Canonical rectangle observations for matrix gradients: fixed weak spaces and checkerboard counterexamples

Lluis Eriksson
Independent researcher

Source preparation: 11 November 2026
Revision 250

Abstract

We study finite positive matrix-weighted gradients on a periodic lattice, with observations restricted to normalized rectangular means in a fixed internal basis. For a spatially constant rank-q weak projector in $\mathbb{R}^g$ satisfying $P_{\alpha\alpha} \le \Lambda q/g$, the minimum actual restoring rank on the flat coexact domain has order $qn^D + gn^{D-2}$, uniformly in q. An explicit choice of original-channel rectangle partitions attains this order. For this constructed map, the least quadratic coefficient has order $c\min(g/q, n^2)$. The construction uses a stated theorem of Marcus–Spielman–Srivastava, augmented to control both Gram matrices and cardinalities, and derived finite Neumann estimates. Spatial constancy is substantive. On an exact positive two-channel checkerboard with the same local spectra and fixed basis, full output counts are $O(n \log n)$ in dimension three and $O(n^2)$ in dimension four, with finite size-dependent penalties. We derive the point estimates, actual strong graphs, integer meshes and original-rectangle realization. A separate plane construction works at every fixed nonzero angle between the two weak lines. These are finite observation-class results; optimal checkerboard rank, practical conditioning and historical priority remain unresolved.

Contents

1 The observation class and the main results **2**
1.1 Actual finite gradients and the coexact domain 2
1.2 Canonical observations and normalization 3
1.3 Sharp order for a fixed weak space 3
1.4 Spatial variation with the same local spectra 4
1.5 The unchanged capped budget and proof roadmap 5

2 Finite spectral and projection facts **5**

3 A fixed coherent weak space **7**
3.1 The primary MSS input and an added count coordinate 7
3.2 Recursive whitening with a noncircular error bound 8
3.3 Stopping, original groups and the full row count 9
3.4 All-field coercivity with one strong-energy budget 10
3.5 Actual-rank lower bound and the cancellation of weak dimension 10
3.6 Least coefficient for the constructed fixed weights 11

4 Actual checkerboard graphs and plane observations **12**
4.1 Selected strong cycles and an exact Dirichlet margin 12
4.2 Every fixed nonzero angle: proof coordinates only 12

5 Distributed point anchors in dimensions three and four **13**
5.1 The full strong graph in original coordinates 14
5.2 Balanced quantiles and both layer norms 15
5.3 Dimension three: the critical logarithmic mesh 15
5.4 Dimension four: a three-dimensional source mesh 16
5.5 Finite completion, constants and actual restricted rank 16

6 Scope, classical inputs and review context **17**
6.1 What the comparison establishes 17
6.2 Primary comparisons and their actual extent 17

A Source correspondence and retained corrections **19**

1 The observation class and the main results

The problem is to restore a fixed coercivity target using few permitted observations. Two distinctions matter: the dimension of a displayed output space versus its actual rank on the stated domain, and a finite penalty versus a penalty bounded independently of the lattice size. We keep both explicit. The proofs concern one positive family of finite operators, rather than a uniform theorem for arbitrary matrix profiles.

1.1 Actual finite gradients and the coexact domain

Fix integers $D \geq 3$, $n \geq 2$ and $g \geq 1$, and write

$$X = (\mathbb{Z}/n\mathbb{Z})^D, \qquad V = n^D, \qquad \mathcal{H}_g = (\mathbb{R}^D \otimes \mathbb{R}^g)^X.$$

All unqualified inner products and norms use counting measure. Put $\partial_i f(x) = f(x+e_i) - f(x)$, retaining every positive forward label, including period-two multiplicities. For $z \in \mathcal{H}_g$ let

$$\mathcal{E}(z) = \sum_{x \in X} \sum_{i,k=1}^{D} \sum_{\alpha=1}^{g} |\partial_i z_{k,\alpha}(x)|^2.$$

For a real one-form v define $(d_1 v)_{ij} = \partial_i v_j - \partial_j v_i$ for $i < j$. The adjoint uses the counting inner products, with each unordered pair $i < j$ counted once. Our domain is

$$\mathcal{U}_0 = \operatorname{im} d_1^*, \qquad \mathcal{U} = \mathcal{U}_0^g \subset \mathcal{H}_g, \qquad d_0 := \dim_{\mathbb{R}} \mathcal{U}_0 = (D-1)(V-1). \tag{1}$$

The dimension identity is proved in Lemma 2.1. No artificial componentwise replacement of this domain is used.

For a positive symmetric $g \times g$ field $B(x)$ set

$$Q_B(z) = \sum_{x \in X} \sum_{i,k=1}^{D} (\partial_i z_k(x))^T B(x) \partial_i z_k(x). \tag{2}$$

The coefficient is evaluated at the actual forward *source* x. The full gradient-stack profile is $\mathcal{A}(x) = I_{D^2} \otimes B(x)$; there are $D^2 g$ local eigenvalues, while the field has $h = Dg$ components. This notation does not assert an identity with an interacting or continuum operator.

1.2 Canonical observations and normalization

For each original internal channel α, choose a partition $\mathcal{P}_\alpha$ of $\{0,\dots,n-1\}^D$ into nonempty integer rectangles $S=\prod_i\{a_i,\dots,b_i\}$. Different channels may have different partitions; within a channel the partition is shared across all spatial components k. The canonical observation map has rows

$$(Cz)_{\alpha,k,S}=V^{-1/2}|S|^{-1/2}\sum_{x\in S}z_{k,\alpha}(x),\qquad N_{\mathrm{out}}=D\sum_\alpha|\mathcal{P}_\alpha|. \tag{3}$$

These are unrotated, fixed-weight component means. No mixed internal row or subsequent dimension-reducing channel mixture is part of this class. Let $\mathcal{R}$ be the counting orthogonal block projection replacing each component by its arithmetic mean on its channel's rectangles. Then

$$\|Cz\|^2=V^{-1}\|\mathcal{R}z\|^2. \tag{4}$$

The *actual rank* is $r(C)=\mathrm{rank}(C|_{\mathcal{U}})$, which need not be N_{out}. We call C restoring at target $c>0$ if some finite $K\ge 0$ gives

$$V^{-1}Q_B(z)+K\|Cz\|^2\ge cV^{-1}\|z\|^2\quad(z\in\mathcal{U}). \tag{5}$$

We separately state when the same inequality holds on all of $\mathcal{H}_g$. On these *same* rows the counting-energy convention uses coefficient VK. Arbitrary row rescaling would change the penalty question and is excluded. The projection $\mathcal{R}$ is not assumed to preserve $\mathcal{U}$.

1.3 Sharp order for a fixed weak space

Let $1\le q\le g$, $\Lambda\ge 1$, and let P be a spatially constant rank-q orthogonal projector satisfying

$$P_{\alpha\alpha}\le\Lambda q/g\quad(1\le\alpha\le g). \tag{6}$$

With $c>0$ fixed put

$$a=\frac{c}{16D},\qquad H=cn^2,\qquad B=HI_g-(H-a)P,\qquad b=\min\{g/q,n^2\}. \tag{7}$$

There is no lower bound on nonzero diagonal entries of P and no prescribed well-conditioned subset of channels.

Theorem 1.1 (Fixed weak space)**.** *For* (7)–(6) *there are coefficient-only canonical partitions such that*

$$N_{\mathrm{out}}\le L_{D,\Lambda}(qV+gV/n^2),\qquad Q_B(z)/V+32cb\|Cz\|^2\ge c\|z\|^2/V\quad(z\in\mathcal{H}_g). \tag{8}$$

One valid constant is

$$C_0=2^{24},\quad A_D=\frac{D2^D(17/16)}{2^{D-2}-1},\quad L_{D,\Lambda}=A_D+4352D+\frac{17}{4}DC_0\Lambda.$$

Every restoring internally channel-separable linear map, even beyond canonical rectangles, has actual rank

$$r\ge\max\left\{qd_0,\frac{gd_0}{2\Lambda(1+8Dn^2)}\right\}. \tag{9}$$

This rank bound only uses the restoring inequality at zero observation; it also holds for any finite penalty $\Phi(Cz)$ with $\Phi(0)=0$ in place of the quadratic penalty. Consequently the minimum canonical restoring actual rank has order $qV+gV/n^2$, with constants depending

only on D, Λ. For the particular constructed canonical map, its least quadratic coefficient on either $\mathcal{U}$ or all fields satisfies

$$\frac{3(D-1)}{16\Lambda L_{D,\Lambda}} cb \le K_{\min} \le 32cb. \tag{10}$$

All constants are independent of q, g, n, c.

Here channel-separable means a direct sum of linear maps, one acting on each original channel's entire D-component field. No lower bound for internally mixed maps is asserted. The coefficient order in (10) concerns a fixed constructed map, not an upper bound for every map of that rank. The construction is a finite existence prescription; neither efficient search nor a practical condition number is established.

1.4 Spatial variation with the same local spectra

For even n let $p(x) = \sum_i x_i \bmod 2$, which is well-defined on the torus. For $g = 2, q = 1, \Lambda = 2$, set

$$P(x) = \begin{cases} e_1 e_1^T, & p(x) = 0, \\ e_2 e_2^T, & p(x) = 1, \end{cases} \qquad B(x) = HI_2 - (H-a)P(x), \tag{11}$$

with the same a, H from (7). Only spatial constancy of P has been removed; local eigenvalues, positivity, upper coherence, observed basis, canonical weights and actual domain are unchanged.

Theorem 1.2 (Orthogonal checkerboard)**.** *For* (11) *and every even $n \ge 2$:*

(i) For every $D \ge 3$ a canonical map with $N_{\text{out}} = 2D(n^{D-1}+1)$ restores c on all fields with a finite coefficient.

(ii) For $D = 3$ a canonical map with $N_{\text{out}} \le 2^{19} n(1 + \ln n)$ restores c on all fields.

(iii) For $D = 4$ a canonical map with $N_{\text{out}} \le 2^{25} n^2$ restores c on all fields.

In (ii)–(iii), the non-singleton construction has the sufficient normalized coefficient

$$K = M + \frac{2M^2}{3c}, \qquad M = 4DH + c; \tag{12}$$

the explicit singleton fallbacks use $K = c$. Actual rank on $\mathcal{U}_0^2$ is at most the corresponding full count. These are upper bounds, not optimality, matching lower bounds, logarithmic necessity or uniform-penalty statements.

In particular a universal per-profile $\Omega(V)$ rank lower cannot be extended from constant P merely by retaining the same positive local spectra and pointwise coherence. This does not refute a worst-case lower over a larger class containing constant profiles, or a lower with an additional uniform penalty restriction.

Theorem 1.3 (Fixed-angle plane construction)**.** *Replace the two projectors in* (11) *by any two fixed distinct weak lines with principal angle $\theta \in (0, \pi/2]$, in the original observed basis. Let $\kappa = 1 - |\cos\theta|$ and $m_\theta = \lceil 4/\sqrt{\kappa} \rceil$. For every $D \ge 3$ and even $n \ge 2$ there is an all-field restoring canonical map with*

$$N_{\text{out}} \le 2Dm_\theta(n^{D-1}+1)$$

and a finite penalty. When $n \ge m_\theta$, (12) *suffices; otherwise the singleton fallback has $K = c$.*

For fixed θ, D this is $o(V)$, with constants allowed to diverge as $\theta \downarrow 0$. We prove no n-dependent-angle order and no extension of the point-anchor orders in Theorem 1.2(ii)–(iii) to nonorthogonal lines. That separate question is outside the present theorem list.

1.5 The unchanged capped budget and proof roadmap

To fix the normalization inherited by these examples, define for positive local gradient-stack eigenvalues $\lambda_j(x)$

$$\alpha_D = \left(\frac{3}{4}\right)^D \frac{D}{2}\left(\frac{D}{D-2}\right)^{D-1}, \qquad F_c(\mathcal{A}) = \sum_{x,j} \min\{1, \alpha_D (c/\lambda_j(x))^{D/2}\}. \tag{13}$$

For the fixed-space family, a has multiplicity qD^2 and H has multiplicity $(g-q)D^2$ per site. Its weak cap equals one, so

$$F_c(\mathcal{A}) = D^2 V[q + (g-q)\min\{1, \alpha_D n^{-D}\}] \le D^2(qV + \alpha_D g). \tag{14}$$

Indeed $\alpha_D (16D)^{D/2} \ge (3\sqrt{D})^D > 1$. Thus $h + F_c(\mathcal{A}) = O_D(qV+g)$; we do not identify it with the optimal rank expression $qV + gV/n^2$. For both checkerboards the same formula with $q=1, g=2$ is unchanged. In particular

$$\alpha_3 = 729/128, \quad \alpha_4 = 81/16, \qquad F_c = D^2 V[1 + \min\{1, \alpha_D n^{-D}\}], \quad h = 2D.$$

Every cap and multiplicity is retained.

Section 2 derives all elementary finite spectral and projection facts. Section 3 proves Theorem 1.1; its sole non-elementary input is the explicitly stated primary MSS theorem. Sections 4–5 prove the actual strong-graph, plane and point constructions. Scope, limited primary comparisons and exact source/review context are separated in Section 6.

2 Finite spectral and projection facts

This section derives the local estimates used later; a regional gap is not an additional hypothesis.

Lemma 2.1 (Actual Fourier domain and periodic gap)**.** *The dimension in* (1) *is* $d_0 = (D-1)(V-1)$*. Every element of* $\mathcal{U}_0$ *has mean zero. On every componentwise mean-zero field,*

$$\mathcal{E}(f) \ge \frac{16}{n^2}\|f\|^2, \qquad \textit{while on every field } \mathcal{E}(f) \le 4D\|f\|^2.$$

Proof. Use the counting-orthonormal characters $V^{-1/2}\exp(2\pi i \ell\cdot x/n)$ with $\ell \in \{0,\dots,n-1\}^D$. The forward difference multiplier is $\xi_i(\ell) = \exp(2\pi i \ell_i/n) - 1$. Its energy multiplier is $4\sum_i \sin^2(\pi \ell_i/n)$, which is at most $4D$. At a nonzero frequency at least one term is at least $4\sin^2(\pi/n) \ge 16/n^2$, by concavity of sine on $[0,\pi/2]$; the case $n=2$ gives equality and retains the actual forward labels.

The curl symbol is $v \mapsto \xi \wedge v$. When $\xi \neq 0$, its kernel is $\mathbb{C}\xi$, so the adjoint image has complex dimension $D-1$; at zero frequency that image is zero. Conjugate frequencies pair to give twice their complex dimension as a real dimension. At self-conjugate frequencies each ξ_i is real (0 or -2), and the real adjoint image has dimension $D-1$ directly. Summing gives $(D-1)(V-1)$, without an omitted period-two mode. The zero frequency is absent, proving the mean-zero assertion. The upper energy bound also follows from the counting isometry of every shift. □

Lemma 2.2 (Finite counting Neumann basis)**.** *On an* ℓ*-vertex path,* $\ell \ge 2$*, the counting-orthonormal vectors are*

$$\phi_0(t) = \ell^{-1/2}, \qquad \phi_j(t) = \sqrt{2/\ell}\cos\frac{\pi j(t+1/2)}{\ell} \quad (1 \le j < \ell),$$

with eigenvalues $4\sin^2(\pi j/(2\ell))$. *A singleton has only its constant mode. On any product rectangle with side lengths at most* L, *the counting variance is at most* $L^2/4$ *times its internal edge energy, also for vector fields; singleton factors contribute zero variance.*

Proof. For $\theta = \pi j/\ell$, interior second differences of $\cos(\theta(t+1/2))$ multiply it by $2(1-\cos\theta)$. The virtual endpoint values $f(-1) = f(0)$ and $f(\ell) = f(\ell-1)$ give precisely endpoint degree one. For an integer $1 \le r < 2\ell$,

$$\sum_{t=0}^{\ell-1} \cos\frac{\pi r(t+1/2)}{\ell} = \frac{\sin(\pi r)}{2\sin(\pi r/(2\ell))} = 0.$$

The denominator is nonzero in this range. Product-to-sum, with the $r = 0$ sum equal to ℓ, proves the stated orthogonality and normalization. The needed sums $j + j'$ are at most $2\ell - 2$. There are ℓ orthonormal vectors, so this is a complete basis.

Tensor products diagonalize the product-path energy. Since $\sin t \ge 2t/\pi$ on $[0, \pi/2]$, every nonconstant product mode has eigenvalue at least $4/L^2$. Removing the constant mode gives the variance bound by expansion in this finite basis. If the entire rectangle is a singleton there is no positive eigenvalue to invoke and the variance is zero. Sum the scalar inequalities for vector fields. No dimension factor is lost. □

For an integer $j \ge 0$ partition each coordinate chart into the nonempty intervals

$$\{\lfloor rn/2^j \rfloor, \dots, \lfloor (r+1)n/2^j \rfloor - 1\}, \qquad 0 \le r < 2^j,$$

and take their products. Call the partition $\mathcal{P}_j$ and its counting averaging projection Π_j. Every coarse endpoint is an even-index fine endpoint, so the partitions are nested after empty intervals are omitted. There are at most 2^{Dj} cells and their sides have length at most $L_j = \lceil n/2^j \rceil$. In particular Π_0 is global averaging.

Lemma 2.3 (Actual spatial pieces)**.** *Fix* $k \ge 0$. *The operators*

$$\Delta_j = \Pi_{j+1} - \Pi_j \quad (0 \le j < k), \qquad \Delta_k = I - \Pi_k$$

are mutually orthogonal projections with sum $I - \Pi_0$. *For every field* f,

$$\|\Delta_j f\|^2 \le \frac{L_j^2}{4}\mathcal{E}(f) \quad (0 \le j \le k). \tag{15}$$

For $j < k$, Δ_j *is a subprojection of* Π_{j+1}; *the tail is a subprojection of the singleton observation* I.

Proof. Nested averaging projections satisfy $\Pi_i\Pi_j = \Pi_{\min(i,j)}$. Their differences and the displayed tail are therefore orthogonal projections and telescope. Apply Lemma 2.2 on every parent cell of $\mathcal{P}_j$ and sum. Its internal edges are a subset of the original labelled periodic edges, counted once. Other edges add nonnegative energy. This gives $\|(I - \Pi_j)f\|^2 \le L_j^2\mathcal{E}(f)/4$. Each Δ_j is a subprojection of $I - \Pi_j$, proving (15), including the single tail. The subprojection statements follow from the same identities. Wrap edges are never inserted into a Neumann cell as fictitious interior edges. □

Lemma 2.4 (Finite kernel-to-all-field completion)**.** *Let* L *be a positive semidefinite operator on a finite real Hilbert space, let* R *be an orthogonal projection, and suppose* $\langle u, Lu\rangle \ge (c+\delta)\|u\|^2$ *on* $\ker R$, *with* $c, \delta > 0$. *If* $\|L - cI\| \le M$, *then*

$$\langle z, Lz\rangle + \left(M + \frac{2M^2}{\delta}\right)\|Rz\|^2 \ge c\|z\|^2 \quad \textit{for all } z. \tag{16}$$

For (2) *with* $0 < B(x) \le HI$, *one may take* $M = 4DH + c$.

Proof. Write $z = u + v$, $u = (I - R)z$, $v = Rz$, and $F = L - cI$. Then $\langle u, Fu\rangle \geq \delta\|u\|^2$, $\langle v, Fv\rangle \geq -M\|v\|^2$, and

$$2|\langle u, Fv\rangle| \leq 2M\|u\|\|v\| \leq \frac{\delta}{2}\|u\|^2 + \frac{2M^2}{\delta}\|v\|^2.$$

After adding the displayed penalty a nonnegative $\delta\|u\|^2/2$ remains. The final operator bound follows from Lemma 2.1's upper energy bound and $B(x) \leq HI$. Nothing requires R to preserve a later restricted domain. Constants must still be controlled by R or a kernel margin; they are not silently discarded by this lemma. □

Lemma 2.5 (Original-coordinate singleton isolation)**.** *Given A distinct points in the original D-dimensional chart, there is a partition into at most $1 + 2DA$ nonempty integer rectangles in which every specified point is a singleton.*

Proof. Start with the whole chart. In the rectangle containing the next point, cut before and after its first coordinate and retain the two possible side rectangles. Continue only within the middle slice, through the remaining coordinates. At most $2D$ side rectangles and the desired singleton replace one rectangle, a net increase at most $2D$. Omit empty pieces. Disjoint coverage follows from the first coordinate at which a point differs from the anchor. Earlier singleton anchors lie in other rectangles and remain unchanged; an already-singleton point costs zero. Iteration proves the count. This uses no global tensor-grid propagation of cuts or point separation. □

3 A fixed coherent weak space

We prove Theorem 1.1. Write $\mathsf{U} \in \mathbb{R}^{g\times q}$ for an orthonormal basis of $\operatorname{im} P$, so $\mathsf{U}\mathsf{U}^T = P$. One deterministic choice is Gram–Schmidt on the first independent columns of P, in original index order. If $u_\alpha \in \mathbb{R}^q$ is its αth row, then

$$\sum_{\alpha=1}^{g} u_\alpha u_\alpha^T = I_q, \qquad \|u_\alpha\|^2 = P_{\alpha\alpha} \leq d := \Lambda q/g. \tag{17}$$

These coordinates serve the proof and channel selection. The observed fields remain the original $z_{k,\alpha}$.

3.1 The primary MSS input and an added count coordinate

We use precisely the following classical theorem of Marcus–Spielman–Srivastava [1, Theorem 1.4, p. 3], in its version-pinned arXiv statement.

Theorem 3.1 (MSS finite-support theorem)**.** *Let independent finitely supported random vectors $v_1, \ldots, v_m \in \mathbb{C}^N$ satisfy $\sum_i \mathbb{E}(v_i v_i^*) = I_N$ and $\mathbb{E}\|v_i\|^2 \leq \varepsilon$ for every i, where $\varepsilon > 0$. With positive probability,*

$$\|\sum_i v_i v_i^*\| \leq (1 + \sqrt{\varepsilon})^2.$$

We use this published theorem, not a proof or computational verification of it. The simultaneous cardinality assertion needed next is not simply assumed from a frame-partition corollary.

Lemma 3.2 (Simultaneous Gram and count halving)**.** *Suppose* $\sum_{i\in I} w_i w_i^T = I_q$, $m = |I|$, *and* $\delta = \max_i \|w_i\|^2 + 1/m$. *There is a partition* $I = S_1 \sqcup S_2$ *such that, for both* $j = 1, 2$,

$$(\tfrac{1}{2} - e)I_q \le \sum_{i\in S_j} w_i w_i^T \le (\tfrac{1}{2} + e)I_q, \qquad \tfrac{1}{2} - e \le |S_j|/m \le \tfrac{1}{2} + e, \quad e = \sqrt{2\delta} + \delta. \tag{18}$$

All original rows, including zero rows, are counted.

Proof. Independently for every i, choose a uniform sign $\sigma_i \in \{-1, 1\}$ and a uniform label $a_i \in \{1, 2\}$, also independent within each pair. Set $b_i = (w_i, \sigma_i/\sqrt{m}) \in \mathbb{R}^{q+1}$ and $v_i = \sqrt{2}e_{a_i} \otimes b_i \in \mathbb{R}^{2(q+1)}$. The sign removes the off-diagonal covariance with the last coordinate:

$$\mathbb{E}(b_i b_i^T) = \operatorname{diag}(w_i w_i^T, 1/m), \qquad \sum_i \mathbb{E}(v_i v_i^T) = I_{2(q+1)}.$$

Each vector has four possible outcomes and squared norm $2(\|w_i\|^2 + 1/m) \le 2\delta$. Theorem 3.1 therefore supplies one common outcome with norm at most $(1 + \sqrt{2\delta})^2$. In each label block, the sum is $2\sum_{i:a_i=j} b_i b_i^T$. Its leading principal block gives the upper Gram bound, and its last scalar entry gives the upper count bound, both with $(1 + \sqrt{2\delta})^2/2 = 1/2 + e$. Complementing the *same* partition gives both lower bounds. Signs do not alter the original Grams or cardinalities. In particular a zero w_i still has its count coordinate. The finite set of feasible original-channel partitions has a lexicographically first member; this specifies existence without claiming an efficient search. □

3.2 Recursive whitening with a noncircular error bound

Start with $I_0 = \{1, \dots, g\}$ and write $G_t = \sum_{\alpha\in I_t} u_\alpha u_\alpha^T$, $m_t = |I_t|$. When G_t is positive, whiten *only for selection*: $w_\alpha = G_t^{-1/2} u_\alpha$. Apply Lemma 3.2, choose a deterministic child I_{t+1}, and set

$$\delta_t = \max_{\alpha\in I_t} \|G_t^{-1/2} u_\alpha\|^2 + 1/m_t, \qquad \eta_t = 2\sqrt{2\delta_t} + 2\delta_t.$$

Congruence by $G_t^{1/2}$ gives

$$\frac{1-\eta_t}{2} G_t \le G_{t+1} \le \frac{1+\eta_t}{2} G_t, \quad \frac{1-\eta_t}{2} m_t \le m_{t+1} \le \frac{1+\eta_t}{2} m_t. \tag{19}$$

Lemma 3.3 (Uniform conditioned prefixes)**.** *Put* $C_0 = 2^{24}$. *If* $4^k \le g/(C_0 \Lambda q)$, *then* $2k$ *successive halvings exist, and for every* $0 \le t \le 2k$,

$$\frac{15}{16} 2^{-t} I_q \le G_t \le \frac{17}{16} 2^{-t} I_q, \qquad \frac{15}{16} g 2^{-t} \le m_t \le \frac{17}{16} g 2^{-t}. \tag{20}$$

Proof. Assume only the already established bounds at a current stage. Then

$$\delta_t \le \frac{16}{15} 2^t (d + 1/g) \le \frac{32}{15} 2^t d,$$

because $\Lambda q \ge 1$. The current Gram is invertible, justifying that stage's whitening. Over any justified prefix of $t < 2k$, geometric sums give

$$\sum_{t<2k} \eta_t \le 12\sqrt{d4^k} + 5d4^k \le \frac{12}{4096} + \frac{5}{2^{24}} < \frac{1}{256}. \tag{21}$$

Here the displayed full sum denotes a majorant for every justified prefix, not an assumption that future stages exist. Indeed $\sum_{t<2k} 2^{t/2} = (2^k - 1)/(\sqrt{2} - 1)$ and

$$\frac{2\sqrt{2}\sqrt{32/15}}{\sqrt{2} - 1} < 12, \qquad 2(32/15) < 5.$$

For example $\sqrt{2} < 3/2$, $\sqrt{32/15} < 3/2$ and $(\sqrt{2}-1)^{-1} = \sqrt{2}+1 < 5/2$ already bound the first expression by $45/4 < 12$. Also $d4^k \le 2^{-24}$; the last numerator in (21) is $49157 < 65536$ over denominator 2^{24}.

For nonnegative errors with sum $S < 1$, $\prod(1-\eta_t) \ge 1-S$ and $\prod(1+\eta_t) \le 1/(1-S)$. Start from $G_0 = I_q, m_0 = g$. The current loose bounds give the next error bound; the resulting prefix products lie between $255/256$ and $256/255$, strictly within $15/16$ and $17/16$. Applying (19) proves (20) at the next stage. This completes a prefix induction and proves existence of all stages, rather than assuming their conditioning. Every $\eta_t < 1$, and both children have positive Gram and positive count. The case $k = 0$ contains no halving and is immediate. □

3.3 Stopping, original groups and the full row count

Set

$$M_* = \min\left\{\frac{n}{32}, \sqrt{\frac{g}{C_0\Lambda q}}\right\}. \tag{22}$$

If $M_* < 1$, give *all* channels singleton partitions. This priority covers $n < 32$ or $g < C_0\Lambda q$, including zero frame rows. It will be counted and checked separately. Otherwise put $k = \lfloor \log_2 M_* \rfloor$ and run exactly $2k$ halvings. For $0 \le j < k$ let $J_j = I_{2j} \setminus I_{2j+2}$, and let $J_k = I_{2k}$ denote the one tail. These disjoint sets exhaust all original channels. Their *unwhitened* Grams satisfy

$$G_{J_j} \ge \frac{43}{64} 4^{-j} I_q \ (j < k), \qquad G_{J_k} \ge \frac{15}{16} 4^{-k} I_q. \tag{23}$$

The first constant is $(15/16) - (17/64)$, using a lower parent and upper later subset in their Loewner difference. These are matrix bounds, not traces.

Use the actual spatial pieces of Lemma 2.3. Since $2^k \le n/32$, their parent lengths satisfy

$$L_j \le \frac{33}{32} n 2^{-j} \quad (0 \le j \le k).$$

Define $\lambda_j = 5L_j^2/(8n^2)$, including the tail. Then

$$\lambda_j \le \frac{5445}{8192} 4^{-j} < \frac{43}{64} 4^{-j}, \qquad \frac{43}{64} = \frac{5504}{8192}. \tag{24}$$

Thus (23) implies $G_{J_j} \ge \lambda_j I_q$ for every piece, also the tail. Give channels in J_j the child partition $\mathcal{P}_{j+1}$ when $j < k$, and channels in J_k the singleton partition. Every original channel receives exactly one partition. At $k = 0$ all channels belong to the singleton tail; $G_{J_0} = I_q \ge 5I_q/8$ directly.

For $j < k$, $|J_j| \le (17/16) g 4^{-j}$. Hence, with the constant A_D in Theorem 1.1,

$$N_{\text{coarse}} \le D 2^D \frac{17}{16} g \sum_{j=0}^{k-1} 2^{(D-2)j} \le A_D g 2^{(D-2)k} \le A_D g V/n^2.$$

The sum is zero at $k = 0$. Since $2^k > M_*/2$,

$$|J_k| \le \frac{17}{16} g 4^{-k} < \frac{17}{4} \max\{1024 g/n^2, C_0\Lambda q\} \le 4352 g/n^2 + \frac{17}{4} C_0\Lambda q.$$

Multiply by DV for every singleton row to obtain $N_{\text{out}} \le L_{D,\Lambda}(qV + gV/n^2)$. In the priority branch the full count DgV is at most $1024DgV/n^2$ if $n < 32$, or $DC_0\Lambda qV$ if $g < C_0\Lambda q$; the same stated constant pays both. Equalities at either threshold use the $k \ge 0$ branch. No large-size inference is applied at an ineligible input.

3.4 All-field coercivity with one strong-energy budget

First take z with componentwise mean zero. Put $f = \mathsf{U}^T z$ and $s = z - \mathsf{U}f$. These are q weak forms and g strong remainder forms, respectively; $\mathsf{U}^T s = 0$. Because U is spatially constant,

$$\|z\|^2 = \|f\|^2 + \|s\|^2, \qquad Q_B(z) = a\mathcal{E}(f) + T, \qquad T = H\sum_{\alpha} \mathcal{E}(s_\alpha). \tag{25}$$

For each actual group the parent-cell estimate gives

$$H\sum_{\alpha\in J_j} \mathcal{E}(s_\alpha) \geq \frac{4H}{L_j^2}\sum_{\alpha\in J_j} \|\Delta_j s_\alpha\|^2.$$

The channel groups are disjoint: summing these inequalities spends T once. Since $\Delta_j s_\alpha = \Delta_j z_\alpha - u_\alpha \cdot \Delta_j f$, use

$$\|p - v\|^2 \geq \frac{9}{10}\|v\|^2 - 9\|p\|^2.$$

The difference is $10\|p\|^2 - 2\langle p, v\rangle + \|v\|^2/10 \geq 0$. Using the actual Grams and (24) yields

$$\begin{aligned} T &\geq \frac{9}{10}\sum_{j=0}^{k} \frac{4H}{L_j^2}\langle \Delta_j f, G_{J_j}\Delta_j f\rangle - 36\sum_{j=0}^{k}\frac{H}{L_j^2}\sum_{\alpha\in J_j}\|\Delta_j z_\alpha\|^2 \\ &\geq \frac{9}{4}c\|f\|^2 - 36cb\|\mathcal{R}z\|^2. \end{aligned} \tag{26}$$

The positive pieces sum to $\|f\|^2$ by mean zero and their orthogonality. For the negative term, Δ_j is a subprojection of its channel's child observation, or of I at the tail. Also $L_j \geq n2^{-j}$ gives $H/L_j^2 \leq c4^j \leq c4^k$, and stopping implies $4^k \leq b$. Disjoint channel groups therefore bound the negative sum by the displayed single observation norm. No factor q or extra copy of strong energy is used.

Independently Lemma 2.1 gives $T \geq 16c\|s\|^2$. Take 8/9 of (26) and 1/9 of the global-gap inequality $T \geq 16c\|s\|^2$. This convex combination gives

$$T + 32cb\|\mathcal{R}z\|^2 \geq 2c\|f\|^2 + \frac{16}{9}c\|s\|^2 \geq c\|z\|^2.$$

Add $a\mathcal{E}(f) \geq 0$. In the priority branch $\mathcal{R} = I$ and $K = c$ already suffice, so $32cb$ do also.

Finally every channel partition observes every global constant: $\mathcal{R}\Pi_0 = \Pi_0\mathcal{R} = \Pi_0$. Its constant and mean-zero images are orthogonal, Q_B ignores constants, and $32cb \geq c$ pays their full norm. This proves the all-field statement and, by (4), the precise normalization in (8). Strong constants on zero weak-frame rows are included, not assigned a nonexistent spectral gap.

3.5 Actual-rank lower bound and the cancellation of weak dimension

Let $C = \bigoplus_\alpha C_\alpha$ be any internally channel-separable linear map on $\mathcal{U}$. Let S_α be its row-space orthogonal projection *inside* $\mathcal{U}_0$, and $S = \bigoplus_\alpha S_\alpha$. Then $r = \operatorname{rank}(C|_{\mathcal{U}}) = \sum_\alpha \operatorname{rank} S_\alpha$ and $\ker C = \ker S$. The weak embedding $T_0 : \mathcal{U}_0^q \to \mathcal{U}$, $T_0 f = \mathsf{U}f$, is an isometry of real domain dimension qd_0. Constant internal coefficients preserve the actual coexact space. Lemma 2.1 gives

$$Q_B(T_0 f) = a\mathcal{E}(f) \leq \frac{c}{4}\|f\|^2. \tag{27}$$

Taking a real orthonormal weak basis gives

$$\operatorname{tr}(T_0^* S T_0) = \sum_\alpha P_{\alpha\alpha}\operatorname{rank} S_\alpha \leq \frac{\Lambda q}{g} r. \tag{28}$$

Averaging over dimension qd_0 therefore supplies a *unit* f with $t = \|ST_0 f\|^2 \le \Lambda r/(gd_0)$: the two q factors cancel. Set $w = T_0 f$ and $u = (I - S)w$. Then $u \in \mathcal{U}$, $Cu = 0$ and $\|u\|^2 = 1 - t$. Positive-semidefinite polarization gives

$$Q_B(u) \le 2Q_B(w) + 2Q_B(Sw) \le c/2 + 8DHt.$$

The last step uses $Q_B(v) \le 4DH\|v\|^2$ on actual fields. At zero observation the restoring inequality, even with any finite Φ satisfying $\Phi(0) = 0$, implies $c(1 - t) \le c/2 + 8Dcn^2 t$. Thus $t \ge [2(1 + 8Dn^2)]^{-1}$, proving the second bound in (9) by (28).

Independently CT_0 must be injective, since a nonzero weak field in its kernel would violate restoration by (27). Hence $r \ge qd_0$. Since $d_0 \ge (D - 1)V/2$ and $1 + 8Dn^2 \le (1 + 8D)n^2$, their maximum is at least

$$\frac{D - 1}{8\Lambda(1 + 8D)}(qV + gV/n^2).$$

Together with the actual canonical construction this proves the rank order. Both a minimum on $\mathcal{U}$ and a minimum requiring all-field restoration obey these orders, because the lower is already on $\mathcal{U}$ and the upper is all-field. There is no output-dimension substitution in the lower argument.

3.6 Least coefficient for the constructed fixed weights

For a canonical map, $\mathcal{R}_\alpha|_{\mathcal{U}_0}$ is a contraction into the ambient one-channel field space; its squared singular values sum to at most its actual rank. That rank equals $\operatorname{rank}(C_\alpha|_{\mathcal{U}_0})$ by their equal kernels in (4), even if the projection image leaves $\mathcal{U}_0$. Consequently

$$\operatorname{tr}(T_0^* \mathcal{R}^* \mathcal{R} T_0) = \sum_\alpha P_{\alpha\alpha} \|\mathcal{R}_\alpha|_{\mathcal{U}_0}\|_{\mathrm{HS}}^2 \le \frac{\Lambda q}{g} r.$$

Restoration on weak fields and (27) require $K\|\mathcal{R}T_0 f\|^2 \ge (3c/4)\|f\|^2$. Sum a real orthonormal weak basis and cancel q to obtain

$$K \ge \frac{3cgd_0}{4\Lambda r}. \tag{29}$$

For the constructed map $r \le L_{D,\Lambda}(qV + gV/n^2)$, so

$$K \ge \frac{3(D - 1)c}{8\Lambda L_{D,\Lambda}} \frac{g}{q + g/n^2} \ge \frac{3(D - 1)c}{16\Lambda L_{D,\Lambda}} \min\{g/q, n^2\}.$$

The last comparison follows by considering which denominator term is larger. The sufficient all-field $32cb$ proves the opposite bound. A least coefficient exists: the admissible set is a nonempty closed intersection of restoring half-lines in K, bounded below by (29). This works on $\mathcal{U}$ or all fields for the same map. These statements complete Theorem 1.1.

Remark 3.4 (Endpoints and the role of coherence). For $q = g$, $P = I$ and $M_* < 1$; there is no strong complement. The singleton map has actual rank gd_0 and sufficient $K = c$, with $b = 1$. The singleton priority also makes $n = 2$, $n < 32$ and $g < C_0\Lambda q$ explicit, rather than an extrapolation of whitening. At $n = 32$ or $g = C_0\Lambda q$, $k = 0$ is permitted and the one singleton tail is valid. Without coherence a coordinate weak q-plane can use singleton weak channels and one root on each strong channel, with actual rank qd_0 and $K = c$; Lemma 2.1 controls strong mean-zero forms and the roots pay strong constants. But its diagonal entries equal one, so it is inadmissible for fixed Λ when $g > \Lambda q$. In the remaining range the extra gV/n^2 term is controlled by a Λ-dependent multiple of qV. This supplies no contrary coherent example.

4 Actual checkerboard graphs and plane observations

We first prove Theorem 1.2(i) and Theorem 1.3. Throughout this section n is even. The parity convention at torus wraps is part of the family, not an asymptotic simplification.

4.1 Selected strong cycles and an exact Dirichlet margin

For the orthogonal checkerboard, channel one has strong outgoing coefficient H at odd sources and channel two at even sources. Fix one original scalar spatial/channel component and call its strong-source parity A. Select only its strong $+e_1$ and $+e_D$ edges. Each A vertex has its two selected outgoing edges, while each opposite-parity vertex has its two selected incoming edges. The directions are distinct, also at period two. For an A-parity base u_0, define

$$u_t = u_0 + t(e_D - e_1), \qquad v_t = u_t + e_D, \qquad t \in \mathbb{Z}/n\mathbb{Z}.$$

The translation has order exactly n, and parity separates the u and v sets. The selected edges are u_t–v_t and v_t–u_{t+1}, the latter being the original $+e_1$ edge from u_{t+1}. Thus the selected components are disjoint actual $2n$-cycles and cover all vertices. This is a labelled finite graph, not a continuum comparison.

Let $\Sigma = \{x : x_D = 0\}$. Each cycle has exactly one $u_t \in \Sigma$ and one $v_{t-1} \in \Sigma$; they are adjacent. Removing these two zero vertices leaves $2n-2$ free vertices between zero endpoints. For a path with N free vertices and zero endpoints the sine vectors $\sin(\pi jt/(N+1))$, $1 \le j \le N$, diagonalize its Dirichlet matrix with eigenvalues $4\sin^2(\pi j/[2(N+1)])$. Indeed the interior second-difference identity gives the eigenvalue and the endpoint sines vanish; the N eigenvalues are distinct, so the corresponding nonzero orthogonal vectors span the path space. Its first eigenvalue is at least $4/(N+1)^2$ by the same sine inequality as before. A path with no free vertices has zero norm and requires no positive eigenvalue assertion. Consequently every field zero on Σ satisfies

$$Q_B(z) \ge \frac{4H}{(2n-1)^2}\|z\|^2 = (c+\delta_n)\|z\|^2, \qquad \delta_n = \frac{c(4n-1)}{(2n-1)^2} > 0. \tag{30}$$

Sum the selected cycle estimates over both original channels and all D spatial components. All other strong edges and all weak terms are merely discarded as nonnegative. The edge between adjacent zero anchors contributes zero. At $n = 2$ these are actual four-cycles; no parallel-label cancellation or missing endpoint is used.

In each channel observe every site of Σ as a singleton, and the one remaining rectangle $\{0,\dots,n-1\}^{D-1} \times \{1,\dots,n-1\}$. It is nonempty because $n \ge 2$. This is an exact original-chart partition, with $n^{D-1}+1$ cells per channel, shared over spatial components. Its block kernel is zero on Σ. Lemma 2.4, with $M = 4DH + c$ and $\delta = \delta_n$, supplies the finite coefficient $M + 2M^2/\delta_n$. Together with (4) this proves Theorem 1.2(i), including every global constant. This particular one-plane coefficient is not claimed to have the smaller point/plane-mesh penalty order in (12).

4.2 Every fixed nonzero angle: proof coordinates only

Let the two weak lines be spanned by unit vectors v_0, v_1, with principal angle $\theta \in (0, \pi/2]$. At parity ℓ let $P_\ell = v_\ell v_\ell^T$. Choose constant unit strong normals $s_\ell \perp v_\ell$ and, for each spatial component, set $w_\ell = s_\ell \cdot z$. The two-dimensional Gram matrix $s_0 s_0^T + s_1 s_1^T$ has eigenvalues $1 \pm |s_0 \cdot s_1|$, as follows from its trace two and determinant $1 - |s_0 \cdot s_1|^2$. Thus

$$\|w_0\|^2 + \|w_1\|^2 \ge \kappa\|z\|^2, \qquad \kappa = 1 - |\cos\theta| > 0, \tag{31}$$

also after summing over sites and spatial components. These are auxiliary analytic coordinates, not permitted observed rows. The actual energy contains H times differences of w_0 from even sources and H times differences of w_1 from odd sources. All weak terms stay positive. The selected two-direction cycle proof applies separately to each scalar functional with its own strong-source parity; these energy contributions are distinct parts of the actual source-weighted form.

Put $m = \lceil 4/\sqrt{\kappa} \rceil$. If $n < m$, all-channel singletons have $K = c$ and count $2Dn^D \le 2Dmn^{D-1}$, proving the stated bound directly. If $n \ge m$, choose the m distinct planes with coordinates $r_j = \lfloor jn/m \rfloor$, $0 \le j < m$. Their positive cyclic gaps ℓ_j are at most $L = \lceil n/m \rceil \le 2n/m$. Each plane meets each selected cycle in an adjacent pair. Between two successive anchor pairs there are precisely $2\ell_j - 2$ free vertices with zero endpoints. If $\ell_j = 1$ there are none; otherwise the finite Dirichlet estimate gives eigenvalue at least $4/(2\ell_j - 1)^2$. The wraparound segment is included by the cyclic gap, and zero–zero edges add zero. For an original field zero on the plane union, both w_0, w_1 are zero there, hence

$$\begin{aligned} Q_B(z) &\ge \frac{4H}{(2L-1)^2}(\|w_0\|^2 + \|w_1\|^2) \\ &\ge \frac{H\kappa}{L^2}\|z\|^2 \ge \frac{c\kappa m^2}{4}\|z\|^2 \ge 4c\|z\|^2. \end{aligned} \tag{32}$$

All inequalities here are non-strict, also at the zero field or when every vertex is anchored. No nonexistent singleton Dirichlet eigenvalue is used.

In *each original channel*, observe every point of every selected plane as a singleton. For each nonempty interval between planes use the one full cross-section rectangle. With $r_0 = 0$ and $r_m = n$, these intervals are $\{r_j + 1, \dots, r_{j+1} - 1\}$ in the last coordinate; there is no omitted pre-zero interval. Empty slabs are dropped. This is an exact nonwrapping original-chart partition with at most $mn^{D-1} + m$ cells per channel. It is shared across all spatial components and uses only the rows (3). A kernel field is pointwise zero on the plane union. Apply Lemma 2.4 with $\delta = 3c$ to (32); (12) suffices. This proves Theorem 1.3 without mixing observed channels.

At $\theta = \pi/2$, this is a nonoptimal four-plane variant of the preceding one-plane construction. At $\theta = 0$, $\kappa = 0$ and the analytic normal map is singular; the proof does not apply and does not refute the constant-P theorem. Constants depending on a fixed positive θ may diverge as the angle tends to zero. No point-anchor order at a nonorthogonal angle is claimed by this plane argument.

5 Distributed point anchors in dimensions three and four

We prove the two smaller full counts in Theorem 1.2. The finite point estimate uses all constant modes and counting factors. Its dimension dependence is classical in character; the exact inequalities needed here are derived rather than imported from a continuum estimate.

Lemma 5.1 (Finite balanced-box point estimates)**.** *Let $p = 2$ or 3, let $L \ge 1$ be an integer, and let Q be a product of p integer paths with lengths $\ell_i \in [L/2, L]$. Use counting norm and internal edge energy $\mathcal{E}_Q$, each path edge once. At any $b \in Q$, $f(b) = 0$ implies*

$$\|f\|_Q^2 \le \begin{cases} 13L^2(1 + \ln L)\mathcal{E}_Q(f), & p = 2, \\ 113L^3\mathcal{E}_Q(f), & p = 3. \end{cases} \tag{33}$$

Here $\|f\|_Q^2$ denotes $\sum_{y \in Q} |f(y)|^2$.

Proof. Put $N = \prod_i \ell_i$. If $N = 1$ the anchored function is zero, so assume $N > 1$. In particular $L \geq 2$. The tensor basis from Lemma 2.2 has point values and eigenvalues

$$|\phi_\nu(b)|^2 \leq 2^p/N, \qquad \lambda_\nu = 4\sum_i \sin^2 \frac{\pi\nu_i}{2\ell_i} \geq \frac{4|\nu|^2}{L^2} \quad (\nu \neq 0), \qquad 0 \leq \nu_i < \ell_i.$$

A singleton factor forces its index to zero. All indices are nonnegative cosine indices; no signed-frequency multiplicity is suppressed. On the shell $r = \max_i \nu_i \geq 1$, for $p = 2$ there are at most $2r + 1 \leq 3r$ indices, and for $p = 3$ at most $3r^2 + 3r + 1 \leq 7r^2$. Since $|\nu|^2 \geq r^2$ and $r \leq L - 1$,

$$\sum_{\nu \neq 0} |\nu|^{-2} \leq \begin{cases} 3(1 + \ln L), & p = 2, \\ 7L, & p = 3. \end{cases}$$

The first bound follows from $\sum_{r=1}^{L-1} r^{-1} \leq 1 + \ln L$ by integration; the second is the finite sum of at most $L - 1$ terms each bounded by seven. Consequently the nonconstant diagonal inverse satisfies

$$G_Q(b, b) := \sum_{\nu \neq 0} \frac{|\phi_\nu(b)|^2}{\lambda_\nu} \leq \begin{cases} 3L^2(1 + \ln L)/N \leq 12(1 + \ln L), & p = 2, \\ 14L^3/N \leq 112, & p = 3, \end{cases} \tag{34}$$

using $N \geq L^p/2^p$. Expand $f = \sum a_\nu \phi_\nu$ in this counting-orthonormal basis. Cauchy–Schwarz gives $|f(b) - \bar{f}|^2 \leq G_Q(b, b)\mathcal{E}_Q(f)$, where $\bar{f}$ is its arithmetic mean. The constant-mode mass is $N|\bar{f}|^2$, not $|\bar{f}|^2$. If $f(b) = 0$, this mass is at most $12L^2(1 + \ln L)\mathcal{E}_Q$ when $p = 2$, and $112L^3\mathcal{E}_Q$ when $p = 3$, because $N \leq L^p$. The positive spectral gap separately bounds the variance by $L^2\mathcal{E}_Q/4$. Adding the orthogonal constant and variance parts proves (33): $12 + 1/4 < 13$, and $112L^3 + L^2/4 \leq 113L^3$. The $N = 1$ case includes $L = 1$ and any permissible all-singleton product. Thus no nonexistent positive gap or lost normalization is used. Point values were bounded uniformly, so a corner anchor is permitted without an additional boundary trace assumption. □

5.1 The full strong graph in original coordinates

Let $p = D - 1$, with $D = 3$ or 4. Fix an original scalar spatial/channel component and its strong-source parity A. For every residue s of that parity set, with $y \in (\mathbb{Z}/n\mathbb{Z})^p$,

$$x_s(y) = (y_1, \ldots, y_p, s - \textstyle\sum_i y_i) \pmod{n}, \quad u_s(y) = z(x_s(y)), \quad v_s(y) = z(x_s(y) + e_D). \tag{35}$$

The first map is a bijection onto sum layer s; the shifted map is a bijection onto layer $s + 1$. Even n makes residue parity agree with original site parity. The $n/2$ pairs $\{s, s + 1\}$ cover all sum residues disjointly, including $\{n - 1, 0\}$ for odd source parity. Their pair norms therefore sum to the original scalar counting norm exactly.

The actual strong $+e_D$ edge is $u_s(y)$–$v_s(y)$. For $1 \leq i \leq p$, $x_s(y + e_i) + e_D = x_s(y) + e_i$, so the other strong edges are $u_s(y)$–$v_s(y + e_i)$. They are *all* D outgoing strong labels, not a surrogate subset with reassigned endpoint weights. Their exact energy is

$$E_s = \sum_y |u_s(y) - v_s(y)|^2 + \sum_{y,i=1}^{p} |u_s(y) - v_s(y + e_i)|^2. \tag{36}$$

A cross edge followed by a rung connects $u_s(y)$ to $u_s(y + e_i)$; all source translations are generated and every target has a rung, proving connectivity. At $n = 2$ the heads $y, y + e_i$ are distinct. The induced source-torus energy may repeat an undirected neighbor pair, but every original forward label is retained. Translations remain bijections.

Insert $v_s(y+e_i)$ between $u_s(y+e_i)$ and $u_s(y)$ and square with factor two. For the labelled source energy this proves

$$\mathcal{E}_p(u_s) := \sum_{y,i=1}^{p} |u_s(y+e_i)-u_s(y)|^2 \le 2p\sum_y |u_s-v_s|^2+2\sum_{y,i}|u_s(y)-v_s(y+e_i)|^2 \le 2pE_s. \quad (37)$$

The repeated-rung factor is exactly $2p$, also at period two. Summing HE_s over all layer pairs, both channels and all spatial components gives a lower part of the original Q_B; weak-source terms are only discarded as positive.

5.2 Balanced quantiles and both layer norms

For an integer $1 \le m \le n$, partition each source coordinate into

$$I_j = \{\lfloor jn/m \rfloor, \dots, \lfloor (j+1)n/m \rfloor - 1\}, \qquad 0 \le j < m.$$

These are nonempty paths with lengths $\lfloor n/m \rfloor$ or $L = \lceil n/m \rceil$. If $L = 1$ all are singletons; otherwise a shorter side is $L - 1 \ge L/2$. Every product cell therefore meets the exact hypotheses of Lemma 5.1. Anchor its lowest-coordinate corner as a source value $u_s(b) = 0$, at the actual original point $x_s(b)$. Cell internal edges are a subset of (37), counted once; across-cell and wrap edges are discarded as nonnegative. Define

$$X_2 = L^2(1+\ln L), \qquad X_3 = L^3.$$

Lemma 5.1 and (37) give $\|u_s\|^2 \le 52X_2E_s$ for $p = 2$ and $\|u_s\|^2 \le 678X_3E_s$ for $p = 3$. Rungs also give $\|v_s\|^2 \le 2\|u_s\|^2 + 2\sum |v_s - u_s|^2$. As $X_p \ge 1$, both original layers satisfy

$$\|u_s\|^2 + \|v_s\|^2 \le \begin{cases} 158X_2E_s, & p = 2, \\ 2036X_3E_s, & p = 3. \end{cases} \quad (38)$$

Here $3 \cdot 52 + 2 = 158$ and $3 \cdot 678 + 2 = 2036$; no target-layer mass is omitted. These formulas remain valid at singleton cells because the anchored source value is zero and the rung still controls its target.

There are $(n/2)m^p$ distinct source anchors per channel: the cell corners are distinct, (35) is a bijection, and different source layers are disjoint. Both channels together have $A = nm^p$ channel-labelled anchors. Use the *same* anchor sites for all D spatial components in that channel. Although source-coordinate cells need not be original rectangles, only their original *point* sites are observed. Applying Lemma 2.5 independently to the two channels gives actual canonical partitions with

$$N_{\text{out}} \le 2D + 2D^2A = 2D + 2D^2nm^p. \quad (39)$$

The initial rows and every channel/component factor are included. Original chart representatives, not wrapping rectangles or diagonal averages, realize the points. A block-kernel field is zero at every required anchor; extra nonpoint means only narrow this kernel.

5.3 Dimension three: the critical logarithmic mesh

Take $D = 3$, $T_n = 1 + \ln n$, and let m be the least positive integer with $m^2 \ge 4096T_n$. If $m \le n$, then $L \le 2n/m$ and $L \le n$, giving

$$X_2 = L^2(1+\ln L) \le \frac{4n^2}{m^2}T_n \le n^2/1024.$$

Sum (38) over the disjoint pairs, channels and components. Since $H = cn^2$, a field zero at the prescribed anchors satisfies

$$Q_B(z) \geq \frac{H}{158X_2}\|z\|^2 \geq \frac{1024}{158}c\|z\|^2 \geq 4c\|z\|^2, \tag{40}$$

because $4 \cdot 158 = 632 < 1024$. Put $\rho = \sqrt{4096T_n} \geq 1$. Minimality yields $m = \lceil \rho \rceil \leq 2\rho$, so $m^2 \leq 16384T_n$. The complete count (39) is therefore

$$N_{\text{out}} \leq 6 + 18nm^2 \leq 6 + 294912nT_n < 2^{19}nT_n.$$

We used $nT_n \geq 1$ and $294912 + 6 < 524288$, without divisibility or exact-square assumptions.

If $m > n$, minimality gives $(m-1)^2 < 4096T_n$ and the integers satisfy $n \leq m - 1$. The all-channel singleton fallback has

$$N_{\text{out}} = 6n^3 < 24576nT_n < 2^{19}nT_n, \qquad K = c.$$

This covers every ineligible even n, including $n = 2$ and $n = m - 1$. At $m = n$ the mesh branch is allowed and its source cells are singleton; (40) still follows from the derived rung bound. Thus no guessed large-size cutoff, empty mesh cell or unproved singleton gap is used.

5.4 Dimension four: a three-dimensional source mesh

For $D = 4$, first take $n \geq 256$ and let m be the least positive integer with $m^3 \geq 65536n$. The integer n itself is eligible because $n^2 \geq 65536$, so $m \leq n$. Also $L \leq 2n/m$ and

$$X_3 = L^3 \leq 8n^3/m^3 \leq n^2/8192.$$

The other case of (38) now gives, on the actual anchor kernel,

$$Q_B(z) \geq \frac{8192}{2036}c\|z\|^2 \geq 4c\|z\|^2, \qquad 4 \cdot 2036 = 8144 < 8192. \tag{41}$$

For $\rho = (65536n)^{1/3} \geq 256$, $m = \lceil \rho \rceil \leq 2\rho$, hence $m^3 \leq 524288n$. Thus $A = nm^3 \leq 524288n^2$ and

$$N_{\text{out}} \leq 8 + 32A \leq 8 + 2^{24}n^2 \leq 2^{25}n^2.$$

At $n = 256$ one has $m = n$ exactly; singleton source cells and their rungs are valid in (41). Noncubes are covered by the same ceiling argument. For even $n < 256$, use all-channel singletons: $N_{\text{out}} = 8n^4 < 524288n^2 < 2^{25}n^2$ and $K = c$. No three-dimensional source-cell estimate is applied in that fallback.

5.5 Finite completion, constants and actual restricted rank

For either mesh the canonical block kernel obeys $Q_B(u) \geq 4c\|u\|^2$ by (40) or (41). Lemma 2.4 with $\delta = 3c$ and $M = 4DH + c$ proves the all-field inequality with exactly (12). Dividing by V and using (4) gives the stated canonical normalization. On a constant field $Q_B = 0$ and $\mathcal{R}z = z$; the finite coefficient is at least c, so these constants are paid. The conservative mesh coefficient can grow as cn^4; in the fallbacks $\mathcal{R} = I, K = c$ directly. Restrict only after this all-field proof to the original $\mathcal{U}_0^2$. Its actual rank is at most the full count, with no assumed row independence or coexact invariance of block averaging. This completes Theorem 1.2(ii)–(iii).

The same positive stack still has D^2 weak and D^2 strong eigenvalues per site. Formula (14) with $q = 1, g = 2$, including its original $\alpha_3 = 729/128$ or $\alpha_4 = 81/16$, is unchanged and has order V. These improved observations do not reduce the inherited capped budget. Neither the logarithm in this particular finite construction nor its point count is asserted necessary for every permitted partition. The present point proof is for *orthogonal* lines only; it is distinct from the fixed-angle plane result and supplies no claimed solution to its unproved point-anchor analogue.

6 Scope, classical inputs and review context

6.1 What the comparison establishes

For fixed D, Λ, Theorem 1.1 is uniform in weak dimension and channel count under a *constant* weak projector. It neither supplies a basis-free lower nor compares all nonlinear observation classes. Its finite-support MSS input is nontrivial and classical; the extra signed count coordinate and the conditioning/count bookkeeping are displayed as an application, with no historical-priority assertion. The checkerboard results change only spatial constancy, not local positivity, local spectra, upper coherence, original channel basis or canonical weights. They show why that hypothesis cannot be dropped from a per-profile lower of the same form. Strong-graph geometry, rather than local spectra alone, enters the constructions. Their improved row upper bounds are not optimal-rank or uniform-penalty theorems.

The least coefficient assertion for the fixed-space result is deliberately for the *particular economical map* constructed here. It is not a global minimum over all canonical maps: all-channel singletons always permit $K = c$ for a positive form, at the cost of many rows. Conversely no practical conditioning or implementation conclusion follows from finite existence, lexicographic selection or a rank estimate. The large constants and the checkerboard coefficients make those limitations substantial.

6.2 Primary comparisons and their actual extent

Theorem 3.1 is the precise published input from Marcus–Spielman–Srivastava [1]. The retained primary reading concerns Theorem 1.4 and Corollary 1.5 on printed page 3 of arXiv:1306.3969v4, including the opening corollary argument. It is not a new audit of the full published MSS proof. The simultaneous count statement used here is proved with the added coordinate; it is not attributed to an unstated cardinality clause in that source.

Local coarse averages and their kernels are established mechanisms, as the selected reading of Målqvist–Peterseim [4, Sections 2.2–2.3] already shows. Finite graph Poincare estimates are classical; the recorded reading of Levi–Santagati–Tabacco–Vallarino [5, Section 2, Theorem 2.1] does not replace the explicit tensor-path proof given here. Lawler–Limic [2] provides classical Green-function and potential-kernel context. The later bounded reading used an author-hosted 289-page draft of unidentified date, whose pagination differs from the publisher edition: Section 4.3 and Theorem 4.3.1 (pp. 81–82), Section 4.4.1 and Theorem 4.4.4 (p. 88), and the opening of Section 6.7 (pp. 152–153), with selected surrounding excerpts, not complete proofs. The first concerns transient infinite-lattice Green functions in dimensions at least three; the potential-kernel section supplies two-dimensional logarithmic context. Boundary rules, generator normalization and removal of the constant mode must be translated before any constants transfer. Lemma 5.1 derives its finite Neumann constants directly and proves no logarithmic necessity for all allowed observations.

The GenEO work of Spillane et al. [3] is a substantive adaptive coarse-space predecessor. Operator-dependent local eigenmodes and partition-of-unity extensions give robust two-level Schwarz convergence. The retained reading used an author-linked institutional copy of the journal PDF, including selected setup, Lemma 2.11, Definitions 3.7/3.9, Assumptions 2.1/3.12/3.13, Lemma 3.21 excerpts, and Theorem 3.22 with Corollary 3.23 and its size qualification; it was not a full-proof, implementation or numerical-section reading. These adaptive vector functions are not prescribed original-channel rectangle means, and their selected dimension can depend on coefficient geometry and be large. A complete reduction to our domain, exact row class, uniform rank count and fixed-weight penalty has not been established; neither has an impossibility of such a reduction. In particular, failure to establish that implication in a limited pass provides no evidence of originality. Conversely, our admitted

finite large penalties do not establish the contrast-robust practical conditioning studied by that solver literature.

Braides–Chiadò Piat–Piatnitski [6] already studies periodic high-contrast strong/weak lattice networks. The later bounded reading includes abstract/introduction excerpts, Section 2 connectivity assumptions and the statement and preceding decomposition of Theorem 5.1, not its homogenization proof or flows. A finite paired-layer checkerboard needs an actual encoding into that theorem's model assumptions before an implication can be claimed; common terminology is not that encoding. Strong-network geometry is classical. Implication of the exact finite canonical row bounds is unresolved here, not proved impossible. A closer vector/form predecessor needing full comparison is Budninskiy–Owhadi–Desbrun [7]. The institutional primary abstract, not the full article, was read. It describes operator-adapted differential forms and differential constraints, so vector forms and operator adaptation are not themselves available novelty claims. Its complete relation to the fixed canonical measurements, growing g, q, actual restricted ranks and penalty orders here is unestablished. Failed full-text retrieval is not evidence of absent overlap.

The new bounded comparison covered four works: MSS, Lawler–Limic, GenEO and Braides et al.; the other citations have the earlier limited readings just described. The comparison did not receive or audit this manuscript's completed source. These are limited primary readings, not an exhaustive literature or historical-equivalence audit. No external full-text PDF is redistributed in the source packet. Classical Neumann, Green, PSD trace, strong-graph and rectangle-isolation methods are acknowledged. The specific rank/penalty statements require specialist priority scrutiny rather than a claim of a new general principle.

Model assessments and review context

At the source-preparation freeze of 11 November 2026, no PDF of this manuscript had yet been compiled, rendered, read or numerically assessed. The individual source arguments have retained analytical AI source reports; those reports are not a whole-manuscript source audit or an exact-PDF review of this consolidation. A separate integrated-source gate, compilation, page-by-page rendering review and first exact-PDF assessment are later gates. There is no numerical grade for this source preparation.

Lluis Eriksson is the named author. Drafting and consolidation assistance was provided through Codex/OpenAI, exposed as an agent based on GPT-6; the exact served backend and actual reasoning effort were not exposed and are unknown. Earlier source audits used separately assigned AI roles in a shared account and persistent context. No independent, blind, human or fresh-provider review is established by that role separation. The finite proofs are not Lean or other formal certificates, and repository checks do not formalize them.

The exact earlier 232B PDF, SHA-256

`5595e89c1bcbec1d9b1ae35053a76e6a962f67b45d5d1ac483e5b7a60c9b9be0`

received a recorded importance assessment of 4.10/10. That judgment belongs only to those earlier PDF bytes and supplies no grade here. Its complete first report, immutable original opinion, same-round protocol clarification and author response are retained as historical context, without regrading. The three prior material issues remain unresolved context rather than a new assessment of this manuscript:

1. Historical priority and exact implication by predecessors remain unestablished after limited primary lookups.

2. Practical usefulness, conditioning with allowed penalties, applications and efficient implementation remain unestablished.

3. The absence of a fatal error in that earlier reading was not unconditional editorial acceptance. Finite, observation-class, basis/growth and penalty hypotheses must remain explicit; this consolidation does not create a new editorial decision or close that issue.

No score transfers from the earlier article or from source reports. The present scope is finite matrix-gradient observations and the stated actual flat coexact domain; there is no interacting producer, continuum result or physical mass-gap conclusion.

The source supplied to a later gated build is

`4ec599da41521ecfdadd7d2aee69faab8b3819f1`

The default `SOURCE-NOT-BUILT` marks source preparation only; the build must supply its actual full frozen commit. Prepared PDF metadata controls are not evidence of bit reproducibility. That claim requires actual matching PDFs from two fresh clones, which were not produced in this source-preparation stage.

A Source correspondence and retained corrections

The source packet retains the complete grouped 235–248 source/audit snapshot, including rejected variants and adverse reports. Full hashes, exact original commits and paths are in `SOURCE-MAPPING.json` and `SOURCE-INPUTS.json`; the table gives abbreviated raw SHA-256 identities for navigation. The manuscript's proofs replace several earlier specialized selection arguments by the common MSS proof. They do not assert that the new constructed map is identical to each earlier one.

Source	Earlier role	Place in this manuscript	Raw hash prefix
235	Dyadic local rectangles, uniform weak line, $g = V$	Finite Neumann and piece estimates; special case of Theorem 1.1	`4ccc3829df28`
236B	Arbitrary integer sizes; corrected counting normalization	Lemma 2.3, (4) and the explicit VK convention	`b29ea47b3949`
237	Fixed-canonical rank/penalty trace	(29); order recovered for the new constructed map	`2d7a29ed2129`
238	Uniform line, $g \geq n^2$	$q = 1$ large-channel case of Theorem 1.1	`d32783c5ff5d`
239	Uniform line, every $g \geq 1$	$q = 1$ rank order; common replacement construction	`d277c4773cca`
240B	Least coefficient for the specified economical map	(10), qualified for the new fixed map	`d4c37b82e9d9`
241B	Balanced rank-one polarization; original cap restored	$q = 1$ scope, (13); no lower row-weight hypothesis here	`d62cae0af036`
242	Upper-delocalized line without positive minimum row weight	$q = 1$ case; zero rows counted explicitly	`1016e0c7edc1`
243	Two weak directions, actual matrix groups	$q = 2$ case of Theorem 1.1; singleton priority retained	`985dbb4412fa`
244	Uniform growing weak dimension	Theorem 1.1, MSS count augmentation, complete upper and lower proofs	`59498789a14e`
245	Orthogonal spatial checkerboard	Theorem 1.2(i), exact cycle and finite completion	`6e2aad56076a`
246	Arbitrarily small fixed angle, plane mesh	Theorem 1.3; no nonorthogonal point order	`ab9ab2195aff`

Source	Earlier role	Place in this manuscript	Raw hash prefix
247	Orthogonal dimension four, point anchors	Theorem 1.2(iii), (41)	`e2ab6a1db681`
248	Orthogonal dimension three, logarithmic point cost	Theorem 1.2(ii), (40)	`93786fa89262`

The original 236A missing-V counting-penalty defect, the 240A small-input branch defect, and the 241A changed-cap normalization rejection remain bound to their original sources and reports. Corrected B sources do not turn those originals into passes. The nonblocking 242 empty-prefix base-case clarification and 243 singleton-priority clarification are retained too. In this manuscript the whitening induction has an explicit initial stage and the singleton priority applies to all channels. These preserved corrections are source history, not new numerical assessments. The active fixed-angle point-anchor question has not been developed into a source proof here.

References

[1] A. W. Marcus, D. A. Spielman and N. Srivastava, *Interlacing Families II: Mixed Characteristic Polynomials and the Kadison–Singer Problem*, arXiv:1306.3969v4 (2014), Theorem 1.4. `https://arxiv.org/abs/1306.3969v4`.

[2] G. F. Lawler and V. Limic, *Random Walk: A Modern Introduction*. Author-hosted 289-page draft, date unidentified in the inspected version; pagination differs from the publisher edition. `https://www.math.uchicago.edu/~lawler/srwbook.pdf`. Selected reading: Section 4.3/Theorem 4.3.1, Section 4.4.1/Theorem 4.4.4, and the opening of Section 6.7, not complete proofs.

[3] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein and R. Scheichl, *Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps*, Numerische Mathematik **126** (2014), 741–770. `https://doi.org/10.1007/s00211-013-0576-y`. Author-linked institutional journal copy: `https://strathprints.strath.ac.uk/44827/1/10.1007_s00211_013_0576_y.pdf`. Selected setup, definitions and statements, not the full paper.

[4] A. Målqvist and D. Peterseim, *Localization of Elliptic Multiscale Problems*, arXiv:1110.0692v4. `https://arxiv.org/abs/1110.0692v4`. Selected reading: Sections 2.2–2.3.

[5] Levi, Santagati, Tabacco and Vallarino, *Poincare inequalities on graphs*, arXiv:2203.12530v2. `https://arxiv.org/abs/2203.12530v2`. Selected reading: Section 2 and Theorem 2.1.

[6] Braides, Chiadò Piat and Piatnitski, *Discrete double-porosity models*, arXiv:1406.1752v1 (2014). `https://arxiv.org/abs/1406.1752v1`. Selected reading: introductory excerpts, Section 2 assumptions and Theorem 5.1 statement/decomposition; not the full proof.

[7] M. Budninskiy, H. Owhadi and M. Desbrun, *Operator-adapted wavelets for finite-element differential forms*, Journal of Computational Physics **388** (2019), 144–177. `https://doi.org/10.1016/j.jcp.2019.02.018`. Institutional primary abstract: `https://authors.library.caltech.edu/records/8j84b-1ep03`. Only that abstract, not the full article, was read in the retained comparison.