

From the Gibbs Weight to the Spectral Gap: A Complete Machine-Checked Osterwalder–Seiler Chain for the $\mathbb{Z}_2$ Lattice Gauge Chain

The mathematics is textbook; what is verified
is that every interface composes

Lluis Eriksson

Abstract

A mass gap is a statement about the spectrum of an operator, but a lattice gauge theory is given as a measure. The passage between the two — the Osterwalder–Seiler reconstruction — proceeds through reflection positivity, a Gelfand–Naimark–Segal space, a transfer operator, and an *identification* theorem asserting that expectations in the measure are matrix elements of that operator. This paper reports a complete formal verification of that passage, in Lean 4 with `mathlib`, for the $\mathbb{Z}_2$ lattice gauge chain: from the Boltzmann weight $e^{\beta s}$ to exponential decay of a correlation function of the measure, with every arrow a machine-checked theorem and no arrow assumed. The endpoint is the exact identity $\mathbb{E}_n[f(\sigma_0)f(\sigma_n)] = (\tanh \beta)^n$ for the sign observable, a statement in which no operator occurs and whose proof goes through one, and whose rate is exactly the verified non-vacuum transfer eigenvalue $\tanh \beta$; equivalently, the normalised transfer operator has spectral gap $1 - \tanh \beta > 0$, and the corresponding mass is $-\log(\tanh \beta) > 0$ for $\beta > 0$. All three are machine-checked.

We state the limits with the same precision as the results. The system has one $\mathbb{Z}_2$ variable per time slice, so its spatial slice is a point and its transfer operator is a 2×2 matrix; a lattice with spatial extent is not treated. The bound is at fixed finite size and is *not* volume-uniform. The pairing of this system is definite — proved, not assumed — so the Gelfand–Naimark–Segal quotient is the identity and therefore does no work; a system with a degenerate pairing still needs it. The underlying mathematics is textbook: Osterwalder–Seiler is from 1978 and the transfer matrix of the Ising chain is older still. The contribution claimed here is not a new theorem but a verified composition: that the interfaces of the reconstruction fit together with nothing hidden between them, exhibited on the smallest system where all of them are simultaneously present. Nothing in this paper is a claim about $SU(N)$, the continuum limit, or the Yang–Mills mass gap.

1 Introduction

1.1 The gap between a measure and a spectrum

Constructive lattice gauge theory produces a measure. The mass gap is a property of an operator. The bridge between them is the Osterwalder–Seiler reconstruction [1]; the surrounding framework is [2]. Reflection positivity of the measure gives a positive semi-definite pairing on functions of the fields in a half space; quotienting by its null space and completing gives a Hilbert space; translation in the reflected direction descends to a self-adjoint contraction T fixing the vacuum vector Ω ; and the *identification* theorem states that the Euclidean expectation of an observable separated by n time steps from another is the matrix element $\langle A\Omega, T^n B\Omega \rangle$. Only after

that last step does exponential decay of correlations become a statement about the spectrum of T , and only then can “mass gap” mean what it is supposed to mean.

The identification is the step that is easiest to assume and hardest to supply. It is where a formalisation can quietly become circular: if the measure is *defined* through the operator, the identification is a tautology and the chain proves nothing. This paper’s central design constraint — fixed in advance, and verified mechanically rather than by inspection — is that the two sides be defined independently.

1.2 What is proved here

For the $\mathbb{Z}_2$ chain with $n+1$ time slices we define, using exp and the bond sign and nothing else,

$$W_\beta(\sigma) = \prod_t e^{\beta s(\sigma_t, \sigma_{t+1})}, \quad s(i, j) = \begin{cases} +1 & i = j \\ -1 & i \neq j, \end{cases}$$

the partition function $Z_n = \sum_\sigma W_\beta(\sigma)$ and the expectation $\mathbb{E}_n$. Separately, and without reference to any measure, we define the operator

$$T_\beta = a(\beta) \mathbf{1} + b(\beta) S, \quad a = \frac{e^\beta}{e^\beta + e^{-\beta}}, \quad b = \frac{e^{-\beta}}{e^\beta + e^{-\beta}},$$

where S exchanges the two configurations. The results are:

- $Z_n = 2(e^\beta + e^{-\beta})^n$, computed;
- the identification $\mathbb{E}_n[\overline{A(\sigma_0)} B(\sigma_n)] = \langle A\Omega, T_\beta^n B\Omega \rangle$ for every n and every pair of observables;
- the non-vacuum bound $\|T_\beta - |\Omega\rangle\langle\Omega|\| \leq a - b = \tanh \beta$, hence spectral gap $1 - \tanh \beta$, and for $\beta > 0$ mass $-\log(\tanh \beta)$;
- and the endpoint $\mathbb{E}_n[f(\sigma_0)f(\sigma_n)] = (\tanh \beta)^n$.

Verification data for all of the above are collected in Section 9.

Reflection positivity of the same weight was verified separately [5] and the abstract criterion relating a gap to exponential clustering was verified separately [4]; both are consumed here, so the chain is complete end to end within one formal development.

1.3 What is deliberately not claimed

We separate this out because thematic proximity is not mathematical transfer.

1. **The system is the smallest one.** One $\mathbb{Z}_2$ variable per time slice. The spatial slice is a point; T_β is a 2×2 matrix. A lattice with spatial extent has a transfer operator on a space growing with the volume and is not treated here.
2. **The bound is not volume-uniform.** It is a statement at fixed finite size. At this size the question of volume-uniformity does not arise, which is exactly why the result is not evidence about the volume-uniform case.

3. **The Gelfand–Naimark–Segal quotient is absent, not closed.** The pairing of this system is already definite — Remark 6.3 proves it — so the quotient is the identity map. That is an evasion of the difficulty, not a solution of it; a system with a degenerate pairing still requires it.
4. **The mathematics is not new.** The transfer matrix of the Ising chain predates all of this, and the reconstruction is [1]. What is new is that each interface is a machine-checked theorem and that they compose without a gap.
5. **Nothing here concerns $SU(N)$, the continuum limit, or the Clay problem [8].** No reduction from any statement in this paper to the Yang–Mills mass gap is claimed, and none is known to the author.

2 The system and the measure

Definition 2.1 (Configurations and weight). A configuration of the chain with n bonds is a map $\sigma : \{0, \dots, n\} \rightarrow \mathbb{Z}_2$. With $s(i, j) = +1$ if $i = j$ and -1 otherwise, the bond weight is $k_\beta(i, j) = e^{\beta s(i, j)}$ and the path weight is $W_\beta(\sigma) = \prod_{t=0}^{n-1} k_\beta(\sigma_t, \sigma_{t+1})$.

In Lean these are `z2Sign`, `z2Bond` and `z2PathWeight`. The partition function `z2Partition` and the expectation `z2Expect` are built from these.

Remark 2.2 (Non-circularity, verified mechanically). *The definitions above and the two derived from them form a closed dependency set:*

$$\text{z2Expect} \rightarrow \{\text{z2PathSum}, \text{z2Partition}\} \rightarrow \text{z2PathWeight} \rightarrow \text{z2Bond} \rightarrow \{\text{exp}, \text{z2Sign}\} \rightarrow \emptyset.$$

Printing all six definitions and inspecting the printouts shows that no transfer operator, and neither of the coefficients a , b , occurs anywhere on the measure side. Because the closure is closed, this is a complete check rather than a sample. This is the condition that makes the identification of Section 5 a theorem rather than a restatement of a definition.

3 Reflection positivity

Reflection positivity of the $\mathbb{Z}_2$ Wilson weight was established in a companion development [5] and is recalled here only to place the present chain. The route avoids both Bochner’s theorem and the Schur product theorem by formulating the hypothesis as *non-negative combination of characters* rather than as non-negativity of Fourier coefficients; that class is closed under products by indexing over pairs, which is what the reflected pairing requires. For $\mathbb{Z}_2$ at $\beta \geq 0$ the coefficients are $(e^\beta \pm e^{-\beta})/2$, and their non-negativity is monotonicity of $\exp$. The relevant artifacts are `osPairing_nonneg`, `z2Wilson_isPSDKernel` and `z2Wilson_reflectionPositive`, with non-degeneracy at `z2WilsonSystem_nondegenerate`.

Remark 3.1. *For $\mathbb{Z}_N$ with $N > 2$ the corresponding coefficients are discrete Bessel-type sums and their non-negativity is not proved in this development. The chain of this paper is therefore specific to $\mathbb{Z}_2$ and does not extend to $\mathbb{Z}_N$ without that missing input.*

4 The transfer operator, read off the Gibbs weight

Definition 4.1. On $\mathbb{C}^2$ with the Euclidean inner product let S be the exchange of the two configurations and $T_\beta = a(\beta)\mathbf{1} + b(\beta)S$ with a, b as above.

The point of this definition is what it does *not* contain: no eigenvector, no projection, no spectral datum. The coefficients are the normalised Boltzmann weights. In Lean, [z2TransferOp](#).

Theorem 4.2 (The vacuum is fixed, as a consequence). $T_\beta\Omega = \Omega$ for the uniform vector $\Omega = 2^{-1/2}(1, 1)$.

Proof. $T_\beta\Omega = (a + b)\Omega$ and $a + b = 1$ by construction of the normalisation. $\square$

Formally [z2TransferOp_fix](#). We stress that this is a consequence of normalisation rather than a hypothesis; a development in which $T\Omega = \Omega$ is assumed has moved the content elsewhere.

Theorem 4.3 (Non-vacuum spectral bound). For $\beta \geq 0$, $\|T_\beta - |\Omega\rangle\langle\Omega|\| \leq a(\beta) - b(\beta) = \tanh \beta < 1$, where $|\Omega\rangle\langle\Omega|$ denotes the orthogonal projection onto $\mathbb{C}\Omega$.

Proof, as formalised. The projected operator acts componentwise as $(a - b)$ times the projection off the vacuum; this identity is proved from $a + b = 1$ alone. The norm bound then follows from an explicit operator-norm estimate, with no appeal to a spectral theorem. The evaluation $a - b = \tanh \beta$ is proved rather than asserted. $\square$

Artifacts: [z2TransferOp_gap](#), [z2A_sub_z2B_eq_tanh](#), and the packaging [z2TransferOp_vacuumTransfer](#).

Remark 4.4 (Three quantities, kept apart). With the vacuum eigenvalue normalised to 1, three different numbers are in play, and we name them as the literature does.

- $\tanh \beta$ is the non-vacuum eigenvalue: the spectral radius of T_β restricted to the orthogonal complement of the vacuum, equivalently the contraction rate. This is what Theorem 4.3 bounds and what appears as the decay rate in Theorem 6.1.
- $1 - \tanh \beta$ is the spectral gap of the transfer operator, strictly positive for every β : [z2_spectral_gap_pos](#).
- $-\log(\tanh \beta)$ is the mass — the gap of H in $T_\beta = e^{-H}$ — strictly positive for $\beta > 0$: [z2_mass_pos](#).

Earlier drafts of this development, and the identifiers [z2TransferOp_gap](#), [connCorrC_le_of_gap](#) and [sharp_clustering_iff_gap](#), use “gap” for the first of the three. Those names are retained because companion developments already consume them and their permalinks are fixed; an identifier is not a claim, and what each of them states is the non-vacuum bound. In this paper’s prose the three are kept distinct.

5 The identification

Write $A\Omega$ for the vector with entries $A(i)\Omega(i)$.

Theorem 5.1 (Identification). For every $n \in \mathbb{N}$ and all observables $A, B : \mathbb{Z}_2 \rightarrow \mathbb{C}$,

$$\mathbb{E}_n[\overline{A(\sigma_0)} B(\sigma_n)] = \langle A\Omega, T_\beta^n B\Omega \rangle.$$

Proof, as formalised. Split the sum over configurations at the first spin. Writing $\sigma = (s, \tau)$, the first bond factors out of the path weight and the inner sum over s turns the observable A into $A'(i) = \sum_s \overline{A(s)} k_\beta(s, i)$, which is $(e^\beta + e^{-\beta})$ times one application of T_β . Each bond therefore emits one factor of the normalisation and one application of the operator, and closing the induction step requires only the symmetry of the kernel — that is, the self-adjointness of T_β . After n steps,

$$\sum_\sigma \overline{A(\sigma_0)} B(\sigma_n) W_\beta(\sigma) = (e^\beta + e^{-\beta})^n \langle A, T_\beta^n B \rangle,$$

where A, B on the right are read as vectors of $\mathbb{C}^2$. Taking $A = B = 1$ and using $T_\beta \mathbf{1} = \mathbf{1}$ — again $a + b = 1$ — gives $Z_n = 2(e^\beta + e^{-\beta})^n$. The two powers of the normalisation cancel in the quotient, leaving $\frac{1}{2} \langle A, T_\beta^n B \rangle$; and $\langle A\Omega, T_\beta^n B\Omega \rangle = (2^{-1/2})^2 \langle A, T_\beta^n B \rangle$ is the same number. $\square$

The induction step is `z2PathSum_succ`, the partition function `z2Partition_eq`, and the theorem `z2_identification`.

Remark 5.2. *The cancellation of $(e^\beta + e^{-\beta})^n$ between numerator and denominator is not arranged by hand: both sides are computed independently and the factor appears in each. Judge 2 of the pre-registration ([docs/O-BRIDGE-CHARTER.md](#), Amendment 8) required exactly this — that the partition function be computed rather than symbolically cancelled — because a development that divides by an uncomputed Z can hide an error of normalisation that would falsify the identification.*

Remark 5.3 (What in this argument is general, and what is not). *It is worth separating the two, because the next system will reuse one and not the other. The induction above — splitting configurations at the first spin, factoring one bond out of the path weight, and absorbing it into the observable — uses only that the state space is finite and that the bond kernel is symmetric. It does not use that the group has two elements, and it does not use the explicit form of k_β . Everything downstream of it does: the dictionary between observables and vectors, the evaluation $Z_n = 2(e^\beta + e^{-\beta})^n$, the eigenvector computation of the next section, and the operator-norm bound of Theorem 4.3 are all carried out in explicit coordinates on a two-dimensional space. The formalisation reflects this asymmetry in its proofs but not in its types: every declaration is stated at the concrete state space, so the general half of the argument is present but not abstracted. Porting to a larger system therefore means re-typing the induction rather than reproving it, and redoing the rest.*

6 The endpoint

Let $f(0) = +1$, $f(1) = -1$.

Theorem 6.1 (Exact two-point function). *For every n , $\mathbb{E}_n[f(\sigma_0)f(\sigma_n)] = (\tanh \beta)^n$. Moreover $\mathbb{E}_n[f(\sigma_0)] = 0$, so the connected two-point function coincides with the full one.*

Proof. f , read as a vector, satisfies $T_\beta f = (a - b)f$, so $T_\beta^n f = (a - b)^n f$; Theorem 5.1 and $\langle f, f \rangle = 2$ give the value, and $a - b = \tanh \beta$. The one-point function is $\langle f\Omega, \Omega \rangle$, which vanishes because f sums to zero. $\square$

Artifacts: `z2_two_point_tanh`, `z2_measure_clustering` (the same in norm form), and non-vacuity at `z2_two_point_ne_zero`.

This is the statement the whole chain exists to produce. Its left-hand side contains no operator: it is an expectation in a measure defined from $e^{\beta s}$. Its right-hand side is a number that came out of the spectrum of T_β . Its rate is exactly the non-vacuum eigenvalue bounded in Theorem 4.3; the spectral gap and the mass are the derived quantities of Remark 4.4. That the measure-side decay rate equals the non-vacuum transfer eigenvalue, from which the spectral gap and the mass follow, is the content of the reconstruction, and here it is checked rather than asserted.

Remark 6.2 (On what an exact answer does and does not show). *For this system the two-point function can be computed directly, without any of the machinery above; the value $(\tanh \beta)^n$ is classical [3]. The verification reported here is therefore not evidence that the machinery is necessary — at this size it plainly is not. It is evidence that the machinery is correct: an independently constructed operator-theoretic route reproduces the classical answer exactly, which is the check one wants before applying the same route where no direct computation is available.*

Remark 6.3 (What “complete” means here). *The reconstruction quotients by the null space of the pairing, and a reader is entitled to ask whether the operator of this paper is literally the descent of time translation to that quotient or is instead read off the bond kernel and then matched to the measure. It is the latter, and that is deliberate — an independently constructed operator is what makes Theorem 5.1 a theorem rather than a definition. What makes the chain nonetheless complete is that the quotient step is present and is proved to be the identity, not assumed to be. The $n = 0$ instance of Theorem 5.1 is the pairing of this system, and `z2_pairing_nondegenerate` shows its null space is zero; `z2Obs_injective` shows $A \mapsto A\Omega$ is injective, so that map is the reconstruction map. The general- n identification then says that T_β implements the time translation on its image. In this definite finite model the Osterwalder–Seiler quotient is therefore canonically trivial, and the operator obtained from the normalised bond kernel realises the same reconstructed Hilbert-space dynamics. “Complete” means every link of the chain is present and composes; it does not mean that every link was hard.*

7 The abstract criterion, consumed

The companion development [4] proves, for a self-adjoint contraction T with $T\Omega = \Omega$ and $\|\Omega\| = 1$, that a bound $\|T - |\Omega\rangle\langle\Omega|\| \leq r$ is *equivalent* to exponential decay of the connected two-point function at rate r along a family with dense span, with an unconstrained constant per observable. Theorem 4.3 supplies a hypothesis of that criterion from an operator that came from the measure, so the criterion is consumed rather than merely stated: `connCorrC_le_of_gap` and `sharp_clustering_iff_gap`, resting on `norm_le_of_dense_decayDomain`.

Remark 7.1. *Before the present work, every witness of that criterion in the development had the form $r1 + (1 - r)P$ — an operator built from the conclusion. An adversarial audit and an external review both identified this as the criterion’s outstanding weakness. T_β is not of that form; it is read off the Gibbs weight. We record that the objection applied, and that it is answered for this system and for no larger one.*

8 Formalisation architecture

The development is Lean 4 [6] with `mathlib` [7] pinned to an exact commit. Three design decisions are worth recording because they are what made the chain composable.

Separation of the two sides into disjoint definition sets. The measure side and the operator side share no definitions; they meet in exactly one lemma, which states that the normalised bond weight equals the transfer matrix entry. This is what makes Remark 2.2 checkable by printing definitions. A development that shares a “kernel” definition between the two sides cannot make that check.

Observables as plain functions, vectors as a separate type. Observables are $\mathbb{Z}_2 \rightarrow \mathbb{C}$; vectors live in the Euclidean space. A single explicit dictionary maps between them. Identifying the two types would have made the identification theorem read as a triviality.

Induction on the path, not on the operator. The proof of Theorem 5.1 inducts on the number of bonds, splitting configurations with the standard cons equivalence, and never manipulates T^n symbolically. The operator power is assembled by the induction rather than assumed to interact with the sum.

House notes. Two recurrences cost builds and are recorded for re-use. First, goals produced by the self-adjointness/symmetry interface sit over the *linear-map* coercion, so a lemma stated about the continuous linear map does not apply until the coercion is bridged. Second, the two unfolding lemmas for powers act on opposite sides, and pairing an iteration lemma that unfolds on the left with a power lemma that unfolds on the right leaves the two sides of an induction step misaligned in a way whose error message points nowhere near the cause.

9 Reproducibility and verification data

All artifacts are at source checkpoint `bd2c18394ce8a55ec73413c11a9f311bbf002248` on branch `main` of <https://github.com/lluiseiriksson/THE-ERIKSSON-PROGRAMME>. Every link in this paper was resolved against that commit before compilation.

Quantity	Value
Full core build	8422 jobs, success
Oracle commands	2431
with axiom dependencies	2409
axiom-free	22
Distinct axioms across all outputs	<code>propext</code> , <code>Quot.sound</code> , <code>Classical.choice</code>
Occurrences of <code>sorryAx</code>	0
Project axioms	0
Errors	0

The oracle script is [oracle_check.lean](#); the core root is [YangMillsCore.lean](#). The five criteria this work was required to meet — including the non-circularity condition of Remark 2.2 and the computed partition function of Section 5 — were fixed in [docs/O-BRIDGE-CHARTER.md](#), Amendment 8, in a commit preceding the existence of the module they govern; their verdicts are recorded in [docs/VERIFICATION-LEDGER.md](#), Addendum 512.

10 Theorem–artifact map

Statement	Artifact
Bond weight, path weight	z2Bond , z2PathWeight
Expectation in the measure	z2Expect
Transfer operator	z2TransferOp
Theorem 4.2 (vacuum fixed)	z2TransferOp_fix
Transfer data	z2TransferOp_vacuumTransfer
Theorem 4.3 (gap)	z2TransferOp_gap
Rate equals $\tanh \beta$	z2A_sub_z2B_eq_tanh
Spectral gap $1 - \tanh \beta > 0$	z2_spectral_gap_pos
Mass $-\log(\tanh \beta) > 0$	z2_mass_pos
Induction step	z2PathSum_succ
Partition function	z2Partition_eq
Theorem 5.1 (identification)	z2_identification
Theorem 6.1 (endpoint)	z2_two_point_tanh
Endpoint, norm form	z2_measure_clustering
Non-vacuity	z2_two_point_ne_zero
Reconstruction map injective	z2Obs_injective
Quotient is the identity	z2_pairing_nondegenerate
Reflection positivity	z2Wilson_reflectionPositive
Non-degeneracy	z2WilsonSystem_nondegenerate
Positive-definite kernel	z2Wilson_isPSDKernel
Reflected pairing	osPairing_nonneg
Gap $\Rightarrow$ clustering	connCorrC_le_of_gap
Clustering $\Leftrightarrow$ gap	sharp_clustering_iff_gap
Dense-family criterion	norm_le_of_dense_decayDomain

11 Continuation

The next step is not a larger gauge group but a harder pairing. The Gelfand–Naimark–Segal quotient does no work in this paper because the pairing is definite; the first genuine test is a system whose pairing is degenerate, so that the quotient must be constructed and its properties transported through it. After that comes spatial extent — a transfer operator acting on a space that grows with the volume — and only after that does volume-uniformity become a meaningful question. Volume-uniformity is the first point at which any of this could bear on a mass gap, and it is not reached here. We state this ordering because the temptation in this area is to jump to the last item; the intervening steps are where constructions of this kind are expected to fail, and the present one has not yet been asked the question.

References

- [1] K. Osterwalder and E. Seiler, *Gauge field theories on a lattice*, Annals of Physics **110** (1978), 440–471.
- [2] J. Glimm and A. Jaffe, *Quantum Physics: A Functional Integral Point of View*, 2nd ed., Springer, 1987.
- [3] B. Simon, *The Statistical Mechanics of Lattice Gases, Volume I*, Princeton University Press, 1993.
- [4] L. Eriksson, *Clustering and the Transfer-Operator Gap: A Machine-Checked Dense-Family Criterion*, viXra, 2026.
- [5] L. Eriksson, *Reflection Positivity for Abelian Lattice Gauge Theory, Machine-Checked*, viXra, 2026.
- [6] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, CADE-28, Lecture Notes in Computer Science **12699** (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 Prize Problem description, 2000.