Coarse information and finite rank obstructions for flat lattice Hodge forms

Lluis Eriksson

30 October 2026

Abstract

On a finite periodic lattice with unit spacing, a flat Hodge form cannot have a positive coercivity constant uniform under refinement if it is supplemented only by linear observations in a fixed finite-dimensional real space and linear fine-space repairs of uniformly bounded rank. The linear observations may vary arbitrarily with the refinement, and their penalty need only vanish at zero. We prove the same obstruction directly on kernels of uniformly bounded rank constraints. The proof constructs a real physical Fourier subspace of dimension $2(2R+1)^d$ with Hodge energy at most $4\pi^2 dR^2/n^2$ times its squared norm. Its intersection with the unobserved directions gives a necessary rank budget. We also explain why this budget is not sufficient, using a two-dimensional cross-coupling example and a classical complete criterion for a fixed map. The Fourier construction, finite budget, and uniform obstructions have Lean proofs; the fixed-map criterion is an analytic result here. These are finite flat statements, with no interacting or continuum mass-gap claim.

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 always real-linear. The penalty applied to its output may be nonlinear. 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

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

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 $B_y = \{\ell y + r : 0 \leq r_i < \ell\}$. The original line-integral map is

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

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 $\sigma \geq (d-2)/2$ 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 $\ell^{d-2-2\sigma} \leq C_P$. The real-exponent comparison is analytic here.

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

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

forces $C_P(\ell) \geq \lambda_\ell^{-1} \geq \ell^2/(4\pi^2)$. 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 $1 \leq r \leq R \leq \lfloor (\ell-1)/2 \rfloor$, 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 $4\sin^2(\pi R/\ell) \leq 4\pi^2 R^2/\ell^2$. Finite-dimensional min–max gives $\mu_{2R}(H_\ell) \leq 4\pi^2 R^2/\ell^2$. 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 $\Lambda_R = \{-R, \dots, R\}^d$ 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 $A \in S_R$,*

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

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 $|e^{\mathrm{i}t} - 1| \le |t|$ and $|k_i| \le R$, 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*

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

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 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 7.1. *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 $H = \operatorname{diag}(0, 2)$, $c = 1$ 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.

8 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.

9 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 corrected map is `declaration-map.json` and the preserved evidence is `reproducibility/`. Exact source SHAs, not moving branch heads, identify the formal claims. The revised map and manuscript have their own revision SHA in the accompanying package.

The geometric and Fourier modules are ordinary proof dependencies, not assumptions about a newly supplied band. Each development stage has complete same-source fresh-clone Colab evidence. 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.

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.

Model assessments and review context

This text was prepared with the Codex agent in the author's ongoing research chat. The exact served model identifier is not exposed in this preparation record; a different named reviewer or reasoning setting is not inferred. The agent had the prior manuscript objections, source audits and reproduction records, and was not blind to the programme. An earlier eight-page PDF (SHA256 recorded in the accompanying report) was assessed once by a subagent configured as `gpt-6-astra`, with express author authorization. Actual backend identity and reasoning setting were not exposed. The reviewer saw the exact PDF and prior context; the workflow used the same chat, account and provider and was not blind or organizationally independent. Its complete report and six objections are preserved with this revision. This revised artifact incorporates scope and attribution corrections; the earlier assessment does not apply to its new bytes. No numerical score, independent human peer review or editorial acceptance is claimed. The reviewer did not rebuild Lean or independently audit the complete evidence/dependency closure; formal reproduction statements above describe the supplied records. Fresh runs of the same workflow are reproducibility checks, not independent scientific assessments. Previous version scores, if any, do not transfer to a revised artifact.

Lluis Eriksson is both the author and operator of AIRR.science. Any later founder-author editorial decision must be recorded separately from model assessments and from independent review. The existing record [1] is an author-authorized historical import, not assessed under the current AIRR screening. This preparation does not change its public version or publication status. Code and manuscript licensing must be recorded separately; no new license grant for the historical PDF is inferred.

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`.