Reading the Curve, Not Counting the Points:
A Two-Ring Petrov–Galerkin Curvature Force for Discrete Yang–Mills Dynamics on the Poincaré Homology Sphere

Bee Rosa Davis
Davis Geometric
ORCID 0009-0009-8034-4308

November 2026

Abstract. On shape-regular meshes of the Poincaré homology sphere $S^3/2I$, the lowest-order Whitney discretization of the Yang–Mills force, evaluated on the exact edge integrals of a smooth connection and tested against every discrete variation, does not converge: for an Abelian control its relative error stalls at 26% under refinement. We show that the floor comes from how curvature is represented, not from the data. A closed linear curvature recovered on each cell's face-star and tested against the original Whitney variations is consistent but unstable, because the minimal recovery is exactly determined and amplifies rough data. Recovering instead on a two-ring neighbourhood by equality-constrained least squares, exact on closed linear curvature, keeps the amplification mesh-uniform and gives a Petrov–Galerkin force with no growing modes at every level computed; its error falls to 1.45% on the finest mesh. Reading three rings to second order, with each cell's own fluxes held exact, lowers it to 0.158%, with observed orders near two. The force is not the gradient of an energy, and we identify the obstruction. The construction carries over to $SU(2)$: on a Berger metric the force converges where the native one stalls at 25.5%, and its time integration follows the Ricci-flowed metric to within 0.87%. In the spin-$j/2$ sector of the linearised $SU(2)$ Yang–Mills operator at the Berger field, $(j(j+2)-14)/3$ is an exact eigenvalue, and we prove that it is the bottom of its sector in the four low sectors that survive on the Poincaré sphere. There the first sector above the left-invariant one begins $168/3$ above the field's growth mode, one third of the sphere's first scalar eigenvalue.

Keywords. Lattice gauge theory; Yang–Mills; Whitney forms; Petrov–Galerkin methods; curvature recovery; constrained least squares; Poincaré homology sphere; Berger sphere; spectral gap; computer-assisted proof.
MSC 2020. 65N30, 65N12, 81T13, 81T25, 58E15, 58J50, 53C30, 65G30.

1 Introduction

A child with a shape-sorter does not force the square block into the round hole. She turns the block until its outline matches an opening, and when no opening matches, she stops. Nobody has to be taught that this is intelligence. Fitting a shape against a guide, turning it into register, holding it where it matches and refusing where it does not, is among the first things people recognise as understanding, long before anything is measured. This paper grew from that observation, and from the image of a magnet: pieces that match should be drawn into register, not clamped into place.

The archaeologist's rim chart is the same act made quantitative. Given a single rim sherd from a broken pot, she lays its curved edge against a sheet of concentric arcs and turns it until it

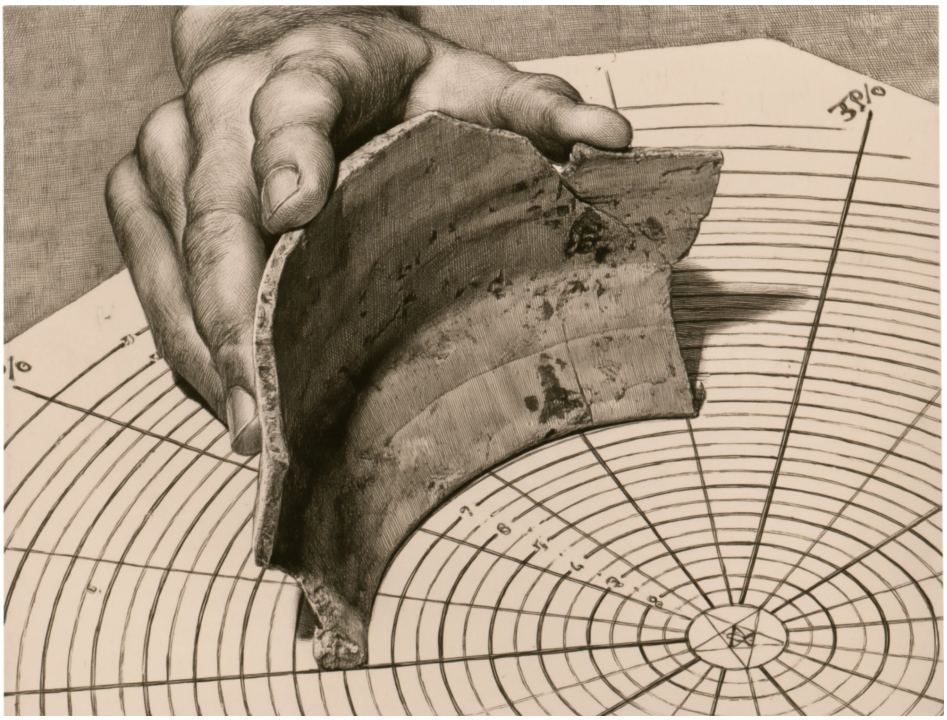


Figure 1. An archaeologist's rim chart in use: a rim sherd held edge-down on a sheet of concentric arcs of known diameter, with radial lines marking the share of the circumference the sherd covers. The arc it sits flush on gives the width of the pot.

sits flush on one of them. The arc it matches gives the width of the pot. Only the edge is used; the body of the pot is gone.

The reading is only as good as the amount of edge the guide gets to check. A sherd short enough to offer only three trustworthy points sits flush against *some* arc no matter what: three points determine a circle exactly, with nothing to spare. A chip in the rim then does not read as a chip. It reads as curvature, and on a short arc a small displacement of one point swings the fitted radius enormously. The fit cannot tell damage from shape, because it has nothing left over to check against. Give it more edge (a longer sherd, or neighbouring sherds that join it) and the match is *overdetermined*. Now a chip is outvoted rather than obeyed. And if the rim really is a circle, the best fit still recovers that circle exactly: the extra edge costs no accuracy on true shape and buys immunity to damage. When no arc fits at all, the honest reading is no reading.

A robot arm driven through a list of waypoints often moves without the curvature the task has: straight segments joined at corners. More waypoints shorten the segments but do not, by themselves, produce a curve, because nothing in the controller measures one. A child fitting a peg to a hole knows nothing about the materials and does no arithmetic, yet reads the curvature of the outline directly, from the relation between the edge under her finger and the edge beside it. The difference is not how much she sees but what she measures. The lowest-order Whitney discretization is the robot: it holds curvature constant on each cell, and refining the mesh, which gives it more data, leaves its force error at 26%. The recoveries in this paper are given exactly the same edge numbers and read from them the curvature that runs between neighbouring cells. Even the second ring adds no data; it changes only how much of the outline each reading

consults.

This paper is that act on a 3-manifold. The pot is a smooth connection on the Poincaré homology sphere. The sherds are the cells of a mesh, and, as with the rim, we are given only numbers on their edges; the curvature between those edges is what we must read. The guide is fixed. Every force in this paper is tested against the original Whitney variations of the mesh, and the recovered curvature is the piece that is turned to fit them.

Two of the guide's moves were already present in the recovery we start from. Each neighbouring cell is turned into register before it is compared, by the exact symmetry that relates the sheets of the cover. A neighbour that does not land in register, or a fit that is ill-posed, is refused rather than forced. What that recovery lacked was edge to spare. Read from each cell and its immediate neighbours, it is exactly determined, like three points on a short arc. It is accurate on smooth data, but it reads every rough pattern as large curvature, and the dynamics built on it are unstable. Read from neighbours of neighbours, the fit is overdetermined: the cell's own faces are held exactly, and every other face casts a vote. The fit stays exact on the shapes it is meant to capture, and the instability disappears. Redundancy, not resolution, is the stabilizer.

A guide does not move to meet the piece, and neither does ours: the force is not symmetric, and so it is not the gradient of an energy. Whether a *reciprocal* version exists, in which each piece's pull on the others matches the pull it receives, is a separate question. Section 11.1 states exactly what it would require.

Lattice gauge theory and geometric flows meet in a single question. Given a smooth connection A on a closed Riemannian 3-manifold, and only its integrals along the edges of a mesh, can a discrete force be computed from that data alone so that it converges to the continuum Yang–Mills force $-\delta F_A$, uniformly over all discrete variations? This question belongs to a programme that combines the heat-kernel methods of spectral geometry with Wilson's lattice gauge theory [26], and asks whether lattice gauge dynamics computed from coarse edge data reproduce the geometric flows they discretize. For homogeneous $SU(2)$ backgrounds a continuum statement is in hand.

Proposition 1.1 (Instantaneous correspondence)**.** *Let $g = \mathrm{diag}(p^2, q^2, r^2)$, $p, q, r > 0$, be a left-invariant metric on $SU(2)$ in a frame with $[e_a, e_b] = 2\varepsilon_{abc} e_c$, and write $(g_0, g_1, g_2) = (p^2, q^2, r^2)$ with indices mod 3. The Levi-Civita connection of g has axial coefficients $x_i = (g_{i+1} + g_{i+2} - g_i)/\sqrt{g_{i+1} g_{i+2}}$, and on diagonal connections y the Yang–Mills energy is $E(y) = \frac{1}{2}\sum_i (y_{i+1} y_{i+2} - 2y_i)^2/(g_{i+1} g_{i+2})$. If g moves by the Ricci flow, $\dot{g}_i = -4\big(g_i^2 - (g_{i+1} - g_{i+2})^2\big)/(g_{i+1} g_{i+2})$, then*

$$\dot{x}_i \;=\; -g_i \left.\frac{\partial E}{\partial y_i}\right|_{y=x}, \qquad i = 0, 1, 2 :$$

the Levi-Civita connection moves by the negative Yang–Mills gradient, with unit clock.

Proof. Both sides are rational functions of p, q, r. Their difference, put over a common denominator, vanishes identically in exact symbolic arithmetic, and the off-diagonal derivatives of the full nine-coordinate energy vanish on diagonal connections, so the diagonal ansatz is invariant (`e2e/probes/ricci_wilson_oracle.py`, `symbolic_certificate`). □

The statement is instantaneous and homogeneous. It is not a general Ricci-flow result: with the background metric held fixed instead of flowing, the trajectories separate in every tested

non-round case. Any lattice realisation of the programme also needs something more basic: a discrete Yang–Mills force that converges to its continuum counterpart. That prerequisite is the subject of this paper: first for an Abelian control, then for an $SU(2)$ connection on a flowing Berger metric.

The difficulty is not approximation in the usual sense. The lowest-order Whitney discretization [25, 5] is a perfectly good finite element method: its Galerkin solutions converge [2, 3, 13]. But the *instantaneous* question is different. The force is evaluated on the exact edge data $I_h A$ and tested against every Whitney variation. Here a curvature error of order h meets a test field whose curl is of order h^{-1}, and the product is of order one. The inconsistency of the combinatorial codifferential on unstructured meshes is known [1]; on our meshes it appears as a relative error of 25.9% to 26.1% that does not decrease under refinement (Table 2).

Prior attempts, recorded in the accompanying engineering reports, took two routes, and each failed one of two requirements:

- *Variational* recoveries reconstruct a better curvature and differentiate its energy. The energy becomes accurate (density 1.987 towards 2), but the gradient is wrong by hundreds of percent, because differentiating the reconstruction replaces the test variations.
- A *Petrov–Galerkin* recovery keeps the original Whitney test variations and replaces only the trial curvature. It is consistent, at 1.17% for $n = 5$, but its operator has eigenvalues with real part as negative as -2.4×10^4.

Related work. Whitney forms [25, 5] and finite element exterior calculus [2, 3, 13] give convergent Galerkin discretizations of the same operators; the failure studied here is the instantaneous consistency of the combinatorial codifferential, known on unstructured meshes [1]. Curvature recovery from patches goes back to the superconvergent patch recovery of Zienkiewicz and Zhu [27], and recovered finite element methods use it inside Galerkin schemes [11]; our recovery is a constrained least-squares fit whose stability under oversampling is in the spirit of [7]. On the lattice side, the chord reading of a plaquette is Wilson's [26], gradient flows of the lattice action are standard in lattice QCD [22], and the time stepping is by Lie-group integrators [23, 14]. Stability of Yang–Mills fields at critical points on spheres and symmetric spaces goes back to Bourguignon–Lawson and Laquer [6, 20], and closed-form fluctuation spectra on the round S^3 are known [18]; Section 10.2 reviews these and goes beyond them.

Contributions.

1. An exact decomposition of the instantaneous force error into a *reconstruction* term and a *test-variation* term (Proposition 4.1). It explains, in one identity, the Whitney floor, the variational divergence, and the Petrov–Galerkin consistency.
2. A diagnosis of the Petrov–Galerkin instability as a *conditioning* defect: the one-ring recovery is exactly determined, so its amplification of rough data is the reciprocal of a small singular value (Section 6).
3. A two-ring, equality-constrained least-squares recovery. It is exact on closed linear curvature, has mesh-uniform amplification, and gives a force that passes the unchanged Criterion C and the stability test. At $n = 1, 2$ its symmetric part is positive definite, so there the semi-discrete flow is dissipative.

4. Mechanism removals and a quantitative account of why symmetrized and variational variants of the same recovery fail (Section 8).
5. A count of the clues each neighbourhood holds, and the reading it calls for: three rings fitted to second order, with the cell's own fluxes kept exact. Its force reaches 0.158% at $n = 5$ with observed orders near two, stays stable and dissipative, and passes Criterion C (Section 9).
6. The same construction for $SU(2)$: a covariant recovery that converges at second order, a force that passes Criterion C on a Berger metric where the native force stalls at 25%, and a time integration that follows the Ricci-flowed metric to within 0.87% at $n = 3$ (Section 10).
7. A closed form for the bottom of the linearised $SU(2)$ spectrum at the Berger field, $(j(j+2)-14)/3$ in the sector of spin $j/2$, an exact eigenvalue on an explicit eigenvector in every sector, proved to be the bottom of its sector for $j = 0, 12, 20, 24$, with every higher surviving sector bounded above 142 (Theorem 10.2); minimality for general j is not proved. The companion paper [9] develops this as a self-contained spectral theorem, the Floor Theorem. On the Poincaré sphere the first sector above the left-invariant one therefore begins 168/3 above the field's growth mode, one third of the sphere's first scalar eigenvalue, and both lattice forces resolve it (Proposition 10.1).

A question for a companion paper. Every face of the mesh carries a turn, its holonomy, and there are three honest ways to read how far it turned: straight across, by the chord $U - 1$ that lattice gauge theory has used since Wilson [26]; along the way it went, by the arc $\log U$ that the continuum uses; and out and back along the tangents, whose leg is the Cayley parameter. This paper reads the same faces two of those ways, the arc in the recovery tests and the chord in the force, and Section 10.2 measures the difference in the force itself: far too small to matter here. Why the three readings miss one another in a fixed ratio of two to one, why the same ratio comes back in energy and in time stepping, and which of them the flow should step with, is the subject of the companion paper *Three Ways Across a Turn* [10].

Outline. Section 2 fixes the manifold, meshes, data and the convergence criteria, and Section 3 the computational setting. Section 4 decomposes the force error, Section 5 introduces the Petrov–Galerkin curvature, and Section 6 explains its instability and its two-ring remedy. Section 7 reports convergence and stability, Section 8 the mechanism removals and the obstruction to a variational force, and Section 9 the second-order reading. Section 10 carries the construction to $SU(2)$: the recovery, the force, the exact low spectrum at the Berger field, and the flow in time. Section 11 discusses the results and Section 12 concludes.

2 Setting

The manifold. Let $S^3 \subset \mathbb{H}$ be the unit quaternions with the round metric, and let $2I \subset S^3$ be the binary icosahedral group (order 120), acting by left multiplication. The quotient $M = S^3/2I$ is the Poincaré homology sphere, of volume $2\pi^2/120 = \pi^2/60$. Since $H^1(M) = H^2(M) = 0$, every closed 2-form on M is exact, and the only discrete kernel of the curl is pure gauge.

Meshes. We use the four shape-regular triangulations $\mathcal{T}_n$, $n \in \{1, 2, 3, 5\}$, frozen in the programme's evidence. They have $120n^3$ tetrahedra each: 120, 960, 3240 and 15000 cells, with 142,

1124, 3786 and 17510 edges, and maximal edge lengths $h_{\max} = 0.388, 0.202, 0.135, 0.082$. Physical cells are radial projections of chordal tetrahedra in $\mathbb{R}^4$. Neighbouring cells may be stored on different sheets of the cover $S^3 \to M$; they are related by left multiplication by an element of $2I$, to residual 3×10^{-16}.

Data and control. The discrete unknown is a 1-cochain $a \in \mathbb{R}^E$. The data are the exact signed edge integrals $a = I_h A$ of a left-invariant Maurer–Cartan component $A = \theta^1$, an Abelian $U(1)$ connection. On the unit sphere $\delta dA = 4A$, so the continuum Yang–Mills velocity is $\partial_t A = -\delta dA = -4A$, and the discrete target velocity is $-4a$. The curvature $F = dA$ is smooth and closed, and its face fluxes are exactly Da by Stokes' theorem, where D is the face–edge incidence matrix.

Whitney spaces and masses. Let W_e and W_f be the lowest-order Whitney 1- and 2-forms, with $dW_\eta = W_{D\eta}$. The native mass is $B_h = (\langle W_e, W_{e'} \rangle)$, in the L^2 product of the round metric. We write $\langle \cdot, \cdot \rangle$ for that product on 2-forms and $jf = W_f$ for the Whitney field with fluxes f.

Force laws and the convergence criterion. A discrete force law assigns to the data a covector $g(a) \in \mathbb{R}^E$, and the discrete velocity is $v = -B_h^{-1} g(a)$. The relative force error is

$$\varepsilon_n \;=\; \frac{\|v + 4a\|_{B_h}}{\|4a\|_{B_h}}, \qquad \|x\|_{B_h}^2 = x^\top B_h x / |M|. \tag{1}$$

Criterion C (denoted BW05 in the code) requires the following, with all quantities fixed in advance and none fitted:

- $\varepsilon_5 < 5\%$;
- observed orders $p_i = \log(e_i/e_{i+1})/\log(h_i/h_{i+1})$ of the absolute error $e_n = \|v + 4a\|_{B_h}$, of at least $\frac{1}{2}$ for the final two refinements;
- quadrature ($q4$ vs. $q6$) and solver errors below $e_n/10$ at every level.

This is the convention the gate implements (`e2e/probes/ricci_shape_evidence.py`, `numerical_gate`). The stability test *S3* requires that the generalized spectrum of (K, B_h), for the assembled force operator $g(a) = Ka$, have exactly as many eigenvalues of modulus below one as the Whitney gauge kernel, and all others with strictly positive real part.

3 Implementation in GIGI

The computations in this paper run in GIGI, a source-available geometric database and computation engine, released under the PolyForm Noncommercial License 1.0.0 at `https://github.com/nurdymuny/gigi`. GIGI is not described elsewhere, so this section says what it computes here, how each part was verified, and which parts remain research code. The implementation has three layers.

Native layer (Rust). GIGI's gauge library supplies the geometry: the Levi-Civita transport along each edge as a unit quaternion (`gauge::ricci_link::levi_civita_link`), the metric Whitney masses of each radial spherical tetrahedron (`gauge::ricci_whitney::tetra_mass`), and the cell action and tangent mass (`gauge::ricci_simplicial`). A small executable, `examples/ricci_operator_probe.rs`, runs them over a mesh. It checks every stored edge transport against the transport recomputed from the cell's own vertices, and it refuses the whole input if one disagrees:

```
let expected = parts(levi_civita_link(
    input.metric, cell.vertices[a], cell.vertices[b])?);
if (0..4).any(|j| (local[e][j] - expected[j]).abs() > 2e-12) {
    return Err("orbit identification disagrees with actual local transport".into());
}
let mass   = tetra_mass(cell.vertices, input.metric, input.order)?;
let action = cell_action(local, mass.face_mass)?;
```

The executable used here is identified by a hash in the archived inputs, together with the sources it was built from. Two native tests run against it. The first compares the cell action, masses and every derivative with an independent Python implementation. The second assembles the operator once on the quotient M and once on the entire cover S^3 (120 copies), and requires the quotient's energy, volume, gradient and Hessian blocks to equal the cover's divided by 120, to 10^{-11}. Both pass (rerun, 7.2 s). This is the check that the deck transformations of Section 2 are handled correctly before any curvature is recovered.

Research layer (Python). The recovery, the force and the convergence tests are Python in the same repository: `scripts/ricci_pg_curvature.py` (chart, transport, one-ring fit), `scripts/ricci_pg_tworing.py` (two-ring fit) and `scripts/ricci_pg_tworing_gate.py` (Criterion C and S3). The two-ring neighbourhood carries each new cell onto K's sheet by the left multiplication fixed at the face it was reached through, and refuses if the shared vertices do not coincide. This is the alignment step of the introduction. The fit (6) is then assembled as one linear map from all neighbourhood fluxes to the eleven coefficients, by the null-space method: K's own faces are held exactly, and the rest pull through least squares (the excerpt is lightly condensed from the source):

```
Ap   = np.linalg.pinv(A)                  # own faces: exact
N    = null_space(A)                      # freedom left after the own faces
RNp  = np.linalg.pinv(R @ N)              # neighbours: least squares in that freedom
fmap = np.hstack([Ap - N @ RNp @ R @ Ap, N @ RNp])
```

The code refuses a cell whose flux map has rank below eleven, whose condition number exceeds a fixed cap, or whose own fluxes are not reproduced to 10^{-12}.

Frozen inputs and recorded outputs. The meshes, edge data and native outputs are frozen in an archived checkpoint whose manifest lists every file with its SHA-256 hash, and every program that writes a result refuses to overwrite an earlier one. Criterion C was evaluated on the Abelian control only; the Berger metric is treated by the $SU(2)$ force of Section 10.2, which has its own recorded outputs. The programme's release criterion, a flow test on a merged build, is not yet met. Table 1 maps each table and figure to the program that regenerates it.

Engine verb. The two-ring recovery is also a GIGI verb, `RECOVER_CURVATURE`, implemented natively in Rust (`src/gauge/curvature_recovery.rs`). It reads the mesh and the edge cochain from two bundles and returns, for every cell, the recovered closed linear curvature and the diagnostics reported here (amplification, condition, neighbourhood size):

```
RECOVER_CURVATURE cells FROM edges ON a RINGS 2;
```

Like the research layer, it refuses rather than guesses. It names the cell when the mesh is not

Table 1. Where each result lives in the GIGI repository.

Result	Regenerated by
Native transport and masses	`examples/ricci_operator_probe.rs` (pinned binary)
Native checks	`e2e/probes/test_ricci_operator_native.py`
Force errors, Table 2	`make_data.py`, `make_tables.py`
Criterion C and S3 verdicts	`scripts/ricci_pg_tworing_gate.py`, `diagnostic.json`
Numerical range, Table 3	`numerical_range.py`
Amplification, Figure 5	`make_data.py` → `conditioning.json`
Unit tests of the two-ring fit	`scripts/test_ricci_pg_tworing.py`
$SU(2)$ numerical range, Table 10	`scripts/ricci_su2_pg_numrange.py`
$SU(2)$ time integration, Table 11	`scripts/ricci_su2_pg_flow.py`
Cayley step, $n = 3$ rows	`scripts/ricci_su2_pg_flow_cayley.py`
Slow and fast residual	`scripts/ricci_su2_pg_flow_filter.py`
Engine verb `RECOVER_CURVATURE`	`src/gauge/curvature_recovery.rs`
Verb tests, parity, live run	`tests/recover_curvature.rs`, `e2e/probes/*recover_curvature*`

closed, when a neighbour fails to coincide after transport, when a flux map is rank-deficient or ill-conditioned, or when the cell's own fluxes are not reproduced. On all four meshes and for one and two rings, it agrees with the research code cell by cell: coefficients to 1.0×10^{-14} relative, amplification to 2.1×10^{-14} and condition number to 2.0×10^{-14}, with identical neighbourhoods. On the finest mesh ($\mathcal{T}_5$, 15000 cells) the native recovery takes 8.7 s. A live `gigi-stream` built from this tree, with $\mathcal{T}_1$ ingested over HTTP, returns the same values through both of its query surfaces. One implementation detail matters: `nalgebra`'s singular value decomposition reconstructed the rank-deficient own-face matrix to only 2.6×10^{-3}, which broke the equality constraint, so the verb uses a one-sided Jacobi decomposition, accurate to rounding on these small matrices.

Scope of the code. The transport, masses and curvature recovery are GIGI engine code; the recovery is reachable as a query verb. The Criterion C measure stays an independent script: it is the oracle the engine is judged against, so it is deliberately not engine code.

The $SU(2)$ verbs. The non-Abelian computations of Section 10 are engine verbs too. The covariant reading is `RECOVER_CURVATURE ... GROUP SU2`; `LINK_FORCE` returns the force on every link, two-ring or native, and states in every answer whether the force is the gradient of an energy (only the native one is); `LINK_FLOW` integrates the flow with the Cayley step, refuses a missing metric sample by naming its time, and builds its own Newton preconditioner from the engine's linearisation. Each was held to a criterion fixed before it was built, against the research code on the same meshes:

- the covariant reading, to 1.1×10^{-13} at $n \leq 3$;
- the force, to 1.2×10^{-13} at $n \leq 5$ for two rings, native and the removal without register, and every entry of Table 7 recomputed from it to the printed digit;
- the linearisation, to 2.5×10^{-11} against the research matrices at $n \leq 3$;
- every Cayley run at $n \leq 3$, including the frozen-metric control, to 1.4×10^{-13} in E, with identical Newton iteration counts;

- the same results through a running engine, by both of its query surfaces, from a binary identified by its hash.

Each mechanism, removed, makes its test fail in the same run that passes: without register the reading is off by 0.30 and the force error returns to 39.02%; a finite-difference defect derivative moves the linearisation by 2×10^{-6}; evaluating the Cayley step at its end instead of its midpoint moves E by 6%; a round metric substituted for the Berger one gives 26.81% instead of 36.25%, which is why the engine refuses a missing metric rather than assume one. The engine is also faster: the two-ring force at $n = 5$ takes 10 s against 1,512 s for the research loop, and its linearisation at $n = 3$ takes 105 s against 6,672 s on twelve processes. The preconditioner is a dense factorisation, so the engine flow stops at $n = 3$ and refuses larger meshes by name.

4 Where the error comes from

Every force law in this paper has the form

$$g(a)[\eta] \;=\; \langle\, \mathcal{R}(Da)\,,\; \mathcal{S}(D\eta)\,\rangle, \tag{2}$$

with a *trial* map $\mathcal{R}$ and a *test* map $\mathcal{S}$, both linear from face fluxes to 2-form fields. The exact continuum force tested against a Whitney variation is, by integration by parts on the closed manifold,

$$\ell(\eta) \;=\; \langle \delta F,\, W_\eta \rangle \;=\; \langle F,\, dW_\eta \rangle \;=\; \langle F,\, W_{D\eta} \rangle \;=\; \langle F,\, jD\eta \rangle. \tag{3}$$

Proposition 4.1 (Error decomposition)**.** *Let $\Phi_F = Da$ be the fluxes of F. For every variation η,*

$$g(a)[\eta] - \ell(\eta) \;=\; \underbrace{\langle \mathcal{R}\Phi_F - F,\, \mathcal{S}D\eta \rangle}_{\textit{reconstruction}} \;+\; \underbrace{\langle F,\, (\mathcal{S} - j)D\eta \rangle}_{\textit{test variation}}. \tag{4}$$

Proof. Add and subtract $\langle F,\, \mathcal{S}D\eta \rangle$ in $g(a)[\eta] - \langle F,\, jD\eta \rangle$. □

On a quasi-uniform family of meshes, shape-regular with $h_{\max} \le c\, h_{\min}$, the curl of a Whitney variation obeys the inverse estimate $\|W_{D\eta}\| \le C_{\mathrm{inv}} h^{-1} \|W_\eta\|$ [13]; on a shape-regular family alone it holds with the local sizes h_K in place of h. The meshes used here are quasi-uniform: $h_{\max}/h_{\min} \le 1.14$ at every level. Three consequences follow, matching the measurements in Section 7:

- *Whitney* ($\mathcal{R} = \mathcal{S} = j$). The test-variation term vanishes, but the reconstruction term pairs the $O(h)$ interpolation error $jF - F$ with an $O(h^{-1})$ test field. The bound is of order one, and the measured error indeed floors at 26%.
- *Variational* ($\mathcal{S} = \mathcal{R}$). An accurate $\mathcal{R}$ makes the first term small, but the second becomes $\langle F,\, (\mathcal{R} - j)D\eta \rangle$: the recovery's response to arbitrary, rough variations, paired with the smooth field. Nothing controls it, and the measured error grows under refinement.
- *Petrov–Galerkin* ($\mathcal{S} = j$). The test-variation term vanishes identically, and

$$|g(a)[\eta] - \ell(\eta)| \;\le\; C_{\mathrm{inv}}\, h^{-1}\, \|\mathcal{R}\Phi_F - F\|\, \|W_\eta\|. \tag{5}$$

  Consistency of order h therefore follows from second-order accuracy of the recovery.

5 Petrov–Galerkin curvature

For each cell K we recover a closed linear curvature proxy, $P_K(x) = c + Bx$ with $\operatorname{tr} B = 0$: eleven coefficients, a space we denote $\mathcal{P}_1^{\mathrm{cl}}$, written in a chart of K. The chart is a central projection onto the affine 3-plane through K's vertices, with an orthonormal frame centred at K's projected centroid and scaled by K's diameter. Neighbouring cells are carried onto K's sheet by the left multiplication $p\bar{q}$ fixed by a shared vertex, and every shared vertex must coincide after transport to 10^{-12}; otherwise the method refuses. The flux of $P \in \mathcal{P}_1^{\mathrm{cl}}$ through a straight chart triangle is exactly $P(\text{centroid}) \cdot \frac{1}{2}(z_2 - z_1) \times (z_3 - z_1)$.

Let X_K be the matrix mapping the eleven coefficients to the fluxes through the distinct faces of a neighbourhood $\mathcal{N}(K)$. Split it into K's own four faces, A_K, and the rest, R_K. Then define

$$\mathcal{R}_K f \;=\; \arg\min_x \|R_K x - f_{\mathrm{rest}}\|_2^2 \quad \text{subject to} \quad A_K x = f_{\mathrm{own}}, \tag{6}$$

solved by the null-space method. This is a linear map $\mathcal{R}_K : \mathbb{R}^{|\mathcal{N}(K)|} \to \mathbb{R}^{11}$. The global trial field is $(\mathcal{R} f)|_K = P_K[\mathcal{R}_K f]$, and the Petrov–Galerkin force is

$$K_{\mathrm{PG}} \;=\; D^\top j^\top H \mathcal{R}\, D, \qquad g(a) = K_{\mathrm{PG}}\, a, \tag{7}$$

where H is the physical mass pairing a Whitney field with a cellwise recovered field. Because the recovery consumes only face fluxes, K_{PG} annihilates pure gauge exactly.

Lemma 5.1 (Exactness and accuracy)**.** *If X_K has full column rank, then $\mathcal{R}_K X_K = I_{11}$: the recovery reproduces every closed linear curvature exactly. Let $\Lambda_K = \|\mathcal{R}_K\|_2$. For smooth closed F,*

$$\|P_K[\mathcal{R}_K \Phi_F] - F\|_{L^\infty(K)} \;\le\; C\,(1 + \Lambda_K)\, h^2\, \|F\|_{W^{2,\infty}(\mathcal{N}(K))},$$

where both fields are compared in K's chart, Λ_K is measured in the scaled units of the proof, and C depends only on the shape-regularity constant of the mesh and on the number of rings.

Proof. Exactness. For data $X_K x$, the constraint $A_K x' = f_{\mathrm{own}}$ and the residual $\|R_K x' - f_{\mathrm{rest}}\|$ are both satisfied exactly at $x' = x$, and the constrained minimizer is unique when X_K has full column rank. So $\mathcal{R}_K X_K x = x$.

The chart reads every face as a straight triangle. Each face of the mesh, on K or on a neighbour, is the radial image on S^3 of a flat triangle $\Delta \subset \mathbb{R}^4$ spanned by its three vertices. The left multiplications that carry neighbours onto K's sheet are isometries of $\mathbb{R}^4$, so they map flat triangles to flat triangles. K's chart is the central projection onto the affine hyperplane H_K through K's vertices. The radial image of Δ meets H_K in $\operatorname{cone}(\Delta) \cap H_K$. The cone is a convex set in a three-dimensional linear subspace, and the method refuses any neighbour with a vertex outside the projective chart. So this intersection is the straight triangle on the projected vertices. Let $\tilde{F}$ be the pull-back of F to the chart. By the change of variables, each measured flux is exactly the flux of $\tilde{F}$ through a straight chart triangle. The recovery's input therefore carries no geometric distortion at all. Pulling back preserves closedness, so $\tilde{F}$ is a smooth divergence-free field on the chart.

Accuracy. Write d for K's diameter, so that the chart coordinates scaled by d are of order one on $\mathcal{N}(K)$. Let F_1 be the linear Taylor polynomial of $\tilde{F}$ at K's projected centroid. Its divergence is the divergence of $\tilde{F}$ at that point, which is zero, so $F_1 \in \mathcal{P}_1^{\mathrm{cl}}$. By exactness,

$$P_K[\mathcal{R}_K \Phi_F] - \tilde{F} \;=\; P_K\big[\mathcal{R}_K(\Phi_F - X_K F_1)\big] + (F_1 - \tilde{F}).$$

On a conforming shape-regular mesh, cells that share a vertex have comparable diameters, with a ratio bounded in terms of the shape-regularity constant; so $\mathcal{N}(K)$, for a fixed number of rings, has diameter at most $c_1 d$, and $|F_1 - \tilde{F}| \le \frac{1}{2} c_1^2 d^2 |\tilde{F}|_{W^{2,\infty}}$. That bounds the second term on K. It also bounds each flux defect, $|(\Phi_F - X_K F_1)_f| \le \text{area}(f)\, \frac{1}{2} c_1^2 d^2 |\tilde{F}|_{W^{2,\infty}}$, where every face has area at most $c_2 d^2$ and there are at most c_3 faces. In the scaled units in which Λ_K is measured (fluxes divided by d^2), the defect vector therefore has norm at most $\frac{1}{2} c_1^2 c_2 \sqrt{c_3}\, d^2 |\tilde{F}|_{W^{2,\infty}}$, and $\mathcal{R}_K$ multiplies it by at most Λ_K. A coefficient vector of size δ in these units gives a field of size at most $c_4 \delta$ on K. Together,

$$\|P_K[\mathcal{R}_K \Phi_F] - \tilde{F}\|_{L^\infty(K)} \;\le\; C\,(1 + \Lambda_K)\, d^2\, |\tilde{F}|_{W^{2,\infty}(\mathcal{N}(K))}.$$

The constants $c_1, \dots, c_4$ depend only on shape regularity and the number of rings, and $d \le h$.

Back to the sphere. The inverse chart is smooth. On a region of size $O(h)$ near the point of tangency, its derivatives up to second order are bounded by constants that again depend only on shape regularity. Second derivatives of the pull-back mix derivatives of F of orders zero to two, so the bound is in the full norm, $|\tilde{F}|_{W^{2,\infty}} \le C \|F\|_{W^{2,\infty}}$, and the chart norm on K is equivalent to the metric norm with constants of the same kind. This conversion costs a bounded factor, not an additional power of h. That gives the stated bound. □

Combining Lemma 5.1 with (5) gives a consistency error of order Λh, where $\Lambda = \max_K \Lambda_K$. Consistency requires only that Λ be bounded uniformly over the mesh family. Full column rank on each patch does not give that; it is a hypothesis, measured on the meshes computed here (Section 6), not proved for a general family. The convergence results below are finite-level evidence under it, not a theorem about every refinement. Stability is another matter.

6 Conditioning: why one ring is unstable

This section is the technical core of the paper. The introduction asked for pieces to be *drawn into register, not clamped into place*. Here that image becomes a statement about a linear map: how many independent pieces of information each cell's fit reads, how many it needs, and what happens to the spare ones.

Counting the star. The minimal neighbourhood is K and its four face-neighbours. Five cells have 20 faces between them, and K shares one face with each neighbour, so the star contains 16 distinct faces. (On a general mesh two neighbours of K could also share a face; on our meshes they never do. All 16 faces are distinct for each of the 4320 cells checked at $n = 1, 2, 3$.) The fluxes through these faces are not independent: the net flux out of each cell is zero, which gives one closure relation per cell. The five relations are independent, because each neighbour's relation involves three faces that no other relation contains. So the fluxes of any closed field through the star lie in a space of dimension $16 - 5 = 11$.

That is exactly $\dim \mathcal{P}_1^{\mathrm{cl}}$. The flux map X_K takes the eleven coefficients of a closed linear field to its sixteen face fluxes, and on every cell it has full rank 11. It is therefore a bijection onto the closure space. The edge data are always in that space too: face fluxes computed from edge integrals satisfy every closure relation automatically, because the boundary of a boundary is zero. So the constrained problem (6) has a solution that fits every flux exactly. In the one-ring case it is *exactly determined*. Formally only K's own four faces are hard constraints, but once they are imposed there is no freedom left, and the other twelve fluxes are matched exactly as well. Every flux is effectively clamped.

What clamping costs. On the closure space the fit is the inverse of the flux map, $\mathcal{R}_K = X_K^{-1}$, so its amplification satisfies $\Lambda_K = \|\mathcal{R}_K\|_2 \geq 1/\sigma_{\min}(X_K)$. The amplification counts all flux patterns, including those no exact data could produce, so it can exceed the bound. Measured on every cell, Λ_K lies between 1.03 and 1.13 times $1/\sigma_{\min}(X_K)$. A rough flux pattern that resembles the weakest direction of X_K is read as a closed linear field roughly $1/\sigma_{\min}$ times larger than the pattern itself.

Figure 5 shows how large that is. The one-ring amplification has median 251.8, 125.4 and 125.8 on $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$, with maxima of 252.3, 254.0 and 248.9. The operator inherits it at the rough end of its spectrum. At $n = 2$, 489 of the 961 physical eigenvalues of K_{PG} have negative real part (Table 3), and the symmetric part of K_{PG} is negative on 472 of the 961 physical directions, down to -2.4×10^4, about 2.6 times the largest Whitney eigenvalue of 9248 (Figure 10, Table 3). Restoring the second ring removes every one of those negative directions, and removing it brings them back (Section 8). That is the evidence that the instability is the clamping, not the curvature model.

Why the cell's own faces are held exactly. Nothing is thrown away. A Whitney 2-form on K is determined by its four face fluxes, and the recovered field reproduces those four fluxes exactly (within 10^{-16} on every cell checked). So the Whitney interpolant of the recovered field on K *is* the Whitney field. The recovery keeps every number the Whitney discretization had, and adds only what Whitney cannot represent: the variation of the curvature across the cell, read from the neighbours. Holding the own faces exactly is therefore not a tuning choice. It is what makes the recovery a refinement of the Whitney field rather than a replacement for it.

Drawn into register, not clamped into place. The remedy is the one the rim chart suggests, and it comes from approximation theory rather than damping. Interpolation on a minimal set of samples can have a large Lebesgue constant. Least squares on an *oversampled* set is the classical remedy, and it keeps exactness on the model space; Cohen, Davenport and Leviatan [7] quantify when it is stable. Taking $\mathcal{N}(K)$ to be every cell within two face-steps of K gives neighbourhoods of 14 or 15 cells and 40 to 44 distinct faces. Their fluxes span a closure space of dimension 26 to 29, against the eleven unknowns, so the fit is oversampled by a factor between 2.4 and 2.6.

Three steps make the two-ring fit, and each answers a different part of the image.

- *Register.* Every neighbour is first turned onto K's sheet of the cover by the exact deck transformation fixed at the face it was reached through, and refused if its shared vertices do not then coincide. Pieces are compared only once they are in register.
- *Seat.* K's own four fluxes are held exactly, as above. The piece rests on the guide where it is known to touch.
- *Draw.* Every other face pulls on the fit through least squares, all with equal weight. Many weak pulls replace a few rigid clamps. A rough flux is outvoted by its neighbours rather than obeyed.

The fit stays exact on its model space: $\mathcal{R}_K X_K = I$ to within 10^{-13} on every cell checked. Lemma 5.1 therefore holds unchanged, and the curvature still converges at second order (Figure 9). The amplification, however, falls to a median of 30.5, 35.2 and 34.6, with every cell between 22.8 and 47.2 on $\mathcal{T}_2$ and $\mathcal{T}_3$. It is *mesh-uniform*: its distribution barely moves from $\mathcal{T}_2$ to $\mathcal{T}_3$. The bound $\Lambda_K \geq 1/\sigma_{\min}(X_K)$ still holds. With two rings, Λ_K is 1.07 to 2.17 times that bound, because a least-squares map also answers to the spare, inconsistent directions, but both are far

smaller than with one ring. In the language of superconvergent patch recovery [27], the patch is redundant, and the redundancy is what absorbs roughness. Recovery operators are also the building block of the recovered finite element methods of Georgoulis and Pryer [11], which apply recovery to discontinuous spaces to construct Galerkin schemes. Here the recovery acts only on the trial curvature of a fixed Whitney discretization, and the test space is left unchanged.

Decisions, and why.

- *Two rings.* The star (one face-step) is exactly determined; two whole face-steps is the smallest radius that oversamples, and it grows the neighbourhood symmetrically, one face-step in every direction. For the linear fit only radii one and two were ever run; no radius was tuned. The second-order reading of Section 9 uses three rings because its count requires them.
- *Equal weights.* Every voting face counts the same, so the fit has no parameter. Distance-weighted variants are possible and were not tested.
- *Chart scaled by the cell's diameter.* The coefficients, and so Λ_K, are dimensionless. Amplifications on different meshes can then be compared, which is what "mesh-uniform" means.
- *Refusal over repair.* A rank-deficient or ill-conditioned flux map, or a neighbour that does not land in register, stops the fit rather than being patched. None of these refusals fired on the meshes reported here.

The language of drawing and clamping names the mechanism. It does not claim that a magnetic field enters the equations; Section 11.1 returns to what a fully reciprocal pull would require.

Seeing the amplification. Figure 7 shows what an amplification of 125 means for a single cell. One edge value in a neighbouring cell is changed by 5% of the median edge value, a chip in the rim, and the recovered curvature inside K is compared before and after. The one-ring curvature swings; the two-ring curvature barely moves. Over 300 randomly chosen cells and chips on $\mathcal{T}_2$, the one-ring response is a median 14.5 times the two-ring response (quartiles 5.8 and 25.3). That is larger than the ratio of amplifications, because least squares does two things at once: it bounds the amplification, and it dilutes a single bad value among many good ones. The ratio is not always above one (the smallest is 0.84), but the typical cell is far steadier. Figure 6 shows the same contrast for every cell of $\mathcal{T}_2$ at once.

The price is size. The one-ring fit reads 16 face fluxes per cell; the two-ring fit reads 40 on $\mathcal{T}_1$ and 40 to 44 on $\mathcal{T}_2$ and $\mathcal{T}_3$. Each per-cell problem is still an 11-coefficient constrained least-squares solve, and the assembled operator couples each cell to its two-ring neighbourhood rather than its one-ring star.

Not all of that size is needed. Figure 4 draws each ring in three dimensions and marks the second-ring cells that a greedy selection keeps: the ones that most lower the amplification, added until it is within 1.25 times the full two-ring value. A companion experiment rebuilds the force on those smaller stencils. With two to four of the roughly ten second-ring cells per cell on average, it passes Criterion C (1.57% at $n = 5$) and S3, while the same number of cells chosen at random is unstable, with 26 and then 48 runaway modes. The selection matters,

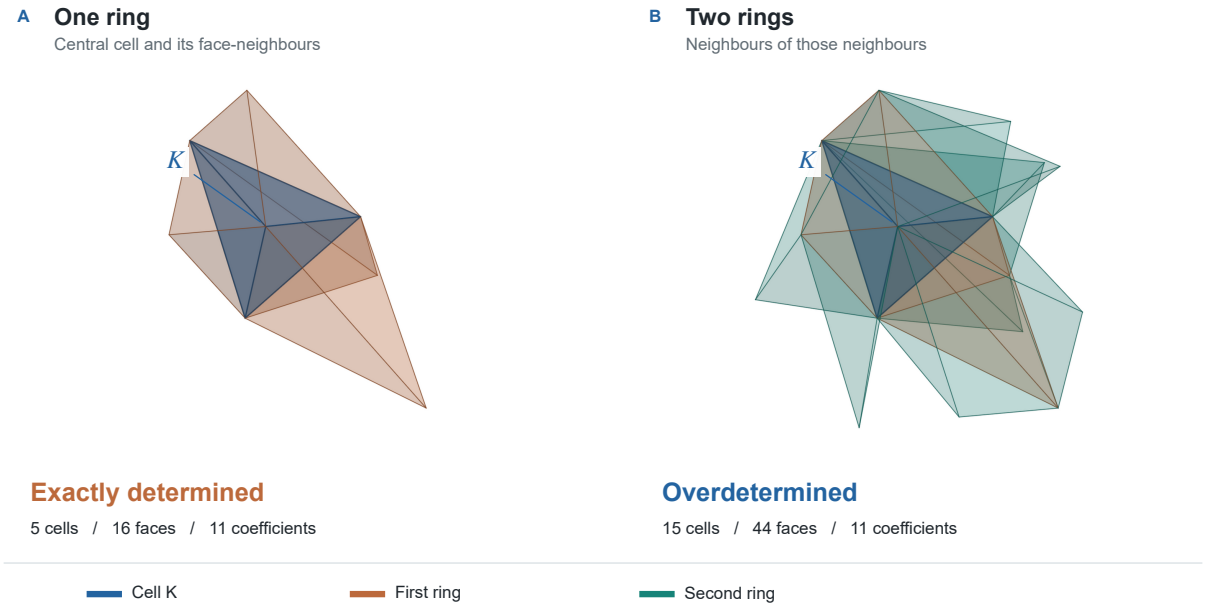


Figure 2. Recovery neighbourhoods of cell $K = 0$ of $\mathcal{T}_2$, in K's chart after every neighbour has been carried onto its sheet. Both views use the same camera and scale, with no displacement. Blue marks K, copper its face-neighbours, and teal the second ring. Translucent faces and edges are composited in actual depth order, with shared faces counted once; colour and opacity are drawing conventions, not measured fields. **(A)** One ring: 5 cells and 16 faces. Five closure relations leave exactly the eleven independent fluxes the closed linear fit needs. **(B)** Two rings: 15 cells and 44 faces, overdetermining the same eleven coefficients. Figure 3 isolates K for its measurements.

not just the count. The folded stencil is smaller but not more accurate than two rings, and its saving shrinks under refinement, so the paper keeps the uniform two-ring rule (repository: `docs/RING_FOLD_FORCE_RESULTS.md`).

7 Results

All numbers below were recomputed from the frozen inputs by the scripts accompanying this paper. The original Whitney and one-ring numbers reproduce the earlier reports to every printed digit; for example 29.8111% and 4.8414% at $n = 1$.

Convergence. Table 2 and Figure 8 give the relative force error.

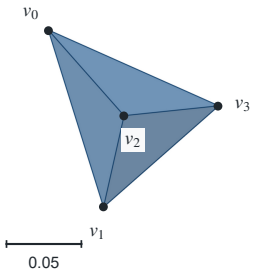
- The two-ring Petrov–Galerkin force falls from 6.22% to 1.45%, with observed orders 0.73, 0.80 and 1.29, and passes the unchanged Criterion C.
- The $q4/q6$ velocity differences are 2.3×10^{-7}, 8.2×10^{-9}, 1.4×10^{-9} and 1.0×10^{-10}, at least five orders below the error.

The cell, measured.

Four orthographic views. One scale. Measurements refer to the chordal tetrahedron.

A **Oblique view**

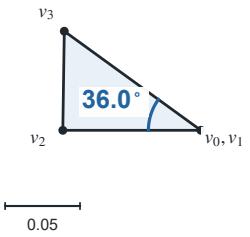
Six true edge lengths



v_0v_1	0.153	v_1v_2	0.092
v_0v_2	0.176	v_1v_3	0.113
v_0v_3	0.186	v_2v_3	0.066

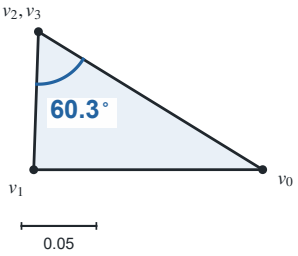
B **Along edge 01**

The two endpoints coincide



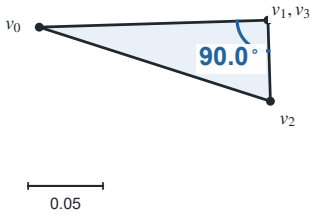
C **Along edge 23**

The two endpoints coincide



D **Along edge 13**

The two endpoints coincide



Lengths: chordal units on the unit three-sphere. Rear edges remain visible through the translucent faces.

Figure 3. The chordal tetrahedron K of Figure 2, isolated in four orthographic views at one scale; each bar is 0.05 in chordal length on the unit three-sphere. **(A)** An oblique view with translucent faces, followed by the six edge lengths keyed to the labelled vertices. These are spatial lengths, not projected lengths. **(B)–(D)** Looking along edges v_0v_1, v_2v_3 and v_1v_3: the endpoints coincide in projection, and the two incident faces appear as lines whose angle is the internal dihedral, respectively 36.0°, 60.3° and 90.0° to the displayed precision. Drawing vertices use the saved six-decimal chart coordinates; measurement labels come from `data/cell_views.json`, which also retains the other dihedrals, face angles and volume.

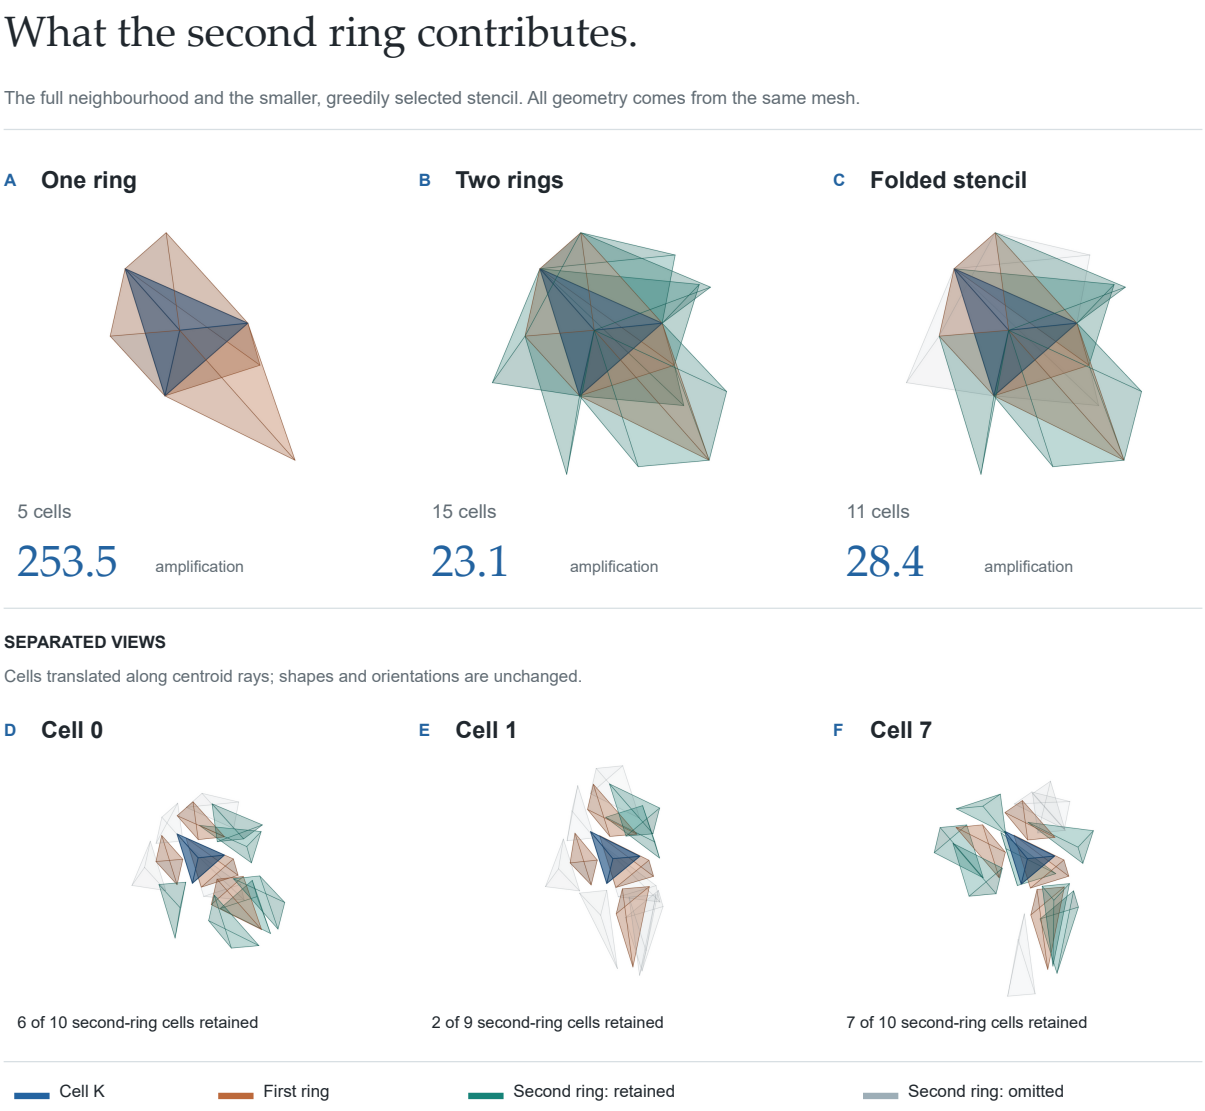


FIGURE 4. Full and selected neighbourhoods on $\mathcal{T}_2$, with depth-ordered transparency. **(A)–(C)** Cell 0 with one ring, two rings, and the folded stencil; each panel reports its cell count and fit amplification $\|\mathcal{R}_K\|_2$. The greedy selection retains second-ring cells until amplification is within 1.25 times the full two-ring value. Omitted cells remain as faint grey context and are excluded from the folded count. **(D)–(F)** Separated views of cells 0, 1 and 7, retaining respectively 6 of 10, 2 of 9, and 7 of 10 second-ring cells. Neighbours are translated along centroid rays without changing their shapes or orientations; the gaps are illustrative, not physical distances. Each cell is drawn in its own chart, with a common camera and scale within each row. The saved geometry is in `data/rings_3d.json`; `make_polished_plates.py` renders Figures 2–4 from the saved geometry and measurement records.

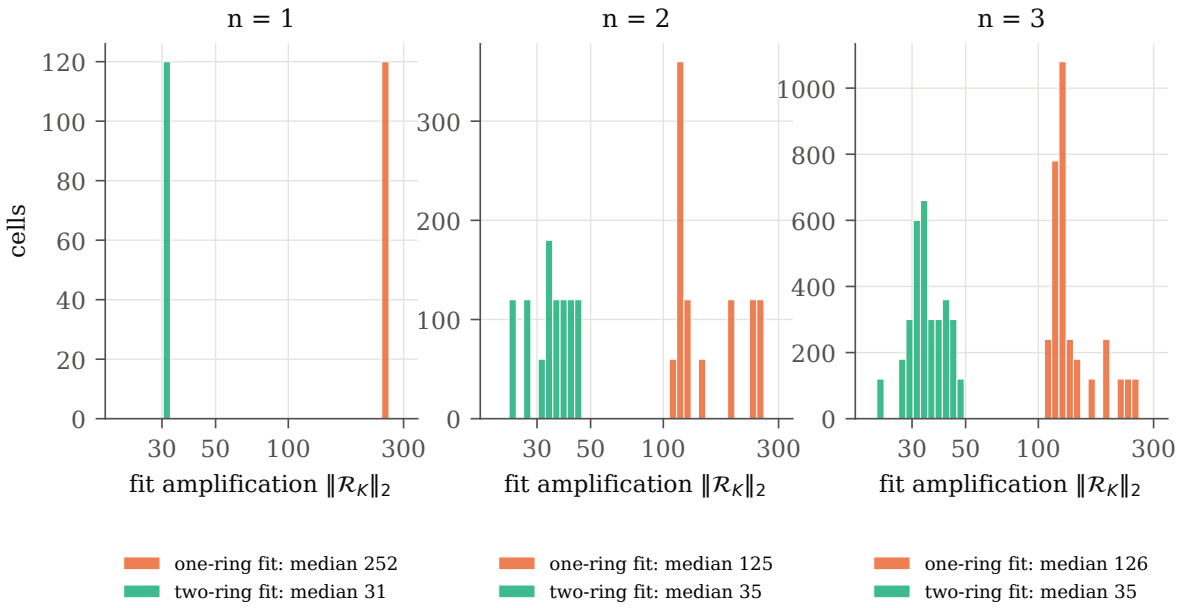


Figure 5. Per-cell amplification $\|\mathcal{R}_K\|_2$ of the curvature recovery, from neighbourhood fluxes to the eleven coefficients. The one-ring recovery is exactly determined; the two-ring recovery is overdetermined and its amplification is lower and mesh-uniform.

- The solver corrections are of order 10^{-12}.

The two-ring error is slightly above the one-ring error at every level (1.45% vs. 1.17% at $n = 5$), a modest price for stability. The recovered curvature itself converges at second order (Figure 9), reaching 0.155% at $n = 5$, against Whitney's first-order 1.53%.

Stability. Table 3 and Figure 10 give the two stability measures.

- *S3.* With two rings, the number of small eigenvalues equals the Whitney gauge dimension (21, 163, 545), and every physical eigenvalue has positive real part. The smallest real parts are 4.001, 3.994 and 3.997, the control's eigenvalue $\delta dA = 4A$; Whitney's are 4.017, 4.005 and 4.002.
- *Numerical range.* We also report the stronger certificate: the smallest eigenvalue of the symmetric part $\frac{1}{2}(K_{\mathrm{PG}} + K_{\mathrm{PG}}^{\top})$ against B_h, restricted to the physical (non-gauge) space.

Proposition 7.1 (Dissipativity). *If $\frac{1}{2}(K + K^{\top}) \geq \mu B_h$ on the physical space with $\mu > 0$, then every solution of $B_h \dot{a} = -Ka$ satisfies $\|a(t)\|_{B_h} \leq e^{-\mu t}\|a(0)\|_{B_h}$ there. In particular there is no transient non-normal growth.*

Proof. $\frac{d}{dt}\frac{1}{2}\|a\|_{B_h}^2 = -a^{\top}Ka = -a^{\top}\frac{1}{2}(K + K^{\top})a \leq -\mu\|a\|_{B_h}^2$. □

With two rings, $\mu = 3.998$ at $n = 1$ and $\mu = 3.969$ at $n = 2$. With one ring, the symmetric part is indefinite, with 70 and 472 negative directions and minima of -1.35×10^4 and -2.42×10^4.

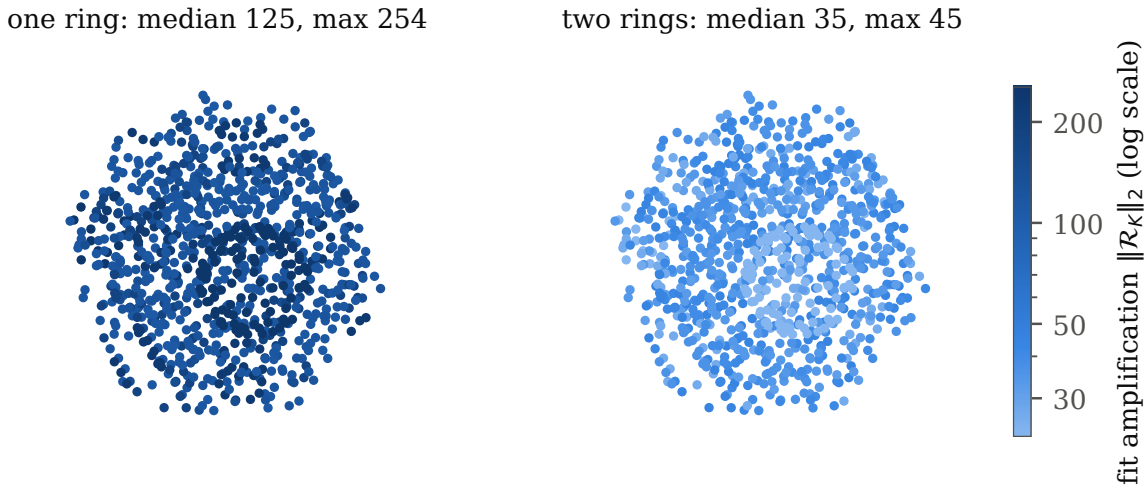


Figure 6. Fit amplification $\|\mathcal{R}_K\|_2$ for every cell of $\mathcal{T}_2$, one dot per cell at its centroid, all 960 cells carried onto one sheet of the cover and drawn in the chart of one of them. Same logarithmic colour scale in both panels.

Table 2. Relative instantaneous force error ε_n (%, $q6$). Orders are observed orders between consecutive levels, computed as in Criterion C from the absolute mass-norm error. The two-ring Petrov–Galerkin column satisfies Criterion C, evaluated by the unchanged criterion.

n	$h_{\max}$	Whitney	PG 1-ring	order	**PG 2-ring**	**order**	Var. 1-ring	Var. 2-ring
1	0.3881	29.81	4.84		**6.22**		58.2	45.4
2	0.2017	25.98	3.55	0.46	**3.83**	**0.73**	663.4	276.3
3	0.1355	25.86	2.30	1.09	**2.79**	**0.80**	778.3	330.2
5	0.0816	26.13	1.17	1.33	**1.45**	**1.29**	—	—
Two-ring: Criterion C satisfied; $q4/q6$ differences 2.3e-07–1.0e-10; solver corrections ≤ 4.1e-12.								

8 Mechanism removals, and the variational question

Each ingredient was removed in turn, on the same data and with the same measurement (Table 4):

- Removing the second ring restores the exactly determined fit, and S3 fails.
- Replacing the recovered trial by the Whitney trial restores the 26% floor.
- Replacing the Whitney test by the recovered test gives the variational force, which diverges (45% and 276% at $n = 1, 2$ with two rings; 58%, 663% and 778% with one).

The natural symmetric repair is the operator $K_{\mathrm{PG}} + K_{\mathrm{PG}}^{\top} - K_W$, the Hessian of the energy linearized in the recovery. It equals K_{PG} plus the transpose term $D^{\top}(\mathcal{R} - j)^{\top} H j D$. By (4), that term is precisely a test-variation defect, and Table 4 shows it growing under refinement for both ring counts. A force derived from an energy therefore needs a recovery whose *transpose*

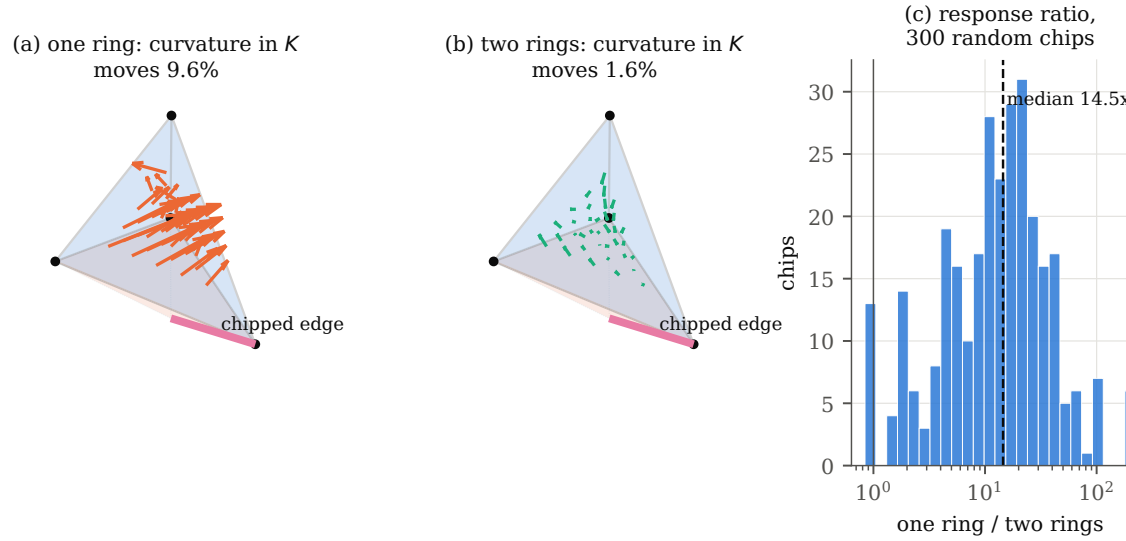


Figure 7. A chip in the rim. The same small change to one neighbouring edge value (marked, 5% of the median edge value), and the resulting change in the recovered curvature at points inside K, drawn to the same arrow scale in panels (a) and (b). **(c)** The ratio of the one-ring to the two-ring response for 300 random cells and chips on $\mathcal{T}_2$: median 14.5 (dashed), quartiles 5.8 and 25.3. The ratio is below one only rarely (the smallest is 0.84; the thin line marks one).

Table 3. Stability. S3: generalized spectrum of (K_{PG}, B_h); "gauge" counts eigenvalues of modulus below one (Whitney / PG). μ: smallest eigenvalue of the symmetric part against B_h on the physical space (Proposition 7.1). Dashes: not decomposed.

		min Re λ		# Re $\lambda < 0$		μ	
n	gauge	1-ring	**2-ring**	1-ring	**2-ring**	1-ring	**2-ring**
1	21/21	-13479	**4.001**	70	**0**	-13512	**3.998**
2	163/163	-23940	**3.994**	489	**0**	-24227	**3.969**
3	545/545	—	**3.997**	—	**0**	—	—

preserves the Whitney moments of smooth fields:

$$\mathcal{R}^{\top} H c_F \approx j^{\top} H c_F.$$

That is a column-wise condition. Imposing it exactly, together with row-wise exactness (6), couples every face to every cell whose neighbourhood contains it. The coupled system can be solved, to 10^{-8} at $n = 1$ but only to 7×10^{-5} at $n = 2$, and the next lemma shows why it cannot be imposed exactly on the fields it most needs.

Lemma 8.1 (Rows and columns conflict)**.** *Work in a chart, as Lemma 5.1 does. Suppose $\mathcal{R}$ reproduces every closed linear field from its fluxes, $\mathcal{R} I^2 p = p$, and also satisfies the column condition exactly, $\langle F, \mathcal{R} D\eta \rangle = \langle F, jD\eta \rangle$ for all η, for some field F. Then $\langle F,\ p - \Pi p \rangle = 0$ for every closed linear p, where $\Pi = jI^2$ is Whitney interpolation of 2-forms.*

Proof. Choose a potential A_p with $dA_p = p$, and take $\eta = I^1 A_p$. Integration commutes with d, so

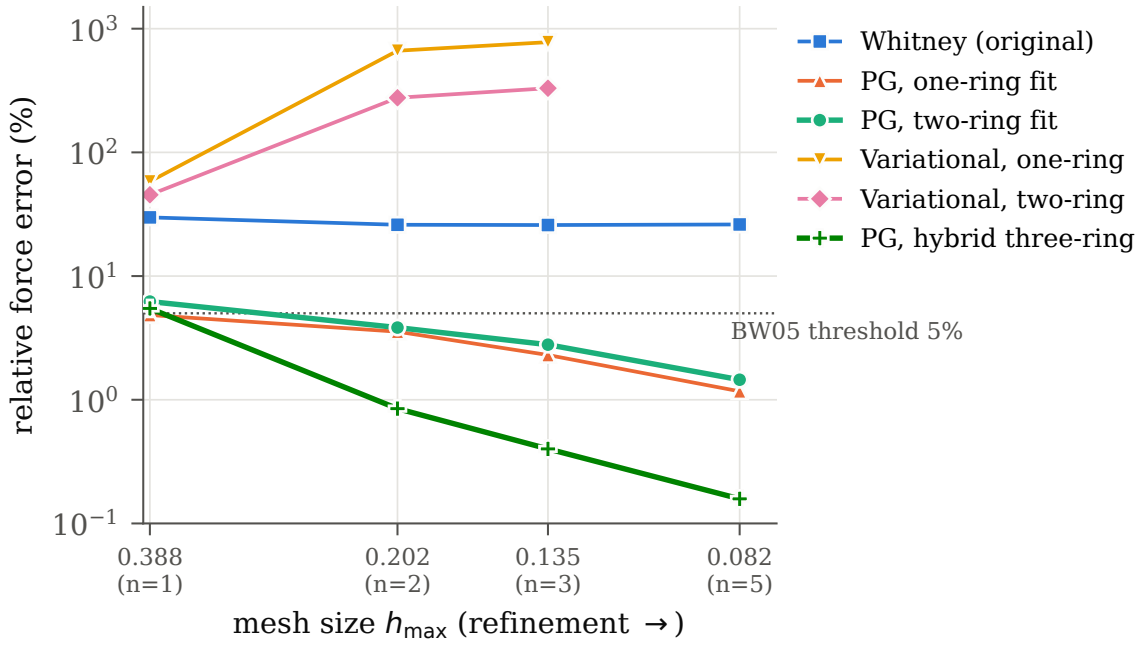


Figure 8. Relative instantaneous force error under refinement. Whitney floors near 26%. Every Petrov–Galerkin force converges, the hybrid three-ring reading (Section 9) at an order near two, and the variational force of the same recovery diverges.

$D\eta = I^2 p$. The column condition then reads $\langle F, \mathcal{R} I^2 p \rangle = \langle F, j I^2 p \rangle$, and row exactness turns the left side into $\langle F, p \rangle$. □

Whitney interpolation keeps each face's total flux, not where on the face it passes, so $\langle F, p - \Pi p \rangle$ is generically of relative order h. On the sphere, the chart's $O(h^2)$ distortion does not change this. Exact rows and exact columns on linear fields therefore ask for something that is false to first order. In the language of the guide: every piece holds its own edge exactly, and must also push back on the guide exactly as the Whitney tents would. The tents keep each wall's total but not its tilt, so the two demands disagree. The disagreement has to go somewhere. In our exploratory construction, on two rings, the only free directions are the field's own second-order part (Lemma 9.1). The reconstruction error grew to 432% at $n = 1$ and 62% at $n = 2$. Imposing the column condition on constant fields only removes the conflict, and the reconstruction term returns to 5.2% and 3.8%. But the test-variation term then stays at 58% at $n = 2$.

What remains is the poke itself. Each cell refits on its own, so the response to a change in one edge value is a set of local readings that disagree where cells meet. The defect pairs with $D\eta$, which costs a factor of h^{-1}. A reading in which every cell agrees on the response, a single tangentially continuous 1-form per edge value whose curl is the curvature reading, removes that factor: integration by parts moves the derivative from the variation onto the smooth field, and the test-variation term becomes $\langle \delta F, (\mathcal{R}^1 - W^1)\eta \rangle$. Each shared response then only has to carry the same total as its Whitney tent. This is the magnet image taken literally: one field that both sides feel, not a private opinion on each side. The mass that defines what a variation weighs has to come from the same reading. Building such a reading, with exact gauge invariance, is the

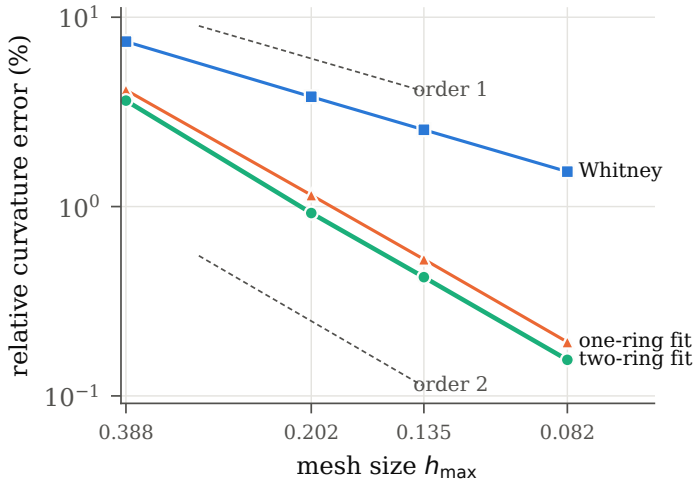


Figure 9. Relative curvature error of the trial field. Whitney converges at first order; both recoveries at second order, as Lemma 5.1 predicts.

Table 4. Mechanism removals: relative force error (%) at $n = 1, 2$ (and $n = 5$ where measured). "Transpose term" is $D^\top(\mathcal{R} - j)^\top H j\, Da$, measured in the same norm.

Variant	$n = 1$	$n = 2$	$n = 5$	Fails
Two-ring Petrov–Galerkin (this work)	6.22	3.83	1.45	—
− second ring (one-ring PG)	4.84	3.55	1.17	S3
− recovered trial (Whitney)	29.81	25.98	26.13	Criterion C floor
− Whitney test (variational, 2-ring)	45.4	276.3	—	Criterion C, diverges
− Whitney test (variational, 1-ring)	58.2	663.4	—	Criterion C, diverges
Symmetrized $K_{\mathrm{PG}} + K_{\mathrm{PG}}^\top - K_W$, 2-ring	23.0	276.5	—	Criterion C
transpose term alone, 2-ring	19.7	275.6	—	
Symmetrized, 1-ring	64.8	671.3	—	Criterion C
transpose term alone, 1-ring	61.2	669.8	—	

open problem this paper leaves.

9 A third ring: reading to second order

The two-ring force converges, but slowly: its observed orders are 0.73, 0.80 and 1.29, while its recovered curvature converges at second order. This section asks what the remaining error is made of, and answers by counting again.

Counting the clues. Section 6 counted the star: 16 faces, 11 independent fluxes, exactly $\dim \mathcal{P}_1^{\mathrm{cl}}$. The same count works one order up.

Lemma 9.1 (Unknowns one order up)**.** *In $\mathbb{R}^3$, the divergence-free vector fields whose components are polynomials of degree at most one form a space of dimension* 11*. Those whose components are homogeneous quadratics form a space of dimension* 15*. A closed field read to second order therefore has* 26 *unknowns.*

Proof. A field $c + Bx$ has $3 + 9$ coefficients and divergence $\operatorname{tr} B$, one condition: 11. A homogeneous quadratic field has $3 \times 6 = 18$ coefficients, and its divergence is a linear function, three conditions.

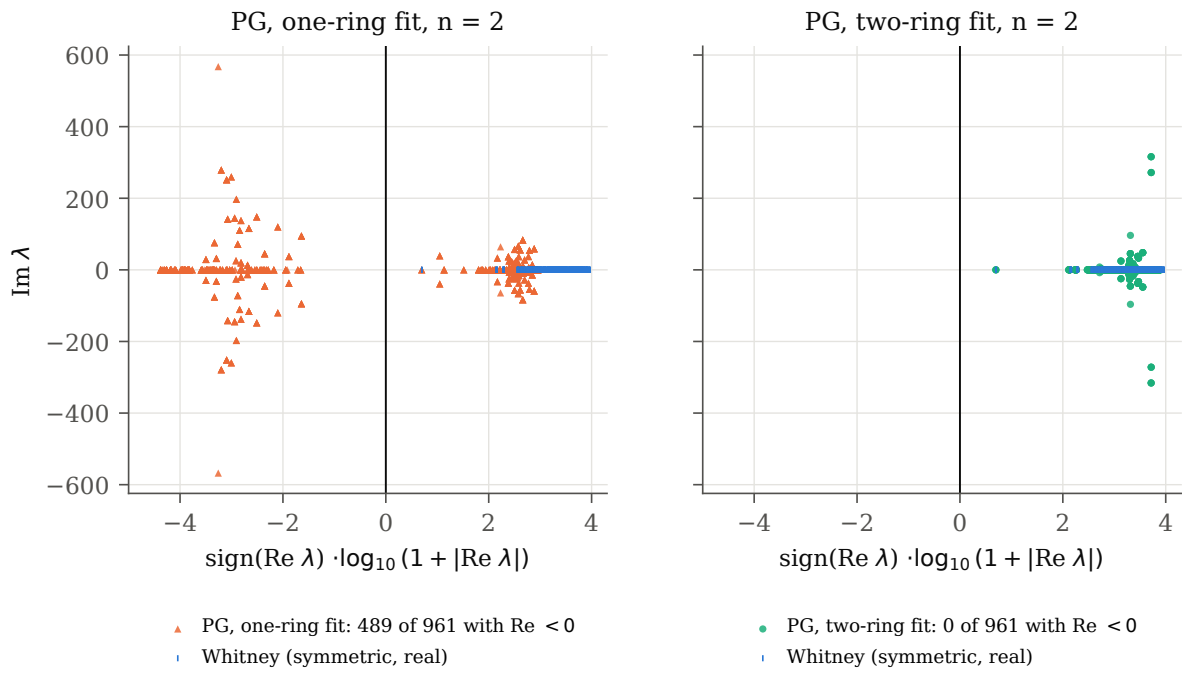


Figure 10. Generalized spectra of (K_{PG}, B_h) at $n = 2$; the real axis is signed-logarithmic. One ring: many rough modes with large negative real part. Two rings: every physical eigenvalue in the right half-plane, and the real extent matches the Whitney spectrum (ticks).

The divergence is onto the linear functions (the field $x_i x_j\, e_k$ alone already reaches every one of them), so it removes exactly three: 15. □

Now count clues: the rank of the flux map from edge data on each neighbourhood. Two rings give 26 at $n = 1$ (29 at $n = 2$), and three rings give 50 (64). Every sampled cell gives the same numbers. On every sampled cell at $n = 1$ the face fluxes of the 26 fields of Lemma 9.1 have rank 26, and they span the two-ring clue space exactly. So the two-ring fit's spare clues are not noise directions: they are the second-order part of the field.

The rim chart says what this means. A circle is fixed by three points, and the two-ring fit has, in effect, enough edge to fit a circle with a vote to spare. But a real mouth is never quite round. Over a long sherd its flare bends the edge, and a circle laid against that much edge reads the flare as a wrong radius. How wrong depends on how much far edge there happens to be, on one side or the other, so the error changes from sherd to sherd. Fitting an oval instead needs five points. A sherd that offers exactly five is short again: every point is clamped, and a chip reads as shape. That is the two-ring fit asked to read to second order, 26 unknowns from 26 clues, and its median amplification is about 1,490. A third ring gives about twice as many clues as second-order unknowns. That is the same vote the two-ring fit gives the linear one, now with room to read the flare (Figure 11).

The hybrid reading. Fit the 26 closed linear and quadratic coefficients over three rings by the constrained least squares of (6), K's own faces exact; call the map $\mathcal{R}_K^{(26)}$. The trial field stays in $\mathcal{P}_1^{\mathrm{cl}}$, so that the test pairing, the assembly and every other part of the force are unchanged.

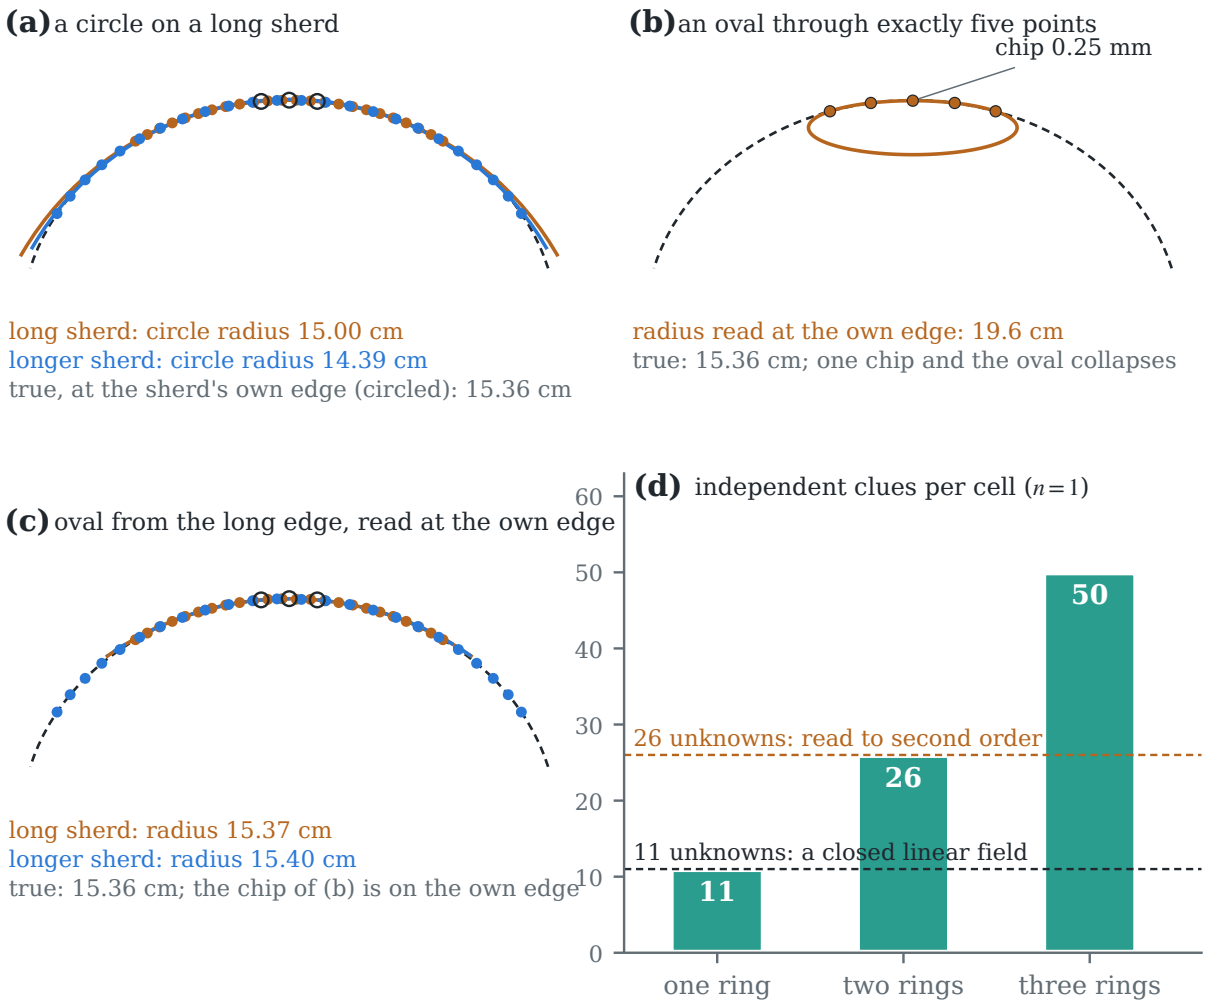


Figure 11. The rim chart one order up (a drawn illustration, computed). The true rim is an oval, semi-axes 13 and 11 cm (dashed); the reading wanted is its radius at the sherd's own edge (circled), 15.36 cm. (a) A least-squares circle reads the flare as a wrong radius, and how wrong depends on how much far edge the sherd has: 15.00 and 14.39 cm. (b) An oval through exactly five points has nothing to spare; one 0.25 mm chip collapses it, and it reads 19.6 cm. (c) The hybrid: an oval fitted to the long edge, read at the sherd's own edge, gives 15.37 and 15.40 cm with the same chip. (d) Independent clues per cell at $n = 1$ against the unknowns of Lemma 9.1. Generated by `make_figure_rimchart_oval.py`.

Keeping only the linear part of $\mathcal{R}_K^{(26)}$ would give up K's own fluxes, since a quadratic field has flux through them (a 4.8% defect, measured). Instead, add back the closed linear field of least norm that carries the quadratic part's own-face flux:

$$\mathcal{R}_K^{\mathrm{hyb}} \;=\; \mathcal{R}_K^{(26)}\big|_{\mathcal{P}_1^{\mathrm{cl}}} \;+\; A_K^{+} A_K^{\mathrm{q}} \mathcal{R}_K^{(26)}\big|_{\mathrm{quad}}, \tag{8}$$

where A_K and A_K^{q} are the own-face flux rows of the linear and quadratic fields and A_K^{+} is the pseudoinverse.

Proposition 9.2 (What the hybrid reading keeps)**.** *For all edge data a, with fluxes $\Phi = Da$:*

1. *$A_K \mathcal{R}_K^{\mathrm{hyb}} \Phi = \Phi_{\mathrm{own}}$: K's own fluxes are reproduced exactly;*
2. *$\mathcal{R}_K^{\mathrm{hyb}} X_K p = p$ for every $p \in \mathcal{P}_1^{\mathrm{cl}}$, so Lemma 5.1 applies unchanged;*
3. *a closed homogeneous quadratic field q is read as $A_K^{+} A_K^{\mathrm{q}} q$, the closed linear field of least norm with q's fluxes through K's own faces. It depends on q and on K alone, not on the outer rings.*

Proof. The fluxes of edge data through the four faces of a cell sum to zero, because the boundary of a boundary is empty. They therefore lie in the three-dimensional closure space of K, which is the range of A_K (A_K has rank 3). The quadratic fields are closed, so the range of A_K^{q} lies in the same space. The 26-field fit holds the own faces exactly: $A_K \mathcal{R}_K^{(26)}|_{\mathcal{P}_1^{\mathrm{cl}}} \Phi + A_K^{\mathrm{q}} \mathcal{R}_K^{(26)}|_{\mathrm{quad}} \Phi = \Phi_{\mathrm{own}}$. Since $A_K A_K^{+}$ is the identity on the range of A_K, applying A_K to (8) gives exactly this sum, which is (1). The fit is exact on its own 26 fields: $\mathcal{R}_K^{(26)}$ returns $(p, 0)$ on $X_K p$ and $(0, q)$ on the fluxes of q. Substituting these into (8) gives (2) and (3). □

In the rim chart, the hybrid reads the oval from the long edge and reports what that oval does at the sherd's own edge (Figure 11c). The flare is read from far away, but it can reach the reading only through the sherd's own edge. That is why, unlike the two-ring circle, the result no longer depends on how much far edge each sherd happens to have.

The construction came out of the variational investigation of Section 8. It was probed first, then measured against criteria fixed before the run (repository: `docs/RICCI_PG_HYBRID_SPEC.md`).

Result. Table 5 and Figure 8 give the result.

- *Consistency.* The force error falls to 5.47%, 0.848%, 0.401% and 0.158% at $n = 1, 2, 3, 5$, with observed orders 2.84, 1.88 and 1.83. It passes the unchanged Criterion C, and it is below the two-ring force's at every level: 9.2 times below at $n = 5$.
- *Stability.* S3 holds at $n \le 3$: the small-eigenvalue counts equal the gauge dimensions, and the smallest physical real parts are 4.016, 4.005 and 4.002.
- *Dissipativity.* The symmetric part is positive definite at $n = 1, 2$, with $\mu = 4.010$ and 4.002. These sit closer to the continuum value 4 than the two-ring force's 3.998 and 3.969.
- *Controls.* The $q4/q6$ velocity differences are at most 2×10^{-7}. The final values reproduce the exploratory probe to 10^{-12}.

The recovered curvature still converges at second order (orders 2.13, 2.08, 2.05), but with a constant 2.8 times smaller at $n = 5$ (0.056% against 0.155%).

Table 5. The hybrid three-ring reading (Section 9) against the published two-ring reading: relative force error (%, $q6$) and observed orders, computed as in Table 2; curvature error of the trial field; largest condition number of the per-cell flux map; S3 and the dissipativity constant μ of Proposition 7.1. Dashes: not computed.

	two-ring		**hybrid, three rings**					
n	PG	order	**PG**	**order**	curvature	cond.	min Re λ	μ
1	6.22		**5.474**		1.459	68	4.016	4.010
2	3.83	0.73	**0.848**	**2.84**	0.361	106	4.005	4.002
3	2.79	0.80	**0.401**	**1.88**	0.158	136	4.002	—
5	1.45	1.29	**0.158**	**1.83**	0.056	155	—	—
Hybrid: Criterion C satisfied; $q4/q6$ differences 2.0e-07–6.1e-11; solver corrections ≤ 4.1e-12.								

What does the work. Two removals, from the probe, isolate the mechanism.

- *The third ring alone does not help.* The linear fit over three rings gives 4.02%, 5.63% and 3.80% at $n = 1, 2, 3$: better than two rings at $n = 1$, worse at $n = 2$ and 3. More clues read with a first-order model bring in more second-order bias.
- *Reading to second order does,* and the own-face term keeps it honest. Keeping the linear part of $\mathcal{R}_K^{(26)}$ without that term gives 4.83%, 1.19% and 0.56%, at the price of K's own fluxes. The hybrid gives 5.47%, 0.85% and 0.40% with them exact.

The amplification also falls, to a median of 26.8, 16.2 and 15.5 at $n = 1, 2, 3$ (two rings: 30.5 to 35.2).

What is not shown.

- The bound (5) promises order one from a second-order recovery. The measured order of nearly two means the inverse estimate is not attained on this reconstruction error. A plausible reading is that the hybrid's error varies smoothly from cell to cell, where a linear fit's neighbourhood-dependent bias does not. We have not proved this.
- The cost is a larger stencil: 80 to 100 faces per cell instead of 40 to 44. A build at $n = 5$ takes 559 s, against 200 s for two rings.
- Everything in this section is Abelian. The $SU(2)$ results of Section 10 use the two-ring reading, and a covariant hybrid is open.
- Like the two-ring force, the hybrid force is not the gradient of an energy.

10 Beyond the Abelian control: covariant recovery

The results above are Abelian. This section asks whether the same reading of curvature, seated on the cell's own faces with the neighbours voting, survives when the field is a genuine $SU(2)$ connection: first the recovery, then the force (Section 10.2), then the flow in time (Section 10.3).

10.1 The covariant recovery

What changes. An $SU(2)$ field on the mesh is a link U_e on every edge, and there is no edge cochain to add up. What a face carries is its holonomy $U_f = U_{ij}U_{jk}U_{ki}$, and its logarithm

$X_f = \log U_f \in \mathfrak{su}(2)$ is the non-Abelian face flux, to leading order the integral of F over the face. Each X_f is expressed in the frame of the vertex where its loop starts. Fluxes from different faces can be compared only after each is *turned into register* in the cell's frame,

$$X_f \;\mapsto\; W_f X_f W_f^{-1},$$

where W_f is the link transport from the face's base vertex to the cell's vertex 0 along a fixed breadth-first tree of the neighbourhood. In the Abelian case, register was only the deck transformation between sheets; here it also carries the gauge frame, as the covariant Bianchi identity $D_A F = 0$ requires.

Why the same fit applies, to leading order. In Fock–Schwinger gauge at the cell's centre, $A(0) = 0$ and $dF = -[A \wedge F]$ vanishes at the centre, entering only at the next order. To the order of a linear model, each $\mathfrak{su}(2)$ component of F is therefore a closed linear field. The paper's fit map $\mathcal{R}_K$ depends only on geometry, so it is applied unchanged to the three components of the registered fluxes. Amplification and conditioning are exactly those of Section 6. This is a leading-order argument, and it is used as one. Register runs along a fixed breadth-first tree, which agrees with radial Fock–Schwinger transport only to leading order.

The own faces in $SU(2)$. In the Abelian case the cell's own faces are held exactly. In $SU(2)$ they cannot be. The four own-face fluxes of a closed linear field sum to zero (its flux map has rank three), but the Bianchi identity of a lattice field is multiplicative: neither the logarithms nor the vector parts of the registered face holonomies sum to zero on a finite cell. No closed linear field matches all four readings, and the fit, whose own-face step is a pseudoinverse, matches them in least squares instead. We measured the residual on every cell (`scripts/ricci_su2_ownface_residual.py`). It equals, to rounding, the smallest residual any closed linear field can achieve, the component of the readings along the flux map's null direction, and it falls like h^2: its median over cells is 2.53%, 0.60% and 0.27% of the own readings at $n = 1, 2, 3$ (maxima 4.3%, 0.82%, 0.37%), the same for both readings of the face. So in $SU(2)$ the own faces are held exactly only in the Abelian sense; the non-Abelian readings are held up to this residual, a relative perturbation of order h^2 in the data. The convergence results below include it; they do not rest on exact preservation.

Tests. The prototype (`scripts/ricci_su2_recovery.py`) passes three tests.

- *Embedded U(1).* The paper's Abelian data, wrapped as $SU(2)$ links along one axis, return the Abelian coefficients to 1.4×10^{-15}.
- *Gauge covariance.* A random gauge rotation at every vertex rotates the recovered field exactly, to 1.3×10^{-14}. Without register, covariance fails by 14.
- *Exact non-Abelian fields.* These are the Levi-Civita spin connections of the round metric and the Berger metric $(3, 1, 1)$, with links from the programme's exact transport. The target is built from that same exact transport around small triangles (2×10^{-5} in chart units) inside each cell, carried straight back to the frame vertex, so it needs no frame convention of its own. The error is the relative L^2 error of the recovered flux over 24 small triangles in each of 60 cells per mesh, drawn with a fixed seed (Table 6).

Past the coarsest mesh, the covariant recovery converges at second order on both fields, the same order as the Abelian recovery. The non-Abelian Stokes correction does not cost an order

Table 6. Covariant two-ring recovery on exact $SU(2)$ fields: relative error (%) of the recovered flux. The last column gives the observed orders over the last two refinements. Mechanism removals: *no register* uses each flux in its own base frame; *abelianised* replaces the logarithm of the holonomy by the sum of the edge logarithms.

Field and reading	$n = 1$	$n = 2$	$n = 3$	$n = 5$	orders
Berger, covariant	28.75	2.51	1.12	0.40	2.03, 2.03
Berger, no register	27.30	9.48	6.04	3.18	1.13, 1.27
Berger, abelianised	51.97	44.87	43.96	44.21	0.05, -0.01
Round, covariant	17.21	0.74	0.35	0.13	1.90, 1.93
Round, no register	17.21	6.18	3.82	2.04	1.21, 1.24
Round, abelianised	103.19	100.04	99.84	99.84	0.01, 0.00

here, so no midpoint correction of the face flux is needed. Without register the error converges at only about first order, so register buys an order. The abelianised flux stalls (at about 44% on the Berger field and 100% on the round field) because summing edge logarithms cannot see the commutator part of the curvature. That removal reproduces Abelian data exactly, which is why a non-Abelian test is needed to expose it. The coarsest mesh is pre-asymptotic: the two-ring stencil on 120 cells is large relative to the space. On the round field at $n = 1$, the covariant and no-register readings agree to every printed digit; at every finer level they differ widely. We record this and do not explain it.

10.2 The non-Abelian force

The recovery is only half of the paper's move; the other half is the force. Here the contract is the programme's own non-Abelian test, unchanged. The meshes, Berger metric $(3, 1, 1)$, exact links, native $SU(2)$ tangent mass and continuum target are those of the earlier shape-regular experiment, which found the native force stalled near 25%.

The force. Per cell, the native action is $E = 2\sum_{ij} M_{ij}\langle D_i, D_j\rangle$. Here $D_i = U_i - 1$ is the defect of face i's holonomy registered at vertex 0, and M is the metric Whitney face mass. Its derivative pairs a test variation $\partial D_i/\partial\eta$ with a trial curvature $h_i = \sum_j M_{ij} D_j$. The Petrov–Galerkin force keeps every test variation exactly and replaces the trial's vector part by $h_i = \int_K \langle P_K, W_i\rangle_g$, where P_K is the covariant two-ring recovery of the registered defects and W_i the Whitney face function.

Which reading of the face. In the force both steps read each face by its chord, Wilson's $U_f - 1$: the action through the defect, and the recovery through the defect's vector part (`flux="defect"` in the code). The recovery tests of Section 10.1 read the arc, $\log U_f$. The companion paper [10] works out what separates the readings, together with a third, the tangent. On these meshes the face angles are at most 0.094, and the chord–arc energy gap $a^2/12$ is below 0.08% at every face. A small energy gap does not by itself bound the gap between forces, which are derivatives (for one face the ratio of the two derivatives is $\sin a/a$, not the energy ratio), so we measured the force directly (`scripts/ricci_su2_force_log_vs_defect.py`). With the arc in the recovery instead of the chord, the two-ring force changes by 0.090%, 0.0084% and 0.0021% of its norm at $n = 1, 2, 3$, and its error by at most 0.011 percentage points (36.25%, 7.43%, 4.69% either way). The choice of reading is not a source of the measured error.

Table 7. $SU(2)$ force on the Berger field: relative full-flow error (%), $q6$. The last column gives the observed orders over the last two refinements, computed from the absolute error as in Criterion C.

Force	$n = 1$	$n = 2$	$n = 3$	$n = 5$	orders
Native $SU(2)$ Whitney	26.25	24.74	24.96	25.51	-0.04, -0.05
Two-ring Petrov–Galerkin	36.25	7.43	4.69	2.51	1.14, 1.23
removal: one ring	33.70	9.96	5.70	2.82	1.39, 1.38
removal: no register	39.02	28.72	28.95	29.19	-0.03, -0.02

Controls. Before any result was read, a Python replica reproduced the native energy to 1.5×10^{-16} at $n = 1$ (7.3×10^{-16} at $n = 2$) and the native gradient to 4.5×10^{-10} (4.8×10^{-10}); with the face defect differentiated in closed form (Section 10.3) the gradient agrees to 8×10^{-15} (1.7×10^{-14}). Feeding the cell's own Whitney field through the new pairing reproduced the native gradient exactly, and the native error was recomputed to the recorded 26.2459% at $n = 1$.

Result. The two-ring Petrov–Galerkin force passes Criterion C on the Berger field. Its final error is 2.51% and its final orders are 1.14 and 1.23. Its $q4$ and $q6$ errors agree to within 1.4×10^{-7}, and solver residuals are below 10^{-12} (Table 7, Figure 12). The native force does not pass: its error stays at 25% and its orders are negative. Register is essential. Without it the force stays at about 29%, no better than native. One ring is accurate too, as in the Abelian case; there, stability is what separated one ring from two. The coarsest mesh is again pre-asymptotic for every recovery.

Stability. As in Section 6, the object is the generalized spectrum of (K, B). Here K is the linearised force at the Berger field, one column per right-trivialised link direction by central differences, and B is the native $SU(2)$ tangent mass. The thresholds were fixed together with these results rather than before them, so the record cannot show the order; they do not decide anything here, since every cutoff between the two-ring scale (about -4.5) and the one-ring scale (below $-8{,}000$) gives the same verdict; the thresholds for $n = 3$ were committed before that level was measured. Every operator is built with the closed-form defect derivative. Rebuilt with column steps 10^{-5}, 10^{-6} and 10^{-7}, each operator gives the same count of near-zero eigenvalues (62 for every force at $n = 1$; 483, 481 and 483 for native, two-ring and one-ring at $n = 2$), the nearest eigenvalue to the unit circle sits at least 0.009 away, and the left edge of the numerical range below agrees to about 10^{-10}. Against the first build, which used a finite-difference defect derivative, every eigenvalue moves by at most 9×10^{-6} relative and no count or verdict changes. A rough mode is an eigenvalue with $\mathrm{Re}\,\lambda < -10$; the continuum scale is of order 4, and the Abelian one-ring instability sat above 10^4 in magnitude. As controls, the cell-by-cell gradient reproduces the force of Table 7 to 3×10^{-14}, and a cell-local Jacobian column matches a full-gradient column to 7×10^{-12}.

The two-ring force has no rough modes at any of the three levels (Table 8, Figure 13). The one-ring force has 207, then 1,451, then 4,437, and its most negative eigenvalue grows under refinement ($-8{,}621$, $-19{,}829$, $-45{,}543$), the same signature as the Abelian one ring. At $n = 3$ the absence of rough modes was also certified independently, from a sparse factorization of the separately built (batched) operator rather than a dense eigensolver: by Sylvester's law of inertia, no point of the numerical range of the two-ring operator, and hence no eigenvalue, lies below -10. The same factorization gives the left edge of that range as -4.665611, agreeing with the

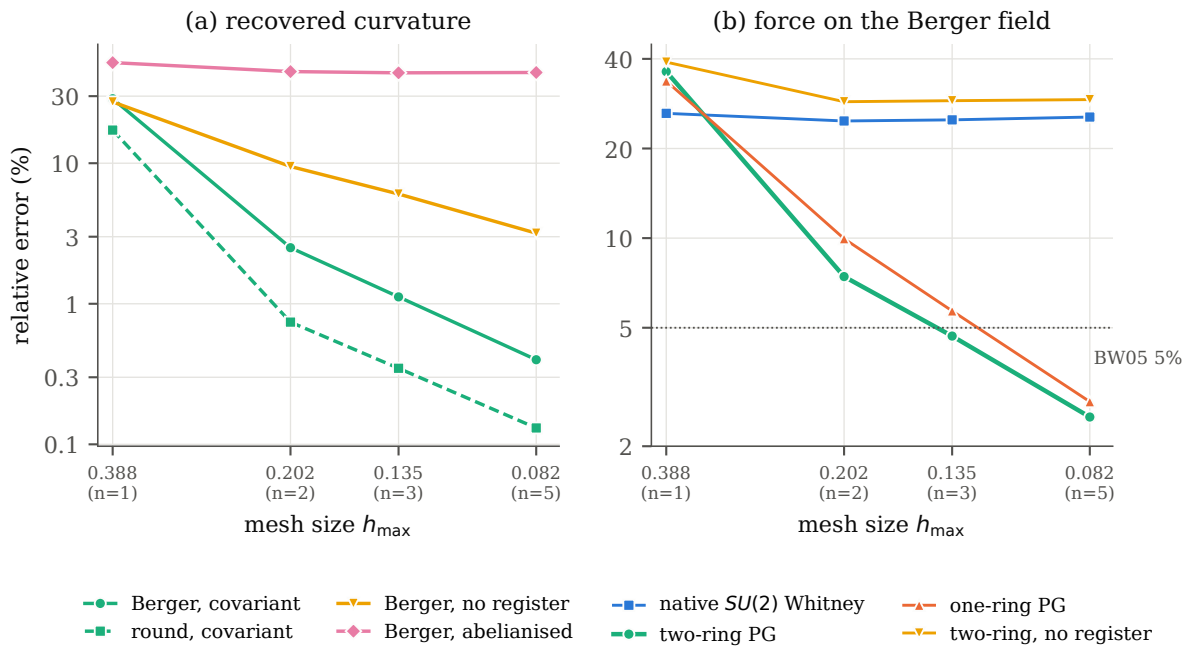


Figure 12. $SU(2)$ under refinement. (a) The covariant two-ring recovery on the exact Levi-Civita fields of the round and Berger metrics converges at second order; without register it loses an order, and the abelianised flux stalls. (b) The force on the Berger field: the native force stays near 25%, the two-ring Petrov–Galerkin force falls below the Criterion C threshold, and without register it stalls near 29%.

dense value to about 10^{-11}, and the same count of growing directions (772; native 774). The few negative eigenvalues of the two-ring operator belong to the field, not to the method. The Berger field is not at rest, so its linearisation has some directions of growth. At $n = 2$, the five most negative real parts are -4.632, -4.632, -2.503, -2.503, -1.390 for the two-ring operator and -4.650, -4.650, -2.504, -2.504, -1.366 for the native one. The two-ring operator is strongly non-normal: its imaginary parts reach 104, 150 and 418 at $n = 1, 2, 3$, where the native operator's stay below 1. That extent grows with refinement; the numerical range below shows that it stays harmless at every level measured. That is the cost of a force that is not the gradient of an energy. Gauge directions are not exact zero modes at a field that is not at rest, so the count of small eigenvalues was recorded but not used as a criterion; the exact check below says which they are.

Exact spectra on three-spheres: what is known. Exact spectral results for Yang–Mills fields on three-spheres exist in three forms. None of them reaches the operator studied here.

- *Stability at critical points.* Bourguignon and Lawson [6] studied the second variation of the Yang–Mills functional at Yang–Mills fields over the round sphere and other locally homogeneous spaces. Laquer [20] determined the index and nullity of the canonical connection over every simply connected compact irreducible symmetric space; for a group manifold such as $S^3 = SU(2)$ the index is one (see also [24]). These results concern critical points and symmetric metrics. The Berger field is neither: it is not at rest, and the Berger metric is not symmetric.
- *Fluctuation spectra on the round sphere.* Kumar, Lechtenfeld and Picanço Costa [18] diago-

Table 8. Linearised $SU(2)$ force on the Berger field: generalized spectrum against the native tangent mass. "Rough" counts eigenvalues with $\operatorname{Re}\lambda < -10$.

Force	n	directions	rough	most negative $\operatorname{Re}\lambda$	$\max \lvert\operatorname{Im}\lambda\rvert$
Native $SU(2)$ Whitney	1	426	0	−4.48	0.80
Two-ring Petrov–Galerkin	1	426	**0**	−4.39	104
removal: one ring	1	426	207	−8,621	345
Native $SU(2)$ Whitney	2	3,372	0	−4.65	0.63
Two-ring Petrov–Galerkin	2	3,372	**0**	−4.63	150
removal: one ring	2	3,372	1,451	−19,829	663
Native $SU(2)$ Whitney	3	11,358	0	−4.66	0.64
Two-ring Petrov–Galerkin	3	11,358	**0**	−4.65	418
removal: one ring	3	11,358	4,437	−45,543	1,306

nalized the Yang–Mills fluctuation operator around homogeneous backgrounds on the round S^3, sector by sector, in closed form. The meron's spectrum there begins $-2, 1, 6, 7, 10$. There is no squashing and no quotient.

- *Closed forms on Berger spheres for other operators.* On Berger spheres, exact spectra are known for the scalar Laplacian [21], the Dirac operator [4, 17] and, recently, the Hodge Laplacian on 1-forms [12]. None of them is the Yang–Mills Hessian, and none is taken at a connection.

On the Poincaré sphere, the selection rule that keeps only the spins $j/2$ with icosahedral invariants goes back to Klein [16]; its first nonzero scalar eigenvalue, 168, is standard in the literature on cosmic topology [19]. We found no exact result for the Yang–Mills Hessian on a squashed three-sphere, nor at a connection that is not a Yang–Mills field. The rest of this section supplies one, and checks it against both results above that it can be compared with: the round-sphere spectrum of [18] and Laquer's index.

An exact check of the spectrum. The Berger field, its metric and the energy are all invariant under left translations of S^3, so the linearised Yang–Mills operator splits by representation, and each piece is a finite matrix that can be computed exactly. The smallest piece is the left-invariant sector: $\mathfrak{su}(2)$-valued 1-forms that are constant in the left-invariant frame, nine numbers in all. On the Poincaré sphere, a quotient by left translations, a sector of spin $j/2$ survives only as often as the binary icosahedral group has invariants in its spin-$j/2$ representation, which the Molien series $(1 + t^{30})/\big((1 - t^{12})(1 - t^{20})\big)$ counts [16]. After spin 0, the next sectors are spin 6, 10, 12 and 15. We computed the Hessian of $E = \frac{1}{2}\|F\|^2$ against the L^2 mass on the spin-0 sector exactly, and on the higher sectors as exact finite matrices of size $9(j + 1)$. As controls, spin 0 is reproduced to 8×10^{-15} by the general construction, and on the round metric without the commutator the construction returns exactly the coexact 1-form spectrum of the unit S^3 (for spin 6: 144, 196 and the gauge zeros, with the known multiplicities). On the round metric with the commutator, the construction gives -2, 1, 6, 7 and 10 in spins 0 to 2, the published fluctuation spectrum of the meron on the round S^3 [18]; at the Berger metric the same method goes beyond what is in the literature. These exact values were computed after the lattice spectra had been measured, so they are a check, not a criterion fixed in advance. Table 9 compares them with the measured eigenvalues.

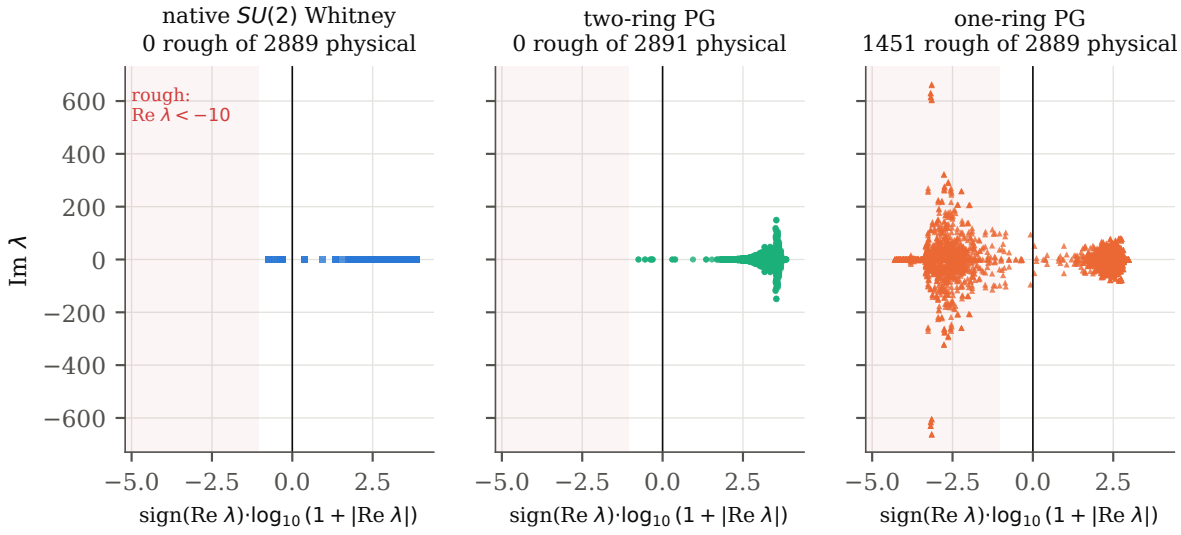


FIGURE 13. Generalized spectra of the linearised $SU(2)$ force at $n = 2$, against the native tangent mass (signed-logarithmic real axis). The shaded band is the "rough" region, $\mathrm{Re}\,\lambda < -10$.

Every value of the spin-0 sector is matched by both forces, and the real parts approach the exact values from $n = 1$ to $n = 3$. At $n = 1$ and 2 the matches were checked by eigenvector, not only by value. That check corrected one entry: at $n = 1$ the nearest real eigenvalue to $16 - 2\sqrt{73}$ (-0.536) has overlap 0.15 with the sector, and the sector mode is the complex pair at $-0.674 \pm 0.80i$. The spin-6 sector matches the measured band that begins near 50: the native force approaches it from above, and the two-ring force lies a little below it at $n = 3$. No rate is claimed from three levels. Three facts follow:

- *The field's growth directions are exactly three:* $-14/3$ (twice) and $16 - 2\sqrt{73} = -1.088$, each exactly orthogonal to the gauge directions. The background is not at rest, so gauge directions are tilted rather than zero, and the left-invariant sector contains three of them, $b_i = [M_i, \xi]$ for constant ξ, on which the Hessian takes the values 8, 12 and 12. The value 8 is a pure gauge direction, and the pairs at -2.558 and 21.891 are 40% and 60% gauge. On the gauge-orthogonal slice the sector reads $-14/3$ (twice), -1.088, $22/3$ (twice) and 33.088. All of the growth comes from the field acting on itself: without the commutator, the same sector on the Berger metric is $4/3$ (six times) and 12 (three times), all positive.
- *The other negative eigenvalues are gauge-type.* Each higher sector of spin $j/2$ carries $j + 3$ negative eigenvalues, gauge directions tilted by the background, shrinking like $1/j$ (-2.35 at spin 6, -1.73 at spin 10). Below -1 the exact count is 41; the lattice shows 9 at $n = 2$ and 13 at $n = 3$, resolving the sectors one by one. On gauge directions the lattice Jacobian differs from the Hessian by a term proportional to the nonzero gradient, so this comparison is indicative, not a criterion.
- *The physical spectrum has an exact gap.* On the gauge-orthogonal slice, the spin-0 sector ends at 33.088, spin 6 begins at $154/3 = 51.33$, and nothing else lies below 142, where spin 10 begins (spin 12 begins at $610/3$). In every higher sector the negative eigenvalues

Table 9. The exact low spectrum of the linearised $SU(2)$ Yang–Mills operator at the Berger field on the Poincaré sphere, against the real part of the measured lattice eigenvalue. Spin 0 is the left-invariant sector, spin 6 the next allowed one. At $n = 1, 2$ every spin-0 entry is identified by its eigenvector, as the eigenvalue whose eigenvector has the largest overlap (in the native mass) with the nine discrete left-invariant directions: 0.76 to 0.999, except for $16 - 2\sqrt{73}$. The lattice splits that mode: at $n = 1$ it is a complex pair (0.44 each), and at $n = 2$ it is shared between a real eigenvalue (0.50) and a complex pair at $-1.087 \pm 0.18i$ (0.34 each). The entries at $n = 3$, and the spin-6 entries, are matched by nearest value. The measured eigenvalues carry imaginary parts of up to 0.80 (the split mode at $n = 1$) and 0.42 (the 21.891 pair at $n = 2$): the lattice linearises the force in link coordinates at a field that is not at rest, where the Jacobian is not symmetric. Dashes: not resolved at that level. Gauge content of the exact eigenvectors: $^{\mathrm{g}}$ a pure gauge direction; $^{\mathrm{m}}$ a mixture (40% gauge for the negative pair, 60% for the positive one); unmarked spin-0 values are exactly orthogonal to the gauge directions. Spin-6 values are those of the full Hessian; on the gauge-orthogonal slice they are 154/3 and 59.504.

Exact	mult.	native $n = 1$	native $n = 2$	native $n = 3$	two-ring $n = 3$
spin 0					
$-14/3 = -4.6667$	2	-4.475	-4.651	-4.661	-4.651
$(29 - \sqrt{1345})/3 = -2.5581^{\mathrm{m}}$	2	-2.191	-2.504	-2.529	-2.529
$16 - 2\sqrt{73} = -1.0880$	1	$-0.674 \pm 0.80i$	-1.041	-1.088	-1.084
8^{g}	1	7.083	7.869	7.954	7.931
$(29 + \sqrt{1345})/3 = 21.891^{\mathrm{m}}$	2	20.71	21.77	21.86	21.80
$16 + 2\sqrt{73} = 33.088$	1	32.64	33.44	33.29	33.16
spin 6					
51.33	2	—	53.41	52.27	49.27
59.52	2	—	61.92	60.69	57.50

are gauge-dominated, with overlap at least 0.999 with the gauge directions; they are not exactly gauge, since off a critical point H does not preserve the slice. On this manifold and this background the low spectrum is therefore fixed by symmetry. That makes it an exactly solvable control for any future spectral observable, not a test of one.

Notation for the sectors. On $S^3 = SU(2)$ let e_0, e_1, e_2 be the left-invariant frame with $[e_a, e_b] = 2\varepsilon_{abc}e_c$, and let the Berger metric be diagonal in it with eigenvalues $g_0 = 3$, $g_1 = g_2 = 1$. Then $\hat{e}_i = e_i/\sqrt{g_i}$ is orthonormal, with coframe $\hat{\theta}^i$. The *Berger field* is the Levi-Civita (spin) connection of this metric, written as the $\mathfrak{su}(2)$-valued 1-form $A = \sum_i M_i\hat{\theta}^i$ with constant $M_i \in \mathfrak{su}(2)$, where τ_0, τ_1, τ_2 is a basis of $\mathfrak{su}(2)$ with the same brackets as the frame, τ_0 along the long axis. On the spin-$j/2$ representation V_j (dimension $j + 1$), the frame vector e_i acts by an anti-Hermitian ρ_i with $[\rho_a, \rho_b] = 2\varepsilon_{abc}\rho_c$, and $J_0 = \frac{i}{2}\rho_0$ has the weights $-\frac{j}{2}, \dots, \frac{j}{2}$. A variation in the sector is $b = \sum_i b_i\hat{\theta}^i$ with $b_i \in V_j \otimes \mathfrak{su}(2)$. H is the Hessian of $E = \frac{1}{2}\|F\|^2$ there, against the L^2 mass. The gauge map is $G\xi = (\hat{e}_i\xi + [M_i, \xi])_i$ for $\xi \in V_j \otimes \mathfrak{su}(2)$, with L^2 adjoint G^*.

The floor formula. The spectrum has a closed form at its floor in the sectors that matter. Write $f \in V_j$ for the top weight along the long axis of the Berger metric, $J_0 f = \frac{j}{2}f$, and $\hat{\theta}^i$ for the orthonormal coframe.

Proposition 10.1 (The floor formula)**.** *At the Berger metric $g = (3, 1, 1)$, in the sector of spin $j/2$, the*

field

$$b \;=\; f\,(\hat\theta^1 + i\,\hat\theta^2)\otimes(\tau_1 + i\,\tau_2)$$

is an eigenvector of the Hessian of $E = \frac{1}{2}\|F\|^2$ *against the* L^2 *mass, orthogonal to every gauge direction, with eigenvalue*

$$\lambda(j) \;=\; \frac{j(j+2)-14}{3}.$$

Derivation. At the top weight, $\hat e_0 f = -ijf/\sqrt{3}$, $J_+ f = 0$, $\rho_2 f = i\rho_1 f$ and $\|\rho_1 f\|^2 = j$, so every term of the sector Hessian involves only f and $u = \rho_1 f/\sqrt{j}$, with coefficients polynomial in j. On the squashed plane, frame and Lie indices $1,2$ in the order $(1,1),(1,2),(2,1),(2,2)$, the Hessian is

$$\Big(\tfrac{j^2}{3}+j\Big)I + \frac{1}{3}\begin{pmatrix} 14 & 2ij & -ij & 28 \\ -2ij & 5 & -19 & -ij \\ ij & -19 & 5 & 2ij \\ 28 & ij & -2ij & 14 \end{pmatrix},$$

with eigenvalues $\frac{1}{3}(j^2+2j-14)$, $\frac{1}{3}(j^2+4j-14)$ and $\frac{1}{3}\big(j^2+3j+33\pm 3\sqrt{j^2+9}\big)$. The first has eigenvector $(-1,-i,-i,1)$, which is b up to a factor. Three checks close the argument:

- its coupling to the only other top-weight direction it could reach, the τ_0 component of b_0, vanishes;
- its covariant derivative has no component along u, because the two circular factors cancel, so the Hessian sends it nowhere outside the top weight;
- its overlap with the gauge directions $\hat e_i\xi + [M_i,\xi]$ vanishes at every weight.

The computation is symbolic in j (`scripts/ricci_su2_sector_formula.py`). □

Proposition 10.1 gives an eigenvalue in every sector. The next result shows that it is the bottom of its sector in the four low sectors that survive on the Poincaré sphere, and that this makes the gap exact; for general j minimality is not proved.

Here the *gauge-orthogonal spectrum* of a sector is the spectrum of the Hessian compressed to the slice, PHP with P the L^2 projection onto $\ker G^*$: the field is not at rest, so H does not preserve the slice, and this compression is the object the theorem is about. The gap is a spectral gap of this classical linearised operator; it is not a statement about the quantum theory or a Yang–Mills mass gap.

Theorem 10.2 (The Floor Theorem)**.** *At the Berger metric* $g = (3,1,1)$*:*

1. *every gauge-orthogonal eigenvalue in the sector of spin* $j/2$ *is at least* $\big(\sqrt{j(j+2)/3}-\omega\big)^2 - 12$*, with* $\omega = \sqrt{76/3} = 5.0332\ldots$*, whenever* $\sqrt{j(j+2)/3} \ge \omega$*;*
2. *for* $j = 0, 12, 20, 24$*,* $\lambda(j)$ *is the lowest gauge-orthogonal eigenvalue of its sector, and every other eigenvalue of* $H + GG^*$ *is at least* $\lambda(j) + 1/3 - 10^{-53}$*; numerically the same picture holds for every* $j \le 120$*;*
3. *on the Poincaré sphere, the gauge-orthogonal spectrum below* 142 *consists of the spin-*0 *values* $-14/3$ *(twice),* $16 - 2\sqrt{73}$*,* $22/3$ *(twice) and* $16 + 2\sqrt{73}$*, and the spin-*6 *values from* $154/3$ *upward; spin* 10 *begins exactly at* 142*.*

Proof. On the slice $G^*b = 0$, so $\langle b, Hb\rangle = \langle b, (H + GG^*)b\rangle$: adding the gauge term changes nothing there.

The Weitzenböck form. The operator $H + GG^*$ is a Weitzenböck operator,

$$H + GG^* \;=\; \sum_i X_i^* X_i + R_0, \qquad X_i = \rho_i/\sqrt{g_i} + \Omega_i,$$

where X_i is the full covariant derivative (Ω_i acts on the Lie and frame indices) and R_0 is a constant 9×9 matrix whose smallest eigenvalue is exactly -12. The identity holds in the universal enveloping algebra $U(\mathfrak{su}(2))$, so in every sector at once: both sides are 9×9 matrices over $U(\mathfrak{su}(2))$ of filtration degree at most two, and after putting every product in Poincaré–Birkhoff–Witt order (with $[\rho_a, \rho_b] = 2\varepsilon_{abc}\rho_c$ and $\rho_i^* = -\rho_i$) every coefficient of degree one and two vanishes and the constant term is R_0, in exact arithmetic with $\sqrt{3}$ symbolic. The same computation gives $\omega^2 = \sum_i \|\Omega_i\|^2 = 76/3$ exactly (`scripts/ricci_su2_gap_certificates.py`).

Part (1). Two inequalities combine:

- since $g_0 = 3 \geq g_1 = g_2 = 1$ and $-\sum_i \rho_i^2 = j(j+2)$ on V_j, $\sum_i \|\rho_i v\|^2/g_i \geq j(j+2)\|v\|^2/3$;
- each term satisfies $\|X_i v\| \geq \big|\|\rho_i v\|/\sqrt{g_i} - \|\Omega_i v\|\big|$.

The reverse triangle inequality in $\mathbb{R}^3$ then gives $\sum_i \|X_i v\|^2 \geq (\sqrt{j(j+2)/3} - \omega)^2\|v\|^2$, with $\omega^2 = \sum_i \|\Omega_i\|^2$.

Part (2). $\lambda(j)$ is an eigenvalue of $H + GG^*$ of multiplicity at least two: the eigenvector of Proposition 10.1 at the top weight, and its image at the bottom weight under the real structure of the integer-spin representation, which commutes with $H + GG^*$ because every coefficient is real. For $j = 0, 12, 20, 24$ the rest is a rigorous count of the eigenvalues below $\mu = \lambda(j) + 1/3$. The matrix $H + GG^* - \mu$ is enclosed in interval arithmetic (mpmath's `iv` context: outward-rounded endpoints, 200 bits, $\sqrt{3}$ and every entry enclosed) and realified, which doubles every eigenvalue. The symmetric midpoint matrix M of its lower triangle lies within ρ of the exact matrix; M is factored $LDL^\top = M + \Delta M$ without pivoting, with $|\Delta M| \leq \gamma_N |L||D||L^\top|$ [15, Thm. 9.3]. Since $LDL^\top$ is congruent to D, Sylvester's law gives the exact inertia of $M + \Delta M$, and Weyl's inequality with $\varepsilon = \rho + \|\gamma_N |L||D||L^\top|\|_F$ transfers it: k negative pivots allow at most $k/2$ eigenvalues of $H + GG^*$ below $\mu - \varepsilon$. In every case D has four negative pivots, the smallest pivot has absolute value at least 0.63, and $\varepsilon < 3 \times 10^{-54}$. So at most two eigenvalues lie below $\lambda(j) + 1/3 - \varepsilon$; both copies of $\lambda(j)$ are among them, so $\lambda(j)$ is the lowest eigenvalue, twice, and every other is at least $\lambda(j) + 1/3 - \varepsilon$. At $\mu = \lambda(j) - 1/3$ no pivot is negative, a consistency check. Every pivot is recorded, and the script refuses to certify on a zero pivot, a different count, or $\varepsilon \geq 1/3$; the companion paper [9] tabulates the data. The eigenvector of Proposition 10.1 lies in the slice, so $\lambda(j)$ is the minimum there. For every $j \leq 120$ the same picture holds in double precision, with the next eigenvalue at least $\lambda(j) + 2/3$; that wider range is evidence, and the theorem does not rely on it.

Part (3). The bound of (1) exceeds 142 for every $j \geq 30$: at $j = 30$ it is 153.26 (exactly, $\sqrt{320} - \sqrt{76/3} > \sqrt{154}$, with margin 0.45 in that form). Part (2) covers the surviving sectors $j = 0, 12, 20, 24$, and the spin-0 slice was computed in exact arithmetic (`scripts/ricci_su2_sector_minimality.py`). Both certificates are in `research/su2-left-invariant-sector-20261119/gap_certificates.json`. □

For $j > 120$, minimality is not needed for the gap and is not proved. A companion paper [9] develops these results as a self-contained spectral theorem.

In one line, with $\lambda_j(\Delta_{S^3}) = j(j+2)$ the eigenvalue of the scalar Laplacian of the round unit sphere on the sector of spin $j/2$, the floor formula reads

$$\boxed{\lambda_{\text{floor}}(j) \;=\; \frac{\lambda_j(\Delta_{S^3}) - 14}{3} \;=\; \frac{j(j+2) - 14}{3}}$$

It is an eigenvalue in every sector (Proposition 10.1); the Floor Theorem shows that it is the floor of every sector below $j = 30$ that survives on the Poincaré sphere, and bounds every higher sector. On the Poincaré sphere the floors of the first two surviving sectors differ by $\lambda_1(\Delta_{S^3/2I})/3 = 168/3$.

Taken together, the field's strongest growth direction and the bottom of each surviving sector below $j = 30$ are one formula (Figure 14). At $j = 0$ it gives $-14/3$, the growth mode. The first sector that survives on the Poincaré sphere is $j = 12$, where $j(j+2) = 168$ is the first nonzero eigenvalue of the scalar Laplacian on the Poincaré sphere [19]. So the first surviving sector above the left-invariant one begins $168/3 = 56$ above the growth mode, one third of the sphere's first scalar eigenvalue. The spin-0 values $16 - 2\sqrt{73}$, $22/3$ and $16 + 2\sqrt{73}$ lie in between; the spectral gap between the two sectors runs from $16 + 2\sqrt{73} = 33.09$ to $154/3 = 51.33$. special to this squashing: at $(2,1,1)$, $(4,1,1)$ and the round metric, the bottom of a sector is not of this form.

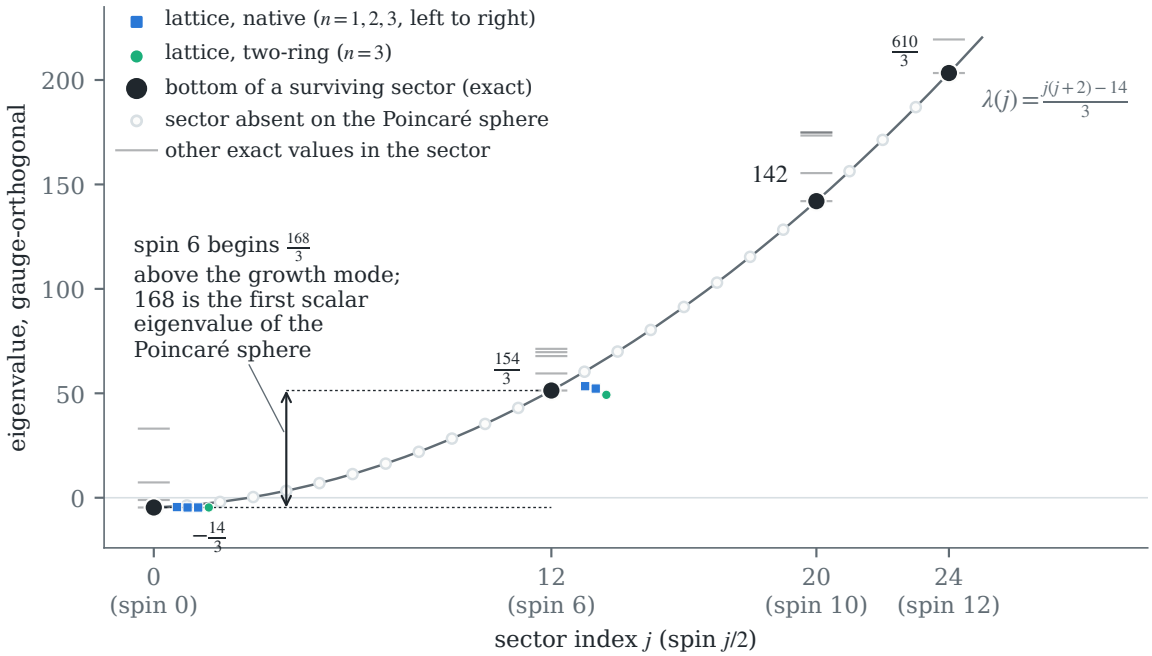


Figure 14. The physical spectrum of the linearised $SU(2)$ operator at the Berger field, sector by sector. Curve: $\lambda(j) = (j(j+2) - 14)/3$ (Proposition 10.1). Filled: the bottom of each sector that survives on the Poincaré sphere; hollow: sectors the icosahedral group forbids; grey ticks: the other exact gauge-orthogonal values of a surviving sector. Squares and circles: measured lattice eigenvalues near the bottom of spins 0 and 6 (native $n = 1, 2, 3$ and two-ring $n = 3$; spin 0 identified by eigenvector at $n \le 2$, spin 6 by value). Generated by `make_figure_spectrum_ladder.py`.

Numerical range. Eigenvalues alone do not rule out transient non-normal growth, so we also measured the numerical range $W(\hat{G})$ of the force whitened by the native mass, $\hat{G} = L^{-1}KL^{-\top}$

Table 10. Numerical range of the linearised $SU(2)$ force, whitened by the native mass, on the Berger field. "Left edge" is $\min \operatorname{Re} W(\hat{G})$; $\eta = \|\hat{G}\hat{G}^\top - \hat{G}^\top\hat{G}\|/\|\hat{G}\|^2$ measures non-normality.

Force	n	left edge	slowest $\operatorname{Re}\lambda$	numerical radius	η
Native $SU(2)$ Whitney	1	−4.634	−4.475	1,913	0.0014
Two-ring Petrov–Galerkin	1	**−4.563**	−4.394	1,652	0.134
removal: one ring	1	−8,774	−8,621	8,774	0.116
Native $SU(2)$ Whitney	2	−4.682	−4.650	7,339	0.0003
Two-ring Petrov–Galerkin	2	**−4.664**	−4.632	7,027	0.142
removal: one ring	2	−19,993	−19,829	19,993	0.176
Native $SU(2)$ Whitney	3	−4.676	−4.661	—	—
Two-ring Petrov–Galerkin	3	**−4.666**	−4.651	—	—
removal: one ring	3	−45,922	−45,543	—	—

Dashes: not measured at $n = 3$.

with $B = LL^\top$. Its left edge $\min \operatorname{Re} W(\hat{G}) = -m$ bounds the linearised flow, $\|e^{-t\hat{G}}\| \le e^{mt}$, with constant one. The criterion, fixed together with the first-level results, asked for the two-ring left edge to be at least twice the native one and at least −10. Table 10 gives the result.

The two-ring force passes at all three levels. Its left edge lies 0.17 to the left of its slowest eigenvalue at $n = 1$, 0.03 at $n = 2$ and 0.015 at $n = 3$, and matches the native operator's, so the only growth is the field's own motion. For the operator frozen at the initial Berger field, with the mass held fixed, the envelope over the window $[0, 0.02]$ of the time-integration experiment is $e^{4.664\times 0.02} \approx 1.098$ for both forces. This is a bound for the frozen linearization only. Along the actual flow the mass $B(t)$ changes with the metric, and even the linear energy identity gains a term, $\frac{d}{dt}\frac{1}{2}\xi^\top B\xi = -\xi^\top \operatorname{sym}(K)\xi + \frac{1}{2}\xi^\top \dot{B}\xi$; linearizing along the moving trajectory adds others. The envelope therefore does not certify the nonlinear moving-metric run; the agreement measured over that window (Section 10.3) is a separate, empirical result. The one ring fails, as the removal requires, with a left edge of −19,993 at $n = 2$ and −45,922 at $n = 3$. The two-ring operator is non-normal (η about 0.14, against 3×10^{-4} native), but its numerical radius stays within 1% of its norm, and at $n = 2$ the left edge sits on its slowest eigenvalue: the lopsidedness lives in fast, strongly damped modes, not where growth would begin. Because the Berger field moves, the symmetric part is indefinite for every operator, so the dissipativity certificate of Proposition 7.1 does not apply here, and the frozen envelope above is the honest bound. For a time stepper, Crouzeix–Palencia [8] bounds k steps of any rational rule r by $(1+\sqrt{2})\sup_W |r(hz)|^k$. The explicit-Euler limit it implies at $n = 2$ is $h \approx 2.85 \times 10^{-4}$; the flow at $n = 2$ used the fourth-order rule with $h = 3.125 \times 10^{-4}$, inside that rule's larger limit of about 3.8×10^{-4}.

10.3 The flow in time

Time integration. The last test lets the lattice move, under a lattice Yang–Mills gradient flow of the kind used in lattice QCD [22], here with a moving metric. Starting from the exact links at $t = 0$, the links evolve under $B(U, g(t))\,\dot{\xi} = -a(U; g(t))$, where the metric $g(t)$ follows the Ricci flow of the Berger metric, by a fourth-order Runge–Kutta–Munthe-Kaas step [23, 14] that keeps every link on $SU(2)$. At $t = 0.02$ the links are compared with the exact links of the flowed metric: E is the mass-weighted distance $\sqrt{\delta^\top B\delta/\mathrm{Vol}}$ between them, D is the same distance travelled by

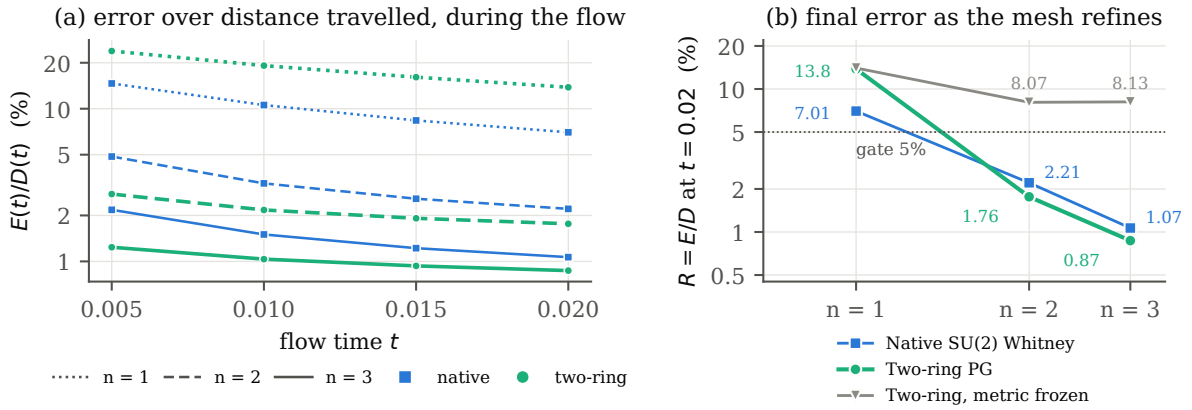


FIGURE 15. The $SU(2)$ flow in time. (a) The error over the distance travelled, $E(t)/D(t)$, at the four sample times of each run: native (squares) and two-ring (circles) at $n = 1$ (dotted), 2 (dashed) and 3 (solid). (b) The final value R against the mesh. The two-ring force overtakes native at $n = 2$ and stays ahead at $n = 3$; the frozen-metric control stops improving after $n = 2$. Every plotted value is in Table 11 or its recorded outputs; $n \leq 2$ uses RKMK4 and $n = 3$ the Cayley step.

the exact field, and $R = E/D$ is the error over the distance actually moved. The criteria, fixed before measurement, asked for $R < 5\%$ at the finest level with an observed order of at least 0.5, and for halving the step, or changing the quadrature, to move E by less than a tenth. Two negative controls were corrected after the $n = 1$ results, beside the unchanged original criteria. As first written, both controls asked for $R \geq 10\%$ at every level: the frozen-metric control fails that at $n = 2$ (8.07%) and $n = 3$ (8.13%), and the native control fails it at every level. The corrected controls ask instead whether the frozen run's gap to the moving run is resolvable and fails to converge, and whether the native force is overtaken and stalls; both versions are reported in the results document. The runs were first made by a Python replica of the programme's flow test. The engine's Cayley flow reproduces every Cayley run at $n \leq 3$, to 1.4×10^{-13} in E with identical Newton iteration counts; the fourth-order runs at $n \leq 2$ remain script runs. Table 11 and Figure 15 give the result.

The two-ring force passes. At $n = 2$ it ends 1.76% off, and E falls by a factor of 7.65 from $n = 1$: an observed order of 2.94 with the mesh size halving per level. With two levels that is one ratio, not a measured rate. Halving the step changes E by 2.6×10^{-10} relative at $n = 2$, and $q4$ against $q6$ by 4.6×10^{-8} at $n = 1$, so the error is the lattice's, not the stepper's or the quadrature's. The crossover of the instantaneous force reappears in time: native is better at $n = 1$, two rings are better at $n = 2$, in the raw and the gauge-invariant measure alike. The frozen-metric control, the programme's known failure, becomes informative at $n = 2$. Its gap to the moving run, $\Delta \geq |E_f - E_m| = 0.032$, exceeds twice the moving run's own error, and its error grows faster in each quarter of the window (0.006, 0.015, 0.027, 0.041) while the moving run's levels off (0.004, 0.006, 0.008, 0.009). The lattice sees the moving background, and the two-ring force follows it. The one ring fails at the two-ring step, as its spectrum predicts: its error reaches $E = 8.8$ in the first quarter of the window and 12.7 at the end, 25 times the distance the exact field travels.

One prediction failed, and the reason was then measured. The corrected native control expected the native error to stall in time, as its instantaneous error stalls near 25%; instead it

Table 11. Time integration of the $SU(2)$ force on the Ricci-flowed Berger metric over $t \in [0, 0.02]$, $q6$. $R = E/D$ is the final error over the distance the exact field travelled; the last column measures it on gauge-invariant data. Rows at $n \leq 2$ use the explicit RKMK4 step; rows at $n = 3$ use the Cayley step with 8 implicit steps.

Force	n	steps	$E(0.02)$	R (%)	gauge-invariant R (%)
Native $SU(2)$ Whitney	1	32	0.0349	7.01	3.10
Two-ring Petrov–Galerkin	1	32	0.0689	13.81	10.80
control: metric frozen at $g(0)$	1	32	0.0699	14.02	9.39
removal: one ring	1	32	12.7	2541	–
Native $SU(2)$ Whitney	2	64	0.0113	2.21	2.10
Two-ring Petrov–Galerkin	2	64	0.0090	**1.76**	1.82
control: metric frozen at $g(0)$	2	64	0.0412	8.07	5.70
Native $SU(2)$ Whitney	3	8	0.0055	1.07	1.22
Two-ring Petrov–Galerkin	3	8	0.0045	**0.87**	0.77
control: metric frozen at $g(0)$	3	8	0.0417	8.13	5.30

falls, with observed order 1.63. To first order the trajectory error obeys $y' = -\hat{G}y + \hat{r}(t)$, where $\hat{r} = L^\top r$ is the whitened residual, the lattice velocity at the exact links minus the exact velocity. Integrated exactly from residuals sampled along the exact path, this linear filter predicts every error in Table 11, for both forces at both levels, to within 6.1%. The split of the residual between slow directions of $\hat{G}$ (rate below $1/T = 50$) and fast ones then explains the failure. At $n = 2$ only 0.85% of the native residual lies in slow directions, against 7.6% for two rings: the native force's 25% error is real, but it is stiff, and the dynamics removes it within the window. Refinement makes it stiffer, so the fraction of the undamped error that survives falls from 0.29 to 0.095 for native between $n = 1$ and $n = 2$. The stall prediction was therefore about the wrong quantity. Over a window this short, what carries into the trajectory is the slow part of the error, 2.28% for native against 2.05% for two rings at $n = 2$, falling with observed orders 1.17 and 2.62. On a longer window more directions count as slow, and the native force's stiff error begins to matter: on $[0, 0.1]$ its order between $n = 1$ and 2 falls from 1.63 to 1.32 (Limitations).

The third level. At $n = 3$ (11,358 unknowns) the explicit step needs too many force evaluations, so the flow is stepped by the Cayley map of the companion paper [10]: $U_{k+1} = U_k \operatorname{cay}(\Omega)$, with Ω solving $\Omega = h\,\xi(U_k \operatorname{cay}(\Omega/2), t_k + h/2)$, the Lie form of the implicit midpoint rule. Its linearisation is Crank–Nicolson, which is A-stable: stable for any step on the dissipative modes. On the field's few growing modes ($\dot{a} = -\lambda a$ with $\lambda < 0$) its factor $(1 - h\lambda/2)/(1 + h\lambda/2)$ is singular at $h = -2/\lambda$, so the step must stay well below that. For the strongest growing mode, $\lambda = -4.66$, that is $h = 0.43$; the step used, $h = 0.0025$ (eight steps over the window), is about 170 times smaller. Every update stays exactly on $SU(2)$, which is group membership, not a statement about accuracy or stability. Each step is a quasi-Newton solve with one iteration matrix factored per run. The step was checked before use. Its distance from the exponential map falls as $|v|^3/12$, the phase lag of the companion paper. Its observed time order is 1.99. At $n = 1$ and $n = 2$, with 8 steps, it reproduces the RKMK4 errors of Table 11 to within 1.5% for both forces, using between 0.17 and 0.36 of the force evaluations. The rule fixed in advance takes the smallest step count passing at both levels, which is 8. Four steps fail at $n = 2$, and a later experiment, with

its predictions committed first, traced the failure to stiff modes. Crank–Nicolson is A-stable but not L-stable: it flips the sign of a stiff component instead of damping it. Replacing the first step by an L-stable second-order rule (two-stage SDIRK, $\gamma = 1 - 1/\sqrt{2}$) cuts the first-step error at $n = 2$ from 22.4% to 8.9% (native) and from 13.1% to 6.1% (two rings), and changes the final error of the eight-step runs by at most 0.05% of itself. A first-order damping step (backward Euler) did not separate the two effects: its own error was larger than the one it removed. The reported runs use no start-up. The $n = 3$ runs were started with both 8 and 16 steps before that verdict was committed, so the rule could not depend on them; the two differ by 0.75%. One criterion was relaxed after a failure, and the record keeps both: the first attempt asked Newton for a relative increment of 10^{-10}, and every native run at $n = 1$ failed at its first step, because the force itself, with a finite-difference defect derivative, was only computable to about 3.5×10^{-10}. The tolerance was amended to 10^{-8}, about thirty times above that floor and far below every comparison the criteria make; the failed runs are kept. The defect derivative was then put in closed form, which lowers the force's noise to about 10^{-13}; every reported run uses both changes.

At $n = 3$ every step converged, in at most eight iterations, and every criterion set before the run holds except the corrected native control's stall, discussed below. The two-ring error is 0.87% of the distance moved, falling by a factor of 2.03 from the $n = 2$ Cayley run ($E = 0.00904$; against the RKMK4 value of the table, 2.02). The mesh size shrinks by 3/2 between these levels (the link count grows by 1.5^3), so the observed order is 1.74. Orders here use the nominal mesh ratios 2 and 3/2; with the measured largest edge, as in Tables 6–7, they would be slightly higher (3.11 and 1.78). Halving the step moves E by 0.75%. The gap between the moving and frozen-metric runs, now measured exactly from saved final states, is $\Delta = 0.042$, 9.4 times the moving run's error. The frozen run's error does not converge at all (0.0412 to 0.0417, order −0.03). The two-ring force stays ahead of native, at 0.82 of its error (0.80 at $n = 2$). The native error still falls (order 1.80), so the stall asked for by the corrected control is again absent, as the filter model explains. That model also forecast, without a criterion, that the two-ring error would fall faster than the native one. At this refinement it did not: the two orders are 1.74 and 1.80.

R divides by the whole distance the exact field travels, and most of that motion is the Yang–Mills flow at the initial metric, which the frozen run reproduces too. The part of the motion that is due to the metric moving is Δ. Measured against it, the two-ring error is $0.00446/0.0420 = 10.6\%$ at $n = 3$; at $n = 2$, where Δ is only bounded ($0.032 \le \Delta \le 0.050$), it lies between 18% and 28%. So the response to the moving metric is tracked to about a tenth at the finest level, and that fraction is falling; R alone would suggest about a hundredth.

Cost, and a second check of every flow result. The force was first written cell by cell; at $n = 3$ one two-ring evaluation took 89 s. It was then rewritten to do the same arithmetic for every cell at once, under criteria fixed before any of the new code ran. They make the cell-by-cell code the reference and test the rewrite for parity, at $n = 1, 2, 3$, for both forces, at two link states and two metrics. The native gradient comes out identical to the last bit; the two-ring gradient agrees to 5.4×10^{-14}, the mass to better than 10^{-15} and the velocity to 1.9×10^{-13}. An amendment, written after the first timings and before the change, caches the metric integrals, which depend on g alone and are the same for every Newton iteration of a step; it changes rounding only, and every parity test was run again with it. The Cayley runs were then repeated on the new code: the four at $n = 1, 2$ with 8 steps, before and after the amendment, and the three at $n = 3$. Each reproduces

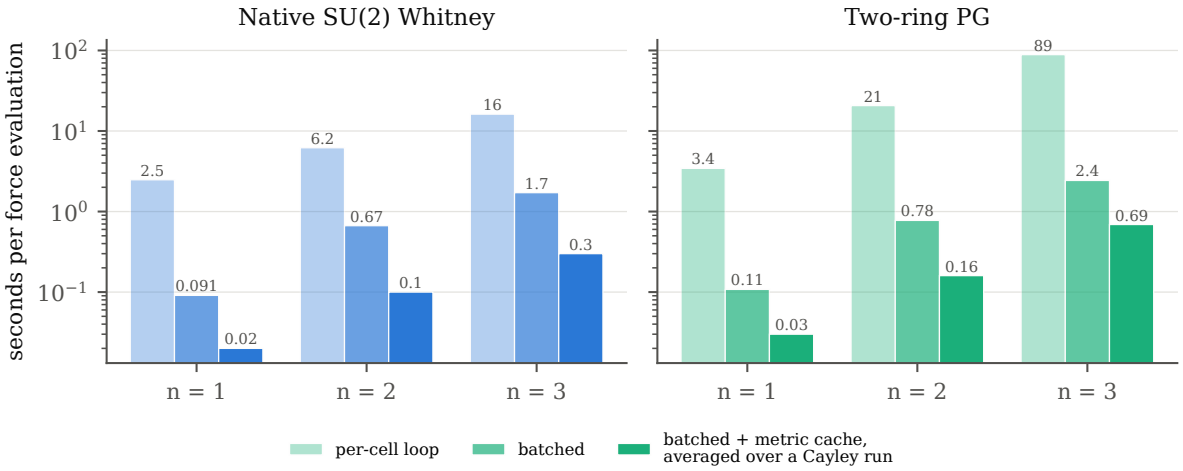


Figure 16. Seconds per force evaluation (gradient, mass and the mass solve) for the cell-by-cell code, the batched rewrite, and the batched rewrite with the metric integrals cached, averaged over a whole eight-step Cayley run (which recomputes the metric once per step). One process on one laptop CPU, nothing else running for the first two series. The three versions give the same numbers to rounding.

its committed final error to 5×10^{-14} or better (the native run at $n = 3$ to the last digit), with the same number of Newton iterations at every step. A full two-ring run at $n = 3$ now takes 133 s, of which 74 s is setup; Figure 16 gives the cost per evaluation. These are timings of research scripts on one laptop, and they matter here for one reason: the flow can now be run again by anyone, and longer windows and finer levels are within reach.

Scope. The covariant recovery and the non-Abelian force are engine verbs that reproduce the research scripts to 10^{-13}; the fourth-order flow runs at $n \le 2$ and the eigenvalue solves remain scripts. The force result is the instantaneous Criterion C on one metric, with the spectrum of Table 8. The time integration is a research replica on one metric and one window, at $n \le 3$, and stability and the numerical range are measured at $n \le 3$. A reciprocal version remains the open problem of Section 11.1.

11 Discussion

The central observation is simple to state. *The Petrov–Galerkin force was never inconsistent; it was ill-conditioned.* The one-ring recovery is the minimal-information choice, and it has no choice but to read every rough flux pattern as a linear field. Once the recovery is given more information than it strictly needs, rough patterns are averaged rather than interpolated. Exactness on the model space is untouched, so consistency survives. Stability follows without any damping parameter, filter or regularization: the only change is a neighbourhood radius, and for the linear fit only radii one and two were run. The same count, taken one order higher, gives the hybrid reading of Section 9: three rings, read to second order, raise the observed order of the force from about 0.8 to about 1.85.

11.1 Fitting to a guide

The introduction described the method as fitting a shape to a guide. Each part of that description is a property of the force that can be checked, and Table 12 checks it for every force law in this paper.

Table 12. Fitting to a guide, scored for each force law. "Fixed guide": tested against the original Whitney variations. "Stable": no physical mode with negative real part (S3); for two rings the stronger certificate $\mu > 0$ holds at $n = 1, 2$. "Reciprocal": symmetric, so the force is the gradient of an energy.

Force law	Fixed guide	Stable	Accurate (Criterion C)	Reciprocal
Whitney	yes	yes	no (26% floor)	yes
Variational recovery	no	yes	no (diverges)	yes
Symmetrized Petrov–Galerkin	partly	no	no (grows)	yes
Petrov–Galerkin, one ring	yes	no	yes	no
Petrov–Galerkin, two rings	**yes**	**yes**	**yes**	**no**

One pattern in the table is worth stating plainly. *Every accurate force keeps the guide fixed, and every force that moves the guide is inaccurate.* A fixed guide is not sufficient, as Whitney shows, but on these meshes it has been necessary.

Turning into register. A neighbouring cell stored on another sheet of the cover is not compared with K as it lies. It is first turned onto K's sheet by the exact deck transformation, and only then may its fluxes act. Both recoveries do this, and without it the fit's condition number reaches 2.4×10^{14}. Turning into register is what makes the comparison meaningful at all; it is not what restores stability.

Refusing what does not fit. The recovery refuses rather than forces. If a transported neighbour's shared vertices do not coincide with K's to 10^{-12}, or if the flux map is rank-deficient or worse conditioned than 10^{12}, no curvature is produced. None of these refusals was triggered on the meshes reported here; they are the method's way of saying "this piece does not belong here".

Edge to spare. This is the one change. In the one-ring fit every flux is a hard constraint, so a rough flux has nowhere to go but into the curvature. In the two-ring fit only K's own faces are hard; the rest are fitted in least squares, and a rough flux is outvoted rather than obeyed. The median amplification falls from 125–252 to a mesh-uniform 30.5–35.2 (Figure 5), and stability follows.

Reciprocity. The missing property is the one a guide does not have. Our force is not symmetric, and Section 8 shows why the obvious symmetrizations fail: they make the pull reciprocal by letting the guide move, and the accuracy goes with it. The condition $\mathcal{R}^\top H c_F \approx j^\top H c_F$ says what a reciprocal version must achieve: on smooth fields, the piece and the guide must agree from both sides. An earlier probe in the accompanying reports placed a symmetric projection layer between the two. It cut the Whitney floor from 26% to about 6% but still floored. A force that is accurate, stable and reciprocal is the open problem this paper leaves.

11.2 Where the ideas came from

This paper describes its method in everyday language, a guide, a magnet, a rim chart, a child's reach, because that is where the method started. Each picture came first as a physical intuition and was made precise afterwards; the mathematics then had to stand on its own.

The magnet. The idea that a piece should lock to its guide without being forced, acting through a layer from both sides, is the magnet idea of Section 11.1. Its literature and the two probes it led to, including the symmetric projection layer that cut the floor to about 6%, are recorded in `docs/RICCI_PROJECTION_LAYER_RESULTS.md`.

Chord, arc and tangent. Building the $SU(2)$ force put two readings of one turn side by side, the arc $\log U$ in the recovery tests and the chord $U - 1$ in the force. Asking what separates them, and what the third, outside path means, became the companion paper [10]. It finds the chord and the outside path missing the arc in a fixed ratio of two to one, in length, in energy and in time stepping. For this paper it settles two points: the chord–arc gap on our meshes is too small to matter, measured in the force itself (Section 10.2), and the Cayley map is the integrator that carried the flow to $n = 3$, where the explicit step became unaffordable.

What grew from this work. Two research notes in the repository record further ideas that came out of these results, each with every link tagged as structural, a resemblance only, or open. The first connects the numerical range of Proposition 7.1, the author's work on the Crouzeix conjecture and on graph folds to the roughness that one ring amplifies. The second continues from the chord–arc–tangent picture to the curvature of a single unit. Neither is used as evidence in this paper.

Limitations.

- The Abelian results concern a single control on a single round quotient.
- The covariant recovery of Section 10.1 is tested on two exact left-invariant fields with sampled cells; it is an engine verb, equal to the research script to 1.1×10^{-13}. The $SU(2)$ force passes Criterion C on the Berger metric only, and its stability and numerical range are measured at $n \le 3$; the imaginary extent of its spectrum grows with refinement (104, 150, 418), and whether it keeps growing at $n = 5$ is not measured.
- The time integration covers $n \le 3$ over $t \in [0, 0.02]$. Its Cayley runs are reproduced by the engine; the fourth-order runs at $n \le 2$ are scripts; and none of it is yet the programme's release flow test on a merged build. The engine flow stops at $n = 3$ until a sparse preconditioner is built. The predicted stall of the native force did not appear at $n = 2$ or $n = 3$; a first-order filter model explains why. The two-ring force keeps its lead over native (at 0.80, then 0.82 of its error) but does not widen it. R measures the error against all the motion; against the part due to the moving metric alone, the two-ring error is about 10.6% at $n = 3$. On a five-times-longer window, $[0, 0.1]$ at $n \le 2$, the two-ring force stays ahead of native at every sample time, but its lead shrinks from 20% to 2%. By $t = 0.1$ the two error vectors point the same way (cosine 0.988 in the mass norm at $n = 2$), so the two flows share an error direction. Nearly parallel error vectors show a shared direction, not its cause: the dynamics can map different force errors into the same slow direction, and shared target, mass, time or geometric errors can align them too. Our hypothesis is that

the shared Whitney mass and the target it defines dominate; a mass built from the same reading as the force would test it, and that test has not been run.

- The S3 spectrum is decomposed at $n \leq 3$ and the numerical range at $n \leq 2$; the finest level is not decomposed.
- For the two-ring force, the observed orders 0.73 and 0.80 before the last refinement are only moderately above the threshold. The hybrid reading (Section 9) gives 1.88 and 1.83 on the same meshes, but it is Abelian only, and its order exceeds what the bound (5) explains.
- The force is not the gradient of an energy. Whether a dissipative, consistent, non-variational force suffices for the programme's purposes is a modelling decision, not a numerical one.
- No Ricci-flow claim is made. The targets are the instantaneous Yang–Mills force and, for $SU(2)$, one short window of the flow on one metric; the results are a prerequisite for a lattice correspondence between geometric flow and gauge dynamics, not an instance of one.

12 Conclusion

The Whitney floor of the instantaneous Yang–Mills force is a representation defect, and it is removed by reading curvature from a neighbourhood rather than from a single cell. The minimal reading is exactly determined and therefore unstable; a two-ring reading that holds each cell's own faces exactly and lets the rest vote is consistent and, at every level computed, stable, and a three-ring reading to second order raises the observed order to about two. The same reading carries over to $SU(2)$, with the own faces held up to a non-Abelian residual that falls like h^2, where it converges on a Berger metric that defeats the native force and follows the Ricci-flowed metric in time. At that Berger field the low spectrum of the linearised operator is exact where it matters: $(j(j+2)-14)/3$ is an eigenvalue in every sector and, by certified computation, the bottom of each low sector that survives on the Poincaré sphere, where the gap is set by the sphere's first scalar eigenvalue (the Floor Theorem; see the companion paper [9]). Two problems remain open: a force that is also the gradient of an energy, which Lemma 8.1 shows cannot come from exact moment conditions on a per-cell reading, and the mass that a flow built on such a force would need.

Funding and competing interests

This work received no funding. The author declares no competing interests.

Code and data availability

All code, frozen inputs and results are available in the GIGI repository, `https://github.com/nurdymuny/gigi`; this version, with its source and the code and data behind it, is archived on Zenodo under the concept DOI `https://doi.org/10.5281/zenodo.24328085`. Every number and figure in this paper is regenerated from the frozen, hash-checked inputs of the GIGI Ricci checkpoint, by the following programs.

- Two-ring recovery: `scripts/ricci_pg_tworing.py`; its focused tests: `scripts/test_ricci_pg_tworing.py`.

- Criterion C and stability tests: `scripts/ricci_pg_tworing_gate.py`.
- Hybrid three-ring reading: `scripts/ricci_pg_hybrid.py`; its Criterion C, S3 and numerical-range tests: `scripts/ricci_pg_hybrid_gate.py`; design document `docs/RICCI_PG_HYBRID_SPEC.md`; recorded output in `research/ricci-pg-hybrid-20261119/`; the exploratory probe's recorded outputs in `research/variational-baseprobe-20261119/`.
- Figures, data and tables: `make_data.py`, `numerical_range.py`, `make_figures.py` and `make_tables.py` in the paper's directory; `make_polished_plates.py` renders Figures 2–4 from the saved drawing coordinates and measurement recorded outputs. Its focused rendering checks are in `test_polished_plates.py`.
- Engine verb: `src/gauge/curvature_recovery.rs`, with parity and live-run recorded outputs in `research/recover-curvature-parity-20261117/`.
- Covariant $SU(2)$ recovery: `scripts/ricci_su2_recovery.py`; its tests: `scripts/test_ricci_su2_recovery.py` and `scripts/ricci_su2_g2_probe.py`; recorded outputs in `research/recover-curvature-su2-20261117/`.
- $SU(2)$ force: `scripts/ricci_su2_pg_force.py`; recorded outputs in `research/su2-pg-force-20261117/`.
- $SU(2)$ stability: `scripts/ricci_su2_pg_stability.py`; recorded outputs, with every eigenvalue, in `research/su2-pg-stability-20261117/`; at $n = 3$, `scripts/ricci_su2_pg_spectrum.py` with recorded outputs in `research/su2-spectrum-n3-20261119/`; the exact-derivative rebuild and the sparse certificate, `scripts/ricci_su2_exact_stability.py` with recorded outputs in `research/su2-exact-stability-20261119/`; the exact sectors, `scripts/ricci_su2_left_invariant_sector.py` and `scripts/ricci_su2_spin_sector.py`, their gauge content `scripts/ricci_su2_sector_gauge_check.py`, the eigenvector identification `scripts/ricci_su2_sector_eigvec_check.py`, with recorded outputs in `research/su2-left-invariant-sector-20261119/`.
- $SU(2)$ engine verbs, the design document `docs/YANG_MILLS_ENGINE_SPEC.md`; the code `src/gauge/curvature_recovery.rs`, `src/gauge/ricci_pg_force.rs` and `src/gauge/ricci_link_flow.rs`; the parity, live and removal harnesses in `e2e/probes/`; recorded outputs in `research/ym1-parity-20261119/` through `research/ym7-spectrum-20261119/`.
- Cayley start-up experiments: `scripts/ricci_su2_pg_flow_cayley.py (-startup)`; recorded outputs in `research/su2-startup-20261119/` and `research/su2-startup-sdirk-20261119/`.

- $SU(2)$ numerical range:
  `scripts/ricci_su2_pg_numrange.py`;
  recorded outputs in `research/su2-numrange-20261118/`.
- $SU(2)$ flow: `scripts/ricci_su2_pg_flow.py` (RKMK4), `scripts/ricci_su2_pg_flow_filter.py` (slow and fast residual) and `scripts/ricci_su2_pg_flow_cayley.py` (Cayley step); recorded outputs in `research/su2-flow-20261118/`, `research/su2-filter-20261118/` and `research/su2-cayley-20261118/`.
- Batched $SU(2)$ force and its parity tests: `scripts/ricci_su2_batched.py`; recorded outputs, including the repeated Cayley runs, in `research/su2-batched-20261119/`. Figures 15 and 16: `make_figures_su2flow.py`, which writes `data/su2_flow.json` and `data/su2_cost.json` from those recorded outputs.
- The native $SU(2)$ baseline of Table 7: `output/ricci-shape-regular-20261117/berger/diagnostic.json`.

References

[1] D. N. Arnold, R. S. Falk, J. Guzmán, G. Tsogtgerel, On the consistency of the combinatorial codifferential, *Trans. Amer. Math. Soc.* 366 (2014) 5487–5502.

[2] D. N. Arnold, R. S. Falk, R. Winther, Finite element exterior calculus, homological techniques, and applications, *Acta Numerica* 15 (2006) 1–155.

[3] D. N. Arnold, R. S. Falk, R. Winther, Finite element exterior calculus: from Hodge theory to numerical stability, *Bull. Amer. Math. Soc.* 47 (2010) 281–354.

[4] C. Bär, The Dirac operator on homogeneous spaces and its spectrum on 3-dimensional lens spaces, *Arch. Math.* 59 (1992) 65–79.

[5] A. Bossavit, Whitney forms: a class of finite elements for three-dimensional computations in electromagnetism, *IEE Proc. A* 135 (1988) 493–500.

[6] J.-P. Bourguignon, H. B. Lawson, Stability and isolation phenomena for Yang–Mills fields, *Comm. Math. Phys.* 79 (1981) 189–230.

[7] A. Cohen, M. A. Davenport, D. Leviatan, On the stability and accuracy of least squares approximations, *Found. Comput. Math.* 13 (2013) 819–834.

[8] M. Crouzeix, C. Palencia, The numerical range is a $(1+\sqrt{2})$-spectral set, *SIAM J. Matrix Anal. Appl.* 38 (2017) 649–655.

[9] B. R. Davis, The Floor Theorem for the Yang–Mills Hessian at the Berger field: the low spectrum on the Poincaré homology sphere, 2026, `https://doi.org/10.5281/zenodo.24386614`.

[10] B. R. Davis, Three ways across a turn: chord, arc and tangent, Zenodo, 2026, `https://doi.org/10.5281/zenodo.24375338`.

[11] E. H. Georgoulis, T. Pryer, Recovered finite element methods, *Comput. Methods Appl. Mech. Engrg.* 332 (2018) 303–324.

[12] J. Henkel, E. A. Lauret, Hodge Laplacian on 1-forms of homogeneous 3-spheres, 2026, arXiv:2605.05406.

[13] R. Hiptmair, Finite elements in computational electromagnetism, *Acta Numerica* 11 (2002) 237–339.

[14] A. Iserles, H. Z. Munthe-Kaas, S. P. Nørsett, A. Zanna, Lie-group methods, *Acta Numer.* 9 (2000) 215–365.

[15] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM, Philadelphia, 2002.

[16] F. Klein, *Vorlesungen über das Ikosaeder und die Auflösung der Gleichungen vom fünften Grade*, Teubner, Leipzig, 1884.

[17] J. Kling, D. Schueth, On the Dirac spectrum of homogeneous 3-spheres, arXiv:2204.12990.

[18] S. Kumar, O. Lechtenfeld, G. Picanço Costa, Instability of cosmic Yang–Mills fields, *Nucl. Phys. B* 973 (2021) 115583, arXiv:2102.08401.

[19] M. Lachièze-Rey, Eigenmodes of dodecahedral space, 2004, arXiv:gr-qc/0402035.

[20] H. T. Laquer, Stability properties of the Yang–Mills functional near the canonical connection, *Michigan Math. J.* 31 (1984) 139–159.

[21] E. A. Lauret, The smallest Laplace eigenvalue of homogeneous 3-spheres, arXiv:1801.04259.

[22] M. Lüscher, Properties and uses of the Wilson flow in lattice QCD, *J. High Energy Phys.* 08 (2010) 071.

[23] H. Munthe-Kaas, Runge–Kutta methods on Lie groups, *BIT* 38 (1998) 92–111.

[24] Y. Ohnita, Y. L. Pan, On weakly stable Yang–Mills fields over positively pinched manifolds and certain symmetric spaces, preprint MPI/89-49, Max-Planck-Institut für Mathematik, Bonn, 1989.

[25] H. Whitney, *Geometric Integration Theory*, Princeton University Press, 1957.

[26] K. G. Wilson, Confinement of quarks, *Phys. Rev. D* 10 (1974) 2445–2459.

[27] O. C. Zienkiewicz, J. Z. Zhu, The superconvergent patch recovery and a posteriori error estimates, *Int. J. Numer. Methods Eng.* 33 (1992) 1331–1364.