Coverage and coexactness in checkerboard slice observations

Lluis Eriksson
THE ERIKSSON PROGRAMME

12 November 2026

Abstract

We study original-channel rectangular means for a finite three-dimensional checkerboard gradient energy with weak coefficient $c/48$ and strong coefficient cn^2. Each partition has at most $36n$ rows and is specified by one cut per slice. If the cut diagonal is a permutation, actual singleton observations give an explicit logarithmic coercivity bound and a finite quadratic penalty, valid on all fields and therefore on the actual curl-adjoint image. For arbitrary cuts, the full means exclude every nonzero field of exactly zero strong energy; this qualitative fact does not restore the target coefficient c. Conversely, at an explicit, extremely large dyadic threshold, every cut diagonal with image of size at most $n/32$ admits a real, nonzero, genuinely coexact invisible field of energy below $c/2$ times its squared norm. We also prove a contrast-uniform factor-four energy bound for the true coexact projection, give an explicit counterexample to its preservation of these mean kernels, and quantify a necessary correction cost. The last estimate concerns the correction alone, not the energy of the corrected field. The permutation case at the original target for the nonlinear example $F(i) = i + 2i^2$, historical priority, and a matching logarithmic rate remain open; some affine permutation subfamilies already have counterexamples at that target. These are finite analytical results, not formal verification or a physical Yang–Mills spectral-gap theorem.

1 Problem, scope and principal conclusions

Low-dimensional observations of a high-contrast energy depend both on the observed functionals and on the domain on which a restoring inequality is required. Here the domain distinction is concrete: a field in the actual curl-adjoint image need not remain there after rectangular averaging, while projection into that image need not preserve rectangular mean zeros. We keep these two operations separate throughout.

The purpose is a consolidated account of the programme's finite analytical sources 266–273. Their complete proofs and source-specific reports accompany the reproducibility package. The preceding checkerboard working paper [3] treats broader observation-count and tagged-window questions. We do not claim a new general theory of capacity or Hodge decomposition. Logarithmic point estimates in two dimensions are classical [1]; stability of projections in differential complexes is also an established subject [2]. Every estimate needed for this particular coefficient family is derived below. A bounded comparison with these primary sources does not establish historical originality or wider practical relevance.

The results describe three distinct situations. Permutation diagonals yield a paid lower bound of order $c/(1 + \log n)$. Arbitrary diagonals still exclude exactly strong-null invisible fields. Low-coverage diagonals nevertheless fail at target c through fields of positive, small strong energy. Thus qualitative injection on the strong kernel is insufficient, and coverage alone is not proved sufficient. We establish no condition-number or solver guarantee.

1.1 Finite model and actual observations

Let $n \geq 4$ be even, $X = (\mathbb{Z}/n\mathbb{Z})^3$, $V = n^3$, and $c > 0$. Set

$$a = c/48, \qquad H = cn^2, \qquad S(x) = x_1 + x_2 + x_3 \pmod{n}.$$

The two original channels are orthogonal. Each has three spatial components. For $\partial_j f(x) = f(x + e_j) - f(x)$, channel 1 has coefficient H at odd source parity and a at even source parity; channel 2 has the reverse assignment. The coefficient is scalar across spatial directions and components within a fixed source and channel. Write

$$Q(u) = Q_a(u) + Q_H(u) = \sum_{\ell=1}^{2} \sum_{x,j,r} b_\ell(x) |\partial_j u_{\ell,r}(x)|^2.$$

All norms are counting norms. All forward labels are retained.

For a one-form v, the forward curl is $(d_1 v)_{jl} = \partial_j v_l - \partial_l v_j$, $j < l$. The domain of interest is

$$U = (\operatorname{im} d_1^*)^2.$$

This is the actual image with counting adjoints; constants are excluded. We never replace it by a divergence-free space without removing harmonic modes.

Choose any $t_i \in \{0, \ldots, n-1\}$. On the slice $x_1 = i$, partition x_2 into the nonempty intervals

$$[0, t_i - 1], \quad \{t_i\}, \quad [t_i + 1, n - 1],$$

and partition x_3 into $\{0\}$ and $[1, n-1]$. Every original spatial component in each channel is observed on every resulting cell. A row is $V^{-1/2}|B|^{-1/2} \sum_{x \in B} u_{\ell,r}(x)$. Call the full map C_t, and put $F(i) = i + t_i \pmod{n}$. If Z_0 and Z_* count cuts equal to 0 and $n-1$, respectively, then the raw row count is

$$36n - 12(Z_0 + Z_*), \qquad 24n \leq \#\text{rows} \leq 36n.$$

Every cell is a nonwrapping chart rectangle. No cyclic-union mean or mixed channel replaces these rows. In particular $(i, t_i, 0)$ is always an actual singleton. The target is

$$Q(u)/V + \Phi(C_t u) \geq c\|u\|^2/V, \qquad \Phi(0) = 0, \tag{1}$$

on U, with a finite observation penalty. Our counterexamples also defeat quadratic penalties. We do not assert failure for every canonical partition.

2 The strong graph and a paid point estimate

Fix a scalar component in a channel with strong-source parity p. Its strong components are

$$G_s = \{x : S(x) \in \{s, s+1\}\}, \qquad s \equiv p \pmod{2}.$$

Indeed the undirected strong path $x \to x + e_j \to x + e_j - e_k$ connects source sites by zero-sum generators. The vectors $e_1 - e_3, e_2 - e_3$ generate every zero-sum displacement modulo n; each target has a rung. No strong edge leaves a paired layer. Each component therefore has $2n^2$ vertices.

Parametrize its source by $x_s(y) = (y_1, y_2, s - y_1 - y_2)$ and its target by $x_s(y) + e_3$, $y \in (\mathbb{Z}/n\mathbb{Z})^2$. For values $b(y), d(y)$ the exact component energy is

$$E_s = \sum_y |b(y) - d(y)|^2 + \sum_{y,j=1,2} |b(y) - d(y + e_j)|^2.$$

Thus $Q_H = H \sum E_s$, summing channels and spatial components. In the unitary two-dimensional Fourier transform, with normalization n^{-1}, the Hermitian matrix has diagonal 3 and off-diagonals $-A, -\overline{A}$, where $A = 1 + \mathrm{e}^{i\theta_1} + \mathrm{e}^{i\theta_2}$. Its eigenvalues are $3 \pm |A|$, and

$$9 - |A|^2 = |1 - \mathrm{e}^{i\theta_1}|^2 + |1 - \mathrm{e}^{i\theta_2}|^2 + |\mathrm{e}^{i\theta_1} - \mathrm{e}^{i\theta_2}|^2. \tag{2}$$

At zero frequency the eigenvalues are $0, 6$; only the constant on both layers is omitted below.

Lemma 2.1. *Put $L = 1 + \log n$. On one actual strong component, for any f of total mean zero and any vertex v,*

$$\|f\|^2 \le (3n^2/8)E_s(f), \qquad |f(v)|^2 \le G_n E_s(f), \qquad G_n \le 2L,$$

where the actual inverse diagonal is

$$G_n = \frac{1}{n^2}\left[\frac{1}{12} + \sum_{k \neq 0} \frac{3}{9 - |A(k)|^2}\right]. \tag{3}$$

Proof. Choose centered residues $r_j \in [-n/2, n/2 - 1]$. The finite sine bound gives $|1 - \mathrm{e}^{2\pi i r_j/n}|^2 \ge 16 r_j^2/n^2$. Equation (2) and $3 + |A| \le 6$ imply $3 - |A| \ge 8/(3n^2)$ at nonzero frequency. The zero-frequency antisymmetric mode also obeys this gap, proving the norm bound.

Diagonalize the finite positive-energy coordinates and apply weighted Cauchy–Schwarz to point evaluation. At nonzero frequency the inverse matrix diagonal is $3/(9 - |A|^2)$. The remaining zero-frequency mode contributes $(1/2)/6 = 1/12$. Squared Fourier normalization is $1/n^2$, proving (3). Discarding only the third nonnegative term in (2) gives

$$G_n \le \frac{1}{12n^2} + \frac{3}{16} \sum_{r \neq 0} (r_1^2 + r_2^2)^{-1}.$$

The shell $\max(|r_1|, |r_2|) = j$ has at most $8j$ points and denominator at least j^2. Consequently the sum is at most $8 \sum_{j=1}^{n/2} j^{-1} \le 8(1 + \log(n/2))$, and $G_n \le 2L$ follows. The possibly smaller Nyquist shell only improves this finite upper bound. □

3 Permutation coverage gives logarithmic coercivity

Theorem 3.1. *If F is a permutation, then for every ambient field u,*

$$Q(u)/V + \frac{4cn^2}{9L}\|C_t u\|^2 \ge \frac{c}{9L}\|u\|^2/V. \tag{4}$$

On the actual observation kernel there is the stronger estimate

$$C_t u = 0 \quad \Longrightarrow \quad Q(u) \ge \frac{c}{5L}\|u\|^2. \tag{5}$$

Both hold on U by restriction, without invariance of rectangular averaging.

Proof. For each channel and strong component select the unique actual singleton $v_s = (i_s, t_{i_s}, 0)$ for which $F(i_s) = s$ is its source residue. Let a scalar component restriction have mean m and residual $r = f - m$. If the observations vanish, $m = -r(v_s)$. Lemma 2.1 and orthogonal counting decomposition yield

$$\|f\|^2 = \|r\|^2 + 2n^2|m|^2 \le [3/8 + 4L]n^2 E_s(f) \le 5n^2 L E_s(f).$$

Sum the original scalar components and use $Q \geq Q_H = cn^2 \sum E_s$.

For an arbitrary field, $m = f(v_s) - r(v_s)$, so

$$\|f\|^2 \leq 4n^2|f(v_s)|^2 + [3/8 + 8L]n^2 E_s(f) \leq 4n^2|f(v_s)|^2 + 9n^2 L E_s(f).$$

The selected original singleton rows are distinct, and their squared unnormalized values sum to at most $V\|C_t u\|^2$. Therefore $\|u\|^2 \leq 4n^2 V\|C_t u\|^2 + (9L/c)Q(u)$, which is (4). This pays the finite-penalty step directly on all fields. □

For dyadic n, the nonlinear example $t_i = 2i^2 \pmod n$ has permutation diagonal: $F(i) - F(j) = (i-j)[1 + 2(i+j)]$, and the bracket is odd and invertible. The map is nonaffine for $n \geq 8$; this elementary example is not claimed new. Theorem 3.1 does not attain target c, prove logarithmic necessity, or decide (1) for this example.

4 Full means exclude exactly strong-null invisible fields

Lemma 4.1. *If every adjacent midpoint of a finite real cyclic sequence $A_0, \ldots, A_{M-1}$ belongs to its range, the sequence is constant.*

Proof. The case $M = 1$ is immediate. Otherwise choose $\sigma(j)$ with $2A_{\sigma(j)} - A_j - A_{j+1} = 0$, cyclically. The integer coefficient matrix has row sums zero. Modulo 2 it is $I + T$, where T is the next-coordinate shift. Keep rows $0, \ldots, M-2$ and columns $1, \ldots, M-1$: the resulting matrix modulo 2 is lower bidiagonal with diagonal entries 1. Its integer minor is odd and nonzero, so real rank is at least $M - 1$. Constants already belong to its kernel, so that kernel consists only of constants. This does not require rational values or an injective selection σ. □

Theorem 4.2. *For every cut map t in the stated slice family, $\ker C_t \cap \ker Q_H = \{0\}$ on the whole real ambient space, hence on U. No parity, permutation or motion assumption on the cuts is required.*

Proof. Fix a scalar component in a channel of strong-source parity p. Zero strong energy makes it $f(S)$ with $f(s) = f(s+1)$ for $s \equiv p$, by the actual graph connectivity in Section 2. All cell sums zero give global sum zero and hence $\sum_s f(s) = 0$. Choose a periodic primitive $G(s+1) - G(s) = f(s)$. Strong locking implies $G(s+1) = [G(s) + G(s+2)]/2$ when $s \equiv p$.

The actual singleton gives $f(F(i)) = 0$. The actual left cell, or the identity when $t_i = 0$, gives

$$0 = \sum_{j=0}^{t_i - 1} f(i+j) = G(F(i)) - G(i).$$

This remains valid at $t_i = n - 1$; no wrapping cell is introduced. For any $i \equiv 1 - p$, if $F(i) \equiv p$, then $G(i)$ is already a value on parity p. Otherwise $f(F(i)) = 0$ makes $G(F(i)) = G(F(i) + 1)$, also a value on parity p. Thus every value of G on parity $1 - p$ belongs to its range on parity p. These are precisely the adjacent midpoints of $A_j = G(p + 2j)$, $0 \leq j < n/2$. Lemma 4.1 makes G, and then f, constant and zero, respectively. Apply this to every original component. □

This theorem excludes one exact-zero mechanism, not a small positive energy. It supplies no quantitative size-uniform bound and does not remove the permutation assumption from Theorem 3.1.

5 Low coverage admits a true coexact invisible field

Theorem 5.1. *Let $n = 2^m \geq 1024$, $K = |\operatorname{im} F|$, and suppose*

$$K \leq n/32, \qquad T = \log(n/64) \geq 2^{34}.$$

There is a real nonzero $z \in U$ with $C_t z = 0$ and $Q(z) < (c/2)\|z\|^2$. Consequently no finite penalty vanishing at zero restores target c for this map. Within this slice family and at these sizes, target c requires $K > n/32$; neither sufficiency nor optimality is asserted.

We give the full construction. Write $\delta_j = I - \tau_{-e_j}$, and let $\delta = I - \tau_{-1}$ on a profile of S. For a scalar potential Ψ set

$$\mathcal{P}\Psi = (0, \delta_2\delta_3^2\Psi, -\delta_2^2\delta_3\Psi).$$

This has the actual two-form preimage $A_{23} = \delta_2\delta_3\Psi$, $A_{12} = A_{13} = 0$, because $(d_1^* A)_2 = \delta_3 A_{23}$ and $(d_1^* A)_3 = -\delta_2 A_{23}$. Use this field in channel 1 and zero in channel 2. Every original rectangular sum telescopes to four endpoint values of $\delta_3\Psi$ or $-\delta_2\Psi$. It suffices to annihilate both gradients at

$$x_1 = i, \qquad x_2 \in \{-1, t_i - 1, t_i\}, \qquad x_3 \in \{-1, 0\}. \tag{6}$$

Duplicated tags at endpoint cuts only remove requirements.

5.1 A derived profile kernel with a uniform inverse

Put $B = \operatorname{im} F$, $Z = B + \{-3, -2, -1, 0\}$, so $|Z| \leq 4K$. Use unitary one-dimensional characters $e_p(s) = n^{-1/2}\mathrm{e}^{2\pi i p s/n}$. For $r = 0, \ldots, 4K$, let $p_r = n/4 + r$, $h = n/2$, $\zeta_r = \mathrm{e}^{2\pi i p_r/n}$ and $b(\zeta) = 1 - \zeta^{-1}$. Define

$$\phi_r = (1 + \zeta_r)e_{p_r} + (1 - \zeta_r)e_{p_r+h},$$
$$\psi_r = \frac{1 + \zeta_r}{b(\zeta_r)^3}e_{p_r} + \frac{1 - \zeta_r}{b(-\zeta_r)^3}e_{p_r+h}.$$

The $4K+1$ first frequencies lie in $[n/4, 3n/8]$, including the endpoint when $K = n/32$; partners lie in $[3n/4, 7n/8]$. They are distinct and nonzero. The phases lie in $[\pi/2, 3\pi/4]$, so both multipliers have modulus at least $1/2$. Disjoint supports give independence and, for every combination,

$$\phi = \delta^3\psi, \qquad \|\psi\|_S^2 \leq 64\|\phi\|_S^2. \tag{7}$$

Evaluation on Z has at most $4K$ outputs. Rank–nullity therefore supplies a nonzero ψ in this $4K+1$ dimensional space with $\psi|_Z = 0$. The nonzero multipliers make $\phi \neq 0$. No conditioning or coefficient bound for the evaluation matrix is assumed.

For odd s, each ϕ_r has zero forward difference, since its parity-pair bracket is $(\zeta - 1)(1 + \zeta) - (-\zeta - 1)(1 - \zeta) = 0$. Thus $z_0 = (0, \phi(S), -\phi(S)) = \mathcal{P}[\psi(S)]$ is exactly strong locked, and

$$\|z_0\|^2 = 2n^2\|\phi\|_S^2 > 0. \tag{8}$$

5.2 Finite cutoff and every original mean

For $y = (x_2, x_3)$ put $r(y) = \max(d_n(x_2, 0), d_n(x_3, 0))$, with cyclic integer distance, and $R = n/8$. Define

$$\chi(y) = \begin{cases} 0 & r \leq 8, \\ \log(r/8)/T & 8 < r < R, \\ 1 & r \geq R. \end{cases}$$

Let $E_\chi = \sum_{y,j=2,3} |\chi(y+e_j) - \chi(y)|^2$. A nonzero edge crosses radii $j, j+1$, $8 \le j < R$, with difference at most $1/(Tj)$. The shell has $8j$ points. Counting the two outgoing labels from both shells bounds relevant edges by $16j + 16(j+1) \le 48j$. Hence

$$E_\chi \le \frac{48}{T^2} \sum_{j=8}^{R-1} \frac{1}{j} \le \frac{96}{T}, \quad \sum_y \chi^2 \ge \frac{3n^2}{4}, \quad \|\delta^\beta \chi\|_{2D}^2 \le 4^{|\beta|-1} E_\chi \quad (|\beta| \ge 1). \tag{9}$$

The last bound retains one first difference and bounds each other difference by operator norm 2. Backward and forward first norms agree by reindexing. The mass estimate discards the square of radius R, whose size $(2R+1)^2 \le n^2/4$. Torus seam edges are outside the varying region.

Take $\Psi = \psi(S)\chi$ and $z = \mathcal{P}\Psi$. At fixed tags $x_2 = -1, x_3 = -1, 0$, all potential values used by the two first gradients have transverse radius at most 2, so vanish. At moving tags $x_2 = t_i + u, x_3 = v$, $u, v \in \{-1, 0\}$, the values and their backward neighbors have sum coordinates in $F(i) + \{-3, -2, -1, 0\} \subset Z$. They vanish irrespective of the cutoff. Thus (6) holds and every original normalized cell mean is zero. The displayed preimage already pays actual coexactness; no divergence converse is used.

5.3 Nonzero norm and original source energy

Write $z = \chi z_0 + E$. For either order-three derivative, the exact backward Leibniz formula has terms

$$\binom{\alpha}{\beta} [\tau_{-\beta} \delta^{\alpha-\beta} \psi(S)] \delta^\beta \chi.$$

Beyond $\beta = 0$ their total multiplicity is 7 per nonzero component. At fixed transverse y, the x_1 sum of any shifted profile square is its one-dimensional norm. A term with $b = |\beta| \ge 1$ has squared norm at most $4^{3-b}\|\psi\|_S^2 \, 4^{b-1} E_\chi = 16\|\psi\|_S^2 E_\chi$. Seven-term Cauchy and two components, then (7)–(8), give

$$\|E\|^2 \le 2^{11} \|\psi\|_S^2 E_\chi \le \frac{2^{16} E_\chi}{n^2} \|z_0\|^2. \tag{10}$$

This uses the norm of the entire profile, not a sum of absolute Fourier coefficients; there is no factor depending on K. The mass bound gives $\|\chi z_0\| \ge 3\|z_0\|/4$. Our threshold implies $E_\chi < 2^{-26} < 1$; $n \ge 1024$ then gives $\|E\| \le \|z_0\|/4$ and

$$\|z\|^2 \ge \|z_0\|^2/4 > 0. \tag{11}$$

At odd sources the strong difference of χz_0 is just a cutoff difference times $z_0(x)$. Summing profiles in x_1 and dropping only the odd restriction yields $Q_H(\chi z_0) \le 2H\|\phi\|_S^2 E_\chi = cE_\chi \|z_0\|^2$. The full unweighted three-direction gradient is at most 12 times counting norm, so $Q_H(E) \le 12H\|E\|^2 < 2^{20} cE_\chi \|z_0\|^2$. The ambient positive quadratic form and (11) give

$$Q_H(z) < 2^{22} cE_\chi \|z_0\|^2 \le 2^{24} cE_\chi \|z\|^2 < (c/4)\|z\|^2.$$

Also $Q_a(z) \le 12a\|z\|^2 = (c/4)\|z\|^2$. Thus the complex field has quotient below $c/2$. Real and imaginary potential parts are actual coexact invisible fields; their energies and norm squares add. Some nonzero real part has no larger quotient. This proves Theorem 5.1. Theorem 4.2 implies its strong energy is strictly positive; it was not used to construct the witness or its upper bound.

The threshold is enormous and conservative, with no practical-size claim. For example $t_i = ((i \bmod 2) - i) \bmod n$ has $F(i) = i \bmod 2$, is nonaffine, and has exactly $36n - 24$ full rows. Only $i = 0, 1$ have zero cuts, and none has cut $n - 1$. It is covered by Theorem 5.1. Permutations, including $i + 2i^2$, are outside its low-coverage hypothesis.

6 Energy stability of the actual coexact projection

Let P be counting orthogonal projection onto U, independently in the two original channels. At nonzero physical Fourier frequency k put $g_j = \mathrm{e}^{2\pi i k_j/n} - 1$. Its symbol is

$$P_g = I - gg^*/|g|^2, \qquad P_0 = 0. \tag{12}$$

Indeed the curl symbol is $g\wedge$, and for $g^*v = 0$ the actual two-form $A_{jl} = (g_j v_l - g_l v_j)/|g|^2$ has adjoint v. Conversely antisymmetry gives $g^* d_1^* = 0$. At zero frequency the image is zero. Symbols conjugate at opposite frequencies, so P preserves real fields.

Theorem 6.1. *For every ambient field, $Q_H(Pu) \le 4Q_H(u)$ and $Q_a(Pu) \le 4Q_a(u)$. The same conclusion holds for any nonnegative scalar parity weights a, H, with constant independent of contrast and n.*

Proof. For nonzero $g, h \in \mathbb{C}^3$ and spatial vectors A, B, write $R_j = g_j A - h_j B$ and $s^2 = 1 - |g^*h|^2/(|g|^2|h|^2)$. The difference of the two rank-one complementary projections has norm s. To verify this over $\mathbb{C}$, adjust the phase of $h/|h|$ so its inner product with $g/|g|$ is $\alpha \ge 0$, and write it $\alpha e + sv$. On their two-dimensional span the rank-one projection difference is $\left(\begin{smallmatrix} s^2 & -\alpha s \\ -\alpha s & -s^2 \end{smallmatrix}\right)$, whose square is $s^2 I$; it is zero on the orthogonal complement. Parallel vectors give zero difference directly.

Completing the square, with the Hermitian scalar g^*h, gives

$$\sum_j \|R_j\|^2 = |g|^2 \|A - (g^*h)B/|g|^2\|^2 + |h|^2 s^2 \|B\|^2.$$

The projected residual is $R'_j = P_g R_j + h_j(P_g - P_h)B$. Each tuple on the right has norm at most $\|R\|$, so $\|R'\| \le 2\|R\|$. Replacing B by $-B$ gives the same estimate for the plus residual. If either symbol is zero, the actual $P_0 = 0$ reduces the estimate to norm contraction of the other projection.

Pair physical frequencies $k, k+\eta$, $\eta = (n/2, n/2, n/2)$, once each. Their forward symbols are g, h, with $h_j = -\mathrm{e}^{2\pi i k_j/n} - 1$. On an odd source their gradient tuple is $e_k(x)(g_j A - h_j B)$; on even sources it is $e_k(x)(g_j A + h_j B)$. Distinct unordered pairs are orthogonal on each parity sublattice, since its character sum vanishes unless the frequency difference is 0 or η. The restricted squared character norm is exactly $1/2$. Each parity energy is therefore half the sum of these squared tuples. Apply the residual estimate and the nonnegative original weight. The reversed channel only exchanges the two signs. □

Norm contraction supplies no lower bound for $\|Pu\|$. The theorem also does not preserve any observation row or assert energy commutation. The next example makes the remaining mean obstruction explicit.

7 Loss of means and a necessary correction cost

Proposition 7.1. *For $4 \mid n$, the real ambient field in channel 1*

$$u(x) = -2\sin(\pi x_1/2)\sin(\pi x_3/2)e_1$$

has $C_t u = 0$ for every cut map, but $(Pu)_3(0, t_0, 0) = -1/2$. Thus its true coexact projection is visible for every map in the specified full slice family.

Proof. The field is independent of x_2, vanishes at $x_3 = 0$, and its sine sum over $[1, n-1]$ is zero. Every original cell mean therefore vanishes. Use the complex precursor $w = \mathrm{e}^{i\pi x_1/2}(\mathrm{e}^{i\pi x_3/2} - \mathrm{e}^{-i\pi x_3/2})e_1$, whose real part is u. At its two frequencies the symbols are $g_+ = (i-1, 0, i-1)$

and $g_- = (i-1, 0, -i-1)$, both squared norm 4. At $x_3 = 0$, the first two components of $P_{g_+}e_1 - P_{g_-}e_1$ vanish, and the third is

$$-[(i-1) - (-i-1)]\overline{(i-1)}/4 = (-1+i)/2.$$

The operator P is real, so at $x_1 = 0$ its real third value is $-1/2$, independent of x_2. The actual normalized singleton row is $-1/(2\sqrt{V}) \neq 0$, including endpoint cuts. □

Theorem 7.2. *For any permutation diagonal and any $w \in U$, with channel-1 components w_1, w_2, w_3, select v_s as in Theorem 3.1 for odd s. Then*

$$H \sum_{s \text{ odd}} |w_1(v_s) + w_2(v_s) + w_3(v_s)|^2 \leq 6LQ_H(w). \tag{13}$$

For the input in Proposition 7.1, every coexact r satisfying $C_t(Pu + r) = 0$ obeys

$$Q_H(r) \geq \frac{c\|u\|^2}{48L}, \qquad \|u\|^2 = n^3. \tag{14}$$

Equation (14) *concerns r only, not the final field $Pu + r$.*

Proof. For a true coexact field, $d_0^* w = 0$ follows from $d_1 d_0 = 0$, and every spatial component has global sum zero by its actual backward-difference adjoint formula. Let $W_j(s) = \sum_{S(x)=s} w_j(x)$ and $T(s) = \sum_j W_j(s)$. Summing divergence on a layer yields $T(s) - T(s-1) = 0$. Its global sum is zero, hence $T(s) = 0$. Thus $f = w_1 + w_2 + w_3$ has zero mean on every individual sum layer and on every strong pair. No converse divergence characterization is needed.

Lemma 2.1 gives $|f(v_s)|^2 \leq G_n E_s(f)$, and edgewise Cauchy gives $E_s(f) \leq 3\sum_j E_s(w_j)$. The selected points are in distinct strong components. Summation and $G_n \leq 2L$ give (13). The other original channel contributes nonnegative energy and is not mixed into the trace. The spatial sum is a derived combination of three actual rows, with its factor 3 paid.

At every $x_1 = i, x_3 = 0$, the first two components of Pu vanish and its third component is $\mathrm{Re}[e^{i\pi i/2}(-1+i)/2] \in \{-1/2, 1/2\}$. Any repair must have $r(v_s) = -Pu(v_s)$, separately in each original spatial row. There are $n/2$ selected points, so the squared spatial-sum trace of r totals $n/8$. Equation (13) gives $Q_H(r) \geq Hn/(48L) = cn^3/(48L)$. Finally $\|u\|^2 = 4n(\sum_i \sin^2(\pi i/2))^2 = n^3$. □

This excludes a strong-null coexact correction for this input. It excludes neither paid corrections nor cancellation of their energy against Pu. For example $r = -Pu$ is admissible and leaves final field zero. We supply no nonzero corrected invisible witness here.

8 Remaining questions and review context

The original coefficient c for the nonlinear permutation $F(i) = i + 2i^2$ remains undecided. Theorem 3.1 is a logarithmic lower bound; Theorem 5.1 is a counterexample for a different coverage regime. Neither establishes necessity of a logarithmic observation count or a matching logarithmic kernel rate for permutation diagonals. The projection and repair results explain why an ambient construction cannot simply be transferred into the actual domain, but they do not close that gap.

All results concern finite counting models. There is no continuum limit, physical spectral-gap assertion, practical threshold, optimal coefficient, solver guarantee or claim that the finite estimates solve Yang–Mills. The programme's separate PR29 record remains 20/41 with TermSource = 0.

Model assessments and review context

The mathematical source notes received complete, exact-source analytical checks by an involved Codex/OpenAI assistant with shared author/programme and prior-audit context. The exact served backend, version and effort were not exposed and are not inferred. These checks are not blind, independently established peer review, formal proofs or numerical experiments. Their full reports, assignments, objections and read scopes are retained in the source package; absence of a detected defect is not an originality finding.

An exact-PDF assessment, if supplied in the deposit metadata, is a separate artifact tied to that PDF's hash. No numerical assessment of an earlier working paper is transferred to this manuscript. Source checks, PDF/model assessment, deposit, working-paper publication, human peer review and editorial acceptance are distinct statuses. The author and programme control the manuscript. The author is also the owner, creator and administrator of AIRR, the intended working-paper destination; this author–operator role is disclosed and does not constitute independent editorial or human peer review.

Source and reproducibility

The production source commit is `bec9892206b0e9d82bc96ed247895d68d9c25e48`. The accompanying package retains original sources 266–273 and their full reports with raw-byte SHA-256 manifests. This manuscript contains the finite definitions and proofs needed for its statements. Linux production logs and exact-PDF/render hashes, when generated, are separate evidence; a successful existing-core Lean/oracle CI run does not formally verify these plain-text analytical proofs. No source-stage claim of numerical or formal verification is made.

References

[1] G. F. Lawler and V. Limic, *Random Walk: A Modern Introduction*, Cambridge University Press, 2010. Author-hosted draft: `https://www.math.uchicago.edu/~lawler/srwbook.pdf`. Classical background, particularly two-dimensional potential theory; our finite Green normalization is derived in Section 2.

[2] D. N. Arnold, R. S. Falk and R. Winther, Finite element exterior calculus: from Hodge theory to numerical stability, *Bulletin of the American Mathematical Society* 47 (2010), 281–354. `https://arxiv.org/abs/0906.4325`. The cochain projections there are not identified with the projection in Section 6.

[3] L. Eriksson, *Canonical and unrestricted observations in a finite checkerboard gradient model*, AIRR working paper ARR-2026-597T88D7WZ98HR0Z, version 2, 2026. `https://airr.science/papers/ARR-2026-597T88D7WZ98HR0Z/versions/v2/`. Working-paper status and its own assessments do not imply acceptance or carry over to the present PDF.