangle$ on S_n. This comparison is needed only on the actual packet image, and is not asserted on every cochain or every divergence-free field.

Write $L_n = DC_n(0)$. The same shrinking-ray proof as in Section 9, now followed by the proved envelope comparison, gives $c\|v\|^2 \le \langle v, P_n v\rangle$ on $S_n \cap \ker L_n$. Its dimension is at least $H - \dim F_n$. Complete an orthonormal basis of that kernel within S_n. Positivity of P_n controls the complementary trace terms. Thus $c(H - \dim F_n) \le \operatorname{tr}(P_n|_{S_n})$, where the restriction means the genuine orthogonal compression.

For an individual normalized packet $U_{k,q}$, its squared component amplitudes are constant in x and are $V_n^{-1}a_{k,l}$ for $l = 0, 1$, where $a_{k,l} \ge 0$ and $a_{k,0} + a_{k,1} = 1$. Consequently

$$\langle U_{k,q}, P_n U_{k,q}\rangle = V_n^{-1} \sum_{i=0}^{D-1} \sum_{l=0}^{1} |\chi_k(e_i) - 1|^2 a_{k,l} \sum_x v_n(i,(x,l)) \le m \sum_i |\chi_k(e_i) - 1|^2 = m\Lambda(k).$$

Sum over all H basis vectors. Finally every signed integer frequency satisfies $|k_i| \le R$, and $|\sin t| \le |t|$ gives $\Lambda(k) \le 4\pi^2 DR^2/n^2$. All D derivative directions remain in this sum even though only two field components are occupied. □

Theorem 11.2 (One positive density on the full integer tail)**.** *Fix D, N_c, $c > 0$ and $m \ge 0$. For each integer n under consideration, allow its own $F_n, C_n, L_n, \Phi_n, K_n, \delta_n, r_n, w_n$ satisfying the original local hypotheses and* (12)*, with these same c, m. There are $k \ge 3$ and $\varepsilon > 0$, depending only on D, m, c, such that*

$$\dim_{\mathbb{R}} F_n \ge \varepsilon n^D \qquad (n \ge k)$$

whenever those finite-side hypotheses hold. If the original family hypotheses hold only for $n \ge n_0$, the conclusion applies for $n \ge \max\{n_0, k\}$. One may take $\varepsilon = 1/(2k^D)$ for any integer $k \ge 3$ satisfying $2m\,4\pi^2 D/k^2 \le c$. No common quadratic penalty constant, differentiability neighborhood or coercivity radius is required.

Proof. Choose such a k by the Archimedean property. For each $n \ge k$ put $R = \lfloor n/k \rfloor$. Integer division yields

$$R \ge 1, \qquad kR \le n < k(R+1) \le 2kR, \qquad 2R < kR \le n.$$

Hence the actual packet is admissible and $R/n \le 1/k$. The finite theorem gives $c(H - \dim F_n) \le (c/2)H$, so $H \le 2\dim F_n$. Also $H \ge (2R)^D \ge (n/k)^D$. Combining these gives the displayed density. Each shrinking-ray argument was applied at a fixed n before using these integer inequalities. No limit of derivatives or local neighborhoods is taken. □

Proposition 11.3 (Exact trace average and a sharper analytic constant)**.** *For $R \ge 1$, put $A_R(x) = \sum_{r=1}^{R} \sin^2(rx)$. The actual multiplier average in* (13) *is*

$$B_{D,R,n} = \frac{1}{H} \sum_{k,q} \Lambda(k) = 4A_R(\pi/n) \frac{2DR+1}{R(2R+1)},$$

$$A_R(x) = \frac{R}{2} - \frac{\sin(Rx)\cos((R+1)x)}{2\sin x},$$

$$B_{D,R,n} \le \frac{2\pi^2}{3n^2}(R+1)(2DR+1).$$

For fixed $0 < \alpha < 1/2$ and $R = \lfloor \alpha n \rfloor$, the original family hypotheses therefore imply the analytic bound

$$\liminf_{n\to\infty} \frac{\dim F_n}{n^D} \geq (2\alpha)^D \left[1 - \frac{m}{c}\, 2D \left(1 - \frac{\sin(2\pi\alpha)}{2\pi\alpha}\right)\right].$$

Its elementary quadratic lower estimate replaces the bracket by $1 - 4\pi^2 D m \alpha^2/(3c)$. This improves the low-frequency quadratic coefficient of the maximum-coordinate estimate by a factor three.

Proof. The positive first coordinate contributes $4A_R/R$ to the average; each of the $D-1$ centered coordinates contributes $8A_R/(2R+1)$. Both real quadratures occur in numerator and denominator. The cosine sum follows by summing $2\sin x \cos(2rx) = \sin((2r+1)x) - \sin((2r-1)x)$ and telescoping; $\sin^2(rx) = (1-\cos(2rx))/2$ gives the formula for A_R. Bounding $\sin^2(rx) \leq r^2x^2$ and summing r^2 gives the finite upper bound, including its lower-order terms. Floor inequalities give $H/n^D \to (2\alpha)^D$, and $\sin x/x \to 1$ gives $B_{D,R,n} \to 2D[1 - \sin(2\pi\alpha)/(2\pi\alpha)]$. Divide the finite rank bound by n^D and take the lower limit. Negative right sides are merely weak bounds. For $m = 0$, taking $\alpha \uparrow 1/2$ gives lower density one. □

These are classical Fourier and finite trace arguments, with the standard discrete-torus spectrum already present in the literature [10]. The exact average and its improved constant are analytic calculations here, not extra formal endpoints or sharp optimal observation-rank claims. They do not identify an interacting background or a continuum Yang–Mills spectral gap.

12 Why dimension is not a sufficiency criterion

The following analytic statement is included to prevent the reverse inference from a necessary rank budget. It is not a new Lean-verified endpoint here. The proof below is the classical singular-Schur range condition [3, 4], expressed as boundedness below on each measurement fiber; we claim no new general matrix criterion. Let E be a finite-dimensional real inner-product space, $H = H^*$ and $B = H - cI$. Fix linear $C : E \to F$, and decompose $E = N \oplus W$ with $N = \ker C$, $W = N^\perp$. In this decomposition write

$$B = \begin{pmatrix} A & D \\ D^* & G \end{pmatrix}.$$

Proposition 12.1. *A finite real-valued function Φ with $\Phi(0) = 0$ can make $\langle x, Bx\rangle + \Phi(Cx) \geq 0$ for every x if and only if $A \succeq 0$ and $\ker A \subseteq \ker D^*$. In that event a quadratic penalty suffices.*

Proof. Taking $x \in N$ proves $A \succeq 0$. For $z \in \ker A$ and fixed $w \in W$, the value at $tz + w$ contains $2t\langle z, Dw\rangle$, while $C(tz + w) = Cw$ is fixed. Boundedness below for all real t forces this coefficient to vanish.

Conversely, $\ker A \subseteq \ker D^*$ implies $D(W) \subseteq \operatorname{ran} A$. Let $A^\dagger$ be the inverse on $\operatorname{ran} A$ and zero on $\ker A$. Completing the square gives

$$\langle n + w, B(n + w)\rangle = \langle n + A^\dagger Dw, A(n + A^\dagger Dw)\rangle + \langle w, (G - D^*A^\dagger D)w\rangle.$$

The restriction $C|_W$ is injective; extend its inverse on $C(W)$ to a linear map $T : F \to W$. The quadratic function $\Phi(y) = \langle Ty, (D^*A^\dagger D - G)Ty\rangle$ cancels the second term and vanishes at zero. If F has a chosen Euclidean norm, a sufficiently large nonnegative multiple of $\|y\|^2$ also suffices, by finite-dimensional norm bounds. □

For a concrete counterexample let $H = \mathrm{diag}(0, 2)$, $c = 1$ and $C(x, y) = x + y$. On $\ker C$, $\langle x, (H - I)x\rangle = 0$, and C detects the negative eigendirection of $H - I$. Nevertheless, on $(t, 1 - t)$ the repaired form is $1 - 2t + \Phi(1)$, unbounded below for every finite $\Phi(1)$. Kernel nonnegativity and low-mode detection alone do not suffice; cross-coupling must be checked.

13 Operator scaling and the actual consumer

In the repository, the flat fixed-volume covariance endpoint uses $K_0 + aQ^*Q$ with the physical counting ℓ^2 norm and the same unscaled flat Hodge operator. More precisely, the declared precision shell includes $-\Sigma$. Specializing every perturbation and its summable majorant to zero and taking $a > 0$ gives the available certified coercivity lower bound $\min(1, a)/C_P$, with $C_P > 0$ from the full-periodic block-Poincare estimate. This is not a claim that this value is optimal. At fixed d, N_c, N' the block codomain is fixed and finite. Inspection therefore identifies the unscaled flat base route with the setting of Theorem 6.1, with $\Phi_L(y) = a_L\|y\|^2$ and no fine repair. Even allowing arbitrary scalar a_L or rescaling Q_L does not yield a positive uniform full-space constant. This is a source-level specialization in this text; no new Lean adapter to that consumer is claimed. At the new source anchor below, the module `PhysicalGaugeFixedPrecision` defines `flatGaugeFixedPrecisionCLM`; its explicit-Poincare coercivity theorem is recorded by its full name and specialization in the linked declaration map.

The inspected endpoint is explicitly fixed-volume. Its positive rate may depend on L, N' and C_P, so there is no contradiction. Nor does the obstruction exclude all possible uniform Combes–Thomas arguments. Its locality radius and matrix-entry bounds also vary with L; coercivity alone would not settle the rate. An interacting Hessian, a growing measurement budget, a different sector norm or a physically rescaled operator must each be audited separately.

In particular, writing $\ell = L + 1$ so that $n_L = \ell N'$, replacing K_0 by $\ell^2 K_0$ changes the band threshold exactly to $4\pi^2 dR^2/(N')^2$. The proof above no longer drives it below every fixed $c > 0$. This observation is an escape from this proof, not an established coercivity theorem for the programme's physical operator. No ultraviolet activity estimate, continuum limit, Wilson-Hessian identification or Yang–Mills mass gap is obtained.

14 Formal evidence and publication context

The source-tree anchor for the new physical-band and uniform-rank results is

`e31b428917aa767b4a1aad097db5cf92fb9cd23e`.

The original factorized-kernel and repair development in Section 3 instead has the distinct tested source anchor

`7b25bfb39357406cf61880e3ac0428d7d1ba3449`.

Its exact modules are not present in the new main tree merely because a reproduction supplement is present there. The two full preserved focal runs at that older SHA agree in compiled object, oracle log and source manifest; their six named reports use only the three standard axioms. The reproduced factorized endpoint is dimension four. The general-dimensional witness calculation, trigonometric closed form, spectral multiplicity, liminf and witness-line statements are analytic here; they are not all silently included in that formalization claim. The external declaration map keeps both anchors. The public source repository is THE-ERIKSSON-PROGRAMME on

GitHub. The preserved parent v5 map is `declaration-map.json` and the preserved evidence is `reproducibility/`. Exact source SHAs, not moving branch heads, identify the formal claims. That parent map does not certify the new full volume extension. The map packaged with this manuscript records each added anchor and the separately reproduced full-tail extension. The accompanying exact-artifact report binds the final source and PDF hashes.

The geometric and Fourier modules are ordinary proof dependencies, not assumptions about a newly supplied band. Each previously verified stage identified below has complete same-source fresh-clone Colab evidence; the new extension boundary is stated separately. For the uniform endpoints the tested source is

`17ca395d4e9de992b1c575454a112da9e6f97e86`.

Its two runs agree in compiled object, raw oracle log and source-manifest hashes. The standalone budget lemma uses `propext, Quot.sound`; the other four uniform-module reports also use `Classical.choice`. No extra axiom is reported. Actual-log acceptance and three real-field rejection tests pass under normal and optimized Python. These tests concern evidence handling, not mathematical independence.

The additional whole-space gradient/Hodge identity and constructed weighted repair bounds use tested source

`dfd49d0b0730a4d36f58a4cbb7a17d6a41fb60bc`.

The full-space tangent endpoints in Section 7 use tested source

`81cf07b2e0d30543682c4073a4a2d82950d638b7`.

Each has two complete same-source fresh Colab Linux reconstructions with matching compiled object, raw oracle and LF-source-manifest hashes. For the tangent source, all three named endpoints use only `propext`, `Classical.choice` and `Quot.sound`. Its oracles and pinned Colab reproduction drivers are separate from the manuscript render. Failed source candidates and their full logs are retained in the programme evidence archive. Later, separately reproduced sources formalize the actual implicit-function curve construction at

`ef81f028ce6e7e8aaca7e9229eb00aa683752a41`.

The local quadratic-energy regular-sector obstruction uses

`9bb90cfd77ec5cea5c4447667fe8baf8c21429a1`.

The displayed sector proofs remain readable analytic arguments; their uniform consequence is not silently included in those four-endpoint reports. No theorem about a nonlinear sector is inferred from a full-space no-go alone.

Short labels and an external declaration map replace long wrapped Lean names in this manuscript. The companion source map records exact names, source SHAs, and evidence directories. Reproduction commands remain in the existing copyable notebooks rather than being split into misleading line fragments. The fixed-map Schur criterion, its counterexample calculation and consumer specialization above are analytic or source-inspection arguments; the whole manuscript is not described as machine checked.

The actual mean-weight finite/asymptotic theorem of Section 9 uses tested source

`ac60c0df885d8866329a839e2b4e73c90595d806`.

Its two own fresh Colab Linux clones 076/076b pass six exact standard-axiom reports and all nine stages. Compiled module SHA-256 is

`28fcc004795a1d8ab74ca87b5e67cdf90f557551bc02b1bb912da38324c26e50`.

Raw oracle and LF-source-manifest hashes also match; full records and the strict mutation-tested validator are preserved in the mean-weight reproduction package. The proof composes the actual character strip, its exact trace, positivity, local derivative-kernel bound and multiplier estimate; these are derived lemmas, not desired conclusions supplied as hypotheses. Nonnegative weights, fixed positive c and the quadratic penalty upper bound remain. The constructed quadratic-weight mean observation instead uses tested source

`7f37d6437a8f69a1d3fddf3dc0bcfad74c6bf9b6`.

It has two same-source fresh-clone reproductions 067/067b with six exact standard-axiom reports and matching compiled, raw-oracle and LF-source hashes. The formal output is an actual fixed finite-dimensional type; the displayed numeric dimension and nonoptimal path explanation are not additional oracle endpoints. The penalty examples and regularity tradeoff are analytic only. Classical quadratic-form, Fourier and trace mechanisms, as discussed in [7], establish no new general principle or priority here.

The signed-weight results in Section 10 have distinct exact formal anchors. The positive-part construction, actual energy order and finite/asymptotic sufficient condition use

`810785b17b2f1bb14dd6caa5435865eb51a86ef3`.

Its fresh own Colab Linux clones 077/077b pass six exact standard-axiom reports. The actual zero-signed-mean fixed-output family instead uses

`497e3b4ef4046cee25b1069b8aa0fd8c0cfaa178`.

Its clones 079/079b pass ten exact reports, including actual Hodge coercivity and the constructed family on one fixed finite output type. The individual square-cycle order, selected positive-part mean and rank consequence use

`4c5728d147721523f5a521011351ac4bf2d3cf86`.

Its clones 081/081b pass ten exact reports. In every pair the compiled module, raw oracle and LF-source-manifest hashes match, with real mutation controls under normal and optimized Python. The declaration map records exact names and hashes, and links the complete reproduction packages. Failed probes 078 and 080, their sources, full diagnostics and corrections remain preserved; no passing result is inherited across source changes.

The individual cycle has zero positive-part means in other pairs. Its formal selected-pair identity is not an all-directions counterexample. The finite sum covering all off-diagonal pairs, its displayed mean calculation, the negative hyperplane diagnostic, numerical output dimension, normalized critical-mean limit interpretation and full-state identity-observation non-vacuity check are analytic here. They are not additional oracle endpoints. No density conclusion from ordinary signed spatial means alone has been established. The code indexes positive side length by $k+1$, equivalent to the positive integer indexing used in the text.

14.1 Evidence boundary for the full volume extension

The actual normalized image and genuine compression trace are separately reproduced at source `9a8946b7`, the whole-image support and constructed positive total envelope at `88445ac6`, and the original local nonlinear finite rank composition at `3a2d5305`. The actual numerical bound and constructed uniform dilation density have two matching own Colab reproductions at `031ca274`, with five exact standard-axiom reports and full manifests. These are distinct source artifacts, not grades inherited by this manuscript.

The full integer-tail source `067fed75` now has two fresh own Colab reproductions, 094/094b, with all nine stages exiting zero and five exact standard-axiom reports. Their compiled module, raw oracle and LF-source manifest hashes match. Real mutation controls pass under both normal and optimized Python. The earlier browser loss and absent saved output are preserved separately; no outcome is reconstructed for that earlier event. The live session manager reported no active sessions before and after these fresh runs. This is execution reproducibility, not independent scientific review. The exact averaged multiplier and sinc limit remain analytic calculations. Preparation and rendering alone do not constitute a completed exact-artifact review.

Historical model assessments and review context

This candidate was prepared by the Codex agent in the author's continuing programme chat, with access to earlier objections, source audits and supplied reproduction records. The served backend model identifier and reasoning setting are not exposed, and no technical isolation or scientific independence is claimed for this preparation.

The public v2 of this record has canonical PDF SHA-256

`b37f2fe86b7a8a526ad02f8b28491c300a9ff72d5653536bc5532e8c9468ed8f`.

A subagent explicitly authorized by the author and configured as `gpt-6-astra` assessed that exact eight-page artifact after first preserving the earlier candidate's objections. Actual backend identity and reasoning setting were not exposed. The workflow used the same account, provider and programme context; it was not blind or organizationally independent. The complete earlier and corrected-artifact reports remain in the v2 package. The corrected-artifact report regarded O1–O5 as resolved and retained O6: no independent complete Lean rebuild or dependency-closure audit. No numerical score or human peer review was supplied.

The immediately preceding public artifact is AIRR v5, a 33-page Working paper released on 4 November 2026 under CC BY 4.0, with CC0 metadata. Its canonical PDF SHA-256 is

`69d802bfea8f8f23bfe6aba03cdc4835feac777e024c939040c7f099a11462c8`.

It is public and citable, but has not completed AIRR screening and is not editorially accepted or human peer reviewed. Earlier versions and their complete reports remain preserved. The author is also the AIRR founder/operator; no independent editorial decision is inferred from publication.

The historical revision136 corrected verified weight-symbol and cross-reference defects in that artifact; those changes alone were editorial. This consolidated successor additionally supplies the classical local and geometric proofs, corrects a false one-cluster-per-leaf inference, and incorporates the separately scoped actual physical embedding and exact finite penalty calculation. A newly rendered PDF will be a different artifact: its hash, visual review, assessment and any necessary deposit permission must be recorded separately. No public version number, private submission number, license grant, historical review score or acceptance is assigned by this source.

15 From necessary rank budgets to sufficient observations

The preceding sections establish necessary information budgets and their formal and analytic boundaries. The continuation below asks when actual coefficient distributions suffice to construct observations of the same order. It is integrated into this revision, not a second paper or deposit. For these sections set $D = d \geq 3$ and $g = N_c^2 - 1$, $h = Dg$. The symbol $\mu \geq 1$ denotes the scalar coefficient mean budget here; it does not replace the earlier symbol $m = N_c^2 - 1$ for internal dimension. The field is denoted U instead of A. All norms and original forward-edge occurrences use the same counting, unit-spacing flat convention. Restrictions on scalar coefficients and the distinction between spectral coordinates and rectangular means are stated explicitly below.

Historical main-text and continuation artifacts were assessed on different PDF bytes. The immediately preceding integrated v5 now has a public exact-artifact AI report, described in Section 26. That report is historical evidence for v5, not an assessment of this changed candidate. A first exact-PDF review must follow rendering and inspection of this candidate.

16 Actual objects and the original observation problem

Fix $D \geq 3$, $N_c \geq 2$, $c > 0$, and the unit-spacing torus $X_n = (\mathbb{Z}/n\mathbb{Z})^D$. The actual real adjoint coefficient space has dimension $g = N_c^2 - 1$; there are D field components, so write $h = Dg$. For the actual forward difference $d_i u(x) = u(x + e_i) - u(x)$, take one scalar total coefficient $a(x) = 1 + w_{i,x,j} > 0$ shared by all derivative directions and field components. Use counting norms and the actual componentwise flat energy

$$Q_w(U) = \mathcal{E}_a(U) = \sum_{x,i,j} a(x) \|d_i U_j(x)\|^2. \tag{14}$$

The existing flat Hodge identity identifies these forms; its formal source is listed in Section 25. All continuation adapters remain analytic. There is no substitution of a covariant or interacting derivative in (14).

The original local data consist of a differentiable map C at zero, with $C(0) = 0$, into an actual finite-dimensional real normed space F, and a penalty satisfying $\Phi(y) \leq K\|y\|^2$ near zero, such that

$$c\|U\|^2 \leq Q_w(U) + \Phi(C(U)) \tag{15}$$

for sufficiently small actual divergence-zero fields. "Local" here means a configuration-space neighbourhood; spatial support of C is an additional condition. The constant c is fixed, while observation maps and neighbourhoods may depend on n.

A useful necessary condition follows directly. If an actual J-dimensional divergence-zero subspace has $Q_w(U) \leq \lambda\|U\|^2$ with $\lambda < c$, then $\dim F \geq J$. Otherwise a nonzero U in that subspace lies in $\ker DC(0)$. Along tU, $C(tU) = o(t)$ and the penalty is at most $o(t^2)$. Dividing (15) by t^2 and taking $t \to 0$ contradicts $\lambda < c$. No rank or band certificate is supplied as an extra hypothesis in this argument.

17 The critical global moment theorem

Put $p = D/2 > 1$. For $\mu \geq 1$ suppose that

$$n^{-D} \sum_{x \in X_n} (\mu / a(x))^p \leq M < \infty. \tag{16}$$

Theorem 17.1. *Under* (16)*, there is an actual real linear spectral observation C such that*

$$c\|U\|^2 \le Q_w(U) + c\|C(U)\|^2 \quad \textit{for every actual cochain } U,$$

and

$$\dim F \le C_* \max\{1, (n/\sqrt{1+\mu})^D\}, \tag{17}$$

where, with $I_p = (p-1)^2/[p(p+1)]$ and $\Gamma_p = \frac{p}{p-1} I_p^{-(p-1)}$, one safe choice is

$$C_* = h\left[1 + \Gamma_p (3/4)^D D^{p+1} c^p M 2^p\right].$$

Thus $\Phi(y) = c\|y\|^2$, $K = c$ and any positive neighbourhood radii supply the original local data. In particular the assertion holds on the actual divergence-zero subspace.

The original subclass $1 \le a \le 2\mu$ with exact arithmetic mean μ is included, but those extra normalizations are not needed for sufficiency. The distribution condition (16) remains additional; it is not derived from an arithmetic mean.

17.1 A finite normalized Fourier lemma

For a finite abelian group of size V, let $\mathcal{F}$ be its counting-unitary Fourier matrix, whose entries have modulus $V^{-1/2}$. For a frequency multiplier g put

$$L = V^{-1} \sup_{t>0} t^p \#\{k : |g(k)|^2 > t\}.$$

For a nonnegative matrix B, write $\mathsf{n}(s, B)$ for its number of eigenvalues strictly greater than s, counted with ordinary multiplicity.

Lemma 17.2. *For $K = M_f \mathcal{F}^* M_g \mathcal{F}$ and $s > 0$,*

$$\mathsf{n}(s, K^*K) \le \Gamma_p L s^{-p} \sum_x |f(x)|^{2p}.$$

The trace and eigenvalue counts are not volume-normalized.

Proof. Remove phases by unitaries and set $A = \mathcal{F}^* M_{|g|^2} \mathcal{F}$. If $L = 0$ the assertion is immediate. Otherwise let $B_0 = pL/(p-1)$ and let A^+ be the inverse on $\operatorname{ran} A$, zero on its kernel. For $0 \le \gamma \le A$ write $\rho(x) = \langle \delta_x, \gamma \delta_x \rangle$. Positivity makes γ annihilate $\ker A$. On the range let $P_E = \mathbf{1}_{\{A^+ \le E\}}$ and $R_E = \mathbf{1}_{\{A^+ > E\}}$. The weak-frequency tail gives

$$\langle \delta_x, P_E A P_E \delta_x \rangle = V^{-1} \sum_{|g(k)|^2 \ge 1/E} |g(k)|^2 \le B_0 E^{p-1}.$$

Indeed the tail above t is t times its count plus the integral of counts above t; the resulting bound is $VLp\,t^{1-p}/(p-1)$. Non-strict thresholds follow by limits. The triangle inequality for $\gamma^{1/2}\delta_x$ split by $P_E + R_E$ yields

$$\langle \delta_x, R_E \gamma R_E \delta_x \rangle \ge (\sqrt{\rho(x)} - \sqrt{B_0} E^{(p-1)/2})_+^2.$$

Exact spectral layer integration consequently gives

$$\operatorname{Tr}(A^+ \gamma) \ge I_p B_0^{-1/(p-1)} \sum_x \rho(x)^{p/(p-1)}. \tag{18}$$

The scalar integral reduces to $\int_0^1 (1-t^{(p-1)/2})^2 dt = I_p$. Now take P to be the spectral projection of $A^{1/2} M_{|f|^2} A^{1/2}$ above s, with rank $\nu > 0$, and set $\gamma = A^{1/2} P A^{1/2}$. Then $0 \le \gamma \le A$, $\mathrm{Tr}(A^+\gamma) \le \nu$, and

$$s\nu < \sum_x |f(x)|^2 \rho(x) \le \Big(\sum_x |f(x)|^{2p}\Big)^{1/p} \Big(\sum_x \rho(x)^{p/(p-1)}\Big)^{(p-1)/p}.$$

Insert (18) and raise to p. The constant becomes $B_0 I_p^{-(p-1)} = \Gamma_p L$. This also explicitly handles a zero-frequency kernel. □

The density method is classical and follows the Rumin–Frank approach [13, Section 2]. The proof above supplies its finite normalization and kernel adapter; it is not a claim of a new Cwikel mechanism or independent discovery.

17.2 Actual lattice frequencies and the weighted inverse

For $n \ge 2$ let $L_0 = d^*d$ and work first on the actual counting-mean-zero scalar space H_0. Its Fourier symbols are

$$z_i(k) = e^{2\pi i k_i/n} - 1, \qquad \lambda(k) = 4 \sum_i \sin^2(\pi k_i/n).$$

Choose centred integer representatives r_i with $|r_i| \le n/2$, using only one representative at an even endpoint. Concavity of sine implies

$$\lambda(k) \ge 16|r|^2/n^2, \qquad |g_i(k)| := |z_i(k)/\lambda(k)| \le n/(4|r|), \quad k \ne 0.$$

Set $g_i(0) = 0$. If $|g_i| > \tau$, then $0 < |r| < R = n/(4\tau)$. For $R \le 1$ there are none, and for $R > 1$ there are at most $(3R)^D$ representatives. Thus

$$n^{-D} \sup_{t>0} t^{D/2} \#\{|g_i|^2 > t\} \le (3/4)^D. \tag{19}$$

This is uniform in n, including $n = 2$, not an imported infinite-volume bound.

Let $L_a = d^* M_a d$. On H_0 it is positive definite. For $f \in H_0$, the flow $j_* = M_a d L_a^{-1} f$ minimizes $\langle j, M_{a^{-1}} j\rangle$ among $d^*j = f$: write $j = j_* + k$, $d^*k = 0$, and the cross term is $\langle dL_a^{-1} f, k\rangle = 0$. Inserting the explicit feasible flow $j = dL_0^{-1} f$ gives the actual quadratic-form comparison

$$L_a^{-1} \le \sum_{i=1}^D T_i^* T_i, \qquad T_i = M_{a^{-1/2}} d_i L_0^{-1}. \tag{20}$$

Extend L_0^{-1} by zero on constants. Lemma 17.2 and (19) bound each $\mathsf{n}(s, T_i^* T_i)$ by

$$\Gamma_p (3/4)^D s^{-p} \sum_x a(x)^{-p}.$$

For nonnegative matrices, the count of their sum above t is no larger than the sum of their counts above t/D: intersect the complements of their high spectral subspaces and apply min–max. No commutation is required. With $t = 1/c$ in (20), inversion on H_0 and the single scalar constant mode yield

$$\#\{\lambda(L_a) < c\} \le 1 + \Gamma_p (3/4)^D D^{p+1} c^p M n^D/\mu^p. \tag{21}$$

No spectral-band dimension is assumed in this derivation.

Completion of Theorem 17.1. Choose an actual real counting-orthonormal eigenbasis of the real symmetric matrix L_a. Retain the coordinates with eigenvalue $< c$, separately in each of the h real field coordinates. Complexification used in the Fourier proof preserves real eigenvalue multiplicities; it does not double them. Energy controls coefficients at eigenvalues $\geq c$, while $c\|C(U)\|^2$ controls the rest. Equation (21), $p = D/2$, and $(1+\mu)/\mu \leq 2$ give (17). For $n = 1$ take the identity on all h coordinates. This constructs the actual global spectral observation. □

18 Every lower positive inverse exponent is insufficient

Theorem 18.1. *For every fixed $0 < r < D/2$, the analogue of Theorem 17.1 is false uniformly under the conditions*

$$\operatorname{mean}((\mu/a)^r) \leq 2, \qquad 1 \leq a \leq 2\mu, \qquad \operatorname{mean}(a) = \mu.$$

Proof. Choose a fixed integer $\ell > 3D/c$, and put $P = 2\ell + 2$, $B = (2\ell+1)^D$. The triangular window $t(s) = (1 - |s|/\ell)_+$ has finite sums

$$T = \sum t^2 = (2\ell^2+1)/(3\ell), \quad V = \sum (d^- t)^2 = 2/\ell, \quad W = \sum (d^+ d^- t)^2 = 6/\ell^2.$$

For a product window $\psi_y(x) = \prod_i t(x_i - y_i)$ and nonzero actual adjoint H, the actual curl $U_1 = d_2^- \psi_y H$, $U_2 = -d_1^- \psi_y H$, other components zero, has backward divergence zero. Its field support and all nonzero forward-gradient start sites lie in the box $y + [-\ell, \ell]^D$. Direct separable sums give

$$\|U\|^2 = 2VT^{D-1}\|H\|^2, \qquad \frac{\mathcal{E}_0(U)}{\|U\|^2} = \frac{W}{V} + (D-1)\frac{V}{T} = \frac{3}{\ell} + \frac{6(D-1)}{2\ell^2+1} \leq \frac{3D}{\ell} < c.$$

Take sufficiently large multiples n of P, set $\mu = n^2$ and $J = \lfloor n^{D-2r}/B \rfloor$. These J boxes fit on the separated P-grid eventually, since $P^D \leq Bn^{2r}$. They satisfy $J \geq n^{D-2r}/(2B) \to \infty$. Put $s = BJ/n^D \leq \mu^{-r}$ and set $a = 1$ on their union, $a = A = (\mu - s)/(1 - s)$ elsewhere. Eventually $s \leq 1/2$, so

$$\operatorname{mean}(a) = \mu, \quad 1 \leq a \leq 2\mu, \quad A \geq \mu, \qquad \operatorname{mean}((\mu/a)^r) \leq s\mu^r + 1 \leq 2.$$

Each curl's gradients start only in its weak box. The actual J-dimensional span has that same Rayleigh quotient, because both field and gradient supports are disjoint. Section 16 forces $\dim F \geq J$, whereas the target maximum at $\mu = n^2$ is 1. This proves the refutation using original nonlinear data. □

Together the two theorems give the sharp inverse-moment exponent for this specified family of global scalar conditions. They do not identify all sufficient coefficient criteria, prove a physical coefficient hypothesis, or impose spatial locality on C.

19 Spatial means and the critical dyadic obstruction

For a partition $\mathcal{P}$ into integer rectangles, its actual normalized means are

$$C_{\mathcal{P}}(U)_{R,j} = \sqrt{|R|}\operatorname{mean}_R U_j, \qquad \dim F = h\#\mathcal{P}.$$

A full dyadic cube tree splits each selected cube into all 2^D equal children and retains unsplit siblings. The origin of the torus coordinate cube may be translated. This class excludes arbitrary anisotropic rectangle cuts and compressed observations.

Proposition 19.1. *There are fixed critical-moment profiles with target maximum* 1 *for which every coercive full dyadic cube-tree mean partition has* $\Omega(\log n)$ *leaves, even after choosing a translated origin. The same profiles admit an explicit constant-size rectangular mean partition satisfying the global inequality for every actual cochain.*

Proof. Keep $\ell > 3D/c$ fixed. Choose a dyadic $L = 2^q \geq 4\ell + 2$, put $\gamma = \max(1, c/4)$, $n = 2^J$ with $J \geq q + 1$, $\mu = \gamma n^2$, and $S = \{0, \dots, L-1\}^D$. Set $a = 1$ on S and $a = A = (\mu - s)/(1 - s)$ elsewhere, $s = L^D/n^D$. Then $1 \leq a \leq 2\mu$, its mean is exactly μ, and

$$\operatorname{mean}((\mu/a)^{D/2}) \leq 1 + L^D \gamma^{D/2} =: M_0.$$

All these parameters except n are fixed, and the target maximum is 1.

In the hypothetical complete level $j_* = J - q$ of any translated dyadic grid, cells have side L. Each coordinate interval of S meets at most two such cyclic grid intervals, so one intersection has integer length at least $L/2 \geq 2\ell + 1$. Taking these intersections independently places a whole gradient-start box for a compact curl inside S and inside one level-j_* cell. If the proposed partition had maximum depth at most j_*, that cell would lie in one actual leaf. The curl then has zero means on every leaf, since each component has total sum zero, yet its weighted Rayleigh quotient is $< c$. This contradicts coercivity. A single zero-extended window needs no packing divisibility assumption on n; $n \geq 2L$ suffices here. The field may depend on the partition, since coercivity is asserted for every field. Thus maximum depth is at least $J - q + 1$, and the full-arity tree identity gives

$$\#\mathcal{P} \geq 1 + (2^D - 1)(J - q + 1). \tag{22}$$

For a successful unrestricted partition, isolate all L^D vertices of S as singletons, and use the D disjoint complement rectangles

$$R_i = [0, L-1]^{i-1} \times [L, n-1] \times [0, n-1]^{D-i}.$$

Coverage follows from the unique first coordinate outside $[0, L-1]$. Every internal edge starts at a strong vertex with coefficient A. On a length-l interval, counting variance and path Cauchy–Schwarz give

$$\sum \|u - \bar{u}\|^2 \leq \frac{l(l-1)}{4A} \sum A \|du\|^2.$$

Indeed the path resistance is at most $(l-1)/A$, and a cut at t is crossed by $t(l-t) \leq l^2/4$ pairs in the variance identity. Product-variance tensorization on original fibres yields a rectangle constant $n^2/(4A)$, with no extra factor D. Singletons have zero residual, internal edge sets are disjoint, and omitted edges have positive energy. Since $cn^2/(4A) \leq c/(4\gamma) \leq 1$, exact orthogonal decomposition proves

$$c\|U\|^2 \leq Q_w(U) + c\|C_{\mathcal{P}}(U)\|^2, \qquad \#\mathcal{P} = L^D + D.$$

These rectangles have potentially unbounded aspect ratio; they leave the cube-tree class. □

The general critical moment question for unrestricted rectangular means remains open. Changing only the stopping rule within full cube trees cannot resolve it. Classical weighted Sobolev inequalities [14] also yield block observations for inverse moments of finite order strictly greater than $D/2$; the dossier contains the explicit balanced-box proof. Proposition 19.1 explains why its critical tree-counting loss cannot be removed merely by choosing a better cube-tree stopping criterion.

20 What a spatial observation proof must still establish

The preceding positive spectral result does not construct a rectangular mean map. Three further checks prevent substitutions for that missing construction. They concern local geometry, the penalty coefficient, and support of observations, respectively. They are collected here as limitations of the same dossier, not as additional papers.

20.1 Critical inverse mass does not remove rectangle geometry

For a rectangle R of side lengths l_i, set $V = \prod_i l_i$, $L = \max_i l_i$, $\kappa = L^2/V^{2/D}$ and $\eta_R = \sum_R a^{-D/2}$. The endpoint Sobolev estimate on balanced cubes, with $q = 2D/(D+2)$, transfers to arbitrary rectangles by repeated coordinate reflection and gives

$$\sum_R |u - \bar{u}_R|^2 \le A_D D^{2/D} \kappa \eta_R^{2/D} \mathcal{E}_{a,R}(u). \tag{23}$$

Here the energy uses original internal edges and their actual starting-vertex weights. The full reflection multiplicities and graph adapters to the classical estimate [14, Section 2.2] are proved here in Section 22; this is an analytic adapter, not a new Lean theorem.

The factor cannot be discarded even for uniform coefficients. With $u(x) = x_i - (L-1)/2$ along a longest side, direct sums give

$$\sum_R u^2 = V(L^2-1)/12, \qquad \mathcal{E}_{a,R}(u) = aV(L-1)/L.$$

Their ratio is $L(L+1)/(12a) \ge \kappa \eta_R^{2/D}/12$. On the thin rectangle of lengths $(L, 1, \dots, 1)$, choosing $a = L^{2/D}$ makes $\eta_R = 1$ while this ratio diverges. Thus compressing a cube-tree chain into thin rectangles cannot simply reuse a shape-independent inverse-mass certificate. This prescribed-block test does not defeat every global partition: omitted crossing edges and a different partition may help. The successful partition in Proposition 19.1 is not contradicted.

20.2 Fixed and free penalties are different obligations

Let C be normalized means on a fixed partition of a finite real Euclidean space. Then $CC^* = I$; $P = C^*C$ is the orthogonal block-mean projection and $Q = I - P$. For the actual positive stiffness L_a, write the blocks on $\operatorname{ran} P \oplus \operatorname{ran} Q$ as

$$A = PL_aP, \qquad B = PL_aQ, \qquad D_0 = QL_aQ.$$

Positivity implies $\ker A \subseteq \ker B^*$: otherwise the energy of $tp + v$ would have a nonzero linear coefficient with zero quadratic coefficient in t. Completing the square with the finite spectral pseudoinverse $A^\dagger$ shows

$$L_a + cP \ge cI \quad \Longleftrightarrow \quad D_0 - B^*A^\dagger B \ge cI \text{ on } \operatorname{ran} Q. \tag{24}$$

Indeed the first inequality is exactly $L_a - cQ \geq 0$, and its quadratic form is $\langle p + A^\dagger Bv, A(p + A^\dagger Bv)\rangle + \langle v, (D_0 - cI - B^* A^\dagger B)v\rangle$. This is classical shorted-operator algebra [12]; it supplies no spatial partition or target count. In particular $D_0 \geq cI$ alone is insufficient.

Proposition 20.1. *On the four-cycle with edge weights* $(1, 6, 4, 1)$ *and blocks* $\{0, 1\}, \{2, 3\}$*, coercivity holds on zero block sums with* $c = 19/5$*, but* $L_a + cP \geq cI$ *fails.*

Proof. For $v = (x, -x, y, -y)$, with the counting Euclidean norm,

$$\mathcal{E}_a(v) - \frac{19}{5}\|v\|^2 = \frac{17}{5}\left(x + \frac{35}{17}y\right)^2 + \frac{84}{85}y^2 \geq 0.$$

The actual normalized observation is

$$Cu = \big((u_0 + u_1)/\sqrt{2},\ (u_2 + u_3)/\sqrt{2}\big).$$

For $u = (19, -9, -5, -5)/14$, direct exact evaluation gives

$$\mathcal{E}_a(u) = 52/7, \quad \|u\|^2 = 123/49, \quad \|Cu\|^2 = 25/49,$$

and hence $\mathcal{E}_a(u) + (19/5)\|Cu\|^2 - (19/5)\|u\|^2 = -6/35$. The weights have mean 3 and lie between 1 and 6. Ten focal Lean theorems preserve the constructed weights, normalized observations, projection identities, kernel square identity and exact negative witness; Section 25 identifies the successful immutable source. Those ten endpoints concern the scalar four-cycle. Section 24 separately derives its actual physical embedding and the all-field kernel premise; only its stated fixed-penalty endpoints have paired formal evidence. □

For the original class allowing an arbitrary finite K, put $J = D_0 - cI$. The necessary and sufficient condition is instead

$$\exists K \geq 0 :\ L_a + KP \geq cI \quad \Longleftrightarrow \quad J \geq 0 \text{ and } B \ker J = 0. \tag{25}$$

Necessity follows by testing $v \in \ker C$, then $v + tp$ with $v \in \ker J$. For sufficiency, $B \ker J = 0$ implies $B^* p \in \operatorname{ran} J$. Completing the square in $v + J^\dagger B^* p$ leaves $A + (K - c)I - BJ^\dagger B^*$, so $K = c + \|BJ^\dagger B^*\|$ suffices. In particular a proved all-cochain margin $D_0 \geq (c + \epsilon)I$ gives $K = c + \|L_a\|^2/\epsilon$. For actual $a \leq 2\mu$, the elementary edge estimate $|r - s|^2 \leq 2|r|^2 + 2|s|^2$ gives $\|L_a\| \leq 8D\mu$, including the parallel edge occurrences at $n = 2$. A margin $2c$ therefore permits $K = c + 64D^2\mu^2/c$, without increasing the number of observations. No margin or partition is proved by this conditional implication. Ordinary block averaging need not preserve divergence, so an estimate only on the physical kernel cannot silently replace the all-cochain hypothesis in this decomposition.

Even free K needs care at equality. On the same cycle let $s = \sqrt{85}$, $c_* = (17 - s)/2$ and $v = (7, -7, 6 - s, s - 6)$. Its kernel defect at c_* is zero, since for general $(x, -x, y, -y)$ it equals $(s - 6)[x + 7y/(s - 6)]^2$. For $p = (\delta/2, \delta/2, -\delta/2, -\delta/2)$ the full defect is

$$\mathcal{E}_a(v + p) + K\|C(v + p)\|^2 - c_*\|v + p\|^2 = (7 + K - c_*)\delta^2 - 10(13 - s)\delta.$$

For each finite $K \geq 0$ a sufficiently small positive δ makes it negative. This additional equality example is exact analytic algebra and finite diagnostic evidence, not one of the ten Lean theorems of Proposition 20.1.

20.3 Sparse point samples cannot replace stable averages

For uniform $a = \mu$, if an observation depends only on vertex values on a set S, $|S| = s$, and $C(0) = 0$, then the scalar field $u = 1$ off S and 0 on S has zero observation. At most $2Ds$ forward-edge occurrences cross S, so $\mathcal{E}_\mu(u) \le 2D\mu s$. Consequently all-cochain coercivity requires

$$s \ge \frac{cn^D}{2D\mu + c}.$$

For an actual divergence-zero test, fix a component direction j and let the field be constant along every j-line, equal to zero on each line meeting S and one on the other lines, in one fixed internal coordinate. There are at most s hit lines. Its backward divergence is zero, its norm squared is at least $n(n^{D-1} - s)$, and transverse crossing energy is at most $2(D-1)\mu ns$. Thus necessarily

$$s \ge \frac{cn^{D-1}}{2(D-1)\mu + c}.$$

If every transverse line is hit, the second bound is immediate without a nonzero test. These bounds give the required obstruction: at $a = \mu = n^2$ and critical moment $M = 1$, bounded samples fail for all cochains when $D \ge 3$ and for these physical tests when $D \ge 4$. For the same uniform coefficients one global mean succeeds for $0 < c \le 16$, by the nonzero Fourier gap $n^2\lambda(k) \ge 16$. This comparison rules out a point-sampling shortcut, not rectangular averages or singletons supplemented by complementary block means. A nonlinear local observation is tested on sufficiently small multiples of these fields.

21 Constructive rectangular bounds and their limits

Retain $\mu \ge 1$, $1 \le a \le 2\mu$, exact mean μ, (16), and $p = D/2$. All assertions in this section remain analytic.

21.1 A general threshold construction

For an interval of length l, the identity $\mathrm{Var}(u) = l^{-1}\sum_{x<y}|u(x) - u(y)|^2$ and path Cauchy–Schwarz give coefficient at most $l(l-1)/4$ per edge: each crossing pair has distance at most $l - 1$, and at most $l^2/4$ pairs cross any edge. Product variance tensorization therefore yields, for a rectangle of maximal side L,

$$\sum_R |u - \mathrm{mean}_R\, u|^2 \le \frac{L(L-1)}{4} \sum_{\text{internal forward edges of } R} |du|^2. \tag{26}$$

There is no additional factor D in this tensorization.

Fix $\tau \ge 1$. If $\tau < c/4$ use singletons. Otherwise set $\ell = \min\{n, \lfloor 2\sqrt{\tau/c}\rfloor\}$ and start with a grid whose sides have length at most ℓ. Isolate each vertex with $a < \tau$ as a singleton. Its current rectangle minus that vertex splits into at most $2D$ rectangles by the first coordinate differing from that vertex. Delete empty pieces and preserve earlier singletons. Remaining rectangles have $a \ge \tau$ and side at most ℓ. Equation (26), nonnegative omitted edges and normalized-mean orthogonality give

$$c\|U\|^2 \le \mathcal{E}_a(U) + c\|C_{\mathcal{P}}U\|^2, \quad (C_{\mathcal{P}}U)_{R,j} = \sqrt{|R|}\,\mathrm{mean}_R\, U_j \tag{27}$$

for ALL cochains. Singleton residuals vanish, including $n = 1$; parallel edge occurrences for $n = 2$ are counted separately, and wrap edges can be omitted positively. With $B_c = (1 + \sqrt{c})^D$,

$$|\mathcal{P}| \le B_c \max\{1, n^D/\tau^p\} + 2D\#\{a < \tau\}, \quad \#\{a < \tau\} \le Mn^D\tau^p/\mu^p.$$

The singleton branch satisfies this count since $B_c\tau^{-p} > 1$ there. Taking $\tau = \sqrt{\mu}$ yields the weaker general bound

$$\dim F \le h(B_c + 2DM) \max\{1, n^D/\mu^{D/4}\}.$$

21.2 Every global threshold can lose the target order

Let $\gamma = \max\{1, c/4\}$, $n = t^2$, $t \ge 2$, $\mu = \gamma n^2$, $b = \gamma n$, $S = \{0, \dots, t-1\}^D$, and $f = n^{-p}$. Set $a = b$ on S, $a = A = (\mu - fb)/(1 - f)$ outside. Then $1 \le a \le 2\mu$, the mean is μ, and the critical moment is less than $2 - f$. The target maximum is one. For any global $\tau \ge 1$, consider the preceding algorithm with any translated initial grid or excision order, retaining grid boundaries and every excised singleton. If $\tau \le b$, no vertex is weak and the grid volume bound gives

$$|\mathcal{P}| \ge n^D/(2\sqrt{b/c})^D = (c/(4\gamma))^p n^{D/2},$$

or the singleton branch is larger. For large n the grid side is less than n. If $\tau > b$, all $t^D = n^{D/2}$ weak vertices remain distinct singletons. Every such threshold therefore loses the target order.

For the SAME profile retain the whole weak cube and split its complement by the first coordinate outside it, giving $D + 1$ rectangles. On S, $ct^2/(4b) = c/(4\gamma) \le 1$; outside, $A \ge \mu$ and $cn^2/(4A) \le 1$. Equation (26) proves (27) for this constant-size partition. This is a mechanism obstruction, not a counterexample to every rectangular partition.

21.3 Supplied scaled cubes attain the target order

Suppose additionally that disjoint integer cubes S_k of side L_k, nonwrapping in one common root, satisfy for fixed $\beta \ge 1$

$$\begin{aligned} cL_k(L_k - 1)/4 \le a(x) \le \beta L_k^2 \quad (x \in S_k), \\ a(x) \ge cn(n-1)/4 \quad (x \notin \cup_k S_k). \end{aligned} \tag{28}$$

These are extra coefficient and geometric assumptions. Each cluster charges at least β^{-p} to $\sum_x a(x)^{-p}$; thus $J \le \beta^p Mn^D/\mu^p$.

Section 23 supplies a self-contained construction for precisely these disjoint integer cubes. Use quarter-offset solid cubes, then a compressed corner-potential subdivision. Its at most $2^{D+1}(2^D - 1)J$ terminal rectangles have no cube corner in their interiors. This does not imply that a terminal rectangle meets only one cube: the proper-intersection and covering bounds give at most $T_D = 2^{2D-1}$ cubes per terminal rectangle. Refining by all their facet coordinates gives at most $(2T_D + 1)^D$ subrectangles per terminal rectangle. Consistent boundary assignment followed by taking integer points yields an actual partition; every surviving rectangle is inside one supplied cube or outside all supplied cubes. Fragmentation is permitted. Consequently, with the explicit dimension-only constant

$$A_D = 2^{D+1}(2^D - 1)(2^{2D} + 1)^D, \qquad |\mathcal{P}| \le A_D \max\{1, J\},$$

$$\dim F \le hA_D(1 + 2^p\beta^p M) \max\{1, n^D/(1 + \mu)^p\}.$$

The earlier one-cluster-per-leaf inference is not used. The multiple-object stress family and the boundary adapter are included in the cited section. Inside a cluster fragment maximal side is at most L_k, outside at most n. Equations (28) and (26) prove (27). Projection separation and a high-budget restriction are unnecessary. This uses classical geometry; it does not derive the supplied cubes from the original moment.

21.4 The exact cube decomposition is not universal

For fixed c, β take arbitrarily large integer $n \geq 3$, $\mu = n$, $W = \{e_1, -e_1\}$, $f = 2/n^D$, with $a = 1$ on W and $a = A = (n-f)/(1-f)$ elsewhere. Then $n < A \leq 2n$, mean $a = n$, and the critical moment is less than $2n^{-p} + 1 - f \leq 2$. Choose also $n > \max\{\beta, 8\beta/c, 1 + 8/c\}$ and $n \geq c$. Every vertex has $a < cn(n-1)/4$, so any family satisfying (28) must cover zero. A side-one cube containing zero would require $A \leq \beta < n$. A cube of side $L \geq 2$ contains at least one of $e_1, -e_1$, in every translated root. Its lower bound forces $L^2 \leq 8/c$, while its upper bound at zero requires $A \leq \beta L^2 \leq 8\beta/c < n < A$. Even allowing overlap cannot repair this exact decomposition.

The SAME profile nevertheless admits a strong grid with $\ell = \min\{n, \lfloor 2\sqrt{A/c} \rfloor\}$ and two singleton excisions. Equation (27) holds and

$$\dim F \leq h\big[2^p(1+\sqrt{c})^D + 4D\big] \max\{1, n^D/(1+\mu)^p\}.$$

There are at most $\lceil n/\ell \rceil^D + 4D$ blocks. If $\ell < n$, $\ell \geq \sqrt{A/c}$; if $\ell = n$, there is one grid box. Remaining rectangles have coefficient A and weak residuals are singletons. The universal supplied-cube route is refuted. The original unrestricted rectangular question is neither proved nor refuted.

Self-contained proof continuations. The next three sections complete the local and supplied-geometry arguments and exhibit one actual physical observation map. The constants denoted A_D are local to their sections: the Sobolev and partition constants differ. Counting measure and unit lattice spacing are used throughout. These continuations do not derive supplied geometry from the inverse moment.

22 A finite proof of the critical rectangle estimate

This supplement derives the balanced-cube endpoint used by the rectangle adapter, rather than importing a local Sobolev estimate as an external input. The method is classical: finite entropy proves the projection inequality, which yields relative edge isoperimetry, layer sets, and a power argument. We claim neither a new general Sobolev principle nor a solution of the open critical rectangular-partition problem. This is part of the existing dossier.

Fix $D \geq 3$, and put

$$p = D/2, \quad q = \frac{2D}{D+2}, \quad q' = \frac{2D}{D-2},$$
$$\alpha = \frac{2(D-1)}{D}, \quad c_D = D(1 - 2^{-1/D}), \quad K_D = \frac{2\alpha D^{1/q'}}{c_D}.$$

An internal axis edge is an unordered adjacent pair, counted once. For its positive-axis orientation $e = (x, x + e_i)$, write $du(e) = u(x + e_i) - u(x)$.

22.1 Finite projection inequality, with proof

For a nonempty finite $A \subset \prod_{j=1}^{D} X_j$, let π_i delete coordinate i. Then

$$|A|^{D-1} \le \prod_{i=1}^{D} |\pi_i A|. \tag{29}$$

Here is a finite proof, including the entropy facts needed for this use. For a finite probability vector r, define $H(r) = -\sum r \log r$ with $0 \log 0 = 0$. Concavity follows from $(-t \log t)'' = -1/t < 0$ and continuity at zero. Consequently mixing conditional distributions cannot decrease their average entropy: $H(Y \mid Z) \ge H(Y \mid Z, W)$. Factoring finite joint probabilities proves the chain rule $H(X) = \sum_j H(X_j \mid X_1, \dots, X_{j-1})$. Also $H(Y) \le \log |\operatorname{supp} Y|$: Jensen's inequality for log gives $\sum_y r_y \log(1/r_y) \le \log \sum_y r_y / r_y$ on the positive support. For X uniform on A, $H(X) = \log |A|$, while

$$\begin{aligned} \sum_i H(X_{-i}) &= \sum_i \sum_{j \ne i} H(X_j \mid X_{<j} \setminus X_i) \\ &\ge \sum_i \sum_{j \ne i} H(X_j \mid X_{<j}) = (D-1) H(X). \end{aligned}$$

The notation in the conditioning set means omit X_i only when $i < j$. Bounding each $H(X_{-i})$ by $\log |\pi_i A|$ and exponentiating proves (29). No projection inequality is an unproved hypothesis in the sequel.

22.2 Relative isoperimetry on a finite cube

Let $Q_s = \{0, \dots, s-1\}^D$, $s \ge 2$, and let $\partial_E A$ consist of internal cube edges crossing A and its complement. If $0 < |A| \le |Q_s|/2$, then

$$|\partial_E A| \ge c_D |A|^{(D-1)/D}. \tag{30}$$

Indeed, put $a = |A|$, $P_i = |\pi_i A|$. At most a/s occupied fibers in direction i are full. Every occupied nonfull fiber contains an internal crossing edge, because the fiber is a connected interval. Edges in distinct fibers/directions are different. Thus projection inequality and AM–GM give

$$|\partial_E A| \ge \sum_i (P_i - a/s) \ge D a^{(D-1)/D} - Da/s \ge D(1 - 2^{-1/D}) a^{(D-1)/D}.$$

Empty A satisfies the same inequality trivially. The boundary is the EDGE boundary, not the vertex boundary in a separate graph theorem.

22.3 Endpoint counting-norm Sobolev–Poincare inequality

For every real u on Q_s, including $s = 1$,

$$\sum_{Q_s} |u - \bar{u}_{Q_s}|^2 \le K_D^2 \left(\sum_{e \subset Q_s} |du(e)|^q \right)^{2/q}. \tag{31}$$

For $s = 1$ both sides vanish. For a nonnegative v supported on at most half the cube, layer-set decomposition and the triangle inequality in $\ell^{D/(D-1)}$, followed by (30), imply

$$\|v\|_{D/(D-1)} \le \int_0^\infty |\{v > t\}|^{(D-1)/D} \, dt \le c_D^{-1} \int_0^\infty |\partial_E \{v > t\}| \, dt = c_D^{-1} \sum_e |dv(e)|.$$

The integrals are finite sums of step lengths between values of v; the last equality follows edgewise from the length of the interval between its two endpoint values. Apply this inequality to v^α. For $a, b \geq 0$, the fundamental theorem of calculus gives

$$|a^\alpha - b^\alpha| \leq \alpha |a-b|(a^{\alpha-1} + b^{\alpha-1}).$$

Since $\alpha D/(D-1) = 2$ and $(\alpha-1)q' = 2$, Holder and the degree bound $\deg_{Q_s}(x) \leq 2D$ give

$$\begin{aligned}\|v\|_2^\alpha &\leq \frac{\alpha}{c_D}\Big(\sum_e |dv(e)|^q\Big)^{1/q}\Big(\sum_{e=\{x,y\}} (v(x)^{\alpha-1} + v(y)^{\alpha-1})^{q'}\Big)^{1/q'} \\ &\leq \frac{2\alpha D^{1/q'}}{c_D}\Big(\sum_e |dv(e)|^q\Big)^{1/q}\|v\|_2^{\alpha-1}.\end{aligned}$$

The second inequality uses $(a+b)^{q'} \leq 2^{q'-1}(a^{q'} + b^{q'})$ and $\sum_e (v(x)^2 + v(y)^2) \leq 2D\|v\|_2^2$. If $v = 0$ the conclusion is immediate; otherwise division proves $\|v\|_2 \leq K_D \|dv\|_q$.

Choose a counting median m of u and set $v_+ = (u-m)_+$, $v_- = (m-u)_+$. Both supports have size at most half the cube. For each edge,

$$|dv_+(e)|^q + |dv_-(e)|^q \leq |du(e)|^q.$$

For endpoint values on the same side of m this is equality in the nonzero part; on opposite sides it is $a^q + b^q \leq (a+b)^q$. Since $2/q \geq 1$, addition of the two squared estimates gives

$$\|u - \bar{u}_{Q_s}\|_2^2 \leq \|u-m\|_2^2 \leq K_D^2(\|dv_+\|_q^q + \|dv_-\|_q^q)^{2/q} \leq K_D^2\|du\|_q^2.$$

This proves (31) without an external graph, ball-adapter, coefficient, or boundary-condition hypothesis.

22.4 Reflection with all aspect-ratio multiplicities

Let $R = \prod_i \{0, \dots, l_i - 1\}$, $V = \prod_i l_i$, $L = \max_i l_i$, and $\kappa = L^2/V^{2/D}$. If $V = 1$ the desired variance bound is zero. Otherwise $L \geq 2$. Put $s = 2L+1$, $r = s^D/V$ and fold Q_s coordinatewise onto R by the endpoint-repeating pattern of period $2l_i$,

$$0, 1, \dots, l_i - 1, l_i - 1, \dots, 1, 0, 0, 1, \dots.$$

For $l_i = 1$ the fold is constant. Since $s/(2l_i) > 1$, each original coordinate value has at least $2\lfloor s/(2l_i)\rfloor \geq s/(2l_i)$ preimages and at most $s/l_i + 2 \leq 3s/l_i$ preimages. An original internal edge has at most the same upper count in its own direction; each full period traverses it twice. Adjacent folded values either coincide or are an original internal edge, possibly with reversed orientation. Thus for the unweighted extension $\widetilde{u}$,

$$\mathrm{Var}_{Q_s}(\widetilde{u}) \geq 2^{-D} r \,\mathrm{Var}_R(u), \qquad \sum_{e \subset Q_s} |d\widetilde{u}(e)|^q \leq 3^D r \sum_{e \subset R} |du(e)|^q.$$

The variance inequality holds by applying the preimage lower bound at the mean of $\widetilde{u}$, and then minimizing over constants on R. Apply (31). Since $2/q - 1 = 2/D$ and $s \leq 3L$,

$$\mathrm{Var}_R(u) \leq A_D \kappa \left(\sum_{e \subset R} |du(e)|^q\right)^{2/q}, \qquad A_D = K_D^2 2^D 3^{2D/q+2}. \tag{32}$$

The factor r is not bounded independently of aspect ratio; it was kept until it became $s^2/V^{2/D}$. Singleton coordinate sides cause no problem.

22.5 Positive original edge weights and field coordinates

For arbitrary positive ORIGINAL edge coefficients b_e, Holder on R with exponents $2/q$ and $2/(2-q)$ gives

$$\Big(\sum_e |du|^q\Big)^{2/q} \le \Big(\sum_e b_e|du|^2\Big)\Big(\sum_e b_e^{-q/(2-q)}\Big)^{(2-q)/q}.$$

Here $q/(2-q) = D/2 = p$. Consequently

$$\mathrm{Var}_R(u) \le A_D\kappa\Big(\sum_{e\subset R} b_e^{-D/2}\Big)^{2/D}\sum_{e\subset R} b_e|du(e)|^2. \tag{33}$$

No reflected weights were introduced: reflection was completed before this Holder step. For $b_{(x,x+e_i)} = a(x) > 0$, each vertex supplies at most D internal starting edges. Thus the inverse edge mass is at most $D\sum_R a^{-D/2}$, proving exactly the scalar-start estimate

$$\mathrm{Var}_R(u) \le A_D D^{2/D}\kappa\Big(\sum_R a(x)^{-D/2}\Big)^{2/D}\mathcal{E}_{a,R}(u).$$

For any fixed finite real internal/field dimension, apply this scalar inequality to each coordinate and sum; its coefficient factor is common. There is no additional internal-dimension factor and no replacement of original edges by enlarged-region energy.

22.6 Necessity of geometry and status

For uniform $a > 0$ and a longest coordinate of length $L \ge 2$, put $u(x) = x_i - (L-1)/2$. Direct sums give

$$\mathrm{Var}_R(u) = V(L^2-1)/12, \qquad \mathcal{E}_{a,R}(u) = aV(L-1)/L,$$

so their ratio is $L(L+1)/(12a) \ge L^2/(12a)$. Since $\kappa(\sum_R a^{-D/2})^{2/D} = L^2/a$, the geometric dependence cannot be uniformly dropped even for uniform weights. This is not a global counterexample to a partition algorithm or to the physical divergence-zero observation problem.

For attribution and comparison see S. Andres, J.-D. Deuschel and M. Slowik, *Harnack inequalities on weighted graphs and some applications to the random conductance model*, arXiv:1312.5473v6, Section 2.2, (2.5)–(2.8). Their local normalized inequalities are classical prior context; this proof does not import their theorem, graph assumptions or Harnack conclusions. The finite projection/entropy, coarea and power steps above are classical, not a new general method. The gain is a complete finite derivation of the specific input previously left external. No Lean execution, independent review, original-priority claim, new physical hypothesis removal or new AIRR paper follows from this supplement. The frozen PDF candidate and its actual reports/permissions remain unchanged; integration would create a new exact artifact and must not inherit their assessments.

23 A finite rectangular partition of disjoint lattice cubes

The linear partition for supplied disjoint cubes can be derived directly, without importing a general fat-object or low-density theorem. We use the classical compressed cube-splitting strategy of de Berg, UU-CS-1998-12, Section 2.1 and Lemma 2.1, and prove the required special-case

geometry and all discrete/energy adapters below. This is not an original general BSP method. No time-complexity, measured implementation or Lean claim is made.

Fix $D \geq 3$ and an integer $n \geq 1$. In the common root $X = \{0, \ldots, n-1\}^D$, take pairwise disjoint nonempty integer cubes

$$S_k = \prod_i \{b_{ki}, \ldots, b_{ki} + L_k - 1\}, \qquad 1 \leq k \leq J.$$

Set

$$Q_k = \prod_i [b_{ki} - 1/4, b_{ki} + L_k - 3/4].$$

These are positive-side closed solid cubes, including $L_k = 1$. They are strictly disjoint: disjoint integer intervals in a separating coordinate leave a real gap at least $1/2$. Their integer points are exactly S_k, and each such point is interior to Q_k. All corners of the Q_k lie strictly inside $[-1, n]^D$.

23.1 Proper-intersection corner bound

Let C be any axis-aligned cube of side $s > 0$ with no corner of any Q_k in its interior. Then at most 2^D of the Q_k meet $\operatorname{int} C$. Indeed, if Q_k of side $a < s$ properly meets C, each of its coordinate intervals has an endpoint in the interior of the corresponding interval of C. Choosing these endpoints gives a corner of Q_k in $\operatorname{int} C$, a contradiction. Thus every properly meeting Q_k has $a \geq s$. In each coordinate its interval then contains at least one endpoint of the interval of C; otherwise it would lie strictly inside that shorter interval. Hence Q_k contains a corner of C. Strict disjointness makes these corner choices injective, proving the bound.

The qualifier "properly" is essential. Closed boundary-only contact must not be counted here. For $C = [0, 1]^D$ and any integer $K \geq 1$, put $\epsilon = 1/(4K + 4)$ and

$$Q_j = [-\epsilon, 0] \times [(2j+1)\epsilon, (2j+2)\epsilon] \times [\epsilon, 2\epsilon]^{D-2}, \qquad 0 \leq j < K.$$

These K strictly disjoint small cubes touch C but none meets its interior or has a corner in its interior. Thus the corresponding closed-contact bound would be false. This boundary guard does not assert that de Berg's theorem is false; our construction below explicitly uses open-cell intersections and interior lattice points.

23.2 A multiple-object stress family

The corner-free property does not imply that a query cube meets only one input cube. This can fail even for the quarter-offset solids of equal-side integer cubes, not merely for arbitrary real objects. For $D \geq 2$, put $n = 4$, $C = [1, 2]^D$, and, for $\sigma \in \{0, 1\}^{D-1}$, set

$$b_\sigma = (2\sigma_1, \ldots, 2\sigma_{D-1}, 1), \quad L_\sigma = 2, \qquad S_\sigma = \prod_{i=1}^{D} \{b_{\sigma i}, b_{\sigma i} + 1\}.$$

The 2^{D-1} integer cubes are pairwise disjoint: distinct binary words have a coordinate with the disjoint integer intervals $\{0, 1\}$ and $\{2, 3\}$. Their solids have, in that coordinate, intervals $[-1/4, 5/4]$ and $[7/4, 13/4]$, with a gap $1/2$. In the last coordinate every solid has interval $[3/4, 9/4]$. Consequently every solid properly meets C, while no solid corner is interior to C: every such corner has last coordinate $3/4$ or $9/4$. Thus a universal proper-occupancy bound for

corner-free query cubes in this exact lattice class must be at least 2^{D-1}. This neither proves sharpness of 2^D nor asserts that this particular C occurs as a leaf of our root-based recursion. It rules out replacing the occupancy argument by a one-object assumption. The finite unit139 baseline leaves all had occupancy at most one, so those runs did not test this phenomenon.

23.3 Constructing linearly many vertex-free rectangles

Let V be the set of all corners of the Q_k, so $N = |V| = 2^D J$. For $J = 0$ use the single root rectangle and skip this construction. For $J > 0$ begin with the cube $[-1, n]^D$. At a current cube C keep exactly $V_C = V \cap \operatorname{int} C$. If V_C is empty, stop at that cube. Otherwise consider its 2^D open half-side children.

If the points are not all in one open child, perform all midpoint cuts and recurse on the nonempty child point sets. Points on the cuts are removed from subsequent interior sets. If all points lie in one open child, reflect and translate coordinates so the common corner is zero and $C = [0, s]^D$. Put

$$a = \max\{x_i : x \in V_C,\ 1 \le i \le D\}, \qquad C' = [0, a]^D.$$

Then $0 < a < s/2$. Successively cut at the D upper facet planes of C', producing C' and the D boxes

$$E_i = [0, a]^{i-1} \times [a, s] \times [0, s]^{D-i}, \qquad 1 \le i \le D.$$

Use disjoint boundary ownership for these boxes. Only C' can contain points of V in its interior, and at least one point of V_C lies on its boundary. Stop at the other boxes and recurse inside C' if needed. Every intermediate or terminal cell is full-dimensional; a boundary ownership convention never creates a separate hyperplane cell.

Here is a direct finite termination and size proof. Among active cells, let $n_C = |V_C| > 0$ and use the nonnegative integer potential

$$\Psi = \sum_{C \text{ active}} (2n_C - 1).$$

Initially $\Psi = 2N - 1$. If a processed cell loses r points to new cut boundaries and has t nonempty child point sets, the decrease is $2r + t - 1$. In the shrinking step $r \ge 1$. In a midpoint step either $r \ge 1$ or $t \ge 2$. Thus every event decreases Ψ by at least one, so there are at most $2N - 1$ events. Each event creates at most $2^D - 1$ additional leaves, since a shrinking step creates only D extra leaves. The number M of terminal rectangles satisfies

$$M \le 1 + (2N - 1)(2^D - 1) \le 2^{D+1}(2^D - 1)J. \tag{34}$$

Tracking exactly $V \cap \operatorname{int} C$ ensures that every terminal rectangle has no corner of an input cube in its interior. This argument is a construction and a decreasing integer potential, not a supplied partition certificate. It does not prove a fast implementation bound.

23.4 A uniform occupancy bound for the terminal rectangles

Terminal cubes meet at most 2^D input cubes properly, by the corner bound. A terminal noncube has the form E_i above after a coordinate reflection. It is contained in the slab $[0, s]^{i-1} \times [a, s] \times [0, s]^{D-i}$. Cover this slab by 2^{D-1} closed cubes of side $s - a$: in coordinate i use $[a, s]$, and in every other coordinate choose either $[0, s - a]$ or $[a, s]$. These intervals cover $[0, s]$ since $a < s/2$. Every covering cube lies inside the parent cube and outside the interior of

C', and its interior contains no point of V. The corner bound therefore applies separately to each covering cube.

If an input cube properly meets E_i, the intersection contains an open set. A point of that set avoiding the finitely many covering-facet planes belongs to the interior of a covering cube. Thus all proper intersections are counted by the cover, even if some particular point lies on a cover boundary. Consequently each terminal rectangle properly meets at most

$$T_D = 2^{D-1}2^D = 2^{2D-1} \tag{35}$$

input cubes. The covering argument, including its boundary qualification, is proved here; no arbitrary-box unclutteredness is assumed.

23.5 From open geometry to an actual integer partition

For each terminal rectangle list the at most T_D cubes meeting its interior. Refine by all their facet coordinates. There are at most $2T_D$ cut coordinates per axis, so at most $(2T_D+1)^D$ refined rectangles per terminal cell. All cuts are axis-parallel. Now assign every splitting hyperplane to one child consistently, intersect with X, and delete empty integer sets. The result is a disjoint cover of X by integer rectangles. Taking integer points and deleting empty sets does not increase the count. More explicitly, use the uniform convention $[l_i, h_i)$ in every axis of every cell. A coordinate split at t then assigns the plane to the upper child. The two intervals $[l, t)$ and $[t, h)$ partition $[l, h)$ exactly, even when t is an integer. By induction through every cut, these half-open cells partition $[-1, n)^D$; intersecting with X gives

$$R = \prod_i (\{\lceil l_i \rceil, \ldots, \lceil h_i \rceil - 1\} \cap \{0, \ldots, n-1\}).$$

This formula is also valid for negative or integer endpoints. An empty axis makes the entire integer cell empty; no lattice point is assigned twice or lost when those cells are deleted. The outer upper face at n contains no point of X.

To check membership rigorously, suppose an assigned integer point x lies in S_k. It is interior to Q_k, while its assigned terminal rectangle has x in its closure. A small ball about x lies in Q_k and meets the interior of that full-dimensional terminal rectangle. Therefore Q_k was in its list. None of the new facet coordinates is an integer. Hence each refined integer rectangle is entirely inside a listed S_k or disjoint from every listed S_k. An unlisted cluster cannot contain one of its integer points, by the same closure argument. Strict disjointness then implies that each integer rectangle is inside one cluster or outside every cluster.

In particular, with the explicit dimension-only constant

$$A_D = 2^{D+1}(2^D - 1)(2^{2D} + 1)^D, \qquad |\mathcal{P}| \le A_D \max\{1, J\}, \tag{36}$$

we have derived the required partition. The constant is deliberately loose and no sharpness or algorithm running-time claim is made. For $n = 1$ and for singleton input cubes all constructions remain valid.

23.6 The actual coefficient and fixed-penalty consequence

Suppose, in addition, that for fixed $c > 0$ and $\beta \geq 1$ the actual positive scalar starting coefficients obey

$$a(x) \geq cL_k(L_k - 1)/4 \quad (x \in S_k),$$
$$a(x) \geq cn(n-1)/4 \quad (x \notin \cup_k S_k),$$
$$a(x) \leq \beta L_k^2 \quad (x \in S_k).$$

These coefficient/cube assumptions are retained, not inferred from the critical moment. On an interval of l integer vertices, pairwise variance, path Cauchy–Schwarz and at most $l^2/4$ pairs through one cut give

$$\sum |u - \bar{u}|^2 \leq \frac{l(l-1)}{4} \sum_{\text{edges}} |du|^2.$$

For completeness, product variance tensorization follows from the commuting coordinate averages $P_i u = l_i^{-1} \sum_{t=0}^{l_i - 1} u(x_1, \ldots, t, \ldots, x_D)$: the telescoping orthogonal decomposition of $I - \prod_i P_i$ gives a sum bounded by $\sum_i \|u - P_i u\|_2^2$. Applying the interval bound on its fibers proves

$$\mathrm{Var}_R(u) \leq \frac{L(R)(L(R)-1)}{4} \sum_{e \subset R} |du(e)|^2.$$

Inside a cluster $L(R) \leq L_k$, and outside all clusters $L(R) \leq n$. The displayed coefficient lower bounds therefore give $c\,\mathrm{Var}_R(u) \leq \mathcal{E}_{a,R}(u)$. Singleton residuals are zero, so no zero lower bound is divided by. Sum over all field/internal coordinates and all partition cells. Original internal positive-axis edge occurrences are disjoint; omitted inter-cell and periodic-wrap occurrences have nonnegative energy. Thus the ACTUAL normalized means

$$(C_{\mathcal{P}} U)_{R,j} = \sqrt{|R|}\ \mathrm{mean}_R\, U_j$$

satisfy

$$c\|U\|^2 \leq \mathcal{E}_a(U) + c\|C_{\mathcal{P}} U\|^2$$

for every cochain, with output dimension $h|\mathcal{P}|$, $h = D(N_c^2 - 1)$. This uses an all-cochain variance estimate, not the invalid inference from a kernel-only margin to fixed penalty c. Means need not preserve the divergence constraint. The inequality consequently also holds on that constraint subspace without changing $C_{\mathcal{P}}$.

If $p = D/2$, $\mu \geq 1$, and the original critical mass bound is $\sum_X a^{-p} \leq M n^D/\mu^p$, then each supplied cube contributes at least $L_k^D/(\beta L_k^2)^p = \beta^{-p}$ to this mass. Hence $J \leq \beta^p M n^D/\mu^p$, and (36) gives

$$\dim F \leq h A_D (1 + 2^p \beta^p M) \max\{1, n^D/(1+\mu)^p\}.$$

This proves the restricted supplied-cube consequence with an explicit constant and without a fat-object BSP theorem as a black box. It does not construct the supplied cubes from the critical moment, remove either scale bound or the strong complement bound, resolve the unrestricted critical rectangular question, or establish a covariant/continuum gap. No new originality or independent-discovery claim is made: the compressed splitting strategy is attributed classical geometry, here derived and adapted completely for this special class. The frozen reviewed PDF and its full reports and pending permission remain unchanged.

24 Physical cycle obstructions for an actual observation map

This part closes a limitation of the scalar four-cycle example: the obstruction is now realized inside the actual divergence-zero flat physical cochain domain. The kernel estimate below is proved for all actual cochains, not just the special fields used as negative witnesses. The conclusion concerns this observation and this fixed penalty. It does not obstruct all partitions or the optimal-rank question from the preceding parts.

24.1 The actual form and rectangular observations

Fix $D \geq 3$, $N_c \geq 2$ and $X = (\mathbb{Z}/4\mathbb{Z})^D$. Write $V = \mathbb{R}^{N_c^2-1}$ with its standard Euclidean norm and $A = (A_\ell(x))_{x\in X, 1\leq \ell\leq D}$ for an actual one-cochain. All norms use counting measure. Choose the actual total coefficients

$$\alpha = (1, 6, 4, 1), \qquad a(x) = \alpha_{x_1}, \qquad w(i, x, \ell) = \alpha_{x_1} - 1$$

for every derivative direction i and physical component ℓ. Define actual forward and backward differences by

$$D_i^+ f(x) = f(x + e_i) - f(x), \qquad D_i^- f(x) = f(x) - f(x - e_i).$$

The flat curl and backward divergence, with one curl for each unordered pair, are

$$(\operatorname{curl} A)_{ij} = D_i^+ A_j - D_j^+ A_i \quad (i < j), \qquad \operatorname{div}^- A = \sum_i D_i^- A_i.$$

Their counting quadratic form and the original componentwise repair are

$$H_0(A) = \sum_{i<j,x} \|D_i^+ A_j(x) - D_j^+ A_i(x)\|^2 + \sum_x \Big\| \sum_i D_i^- A_i(x) \Big\|^2,$$
$$R_w(A) = \sum_{i,x,\ell} w(i, x, \ell) \|D_i^+ A_\ell(x)\|^2.$$

In operator notation the repair is $\langle A, G^* W G A\rangle$ for the actual stacked gradient G and diagonal multiplication W. These formulas fix the energy on paper; no comparison certificate is supplied.

Deriving the flat identity. Periodic reindexing gives

$$\sum_x \langle D_i^+ f, h\rangle = -\sum_x \langle f, D_i^- h\rangle.$$

Coordinate shifts commute. Applying this formula twice and real symmetry therefore yields

$$\sum_x \langle D_i^+ f, D_j^+ g\rangle = \sum_x \langle D_j^- f, D_i^- g\rangle.$$

Also $\sum_x \|D_i^- f\|^2 = \sum_x \|D_i^+ f\|^2$ by translation. Expand each curl square with $f = A_j$, $g = A_i$: it supplies the two off-diagonal gradient norms and the negative mixed terms. Expansion of the divergence square supplies the diagonal norms and the opposite mixed terms. Cancellation for each $i < j$ proves, for every cochain,

$$H_0(A) = \sum_{i,x,\ell} \|D_i^+ A_\ell(x)\|^2.$$

No divergence condition or lift ansatz was used. Define $H = H_0 + R_w$. Adding the displayed repair, with $1 + w(i, x, \ell) = \alpha_{x_1}$, gives

$$H(A) = \sum_{x \in X} \sum_{i=1}^{D} \sum_{\ell=1}^{D} \alpha_{x_1} \|A_\ell(x + e_i) - A_\ell(x)\|^2. \tag{37}$$

The cancellation was unweighted; the varying repair is added afterwards. We do not move α through summation by parts or assert a weighted curl–divergence identity. This is the paper counterpart of `finite_flat_curl_div_energy` and `componentwiseWeightedRepair_energy` in the pinned formal source. Thus the coefficient is the original shared scalar total coefficient, not a substituted energy or an interacting Hessian.

Put $I_0 = \{0, 1\}$ and $I_1 = \{2, 3\}$. For every transverse $y \in (\mathbb{Z}/4\mathbb{Z})^{D-1}$ the blocks $I_b \times \{y\}$ are nonempty axis-aligned rectangles. They are disjoint and cover X. For each physical component ℓ and each standard internal coordinate k, define the actual normalized means

$$(CA)_{b,y,\ell,k} = \frac{1}{\sqrt{2}} \sum_{t \in I_b} A_\ell(t, y)_k. \tag{38}$$

The output includes *all* these coordinates. If $h = D(N_c^2 - 1)$, its dimension is $2h4^{D-1}$; no optimality is asserted.

The spatial mean of a is $\mu = 3$, and $1 \le a \le 2\mu$. For $p = D/2$ its original normalized inverse moment is

$$\frac{1}{|X|} \sum_{x \in X} \left(\frac{\mu}{a(x)} \right)^p = \frac{2 \cdot 3^p + (1/2)^p + (3/4)^p}{4} =: M_D < \infty. \tag{39}$$

This is an actual member of the stated bounded-moment coefficient class when the allowed bound is at least M_D. Its fixed finite size and the large output dimension do not disprove an asymptotic optimal-rank order.

24.2 A fixed-penalty counterexample with a proved kernel premise

For a real cyclic vector $u = (u_0, u_1, u_2, u_3)$ set

$$E(u) = (u_1 - u_0)^2 + 6(u_2 - u_1)^2 + 4(u_3 - u_2)^2 + (u_0 - u_3)^2,$$
$$|u|^2 = \sum_{t=0}^{3} u_t^2, \qquad O(u) = \frac{(u_0 + u_1)^2 + (u_2 + u_3)^2}{2}.$$

Choose a unit internal coordinate vector T. The actual lift is

$$(Ju)_2(x) = u_{x_1} T, \qquad (Ju)_\ell(x) = 0 \quad (\ell \neq 2).$$

Its backward divergence is zero: its only nonzero component is constant in coordinate 2. It is injective and direct summation gives

$$\|Ju\|^2 = 4^{D-1}|u|^2, \quad H(Ju) = 4^{D-1}E(u), \quad \|CJu\|^2 = 4^{D-1}O(u). \tag{40}$$

The last equality uses the entire output in (38), not only a selected nonzero coordinate.

Kernel premise on all actual fields. For $v = (x, -x, y, -y)$, direct algebra yields

$$E(v) = 11x^2 + 14xy + 23y^2, \qquad |v|^2 = 2x^2 + 2y^2,$$
$$E(v) - \frac{19}{5}|v|^2 = \frac{17}{5}\left(x + \frac{35}{17}y\right)^2 + \frac{84}{85}y^2 \geq 0. \tag{41}$$

If $CA = 0$, both pair sums vanish for every actual scalar profile $t \mapsto A_\ell(t, y)_k$. Apply (41) to each such profile and sum over all y, ℓ, k. These profiles account for the full norm of A. Their energies are exactly the $i = 1$ term of (37); every other derivative term is nonnegative because $\alpha_t \geq 1$. Therefore

$$CA = 0 \quad \Longrightarrow \quad H(A) \geq \frac{19}{5}\|A\|^2 \qquad \text{for every actual cochain } A. \tag{42}$$

No divergence condition, scalar-lift ansatz or supplied coercivity certificate was used in this proof.

Failure on an actual divergence-zero field. For $u = (19, -9, -5, -5)/14$ one has

$$E(u) = \frac{52}{7}, \qquad |u|^2 = \frac{123}{49}, \qquad O(u) = \frac{25}{49}.$$

Consequently the actual divergence-zero field Ju satisfies

$$H(Ju) + \frac{19}{5}\|CJu\|^2 - \frac{19}{5}\|Ju\|^2 = -\frac{6}{35}4^{D-1} < 0. \tag{43}$$

Equations (42) and (43) prove the complete fixed-map statement: kernel coercivity does not imply the same coefficient as a fixed coarse penalty, even on the actual flat physical divergence-zero domain. They do not rule out larger penalties or different observations.

24.3 The exact finite penalty at the subsharp target

The failure just proved concerns $K = 19/5$. Here we determine exactly which larger penalties work for this same observation and target. This subsection is a paper proof with an exact rational coefficient check; it is not part of the paired Lean theorem below.

Every $u \in \mathbb{R}^4$ has the unique decomposition

$$u = m(1, 1, 1, 1) + (x, -x, y, -y) + (\delta/2, \delta/2, -\delta/2, -\delta/2),$$

where $m = (\sum_t u_t)/4$, $x = (u_0 - u_1)/2$, $y = (u_2 - u_3)/2$ and $\delta = (u_0 + u_1 - u_2 - u_3)/2$. Direct expansion gives

$$|u|^2 = 4m^2 + 2x^2 + 2y^2 + \delta^2, \qquad O(u) = 4m^2 + \delta^2,$$
$$E(u) = 11x^2 + 14xy + 23y^2 + 7\delta^2 - 10\delta(x + y).$$

Set $X = x - 25\delta/2$ and $Y = y + 75\delta/14$. For every real K,

$$E(u) + KO(u) - \frac{19}{5}|u|^2 = 4\left(K - \frac{19}{5}\right)m^2 + \frac{17}{5}\left(X + \frac{35}{17}Y\right)^2 + \frac{84}{85}Y^2 + \left(K - \frac{1138}{35}\right)\delta^2. \tag{44}$$

Since $1138/35 > 19/5$, this proves sufficiency of $K \geq 1138/35$ on every scalar profile. Necessity follows from

$$u_* = (13, -12, -41/7, 34/7), \qquad O(u_*) = 1, \quad |u_*|^2 = \frac{18174}{49}, \quad E(u_*) = \frac{67468}{49}.$$

It has $m = 0$, $\delta = 1$ and $X = Y = 0$, hence defect $K - 1138/35$.

Apply the scalar inequality to every profile $t \mapsto A_\ell(t, y)_k$ and sum over all y, ℓ, k. Their squared norms and observations exhaust the actual full norms, and their energies exhaust the first-direction contribution in (37). The remaining directions are nonnegative. Conversely Ju_* is actual divergence zero and multiplies the scalar defect by 4^{D-1}. Therefore the following are equivalent:

1. $K \geq 1138/35$;
2. $H(A) + K\|CA\|^2 \geq (19/5)\|A\|^2$ for every actual cochain;
3. the same inequality for every actual divergence-zero cochain.

This is a necessity-and-sufficiency statement for the fixed map, not an obstruction to different partitions. In particular, a larger finite penalty repairs the subsharp target even though $K = 19/5$ fails. The sharp kernel target in the next subsection behaves differently.

A second calculation records the classical block elimination behind the constant. The kernel residual matrix and mixed vector are

$$M = \begin{pmatrix} 17/5 & 7 \\ 7 & 77/5 \end{pmatrix}, \quad \det M = 84/25, \quad b = (5, 5), \quad M^{-1}b = (25/2, -75/14).$$

Thus $b^T M^{-1} b = 250/7$ and the residual coefficient is $K + 16/5 - 250/7 = K - 1138/35$, agreeing with (44). No new Schur-complement principle is claimed.

The unit151 exact verifier constructs E, O and $|u|^2$ from the actual four-coordinate linear forms and compares every coefficient of the polynomial in (m, x, y, δ, K), rather than sampling values. Normal and optimized Python in one own Colab Linux runtime each accepted five checks and rejected three real coefficient mutations; both exits were zero and their JSON outputs were identical. The 8794-byte archive has ten members and nine checked inner hashes. It and the full hand derivation are included in the reproduction package. The driver took about 0.56 seconds; the runtime was released after recovery and exact billed consumption is unknown. This check is neither a new Lean theorem nor independently authored verification. Its source is

`564b84230478cb8d40505cb67985470e9bae699d.`

24.4 Sharp free-penalty extension and its distinct status

The following is an analytic continuation, outside the paired Lean result. Put $s = \sqrt{85}$ and $c_* = (17 - s)/2$. Since $9 < s < 10$, the exact identity

$$E(v) - c_*|v|^2 = (s - 6)\left(x + \frac{7y}{s - 6}\right)^2$$

holds on the same block kernel. The nonzero vector $v_* = (7, -7, 6 - s, s - 6)$ attains equality. The all-profile argument above also establishes the sharp kernel constant for the actual form in paper proof, and the lift establishes sharpness.

For $p_\delta = (\delta/2, \delta/2, -\delta/2, -\delta/2)$, the two vectors are orthogonal, $Cv_* = 0$, and $O(p_\delta) = |p_\delta|^2 = \delta^2$. For any finite $K \geq 0$, direct expansion gives

$$\begin{aligned} E(v_* + p_\delta) + KO(v_* + p_\delta) - c_*|v_* + p_\delta|^2 \\ = (7 + K - c_*)\delta^2 - 10(13 - s)\delta. \end{aligned}$$

The positive number $q_K = 7 + K - c_*$ permits the choice $\delta = 5(13 - s)/q_K$, giving the strictly negative value $-25(13 - s)^2/q_K$. Lifting multiplies this complete defect by 4^{D-1} and preserves actual divergence zero. Thus no finite quadratic penalty repairs this same observation at the sharp kernel endpoint. A penalty $\Phi(z) \leq K\|z\|^2$ only near zero is also defeated: scale the negative witness by a sufficiently small nonzero real number. Both radius conditions can be met, and its defect remains a negative multiple of the square of that scale. This is not a strict-margin claim.

These exact analytic identities and arguments are supplied here, but their complete physical sharp-endpoint/free-penalty statement has *not* been verified in Lean in the paired execution described next. Finite checks in dimensions three and four, including $K = 0, 1, 100$, are diagnostics only, not the universal theorem or independent review.

24.5 Formal scope and the paired reproduction record

The fixed-penalty result, including the actual full-kernel premise and all transport/comparison identities used to connect it, was compiled in two separately allocated Colab Linux CPU sessions with high RAM. Each started with a fresh clone at the same immutable source

`0c706827dcbb11859e56bb6c7dc981eb37a02e69.`

The final theorem is named `physicalCycleFour_kernel_fixed_penalty_counterexample` in `PhysicalCycleFullForm.lean`. It asserts the proved all-cochain kernel bound together with failure of the fixed-penalty inequality on the actual divergence-zero domain. Four preceding printed endpoints identify longitudinal energy, the full actual energy sum, its comparison, and the kernel estimate. They do not assume transport certificates or posit a structure containing the conclusion.

All five actual dependency readouts in each clone contain only `propext`, `Classical.choice` and `Quot.sound`. The normal and optimized validator in each clone accepted all five readouts and rejected three actual oracle mutations. All actual decimal exit sentinels were zero. Lean was version 4.29.0-rc6 with Mathlib pinned in the included manifest. The recovered theorem source, oracle source, compiled binary, oracle output, validator source, normal and optimized outputs and driver have identical hashes across the two fresh executions. The raw per-run reports still say first-clone pass and paired false: the identical driver executes one clone per session. They remain unmodified; the separate aggregate records the completed pair.

The first successful driver took about 374.31 seconds and the second 400.28 seconds. These are not billed connection times. Both environments were disconnected after complete evidence recovery. No Lean compilation, oracle or numerical campaign ran on Windows. Separate fresh executions are not independently authored proofs or independent scientific review. The previous failed attempt, its exact source and full compiler log are preserved: mixed finite-index inference and nested sum type inference were corrected without changing the theorem statements.

The package includes both public evidence archives. Their comparison is recorded in the aggregate file:

`paired-full-form148.json`

Full hashes and source bindings are in `inputs.json`. Each archive has 45 checked inner manifest members and excludes third-party installers and downloaded binary assets. Their complete originals, including those assets, are retained in the private evidence record. The compiled binary hash begins `b0f8ff8654d79d02`, and the full value is in the paired comparison file.

Attribution and scope of novelty. The scalar square identities and negative vectors predate this physical embedding in the dossier. The present consolidation adds the explicit actual-domain embedding and the formally derived all-field comparison; it claims no new Schur-complement principle, priority or exhaustive novelty finding. Carlson, Haynsworth and Markham's classical paper, *A Generalization of the Schur Complement by Means of the Moore–Penrose Inverse*, SIAM Journal on Applied Mathematics 26 (1974), 169–175, doi:10.1137/0126013, supplies algebraic context. The publisher abstract and bibliography were inspected in the source investigation; the full paywalled article was not read. The direct finite proof here carries the mathematical claim.

25 Evidence provenance and verification limits

The analytic sources, their exact revisions and the executable finite diagnostics are collected in a hash-allowlisted companion package. Its map gives full SHAs and SHA256 values; short forms below are identifiers, not substitute hashes.

Item	Source revision	Verification status
Actual flat Hodge identity	`dfd49d0b0730`	earlier formal source
Critical spectral theorem	`527fcdcd119a`	analytic adapters
Dyadic spatial separation	`fb65a67af360`	analytic adapters
Higher moment block proof	`6dcb893155e5`	analytic adapters
Exact finite fixture code	`4bb460e1537e`	own Colab109
Full diagnostic report	`602ae3f0b59a`	finite rational batch

Colab109 checked four actual profiles on $(\mathbb{Z}/2\mathbb{Z})^4$ using rational arithmetic. Each normal and optimized Python mode accepted four profiles and rejected 16 real matrix, inverse or flow mutations, with matching outputs and exit code zero. It checked the constant mode, forward symbols, actual pseudoinverse identities, feasible/minimizing flow energy, comparison inertia, strict low-spectrum multiplicities and noncommuting counts. There are 12 distinct density examples, evaluated again for each profile: 48 evaluations per mode, not 48 independent examples.

All eight files in the diagnostic archive's inner manifest were rehashed. Runtime execution was recorded at 17:13:06.300923–17:13:16.201862 UTC on 2 November2026, with automatic runtime release and zero sessions verified afterwards. Rounded resource units do not prove zero cost. No local Windows mathematical execution occurred. This single embedded-SHA finite batch is not two fresh matching clones, a Lean proof, a theorem about arbitrary dimensions or odd lattices, or independent mathematical review. This manuscript assembly adds no new Lean execution. Operational repository checks on unchanged Lean files do not verify these new analytic proofs.

There is also subsequent, strictly focal formal evidence. Source `41f8492c6190` supplies the ten four-cycle theorems and their oracle. Own Colab118 compiled that same source in two fresh clones; both actual oracle logs and compiled module hashes agree. Each of four normal/optimized validator runs accepted ten readouts and rejected three real corruptions. All readouts use only

`propext`, `Classical.choice`, and `Quot.sound`. The full archive has 69 manifest files and is included with the failed Colab116/117 compiler attempts. Source hashes, complete logs and exact reports are allowlisted. This verifies Proposition 20.1's finite identities, not the geometric, shorted-operator, free-penalty or general rectangular adapters. Own Colab115 separately checked 18 finite diagnostics and eight real mutations in each mode of two fresh clones. Finite examples do not establish their general claims. Section 21 consolidates analytic sources at `0fa0d91000f6`, `312f476ea6e5`, `f52ae344f61f`, `1850be2aa6d5` and `b348f4488b3a`. The original spatial proofs and the self-contained Sobolev and cube arguments remain analytic. The physical fixed-penalty endpoints have the separately pinned paired formal record in Section 24. The exact subsharp threshold has an exact rational coefficient check, not a Lean certificate; the sharp extension is analytic. Rendering does not upgrade any verification status. Earlier exact PDFs and assessments remain preserved.

26 Model assessments and review context

The prior exact v5 received a preparing-agent self-assessment of 6/10 on a separate rubric, with eight retained objections. Its later Google Gemini report used the AIRR assessment scale and scored 3.4/10 with a major-revision recommendation. These scales are not pooled or converted. Neither number rates this changed candidate. The full prior report, request, findings, exact-artifact hash and operator response are preserved in the v5 model-assessment record.

The interface displayed Gemini 4 Argon with High reasoning. The operator observed that display; no backend API attestation was available. The request supplied the exact v5 PDF and withheld earlier reports and scores. Prior participation, technical history isolation and memory isolation remain unknown. The model reported a PDF-only review without external browsing, code execution or inspection of the repository and Lean/Colab evidence. Its asserted proof checks are model statements, not an operator correctness certificate. This is neither established independent review nor human peer review.

The sole listed material objection alleged a missing factor c in v5 Equation (28). Exact TeX, PDF text and the hash-matched rendered page show that c was already present. A separate preparing-agent response records this contradiction without editing the report or its grade. The claim that Proposition 11.3 omitted the endpoint $R = 1$ also contradicts the exact artifact. Other reported defects were real: this candidate corrects stale weight symbols m to μ, adds a stable reference to Section 16, and corrects an additional stale weight symbol in the physical line-test energy. The original v5 PDF and all its reports remain unchanged. No repeat review of unchanged v5 was requested to improve its grade.

Reports on distinct frozen successors. The 33-page revision136, SHA256 prefix `eb4df3`, received its first Google Gemini report of 3.5/10 and minor revision. The operator observed Gemini 4 Argon High; the raw backend identifier is unknown. The review was non-blinded, with historical assessment context in the PDF. Its five retained objections concern uninspected formal/Colab evidence, non-Lean analytic continuations, primary sources not inspected by that reviewer, the open unrestricted rectangular question, and the absent interacting or continuum connection. The full report, request and context are included separately as `review136-*` in the reproducibility package.

The distinct 14-page companion, SHA256 prefix `5c2774`, received one first external Gemini report of 6.5/10 and minor revision. The selected interface displayed Gemini 4 Argon High; the

served backend is unexposed and the report records unknown. No target grade or regeneration was used. Historical context was visible in that PDF, so the review was non-blinded; independence and context isolation are not established. All five objection groups remain in the complete `review150-*` files: incomplete energy definitions and uninspected execution archives; the local-to-global and supplied-cube gap; distinct fixed, subsharp and sharp penalties; alleged technical defects; and uninspected primary sources and review context. Four technical allegations were contradicted by that exact artifact: $a \geq s$ and $a, b \geq 0$ were already printed, $2\,3^p$ was a product, and $O(p_\delta) = |p_\delta|^2 = \delta^2$ was complete across a line break. The original report and grade, including inconsistent score subtotals, remain unchanged. A separate operator disposition retains the evidence.

This successor supplies the missing energy definitions, cancellation and exact finite subsharp threshold. Those additions do not make either model inspect the archives, primary sources or this changed PDF. The 16-page companion revision152 was an unreviewed consolidation, with no inherited score. None of these historical scores rates the present combined artifact. Model evaluations, public working-paper deposits, editorial acceptance and human peer review remain distinct; no new acceptance or peer review is claimed. This candidate is prepared by the same continuing Codex programme agent. The served backend identity and effort for preparation are unexposed; no technical isolation or independence is claimed. No numerical score or completed external exact-artifact assessment is assigned to this changed candidate in the manuscript. The exact resulting PDF must receive a separately retained first assessment after every page is inspected; no earlier score or resolved-objection status transfers automatically.

The principal limitations and objections are these. The new finite Fourier, inverse, field and partition adapters are analytic rather than kernel-verified. The classical Rumin–Frank method and elementary refinement geometry limit any novelty claim; a complete priority comparison is unfinished. The sufficient theorem uses one shared scalar coefficient profile, and its distribution condition is not proved from arithmetic means. Spectral coordinates are global; unrestricted critical rectangular means remain unresolved. Exact finite diagnostics cover a small selected family, and some count bounds are loose. The constants are safe rather than sharp. This manuscript makes no statement about an interacting or continuum physical mass gap. The remaining programme obligations are not reduced by this manuscript; no continuum closure is claimed.

27 The single remaining bounded question

For fixed $D \geq 3, c > 0, N_c \geq 2, M$, do all actual scalar profiles $1 \leq a \leq 2\mu$, exact mean μ, satisfying (16), admit (17) using normalized means on an unrestricted axis-aligned integer rectangular partition? The mechanism must leave full cube trees, for example through anisotropic cuts with a proved variance bound. A falsifier must defeat every such rectangular partition, not merely every dyadic cube tree. After three substantive units without a proved reduction or useful falsifier, the route must be changed or its concrete blocker identified. The common-refinement route remains stopped at its recorded three-unit limit. Formalization and consolidation do not reset that counter or create a second active question. PR29 remains 20/41 with TermSource zero.

References

[1] L. Eriksson, *Exact-Kernel Fourier Families Obstruct Every Penalty Factored Through a Block Map in a Flat Lattice Gauge Form*, ai.viXra 2608.0015v1 (2026), AIRR ARR-2026-7EPM3V3DP796BBRR v1. `https://airr.science/papers/ARR-2026-7EPM3V3DP796BBRR/versions/v1/`. Historical import, not peer reviewed or assessed by AIRR; checked 30 October 2026.

[2] J. Dodziuk, *Finite-difference approach to the Hodge theory of harmonic forms*, American Journal of Mathematics **98** (1976), 79–104. `https://doi.org/10.2307/2373615`. Institutional source: `https://academicworks.cuny.edu/gc_pubs/180/`.

[3] A. Albert, *Conditions for Positive and Nonnegative Definiteness in Terms of Pseudoinverses*, SIAM Journal on Applied Mathematics **17** (1969), 434–440. `https://doi.org/10.1137/0117041`.

[4] D. Carlson, E. Haynsworth and T. Markham, *A Generalization of the Schur Complement by Means of the Moore–Penrose Inverse*, SIAM Journal on Applied Mathematics **26** (1974), 169–175. `https://doi.org/10.1137/0126013`.

[5] J. Dimock, *Covariant Axial Gauge*, arXiv:1410.2798v2 (2014). `https://arxiv.org/abs/1410.2798v2`.

[6] Y. Kudryashov and Mathlib contributors, *Implicit function theorem*, Mathlib source at commit `07642720480157414db592fa85b626dafb71355b`. Pinned primary formal source, particularly `HasStrictFDerivAt.to_implicitFunction` and its exact level-set identity; inspected 30 October 2026. This attributes the classical parametrization; the separately pinned sector formalizations are identified in the evidence section.

[7] D. A. Spielman, *Graphs, Vectors, and Matrices*, Bulletin of the American Mathematical Society **54** (2017), 45–61. Author-hosted primary exposition: `https://cs-www.cs.yale.edu/homes/spielman/PAPERS/bull1557.pdf`.

[8] D. A. Levin and Y. Peres, with contributions by E. L. Wilmer, *Markov Chains and Mixing Times*, second edition, American Mathematical Society (2017), Chapter 13. Author-hosted text: `https://pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf`.

[9] D. Zelazo and M. Bürger, *On the Definiteness of the Weighted Laplacian and its Connection to Effective Resistance*, arXiv:1408.2187v1 (2014), Theorem III.3. `https://arxiv.org/pdf/1408.2187`.

[10] Y. Lin, S. Wan and H. Zhang, *Connection Laplacian on discrete tori with converging property*, arXiv:2403.06105v1 (2024), Theorem 1.2. `https://arxiv.org/html/2403.06105v1`. The sin-squared discrete-torus spectrum is attributed prior work, not a new general spectral principle here.

[11] M. de Berg, *Linear Size Binary Space Partitions for Uncluttered Scenes*, UU-CS-1998-12 (1998), Section2.1, Lemmas2.1 and2.4, Theorem2.5 and Corollary2.9. `https://ics-archive.science.uu.nl/research/techreps/repo/CS-1998/1998-12.pdf`.

[12] W. N. Anderson, Jr. and G. E. Trapp, *Shorted Operators. II*, SIAM J. Appl. Math. 28(1) (1975), 60–71. `https://doi.org/10.1137/0128007`.

[13] R. L. Frank, *Cwikel's theorem and the CLR inequality*, arXiv:1206.3325v1 (2012), Section2, Theorem2.1, Corollary2.2, Lemmas2.4–2.5. `https://arxiv.org/abs/1206.3325`.

[14] S. Andres, J.-D. Deuschel and M. Slowik, *Harnack inequalities on weighted graphs and some applications to the random conductance model*, arXiv:1312.5473v6 (2019), Section2.2, equation(2.5) and Proposition2.1. `https://arxiv.org/abs/1312.5473`.