

Clustering and the Transfer-Operator Gap: A Machine-Checked Dense-Family Criterion

Observable-dependent prefactors are harmless
at a common exponential rate

Lluis Eriksson

July 27, 2026 (version 1.2)

Abstract

Inside a Lean 4 formalization programme for four-dimensional $SU(N_c)$ lattice Yang–Mills, we machine-check the operator-theoretic criterion that stands between exponential decay of a Euclidean correlator and a spectral gap of a transfer operator. Let T be a bounded self-adjoint operator on a Hilbert space and Ω a unit vector *fixed by* T , so that $T\Omega = \Omega$, and put $S = T - |\Omega\rangle\langle\Omega|$. Exponential decay at rate r of the *connected* two-point function $\langle v, T^n v \rangle - |\langle \Omega, v \rangle|^2$ at every v is *equivalent* to the operator-norm bound $\|S\| \leq r$.

The substantive part is the dense-family criterion. Writing $\mathcal{D}_r = \{v : \exists C, \forall n, \|S^n v\| \leq Cr^n\}$, we prove that $\mathcal{D}_r$ is a linear subspace and that its density alone forces $\|S\| \leq r$, the constants being entirely unconstrained: a family of observables whose span is dense, each carrying its own finite constant, suffices. Consequently prefactors that grow with the support of the observable — the shape cluster expansions produce — do *not* obstruct the gap, provided the exponential *rate* is common to the family and the family spans densely. Those two provisos are essential; without them the statement is false.

No mathematical novelty is claimed for the criterion itself, which we expect to be known in the language of local spectral theory; what is offered is its mechanization, its packaging for families of observables, and the consequence for prefactors. We also record what the formalization does *not* contain: no Osterwalder–Seiler Hilbert space for any gauge theory, no reflection positivity of the Wilson measure, no identification of a Euclidean correlator with a matrix element. Nothing here is a claim about the continuum limit or about the Clay problem. All results are machine-checked with no **sorry** and no project axioms.

1 What is proved, and what is not

The gap in the chain

A lattice mass gap, in the sense of Osterwalder and Seiler [1], is a statement about the spectrum of an operator: the transfer operator T obtained from the Euclidean measure by reflection positivity satisfies $0 \leq T \leq 1$, fixes the vacuum vector Ω , and has a spectral gap on $\Omega^\perp$. Exponential decay of a Euclidean correlator is, by itself, a statement about a real-valued function of the separation. The two are connected by a theorem, not by terminology.

Within the formalization programme this paper belongs to, that theorem was absent. The terminal statement of the lattice assembly concludes about a function $\text{cov} : \mathbb{N} \rightarrow \mathbb{R}$; the programme contains no Hilbert space of states, no transfer operator, no reflection positivity, and no spectrum. This paper supplies the *autonomous* half of the missing link: the operator-theoretic

equivalence, at the level of generality where it is a theorem, machine-checked, together with the sharpening that determines what an application must supply.

Scope, stated in advance

The following are *not* established in this paper, nor anywhere in the Lean development accompanying it: reflection positivity of the Wilson gauge measure; the Gelfand–Naimark–Segal quotient; the existence of a transfer operator for any gauge theory; the identification $\mathbb{E}[A \cdot \theta_n B] = \langle A\Omega, T^n B\Omega \rangle$; density of any concrete family of gauge-invariant observables in the fluctuation sector of any such space. Consequently nothing in this paper is a claim about Yang–Mills, about the continuum limit, or about the Clay problem, and no statement below should be read as transferring importance from the subject matter of the programme to the theorems proved here. The equivalence itself is classical; see Section 9.

2 Setting

Let H be a Hilbert space over $\mathbb{R}$ (Sections 3) or over $\mathbb{C}$ (Sections 4–5).

Definition 2.1 (transfer data). A pair (T, Ω) with T a bounded operator on H and $\Omega \in H$ is *transfer data* when T is symmetric, $\|\Omega\| = 1$, and $T\Omega = \Omega$. We write $P = |\Omega\rangle\langle\Omega|$ for the rank-one projection $Pv = \langle\Omega, v\rangle\Omega$ and

$$S = T - P$$

for the *vacuum-projected* operator. These are the properties that reflection positivity is used to produce; here they are hypotheses.

Definition 2.2 (connected two-point function). For $v \in H$ and $n \in \mathbb{N}$,

$$G_v(n) = \langle v, T^n v \rangle - \langle v, \Omega \rangle \langle \Omega, v \rangle.$$

The elementary reason the *truncated* correlator, and not the full one, is the object that sees a gap is the following identity, which is proved by induction from $S\Omega = 0$ and $\langle\Omega, Sv\rangle = 0$.

Lemma 2.3 (structural identity). *For transfer data and every $n \geq 0$, $S^{n+1}v = T^{n+1}v - \langle\Omega, v\rangle\Omega$, and hence $G_v(n+1) = \langle v, S^{n+1}v \rangle$.*

3 The equivalence

Theorem 3.1 (clustering $\iff$ gap). *Let (T, Ω) be transfer data on a real Hilbert space and $r \geq 0$. Then*

$$\|S\| \leq r \iff \forall v \in H, \exists C, \forall n, |G_v(n)| \leq C r^n.$$

Two features of the statement are deliberate. First, the conclusion is an operator-norm bound and not an assertion of the form $\exists m > 0$; the latter shape admits vacuous witnesses, and this programme’s own history contains one. Second, the statement is an equivalence, so it cannot be weakened into vacuity in either direction.

Proof sketch. Forward. For $n \geq 1$, Lemma 2.3 gives $|G_v(n)| = |\langle v, S^n v \rangle| \leq \|v\| \|S^n\| \|v\| \leq \|v\|^2 r^n$; the $n = 0$ term is $\|v\|^2 - |\langle \Omega, v \rangle|^2 \in [0, \|v\|^2]$ by Cauchy–Schwarz.

Backward, the case $r = 0$. Lemma 2.3 at $n = 1$ gives $0 \leq \|Sv\|^2 = G_v(2) \leq C_v \cdot 0^2 = 0$ for every v , so $S = 0$ and $\|S\| \leq 0$.

Backward, the case $r > 0$. For $n \geq 1$ — and only for $n \geq 1$, since $G_v(0) \neq \langle v, S^0 v \rangle$ in general — Lemma 2.3 and self-adjointness give $G_v(2n) = \langle v, S^{2n} v \rangle = \|S^n v\|^2$, so clustering at even times gives $\|S^n v\| \leq \sqrt{C_v} r^n$ pointwise. Banach–Steinhaus applied to $\{r^{-n} S^n : n \geq 1\}$ upgrades this to $\|S^n\| \leq M r^n$; and for symmetric S one has $\|S^2\| = \|S\|^2$ — proved from Cauchy–Schwarz, without the C^* -algebra API — hence $\|S\|^{2^k} \leq M r^{2^k}$ for all k , and $\|S\| \leq r$. $\square$

4 The sharp form: the constants do not matter

Theorem 3.1 assumes clustering at *every* vector. A cluster expansion never delivers that: it bounds a *family* of observables. On a merely dense set, the elementary argument — extend the quadratic-form bound by continuity — requires the constant to be uniformly quadratic, $C_v = K\|v\|^2$. That requirement is a property of that proof. It is not necessary.

Definition 4.1 (decay domain). For a bounded operator S on H and $r \geq 0$,

$$\mathcal{D}_r(S) = \{v \in H : \exists C \geq 0, \forall n, \|S^n v\| \leq C r^n\}.$$

(The Lean definition leaves r and C unconstrained in $\mathbb{R}$ and carries $r \geq 0$ as a hypothesis at each use site, which is why the subspace lemma below holds for either sign; the restriction here is for readability.)

Lemma 4.2. $\mathcal{D}_r(S)$ is a linear subspace of H (a *Submodule* in the formalization).

Proof. $C_{u+v} \leq C_u + C_v$ and $C_{\lambda v} = |\lambda| C_v$. $\square$

Lemma 4.2 contains the algebraic step that passes from a controlled family to its linear span; the substantive mathematics is the functional-calculus argument of Theorem 4.3. What the lemma buys is that per-observable constants, however fast they grow along any exhaustion of the space, close up under linear combination with no uniformity whatsoever.

Theorem 4.3 (sharp form). *Let S be a self-adjoint bounded operator on a complex Hilbert space $H \neq 0$ and $r \geq 0$. If $\mathcal{D}_r(S)$ is dense in H , then $\|S\| \leq r$. Consequently, if $D \subseteq H$ has dense span and every $v \in D$ admits some finite C_v with $\|S^n v\| \leq C_v r^n$ for all n , then $\|S\| \leq r$.*

Proof. Suppose $\|S\| > r$. The norm is attained on the spectrum, so there is $\lambda_0 \in \sigma(S)$ with $|\lambda_0| = \|S\|$; choose $r < c < |\lambda_0|$ and set $f(x) = \max(0, |x| - c)$, so that f vanishes on $\{|x| \leq c\}$ and $f(\lambda_0) > 0$. Define

$$k_n(x) = \frac{f(x)}{(\max(c, |x|))^{2n}}.$$

The denominator is $\geq c^{2n} > 0$ for *every* real x , so k_n is continuous with no case analysis and no division by zero, even when $0 \in \sigma(S)$; and $k_n(x) x^{2n} = f(x)$ identically, the even power letting $|x|$ stand in for x . Hence, in the continuous functional calculus, $f(S) = k_n(S) S^{2n}$, and on the spectrum $\|k_n\|_\infty \leq \|S\|/c^{2n}$. For $v \in \mathcal{D}_r(S)$,

$$\|f(S)v\| = \|k_n(S) S^{2n} v\| \leq \frac{\|S\|}{c^{2n}} C_v r^{2n} = \|S\| C_v \left(\frac{r}{c}\right)^{2n} \xrightarrow{n \rightarrow \infty} 0,$$

so $f(S)$ annihilates $\mathcal{D}_r(S)$, hence all of H by density and continuity. But $\|f(S)\| \geq |f(\lambda_0)| > 0$, a contradiction. The second statement follows from Lemma 4.2, which places the span of D inside $\mathcal{D}_r(S)$. $\square$

Why a functional calculus is needed

Three routes that suggest themselves all fail. Recording why is the main reason a formalization of this criterion is worth reading: it says which architecture survives, not merely that a theorem type-checks.

Banach–Steinhaus does not apply. A dense *subspace* may be meagre — the span of a countable family is — so the non-meagreness hypothesis is not at hand, and pointwise bounds on such a set yield no uniform bound. This is exactly the difference between Theorem 3.1, where the hypothesis is at every vector of H and completeness pays for the uniformity, and Theorem 4.3, where it is not.

Powers do not suffice. Put $B = S^2/\|S\|^2$, a positive contraction with $\|B\| = 1$. It can happen that $B^n v \rightarrow 0$ for *every* v while $\|B^n\| = 1$ for every n — multiplication by x on $L^2[0, 1]$ is exactly that — so the norm level alone yields no contradiction. (For a general self-adjoint S the powers need not tend to 0 at all: if $\pm\|S\|$ is an eigenvalue they converge to the corresponding eigenprojection. What is asserted here is only that the norm-level argument admits a counterexample.)

The resolvent route stalls. A vector $v \in \mathcal{D}_r(S)$ lies in the range of $t - S$ for every $t > r$ by a Neumann series, but with a vector-dependent bound; and a self-adjoint operator with dense range need not be invertible.

The statement is about the spectral measure near the top of the spectrum, and only a functional calculus sees that.

Corollary 4.4 (support-exponential constants are not an obstruction). *Let $\{A_s\}$ be observables whose span is dense in H , indexed so that A_s has support of size s , and suppose $\|S^n A_s\| \leq C(s) r^n$ for all n , with a rate $0 \leq r < 1$ independent of s . Then $\|S\| \leq r$, whatever the growth of C . In particular $C(s) = e^{\kappa s}$ is admissible.*

Corollary 4.4 is the reason the sharp form matters for the programme this paper belongs to. The volume-uniform area law of [5] carries a constant of the form $N_c e^{|\text{supp}| \cdot 4dK}$, exponential in the edge support of the loop, while its decay *rate* is the support-independent polymer activity rate; and the connecting-cluster estimate of the same programme has a support-independent constant but covers only a local, non-total family. A natural reading of that pair is that the two horns cannot be closed simultaneously, so that a total family must carry growing constants and the gap must fail volume-uniformly. Corollary 4.4 shows that reading is false: the growth of the constant is irrelevant, and only the uniformity of the *rate* is load-bearing.

We stress the limited role of that citation. The area-law theorem is used here *only* to motivate the shape of an observable-dependent constant. It does not supply a dense family in any Osterwalder–Seiler space, it does not supply the time-direction correlator, it does not supply a common rate on such a family, and it does not supply the identification with T^n . Nothing in this paper brings a mass gap closer to being deduced from an area law.

5 The bridge in its own objects, and volume-uniformity

Theorem 4.3 is about a bare self-adjoint operator. Composing it with Definitions 2.1–2.2 gives the bridge in the form an application would consume.

Theorem 5.1 (sharp bridge). *Let (T, Ω) be transfer data on a complex Hilbert space $H \neq 0$ with T self-adjoint, and let $r \geq 0$. Then $\|S\| \leq r$ if and only if there is a set $D \subseteq H$ with dense span such that every $v \in D$ admits some finite C with $|G_v(n)| \leq C r^n$ for all n .*

Theorem 5.2 (volume-uniformity). *Let (T_i, Ω_i) be transfer data on complex Hilbert spaces $H_i \neq 0$, i ranging over any index set, and let $0 < r < 1$ be a common rate for which each i admits a family as in Theorem 5.1 — with per-observable constants that need not be related across i or within it. Then, with $m = -\log r > 0$,*

$$\|T_i - |\Omega_i\rangle\langle\Omega_i|\| \leq e^{-m} \quad \text{for every } i.$$

No uniform control of the constants $C_{v,i}$ is assumed. The only cross-volume quantitative input is the common rate r ; once that rate is fixed, Theorem 5.1 yields the same operator-norm bound for every i , its conclusion being a bound whose constant never sees the index. In the intended reading i is a finite volume and m is a volume-uniform mass; we stress that no transfer operator for any gauge theory is constructed here, so the reading is an intention and not a result.

6 Non-vacuity

Every statement above is exhibited with inhabitants in which the fluctuation sector is nonzero, so that none of them is satisfied only in the degenerate one-dimensional case.

Proposition 6.1. *Let Ω, w be orthonormal in H and $0 < r < 1$. Then $T = r \text{id} + (1 - r)P$ is transfer data with*

$$\|T - |\Omega\rangle\langle\Omega|\| = r \quad \text{exactly,} \quad T - |\Omega\rangle\langle\Omega| \neq 0, \quad -\log r > 0.$$

The hypotheses are discharged concretely in $\mathbb{R}^2$ and in $\mathbb{C}^2$ with $\Omega = e_0$, $w = e_1$, so the family is inhabited.

7 What remains

Within the abstract implication isolated here, the remaining bridge obligations are: reflection positivity of the Wilson gauge measure in the time direction; the GNS quotient and the transfer operator with $0 \leq T \leq 1$ and $T\Omega = \Omega$; the identification of the Euclidean correlator with a matrix element; and density, in the fluctuation sector of the resulting space, of a family of gauge-invariant observables for which a clustering estimate at a common rate is available. Each is classical mathematics [1, 2] that has not, to our knowledge, been mechanized; none is established here.

We do not claim that list is exhaustive for any particular notion of “mass gap”. Depending on the formulation one may additionally need uniqueness of the vacuum, control of the rate at all times rather than along a subsequence, the correct gauge-invariant sector, the relation to a Hamiltonian, and a volume or lattice-spacing limit. What we do claim is that the criterion proved above is one obligation that is now discharged, and that the prefactor question is not among the remaining ones.

Remark 7.1 (a boundary of the formalization, not of the mathematics). *Theorem 3.1 is formalized over a real Hilbert space and Theorems 4.3–5.2 over a complex one. The reason is that the library’s continuous functional calculus for self-adjoint elements is derived from the complex case and is unavailable for operators on a real space, and the library contains no complexification of a real inner-product space. The two are separate instantiations of one argument; the transport between them is not formalized here and is not claimed. This is a boundary of the formalization, not of the mathematics.*

8 Formalization architecture

The development is three modules and runs in one direction:

projection identities $\rightarrow$ elementary equivalence $\rightarrow$ decay subspace
 $\rightarrow$ continuous cut-off $\rightarrow$ sharp bridge.

What the library did not provide. Two absences shaped the design. Mathlib has no complexification of a real inner-product space, and its continuous functional calculus for self-adjoint elements is derived from the complex case, so it is unavailable for operators on a real Hilbert space. Theorem 3.1 needs no functional calculus and is therefore stated over $\mathbb{R}$; Theorems 4.3 and 5.1 need one and are stated over $\mathbb{C}$. That split is forced by tooling, not by mathematics, and the transport between the two is not formalized (Remark 7.1).

Two devices that may be reusable. First, the C^* step $\|S^2\| = \|S\|^2$ for symmetric S is proved directly from Cauchy–Schwarz (`norm_sq_of_symm`), so the real-case module depends on no C^* -algebra API at all. Second, the cut-off quotient is defined as $k_n(x) = f(x)/(\max(c, |x|))^{2n}$ rather than $f(x)/x^{2n}$: the denominator is bounded below by $c^{2n} > 0$ at *every* real x , so continuity needs no case analysis and no division by zero can occur even when 0 lies in the spectrum, while the even power makes the factorisation $k_n(x)x^{2n} = f(x)$ hold identically, with no side condition. Both are small, and both removed what would otherwise have been the fragile parts of the argument.

What verification actually costs. Building the three modules and their closure takes 8160 jobs; the full core root takes 8415. The difference is not the interesting quantity: because the modules `import Mathlib` wholesale, essentially the entire cost in both cases is Mathlib, retrieved from cache. The lane’s own content is three module compilations. We record this because a “minimal verification target” would be misleading here — it would not meaningfully reduce the work a referee must do, and the 8415 figure should be read as evidence of integration into the core, not as a measure of this paper’s size.

9 Reproducibility, verification links, and provenance

Freeze

All statements are machine-checked in Lean 4 (toolchain `leanprover/lean4:v4.29.0-rc6`) against Mathlib pinned to commit `07642720480157414db592fa85b626dafb71355b`. The freeze is the source tree at commit `ad9c93d793fa0b082f2373c2aea3ea5b99f9c6b8`; `lake build`

`YangMillsCore` completes with **8415** jobs. Every declaration reachable from the core root is subject to an axiom oracle; the transcripts below record, for **2304** `#print` axioms invocations, that the axiom sets occurring are exactly `[propext, Classical.choice, Quot.sound]`, `[propext, Quot.sound]` and `[propext]` — all subsets of the three permitted — with zero occurrences of `sorryAx` and no project axioms.

Theorem-to-artifact map

The three modules are [YangMills/OS/TransferGap.lean](#) (real case), [YangMills/OS/DenseClustering.lean](#) (sharp form) and [YangMills/OS/SharpBridge.lean](#) (sharp bridge); the column below names the module by initial.

Statement here	Lean name	Mod.
Def. 2.1 (P, S)	<code>vacuumProjection, projectedTransfer</code>	T
Def. 2.2 (G_v)	<code>connCorr</code>	T
transfer data	<code>VacuumTransfer</code>	T
Lem. 2.3	<code>projected_pow_succ, connCorr_eq</code>	T
Thm. 3.1	<code>clustering_iff_gap</code>	T
$\ S^2\ = \ S\ ^2$, proved by hand	<code>norm_sq_of_symm</code>	T
elementary dense form	<code>gap_of_dense_clustering</code>	T
Def. 4.1	<code>decayDomain</code>	D
Lem. 4.2	<code>decaySubmodule</code>	D
cut-off and factorisation	<code>cutOff, cutOffQuot, cutOffQuot_mul_pow, cutOffQuot_le</code>	D
Thm. 4.3	<code>norm_le_of_dense_decayDomain, norm_le_of_span_dense_decay</code>	D
Thm. 5.1	<code>sharp_clustering_iff_gap</code>	S
Thm. 5.2	<code>volumeUniform_sharp_gap</code> (and <code>volumeUniform_gap</code> , real case)	S, T
Prop. 6.1	<code>exists_vacuumTransfer_gap, euclidean_two_dim_vacuumTransfer_gap</code>	T
Prop. 6.1 (complex)	<code>exists_vacuumTransferC_gap, euclidean_two_dim_vacuumTransferC_gap</code>	S

Oracle transcripts and build root

- Core root: [YangMillsCore.lean](#).
- Oracle driver: [oracle_check.lean](#).
- Transcript at the real-case commit: [ORACLE-20260727-34696985.txt](#).
- Transcript at the sharp-form commit: [ORACLE-20260727-a0ce50ea.txt](#).
- Transcript at the sharp-bridge commit: [ORACLE-20260727-e5d76dae.txt](#).
- Governance and scope record: [docs/0-BRIDGE-CHARTER.md](#) and [docs/0-BRIDGE-AUDIT-20260727.md](#).

The transcripts were produced at the three code commits after which they are named; those commits were replayed over two intervening documentation-only commits before publication, and the replay is recorded together with the byte-level equality of the three module blobs.

Provenance and prior art

The equivalence between exponential clustering and a gap of the transfer operator is *classical* and no novelty is claimed for it: it is the standard argument of constructive field theory [2, 3], and the lattice-gauge transfer operator together with its positivity is due to Osterwalder and Seiler [1].

We make no claim of mathematical novelty for Theorem 4.3 either. Its content — that conceptually, $\mathcal{D}_r(S)$ is the spectral subspace $\text{ran } E([-r, r])$ of a self-adjoint S , so that its density forces $\sigma(S) \subseteq [-r, r]$. *That identification is offered as the conceptual reading and is not separately formalized here*; what the Lean development contains is the norm criterion itself, as the theorem-to-artifact map above records. Criteria of this kind are standard material in local spectral theory [4], where $\limsup_n \|S^n v\|^{1/n}$ is the local spectral radius at v and the sets on which it is bounded are the local spectral subspaces. We have not located the precise statement in the form used here and have not performed an exhaustive search; rather than assert priority we state the expectation that it is known, and we invite the closest antecedent. What we do claim is the mechanization, the packaging of the criterion for *families* of observables with per-member constants, and Corollary 4.4, which is the consequence that matters for cluster expansions and which we have not seen drawn.

The volume-uniform area law whose constant shape motivates Corollary 4.4 is [5], the formal antecedent within the same programme; its role here is motivational only, as stated in Section 4.

References

- [1] K. Osterwalder and E. Seiler, *Gauge field theories on a lattice*, Ann. Physics **110** (1978), 440–471.
- [2] J. Glimm and A. Jaffe, *Quantum Physics: A Functional Integral Point of View*, 2nd ed., Springer, 1987; in particular §19.
- [3] B. Simon, *The $P(\phi)_2$ Euclidean (Quantum) Field Theory*, Princeton University Press, 1974.
- [4] K. B. Laursen and M. M. Neumann, *An Introduction to Local Spectral Theory*, London Mathematical Society Monographs New Series **20**, Clarendon Press / Oxford University Press, 2000.
- [5] L. Eriksson, *A volume-uniform area law for lattice Yang–Mills*, machine-checked, [ai.viXra.org:2607.0005](https://arxiv.org/abs/2607.0005).
- [6] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, CADE 28 (2021), 625–635.
- [7] The mathlib Community, *The Lean mathematical library*, CPP 2020, 367–381.
- [8] A. Jaffe and E. Witten, *Quantum Yang–Mills theory*, Clay Mathematics Institute Millennium Problem description, 2000.