Cellulation-Independent Boundary Gauge Averaging and Sharp Class-Sector Gaps in Two-Dimensional Yang–Mills

Compact-group theorem, exact SU(2) constant, and auditable reduction

Lluis Eriksson

with AI-assisted derivation, formalization, and manuscript preparation

August 2026 Version 0.6

Abstract

Let G be a compact connected Lie group equipped with a bi-invariant metric. For two-dimensional G Yang–Mills with heat-kernel action, we prove a cellulation-independent finite theorem that turns a framed annular amplitude into the physical transfer kernel on gauge-invariant boundary states. An explicit conditional-Haar parametrization makes boundary constraints and boundary-edge subdivision unambiguous. Elementary edge and face subdivisions leave the normalized edge integral unchanged; after a finite subdivision made only from these two moves, a radial cut exposes one relative boundary frame, whose normalized Haar average gives

$$Z_t(g,h) = \int_{\mathsf{G}} p_t(gxh^{-1}x^{-1})\ \mathrm{d}x = \sum_{\lambda \in \widehat{\mathsf{G}}} e^{-tc_\lambda} \chi_\lambda(g)\overline{\chi_\lambda(h)}.$$

Here c_λ is the Casimir eigenvalue. Thus boundary gauge averaging is exactly the orthogonal projection of the group heat semigroup onto class functions. Writing $c_*(\mathsf{G}) = \min_{\lambda \neq \mathbf{1}} c_\lambda$, we obtain reflection positivity, the semigroup law, and the sharp mean-zero norm and gap

$$\|T_t\|_{\mathbf{1}^\perp} = e^{-tc_*(\mathsf{G})}, \qquad \mathrm{gap}(T_t) = 1 - e^{-tc_*(\mathsf{G})}.$$

Every irrep of minimum positive Casimir saturates the bound. For the stated SU(2) normalization, $c_* = 3/4$, uniquely at the fundamental representation, so boundary correlations across k equal-area strips decay sharply as $e^{-3kt/4}$. We also give a gap calculus for products and finite central quotients; in particular the induced metric on $\mathrm{SO}(3) = \mathrm{SU}(2)/\{\pm I\}$ removes the fundamental channel and raises the minimum from 3/4 to 2. The ingredients are classical; the contribution is their quantitative end-to-end assembly, including a tree–cotree disk-elimination schedule, subdivision independence, and an explicit machine-checked SU(2) frontier. No claim concerns four-dimensional Yang–Mills or arbitrary local observables.

Keywords. two-dimensional Yang–Mills; compact Lie group; heat kernel; subdivision invariance; transfer operator; spectral gap; Lean 4.
MSC 2020. 81T13, 81T25, 22E30, 68V20.

1 Result and epistemic scope

The character solution and the construction of two-dimensional Yang–Mills on compact surfaces are classical [2, 3, 4]. The point of this paper is not to rediscover them. It is to state and prove as one finite quantitative theorem a chain usually split among subdivision moves, lattice gauge fixing, harmonic analysis, and transfer-operator language:

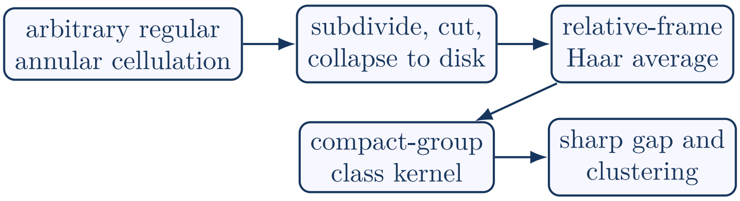

Each arrow has a different logical status. All finite-dimensional and spectral steps are proved below. The disk edge-integral endpoint and the concrete SU(2) character identities have companion Lean proofs. The topology is that of a regular finite cellulation of an annulus; no continuum measure, singular decomposition, or local-field reconstruction is claimed.

Our main result is the following.

Theorem 1.1 (Cellulation-independent compact-group transfer)**.** *Let G be a compact connected Lie group with a bi-invariant metric and Haar measure normalized to one. Let $t > 0$. On every regular finite annular cellulation of total area t, the heat-kernel lattice amplitude with gauge-invariant boundary holonomies represented by $g, h \in \mathsf{G}$ is independent of the cellulation and equals*

$$Z_t(g,h) = \sum_{\lambda \in \widehat{\mathsf{G}}} e^{-tc_\lambda} \chi_\lambda(g) \overline{\chi_\lambda(h)}. \tag{1}$$

The associated integral operator T_t on $L^2_{\mathrm{cl}}(\mathsf{G})$ is a positive self-adjoint contraction satisfying

$$T_t \chi_\lambda = e^{-tc_\lambda} \chi_\lambda, \qquad T_s T_t = T_{s+t}.$$

If $c_*(\mathsf{G}) = \min_{\lambda \neq \mathbf{1}} c_\lambda$, then on $L^2_{\mathrm{cl}}(\mathsf{G}) \ominus \mathbb{C}\mathbf{1}$,

$$\|T_t\| = e^{-tc_*(\mathsf{G})}. \tag{2}$$

Equivalently the bottom of the spectrum of $I - T_t$ on that subspace is $1 - e^{-tc_*(\mathsf{G})}$. Equality is attained by every character whose irrep has Casimir $c_*(\mathsf{G})$.

Compactness and connectedness imply that the Laplacian has discrete spectrum, the constants are its only zero modes, and $c_*(\mathsf{G}) > 0$. The area t already includes the coupling and metric convention. Replacing the heat generator by $\alpha\Delta$ replaces t everywhere by αt. Stating the convention is essential: numerical gap constants from different normalizations cannot otherwise be compared.

2 Compact-group harmonic analysis

Let $\widehat{\mathsf{G}}$ be the unitary dual. For $\lambda \in \widehat{\mathsf{G}}$, write d_λ, χ_λ, and $c_\lambda \geq 0$ for its dimension, character, and Casimir eigenvalue. Peter–Weyl and elliptic heat-kernel theory give

$$p_t(u) = \sum_{\lambda \in \widehat{\mathsf{G}}} d_\lambda e^{-tc_\lambda} \chi_\lambda(u). \tag{3}$$

For each fixed $t > 0$ the series and all character series used below converge absolutely and uniformly; this also follows directly from the smoothness of the heat kernel and standard spectral estimates on a compact manifold. The characters form an orthonormal basis of $L^2_{\mathrm{cl}}(\mathsf{G})$ and satisfy

$$\int_{\mathsf{G}} \chi_\lambda(u) \overline{\chi_\mu(u)} \, \mathrm{d}u = \delta_{\lambda\mu}, \qquad \chi_\lambda(u^{-1}) = \overline{\chi_\lambda(u)}. \tag{4}$$

The key operation is orbital averaging.

Lemma 2.1 (Orbital character product)**.** *For $a, b \in \mathsf{G}$ and every $\lambda \in \widehat{\mathsf{G}}$,*

$$\int_{\mathsf{G}} \chi_\lambda(xax^{-1}b)\ \mathrm{d}x = \frac{\chi_\lambda(a)\chi_\lambda(b)}{d_\lambda}. \tag{5}$$

Proof. Let ρ_λ be the irreducible unitary representation of dimension d_λ. Schur's lemma applied to $A = \int_{\mathsf{G}} \rho_\lambda(x)\rho_\lambda(a)\rho_\lambda(x)^{-1}\ \mathrm{d}x$ gives $A = (\chi_\lambda(a)/d_\lambda)I$. Taking the trace after multiplication by $\rho_\lambda(b)$ proves (5). □

Remark 2.2. The factor d_λ^{-1} in (5) cancels the dimension in (3). Missing this cancellation is precisely what confuses the framed heat kernel with the gauge-invariant cylinder kernel. The conjugating variable must occur as xax^{-1} in (5); the superficially similar word axa^{-1} does not satisfy this identity.

For $\mathsf{G} = \mathrm{SU}(2)$ with the normalization used in the machine-checked specialization, irreps are indexed by $n \in \mathbb{N}_0$ and

$$d_n = n + 1, \qquad c_n = \frac{n(n+2)}{4}, \qquad \chi_n(\theta) = \frac{\sin((n+1)\theta)}{\sin\theta}. \tag{6}$$

The characters are real and Weyl integration reads $\int_{\mathrm{SU}(2)} F = (2/\pi)\int_0^\pi F(\theta)\sin^2\theta\ \mathrm{d}\theta$.

3 From a physical edge integral to the annulus

We state the finite model explicitly. Let Γ be an oriented regular finite cellulation with edge set E and faces F. An edge configuration is $U \in \mathsf{G}^E$. If U_f is the ordered face holonomy and $a_f > 0$ its area, the heat-kernel density is

$$\rho_\Gamma(U) = \prod_{f\in F} p_{a_f}(U_f), \qquad t = \sum_{f\in F} a_f. \tag{7}$$

All edge integrations use product probability Haar measure. There is no formal path integral in this definition.

Definition 3.1 (Boundary-conditioned amplitude)**.** *Choose one marked vertex on each boundary and fix representatives $g, h \in \mathsf{G}$ of the two oriented boundary holonomies. Reorient the edges of an oriented boundary cycle in its direction and write them as $b = (e_1, \dots, e_r)$ from the marked vertex. Its conditional Haar integral at holonomy g is*

$$\int_{\mathrm{Hol}_b = g} \Phi\ \mathrm{d}\nu_g := \int_{\mathsf{G}^{r-1}} \Phi\Big(u_1, \dots, u_{r-1}, (u_1 \cdots u_{r-1})^{-1} g\Big) \prod_{j=1}^{r-1} \mathrm{d}u_j. \tag{8}$$

If the ambient orientation gives an inverse edge, the corresponding coordinate is inverted; inversion preserves Haar probability. The amplitude $Z_\Gamma(g, h)$ is the integral of (7) against product Haar on internal edges and the two measures (8). Thus no delta distribution or omitted gauge-volume normalization is hidden in the notation. Independent changes of frame at the two marked vertices conjugate g and h, so the physical amplitude is a class function in each variable.

Lemma 3.2 (Canonical boundary fibre)**.** *The value in (8) is independent of which boundary edge is solved for. If a boundary edge is replaced by two consecutive edges, the multiplication map from their conditional Haar coordinates to the old edge coordinate pushes the new fibre measure to the old one.*

Proof. It suffices to compare adjacent eliminated coordinates. Freeze every other edge, absorb their ordered product and g into a fixed $C \in \mathsf{G}$, and call the two adjacent coordinates v, y. The constraint has the form $vy = C$. Solving for v parametrizes the fibre by $y \mapsto (Cy^{-1}, y)$; solving for y parametrizes it by $v \mapsto (v, v^{-1}C)$. The transition $y \mapsto Cy^{-1}$ is an inversion followed by a translation and hence preserves normalized Haar measure. Adjacent exchanges connect all choices.

For a subdivision write the old coordinate as $u = ab$. Haar invariance gives

$$\int_{\mathsf{G}^2} F(ab)\ \mathrm{d}a\,\mathrm{d}b = \int_{\mathsf{G}} F(u)\ \mathrm{d}u.$$

On a fixed-holonomy fibre, one of a, b is solved from the same ordered-product constraint; the remaining coordinate integrates to one. This proves the stated pushforward, including the case in which the old edge was the eliminated coordinate. □

The next lemma is the finite mechanism behind cellulation independence.

Lemma 3.3 (Exact elementary subdivision moves)**.** *The boundary-conditioned amplitude is unchanged by either of the following moves, with the boundary data held fixed:*

(i) subdivision of an edge into two consecutive edges;

(ii) subdivision of a face of area $a + b$ by a new edge into faces of positive areas a and b.

Consequently it is unchanged under any finite sequence of such subdivisions or their inverses whenever the inverse cellulation remains regular.

Proof. For (i), every occurrence of the old variable u is replaced by u_1u_2. For any integrable F, Haar invariance and normalization give

$$\int_{\mathsf{G}^2} F(u_1u_2)\ \mathrm{d}u_1\,\mathrm{d}u_2 = \int_{\mathsf{G}} F(u)\ \mathrm{d}u.$$

The extra vertex therefore contributes exactly one, not a gauge-volume factor. If the edge is on a conditioned boundary, the same conclusion is precisely the fibre-pushforward statement of theorem 3.2; hence this proof does not silently exclude subdivision of the edge used to impose holonomy. For (ii), orient the new edge so that the two new face words are Au and $u^{-1}B$. The heat-semigroup convolution identity gives

$$\int_{\mathsf{G}} p_a(Au)p_b(u^{-1}B)\ \mathrm{d}u = p_{a+b}(AB),$$

which is precisely the old face factor. Fubini proves both statements inside the full edge integral. □

Lemma 3.4 (A radial cut after finite subdivision)**.** *Every regular finite cellulation of an annulus has a finite regular subdivision whose one-skeleton contains a simple path joining the two boundary components. Cutting along this path produces a disk cellulation; its exterior word consists of the first physical boundary, one copy of the path, the oppositely oriented second boundary, and the reverse copy of the path.*

Proof. Choose a properly embedded PL arc α between the boundary components and put it in general position with respect to the original one-skeleton: its endpoints lie in boundary edges, it avoids old vertices, and its finitely many interior intersections with edges are transverse. First subdivide every edge at an intersection point (and at the two endpoints). These are moves (i) of theorem 3.3. In each old face the components of α are now pairwise disjoint polygonal arcs with endpoints at vertices of the subdivided boundary. Insert these arcs one at a time. Each insertion divides one disk face into two disk faces, so it is move (ii); choose positive daughter areas whose sum is the old face area. After finitely many insertions α is a simple edge path in the new regular cellulation. This gives directly the required subdivision using only the two analytically invariant moves, rather than appealing to an unmatched stellar operation. Cutting an annulus along a proper boundary-to-boundary arc produces a disk, and its oriented boundary is the first physical boundary, one copy of α, the oppositely oriented second boundary, and the reverse copy of α. $\square$

Let D be the disk obtained from the chosen cut. We write $A_D(q)$ for its conditioned edge integral with exterior holonomy fixed to $q \in \mathsf{G}$; all other edge variables are integrated against product Haar probability.

Lemma 3.5 (Exact disk collapse). *For every regular disk cellulation with positive face areas a_f and exterior holonomy q,*

$$A_D(q) = p_{\sum_f a_f}(q). \tag{9}$$

Proof. Choose the boundary edge eliminated by (8) and call it e_0. The remaining boundary edges form a path, so they extend to a primal spanning tree T containing every boundary edge except e_0. Root T at the marked boundary vertex. Successively changing from a tree-edge variable to the group element at its child vertex is triangular: every step is a left or right Haar translation. Gauge covariance of face holonomy and centrality of p_{a_f} then set every tree edge to the identity, with each tree-coordinate integral contributing exactly one.

Euler's identity $|E| - |V| + 1 = |F|$ leaves one chord per bounded face. Elementary planar tree–cotree duality says that the duals of $E \setminus T$ form a tree on the bounded faces and the exterior face. Because T contains all boundary edges except e_0, the dual of e_0 is its unique edge incident to the exterior. Remove that dual edge and root the resulting tree of bounded faces at the face adjacent to e_0.

Now eliminate bounded faces from the leaves toward the root. For a non-root leaf, its parent chord occurs once in its face word and once with opposite orientation in the parent word. After absorbing the other fixed factors into $A, B \in \mathsf{G}$, the relevant integral is exactly

$$\int_{\mathsf{G}} p_a(Au)p_b(u^{-1}B)\ \mathrm{d}u = p_{a+b}(AB).$$

Thus integrating the parent chord deletes that leaf and merges its positive area into the parent. This is a finite, explicitly ordered elimination; it continues until one face remains. Its boundary word is the fixed exterior holonomy q and its area is $\sum_f a_f$, proving (9). The argument uses only normalized Haar invariance, centrality, planar tree–cotree duality, and heat-kernel convolution, and is valid for every compact group in the hypotheses. $\square$

Lemma 3.6 (Radial-cut disintegration). *For a subdivided annular cellulation as in theorem 3.4 and fixed oriented boundary holonomies g, h, cutting and regluing give the exact finite-dimensional identity*

$$Z_\Gamma(g, h) = \int_G A_D(gxh^{-1}x^{-1})\ \mathrm{d}x. \tag{10}$$

No heat-kernel expansion is used in this step.

Proof. The boundary representatives g and h are held fixed throughout; only internal coordinates are changed. Orient the two copies of the cut oppositely and choose a rooted spanning tree of the cut-open disk containing all but one of the cut-path links. Introduce a vertex potential $k_v \in \mathsf{G}$ recursively along the tree and replace every edge variable by

$$U_e \longmapsto k_{s(e)}^{-1} U_e k_{t(e)}.$$

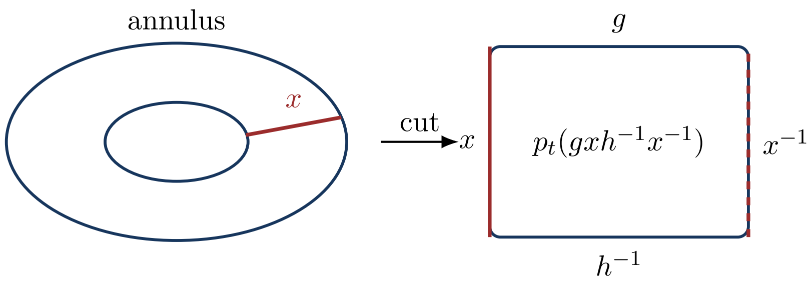

Figure 1: A radial cut exposes a relative boundary frame x. The disk edge integral collapses to one heat kernel; gluing Haar-averages x.

For fixed values of the other coordinates this is a composition of left and right translations on G; normalized Haar measure is therefore unchanged. Every tree edge becomes the identity. Centrality of each p_{a_f} makes every face factor invariant under the simultaneous vertex conjugations, so no Jacobian or residual vertex-volume factor appears.

No boundary representative is integrated and no boundary gauge is divided out in this coordinate change. The only coordinate not removed by the internal tree gauge is the transport x from the frame of the first boundary to that of the second. Reading the exterior word of the disk from the first boundary, the four surviving blocks are g, x, h^{-1}, and x^{-1}, in that order. Thus the conditional integral at fixed x is $A_D(gxh^{-1}x^{-1})$. Regluing identifies the two copies of the cut but does not fix their relative frame; Fubini disintegrates the product Haar integral over this last coordinate. Integrating x proves (10). $\square$

Proposition 3.7 (Cellulation-independent annular reduction)**.** *For every regular finite annular cellulation of total area t,*

$$Z_t(g,h) = \int_G p_t(gxh^{-1}x^{-1})\,\mathrm{d}x. \tag{11}$$

Proof. Use theorem 3.4 to choose a finite subdivision containing a radial edge path. By theorem 3.3, the original and subdivided boundary amplitudes are identical. The fixed-boundary disk theorem gives $A_D(q) = p_t(q)$ by the explicit tree–cotree schedule in theorem 3.5; no gauge-fixing or face-elimination choice is left implicit. Substitute this disk identity into (10). The resulting right side depends only on t, g, h, proving cellulation independence as well. $\square$

Theorem 3.8 (Boundary gauge projection)**.** *For $t > 0$ and $g, h \in \mathsf{G}$,*

$$\int_{\mathsf{G}} p_t(gxh^{-1}x^{-1})\,\mathrm{d}x = \sum_{\lambda\in\widehat{\mathsf{G}}} e^{-tc_\lambda}\chi_\lambda(g)\overline{\chi_\lambda(h)}. \tag{12}$$

The series converges absolutely and uniformly on $\mathsf{G}\times\mathsf{G}$ for fixed $t > 0$.

Proof. Absolute uniform convergence of the heat-kernel character expansion justifies termwise integration. Insert (3), cyclically permute inside the character, and apply theorem 2.1 with $a = h^{-1}$ and $b = g$. Then $\chi_\lambda(h^{-1}) = \overline{\chi_\lambda(h)}$, and the orbital factor d_λ^{-1} cancels the dimension in (3). Uniform convergence of the resulting series follows from $|\chi_\lambda| \le d_\lambda$ and $\sum_\lambda d_\lambda^2 e^{-tc_\lambda} < \infty$, the heat trace of the scalar Laplacian on the compact manifold G. $\square$

Remark 3.9 (Framed versus gauge-invariant boundaries). The framed transition density is $p_t(gh^{-1})$ on $L^2(\mathsf{G})$ and contains all matrix coefficients. The kernel Z_t in (12) acts on $L^2_{\mathrm{cl}}(\mathsf{G})$ and contains one character per irreducible representation. They are related by orthogonal projection, but are not pointwise equal.

4 Transfer operator and reflection positivity

Define conjugation averaging on $L^2(\mathsf{G})$ by

$$(\mathsf{P}_{\mathrm{cl}} f)(g) = \int_{\mathsf{G}} f(xgx^{-1})\ \mathrm{d}x. \tag{13}$$

Proposition 4.1 (Boundary averaging is an orthogonal projection)**.** *The operator P_{cl} is a self-adjoint idempotent contraction whose range is exactly $L^2_{\mathrm{cl}}(\mathsf{G})$. It preserves constants and positivity. If S_t is the full group heat semigroup with kernel $p_t(gh^{-1})$, then*

$$\mathsf{P}_{\mathrm{cl}} S_t = S_t \mathsf{P}_{\mathrm{cl}}, \qquad T_t = S_t|_{L^2_{\mathrm{cl}}(\mathsf{G})} = \mathsf{P}_{\mathrm{cl}} S_t \mathsf{P}_{\mathrm{cl}}|_{L^2_{\mathrm{cl}}(\mathsf{G})}. \tag{14}$$

Thus the physical cylinder operator is a genuine invariant-sector restriction, not an unrelated kernel obtained by deleting dimension factors by hand.

Proof. Haar invariance gives $\mathsf{P}_{\mathrm{cl}}^2 = \mathsf{P}_{\mathrm{cl}}$: after two averages, replace the second conjugating variable y by yx^{-1}. The same changes of variables show $\langle \mathsf{P}_{\mathrm{cl}} f, g\rangle = \langle f, \mathsf{P}_{\mathrm{cl}} g\rangle$, so an idempotent P_{cl} is the orthogonal projection onto its range. Its fixed points are precisely the conjugation-invariant functions. Averaging preserves 1, nonnegativity, and the L^2 norm bound by Jensen's inequality. Finally, p_t is central, so conjugation commutes with convolution by p_t. Restricting the resulting commuting square to the fixed-point space gives (14); its two-point kernel is (12). □

Define, initially on continuous class functions,

$$(T_t f)(g) = \int_{\mathsf{G}} Z_t(g,h) f(h)\ \mathrm{d}h. \tag{15}$$

Uniform convergence and (4) give the diagonal action

$$T_t \chi_\lambda = e^{-tc_\lambda} \chi_\lambda. \tag{16}$$

It follows immediately that T_t extends to a positive self-adjoint contraction on $L^2_{\mathrm{cl}}(\mathsf{G})$, $T_t\mathbf{1} = \mathbf{1}$, and $T_s T_t = T_{s+t}$.

Proposition 4.2 (Reflection form)**.** *For $s > 0$ and $F \in L^2_{\mathrm{cl}}(\mathsf{G})$,*

$$\iint_{\mathsf{G}\times\mathsf{G}} \overline{F(g)} Z_{2s}(g,h) F(h)\ \mathrm{d}g\,\mathrm{d}h = \|T_s F\|_2^2 \geq 0. \tag{17}$$

For a finite family $F_j \in L^2_{\mathrm{cl}}(\mathsf{G})$ and coefficients $z_j \in \mathbb{C}$, the corresponding matrix of reflected pairings is positive semidefinite.

Proof. The semigroup identity gives $Z_{2s}(g,h) = \int_G Z_s(g,k) Z_s(k,h)\ \mathrm{d}k$. Fubini and self-adjointness turn the left side into $\int_{\mathsf{G}} |(T_s F)(k)|^2\ \mathrm{d}k$. Replace F by $\sum_j z_j F_j$ for the matrix statement. □

This is the transfer-kernel form of Osterwalder–Schrader positivity for the one-dimensional time slicing supplied by the cylinder. We do not claim here a reconstruction of a local two-dimensional Wightman field theory; the bounded transfer operator is the precise object proved positive.

5 Exact compact-group gap and sharp boundary clustering

Let $L^2_{\mathrm{cl}}(\mathsf{G})^0 = \{f \in L^2_{\mathrm{cl}}(\mathsf{G}) : \int_{\mathsf{G}} f = 0\}$. The trivial character is $\mathbf{1}$, so $L^2_{\mathrm{cl}}(\mathsf{G})^0$ is the closed span of the nontrivial irreducible characters. Since G is connected and the metric is positive definite, the scalar Laplacian has constants as its entire kernel. Discreteness of its spectrum therefore gives

$$c_*(\mathsf{G}) := \min_{\lambda \neq \mathbf{1}} c_\lambda > 0.$$

Theorem 5.1 (Sharp compact-group mean-zero norm)**.** *For $t > 0$,*

$$\|T_t|_{L^2_{\mathrm{cl}}(\mathsf{G})^0}\| = e^{-tc_*(\mathsf{G})}, \qquad \inf \sigma\big((I - T_t)|_{L^2_{\mathrm{cl}}(\mathsf{G})^0}\big) = 1 - e^{-tc_*(\mathsf{G})}.$$

The equality space is exactly the span of characters with $c_\lambda = c_*(\mathsf{G})$.

Proof. Write $f = \sum_{\lambda \neq \mathbf{1}} a_\lambda \chi_\lambda$. Parseval and (16) give

$$\|T_t f\|_2^2 = \sum_{\lambda \neq \mathbf{1}} e^{-2tc_\lambda} |a_\lambda|^2 \leq e^{-2tc_*(\mathsf{G})} \sum_{\lambda \neq \mathbf{1}} |a_\lambda|^2.$$

Any character with minimum positive Casimir attains equality. The equality condition in the weighted sum also identifies the entire equality space. The spectral statement for $I - T_t$ follows from the same diagonalization. □

Corollary 5.2 (Sharp boundary strip clustering)**.** *Let $f, g \in L^2_{\mathrm{cl}}(\mathsf{G})^0$ and let $k \geq 0$ be an integer. Across k strips of area $t > 0$,*

$$|\langle f, T_t^k g\rangle| \leq e^{-ktc_*(\mathsf{G})} \|f\|_2 \|g\|_2. \tag{18}$$

The exponential rate and prefactor-one operator bound are sharp. Equality is obtained by taking $f = g$ to be any unit character of minimum positive Casimir.

Proof. $T_t^k = T_{kt}$ and theorem 5.1 give the bound by Cauchy–Schwarz; a minimum- Casimir character gives equality. □

Corollary 5.3 (Exact SU(2) constant)**.** *For $\mathsf{G} = \mathrm{SU}(2)$ with (6),*

$$c_*(\mathrm{SU}(2)) = \frac{3}{4}, \qquad \|T_t|_{\mathbf{1}^\perp}\| = e^{-3t/4}, \qquad \mathrm{gap}(T_t) = 1 - e^{-3t/4}.$$

The minimum is unique and the equality space is $\mathbb{C}\chi_1$. Hence (18) becomes the sharp bound $e^{-3kt/4}$.

Proof. For $n \geq 1$, (6) and $n(n+2) - 3 = (n-1)(n+3) \geq 0$ show that $c_n \geq 3/4$, with equality only at $n = 1$. Apply theorems 5.1 and 5.2. □

The compact-group formulation reveals how the boundary gap changes under two basic operations. This is sometimes hidden when one works only with a fixed matrix group.

Proposition 5.4 (Product and central-quotient gap calculus)**.** *Let $\mathsf{G}_1, \mathsf{G}_2$ carry product-normalized bi-invariant metrics.*

(i) For the product metric,

$$c_*(\mathsf{G}_1 \times \mathsf{G}_2) = \min\{c_*(\mathsf{G}_1), c_*(\mathsf{G}_2)\}.$$

If the first minimum is smaller, the minimum sector is the first minimum sector tensored with the trivial representation of the second; the symmetric statement holds if the second is smaller. If they coincide, the two families form the minimum sector.

(ii) If Γ is a finite central subgroup of G *and* G/Γ *has the induced quotient metric, then*

$$c_*(\mathsf{G}/\Gamma) = \min_{\substack{\lambda \neq \mathbf{1} \\ \Gamma \subseteq \ker \rho_\lambda}} c_\lambda \;\geq\; c_*(\mathsf{G}).$$

The inequality is strict exactly when no minimum-Casimir irrep of G *is trivial on* Γ.

Consequently, in the normalization (6),

$$c_*(\mathrm{SO}(3)) = 2, \qquad \mathrm{gap}(T_t^{\mathrm{SO}(3)}) = 1 - e^{-2t}. \tag{19}$$

Proof. Every irrep of a product is $\lambda_1 \boxtimes \lambda_2$, and the product Laplacian gives $c_{\lambda_1 \boxtimes \lambda_2} = c_{\lambda_1} + c_{\lambda_2}$. The smallest nonzero sum is therefore the smaller factor minimum, with the stated equality sector. Irreps of G/Γ are exactly the irreps of G whose kernels contain Γ. The quotient map is a local isometry, so their Casimir eigenvalues are unchanged. Taking the minimum over this restricted dual proves (ii) and its strictness criterion. For $\mathrm{SU}(2)/\{\pm I\}$, precisely the even labels n descend. The first nontrivial one is $n = 2$, and (6) gives $c_2 = 2$. □

Remark 5.5 (Generator gap). The generator H defined spectrally by $H\chi_\lambda = c_\lambda \chi_\lambda$ is unbounded, but $T_t = e^{-tH}$ is bounded. Its lowest nonzero energy is $c_*(\mathsf{G})$, equal to $3/4$ in the stated SU(2) normalization. The bounded statement (2), specialized to SU(2), is what the current formal interface certifies. Calling this the solution of the four-dimensional Yang–Mills mass-gap problem would be false: the dimension, dynamics, and target theorem are different.

6 What is machine checked

The companion Lean development uses Mathlib's concrete SU(2), constructs normalized Haar probability, proves the Weyl character formulas and their orthogonality, and proves the actual infinite heat-kernel convolution semigroup. The following endpoints are directly relevant:

Lean endpoint	Mathematical content
`integral_su2Haar_orbit_character_general`	orbital product formula (5)
`integral_su2ClassHeatKernel_mul`	infinite class-kernel semigroup with a genuine Haar integral
`su2Migdal_twoFace_class_merge`	two-face class Migdal merge
`conditionedEdgeModelAmplitude_eq_heatKernel`	fixed-boundary physical disk edge integral equals the heat kernel
`su2ClassTransferMultiplier_le_fundamental`	every non-vacuum multiplier is bounded by $e^{-3t/4}$
`su2ClassTransferMultiplier_bound_sharp`	the fundamental mode attains that bound

The fixed-boundary disk theorem lives on the public boundary-conditioned branch at commit `6dbb8cebc18ab2d65b6ae24af5216347c476df3f`; its provenance file identifies the audited mathematical snapshot `a1fbea97cbe673d383dbb4bc5e2a2fb70dbf190a`. The sharp multiplier lemmas are included with this manuscript as `SU2ClassTransferGap.lean`. The exact target `Lean2dYangMills.SU2ClassTransferGap` was rebuilt in this environment from the pinned manifest: Lake completed all 2,817 jobs with exit code zero. A separate `#print axioms` audit reports only `propext`, `Classical.choice`, and `Quot.sound` for the four new endpoints. The exact boundary-conditioned commit was also rebuilt as the complete `Lean2dYangMills` library: Lake completed all 3,097 jobs with exit code zero, including the physical disk endpoint cited above.

Limitation 6.1 (Formalization boundary). The current Lean corpus does not yet package elementary subdivision invariance, the PL radial-cut reduction, and the disk collapse as one theorem about regular annular combinatorial maps, nor does it construct the L^2 completion and its unbounded generator. Accordingly, the machine-checked claim is the conjunction of the disk integral, orbital character identity, class semigroup, and scalar sharp multiplier bound. The compact-group extension, topological reduction, and Hilbert-space assembly are proved in this paper, but are not advertised as one monolithic Lean theorem.

7 Reproducibility and falsification checks

The deterministic script `repro/replay.py` performs five diagnostics:

1. the removable endpoint values $\chi_n(0) = n + 1$ and $\chi_n(\pi) = (-1)^n(n + 1)$ are tested exactly;
2. Gauss–Legendre integration of the Weyl formula checks the character Gram matrix;
3. multiplier arithmetic checks $e^{-sc_n}e^{-tc_n} = e^{-(s+t)c_n}$;
4. a finite spectral search checks that $n = 1$ is the unique sharp mode;
5. a random complex character polynomial checks positivity of the reflected quadratic form.

These floating-point checks are useful falsifiers for normalization and sign errors, but they are not substitutes for proof. A successful run prints `REPLAY: PASS`.

8 Novelty, priority, and limitations

Migdal's lattice solution, the heat-kernel formulation, and Witten's character expansion make the basic cylinder formula classical [1, 2, 4]. The Markov/gluing structure of the rigorous two-dimensional Yang–Mills measure is also established in the literature [4, 5]. We therefore make no priority claim for (1) or for diagonalizing a compact-group heat semigroup.

The present theorem is therefore an assembly result, not a priority claim for the character solution or for subdivision invariance separately. Its increment through Version 0.5 is mathematically explicit: the restriction to an "admissible" SU(2) cellulation has been removed, and the sharp spectral statement is parametrized by the minimum positive Casimir of an arbitrary compact connected structure group. The contribution claimed here is:

- an explicit finite-edge-to-transfer chain with all normalization factors;
- exact invariance under the two elementary subdivision moves and a finite PL reduction from every regular annular cellulation to a radial-cut model;
- a theorem-level distinction between framed and gauge-invariant boundary kernels;
- the exact compact-group transfer gap $1 - e^{-tc_*(\mathsf{G})}$ and sharp boundary strip-clustering constant in the same convention as the physical edge integral;
- a precise map of which links are already machine checked and which are paper proofs.

Version 0.6 does not enlarge that novelty claim. It closes three proof interfaces that were previously compressed: the conditional Haar measure is given by coordinates and shown independent of the eliminated boundary edge; the PL arc is inserted using exactly the edge and face moves covered by the analytic invariance lemma; and the disk identity is proved by a finite tree–cotree elimination schedule. These changes strengthen self-containment and falsifiability, not historical priority.

The result is special to the exactly soluble two-dimensional heat-kernel model. It does not address four-dimensional continuum Yang–Mills, interacting local observables in four dimensions, or the Clay Millennium problem. Nor does it turn the boundary estimate into clustering for arbitrary local bulk observables or an Osterwalder–Schrader reconstruction. Regular finite cellulations are covered; singular CW decompositions and continuum limits are not.

9 Conclusion

Two elementary Haar identities make the finite annular theory independent of its regular cellulation; a radial cut then exposes boundary gauge averaging as the hinge between the physical edge model and the class-function transfer operator. The dimension factor cancels for every compact connected G, reflection positivity is a semigroup square, and the sharp gap is fixed by $c_*(\mathsf{G})$. For SU(2) the unique fundamental channel saturates the exact $e^{-3t/4}$ decay. The theorem remains deliberately modest about priority and dimension, but the normalization, cellulation reduction, physical integral, spectral optimum, and formal frontier now belong to one checkable chain.

A Coefficient-level proof of positivity

For a finite character polynomial $F = \sum_{\lambda\in\Lambda} a_\lambda \chi_\lambda$, direct insertion into (17) gives

$$\iint \overline{F(g)} Z_{2s}(g,h) F(h)\, \mathrm{d}g\, \mathrm{d}h = \sum_{\lambda\in\Lambda} e^{-2sc_\lambda} |a_\lambda|^2.$$

This formula simultaneously proves positivity and locates its null space. At positive time every coefficient is strictly positive, so the reflected form is strictly positive on nonzero finite character polynomials. Density extends it to all of $L^2_{\mathrm{cl}}(\mathsf{G})$; strict positivity remains, although there is no uniform lower bound because the Casimir spectrum is unbounded.

B SU(2) normalization cross-check

At $n = 0$, $c_0 = 0$ and $T_t\mathbf{1} = \mathbf{1}$. At $n = 1$, $c_1 = 3/4$. At $n = 2$, $c_2 = 2$. Hence the first multipliers are

$$1, \quad e^{-3t/4}, \quad e^{-2t}, \quad e^{-15t/4}, \ldots.$$

The full heat kernel at the identity begins instead as

$$p_t(1) = 1 + 4e^{-3t/4} + 9e^{-2t} + 16e^{-15t/4} + \cdots,$$

because $p_t(1)$ contains $d_n\chi_n(1) = d_n^2$. The two-point class kernel at $(1,1)$ has the same diagonal value, while its expansion at general (g,h) has no dimension coefficient. This is a compact check that all three formulas are consistent rather than interchangeable.

C Exact reproduction protocol

The single release archive contains the manuscript and its complete audit payload. At its root are `paper.pdf`, `manuscript.tex`, `RELEASE_NOTES.md`, and `MANIFEST.sha256`. Its `repro/` directory contains:

- `replay.py`;
- `SU2ClassTransferGap.lean`;
- `AuditClassTransferGap.lean`;
- `requirements.txt` and a command-level `README.md`;
- `reproduce.ps1`;
- `lean-2d-yang-mills.bundle`, a complete Git history containing both exact cited commits;
- `pinned-main/`, containing the main commit's exact Lake manifest, package file, and Lean toolchain declaration.

Thus neither the PDF nor the Lean dependency snapshot has to be obtained as a second attachment. Every preserved file is covered by the root manifest. From `repro/`, the numerical replay alone is

```
python -m pip install -r requirements.txt
python replay.py
```

and the complete two-checkout Lean replay is

```
powershell -ExecutionPolicy Bypass -File .\reproduce.ps1
```

The driver creates fresh directories from the included Git bundle; it never edits an existing clone. It checks out main at

`05c4ec316cb9aa295416670a2578b1c2e77e1c36`

for the scalar gap module and the boundary-conditioned branch at

`6dbb8cebc18ab2d65b6ae24af5216347c476df3f`

for `conditionedEdgeModelAmplitude_eq_heatKernel`. The audited builds completed 2,817 jobs for the scalar-gap checkout and 3,097 jobs for the complete boundary-conditioned checkout; the release driver is designed to repeat both endpoints and terminate with `FULL REPRODUCTION: PASS`. This still does not turn the compact-group/topological paper proof into a monolithic Lean theorem; it makes the formal claims that are made independently reproducible and version-unambiguous.

References

[1] A. A. Migdal, Recursion equations in gauge field theories, *Soviet Physics JETP* **42** (1975), 413–418.

[2] E. Witten, On quantum gauge theories in two dimensions, *Communications in Mathematical Physics* **141** (1991), 153–209. https://projecteuclid.org/euclid.cmp/1104248198.

[3] A. N. Sengupta, Gauge theory on compact surfaces, *Memoirs of the American Mathematical Society* **126** (1997), no. 600. https://doi.org/10.1090/memo/0600.

[4] T. Lévy, Yang–Mills measure on compact surfaces, *Memoirs of the American Mathematical Society* **166** (2003), no. 790. https://arxiv.org/abs/math/0101239.

[5] B. K. Driver, F. Gabriel, B. C. Hall, and T. Kemp, The Makeenko–Migdal equation for Yang–Mills theory on compact surfaces, *Communications in Mathematical Physics* **352** (2017), 967–978. https://arxiv.org/abs/1602.03905.

[6] C. P. Rourke and B. J. Sanderson, *Introduction to Piecewise-Linear Topology*, Springer, 1972.

[7] The mathlib Community, The Lean mathematical library, in *CPP 2020*, 367–381. https://doi.org/10.1145/3372885.3373824.