

A Mechanized, Non-Circular Renormalization-Group Interface for Wilson Lattice Gauge Correlators

Lluis Eriksson*

July 2026

Abstract

For $SU(N_c)$ Wilson lattice gauge theory on d -dimensional periodic tori ($d \geq 2$), we present a machine-checked (Lean 4, pinned Mathlib) renormalization-group *interface* for two-plaquette truncated correlators, in which every structural ingredient is a theorem rather than a postulate: the scale transformation is a concrete decimation map, defined once — measurable, local, and gauge-covariant — and its induced pushforward preserves probability; the effective measures are its literal iterated pushforwards of the Wilson Gibbs measure; the multiscale decomposition of the correlator is proved by telescoping, never carried as data; the terminal scale of the decomposition is a fixed index $k_T(n) = n$ with typed range $1 \leq k_T(n) \leq n$, which excludes, in the type, the circular depth-zero layer in which an infrared clause would hypothesize the bound being sought; the conditional decay conclusion is stated in the physical distance $2^n u$ with a single constant pair (C, m) quantified before every torus base and depth; and the terminal observable is *operationally support-certified*: the infrared object consumed by the interface equals a base-measure integral of an explicitly composed pullback observable whose dependence is contained in a transported support set, for which the separation lower bound $2^n(2u) - (2^n + 1)$, strictly positive on the whole interface window, is proved. The design is deliberately adversarial: four natural naive formulations are presented together with the explicit countermodels that defeat them — scalar relabeling of known decay, sink flows on measures, clamped scales and per-volume constants, and depth-zero circularity — and with the typed repairs that exclude each. The central hypothesis, **PhysicalTerminalScaleWilsonGate**, is an open proposition: no witness is provided, and the infrared/ultraviolet bounds it demands of the actual Wilson measure are exactly the open analytic mathematics (Balaban-type single-scale estimates). The final theorem is conditional: a witness of the gate yields $|\mathcal{C}(2^n u)| \leq C e^{-m \cdot 2^n u}$ with one pair (C, m) , $m > 0$, for every base $M_0 \geq 4$ and every depth $n \geq 1$. No mass-gap claim, no claim of gate satisfiability, and no thermodynamic or continuum limit is made or implied.

1 The specification problem

The renormalization-group analysis of lattice gauge theories in the tradition of Balaban [1, 2, 3, 4] (see Dimock [5, 6, 7] for a modern exposition, and [9, 10] for the probabilistic state of the art) organizes a correlation function of the fundamental measure into an infrared part, carried by an

*This manuscript and parts of the formal development were produced with AI assistance under the repository’s audit protocol. Every statement labelled *exact* below is machine-checked in Lean 4 against a pinned Mathlib with a clean axiom oracle; countermodel analyses that are not formalized are labelled paper-level where they occur. Repository: <https://github.com/lluiseiriksson/THE-ERIKSSON-PROGRAMME>.

effective theory at a coarse terminal scale, plus a sum of per-scale ultraviolet remainders. This paper concerns a question that is logically prior to any estimate: *what does it mean, in a proof assistant, to state that organization faithfully?*

The question is not vacuous. A formalized “multiscale decomposition” of a correlator function $t \mapsto \text{cov}(t)$ of the shape

$$\text{cov}(t) = \text{cov}_{\text{IR}}(t) + \sum_{k < n_{\text{sc}}(t)} R_{t,k}, \quad (1)$$

with exponential bounds on cov_{IR} and per-scale bounds on the $R_{t,k}$, is *gameable* whenever the decomposition data $(\text{cov}_{\text{IR}}, R, n_{\text{sc}})$ is quantified freely: known information can be relabeled into the required slots with zero renormalization-group content, and the resulting “theorems” transport hypotheses to conclusions without ever touching the physics. A second, independent trap is volumetric: on a fixed finite torus a two-plaquette correlator realizes only finitely many separations, so *every* such function satisfies *some* exponential bound at a large enough constant; interfaces that quantify their constants per volume $(\forall N \exists C(N) \dots)$ are therefore contentless in the direction of the thermodynamic regime, and the quantifier order $\exists C \forall N \dots$ must be enforced *in the type* of the interface.

The contribution of this paper is a mechanized interface for the concrete $\text{SU}(N_c)$ Wilson theory that survives its own adversarial audit. The development method was to state each candidate interface, attack it, record the explicit countermodel when the attack succeeded, and repair the interface *in the type* so that the countermodel becomes unstatable or provably excluded — then iterate. Section 2 presents the four rounds of this process as mathematics: each naive interface, its countermodel with the explicit formula, and the typed repair. Sections 3–6 present the surviving architecture: the concrete decimation and its proved properties (§3), the non-circular terminal gate with the physical-distance exponent (§4), the transported support geometry (§5), and the operational certification composing them (§6). Section 7 states the final conditional theorem; Section 8 names what remains open, plainly; Section 9 records the reproducibility data and the theorem-to-artifact map.

Statuses. As in the companion papers of this repository [16, 17, 18, 19], we use a strict status discipline: **exact** means a Lean 4 theorem in the pinned central repository whose axiom oracle prints exactly `[propext, Classical.choice, Quot.sound]`; **paper-level** means ordinary prose, always flagged. Everything in Sections 3–7 is exact unless flagged; the few paper-level items (one geometric counterexample in §5) support no theorem. The development contains no **sorry** and no project axioms; open questions are carried as named *open propositions* — definitions of **Props** with no witness — and are inventoried in §8 and Table 1.

Honest scope, stated at the outset. All objects are finite-torus objects. The interface’s central hypothesis has *no witness in this development*; providing one for the Wilson data is exactly the open analytic problem (the Balaban-type single-scale activity estimate), and nothing below depends on its satisfiability. The conditional conclusion is a mass-gap-*shaped* decay bound on a finite-torus correlator family; no thermodynamic or continuum limit is constructed, and no claim on the Clay Yang–Mills problem is made (§8). On the formalization side, the closest antecedent we know is the Lean 4 formalization of the *free* scalar field and its Glimm–Jaffe axioms by Douglas, Hoback, Mei and Nissim [11]; the present development is complementary in ambition and scope — it formalizes an *interacting-theory renormalization-group interface* (definitions, identities, geometry, and conditional transport for Wilson lattice gauge theory), not an estimate and not a quantum field theory. Other formalized material in the same repository

includes a volume-uniform Wilson-loop area law via a formalized cluster expansion [13], rooted-tree majorants for polymer expansions with holes [14], and exact two-dimensional SU(2) Yang–Mills computations [15]; the strong-coupling clustering input quoted in §7 belongs to that cluster expansion lane.

2 Naive interfaces and countermodels

Throughout, $\text{cov} : \mathbb{N} \rightarrow \mathbb{R}$ is a correlator function — in the intended instance, a distance-indexed truncated two-plaquette correlator of the Wilson Gibbs measure (§3.1). The four interfaces below are presented in the order of their design; each is retained in the development (with corrected documentation) because the repaired interfaces factor through them, but *none of the first three is load-bearing*: they are the audit trail. All countermodel formulas below are recorded in the development’s documentation; the clamp and depth-zero analyses of §2.3–2.4 are additionally *formalized* as the exact theorems cited there.

2.1 The abstract decomposition, and scalar relabeling

Interface 2.1 (abstract bridge; `CorrelatorBridge`, `RegimeCoherentWilsonBridge`). A *bridge* for cov is a tuple $(\text{cov}_{\text{IR}}, R, n_{\text{sc}})$ satisfying the identity (1) for all t . The naive gate asks: there exists a bridge, a marginal-coupling flow g (Definition 4.1 below), and constants such that $|\text{cov}_{\text{IR}}(t)| \leq C_1 e^{-\varepsilon t}$, each $|R_{t,k}| \leq C_2 e^{-c_0 t} g_k^{\kappa_0}$, and $R \neq 0$.

The transport theorem for this interface (`physical_mass_gap_of_bridge`) is sound: any witness yields $|\text{cov}(t)| \leq (C_1 + C_2 \sum_k g_k^{\kappa_0}) e^{-\min(\varepsilon, c_0) t}$. The interface is nevertheless worthless as a statement of RG content:

Countermodel 2.2 (scalar relabeling). Suppose some decay $|\text{cov}(t)| \leq a e^{-t}$ is already known (for the Wilson theory at strong coupling, such a bound is a theorem — see §7). Set

$$R_{t,0} := a e^{-t}, \quad R_{t,k} := 0 \ (k \geq 1), \quad \text{cov}_{\text{IR}}(t) := \text{cov}(t) - a e^{-t}, \quad n_{\text{sc}} := 1.$$

Every clause of Interface 2.1 is satisfied — the decomposition identity trivially, the IR bound by the triangle inequality, the UV bound after absorbing $g_0^{-\kappa_0}$ into C_2 , and the nonvanishing clause by $R_{t,0} \neq 0$ — with *zero* renormalization-group content. The clause “ $R \neq 0$ ” certifies nonvanishing, not ultraviolet provenance.

Remark 2.3 (typed repair: remainders must be produced, not selected). *The repair removes the freedom to choose R : introduce a family of measures $\mu_0, \mu_1, \mu_2, \dots$ on gauge configurations with μ_0 the Wilson Gibbs measure, define the effective correlator cov_k at scale k as the correlator of μ_k — a definition, not data — and produce*

$$R_{t,k} := \text{cov}_k(t) - \text{cov}_{k+1}(t), \quad \text{cov}_{\text{IR}}(t) := \text{cov}_{n_{\text{sc}}(t)}(t).$$

The decomposition (1) is then a theorem (telescoping; `provenant_decomposition`), not a field, and a nonzero remainder certifies that the measure genuinely changed between consecutive scales (`step_moves_of_provenantRsc_ne_zero`). Artifacts: `RGFlowFamily`, `covEff`, `provenantRsc`, `provenantBridge`.

2.2 Measure-level flows, and the sink flow

Remark 2.3 still quantifies over the *generating map* of the family: an `RGFlowFamily` carries a field $T : \text{Measures} \rightarrow \text{Measures}$ with $\mu_{k+1} = T(\mu_k)$. That field is the next attack surface.

Countermodel 2.4 (sink flow). Take $T(\nu) := 0$ for every ν (or a fixed Dirac measure δ_{A_0} , whose connected correlators vanish, so that probability preservation alone would not repair the interface). Then $\text{cov}_0 = \text{cov}$ and $\text{cov}_k = 0$ for $k \geq 1$: the *entire* physical correlator is relabeled as the first-scale ultraviolet remainder $R_{t,0} = \text{cov}(t)$, its known strong-coupling bound converts into the demanded single-scale shape by adjusting C_2 by $g_0^{-\kappa_0}$, and the terminal infrared part is identically zero. This is Countermodel 2.2 one level up: an arbitrary total map of measures can encode *any* prescribed orbit.

Remark 2.5 (typed repair: no map variable). *The repair removes the map variable altogether. The scale transformation is a fixed, defined decimation B_M acting on gauge configurations (Definition 3.1), and the effective measure is its literal pushforward, $\mathcal{E}_M \mu := (B_M)_* \mu$ — by construction, not as data. Nothing in the surviving interfaces quantifies over transformations, so the sink flow is unstatable; and probability preservation is a theorem (`effectiveMeasure_isProbability`), so no stage of any tower can be the zero measure (`effectiveMeasure_ne_zero`). The fixed Dirac is equally unreachable: the only measures occurring in the surviving interfaces are iterated pushforwards of the Wilson Gibbs measure itself. If the honestly blocked measures happened to have vanishing correlator differences, the nonvanishing data clause of the gate would simply be false — the attack can prevent a witness, never fabricate one.*

2.3 Clamped scales and per-volume constants

The decimation halves the torus, so a multiscale tower must be indexed carefully: starting from size $2^n M_0$, after n steps it reaches the base size M_0 and can descend no further. The development makes the tower a *total* function by clamping — beyond depth n the tower stays at its base stage (§3.3). Totality is a formalization convenience; unguarded, it is an attack surface, and it interacts with the volumetric trap of §1.

Countermodel 2.6 (clamp and per-volume constants). (i) In a gate whose scale count n_{sc} is free data, a witness may point the “terminal” infrared object at a clamped index ($n_{\text{sc}}(t) > n$, where the tower is stationary) or at the root ($n_{\text{sc}}(t) = 0$, where no blocking has occurred and the “infrared part” *is* the physical correlator): relabeling reappears through the index arithmetic. Similarly, a depth-0 tower ($n = 0$) produces no remainders at all, so “nontriviality” clauses quantified over all depths are vacuously threatened at $n = 0$. (ii) Independently, on the fixed torus of size N the canonical correlator (§3.1) is supported on finitely many separations, so for *each* N some constant $C(N)$ makes any exponential bound true; a gate of quantifier shape $\forall N \exists C(N)$ is therefore satisfiable without content.

Remark 2.7 (typed repair: clamp theorems and quantifier order). *Against (i), the clamp is defused by theorems: clamped scales produce zero remainder (`concreteRsc_eq_zero_of_le`, `scaledCanonicalRsc_eq_zero_of_le`), so a nonzero produced remainder forces $k < n$ — ultraviolet content certifies an actual decimation step at an actual genuine scale (`lt_of_concreteRsc_ne_zero`, `lt_of_scaledCanonicalRsc_ne_zero`); nonvanishing clauses are restricted to $n \geq 1$; and the terminal index is ultimately pinned by definition with a typed*

range (§4), so the infrared clause cannot reach the clamp region or the root at all. Against (ii), every surviving gate has the quantifier shape $\exists(\text{constants}) \forall(\text{base } M_0) \forall(\text{depth } n) \dots$ — one constant pack before the unbounded volume family (the tower sizes are strictly increasing and cofinal in all sizes: `towerSize_strictMono`, `towerSize_unbounded`); the forbidden shape is unstable against the gate’s type, and the consumer’s conclusion (Theorem 7.1) carries the same order.

2.4 Depth-zero circularity

One further failure mode survives all of the above, and it is the subtlest: a gate can be non-gameable and still contain its intended conclusion *as a hypothesis*, if its clause family includes a layer at which the tower performs no blocking.

Countermodel 2.8 (the $n = 0$ shell). Consider the “literal” gate that demands the scale-corrected bounds at every base $M_0 \geq 4$ and every depth $n \geq 0$ (`LiteralFidelityConcreteRGWilsonGate`). Every torus size $N \geq 4$ occurs in this family as a base at depth zero ($M_0 := N$, $n := 0$), and at depth zero the scaled infrared clause reads: *the physical correlator itself decays exponentially at every separation in the window, uniformly in N* — which is the statement the whole architecture is trying to reach. Consequently the implication “literal gate $\Rightarrow$ decay at all separations and all sizes” (`allSeparations_of_literalFidelityGate`, a theorem) consumes *only* the depth-zero shell and no renormalization-group content whatsoever. The implication is valid; as a reduction of the open problem it is circular. Two exact factorization facts sharpen the diagnosis: an odd separation t lies on the depth- n dyadic separation lattice $\{2^n u\}$ only at $n = 0$ (`odd_dyadic_factorization_trivial`, `odd_separation_needs_base_scale`), and no separation $t < 2^n$ is reachable at depth n at all (`dyadic_lattice_misses_small_separations`); so a gate that controls all separations *must* lean on its depth-zero layer — there is no honest way to reach odd t from positive scales without new interpolation input (§8).

Remark 2.9 (typed repair: the fixed terminal index). *The repair excludes the depth-zero layer in the type. The terminal gate of §4 quantifies only depths $n \geq 1$; its infrared clause is evaluated syntactically at the fixed terminal index $k_T(n) := n$, with the typed range $1 \leq k_T(n) \leq n$ (`kTerm_pos`, `kTerm_le` — both definitional), so the consumed infrared object is the correlator of the maximally blocked measure, never of the unblocked one; and the index $k = 0$ enters the gate only inside the first difference $R_0(u) = \text{cov}_0(u) - \text{cov}_1(u)$ (`rsc_is_difference`, by `rfl`) — a single-scale remainder of the honest ultraviolet shape, not a standalone decay hypothesis on the physical object. The bound on the physical correlator is then derived through the telescoping, not hypothesized (Theorem 7.1).*

3 The concrete gauge renormalization group

3.1 Setting

Fix a dimension $d \geq 2$. For $N \geq 1$ let $\Lambda_N = (\mathbb{Z}/N\mathbb{Z})^d$ be the periodic torus, with oriented edges $(y, i, \pm)$ ($y \in \Lambda_N$, $i \in \{0, \dots, d-1\}$). For a group G (measurable; the physical instance is the compact $\text{SU}(N_c)$ with its Borel structure), a *gauge configuration* is a G -valued function A on oriented edges obeying the reversal constraint $A(e^{-1}) = A(e)^{-1}$; the configuration space is $\mathcal{A}_N = \mathcal{A}_N(G)$, with the product measurable structure of the positive-edge coordinates. A

plaquette $p = (y, i, j)$, $i < j$, has the usual four-edge boundary and holonomy $\text{hol}(A, p) \in G$; gauge transformations $u : \Lambda_N \rightarrow G$ act by $(u \cdot A)(y, i, +) = u(y) A(y, i, +) u(y + e_i)^{-1}$.

The *Wilson Gibbs measure* at coupling $\beta \in \mathbb{R}$ is the tilted product Haar measure

$$d\mu_{N,\beta}^W(A) = Z_{N,\beta}^{-1} e^{-\beta \sum_p E(\text{hol}(A,p))} d\text{Haar}^\otimes(A), \quad E(U) = \text{Re tr}(\mathbf{1}_{N_c} - U), \quad (2)$$

realized as an honest `Measure.tilted (wilsonGibbsMeasure)`; it is a probability measure for every torus size and every real β (`wilsonGibbsMeasure_isProbability`), and its correlators agree exactly, with no hypotheses, with the Z -quotient Boltzmann expressions (`integral_gibbsMeasure_eq_div, rayCorrelatorOfMeasure_wilsonGibbs`).

For a measure ν on $\mathcal{A}_N$ and an observable $f : G \rightarrow \mathbb{R}$, the *truncated two-plaquette correlator* is

$$T_\nu(p, q) = \int f(\text{hol}(A, p)) f(\text{hol}(A, q)) d\nu - \int f(\text{hol}(A, p)) d\nu \int f(\text{hol}(A, q)) d\nu.$$

Plaquettes carry the *touch graph*: $p \sim q$ iff $p \neq q$ and their supports share an edge; dist_N denotes its graph distance. The *canonical plaquette family* kills all selection freedom: for $t \in \mathbb{N}$, P_t^N is the axis-aligned plaquette in the (e_0, e_1) -plane based at the site $t \cdot e_0$ (torus wrap by $t \bmod N$) — a closed-form geometric definition (`canonicalPlaquette`), with exact linear metric growth

$$\text{dist}_N(P_0^N, P_t^N) = t \quad (2t \leq N) \quad (3)$$

(`canonicalPlaquette_dist_eq`; the lower bound is proved by a $\mathbb{Z}/N$ -valued potential argument, wrap-safe, the upper bound by an explicit axis walk). The *canonical correlator* at separation index t is

$$C_\nu(t) := T_\nu(P_0^N, P_{2t}^N),$$

so that in the wrap-free window $1 \leq t$, $4t \leq N$ the pair is distinct at touch-distance exactly $2t$ (`canonicalPair_admissible`). (The development also carries a choice-based “ray” correlator and an all-pairs supremum correlator `supAbsCorrelator`, with domination lemmas connecting the three; the canonical family is the load-bearing one.)

3.2 The decimation

Definition 3.1 (the block map; `blockMap`). For $M \geq 1$, the decimation $B_M : \mathcal{A}_{2M}(G) \rightarrow \mathcal{A}_M(G)$ sends a fine configuration A to the coarse configuration whose value on the positive coarse edge $(y, i, +)$ is the ordered product of the two fine holonomies along its doubled image:

$$(B_M A)(y, i, +) = A(2y, i, +) \cdot A(2y + e_i, i, +),$$

negative edges receiving the inverse through the reversal constraint. Here $y \mapsto 2y$ is the even-sublattice embedding (`coarseSiteEmbed`).

B_M is a fixed definition — there is no transformation variable anywhere downstream (Remark 2.5). Four properties are proved about it as theorems:

Proposition 3.2 (properties of the decimation). (i) *Measurability: the map B_M is measurable when multiplication on G is jointly measurable* (`measurable_blockMap`).

(ii) *Locality (`blockMap_local`): the coarse value at an edge depends only on the fine configuration at its two constituent fine edges — finite range, with the explicit two-edge carrier named in the hypotheses.*

(iii) Gauge covariance (*blockMap_gaugeAct*): $B_M(u \cdot A) = (u \circ \text{embed}) \cdot B_M(A)$ — *blocking a gauge-transformed configuration equals gauge-transforming the blocked configuration by the restriction of u to the even sublattice (the interior vertex of each doubled edge cancels).*

(iv) Scale geometry (*coarseSiteEmbed_shift*): *one coarse step equals two fine steps, torus wrap included.*

Definition 3.3 (effective measure; *effectiveMeasure*). $\mathcal{E}_M \mu := (B_M)_* \mu$, the pushforward. It preserves probability (*effectiveMeasure_isProbability*) and in particular is never the zero measure (*effectiveMeasure_ne_zero*).

Observable transport is proved in exactly the correlator shape the interface consumes: the coarse integral, truncated correlator, and distance-indexed correlator under $\mathcal{E}_M \mu$ equal the fine integrals of the pulled-back observable $A \mapsto f(\text{hol}(B_M A, p))$ under μ (*integral_effectiveMeasure*, *truncatedCorrelator_effectiveMeasure*, *rayCorrelator_effectiveMeasure*), via the change-of-variables identity for pushforwards with the measurability side conditions discharged from Proposition 3.2(i).

3.3 The tower and the scale-corrected telescoping

Fix a base size $M_0 \geq 1$ and define tower sizes by the doubling recursion $\text{sz}(0) = M_0$, $\text{sz}(m+1) = 2 \text{sz}(m)$, so $\text{sz}(m) = 2^m M_0$ (*towerSize*, *towerSize_eq_pow_mul*); the recursion (rather than the closed form) makes the dependently-indexed tower typecheck without casts. For a start measure μ on $\mathcal{A}_{2^n M_0}$, the *tower measure* $\mu^{(k)}$ is the k -fold blocked measure, living on $\mathcal{A}_{2^{n-k} M_0}$, defined by iterating $\mathcal{E}$; at the base scale the tower *clamps* (stays put) — stated openly and defused by the clamp theorems of Remark 2.7. Every stage of a Wilson-anchored tower is a typed probability measure (*towerMeasure_isProbability*, *wilsonTowerMeasure_isProbability*).

Stage k of the tower lives on a lattice whose spacing is 2^k in finest units. The *scale-corrected canonical correlators* therefore evaluate stage k at separation index $2^{n-k} u$:

$$c_k(u) := \mathcal{C}_{\mu^{(k)}}(2^{n-k} u), \quad R_k(u) := c_k(u) - c_{k+1}(u), \quad (4)$$

(*scaledCanonicalCovEff*, *scaledCanonicalRsc*), so that every stage sees the *same* physical separation index $2^k \cdot (2^{n-k} u) = 2^n u$, and the separation-to-volume ratio u/M_0 is scale-invariant: the single window

$$1 \leq u, \quad 4u \leq M_0 \quad (5)$$

serves the whole tower, with the canonical pair realized at touch-distance exactly twice the stage separation index at every stage (*scaledCanonical_pair_dist*). The multiscale decomposition is then a theorem:

Proposition 3.4 (telescoping; *scaledCanonical_decomposition*). *For every scale count n_{sc} and every u ,*

$$c_0(u) = c_{n_{\text{sc}}(u)}(u) + \sum_{k < n_{\text{sc}}(u)} R_k(u),$$

*and $c_0(u) = \mathcal{C}_\mu(2^n u)$, the canonical correlator of the base measure at the physical separation index (*scaledCovEff_zero_physical*, *definitional*).*

Nothing in (4) is selectable: the only inputs are the Wilson data (μ, f) and the tower geometry (M_0, n) .

4 The non-circular terminal gate

Definition 4.1 (marginal-coupling flow). A *marginal flow* is a sequence $g : \mathbb{N} \rightarrow (0, \infty)$ together with $\beta_c > 0$ such that $\beta_c g_k < 1$ and $g_{k+1} = g_k(1 - \beta_c g_k)$ for all k — the discrete logistic recursion mimicking an asymptotically free marginal coupling. For every $\kappa_0 > 1$ the sum $\sum_k g_k^{\kappa_0}$ converges; this is a theorem of the development (`marginal_coupling_pow_summable_of_recursion`), and it is what turns per-scale bounds into a summed constant.

Definition 4.2 (terminal index; `kTerm`). $k_T(n) := n$. The range facts $1 \leq k_T(n)$ (for $n \geq 1$) and $k_T(n) \leq n$ are definitional (`kTerm_pos`, `kTerm_le`): the terminal stage is the maximally blocked measure, the clamp region is unreachable, and the root $k = 0$ is not the infrared object (Remark 2.9).

Definition 4.3 (the physical terminal gate; `PhysicalTerminalScaleWilsonGate`). Fix $d \geq 2$, $N_c \geq 1$, an observable $f : \text{SU}(N_c) \rightarrow \mathbb{R}$, and $\beta \in \mathbb{R}$. The gate holds iff there exist a marginal flow (g, β_c) and constants $C_1 \geq 0$, $C_2 \geq 0$, $\varepsilon > 0$, $c_0 > 0$, $\kappa_0 > 1$ — *all quantified before every base and depth* — such that for every base $M_0 \geq 4$ and every depth $n \geq 1$, writing $\mu = \mu_{2^n M_0, \beta}^W$ and c_k, R_k for the scale-corrected canonical tower correlators (4) of the Wilson-anchored tower:

(P) every tower stage is a probability measure — a clause that is unconditionally provable (`fidelityGate_probability_clause`) and is carried typed so that every admitted instance is a probability tower by its statement;

(IR) for every u in the window (5),

$$|c_{k_T(n)}(u)| \leq C_1 e^{-\varepsilon \cdot 2^n u},$$

where the left-hand side is *syntactically* the terminal object at the index $k_T(n)$ (`terminalGate_IR_index_eq`, by `rfl`: same term, no bridging freedom);

(UV) for every such u and every $k < k_T(n)$,

$$|R_k(u)| \leq C_2 e^{-c_0 \cdot 2^n u} g_k^{\kappa_0};$$

(NZ) some $R_k(u) \neq 0$ with u in the window and $k < k_T(n)$ — a data clause for the given Wilson data, not engineered into the class (a genuine RG may have fixed measures and invariant observables; for constant f all correlators vanish and (NZ) is simply false).

Three design points, each enforced in the type. First, the *decay variable is physical*: the exponents evaluate at $2^n u$ — the cast $((2^n u : \mathbb{N}) : \mathbb{R})$ appears literally in the Lean `Prop` — so a witness must provide a gap that does not shrink with the depth; the variant gate with renormalized exponents $e^{-\varepsilon u}$ (`TerminalScaleWilsonGate`) is strictly weaker, and the implication `physical` $\Rightarrow$ `renormalized` is a theorem (`terminalScaleGate_of_physicalTerminalGate`) whose converse is false pointwise and is neither stated nor consumed. Second, the quantifier order is the one demanded by Countermodel 2.6(ii). Third, the window (5) is wrap-free and depth-independent: it is a condition on (u, M_0) only (`terminal_window_nonempty`), so it does not shrink as n grows.

No witness. No witness of Definition 4.3 is provided anywhere in this development, and nothing below depends on its satisfiability. Producing the (IR)/(UV) clauses for the actual Wilson measure is the open analytic problem — in substance, a Balaban-type single-scale estimate for the blocked Wilson theory [1, 2, 3, 4, 5]; the gate is the typed statement of *what exactly* must be proved, in a form that Sections 2 and 6 certify cannot be discharged by relabeling.

5 Support geometry

The telescoping of Proposition 3.4 is honest algebra, but a scale decomposition of *one physical separation* requires metric control: the selected plaquette pairs at consecutive scales must be the same geometric object at matched offsets, and the transported observables must have supports that genuinely separate. This section records the proved geometry.

5.1 One step: embedding and the exact canonical doubling

The plaquette embedding $\iota : p \mapsto (\text{site } 2y, \text{ same plane})$ (`plaqEmbed`) is injective. The naive expectation $\text{dist}_{2M}(\iota p, \iota q) \leq 2\text{dist}_M(p, q)$ is *false*: in $d = 3$ there are coarse plaquettes at touch-distance 1 (sharing one edge, in transverse planes) whose embedded images sit at fine touch-distance exactly 3 — a single fine plaquette cannot meet both embedded supports, while an explicit length-3 walk through the doubled shared edge exists. (This counterexample is a paper-level analysis recorded in the development’s documentation; no theorem depends on it — it explains why the general theorem has the constant it has.) What is proved:

Proposition 5.1 (metric transport, one step). *(i) For touch-reachable pairs, $\text{dist}_{2M}(\iota p, \iota q) \leq 3\text{dist}_M(p, q)$ (`plaqEmbed_dist_le_three_mul`), with walks transported constructively. No sharpness claim is formalized. The converse direction (coarse distance dominated by embedded fine distance) is a named open proposition, unproved and consumed nowhere (`PlaqEmbedDistLowerBound`).*

(ii) On the canonical family the scale relation is exact with the physical factor: $\iota(P_\tau^M) = P_{2\tau}^{2M}$ (`plaqEmbed_canonicalPlaquette`) and, for $2\tau \leq M$,

$$\text{dist}_{2M}(\iota P_0^M, \iota P_\tau^M) = 2\text{dist}_M(P_0^M, P_\tau^M)$$

(`canonicalPlaquette_dist_doubles`). The interface consumes only this family.

5.2 One step: the support of the pullback observable

The pullback observable $A \mapsto f(\text{hol}(B_M A, p))$ reads the fine configuration only on the eight fine edges of the 2×2 fine loop over p — the set $\text{supp}_B(p)$ (`blockPlaquetteSupport`, $|\text{supp}_B(p)| \leq 8$), with the locality statement proved in the observable shape the correlators consume (`blockMap_holonomy_congr`, `blockObservable_congr`). Every support edge’s source lies in the 2×2 plane box at the doubled anchor (`blockPlaquetteSupport_source_in_ball`), and for the canonical pair at coarse offset τ in the window $1 \leq \tau, 4\tau \leq M$, *any* fine plaquettes whose supports meet the two transported supports are separated by at least $2\tau - 3$ touch-steps (`blockSupport_canonical_separation_walk`, `blockSupport_canonical_separation_dist`) — the slack 3 being the radius-2 support extent plus one carrier-internal offset, proved by the same wrap-safe circular-potential argument as (3).

5.3 k steps: radius, separation, positivity

Iterating, the k -step transported support of a plaquette (`kStepBlockSupport`) is controlled by a master induction:

Theorem 5.2 (*k*-step radius; `kStepBlockSupport_source_in_ball`). *Every edge of the k -step transported support of a plaquette P keeps its direction in the plane $\{\text{dir}_1 P, \text{dir}_2 P\}$, and its source lies in the plane box at the 2^k -scaled anchor (`iterEmbed_coord`: the anchor coordinates scale exactly, no wrap) with direction-correlated offsets $a + [\text{dir} = \text{dir}_1] \leq 2^k$, $b + [\text{dir} = \text{dir}_2] \leq 2^k$. In particular the support radius is at most 2^k , and one step of the recursion is definitionally the one-step object (`kStepBlockSupport_one`).*

Theorem 5.3 (*k*-step separation; `kStepSupport_canonical_separation_walk`, `kStepSupport_canonical_separation_dist`). *Let $1 \leq \tau$ and $2\tau \leq M$ with M a tower size. Any fine plaquettes on the size- $2^k M$ torus whose supports meet the two k -step transported supports of the canonical pair (P_0^M, P_τ^M) are separated by at least*

$$2^k \tau - (2^k + 1)$$

touch-steps, on every touch-walk (hence in graph distance when reachable). At $k = 1$ this is the one-step constant $2\tau - 3$ exactly (`kStepSeparation_one`). On the window $\tau \geq 2$, $k \geq 1$ the bound is strictly positive (`kStepSeparation_pos`) and strictly increasing in k (`kStepSeparation_strictMono`); at $\tau = 1$ it $\mathbb{N}$ -truncates to 0 (true and uninformative, stated), and neither positivity nor sharpness is claimed outside the stated window.

5.4 The terminal instance

The gate of Definition 4.3 consumes the terminal stage $k = k_T(n) = n$, whose infrared object is — definitionally, by `rfl` — the truncated correlator of the n -fold blocked measure at the canonical pair $(P_0^{M_0}, P_{2u}^{M_0})$ on the *base* torus (`terminalIR_is_canonical_pair_correlator`, with `terminal_stage_size` and `terminal_stage_sep_index`). Instantiating Theorem 5.3 at exactly that pair and exactly $k = n$ (coarse offset $\tau = 2u \geq 2$) gives: in the gate's own window (5), the n -step transported supports of the terminal pair separate every pair of touching carriers by at least

$$2^n(2u) - (2^n + 1) > 0 \quad (n \geq 1, u \geq 1), \quad (6)$$

with strict growth in the depth (`terminalSupport_canonical_separation_walk`, `terminalSupport_separation_pos`, `terminalSupport_separation_strict_growth`). The window is a condition on (u, M_0) only — no n appears in the hypotheses.

6 Operational certification

Section 5 is geometry of *sets*; Definition 4.3 consumes *correlators of measures*. The remaining obligation — and the last piece of the architecture — is the operational connection: that the certified geometry is the geometry *of the observable the gate consumes*. This is supplied by one composed map and three theorems about it.

Definition 6.1 (the composed block map; `iteratedBlockMap`). $B_{M_0, n}^{(k)} : \mathcal{A}_{2^n M_0} \rightarrow \mathcal{A}_{2^{n-k} M_0}$ is the k -fold composition of the tower's configuration-level steps, by the *same* recursion as the tower measure (one definition; every theorem below consumes this term). One step on a genuine scale is definitionally B_M (`iteratedBlockMap_one`); the map is measurable (`measurable_iteratedBlockMap`).

Theorem 6.2 (measure identity; `towerMeasure_eq_map_iterated`). *Every stage of the tower is the literal pushforward of the base measure under the composed map: $\mu^{(k)} = (B_{M_0, n}^{(k)})_* \mu$.*

Theorem 6.3 (k -fold pullback locality; `kStepBlockObservable_congr`). *If two fine configurations agree, in positive-edge coordinates, on the k -step pullback support of a stage- k plaquette P (`kStepPullbackSupport` — the support transport matched step-by-step to the composed map’s own recursion), then the composed pullback observables agree: $f(\text{hol}(B^{(k)} A, P)) = f(\text{hol}(B^{(k)} B, P))$. At $k = 1$ this recovers the proved one-step locality on the eight-edge support, and the pullback support object itself agrees with the `kStepBlockSupport` geometry of §5 in every statement, through a bridge theorem quantified over all predicate families — no cast, no `HEq`, no hypothesis (`kStepPullbackSupport_succ_fine`, `kStepPullbackSupport_terminal_bridge`); the separation bound (6) transfers verbatim onto the operational supports (`terminalPullbackSupport_separation_walk`, `terminalPullbackSupport_separation_dist`).*

Theorem 6.4 (three-term integration identity; `truncatedCorrelator_towerMeasure`, `terminalCorrelator_eq_integral_pullback`). *For measurable f , the truncated correlator of the stage- k tower measure equals the truncated-correlator combination of base-measure integrals of the composed-pullback observables — the product integral and both single-observable integrals transported through the pushforward:*

$$T_{\mu^{(k)}}(p, q) = \int f_p^{(k)} f_q^{(k)} d\mu - \int f_p^{(k)} d\mu \int f_q^{(k)} d\mu, \quad f_p^{(k)}(A) := f(\text{hol}(B^{(k)} A, p)),$$

and, instantiated at the gate’s exact objects, the terminal infrared correlator $c_{kT(n)}(u)$ equals the three-term base-measure integral expression at the exact canonical pair.

Theorem 6.5 (operational support certification; `terminal_support_certified`). *Fix $d \geq 2$, $N_c \geq 1$, a measurable f , $\beta \in \mathbb{R}$, a base M_0 , a depth $n \geq 1$, and u in the window (5). Let $\mu = \mu_{2^n M_0, \beta}^W$, let S_0, S_u be the n -step pullback supports of the terminal canonical pair, and let $f_0^{(n)}, f_u^{(n)}$ be the composed-pullback observables. Then, jointly:*

- (1) *the terminal infrared object consumed by Definition 4.3 equals the base-Wilson-measure truncated correlator of $f_0^{(n)}$ and $f_u^{(n)}$ (all three integrals transported);*
- (2, 3) *each of $f_0^{(n)}, f_u^{(n)}$ depends only on the base configuration’s positive-edge values on S_0 , resp. S_u ;*
- (4) *every touch-walk between fine plaquettes whose supports meet S_0 and S_u respectively has length at least $2^n(2u) - (2^n + 1)$;*
- (5) *that bound is strictly positive on the whole gate window.*

Theorem 6.5 is the sense in which the interface is *operational*: the infrared clause of the gate is a bound on a quantity that is *proved* to be a correlation of two explicitly supported, provably separated functionals of the base Wilson field — not on an opaque number attached to a blocked measure. It certifies support and metric, not decay: no decay bound and no gate witness is claimed.

7 The conditional theorem

Theorem 7.1 (conditional physical-distance decay). *Fix $d \geq 2$, $N_c \geq 1$, $f : \text{SU}(N_c) \rightarrow \mathbb{R}$, and $\beta \in \mathbb{R}$. If `PhysicalTerminalScaleWilsonGate` holds for these data (Definition 4.3), then there exist*

$$C = C_1 + C_2 \sum_{k=0}^{\infty} g_k^{\kappa_0} < \infty, \quad m = \min(\varepsilon, c_0) > 0,$$

— one pair, quantified before every base and depth — *such that for every base $M_0 \geq 4$, every depth $n \geq 1$, and every u with $1 \leq u$ and $4u \leq M_0$,*

$$\left| C_{\mu_{2^n M_0, \beta}^W}(2^n u) \right| \leq C e^{-m \cdot 2^n u}.$$

The proof chain is visible in the formal artifact: the physical correlator is definitionally the scale-zero tower object (Proposition 3.4); the telescoping ends at the same fixed index $k_T(n)$ the infrared clause consumes (Remark 2.9); the ultraviolet sum is dominated by $C_2 e^{-c_0 2^n u} \sum_k g_k^{\kappa_0}$ using the flow summability (Definition 4.1); and no clause at $k = 0$ is consumed anywhere — the bound on the physical object is derived, not hypothesized. By Theorem 6.5, the infrared quantity the hypothesis constrains is itself certified as a correlation of two explicitly supported observables of the base Wilson field at provably positive, depth-increasing support separation (6).

Remark 7.2 (exactly what is claimed). *The Lean artifact of Theorem 7.1 is `wilson_physical_mass_gap_of_physicalTerminalGate`. Theorem 7.1 is a conditional statement about a family of finite-torus correlators. It is not asserted that the gate has a witness; it is not asserted that the Wilson measure satisfies the (IR)/(UV) clauses at any coupling; the theorem controls the separations $2^n u$ within the stated windows and no others; and no thermodynamic or continuum limit is constructed. The value of the theorem is architectural: it reduces a physical-distance, volume-family-uniform decay conclusion to a typed, non-circular, relabeling-proof, operationally support-certified pair of bound families — and Section 2 documents why each of those adjectives costs something.*

Remark 7.3 (non-vacuity of the transport machinery). *The consumer machinery itself is not vacuous, and this is audited at the physical measure: in the proved strong-coupling window of the repository's cluster-expansion lane (the clustering theorem `sun_lattice_exponential_clustering`; cf. [13, 8]), the abstract bridge pipeline of §2.1 fires end to end for the actual $\text{SU}(N_c)$ Wilson distance-indexed correlator, through the trivial decomposition, and yields its exponential decay (`wilson_correlator_mass_gap_strong_coupling`). As the module states on its face, that witness carries no ultraviolet content — the strong-coupling decay was already available directly; it demonstrates only that the transport theorems are inhabited and fire at the physical measure. No analogous witness exists for the terminal gate of Definition 4.3, whose strong-coupling discharge is blocked precisely by the volume-uniformity and nonvanishing clauses — the per-size strong-coupling constants do not assemble into one pack, and no claim that they do is made.*

8 The open frontier

Everything in this section is open, named, and carried in the development as definitions without witnesses. Table 1 lists the named open propositions.

1. **The gate witness.** No witness of `PhysicalTerminalScaleWilsonGate` (or of any of the retained weaker gates) is provided. The proposition could in principle be unsatisfiable; nothing in this paper argues either way.
2. **The physical Wilson bounds.** Discharging the (IR) and (UV) clauses for the actual Wilson measure is the open analytic mathematics: a single-scale activity/decay estimate for the blocked gauge theory of Balaban type [1, 2, 3, 4, 5, 6, 7], together with an infrared bound for the effective theory at the terminal scale. The repository’s polymer-expansion substrate [13, 14] is theorem-fed toward such estimates, but no part of them is proved here, and this paper claims no progress on them.
3. **Odd separations.** The gate’s conclusion family $\{2^n u\}$ misses every odd separation at every depth $n \geq 1$ — an exact theorem (`odd_separation_needs_base_scale`), not a gap in exposition. The named interpolation surrogate (`CanonicalDoublingDomination`: doubling-window ratio boundedness) would extend a dyadic bound to all separations at or above the coarsest covered spacing with explicit adjusted constants (`allSeparations_on_tower_of_dyadic_and_domination`, proved); but the clause itself is unproved for the Wilson data, is satisfiable trivially on data with vanishing correlators (`doubling_domination_of_const`), and admits a formal obstruction: a zero followed by a nonzero value inside one doubling window excludes every finite domination constant (`no_finite_domination_of_zero_then_nonzero`, `canonicalDoublingDomination_obstruction`).
4. **Thermodynamic and continuum limits.** All objects are finite-torus; the conclusion of Theorem 7.1 is uniform over an unbounded volume family but no infinite-volume correlator is defined, no limit is shown to exist, and *a fortiori* nothing is said about a continuum theory.
5. **Intermediate-remainder locality.** The operational certification of §6 covers the terminal stage $k = k_T(n) = n$. The supports of the intermediate remainders $R_k(u)$, $k < n$ — differences of correlators at consecutive stages — are not certified; the composed-pullback locality theorem for those differences is the natural next lemma.
6. **Measurability packaging.** The gate of Definition 4.3 does not carry Measurable f in its type; Theorem 6.5 assumes it. Packaging the measurability of the observable into the gate type, so that the certification applies automatically to every admitted instance, is a small named improvement.

Named open proposition	Lean name	File
metric lower transport	<code>PlaqEmbedDistLowerBound</code>	<code>ConcreteGaugeRGFidelity.lean</code>
dyadic $\Rightarrow$ literal uniformity	<code>DyadicImpliesLiteralQuestion</code>	<code>ConcreteGaugeRGFidelity.lean</code>
doubling interpolation clause	<code>CanonicalDoublingDomination</code>	<code>ConcreteGaugeRGLiteralGate.lean</code>
all-separations gate	<code>AllSeparationsCanonicalWilsonGate</code>	<code>ConcreteGaugeRGLiteralGate.lean</code>
terminal gate (renormalized)	<code>TerminalScaleWilsonGate</code>	<code>ConcreteGaugeRGTerminalGate.lean</code>
terminal gate (physical)	<code>PhysicalTerminalScaleWilsonGate</code>	<code>ConcreteGaugeRGPhysicalGate.lean</code>

Table 1: Named open propositions: exact definitions, no witnesses, no satisfiability claims. The gates are hypotheses of conditional theorems only.

Standing honesty. No reduction from anything in this development to the Clay Yang–Mills problem is claimed, and importance is not inherited by thematic proximity: only a proved reduction transfers it. Within formalized quantum field theory, the construction of the *free* field with its axioms [11] and the present interacting-RG *interface* are, in our view, the two complementary kinds of artifact the subject currently supports: constructions that exist, and specifications — machine-checked, adversarially hardened — of what remains to be constructed.

9 Reproducibility

Toolchain: Lean `leanprover/lean4:v4.29.0-rc6`; Mathlib pinned to commit

`07642720480157414db592fa85b626dafb71355b`

(lakefile and manifest agree); `tectonic` for this manuscript.

Recorded build facts: at the mathematics freeze commit

`d75d8952` (*short id; this manuscript is added in a strictly editorial later commit that changes no Lean source*)

`lake build YangMillsCore` completed successfully (**8409 jobs**, zero `sorry`, zero project axioms); the axiom oracle ran **2184 #print axioms** invocations (2181 distinct declarations; 3 historical duplicates), every one printing exactly

`[propext, Classical.choice, Quot.sound]`

with zero `sorryAx` and zero nonstandard axioms. The oracle transcripts are committed under `docs/oracle-transcripts/`; the development history, including the countermodel record of Section 2 and every audit round, is in `docs/C6-BRIDGE-CHARTER.md` — per the repository’s audit protocol it is archival material and plays no role in the mathematics above.

Headline-theorem-to-artifact map. All modules live under `YangMills/RG/`; the File column abbreviates `CorrelatorBridge.lean` (Bridge), `CorrelatorBridgeProvenance.lean` (Provenance), `ConcreteGaugeRGStep.lean` (Step), `ConcreteGaugeRGFidelity.lean` (Fidelity), `ConcreteGaugeRGLiteralGate.lean` (Literal), `ConcreteGaugeRGTerminalGate.lean` (Terminal), `ConcreteGaugeRGPhysicalGate.lean` (Physical), `ConcreteGaugeRGSupport.lean` (Support). Status *exact* = in-core Lean theorem or definition with direct oracle witness.

Result	Lean name	File	Status
abstract bridge + transport	CorrelatorBridge, physical_mass_gap_of_bridge	Bridge	exact
strong-coupling anchor	wilson_correlator_mass_gap_strong_coupling	Bridge	exact
produced remainders (§2.1)	RGFlowFamily, provenant_decomposition	Provenance	exact
Wilson measure + anchor	wilsonGibbsMeasure, rayCorrelatorOfMeasure_wilsonGibbs	Provenance	exact
provenance of nonzero UV	step_moves_of_provenantRsc_ne_zero	Provenance	exact
decimation, Prop. 3.2	blockMap, blockMap_local, blockMap_gaugeAct	Step	exact
effective measure	effectiveMeasure, effectiveMeasure_isProbability	Step	exact
observable transport	truncatedCorrelator_effectiveMeasure	Step	exact
tower + clamp theorems	towerMeasure, lt_of_concreteRsc_ne_zero	Step	exact
volume-uniform gate + consumer	NontrivialConcreteRGWilsonBridge + consumer	Step	exact
factor-3 transport, Prop. 5.1(i)	plaqEmbed_dist_le_three_mul	Fidelity	exact
canonical metric, (3) + (ii)	canonicalPlaquette_dist_eq, ..._dist_doubles	Fidelity	exact
Wilson probability threading	wilsonGibbsMeasure_isProbability (+ tower)	Fidelity	exact
telescoping, Prop. 3.4	scaledCanonical_decomposition	Fidelity	exact
per-base fidelity gate + consumer	FidelityConcreteRGWilsonGate + consumer	Fidelity	exact
literal gate + all-bases consumer	LiteralFidelityConcreteRGWilsonGate + consumer	Literal	exact
odd/small- t diagnosis (§2.4)	odd_dyadic_factorization_trivial etc.	Literal	exact
depth-zero shell implication	allSeparations_of_literalFidelityGate	Literal	exact
one-step support + separation	blockPlaquetteSupport, blockSupport_..._dist	Literal	exact
conditional de-lacunarization	allSeparations_on_tower_of_dyadic_and_domination	Literal	exact
terminal index, Def. 4.2	kTerm, terminalGate_IR_index_eq, rsc_is_difference	Terminal	exact
domination obstruction	no_finite_domination_of_zero_then_nonzero	Terminal	exact
physical gate, Def. 4.3	PhysicalTerminalScaleWilsonGate	Physical	exact
Theorem 7.1	wilson_physical_mass_gap_of_physicalTerminalGate	Physical	exact
gate comparison	terminalScaleGate_of_physicalTerminalGate	Physical	exact
Theorem 5.2	kStepBlockSupport_source_in_ball	Physical	exact
Theorem 5.3 + positivity	kStepSupport_..._dist, kStepSeparation_pos	Physical	exact
terminal tie, (6)	terminalIR_is_canonical_pair_correlator etc.	Physical	exact
composed map + Thm. 6.2	iteratedBlockMap, towerMeasure_eq_map_iterated	Support	exact
Theorem 6.3	kStepBlockObservable_congr (+ bridge)	Support	exact
Theorem 6.4	truncatedCorrelator_towerMeasure (+ terminal)	Support	exact
Theorem 6.5	terminal_support_certified	Support	exact

From a clean clone at the freeze revision (on Windows, clone with `core.autocrlf=false` and `core.longpaths=true`):

```
lake exe cache get
lake build YangMillsCore
lake env lean oracle_check.lean
tectonic -X compile papers/c6-rg-interface/c6_rg_interface.tex
```

References

- [1] T. Balaban, *Renormalization group approach to lattice gauge field theories. I. Generation of effective actions in a small field approximation and a coupling constant renormalization in four dimensions*, Comm. Math. Phys. **109** (1987), 249–301. doi:10.1007/BF01215223.
- [2] T. Balaban, *Renormalization group approach to lattice gauge field theories. II. Cluster expansions*, Comm. Math. Phys. **116** (1988), 1–22. doi:10.1007/BF01239022.
- [3] T. Balaban, *Convergent renormalization expansions for lattice gauge theories*, Comm. Math. Phys. **119** (1988), 243–285. doi:10.1007/BF01217741.
- [4] T. Balaban, *Large field renormalization. II. Localization, exponentiation, and bounds for the R operation*, Comm. Math. Phys. **122** (1989), 355–392. doi:10.1007/BF01238433.

- [5] J. Dimock, *The renormalization group according to Balaban. I. Small fields*, Rev. Math. Phys. **25** (2013), no. 7, 1330010. doi:10.1142/S0129055X13300100. arXiv:1108.1335.
- [6] J. Dimock, *The renormalization group according to Balaban. II. Large fields*, J. Math. Phys. **54** (2013), 092301. doi:10.1063/1.4821275. arXiv:1212.5562.
- [7] J. Dimock, *The renormalization group according to Balaban. III. Convergence*, Ann. Henri Poincaré **15** (2014), 2133–2175. doi:10.1007/s00023-013-0303-3.
- [8] K. Osterwalder and E. Seiler, *Gauge field theories on a lattice*, Ann. Physics **110** (1978), 440–471. doi:10.1016/0003-4916(78)90039-8.
- [9] S. Chatterjee, *Yang–Mills for probabilists*, in: Probability and Analysis in Interacting Physical Systems, Springer Proc. Math. Stat. **283**, Springer, Cham, 2019, pp. 1–16. doi:10.1007/978-3-030-15338-0_1.
- [10] S. Chatterjee, *A scaling limit of SU(2) lattice Yang–Mills–Higgs theory*, Probab. Math. Phys. **7** (2026), 339–381. arXiv:2401.10507.
- [11] M. R. Douglas, S. Hoback, A. Mei, and R. Nissim, *Formalization of QFT*, arXiv:2603.15770, 2026.
- [12] The mathlib Community, *The Lean mathematical library*, in: Proc. CPP 2020, ACM, 2020, pp. 367–381. doi:10.1145/3372885.3373824.
- [13] L. Eriksson, *A machine-checked volume-uniform Wilson-loop area law via a formalized cluster expansion*, same repository; ai.viXra.org:2607.0005, July 2026.
- [14] L. Eriksson, *Machine-checked rooted-tree majorants for polymer expansions with holes*, same repository; ai.viXra.org:2607.0025, July 2026.
- [15] L. Eriksson, *Exact two-dimensional SU(2) Yang–Mills in Lean*, same repository; ai.viXra.org:2607.0035, July 2026.
- [16] L. Eriksson, *One Amos bound, three consumer sites: machine-checked Bessel-ratio calculus for lattice gauge expansions*, companion paper, same repository, release c2-v1.2.1; ai.viXra.org:2607.0033, July 2026.
- [17] L. Eriksson, *A machine-checked proof of the Amos bound for modified Bessel function ratios*, companion paper, same repository, release c3-v1.0.1; ai.viXra.org:2607.0032, July 2026.
- [18] L. Eriksson, *A machine-checked proof of an Amos-type bound for modified Bessel ratios at real order*, companion paper, same repository, release c4-v1.0.1; ai.viXra.org:2607.0030, July 2026.
- [19] L. Eriksson, *The diagonal Amos-type family at real order: a machine-checked quantitative crossing classification*, companion paper, same repository, release c5-v1.0.1; ai.viXra.org:2607.0037, July 2026.