Rectangular observations for finite matrix gradients: critical budgets and the weak endpoint

Lluis Eriksson

10 November 2026 — manuscript 226C, corrected preparation version

Abstract

For a finite periodic lattice and arbitrary finite symmetric matrices acting on its actual forward gradient stack, we construct a partition into integer rectangles whose scaled component means restore coercivity. The output dimension is $O_D(h(1 + F_c(A)))$ for h real components and capped local eigenchannel budget $F_c(A)$. The product order $hF_c(A)$ is necessary in worst-case joint growth for common partitions retaining all component means. The quadratic penalty is finite and may depend on the coefficient profile, including its negative spectral depth. The proof combines positive spectral replacement, an arbitrary-positive-scalar balanced-ring construction, and singleton refinement at every affected start and forward endpoint. We give the full scalar BV, weighted-gap and compressed-tree proof in an appendix. For positive scalar coefficients with exact mean μ, a weak inverse-$D/2$ distribution bound gives a logarithmic upper bound on observation dimension. An explicit family of actual flat coexact fields forces the same worst-case logarithmic order for every centered differentiable observation with a finite locally quadratic upper penalty. The mechanisms are classical; historical priority and non-equivalence to the closest predecessors remain unestablished. All results concern specified finite component-gradient forms.

1 Objects, principal results and scope

A small number of spatial averages can restore coercivity even when local coefficients are very small or have negative channels. Two issues must be kept separate: how many averages are needed, and how large a penalty must be placed on them. We control the first by a capped coefficient budget while allowing the second to depend on the entire finite profile. At the scalar weak endpoint a logarithmic loss in the coefficient scale is unavoidable in the worst case.

1.1 The actual finite form and observation rows

Throughout, $D \geq 3$, $n \geq 1$ and $h \geq 1$ are integers, $c > 0$,

$$X = (\mathbb{Z}/n\mathbb{Z})^D, \qquad V = n^D, \qquad P = D/2, \qquad r = Dh.$$

We represent X by the fixed chart $\{0, \dots, n-1\}^D$. For $u : X \to \mathbb{R}^h$ the forward stack is

$$du(x) = \big(u(x + e_i) - u(x)\big)_{i=1}^{D} \in \mathbb{R}^r.$$

All additions in this expression are periodic. Given a finite real symmetric $r \times r$ matrix $A(x)$ at each original start x, define

$$Q_A(u) = \sum_{x \in X} du(x)^{\mathsf{T}} A(x) du(x), \qquad \|u\|_V^2 = V^{-1} \sum_x \|u(x)\|^2. \tag{1}$$

No common eigenbasis or sign assumption is imposed on $A(x)$. For a scalar profile $a > 0$ we write

$$Q_a(u) = \sum_{x,i} a(x)\|u(x + e_i) - u(x)\|^2.$$

In particular, the coefficient is always evaluated at the forward start; it is never moved to the endpoint or averaged over an edge.

An integer rectangle is a nonempty product of consecutive intervals in the fixed chart. It need not be a periodic rectangle crossing the chart boundary. For such a rectangle B define the vector row

$$q_B(u) = V^{-1/2}|B|^{-1/2} \sum_{x\in B} u(x) = \sqrt{|B|/V}\ \mathrm{mean}_B\, u. \tag{2}$$

For a partition $\mathcal{R}$ put $C_{\mathcal{R}}u = (q_B(u))_{B\in\mathcal{R}}$, with the Euclidean product norm. Its output dimension is $h\#\mathcal{R}$; on the full field space its rows are independent, so its rank equals that dimension. A coefficient-only partition is fixed by the coefficient profile and deterministic geometric conventions before the field u is chosen.

1.2 An arbitrary symmetric matrix theorem

The constants used below are intentionally conservative:

$$\begin{aligned} b = 2^D, \quad J = 2D, \quad C_{\mathrm{ext}} &= 1 + J\{2J(1 + 12J) + 1\}, \\ G_D = 2^{2D+10}D^2C_{\mathrm{ext}}^2, \qquad B_D &= (2b + 3)bG_D^P, \\ \alpha_D = \frac{D}{2}\left(\frac{D}{D-2}\right)^{D-1}&\left(\frac{3}{4}\right)^D. \end{aligned} \tag{3}$$

Here G_D is the scalar weighted-gap constant and B_D is the resulting tree-count constant. Let $\lambda_j(x)$, $1 \le j \le r$, be the eigenvalues of $A(x)$, counted with multiplicity. Let ν_0 count all its nonpositive local eigenvalues over all x. Define

$$F_c(A) = \nu_0 + \sum_{x,j:\lambda_j(x)>0} \min\{1, \alpha_D c^P \lambda_j(x)^{-P}\}. \tag{4}$$

These are local coefficient eigenvalues, not the eigenvalues of the global operator associated with Q_A.

Theorem 1.1 (Capped-budget rectangular observations)**.** *For every profile A in* (1) *there is a coefficient-only partition $\mathcal{R}$ into integer rectangles such that*

$$h\#\mathcal{R} \le 2Dh\left[1 + B_D + \max\left\{\frac{B_D}{\alpha_D}, D + 1\right\} F_c(A)\right]. \tag{5}$$

There is a finite coefficient-dependent K such that, simultaneously for all $u : X \to \mathbb{R}^h$,

$$V^{-1}Q_A(u) + K\|C_{\mathcal{R}}u\|^2 \ge c\|u\|_V^2. \tag{6}$$

One may take $K = c + 4DL$, with L explicitly defined in (22)*. Neither the partition nor K depends on u.*

The budget caps the cost of each sufficiently small positive channel at one and assigns the same unit cost to each nonpositive channel. Its count thus does not grow with negative depth. The penalty K can grow with that depth; (6) makes no assertion that $K = c$ for signed forms. The elementary singleton partition has hV outputs and also admits a finite penalty for every finite symmetric form. One may choose whichever of that partition and Theorem 1.1 has fewer outputs. No optimal dimension, sharp coefficient, or efficient implementation is asserted.

1.3 The cost of a common full-vector partition

The dimension in Theorem 1.1 concerns a common partition with all h component means on each spatial member. The following refinement shows why this measurement class can require a product of h and the budget.

Theorem 1.2 (Necessary product order for common partitions)**.** *Fix $D \geq 3$ and $c > 0$. For every pair of integers $m, h \geq 1$ there is an actual positive diagonal matrix profile on X_n, $n = 3m + 3$, with exactly $2Dm$ weak direction-labelled channels and*

$$2Dm \leq F_c(A) \leq 2Dm + 1.$$

If a common partition $\mathcal{R}$ outputs all h means (2) *and satisfies* (6) *for every general field $u : X_n \to \mathbb{R}^h$ with any finite quadratic penalty coefficient, then*

$$\operatorname{rank} C_{\mathcal{R}} = h\#\mathcal{R} \geq hm \geq \frac{hF_c(A)}{2D + 1}. \tag{7}$$

Thus no bound $C_D(h + F_c(A))$ holds uniformly in h in this class: setting $h = m$ makes the ratio of necessary rank to $h + F_c(A)$ diverge. Together with Theorem 1.1, this gives matching worst-case $hF_c(A)$ order, up to dimensional constants, in joint $h, F_c(A)$ growth. The lower family consists of general vector fields and is not claimed to be coexact or divergence-free. It does not address component-dependent partitions, channel-adapted means, or postprocessed measurement classes.

1.4 The scalar endpoint and the precise lower-bound domain

For a positive scalar profile let

$$\begin{aligned} \mu = V^{-1} \sum_x a(x), \qquad & Y(x) = \mu / a(x), \\ W(a, \mu) = \sup_{\tau > 0} \tau^P \mathbb{P}_X(Y > \tau), \quad & T(n, \mu) = \max\{1, n^D/(1 + \mu)^P\}, \end{aligned} \tag{8}$$

where $\mathbb{P}_X$ is normalized counting probability. The strict superlevel convention in this definition is used throughout.

Theorem 1.3 (Weak endpoint upper bound)**.** *Suppose $1 \leq a \leq 2\mu$ and $W(a, \mu) \leq M_w$, with μ the actual arithmetic mean. For every $h \geq 1$ there is a coefficient-only rectangular partition with*

$$h\#\mathcal{R} \leq 2Dh \left\{1 + 2^P c^P B_D[1 + PM_w \log \mu]\right\} T(n, \mu) \tag{9}$$

and

$$V^{-1} Q_a(u) + c\|C_{\mathcal{R}} u\|^2 \geq c\|u\|_V^2 \quad (u : X \to \mathbb{R}^h). \tag{10}$$

To state the lower bound without an unspecified Hodge domain, fix $g \geq 1$. A flat one-cochain is $U = (U_1, \dots, U_D)$ with $U_j : X \to \mathbb{R}^g$; we regard it as an $h = Dg$ component field in Q_a. For antisymmetric two-cochains $B_{ij} = -B_{ji} : X \to \mathbb{R}^g$, define

$$D_i^- f(x) = f(x) - f(x - e_i), \qquad (\delta B)_j = -\sum_i D_i^- B_{ij}, \qquad \mathcal{H}_{\mathrm{coex}}(X) = \{\delta B : B \text{ antisymmetric}\}. \tag{11}$$

Backward divergence $\sum_j D_j^- U_j$ vanishes on this space by commutation and antisymmetry. The energy on it is precisely the scalar component-gradient $Q_a(U)$, with all its original forward starts; no weighted curl identity is used to define or replace that energy.

Theorem 1.4 (Weak endpoint obstruction and worst-case order)**.** *Fix $D \geq 3$, $c > 0$ and $g \geq 1$. For every integer $m \geq 1$ there are n_m, a profile a_m with actual mean μ_m, and an m-dimensional real subspace $E_m \subset \mathcal{H}_{\mathrm{coex}}(X_{n_m})$ such that*

$$\begin{aligned} 1 \leq a_m \leq 2\mu_m, \qquad & W(a_m, \mu_m) \leq 1, \\ & Q_{a_m}(U) \leq (c/4) \sum_x \|U(x)\|^2 \quad (U \in E_m), \\ & T(n_m, \mu_m) \leq C(D, c), \qquad \mathbb{E}_X(\mu_m / a_m)^P \geq \frac{m}{2^P 16^D}. \end{aligned} \tag{12}$$

Suppose an observation C from a neighborhood of zero in $\mathcal{H}_{\mathrm{coex}}(X_{n_m})$ to a finite-dimensional real normed space F is differentiable at zero and $C(0) = 0$. Suppose also that $\Phi(y) \leq K\|y\|^2$ near zero for some finite K, and that locally

$$c \sum_x \|U(x)\|^2 \leq Q_{a_m}(U) + \Phi(C(U)). \tag{13}$$

Then $DC(0)|_{E_m}$ is injective and $\dim F \geq m$. For this fixed-parameter sequence m is comparable to $1 + \log \mu_m$ for large m. Thus the logarithmic loss in Theorem 1.3 has matching worst-case order. This is not a claim of optimal rank for every individual profile.

The centered, differentiable and local quadratic assumptions are part of the lower theorem. Its penalty need not have a uniform bound as m varies. A common division of energy and squared field norm by V gives the same rank conclusion. The full critical moment grows along the lower family, so this family does not contradict the strong-moment scalar theorem.

1.5 Relationship to earlier work and proof roadmap

This manuscript is a proposed linked successor to the AIRR record `ARR-2026-7EPM3V3DP796BBRR`, version 13. It combines the complete arbitrary-positive-scalar construction from that predecessor with the subsequent matrix and weak-endpoint arguments. It is a new source artifact; prior judgments do not transfer to it.

The closest constructive predecessors are dyadic-ring approximation and compressed-tree arguments in Brudnyi [2], Sections 1–2 and Appendices I–II, and the bivariate BV approximation setting of Cohen–DeVore–Petrushev–Xu [3]. Their approximation spaces and objectives differ from the specified finite coefficient-only observation problem here. This comparison is not a proof that the present statements are absent from, or inequivalent to, the classical theory. In particular, ring decomposition, boundary-touching holes, BV/Sobolev estimates and tree compression are classical principles. The appendix spells out their finite integer implementation, including heavy atoms and the original edge weights.

Frank [1], Theorems 2.1 and 3.1, supplies the relevant classical Cwikel/CLR and operator-valued spectral context. A weak multiplier hypothesis in such a theorem must be distinguished from replacing the strong coefficient moment in the present problem by a weak spatial distribution hypothesis. Our matrix proof uses elementary local spectral calculus and a scalar minimum eigenvalue bound; it does not invoke a global CLR estimate. The particular positive normalization α_D in (3) fixes the cap threshold and numerical constants; the argument works with any fixed positive cap normalization after corresponding changes. Spectral replacement, scalar domination, rectangle refinement, layer cake and disjoint low-mode rank arguments are classical. No new general spectral principle or historical priority is claimed.

Section 2 proves Theorem 1.1 using the fully proved scalar theorem in Appendix A. Section 3 proves the common-partition product lower bound. Section 4 gives the layer-cake upper bound and the complete compact-bump coexact lower construction. The appendix proves the scalar BV extension, actual weighted ring gap, heavy-chain count and rectangular mean inequality without an assumed geometric or local-gap certificate.

All assertions concern finite unit-spacing component-gradient forms. No interacting or covariant coefficient producer, variable-weight curl replacement, continuum theorem or Yang–Mills solution follows. The programme context PR29=20/41, TermSource 0, is unchanged.

2 Positive replacement and exact rectangular refinement

We prove Theorem 1.1. Every step uses the actual matrix profile on the forward stack in (1). In particular, positivity is used to discard wrap terms only after replacing the signed profile.

2.1 The spectral split and scalar inverse moment

Set

$$\sigma = \alpha_D^{1/P} c, \qquad S_x = 1_{(-\infty,\sigma]}(A(x))\mathbb{R}^r, \qquad N = \sum_x \dim S_x, \qquad I = \{x : S_x \neq \{0\}\}.$$

Here $1_{(-\infty,\sigma]}(A(x))$ is the spectral projector. A degenerate eigenspace at the cutoff is selected in full, so no choice of basis within it enters the construction. Every nonpositive channel is selected. A positive selected channel contributes exactly one to the capped budget, and an unselected channel contributes its uncapped inverse power. Consequently

$$F_c(A) = N + \alpha_D c^P \sum_{x,j:\lambda_j(x)>\sigma} \lambda_j(x)^{-P}, \qquad 0 \le N \le F_c(A), \quad |I| \le N. \tag{14}$$

Equality $\lambda = \sigma$ is included in N.

Put $T = \sigma + c(1+N) > \sigma$. Replace every selected eigenvalue by T, leaving all unselected eigenvalues and every local spectral subspace unchanged, and denote the resulting positive definite matrix by $A_T(x)$. For $N = 0$ this changes nothing. With $a_T(x) = \lambda_{\min}(A_T(x)) > 0$, pointwise matrix positivity gives

$$Q_{A_T}(u) \ge Q_{a_T}(u). \tag{15}$$

This controls the complete stacked vector, including all directional and internal mixing. It requires no commutation between distinct matrices. Moreover, the largest eigenvalue of $A_T(x)^{-P}$

is at most its trace, hence

$$\begin{aligned} c^P \sum_x a_T(x)^{-P} &\le c^P \sum_x \operatorname{tr}(A_T(x)^{-P}) \\ &= c^P \sum_{\lambda_j(x)>\sigma} \lambda_j(x)^{-P} + N(c/T)^P \\ &\le \frac{F_c(A) - N}{\alpha_D} + 1. \end{aligned} \tag{16}$$

The last inequality uses $T \ge c(1+N)$ and $N/(1+N)^P \le 1$ for $P \ge 1$. When $N = 0$ the extra one is harmless.

2.2 The scalar partition is derived, not assumed

Theorem A.1 in Appendix A applies to every finite positive a_T, without an arithmetic-mean, lower-envelope or upper-envelope condition. Its partition $\mathcal{R}_0$, combined with (15), satisfies

$$c\|u\|_V^2 \le V^{-1} Q_{A_T}(u) + c\|C_{\mathcal{R}_0} u\|^2, \tag{17}$$

by summing the same scalar inequality over all h real components. If $R_0 = \#\mathcal{R}_0$, its exact count and (16) yield

$$\begin{aligned} R_0 &\le 2D \left[1 + (2b+3) b (cG_D)^P \sum_x a_T(x)^{-P} \right] \\ &\le 2D \left[1 + B_D \left(1 + \frac{F_c(A) - N}{\alpha_D} \right) \right]. \end{aligned} \tag{18}$$

There is no auxiliary local-gap or extension hypothesis in this application: the scalar construction and its gap are proved in the appendix.

2.3 Isolating every affected endpoint

Let

$$Z = I \cup \bigcup_{i=1}^{D} (I + e_i), \qquad |Z| \le (D+1)|I| \le (D+1)N, \tag{19}$$

where the translations are periodic. We refine $\mathcal{R}_0$ so that each $z \in Z$ is a singleton member. Order Z lexicographically in the fixed chart. If z is not already a singleton, let R be its current rectangle. Partition $R \setminus \{z\}$ according to the first coordinate in which a point differs from z, and according to whether it is smaller or larger in that coordinate. Earlier coordinates are fixed to those of z, later coordinates range over their original intervals, and the differing coordinate ranges below or above z. Each nonempty piece is an integer rectangle. There are at most $2D$ such pieces. Adding $\{z\}$ replaces one rectangle by at most $2D+1$, increasing the count by at most $2D$. Earlier isolated points lie in disjoint singleton members and remain intact. Thus the final partition $\mathcal{R}$ satisfies

$$R := \#\mathcal{R} \le R_0 + 2D|Z| \le R_0 + 2D(D+1)N. \tag{20}$$

A wrapping forward endpoint is simply its chart representative in Z; this causes no alteration of the rectangle argument.

For each original parent B and its final pieces B_k,

$$q_B(u) = \sum_k \sqrt{|B_k|/|B|}\, q_{B_k}(u), \qquad \sum_k |B_k|/|B| = 1.$$

Hilbert-space Cauchy with these exact weights proves

$$\|C_{\mathcal{R}_0} u\|^2 \le \|C_{\mathcal{R}} u\|^2, \qquad V^{-1} \sum_{z \in Z} \|u(z)\|^2 \le \|C_{\mathcal{R}} u\|^2. \tag{21}$$

The second assertion holds because each $z \in Z$ contributes the singleton row $V^{-1/2} u(z)$. The remaining rows are still only component means on rectangles; no gradient or nonrectangular observation has been introduced.

2.4 Charging the complete signed correction

Define the finite number

$$L = \begin{cases} 0, & N = 0, \\ \max\limits_{x,j:\lambda_j(x) \le \sigma} (T - \lambda_j(x)), & N > 0. \end{cases} \tag{22}$$

The spectral construction gives $0 \le A_T(x) - A(x) \le L\mathrm{I}$, with zero difference for $x \notin I$. This is an operator inequality on all r coordinates; in particular it bounds all rotated off-diagonal contributions together. It does not assign signs to individual cross terms. Therefore

$$\begin{aligned} 0 \le Q_{A_T}(u) - Q_A(u) &\le L \sum_{x \in I} \|du(x)\|^2 \\ &\le 2L \sum_{x \in I} \left[D\|u(x)\|^2 + \sum_i \|u(x + e_i)\|^2 \right] \\ &\le 4DL \sum_{z \in Z} \|u(z)\|^2. \end{aligned} \tag{23}$$

Indeed each site has its own bad-start multiplicity at most D and at most D incoming direction-labelled starts. All their endpoints belong to Z. This counts parallel edges at period two separately, as required by the actual form. At $n = 1$ every difference vanishes and the estimate remains valid. No periodic term has been dropped from the signed correction.

Insert (23) into (17), divide the point sum by V, and use (21). This proves (6) with $K = c + 4DL$. The count is

$$\begin{aligned} hR &\le 2Dh \left[1 + B_D + \frac{B_D}{\alpha_D}(F_c(A) - N) + (D+1)N \right] \\ &\le 2Dh \left[1 + B_D + \max\{B_D/\alpha_D, D+1\} F_c(A) \right], \end{aligned}$$

using the nonnegativity of both N and $F_c(A) - N$. This completes the proof of Theorem 1.1.

Remark 2.1 (What the cap does and does not control)**.** *The only use of the budget after the local spectral split is the exact identity* (14)*. In particular its control of observation number is compatible with arbitrarily negative selected eigenvalues, since the penalty absorbs the actual finite correction. When $N = 0$, the correction vanishes and $K = c$. For a positive scalar profile the*

sharper direct scalar construction also has $K = c$ without selecting or replacing any channels. For completeness, choose $M = \max\{0, \max_x[-\lambda_{\min}(A(x))]\}$, so $M \geq 0$ and $A(x) \geq -M\mathrm{I}$ for all x. Then $Q_A(u) \geq -4DM \sum_x \|u(x)\|^2$, because $\sum_x \|du(x)\|^2 \leq 4D \sum_x \|u(x)\|^2$. For the singleton partition $\|C_{\mathcal{R}} u\|^2 = \|u\|_V^2$, so $K = c + 4DM$ proves its asserted finite-penalty alternative. The matrix construction establishes an existence bound with explicit constants, rather than a sharp spectral rank formula.

3 Necessary product order for full-vector common partitions

We prove Theorem 1.2 by an explicit positive diagonal family. The proof concerns all vector fields, as in Theorem 1.1, rather than the smaller coexact domain of the endpoint lower theorem. It uses only local fields and finite-dimensional rank-nullity.

3.1 The coefficient family and every local channel

Let $n = 3m + 3$ and choose centers

$$x_j = (3j + 1, 1, \ldots, 1), \qquad 0 \leq j < m.$$

They lie away from the periodic wrap boundary. With $\sigma = \alpha_D^{1/P} c$, define

$$a_0 = \min\{c/(8D), \sigma/2\} > 0, \qquad M = rn^D, \qquad H = \sigma M > \sigma.$$

In the canonical direction/component basis (i, k) of $\mathbb{R}^{Dh}$ let $A(x)$ be diagonal, with entries

$$\lambda_{i,k}(x) = \begin{cases} a_0, & k = 1 \text{ and } x \in \{x_j, x_j - e_i\} \text{ for some } j, \\ H, & \text{otherwise.} \end{cases} \tag{24}$$

At a center all D direction channels for component one are weak; at an incoming start only its specified direction/component channel is weak. All eigenvalues are positive. For each fixed direction the starts x_j and $x_j - e_i$ are distinct. Across centers their first coordinates lie in disjoint sets $\{3j, 3j + 1\}$. Channels with different direction labels remain distinct even when they share a start. There are therefore exactly $2Dm$ weak channels.

Since $a_0 \leq \sigma/2$, each weak channel has capped cost one in (4). Every other channel has capped cost $\alpha_D c^P H^{-P} = M^{-P}$. Thus, including all strong channels,

$$F_c(A) = 2Dm + (M - 2Dm)M^{-P}, \qquad 2Dm \leq F_c(A) \leq 2Dm + 1. \tag{25}$$

Indeed $0 \leq M - 2Dm \leq M$ and $P > 1$, so the tail is at most $M^{1-P} \leq 1$. Increasing h increases H and keeps this strong tail bounded; it introduces no uncontrolled h-sized budget contribution. The choice of H is permitted by the arbitrary finite profile hypothesis.

3.2 The actual low space and the restriction rank

Let $U_j(x) = 1_{\{x_j\}}(x) e_1$, where $e_1 \in \mathbb{R}^h$ is the first canonical unit vector. These fields are orthonormal in counting units. Their only nonzero direction-labelled forward differences are $-e_1$ at x_j and $+e_1$ at $x_j - e_i$, for each direction i. Both original starts are assigned coefficient a_0 by (24); there are no further wrap contributions. Consequently

$$Q_A(U_j) = 2Da_0 \leq c/4.$$

No nearest-neighbor edge joins distinct centers, and their gradient-start supports are disjoint. Hence the quadratic cross terms vanish and

$$L_m = \operatorname{span}_{\mathbb{R}}\{U_0, \dots, U_{m-1}\}, \quad \dim L_m = m, \quad Q_A(U) \le (c/4) \sum_x \|U(x)\|^2 \quad (U \in L_m). \quad (26)$$

These are general vector functions. No divergence or coexact property is asserted for them.

For any common partition $\mathcal{R}$ into R nonempty rectangles, its actual rows on this space are

$$q_B\left(\sum_j t_j U_j\right) = V^{-1/2}|B|^{-1/2}\left(\sum_{j:x_j \in B} t_j\right) e_1.$$

Each rectangle supplies at most one scalar functional on L_m, so $\operatorname{rank}(C_{\mathcal{R}}|_{L_m}) \le R$. If $R < m$, rank-nullity gives a nonzero $U \in L_m$ with $C_{\mathcal{R}}U = 0$. Coercivity with any finite quadratic penalty, after multiplication by V, would imply

$$c \sum_x \|U(x)\|^2 \le Q_A(U) \le (c/4) \sum_x \|U(x)\|^2,$$

a contradiction. Thus every successful common partition has $R \ge m$. The same argument works for any penalty on these linear observations with $\Phi(0) = 0$, without a differentiability assumption on that penalty.

The full field-space rank is exactly hR, not merely a padded output dimension. To see this, for each B choose a site $z \in B$. A field supported at z with value $\sqrt{V|B|}\, e_k$ has observation e_k in row B and zero in every other row. These fields show surjectivity onto all hR row coordinates. Therefore necessary genuine rank is at least hm. Since $F_c(A) \le 2Dm + 1 \le (2D+1)m$, this proves (7).

3.3 Joint growth and the measurement-class restriction

Set $h = m$. The necessary rank is at least m^2, while $h + F_c(A) \le (2D+1)m + 1$. Their ratio tends to infinity for fixed D, c. Thus a dimension-only $C_D(h + F_c(A))$ bound fails for these common full-vector partitions, even with arbitrary finite penalty. The product lower bound and Theorem 1.1 have matching $hF_c(A)$ order along this family, since $F_c(A) \ge 2Dm$ makes $1 + F_c(A)$ comparable to $F_c(A)$. Sharp coefficients, additive constants and per-profile optimality are not determined.

The extra factor h measures the cost of retaining every component mean on every common partition member; it does not count h independent tests of the one-component low space. The argument does not apply to distinct componentwise partitions, projected or channel-adapted rows, or compressed postprocessed outputs. It requires no external theorem about the rank of those broader observation classes. These elementary rank and local-field arguments carry no claim of historical priority, a new general principle, formal verification, or physical transfer.

4 The weak critical endpoint

The positive and negative results use the same scalar original-start energy. The positive estimate holds on every component field. The negative estimate is already forced on the precisely defined coexact subspace (11).

4.1 Finite layer cake and the upper estimate

Since $a \geq 1$, its actual mean obeys $\mu \geq 1$ and $0 < Y \leq \mu$. For each value y in this finite distribution,

$$y^P = P \int_0^\mu \tau^{P-1} 1_{\{\tau<y\}}\, d\tau.$$

Summing over X proves the exact identity

$$\mathbb{E}_X Y^P = P \int_0^\mu \tau^{P-1} \mathbb{P}_X(Y > \tau)\, d\tau \leq 1 + P M_w \log \mu. \tag{27}$$

The contribution on $(0,1)$ is at most one, and on $[1,\mu]$ we use $\mathbb{P}_X(Y > \tau) \leq M_w \tau^{-P}$. At $\mu = 1$ the second integral is zero; in fact the actual mean and $a \geq 1$ then force $a \equiv 1$. Atoms at strict cutoffs change no ordinary integral. The envelope $a \leq 2\mu$ is retained in the theorem but is not needed for this step.

Corollary A.2 in the appendix, with the actual moment $M = \mathbb{E}_X Y^P$, gives

$$h\#\mathcal{R} \leq 2Dh[1 + 2^P c^P B_D M] T(n,\mu),$$

and the normalized inequality (10) with penalty c. Substitution of (27) proves Theorem 1.3. Its geometry and gap are supplied by the complete scalar construction, rather than by a separate assumed certificate.

4.2 A compact product bump on the actual coexact domain

We now prove Theorem 1.4. Fix a unit vector $e \in \mathbb{R}^g$. Define a globally C^1 function

$$\beta(s) = \begin{cases} s^2(1-s)^2, & 0 \leq s \leq 1, \\ 0, & s \notin [0,1]. \end{cases}$$

We use β for the bump to distinguish it from the critical exponent P. Both β and β' vanish at the endpoints. On $[0,1]$,

$$\beta'(s) = 2s(1-s)(1-2s), \qquad \beta''(s) = 2 - 12s + 12s^2.$$

Thus $0 \leq \beta \leq 1$, $|\beta'| \leq 2$, and β' is globally 2-Lipschitz: its interior second derivative has absolute value at most two, and the zero extension has no derivative jump.

For an integer $L \geq 16$ divisible by 16, put

$$\psi_L(x) = \prod_{i=1}^D \beta(x_i/L), \qquad x \in \mathbb{Z}^D,$$

and define the one-cochain

$$U_1 = D_2^- \psi_L\, e, \qquad U_2 = -D_1^- \psi_L\, e, \qquad U_j = 0 \quad (j \geq 3). \tag{28}$$

Taking $B_{12} = \psi_L e$, $B_{21} = -\psi_L e$ and all other $B_{ij} = 0$ in (11) gives exactly (28), including its signs. Therefore $U = \delta B$ and

$$\sum_j D_j^- U_j = D_1^- D_2^- \psi_L e - D_2^- D_1^- \psi_L e = 0.$$

Every component has zero mean by telescoping. The same identities hold exactly on any torus large enough to contain this support away from wraps.

The one-dimensional lattice sequence $f(t) = \beta(t/L)$ is supported at $t = 1, \dots, L-1$, and $D^- f$ is supported at $t = 1, \dots, L$. It follows from the product formula that U is supported in $\{0, \dots, L\}^D$ and vanishes on every lower coordinate face $x_i = 0$. If $D_i^+ U_j(x) \neq 0$, its original start belongs to this same box. Indeed the only potentially external incoming start is $x_i = -1$; its next endpoint is on the zero lower face, so both field values vanish. Upper outgoing starts $x_i = L$ can contribute and are included. In all other coordinates an external coordinate makes both terms zero. This verifies the boundary-start condition needed below with no interior-only energy convention.

Lemma 4.1 (Finite bump constants)**.** *Let $\mathcal{E}_1(U) = \sum_{x,i,j} \|D_i^+ U_j(x)\|^2$ on $\mathbb{Z}^D$. Then*

$$\begin{aligned} \mathcal{E}_1(U) &\le U_D L^{D-4}, & U_D &= 32D\, 2^D, \\ \sum_x \|U(x)\|^2 &\ge \eta_D L^{D-2}, & \eta_D &= \frac{9}{2^{D+12} 256^{2(D-1)}}. \end{aligned} \tag{29}$$

In particular, with $K_D = U_D/\eta_D$,

$$\mathcal{E}_1(U) \le (K_D/L^2) \sum_x \|U(x)\|^2. \tag{30}$$

Proof. The Lipschitz bound on β' bounds each same-coordinate second finite difference by $2/L^2$. For different coordinates the product of two first differences is bounded by $4/L^2$. All unused factors have magnitude at most one. At most $2D$ component/direction pairs are nonzero; each has at most $(L+1)^D \le 2^D L^D$ original starts by the support analysis. The common squared bound $16/L^4$ gives the first line of (29), including the upper outgoing starts.

For the norm, take $x_2 = L/8 + 1, \dots, L/4$, exactly $L/8$ sites, and $x_i \in [L/4, 3L/4]$ for every $i \neq 2$, at least $L/2$ sites each. On each backward step in coordinate two, the mean-value point belongs to $[1/8, 1/4]$, where $\beta' \ge 3/32$. On the other coordinates $\beta \ge 1/256$ on $[1/4, 3/4]$. Hence

$$\|U_1(x)\| \ge \frac{3}{32L} 256^{-(D-1)}$$

on all these sites. Squaring and multiplying by $(L/8)(L/2)^{D-1}$ gives the stated η_D. It is strictly positive, so division yields (30). □

4.3 Geometric scales, exact mean and the weak distribution

Choose an integer $k \ge 0$ such that $\varepsilon = 4^{-k}$ satisfies $\varepsilon K_D \le c/4$. This choice depends on D, c and is fixed while m varies. Put

$$L_0 = 2^{\max\{4,k\}}, \quad L_j = 2^j L_0 \ (0 \le j < m), \quad L_* = L_{m-1}, \quad n = 16 L_*, \quad H = \varepsilon L_*^2. \tag{31}$$

Every L_j is divisible by 16, and $\varepsilon L_0^2 \ge 1$. Place boxes

$$B_j = q_j + \{0, \dots, L_j\}^D, \quad q_{0,1} = 1, \quad q_{j+1,1} = q_{j,1} + L_j + 3, \quad q_{j,i} = 1 \ (i \ge 2).$$

Their first-coordinate intervals have two unused sites between successive boxes. The rightmost endpoint is $\sum_j L_j + 3m - 2$. Since $\sum_j L_j < 2L_*$ and $m \le 2^{m-1} \le L_*/16$, this is less than

$3L_* + 1 < n$. Other coordinates also lie strictly inside the chart. Thus no support or gradient-start support collides periodically. Translate (28) into each B_j, obtaining $U^{(j)}$. Their field supports and their complete original gradient-start supports are disjoint.

Define the actual scalar coefficient by

$$a(x) = \begin{cases} \varepsilon L_j^2, & x \in B_j, \\ H, & x \notin \bigcup_j B_j. \end{cases} \tag{32}$$

The boxes are disjoint and the top box has the same coefficient as the background. In particular $1 \le a \le H$. Define its mean by its actual sum:

$$\mu = V^{-1} \sum_x a(x) = H - V^{-1} \sum_j |B_j|(H - \varepsilon L_j^2). \tag{33}$$

This is not a freely adjusted normalization. With

$$C_{\mathrm{vol}} = \frac{2^D}{1 - 2^{-D}},$$

the geometric series and $(L_j + 1)^D \le 2^D L_j^D$ imply

$$\sum_j |B_j| \le C_{\mathrm{vol}} L_*^D, \qquad V^{-1} \sum_j |B_j| \le C_{\mathrm{vol}}/16^D < 1/2.$$

It follows that

$$H/2 \le \mu \le H, \qquad \mu \ge 1, \qquad 1 \le a \le H \le 2\mu. \tag{34}$$

The separate inequality $\mu \ge 1$ follows directly from $a \ge 1$, including when $H/2 < 1$.

We check $W(a, \mu) \le 1$ with strict cutoffs. If $0 < \tau \le \mu/H$, then $\tau \le 1$ and $\tau^P \mathbb{P}_X(Y > \tau) \le 1$. If $\tau > \mu/H$, the background and top box are absent. The boxes contributing to $Y > \tau$ satisfy $\varepsilon L_j^2 < \mu/\tau$. If there is any such box, let j be the largest qualifying index. They form an initial segment and have total volume at most $C_{\mathrm{vol}} L_j^D$. Consequently

$$\begin{aligned} \tau^P \mathbb{P}_X(Y > \tau) &\le C_{\mathrm{vol}} (\mu/\varepsilon)^P / V \\ &\le C_{\mathrm{vol}} L_*^D / (16 L_*)^D < 1. \end{aligned} \tag{35}$$

If no box qualifies the probability is zero. Equality at a coefficient atom excludes that box under the strict criterion and cannot increase the bound. This proves the weak distribution assertion for every m, including $m = 1$.

4.4 The low space and every finite local quadratic penalty

Every nonzero gradient of $U^{(j)}$ has its original start in B_j, where (32) has coefficient εL_j^2. All wrap terms vanish for these nonwrapping supports. Thus the complete energy, not merely its interior contribution, is

$$Q_a(U^{(j)}) = \varepsilon L_j^2 \mathcal{E}_1(U^{(j)}) \le \varepsilon K_D \sum_x \|U^{(j)}(x)\|^2 \le (c/4) \sum_x \|U^{(j)}(x)\|^2. \tag{36}$$

The norm bound in Lemma 4.1 makes every $U^{(j)}$ nonzero. Disjoint field supports give norm orthogonality; disjoint gradient-start supports give vanishing quadratic-energy cross terms. The same inequality therefore holds on

$$E_m = \operatorname{span}_{\mathbb{R}}\{U^{(0)}, \dots, U^{(m-1)}\} \subset \mathcal{H}_{\mathrm{coex}}(X), \qquad \dim E_m = m.$$

This proves an actual low subspace, with no local eigenfunction or gap assumption. From (31) and (34),

$$T(n, \mu) \le \left(16\sqrt{2/\varepsilon}\right)^D. \tag{37}$$

The right side is independent of m.

For clarity the finite-dimensional obstruction is proved directly. If $DC(0)|_{E_m}$ has a nonzero kernel vector U, differentiability and $C(0) = 0$ imply $C(tU) = o(t)$. For sufficiently small t, the assumed upper bound on the penalty gives

$$\Phi(C(tU)) \le K\|C(tU)\|^2 = o(t^2).$$

Apply (13) to tU, use the low-space inequality, and divide by t^2. Letting $t \to 0$ would imply

$$c\sum_x \|U(x)\|^2 \le (c/4)\sum_x \|U(x)\|^2,$$

a contradiction. Therefore the derivative is injective on E_m and $\dim F \ge \operatorname{rank}(DC(0)|_{E_m}) = m$. No uniformity of K between profiles was used. Positivity or differentiability of Φ is unnecessary; only its stated finite quadratic upper bound is used. Common normalization of energy and field norm changes none of this finite-dimensional argument.

Finally, disjoint boxes, $|B_j| \ge L_j^D$ and $\mu \ge H/2$ give

$$\mathbb{E}_X(\mu/a)^P \ge V^{-1}\sum_{j=0}^{m-1} L_j^D \left(\frac{\mu}{\varepsilon L_j^2}\right)^P = \frac{m(\mu/\varepsilon)^P}{V} \ge \frac{m}{2^P 16^D}. \tag{38}$$

This establishes every part of (12) and explains why the strong-moment theorem is not contradicted.

4.5 The matching order and its quantifiers

For fixed D, c, ε and L_0 remain fixed. Equations (31) and (34) imply

$$\log \mu = 2(m-1)\log 2 + \log(\varepsilon L_0^2) + \theta_m, \qquad -\log 2 \le \theta_m \le 0. \tag{39}$$

Thus $m/(1 + \log \mu) \to 1/(2 \log 2)$. The lower bound just proved forces output dimension at least a positive fixed-parameter multiple of $1 + \log \mu$ for large m. Theorem 1.3, with $h = Dg$, $M_w = 1$ and (37), gives an upper bound of the same order on these profiles. In particular, no bound proportional to $T(n, \mu)\, o(\log \mu)$ can hold uniformly over the weak-bound class for the centered differentiable/local-quadratic observations under discussion. This proves the final assertion of Theorem 1.4.

The conclusion is a worst-case comparison across profiles. It does not state that each profile needs a logarithm or identify its minimal rank. The lower construction also does not cover a nonzero observation baseline with an unrelated positive penalty at that baseline, nondifferentiable encodings, or penalties outside the stated local quadratic class. Those assumptions matter to the argument rather than merely to its presentation.

5 Model assessments and review context

Lluis Eriksson is the named author. This manuscript was drafted with assistance from Codex, exposed as an OpenAI system based on GPT-6. The exact served backend and reasoning-effort setting are unknown. This drafting assistance is distinct from the involved AI source-audit roles described below.

At preparation, this exact manuscript has not yet been numerically assessed. Its predecessor's judgments do not transfer to these source bytes or to a future compiled PDF. The preparation includes the complete scalar source proof and uses the listed prior source audits of the weak-endpoint, capped-matrix and common-partition rank arguments. Those audits involved AI roles in the same agent system, with parent framing and shared-context limits; the true backend identities are unknown. They are not represented as blind, independent or human reviews. The author of this manuscript is not its source auditor or numerical reviewer.

The complete unit 224 original and logarithmic-addendum, unit 225 and unit 227 source-audit reports, and the complete manuscript 226B source rejection report, remain unchanged external companions. These listed reports, their adverse findings and context disclosures are identified by raw-byte hashes in the source-input manifest; no separate original scalar-audit report is identified there. The earlier unit 224, 225 and 227 qualified acceptances concern the named arguments and do not certify priority, practical complexity, physical transfer, optimal constants, all-size execution, or this manuscript's presentation. No numerical grade is inferred from them. This manuscript requires its own review of the exact source and any exact PDF produced from it.

A complete source audit rejected the exact preparation 226B because its optional singleton-partition derivation omitted the necessary condition $M \geq 0$. For example, its former conditional allowed $n = 1$, $A = \mathrm{I}$, $M = -1$ and a nonzero constant field, for which the asserted auxiliary lower bound is false. That report found no load-bearing defect in the principal theorems. Preparation 226C explicitly chooses a nonnegative spectral depth in that auxiliary derivation and narrows the report-identification wording. The original rejection is preserved in full, including its editorial finding; it is not replaced by this author's correction. This exact corrected source awaits a fresh whole-source audit, and its first exact numerical assessment remains absent. No result of a render of 226B is assumed here.

The selected diagnostics recorded for unit 224 used owned Colab Linux runs in normal and optimized modes: six direct tensor-sum cases, one actual $D = 3$, $L = 16$ coexact case, fifteen profiles and three rejected real negative controls. Unit 225 recorded seven finite matrix families, 131 actual fields and four rejected real controls, including endpoint, normalization and full off-diagonal correction mutations. The normal and optimized outputs were recorded as identical with zero actual exits. These are selected diagnostics, not an all-size formal proof or a two-fresh-clone replay of the construction. Their runtimes were released and read back inactive; billed compute units are unknown. The mathematical proofs above and below do not depend on these diagnostic counts.

Source identity and external evidence

The scalar appendix transcribes the complete mathematical portion of `balanced-rings223.tex` at programme commit

`ddd2cf605c1715a03d2e87991e920c0ec760e527`.

Its raw Git-blob SHA-256, without checkout conversion or normalization, is

6877841cde546f8920c39b98a5dfef8a
f385b1a2d5a4842f95a9dd36985b3dc2.

Changes to that mathematical portion are notation and context adaptations: arbitrary h real components are obtained by summing the same scalar inequality; normalized and counting rows are distinguished; cross-references are internal to this manuscript. The original extended operational and assessment history is retained externally rather than interleaved with the proof.

The full subsequent analytical sources are the unit 224 draft and its logarithmic addendum, and the unit 225 capped-matrix draft. Their raw-byte SHA-256 digests, respectively, are

7a94460f2cbbaf6ac291d64981a7ae6e7
9590067f315923b615f842d6a0b211c

0720cedc459361539c9cedf95b7b1868
31a6ab9414ea27c859918203a39fb059

a9c4bcbc3b8764b2e2cf84b7906ae5fe9
ed3fa96c5f5ae22c68eb665a0f7d448.

The related unit 227 draft and complete source-only audit are also included by identity in the external manifest. Its audit returned a qualified acceptance for the stated common-full-vector measurement class. No unit 227 diagnostics are claimed in this preparation; its lower proof uses no external spectral $h + F$ theorem.

The accompanying `SOURCE-INPUTS.json` names the exact source files, the listed audit reports, context disclosures and unit result records with their raw-file hashes. All prior reports remain unchanged. This evidence identifies the inputs; it does not turn their recorded finite checks into an executable proof for every size.

A Coefficient-selected rectangular means for arbitrary scalar profiles

This appendix gives the complete scalar construction used in the main proofs. It uses the original positive-start scalar coefficient and retains the edges crossing rectangular observation pieces inside a ring. It does not identify a variable-weight curl/divergence form with that gradient form. The balanced-ring, BV and tree mechanisms are classical; historical originality of this finite specialization is unestablished.

A.1 Statement and actual observations

Let integers $D \geq 3$, $n \geq 1$, and $c > 0$ be fixed, and put $X = \{0, \dots, n-1\}^D$. For any finite scalar profile $a(x) > 0$ define

$$P = D/2, \quad b = 2^D, \quad J = 2D, \qquad C_{\mathrm{ext}} = 1 + J\{2J(1 + 12J) + 1\}, \qquad G_D = 2^{2D+10} D^2 C_{\mathrm{ext}}^2.$$

Write $S(A) = \sum_{x \in A} a(x)^{-P}$ and $t = (cG_D)^{-P}/b$. Every internal edge below keeps its original positive start, and all sums and norms initially use counting measure. Here $\mathcal{E}_a(f; \Omega)$ denotes the sum over every chart-internal edge with both endpoints in Ω, weighted by its original positive-start value $a(x)$; $\mathcal{E}_1$ has unit weights. Thus the counting row $\widehat{q}_B$ used at the end of this appendix equals $V^{1/2} q_B$ in the normalized notation of (2).

Theorem A.1. *There is a partition $\mathcal{R}$ of X into nonempty axis-aligned integer rectangles, determined by a and fixed geometric conventions alone, such that*

$$\#\mathcal{R} \le 2D\{1 + (2b+3)S(X)/t\},$$

and, simultaneously for every real field f,

$$\sum_{i,x:\ x,x+e_i\in X} a(x)|f(x+e_i) - f(x)|^2 + c\sum_{B\in\mathcal{R}} \frac{1}{|B|}\Big(\sum_{x\in B} f(x)\Big)^2 \ge c\sum_{x\in X}|f(x)|^2. \tag{40}$$

The original periodic energy also satisfies this inequality, by adding its nonnegative wrap terms. For every integer $h \ge 1$ real components the same partition has output dimension $h\#\mathcal{R}$ and penalty $\Phi(y) = c\|y\|^2$.

Corollary A.2. *For any reference value $\mu \ge 1$, if $n^{-D}\sum_x (\mu/a(x))^{D/2} \le M$, the actual normalized rectangular means have output dimension at most $C_{\mathrm{rect}}T(n,\mu)$, with $K = c$ and*

$$C_{\mathrm{rect}} = 2Dh\{1 + (2b+3)b(cG_D)^P 2^P M\}.$$

The exact arithmetic mean and envelope $1 \le a \le 2\mu$ are unnecessary for this sufficiency. The inverse-moment hypothesis is still assumed; it is not deduced from an arithmetic mean.

Proof. $S(X) \le Mn^D\mu^{-P} \le 2^P M\max\{1, n^D/(1+\mu)^P\}$, since $\mu \ge 1$. Insert $1/t = b(cG_D)^P$ into the theorem. Apply its scalar inequality to each of the h real components. Dividing energy and squared norms by $V = n^D$ gives observation rows $V^{-1/2}|B|^{-1/2}\sum_B f = \sqrt{|B|/V}\ \mathrm{mean}_B f$. The all-field inequality restricts to divergence-zero fields, and linear means with the displayed quadratic penalty supply the original local data. □

A.2 Balanced integer nodes and actual collars

Split every coordinate interval with at least two sites into its floor/ceiling halves; a singleton remains unchanged. Use coordinate order $1, \dots, D$, lower before upper, and lexicographic order on children. Each node has at most b children, and all side lengths differ by at most one. At global depth k the nonempty intervals have lengths h_k or $h_k + 1$, until they are singletons, where $h_k = \max\{1, \lfloor n/2^k \rfloor\}$.

Consider a proper descendant I of an ancestor Q, at global depth k, and set $h_0 = h_k$. Every side of I belongs to $[h_0, 2h_0]$. Every positive clearance from I to the boundary of Q is a union of depth-k intervals, so is at least h_0. This uses the global h_0, not the smallest side of this particular I: at odd sizes that substitution can be false. If every side of Q is at least two, a proper descendant is shorter in every direction, so each axis has at least one positive clearance. If the smallest side of Q is one, every side is one or two; each proper descendant is a corner singleton. That small case is treated separately below.

For the large case put $\Omega = Q \setminus I$. Expand I by exactly h_0 on every face with positive clearance, leaving a boundary-touching face unchanged, and call the resulting rectangle B. Its side lengths ℓ_i lie in $[2h_0, 4h_0]$. The collar $A = B \setminus I$ is covered by the actual positive-clearance face slabs: normal width h_0 and all other sides those of B. There are between D and $J = 2D$ slabs, each contained in Ω. For slabs S_i, S_j with distinct normals, direct product counting gives

$$\frac{|S_i|}{|S_i \cap S_j|} = \frac{\ell_j}{h_0} \le 4, \qquad \frac{|S_i|}{|S_j|} = \frac{\ell_j}{\ell_i} \le 2. \tag{41}$$

Choose the first slab as root; attach every distinct-normal slab directly to it, and its opposite slab, if present, through the first slab with a different normal. This tree has depth at most two.

A.3 An unweighted BV extension, with every hole-crossing edge

For a finite product rectangle S of longest side at most $4h_0$, write $V_1(S) = \sum_S |f - \mathrm{mean}_S f|$ and let $\mathrm{TV}_S(f)$ be the sum of $|df|$ over its actual internal edges. Then

$$V_1(S) \le 2h_0 \mathrm{TV}_S(f). \tag{42}$$

Indeed on an interval with m sites, averaging pair paths bounds the coefficient of its rth edge by $2r(m-r)/m \le m/2$. Commuting coordinate-average projections contract L^1; telescoping their product and applying this interval inequality fibrewise proves (42), without an additional factor D.

Let m_i be the mean on S_i and let c_0 be the root mean. Averaging on the actual intersection and using (41) gives $|S_i||m_i - m_j| \le 4\{V_1(S_i) + V_1(S_j)\}$. For a depth-two path i, p, r, the second edge costs the volume ratio $|S_i|/|S_p| \le 2$, so

$$|S_i||m_i - m_r| \le 4V_1(S_i) + 12V_1(S_p) + 8V_1(S_r) \le 12 \sum_j V_1(S_j).$$

Direct paths satisfy the same safe upper bound. Covering A and counting each actual edge at most J times yields, with $P_D = 2J(1 + 12J)$,

$$\sum_A |f - c_0| \le (1 + 12J) \sum_j V_1(S_j) \le P_D h_0 \mathrm{TV}_A(f). \tag{43}$$

At each positive-clearance hole face take the adjacent outside layer of width h_0, with tangential sides those of I. It is contained in A. On a normal fibre its nearest value satisfies

$$|f(\mathrm{first}) - c_0| \le h_0^{-1} \sum_{\mathrm{fibre}} |f - c_0| + \sum_{\mathrm{fibre}} |df|.$$

Sum over all actual faces and use (43). The sum of hole-crossing jumps to the constant c_0 is at most $J(P_D + 1)\mathrm{TV}_A(f)$. Fill I by c_0 and leave f unchanged on Ω, obtaining $\widetilde{f}$ on Q. Internal hole differences vanish; the added variation consists precisely of these crossing faces. Therefore

$$\mathrm{TV}_Q(\widetilde{f}) \le C_{\mathrm{ext}} \mathrm{TV}_\Omega(f). \tag{44}$$

Boundary-touching faces create no exterior edge; at $h_0 = 1$ the fibre gradient sum is zero. No coefficient is extended, reflected, averaged or reassigned. In particular (44) makes no weighted quadratic trace-extension claim.

A.4 Actual critical Sobolev and weighted gap on a ring

Here are the finite details needed to pass from BV to the original energy. For $T \subset \{0, \dots, m-1\}^D$ with $|T| \le m^D/2$, the occupied-fibre argument and the discrete projection inequality give

$$|\partial T| \ge \sum_i |\pi_i T| - D|T|/m \ge D|T|^{(D-1)/D} - D|T|/m \ge c_{\mathrm{iso}} |T|^{(D-1)/D}, \quad c_{\mathrm{iso}} = D(1 - 2^{-1/D}) \ge \tfrac{1}{2}.$$

Each occupied nonfull fibre supplies an edge, while full fibres number at most $|T|/m$. The projection inequality $|T|^{D-1} \le \prod_i |\pi_i T|$ is the finite Loomis–Whitney inequality; it follows from

iterated Hölder. A singleton cube has only zero variance. Median parts, finite coarea and Minkowski consequently imply

$$\inf_z \|u - z\|_{D/(D-1)} \le (2/c_{\mathrm{iso}})\mathrm{TV}(u)$$

on the cube. On a rectangle of aspect at most two, fold a cube of its longest side onto it: in each coordinate first visit the interval in order, repeat its final endpoint and reflect until the prescribed length. Every vertex and edge has at most two preimages per coordinate; the initial product block maps identically onto the rectangle. Restriction and the variation multiplicity at most 2^D give the safe constant $C_{\mathrm{full}} = 2^{D+1}/c_{\mathrm{iso}}$ for that rectangle, still unweighted.

Apply this bound and (44) to an indicator on Ω. For every real z, $\|1_T - z\|_{s,\Omega} \ge \frac{1}{2}\min\{|T|, |\Omega\setminus T|\}^{1/s}$, where $s = D/(D-1)$. Thus Ω has relative edge isoperimetric constant

$$c_\Omega = c_{\mathrm{iso}}/(2^{D+2}C_{\mathrm{ext}}).$$

Coarea gives $\|g\|_s \le c_\Omega^{-1}\mathrm{TV}_\Omega(g)$ for nonnegative g supported on at most half the vertices. With a median m_f of f, take $w = (f - m_f)_+$ or $(m_f - f)_+$ and put

$$\alpha = 2(D-1)/D, \qquad q = 2D/(D+2), \qquad q' = 2D/(D-2).$$

Then $\alpha s = 2$ and $(\alpha - 1)q' = 2$. The power-difference inequality, Hölder and degree at most $2D$ yield

$$\mathrm{TV}(w^\alpha) \le 2\alpha D^{1/q'}\|dw\|_q \|w\|_2^{\alpha-1}.$$

In fact the q'th power of the endpoint factor sums to at most $2^{q'}D\sum_x w(x)^2$. Divide only when $w \ne 0$. The positive and negative truncated q-edge energies sum to at most the original one; since $2/q \ge 1$, their squared bounds combine to give

$$\sum_\Omega |f - \mathrm{mean}_\Omega f|^2 \le K_\Omega^2\Big(\sum_{\Omega\ \mathrm{edges}} |df|^q\Big)^{2/q}, \qquad K_\Omega = 2\alpha D^{1/q'}/c_\Omega. \tag{45}$$

Only now use weights, on the original retained positive-start edges:

$$\Big(\sum |df|^q\Big)^{2/q} \le \Big(\sum_{\Omega\ \mathrm{edges}} a(\mathrm{start})^{-D/2}\Big)^{2/D}\mathcal{E}_a(f;\Omega) \le \{DS(\Omega)\}^{2/D}\mathcal{E}_a(f;\Omega).$$

Here $q/(2-q) = D/2$, and at most D retained edges start at one vertex. Since $\alpha \le 2$, $c_{\mathrm{iso}} \ge 1/2$ and $2/q' + 2/D = 1$, $K_\Omega^2 D^{2/D} \le 2^{2D+10}DC_{\mathrm{ext}}^2 \le G_D$. Full balanced rectangles obey the same safe bound without removing a hole.

For the small nodes with a singleton direction, Q, or Q minus a corner, is a connected graph with at most b vertices whenever nonempty. Pair paths of length at most $b-1$ give variance at most $b^2\mathcal{E}_1$. Each internal edge satisfies $a(\mathrm{start})^{-D/2} \le S(\Omega)$, so $\mathcal{E}_1 \le S(\Omega)^{2/D}\mathcal{E}_a$. As $G_D \ge b^2$, this proves in every case

$$\sum_\Omega |f - \mathrm{mean}_\Omega f|^2 \le G_D S(\Omega)^{2/D}\mathcal{E}_a(f;\Omega). \tag{46}$$

A singleton has zero variance. Empty groups are never output.

A.5 Coefficient-only heavy-chain compression and complete count

Give each vertex mass $a(x)^{-P}$ and call a tree node heavy if its mass is strictly greater than t. If the root is light, output it. Otherwise let L be the number of terminal heavy nodes and B_h the number with at least two heavy children. Terminal heavy nodes are disjoint, so $L \le S(X)/t$. Finite tree counting gives $B_h \le L-1$. After contracting unary paths, there are at most B_h+L-1 skeleton edges and possibly a root-leading chain; hence at most $B_h + L \le 2L-1$ maximal unary chains. The bound is on chains, not their unrestricted total length.

At a unary node $Q \to Q'$ the increment $Q \setminus Q'$ is the union of at most $b-1$ light children, of mass at most $(b-1)t$. Traverse a maximal chain from ancestor toward descendant, closing a group when the cumulative mass first reaches t. A closed group is a balanced ancestor ring $Q_i \setminus Q_j$ of mass in $[t, bt)$; a residual is output only if nonempty, and has mass less than t. Across all chains the disjoint closed groups number at most $S(X)/t$, while residuals number at most $2L-1$. At branching or terminal heavy internal nodes output each light child; at a heavy singleton terminal output the singleton. Such cubes/singletons number at most $b(B_h+L) \le b(2L-1)$.

These groups partition X: a unary ring excludes the endpoint heavy subtree, which is then handled at the endpoint; branching processing recurses into every heavy child and outputs all remaining children. Thus their total number is at most

$$S(X)/t + (2L-1) + b(2L-1) \le 1 + (2b+3)S(X)/t.$$

Each nonsingleton group has mass at most $bt = (cG_D)^{-D/2}$, so (46) gives residual gap at least c. Heavy singleton mass may exceed this threshold, but its variance is zero and its mean captures its entire norm.

A.6 Rectangular rows while retaining intra-ring interfaces

Partition $Q \setminus I$ by the first coordinate outside I: preceding coordinates belong to I, the current coordinate is below or above I, and later coordinates fill Q. This gives at most $2D$ nonempty integer rectangles. Cubes and singletons require one rectangle. With $\widehat{q}_A(f) = |A|^{-1/2}\sum_A f$, Cauchy gives $\widehat{q}_\Omega(f)^2 \le \sum_{B\subset\Omega} \widehat{q}_B(f)^2$. The actual regional gap therefore implies

$$\mathcal{E}_a(f;\Omega) + c\sum_{B\subset\Omega} \widehat{q}_B(f)^2 \ge c\|f\|_\Omega^2.$$

Crucially $\mathcal{E}_a(f;\Omega)$ retains all edges between rectangular pieces of this same ring; no separate gap is asserted for each piece. Distinct regions are disjoint, so summing discards only nonnegative edges between different regions. The full energy dominates their internal sum. The count of rectangles is at most $2D$ times the region count, proving Theorem A.1. The retained intra-ring interfaces are part of the regional energy throughout.

References

[1] R. L. Frank, *Cwikel's theorem and the CLR inequality*, arXiv:1206.3325v1 (2012), Theorems 2.1 and 3.1. `https://arxiv.org/abs/1206.3325v1`.

[2] Yu. Brudnyi, *Nonlinear Piecewise Polynomial Approximation and Multivariate BV spaces of a Wiener–L. Young Type. I*, arXiv:1511.03971v1 (2015). `https://arxiv.org/abs/1511.03971v1`.

[3] A. Cohen, R. DeVore, P. Petrushev and H. Xu, *Nonlinear approximation and the space* $BV(\mathbb{R}^2)$, American Journal of Mathematics **121** (1999), 587–628. Author PDF: `https://people.math.sc.edu/pencho/Publications/BV9707.pdf`.