hexagon:2611.00824v1

SMOOTH ISOMETRIC EMBEDDINGS OF CLOSED MANIFOLDS IN $\mathbb{R}^{s_n+5}$ IN LARGE DIMENSION

WENCHANG YANG AND YI WANG

Abstract. We prove that for $n \geq 10^{62} + 4$, every closed smooth Riemannian n-manifold admits a smooth isometric embedding into $\mathbb{R}^{s_n+5}$, where $s_n = n(n+1)/2$.

Contents

1. Introduction 1
2. Conventions and the compact-model reduction 3
3. Smooth metric correction at a phase-regular map 5
4. Removal of double points with the metric fixed 10
5. From a compact model to the prescribed metric 14
6. The four-section construction at dimension $s_n + 5$ 16
7. The reference map and the boundary class 21
8. The active map on an exceptional link 23
9. Flat matching on the doubled exterior 36
10. Boundary data when the rank-five normal splits 47
11. Integration of the mixed coefficients and exact metric repair 54
12. Characteristic directions of the active operator 61
13. A simultaneous active differential identity 65
14. Relative filling of the hyperplane metric 69
15. A single set of coefficients on the doubled exterior 72
16. Completing the surviving line 75
17. Global assembly 83
Acknowledgements 86
References 86

1. Introduction

Let (M, g) be a smooth Riemannian manifold of dimension n. An isometric embedding of (M, g) into $\mathbb{R}^N$ is a smooth embedding $F : M \to \mathbb{R}^N$ satisfying

$$\langle \mathrm{d}F_x(u), \mathrm{d}F_x(v) \rangle = g_x(u, v) \qquad (x \in M,\ u, v \in T_xM). \tag{1.1}$$

2020 *Mathematics Subject Classification.* 53C42, 53B20, 58J60, 35F50.

Key words and phrases. global isometric embedding, closed Riemannian manifold, phase regularity, Hopf spheres, h-principle, differential right inverse.

In local coordinates this is a system of

$$s_n = \frac{n(n+1)}{2}$$

equations, and $N = s_n$ is the smallest dimension in which one can expect solvability for every metric. We call $N - s_n$ the *surplus*. Nash [N56] proved that every closed Riemannian manifold embeds isometrically in a Euclidean space, and Günther [Gü89, Gü91] obtained the surplus $\max\{2n, n+5\}$. In [WY26] we replaced it by the constant 28. The main result of this paper lowers the surplus to five in large dimension.

Theorem 1.1. *Let $n \geq 10^{62} + 4$. Every closed smooth Riemannian manifold (M^n, g) admits a smooth isometric embedding into $\mathbb{R}^{s_n+5}$. The image can be contained in any prescribed open ball, and the embedding can be chosen phase regular: at each point x its second fundamental form is injective on $\operatorname{Sym}^2 H$ for some hyperplane $H \subset T_x M$.*

No assumption is made on the topology of M. The same dimension was obtained in [YW26b] for closed stably parallelizable manifolds whose components have vanishing Euler characteristic. The lower bound on n enters only through the counting estimate of Lemma 9.2.

Background. A map is *free* if its first and second derivatives are linearly independent. For a free map the linearized metric operator has an algebraic right inverse, which is the mechanism of the classical proofs [N56, G86, Gü89]; but a free map needs $N \geq s_n + n$. We use instead the weaker condition of phase regularity. A covector ξ is *good* for an immersion F if the second fundamental form of F is injective on $\operatorname{Sym}^2(\ker \xi)$, and F is phase regular if every point has a good covector. For phase-regular maps an exact metric correction holds as soon as $N \geq s_n + 1$ (Proposition 3.1), and as in [YW26b, WY26] the global problem reduces to the construction of a single phase-regular partial isometry on a product of Hopf spheres (Theorem 2.2).

The idea of the proof. A generic bundle map $\underline{\mathbb{R}}^4 \to TM$ has rank four outside a closed three-dimensional submanifold P, the *core*, along which it has rank three. Near P a complex structure on $\underline{\mathbb{R}}^4$ splits $TM = H \oplus L$ with L a line, and

$$\operatorname{rank}\Big(H \oplus \operatorname{Sym}^2 H \oplus L\Big) = (n-1) + \frac{n(n-1)}{2} + 1 = s_n.$$

Thus a map that is free along H and immersive along L leaves a normal bundle N_5 of rank five over P. By Proposition 6.2, N_5 splits off three trivial lines unless $n \equiv 2 \pmod 4$, where $N_5 \simeq 3\mathcal{D} \oplus \underline{\mathbb{R}}^2$ for a line bundle $\mathcal{D}$ with possibly $w_1(\mathcal{D})^3 \neq 0$ (for instance on $\mathbb{RP}^n$ with $n \equiv 6 \pmod 8$). The two cases require different boundary models at the core (Sections 8 and 10).

Outside a tube around P we have $TM = E \oplus \underline{\mathbb{R}}^4$ with $\operatorname{rank} E = m = n - 4$, and we write

$$\mathbb{R}^{s_n+5} = \mathbb{R}^{s_n-152} \oplus \mathbb{R}^{157}.$$

The first factor carries an *old* map, free along E; the second carries an *active* map on four Hopf spheres whose derivatives of order at most three are independent. Their

normal spaces leave $3m + 1$ directions, while the mixed metric coefficients between the two blocks give $4m$ equations. We couple the blocks by a high-frequency product term and solve the mixed equations order by order. This requires a finite-order differential right inverse for the operator

$$(\eta, u, b) \longmapsto \Big(\langle C_{i\alpha}, \eta\rangle + d_\alpha u_i + d_i b_\alpha\Big)_{i,\alpha},$$

which we obtain from one flat normal tensor on the doubled exterior (Proposition 9.3) and one active map whose characteristic variety avoids the bad planes of that tensor (Theorem 13.1). Finally the hyperplane metric is filled across the disk bundle over P (Corollary 14.2), and the remaining terms along L are absorbed by six further Hopf factors (Proposition 16.2).

Organization of the paper. Sections 2–5 prove the compact-model reduction: metric correction, removal of double points, and the passage to the prescribed metric. Section 6 studies the core and the class of N_5, and Section 7 constructs the reference map. Sections 8–10 construct the boundary models and the flat tensor, and Sections 11–13 produce the exact metric on the exterior. Section 14 fills the disk, Section 15 combines the two core cases into one set of coefficients, Section 16 completes the line, and Section 17 proves Theorem 1.1.

2. Conventions and the compact-model reduction

In this section we fix our conventions, introduce partial metrics and partial second fundamental forms along integrable subbundles, and define phase regularity. We then associate a commuting Hopf distribution with a Euclidean subbundle and state the compact-model reduction, Theorem 2.2. Its proof occupies Sections 3–5.

2.1. **Partial metrics, second fundamental forms, and phase regularity.** All manifolds are smooth and second countable, and a closed manifold is compact and without boundary. We write $s_n = n(n+1)/2 = \dim \operatorname{Sym}^2 \mathbb{R}^n$. The Euclidean inner product $\langle\cdot,\cdot\rangle$ on $\mathbb{R}^N$ is extended complex-bilinearly to $\mathbb{C}^N$ whenever complex amplitudes occur. For a smooth map $F\colon M \to \mathbb{R}^N$ we write $\gamma(F) = F^*\langle\cdot,\cdot\rangle = (\mathrm{d}F)^*\mathrm{d}F$; an isometric embedding of (M, g) is a smooth embedding F with $\gamma(F) = g$.

Let $D \subset TX$ be an integrable subbundle of rank n. A smooth map $F\colon X \to \mathbb{R}^N$ is a *partial immersion* along D if $\mathrm{d}_D F = \mathrm{d}F|_D$ is injective; its *partial metric* is $\gamma_D(F) = (\mathrm{d}_D F)^*\mathrm{d}_D F$, a section of $\operatorname{Sym}^2 D^*$. Let $\mathcal{N}_F^D = (\mathrm{d}_D F(D))^\perp$ be the partial normal bundle, of rank $N-n$, and Π_F^D the orthogonal projection onto it. The *partial second fundamental form* $B_F^D\colon \operatorname{Sym}^2 D \to \mathcal{N}_F^D$ is defined on local sections U, V of D by

$$B_F^D(U, V) = \Pi_F^D\Big(U(V(F))\Big).$$

It is tensorial and symmetric: the term $(Uf)\,\mathrm{d}_D F(V)$ produced by the Leibniz rule lies in $\mathrm{d}_D F(D)$, and so does $U(V(F)) - V(U(F)) = \mathrm{d}_D F([U, V])$, because D is integrable. Thus the projection removes the connection terms of the second derivative. In a foliation chart (y, z) in which D is spanned by the derivatives $\partial_i = \partial_{y^i}$ one has

$$\partial_i\partial_j F = \Gamma_{ij}^k\,\partial_k F + B_{F,ij}^D, \tag{2.1}$$

where Γ_{ij}^k are the Christoffel symbols of $g_F = \gamma_D(F)$ computed with z fixed. When $D = TM$ we write $\mathcal{N}_F$ and B_F.

Definition 2.1. A partial immersion F is *phase regular* along D if at each point there is a hyperplane $H \subset D$ on which $B_F^D|_{\mathrm{Sym}^2 H}$ is injective. A nonzero covector $\xi \in D^*$ is *good* if $\ker \xi$ is such a hyperplane. The *best phase margin* at a point is the maximum, over unit covectors ξ, of the least singular value of $B_F^D|_{\mathrm{Sym}^2 \ker \xi}$. Here the norms on symmetric squares are induced by a fixed metric on D.

Since $\dim \mathrm{Sym}^2 H = s_n - n$ for a hyperplane H, phase regularity requires $N \geq s_n$. If $N \geq s_n + 1$ and ξ is good, the orthogonal complement of $B_F^D(\mathrm{Sym}^2 \ker \xi)$ in $\mathcal{N}_F^D$ has dimension $N - s_n \geq 1$. The best phase margin depends continuously on the point and on the leaf two-jet of F. Hence on a compact manifold phase regularity is equivalent to a positive lower bound for the best phase margin, and it persists under small C^2 perturbations.

2.2. **Commuting Hopf distributions and the compact-model theorem.** On the unit sphere $S^{2h-1} \subset \mathbb{C}^h$ let Y be the generator of the Hopf action $z \mapsto e^{2\pi i t} z$, $t \in \mathbb{R}/\mathbb{Z}$, so that $Y(z) = 2\pi \mathrm{i} z$. On $X_h = B \times (S^{2h-1})^a$ let Y_ℓ act as Y on the ℓ-th sphere factor and trivially on the other factors. The fields $Y_1, \ldots, Y_a$ commute and are pointwise linearly independent. Let $W \subset B \times \mathbb{R}^a$ be a subbundle of rank n with a bundle metric g_W. Then

$$w \longmapsto Y_w = \sum_{\ell=1}^{a} w_\ell Y_\ell, \qquad w \in W_b,$$

defines a rank-n distribution $D_W \subset TX_h$, isomorphic to $\pi_B^* W$, which we equip with the metric $\pi_B^* g_W$. All derivatives along D_W keep $b \in B$ fixed. If $e_1, \ldots, e_n$ is a local frame of W, the fields $Y_{e_1}, \ldots, Y_{e_n}$ commute, because their coefficient functions depend only on b and are therefore annihilated by every Y_m. In particular D_W is integrable.

Theorem 2.2 (Compact-model reduction)**.** *Let $n \geq 2$, let $N \geq s_n + 2$, and let B be a closed manifold. Fix a rank-n subbundle $W \subset B \times \mathbb{R}^a$ with a bundle metric g_W. Suppose that for every sufficiently large integer h there is a smooth map*

$$S_h \colon B \times (S^{2h-1})^a \longrightarrow \mathbb{R}^N$$

whose partial metric on D_W is $\pi_B^* g_W$ and which is phase regular along D_W. If M is a closed n-manifold and $TM \simeq f^* W$ for a smooth map $f \colon M \to B$, then every smooth Riemannian metric on M admits a phase-regular isometric embedding into $\mathbb{R}^N$ with image in any prescribed open ball.

Remark 2.3. The integer a, the base B, the subbundle W and its metric are fixed before h varies, and each S_h is defined on the whole product $B \times (S^{2h-1})^a$. The proof uses a single value of h, chosen before all other parameters. The derivative bounds and the phase margin of S_h may therefore depend on that value; no uniformity in h is needed.

3. Smooth metric correction at a phase-regular map

We prove a smooth perturbation theorem for partial metrics near a phase-regular map, with constants independent of the number of charts. Since B_F^D is injective only on symmetric squares of good hyperplanes, the linearized equation has no algebraic right inverse. Along each good phase we invert the transverse block algebraically and solve the remaining equations to high order in the wavelength.

Let X be a compact manifold with an integrable subbundle D of rank n. We use buffered foliation charts (y, z), consisting of inner charts covering X and enlarged charts containing them, in which D is spanned by the y-derivatives. Transitions preserve the leaves, so leaf derivatives transform into combinations of leaf derivatives. Norms are maxima of coordinate norms over the enlarged charts and include derivatives in both y and z. The symbol $\mathcal{J}_D^2 F$ denotes the values, first leaf derivatives and second leaf derivatives of F. We write

$$|u|_{j,\ell} = \sum_{b=0}^{j} \ell^b \|u\|_{C^b}, \qquad 0 < \ell \le 1,$$

and use the same notation for tensors and for the components of $\mathcal{J}_D^2 F$. Then $|uv|_{j,\ell} \le C_j |u|_{j,\ell} |v|_{j,\ell}$ and $|\partial u|_{j,\ell} \le \ell^{-1} |u|_{j+1,\ell}$.

Proposition 3.1 (Uniform smooth correction)**.** *Suppose $N \ge s_n + 1$. Fix inner and enlarged chart radii, a bound on the overlap multiplicity of the enlarged charts, and bounds at every fixed order for the transitions and their inverses. Fix a compact interval of positive numbers containing the eigenvalues of the leafwise coordinate matrices of $\gamma_D(F_0)$ and a positive lower bound on the best phase margin of a smooth reference partial immersion $F_0 \colon X \to \mathbb{R}^N$. Given a sufficiently small C^2 tolerance $\varepsilon_0 > 0$, there are a finite derivative order m_0 and a constant $\eta_0 > 0$, depending on bounds for $\|F_0\|_{C^{m_0}}$ and on the data just fixed, such that every smooth positive definite partial metric h with*

$$\|h - \gamma_D(F_0)\|_{C^{440}} < \eta_0$$

is induced by a smooth phase-regular partial immersion F with $\|F - F_0\|_{C^2} < \varepsilon_0$. The constants m_0 and η_0 are independent of the number of charts. If $X = M \times P$ is a foliated product with $D = TM \oplus 0$, the map F is jointly smooth in the points of M and the transverse parameters.

3.1. **Positive increments.** *The mode equations.* Work in an adapted chart with $y^1 = \psi$, where $\mathrm{d}_D\psi$ is good, and write $g_F = \gamma_D(F)$. Decompose a complex amplitude as

$$Z = \mathrm{d}_D F(g_F^{-1} u) + v, \qquad v \perp \mathrm{d}_D F(D),$$

where u is a complex leaf covector. Differentiating $\langle \partial_i F, Z \rangle = u_i$ and $\langle \partial_i F, v \rangle = 0$ with z fixed and using (2.1), we obtain for the linearization $(D\gamma_D)_F Z = (\mathrm{d}_D F)^T \mathrm{d}_D Z + (\mathrm{d}_D Z)^T \mathrm{d}_D F$

$$(D\gamma_D)_F Z_{ij} = \partial_i u_j + \partial_j u_i - 2\Gamma_{ij}^k u_k - 2\langle B_{F,ij}^D, v \rangle. \tag{3.1}$$

Conjugation by $e^{\mathrm{i}\psi/\tau}$ adds $(\mathrm{i}/\tau)(\delta_{i1} u_j + \delta_{j1} u_i)$. Let $\beta(v) = (\langle B_{F,ab}^D, v \rangle)_{a,b \ge 2}$. For a good phase β is onto, since its adjoint is the injective map $B_F^D|_{\mathrm{Sym}^2 \ker \mathrm{d}_D\psi}$. Put $J = \beta^*(\beta\beta^*)^{-1}$

and set

$$T(u)_{ab} = \tfrac{1}{2}(\partial_a u_b + \partial_b u_a) - \Gamma^k_{ab} u_k, \qquad v = v_0 + J\Big(T(u) - \tfrac{1}{2} h_H\Big), \quad v_0 \in \ker \beta,$$

where $h_H = (h_{ab})_{a,b\geq 2}$ is the transverse part of the desired mode amplitude h. This solves the transverse equations exactly. The remaining equations are equivalent to $(\mathrm{i}/\tau)u + \mathcal{L}u = k$, where

$$\begin{aligned}
(\mathcal{L}u)_1 &= \partial_1 u_1 - \Gamma^k_{11} u_k - \langle B^D_{F,11}, JT(u)\rangle,\\
(\mathcal{L}u)_a &= \partial_1 u_a + \partial_a u_1 - 2\Gamma^k_{1a} u_k - 2\langle B^D_{F,1a}, JT(u)\rangle,\\
k_1 &= \tfrac{1}{2} h_{11} + \langle B^D_{F,11}, v_0\rangle - \tfrac{1}{2}\langle B^D_{F,11}, Jh_H\rangle,\\
k_a &= h_{1a} + 2\langle B^D_{F,1a}, v_0\rangle - \langle B^D_{F,1a}, Jh_H\rangle.
\end{aligned}$$

More precisely, an error e in this system contributes $2e_1$ to the 11 entry and e_a to the $1a$ entry of the metric error, and nothing to the transverse entries. For an integer $p \geq 1$ we choose

$$u = \frac{\tau}{\mathrm{i}} \sum_{b=0}^{p-1} \Big(-\frac{\tau}{\mathrm{i}}\mathcal{L}\Big)^b k.$$

The sum telescopes, and the error is $(-1)^{p-1}(\tau/\mathrm{i})^p \mathcal{L}^p k$. The coefficients of $\mathcal{L}$ are smooth functions of the leaf two-jet on the set where the immersion and phase margins are bounded below. In a fixed ℓ-weighted norm each derivative costs ℓ^{-1} and one higher input order. The error therefore carries a factor $(\tau/\ell)^p$ relative to finitely many norms of h and v_0, and Z is bounded in each fixed slow norm by finitely many norms of h and v_0. When $h = 0$ one has moreover $Z - v_0 = O(\tau/\ell)$ in the same relative sense. All formulas preserve the supports of their inputs. If a phase requires several adapted charts, we split its amplitudes by a partition of unity and add the resulting amplitudes with the same phase before forming quadratic products.

Positive squares. Good covectors are open and dense. Indeed, in a Euclidean frame a covector ξ is good exactly when some maximal minor of $B^D_F \circ \mathrm{Sym}^2(|\xi|^2 \,\mathrm{Id} - \xi \otimes \xi)$ is nonzero, because $\mathrm{Sym}^2(|\xi|^2 \,\mathrm{Id} - \xi \otimes \xi)$ maps onto $\mathrm{Sym}^2 \ker \xi$; the sum of squares of these minors is a polynomial in ξ, nonzero somewhere by phase regularity. At a point, let e_i be an orthonormal coframe for the reference metric. The squares of e_i and of $e_i \pm e_j$ span the symmetric forms, and

$$\sum_{i<j} \Big((e_i + e_j)^{\otimes 2} + (e_i - e_j)^{\otimes 2}\Big) = 2(n-1) \sum_i e_i^{\otimes 2}.$$

Hence weight c on each pair square, with $0 < c < 1/(2(n-1))$, and weight $1 - 2c(n-1)$ on each single square represent the metric with positive coefficients. Perturb these covectors into the good set. A right inverse of the resulting square map L corrects the small error of the positive representation. If $Lc_0 = g_0$ is the corrected representation, with c_0 positive, and R is a right inverse, we replace R by $R + (c_0 - Rg_0)\ell_0$, where ℓ_0 is a linear functional with $\ell_0(g_0) = 1$. This is still a right inverse, and it sends g_0 to c_0; so the coefficients are positive linear functions of the tensor near g_0. We extend the

covectors as differentials of coordinate-linear phases on a small chart and use a squared partition of unity. This gives a fixed representation

$$A = \sum_\nu b_\nu(A)^2 \, \mathrm{d}_D\phi_\nu^{\otimes 2}, \qquad b_\nu(A) = \chi_\nu \sqrt{\ell_\nu(A)}, \tag{3.2}$$

valid in a C^0 neighbourhood of the reference metric. On each sufficiently small chart we choose a smooth unit section n_ν of the kernel of the transverse pairing β of ϕ_ν by projecting and normalizing a fixed ambient vector; this kernel has rank $N - s_n \geq 1$. We then rescale the phases by a hierarchy of increasing positive constants and the coefficient functionals by the inverse squares, which leaves (3.2) unchanged. Every sum or difference of two phases with intersecting supports is good, because after division by its larger weight it is uniformly close to the dominant good covector; doubled phases are good as well. All these properties hold in a fixed leaf two-jet neighbourhood of F_0.

The increment. Let H_{tar} lie in a smaller neighbourhood of g_0, let $q_0 \geq 1$, and let $0 < \delta \leq \tau \leq \ell \leq 1$, $\theta = \tau/\ell$. For a trial tensor A we apply the homogeneous mode formula with $p = q_0 + 2$ to the free amplitudes $\sqrt{2}\,\delta\tau\, b_\nu(A) n_\nu$ along the phases ϕ_ν, and we call the real sum of the resulting waves $U_1(A)$. After division by δ, the coefficient of each wave in $\mathrm{d}_D U_1$ is its leading term $\mathrm{i}\sqrt{2}\, b_\nu(A) n_\nu \otimes \mathrm{d}_D\phi_\nu$ plus an error with a factor θ in every slow norm. Distinct labels with intersecting supports have distinct phase gradients, so only each label paired with its own conjugate contributes to the zero-phase part. Since $|n_\nu| = 1$, the leading coefficients give A by (3.2), and the zero-phase part of $(\mathrm{d}_D U_1)^* \mathrm{d}_D U_1$ equals $\delta^2(A + T_0(A))$. Each term of $T_0(A)$ contains a differentiated slow amplitude or a homogeneous solver correction, so T_0 has a factor θ in each fixed weighted norm. Differences of T_0 have the same factor, with a finite loss of derivatives: the chain rule for the positive square root, applied along the segment between two inputs, bounds $b_\nu(A) - b_\nu(\widetilde{A})$ by finitely many norms of $A - \widetilde{A}$. The nonzero modes have coefficients of size $O(\delta^2)$.

Set $A_0 = H_{\mathrm{tar}}$ and $A_{j+1} = H_{\mathrm{tar}} - T_0(A_j)$ for $j < q_0$. For θ below a threshold the trial tensors stay in the positivity neighbourhood. To fix this threshold, the finitely many C^0 conditions are evaluated backwards through the formulas, which requires finitely many weighted derivatives of the input. The difference bound gives $A_{j+1} - A_j = O(\theta^{j+1})$ in each fixed weighted norm. Since $H_{\mathrm{tar}} = A_{q_0} + T_0(A_{q_0-1})$, the mean at A_{q_0} differs from $\delta^2 H_{\mathrm{tar}}$ by $\delta^2(T_0(A_{q_0}) - T_0(A_{q_0-1})) = O(\delta^2\theta^{q_0+1})$. We cancel each remaining quadratic mode by the mode formula with target its negative, $v_0 = 0$ and $p = q_0 + 2$, and call the sum of these corrections U_2. Derivatives falling on the fast exponentials cost τ^{-1} each. The exact quadratic expansion of the metric of $F + U_1 + U_2$ then gives, for $U = U_1 + U_2$,

$$|U|_{j,\tau} \leq C_j \delta\tau, \tag{3.3}$$

$$|\gamma_D(F + U) - \gamma_D(F) - \delta^2 H_{\mathrm{tar}}|_{j,\tau} \leq C_j\{\delta\theta^{q_0} + \delta^3/\tau\}. \tag{3.4}$$

In the second estimate the two linear mode errors and the residual of the mean are bounded by the first term. The cross term $\mathcal{C}_D(U_1, U_2)$ and the square of $\mathrm{d}_D U_2$ are bounded by $C_j\delta^3/\tau$ and $C_j\delta^4/\tau^2$ respectively, since $|U_2|_{j,\tau} = O(\delta^2)$. All operations use leaf derivatives only. The same estimates hold for derivatives in the transverse directions by the product rule, since these fall only on the smooth coefficients and phases.

3.2. **The iteration.** Before each step we smooth the current map and tensor by localized convolution in all chart variables. We use kernels of integral one whose moments of orders 1 through r vanish. Such a kernel is $K = \sum_{m=0}^{r} a_m \lambda_m^{-d} \Phi(\cdot/\lambda_m)$, where Φ is a compactly supported smooth kernel of integral one, the dilations λ_m are distinct, and $\sum_m a_m \lambda_m^s = \delta_{s0}$ for $0 \le s \le r$; this Vandermonde system is uniquely solvable. We combine the convolutions by a partition of unity, with further cutoffs equal to one on the enlarged supports of the localized terms, and write S_ℓ for the resulting smoothing at scale ℓ. Derivatives of order at most two in the leaf jet may be moved onto the localized map. Since transitions preserve the leaf directions, no derivative transverse to D enters this base bound. All further derivatives may be put on the kernel. Consequently

$$|\mathcal{J}_D^2(S_\ell G)|_{j,\ell} \le C_{j,r} \|\mathcal{J}_D^2 G\|_{C^0}, \qquad |S_\ell H|_{j,\ell} \le C_{j,r} \|H\|_{C^0}.$$

Put

$$W_j(t; G, H) = |\mathcal{J}_D^2 G|_{j,t} + |H|_{j,t}, \qquad W_r(t; G, H) \le K_0.$$

Then for $0 < \tau \le \ell \le t \le 1$ the smoothing error satisfies

$$|\mathcal{J}_D^2(G - S_\ell G)|_{j,\tau} + |H - S_\ell H|_{j,\tau} \le C_{j,r}\{K_0(\ell/t)^r + W_j(t; G, H)(\tau/t)^r\}. \tag{3.5}$$

Indeed, a term in which $b \le r$ derivatives fall on the difference has, by Taylor's formula and the vanishing moments, the bound $C\ell^{r-b}\tau^b$ times the C^r norm of the input components. For $b > r$ we use the boundedness of the smoothing at the inherited order and $(\tau/t)^b \le (\tau/t)^r$. Derivatives of cutoffs and frames are bounded and enter only at lower orders. The same arguments apply when X is a product with compact parameter spaces.

Set $g_0 = \gamma_D(F_0)$. At stage $k \ge 10$, with current scale t, map G and normalized defect H, we choose

$$\delta = t^k, \qquad \ell = t^{11/10}, \qquad \tau = t^{6/5}, \qquad q_0 = r = 40(k+1), \qquad H = \delta^{-2}(h - \gamma_D(G)).$$

We apply (3.3)–(3.4) at $S_\ell G$ with target $S_\ell H - (\delta_+/\delta)^2 g_0$, where the next reserve δ_+ equals τ^k (to stay in stage k) or τ^{k+1} (to pass to stage $k+1$). We keep H in a fixed small C^0 neighbourhood of g_0 and $W_r(t; G, H) \le K_0$. The smoothing estimates then bound all finite input norms needed at each step in terms of K_0 and k. Let E be the error in (3.4), put $G_+ = G + U$ and $H_+ = \delta_+^{-2}(h - \gamma_D(G_+))$. Expanding the quadratic metrics of $G + U$ and of $S_\ell G + U$ separately gives the exact identity

$$\delta_+^2(H_+ - g_0) = \delta^2(H - S_\ell H) - E - \mathcal{C}_D(G - S_\ell G, U), \tag{3.6}$$

where $\mathcal{C}_D(V, U)(a, b) = \langle \mathrm{d}_D V(a), \mathrm{d}_D U(b)\rangle + \langle \mathrm{d}_D V(b), \mathrm{d}_D U(a)\rangle$. The square of $\mathrm{d}_D U$ cancels between the two expansions, and the square of the smoothing error does not occur in (3.6).

By (3.3) we have $|\mathrm{d}_D U|_{j,\tau} \le C_j(k)\delta$. For the smaller reserve $\delta_+ = \tau^{k+1}$, the powers of t of the terms in (3.6) after division by δ_+^2 are

term	power
$\delta(\tau/\ell)^{q_0}$	$(13k+8)/5$
δ^3/τ	$(3k-18)/5$
$\delta^2(\ell/t)^r$	$(18k+8)/5$
$\delta^2(\tau/t)^r$	$(38k+28)/5$
$\delta(\ell/t)^r$	$(13k+8)/5$
$\delta(\tau/t)^r$	$(33k+28)/5$.

The two $(\tau/t)^r$ terms multiply $W_j(t; G, H)$. Since $k \ge 10$, all powers are positive, with minimum at least $12/5$. The larger reserve multiplies every term by τ^2 and only improves the estimates. In passing to $W_j(\tau; G_+, H_+)$, every inherited derivative of positive order of G gains at least $\tau/t = t^{1/5}$, and the two-jet increment is $O(\delta/\tau)$. Therefore

$$W_j(\tau; G_+, H_+) \le \|\mathcal{J}_D^2 G\|_{C^0} + \|g_0\|_{C^0} + C_j(k) t^{1/5}(1 + W_j(t; G, H)). \tag{3.7}$$

We fix K_0 larger than four times the sum of the first two terms throughout the chosen C^2 neighbourhood of F_0. At each stage there is a positive threshold for t below which the step preserves the C^0 neighbourhood of the defect and the bound $W_r \le K_0$. We also require $\|U\|_{C^{\lfloor k/2 \rfloor}} \le c_1 t$ for a fixed small $c_1 > 0$; this is possible since $\|U\|_{C^{\lfloor k/2 \rfloor}} \le C(k) t^{k + \frac{6}{5}(1 - \lfloor k/2 \rfloor)}$ and $k + \frac{6}{5}(1 - \lfloor k/2 \rfloor) > 1$.

We start at $k = 10$ with

$$G = \sqrt{1 - t^{20}}\, F_0, \qquad H = g_0 + t^{-20}(h - g_0),$$

so that H is the normalized defect of G. First we choose the initial scale t, using the reference derivative bounds, so small that the C^2 displacement of G and the sum of all future scales are small; since each step replaces t by $t^{6/5}$, this sum is at most twice the initial scale once $t \le 2^{-5}$. Then we choose η_0, a fixed multiple of t^{20}, which controls the initial C^{440} norm of the normalized defect. At fixed k, with t replaced by $t^{6/5}$ at each step, (3.7) bounds any fixed higher norm after finitely many steps. Indeed, the coefficient of the inherited norm tends to zero; once it is at most $1/4$, the recurrence bounds the norm, and its limit superior is at most $K_0/4$. We wait until the next controlled order $40(k+2)$ has norm below K_0 and the finitely many smallness conditions of stage $k+1$ hold. Then we choose $\delta_+ = \tau^{k+1}$ and pass to stage $k+1$. The estimate (3.7) at the higher order keeps its norm below K_0 across the transition, after decreasing the scale further within the current stage if necessary.

The increment bounds keep the maps in the fixed neighbourhood where the phases are good. Each stage ends after finitely many steps, so $k \to \infty$ and, by the bound $\|U\|_{C^{\lfloor k/2 \rfloor}} \le c_1 t$, the increments sum absolutely in each fixed smooth norm. The limit F induces h, because the bounded normalized defect is multiplied by $\delta^2 \to 0$. The initial choices depend on finitely many reference derivatives.

It remains to obtain uniformity. The bounded reference two-jets with the given phase and metric margins form a compact set. We choose a finite collection of positive-square

and projected-normal constructions valid over this set. A fixed subdivision of each enlarged chart makes the chosen data valid on each piece, by the uniform reference C^3 bound. We colour the intersection graph of the resulting supports and assign the hierarchy of weights to the pairs consisting of a colour and a square index. Distinct supports of the same colour are disjoint. Hence the number of required weights and the number of local terms meeting a chart are bounded independently of the size of the atlas. The good differentials of the finitely many phase types yield uniformly buffered adapted coordinates, obtained by shrinking the chart and inverting a leaf derivative that is bounded away from zero. All estimates above then have uniform constants at each fixed order, including the finitely many starting thresholds at $k = 10$. This proves Proposition 3.1.

4. Removal of double points with the metric fixed

The output of Proposition 3.1 need not be injective. We now remove double points without changing the metric or the target dimension. The tool is a smooth finite-parameter family of exact isometric immersions driven by prescribed oscillating normal perturbations.

Let F_0 be an immersion with normal projection Π_0. For an immersion F close to F_0 in C^1, the restriction $P_F = \Pi_0|_{\mathcal{N}_F}$ is an isomorphism onto $\mathcal{N}_{F_0}$, and its polar part $\Phi_F = (P_F P_F^*)^{-1/2} P_F$ is an isometry $\mathcal{N}_F \to \mathcal{N}_{F_0}$. We call $\widehat{B}_F = \Phi_F \circ B_F$ the *identified second fundamental form*. In local coordinates $\widehat{B}_F(\partial_\alpha, \partial_\beta) = \mathcal{Q}_F(\partial_\alpha \partial_\beta F)$, where $\mathcal{Q}_F$ is a smooth function of $\mathrm{d}F$ with $\mathcal{Q}_{F_0} = \Pi_0$. For $\nu \in \mathcal{N}_{F_0}$ and a covector ξ we write $\nu \otimes \xi^{\otimes 2}$ for the map $Q \mapsto Q(\xi, \xi)\nu$.

Lemma 4.1 (Exact parameter families)**.** *Let $F_0 \colon M \to \mathbb{R}^N$ be a phase-regular immersion of a closed manifold, with $N \geq s_n + 1$, and let $g_0 = \gamma(F_0)$. Take finitely many coordinate phases ψ_l whose differentials are good throughout their charts, smooth normal fields ν_l perpendicular to $B_{F_0}(\mathrm{Sym}^2 \ker \mathrm{d}\psi_l)$ on these charts, and smooth cutoffs χ_l compactly supported in the charts. There are positive constants w_l and a ball of parameters $u = (u_{l,1}, u_{l,2})$ about zero with the following property. For each sufficiently small $t > 0$ there is a family $F_{t,u}$, jointly smooth in (x, u), with $\gamma(F_{t,u}) = g_0$ for all u and $F_{t,u} \to F_0$ in C^2 uniformly in u as $t \to 0$. Put*

$$\tau = t^{6/5}, \quad \sigma = t^{44/5}, \quad b = t^{56/5}, \quad (c_{l,1}, c_{l,2}) = \Big(\cos(w_l \psi_l/\tau), \sin(w_l \psi_l/\tau)\Big),$$

and let $\widehat{B}_{t,u}$ be the identified second fundamental form of $F_{t,u}$. Then

$$\begin{aligned} \partial_{u_{l,i}} \widehat{B}_{t,u} &= \sigma\, \chi_l \nu_l \otimes \mathrm{d}\psi_l^{\otimes 2}\, c_{l,i} + \sigma\, O(\eta(t) + |u|), \\ \partial_{u_{l,i}} F_{t,u} &= -b\, w_l^{-2} \chi_l \nu_l c_{l,i} + b\, O(\eta(t) + |u|), \end{aligned}$$

where the remainders are uniform in C^0 and $\eta(t) \to 0$.

Proof. We include the prescribed phases in the hierarchy of weights used for (3.2), so that all doubled phases and all sums and differences of distinct phases with intersecting

supports are good. Set $\delta = t^{10}$, $\ell = t^{11/10}$ and $G_0 = \sqrt{1-\delta^2}F_0$. Constant scaling preserves the normal spaces and the kernels of the transverse pairings. To the homogeneous modes we add the free amplitudes

$$-\delta\tau\chi_l w_l^{-2}\nu_l(u_{l,1} - \mathrm{i}u_{l,2})$$

with the weighted phases $w_l\psi_l$. Since the mode operator is complex linear, the resulting amplitude is $\delta\tau(u_{l,1} - \mathrm{i}u_{l,2})Z_l^{(0)}$, where $Z_l^{(0)}$ is the output for $-w_l^{-2}\chi_l\nu_l$ and differs from it by $O(\theta)$, $\theta = \tau/\ell = t^{1/10}$. The leading quadratic mean of these modes, divided by δ^2, is the tensor

$$H_{\mathrm{p}}(u) = \sum_l \frac{\chi_l^2|\nu_l|^2(u_{l,1}^2 + u_{l,2}^2)}{2w_l^2}\,\mathrm{d}\psi_l^{\otimes 2}.$$

We use the target $H_{\mathrm{tar}} = (1 - (\tau^{10}/\delta)^2)g_0 = (1 - t^4)g_0$ and begin the positive-square formulas for the background at $A_0 = H_{\mathrm{tar}} - H_{\mathrm{p}}(u)$. For a fixed small parameter ball this stays in the positivity neighbourhood. The zero-phase error $T_0(A, u)$ involves only each label paired with its own conjugate, so it is the sum of a term depending only on A and a term quadratic in u. In the finite recursion

$$A_{j+1} = H_{\mathrm{tar}} - H_{\mathrm{p}}(u) - T_0(A_j, u)$$

all background functions are therefore even in u. Their first parameter derivatives are $O(|u|)$ in the required slow norms, since their second parameter derivatives are uniformly bounded by the finite formulas.

We take accuracy $q_0 = 440$ in this first increment and cancel the nonzero quadratic modes as in (3.4). The target is $(\delta^2 - \tau^{20})g_0$, and the bounds (3.3)–(3.4) hold in every fixed weighted joint norm in (x, u). Since $\tau^{20} = t^{24}$ and $\theta^{440} = t^{44}$, the first normalized defect $\tau^{-20}(g_0 - \gamma(G_1))$ at scale τ differs from g_0 by $O(t^{30} + t^{24/5})$ in these norms. The first map G_1 is within $O(\sigma)$ of F_0 in C^2 and within $O(\delta)$ in C^1. The leading terms of its first parameter derivatives are those in the second line of the lemma, because $\delta\tau = b$. Two spatial derivatives falling on the phase give the leading term of the first line, because $b\tau^{-2} = \delta/\tau = \sigma$; terms in which fewer derivatives fall on the phase are $O(\delta) = \sigma\,O(\tau)$. The homogeneous solver contributes a relative error $O(\tau/\ell)$, the quadratic correction a relative error $O(\delta/\tau)$, and the background parameter derivatives a relative error $O(|u|)$. This gives the displayed estimates before the exact correction, with ordinary second derivatives in place of the second fundamental form.

For the correction, we extend the parameters to a compact torus by a smooth map into the small ball that is the identity on a smaller ball. We run the iteration of Proposition 3.1 on the product, with the partial metric along M, starting at scale τ and stage 10, with scale and waiting choices common to all parameters. The foliated estimates give smooth convergence jointly in the parameters and the points of M. To retain the leading responses we impose additional finite waits on entering a stage $k > 10$: for $j = 1, 2, 3$ we require $C_j(k)t_i^{k-10} \le 1$, where t_i is the scale at the step and $C_j(k)$ is the constant in (3.3). Each wait is finite because $k - 10 > 0$. The C^j norm of a later increment is then at most a fixed constant times $t_i^{10+(6/5)(1-j)}$ for $j = 1, 2, 3$. Since the first scale is τ, the sums of these norms are $O(t^{12})$, $O(t^{264/25})$ and $O(t^{228/25})$

respectively. The first is $o(b)$, the second $o(\delta)$ and the third $o(\sigma)$. Hence one exact family has both parameter estimates, including the second-derivative estimate. Its first jets and their first parameter derivatives differ from those of F_0 by $O(\delta)$. Orthogonal projection and polar identification are smooth functions of the first jets on a uniform immersion neighbourhood, and differentiating either costs $O(\delta) = o(\sigma)$ in the estimate of the second fundamental form. This proves the lemma, with $\eta(t) = t^{1/10}$. □

We use the following parameter-avoidance principle.

Lemma 4.2. *Let Σ be a manifold of dimension m, let $\mathcal{U} \subset \mathbb{R}^P$ be open, and let $\Phi\colon \Sigma \times \mathcal{U} \to \mathbb{R}^k$ be smooth. Let $S \subset \Sigma \times \mathcal{U}$ and suppose that $\mathrm{d}\Phi$ has rank at least $r > m$ at every zero of Φ in S. Then the set of $u \in \mathcal{U}$ for which $\Phi(\cdot, u)$ has a zero in S has measure zero. The same holds for countably many such equations.*

Proof. Near a zero we choose r components of Φ with independent differentials. By the implicit function theorem their common zero set is a submanifold of dimension $m+P-r < P$. Its projection to $\mathbb{R}^P$ has measure zero, by Sard's theorem or by covering it with compact coordinate cubes on which the projection is Lipschitz. Countably many such neighbourhoods cover the zeros. □

We also need a lower bound for the span of the normal directions available to the perturbations of Lemma 4.1. We regard a symmetric tensor $P \in \operatorname{Sym}^2 V$ as a linear map $V^* \to V$. Then $P\xi = 0$ if and only if $P \in \operatorname{Sym}^2 \ker \xi$.

Lemma 4.3. *Let V be a vector space of dimension n, let $\mathcal{N}$ be an inner product space of dimension $\mu \geq q + 2$, where $q = s_n - n$, and let $B\colon \operatorname{Sym}^2 V \to \mathcal{N}$ be linear with $\operatorname{rank} B \geq q + 1$. Suppose that the set Ω of good covectors, namely those $\xi \neq 0$ with $B|_{\operatorname{Sym}^2 \ker \xi}$ injective, is nonempty. Then the spaces $(B(\operatorname{Sym}^2 \ker \xi))^{\perp}$, $\xi \in \Omega$, span a subspace of dimension at least $n + 1$.*

Proof. Let $K = \ker B$ and $d = \dim K \leq n - 1$. For $\xi \in \Omega$ the map $K \to V$, $Q \mapsto Q\xi$, is injective, since its kernel lies in $K \cap \operatorname{Sym}^2 \ker \xi = 0$. Let $\mathcal{R}$ consist of the symmetric tensors P with $P\xi \in K\xi$ for all $\xi \in \Omega$. This condition does not depend on the choice of a nonempty open subset of Ω, since it is the polynomial identity obtained by wedging $P\xi$ with a basis of $K\xi$.

The annihilator of the span in question is $B(\mathcal{R})$. Indeed, a vector BP lies in $B(\operatorname{Sym}^2 \ker \xi)$ exactly when $P - Q \in \operatorname{Sym}^2 \ker \xi$ for some $Q \in K$, that is, when $(P - Q)\xi = 0$. Hence $\dim B(\mathcal{R}) = \dim \mathcal{R} - d$, and it suffices to show that $\dim \mathcal{R} - d \leq d(d-1)/2$.

For $d = 0$, $P\xi = 0$ on an open set forces $P = 0$. For $d > 0$ choose a basis Q_i of K and set $a(\xi) = (Q_i(\xi, \xi))_{i=1}^{d}$. Choose $\xi_0 \in \Omega$ with $a(\xi_0) \neq 0$. Since $\mathrm{d}(Q_i(\xi, \xi)) = 2Q_i\xi$ and the vectors $Q_i\xi$ are independent, $J = \mathrm{d}a_{\xi_0}$ is onto. For $P \in \mathcal{R}$ the differential of $p(\xi) = P(\xi, \xi)$ lies in the span of those of the a_i; in local submersion coordinates it follows that $p = f \circ a$ near ξ_0. The differentials $\mathrm{d}p_{\xi_0}$, $P \in \mathcal{R}$, span a d-dimensional space, since they include those of the Q_i. On the kernel of $P \mapsto \mathrm{d}p_{\xi_0}$ we have $\mathrm{d}f_{a(\xi_0)} = 0$ and $\mathrm{d}^2 p = J^* C J$ with $C = \mathrm{d}^2 f_{a(\xi_0)}$. Homogeneity gives $CJ\xi_0 = 0$, because $\mathrm{d}^2 p(\xi_0, \cdot) = \mathrm{d}p_{\xi_0} = 0$, and $J\xi_0 = 2a(\xi_0) \neq 0$. The assignment $P \mapsto C$ on this kernel is injective, since the Hessian

determines a quadratic form. The symmetric forms on $\mathbb{R}^d$ annihilating a fixed nonzero vector form a space of dimension $d(d-1)/2$. Thus $\dim \mathcal{R} - d \le d(d-1)/2$, and the span has dimension at least

$$\mu - \frac{d(d-1)}{2} \;\ge\; (q+2) - \frac{(n-1)(n-2)}{2} = n+1. \qquad \square$$

Finally, if $L(\alpha) = \sum_l (\alpha_{l,1}\cos\theta_l + \alpha_{l,2}\sin\theta_l)v_l$ for vectors v_l in an inner product space, then $LL^*w = \sum_l \langle w, v_l\rangle v_l$ for all angles θ_l. Thus a cosine–sine pair of columns has the same Gram operator as the single column v_l, and the nonzero singular values of L do not depend on the angles.

Proposition 4.4 (Same-metric approximation by embeddings)**.** *Let M be closed, let $N \ge s_n + 2$, and let $F\colon M \to \mathbb{R}^N$ be a smooth phase-regular immersion. For every $\varepsilon > 0$ there is a smooth phase-regular embedding G with $\gamma(G) = \gamma(F)$ and $\|G - F\|_{C^2} < \varepsilon$.*

Proof. Since phase regularity persists under small C^2 perturbations, we may assume that every map within ε of F in C^2 is a phase-regular immersion. Write $q = s_n - n$ and $\mu = N - n \ge q + 2$. By phase regularity, $\operatorname{rank} B_F \ge q$ everywhere.

Removal of the minimal-rank locus. At a point where $\operatorname{rank} B_F = q$, the kernel K of B_F has dimension n, and $U = (\operatorname{im} B_F)^\perp \subset \mathcal{N}_F$ has dimension $\mu - q \ge 2$. In frames adapted to the splittings of $\operatorname{Sym}^2 TM$ and $\mathcal{N}_F$ defined by K and $\operatorname{im} B_F$ at this point, $\widehat{B}$ has a block matrix whose $(1,1)$ block is invertible nearby. Near the point the rank-q equation is the vanishing of its Schur complement in $\operatorname{Hom}(K, U)$, whose differential with respect to $\widehat{B}$ is the K-to-U block. For every good ξ at the point, $B_F(\operatorname{Sym}^2 \ker \xi) = \operatorname{im} B_F$, so U is the space of free normal directions for the phase ξ. Such a perturbation changes the Schur complement to leading order by $Q \mapsto Q(\xi,\xi)\nu$ with $\nu \in U$. The quadratic evaluations at good ξ span K^*, since a quadratic form vanishing on an open set is zero. Hence finitely many such perturbations span $\operatorname{Hom}(K, U)$. We extend each of them by a coordinate-linear phase and a smooth projected normal field, and choose a cutoff equal to one near the point. The rank-q locus is compact, so a finite list suffices. By Lemma 4.1 and the cosine–sine identity, the Schur equations have, for a small parameter ball and small t, parameter derivative of rank $n(\mu - q) \ge 2n > n$ near this locus. Outside the chosen neighbourhoods the $(q+1)$-st singular value of B_F has a positive minimum. Lemma 4.2 gives an immersion F_1 with the same metric, arbitrarily close to F in C^2, with $\operatorname{rank} B_{F_1} \ge q + 1$ everywhere.

Removal of double points. At every point of F_1, Lemma 4.3 provides $n+1$ independent free normal vectors. There is a fixed neighbourhood of the diagonal in $M \times M$ on which every immersion sufficiently C^1-close to F_1 is injective on pairs of distinct points. To see this, we project, in each of finitely many small convex coordinate balls, to a fixed tangent plane and compare the derivative with an invertible constant matrix; integration along segments gives a uniform lower bound for the displacement on each ball. The remaining pairs form a compact set disjoint from the diagonal. At a coincidence (x, y) in this set we choose $n+1$ independent free normal vectors at x, with phases and cutoffs supported away from y. A finite list covers the coincidence set. In the equation $F_{t,u}(x) - F_{t,u}(y) = 0$ the parameter derivatives divided by b in Lemma 4.1 produce these $n + 1$ normal

directions, uniformly in the phase angles. Together with the n independent tangent directions $\mathrm{d}F_1(T_xM)$, which are orthogonal to them, they give rank at least $2n+1 > 2n = \dim(M\times M)$. This inequality persists for a small parameter ball and small t. Away from the chosen neighbourhoods the value difference has a positive minimum on the compact set of pairs. Lemma 4.2 removes every coincidence. Taking both perturbations small gives the required C^2 bound, and by compactness the injective immersion is an embedding. □

5. From a compact model to the prescribed metric

We now prove Theorem 2.2. We wind M at a high frequency T into the sphere product, pull back a single model map S_h, and correct the resulting metric by Proposition 3.1.

Lemma 5.1. *Let (M,g), f and W be as in Theorem 2.2. There are a smooth immersion $\phi\colon M\to\mathbb{R}^d$ and smooth bundle maps $L\colon TM\to M\times\mathbb{R}^a$ and $R\colon M\times\mathbb{R}^d\to M\times\mathbb{R}^a$ such that, for every $x\in M$, L_x is an isometry of (T_xM,g_x) onto $(W_{f(x)},g_W)$ and $R_x\,\mathrm{d}\phi_x = L_x$.*

Proof. Take a finite coordinate cover and compactly supported functions $\beta_i\ge 0$, at least one positive at every point. The coordinates of ϕ are the β_i together with their products by each of the corresponding local coordinates. At a point with $\beta_i>0$, the derivative of ϕ determines all components of a tangent vector, so ϕ is an immersion. We normalize a bundle isomorphism $TM\to f^*W$ by the inverse positive square root of the change of metric, and obtain isometries $L_x\colon T_xM\to W_{f(x)}$ for the prescribed g. Using the inclusion $W\subset B\times\mathbb{R}^a$ we regard L_x as $\mathbb{R}^a$-valued. Let $A_x=\mathrm{d}\phi_x$, let A_x^+ be the smooth left inverse obtained from orthogonal projection onto $\operatorname{im}A_x$, and set $R_x=L_xA_x^+$. Then $R_xA_x=L_x$. □

Lemma 5.2. *In the setting of Theorem 2.2, let ϕ, L and R be as in Lemma 5.1. There is an integer h for which S_h exists, and there are $K_0,T_0\ge 1$, with the following property. For $K\ge K_0$ and $T\ge\max(K,T_0)$ there is a smooth map $z=(z_1,\dots,z_a)\colon M\to(S^{2h-1})^a$ such that*

$$F_{K,T}(x)=S_h\big(f(x),z_1(x),\dots,z_a(x)\big)$$

satisfies the following, in the charts $x=\exp_p(y/T)$, $|y|<4$, centred at a T^{-1}-net:

(i) *the inner charts $|y|<2$ cover M, the overlap multiplicity is bounded, the transitions and the coefficients of T^2g are bounded at every fixed order, and T^2g is uniformly positive, all uniformly in K and T;*

(ii) *$|F_{K,T}|\le\sup|S_h|$, and every fixed-order derivative of $F_{K,T}$ in these charts is bounded uniformly in K and T;*

(iii) *for a constant C independent of K and T,*

$$\|\gamma(F_{K,T})-T^2g\|_{C^{440}_{\mathrm{scaled}}}\le C(K^{-1}+K/T), \tag{5.1}$$

and the best phase margin of $F_{K,T}$ in these charts has a positive lower bound independent of K and T once $K^{-1}+K/T$ is small.

Proof. Take a maximal K^{-1}-separated net $\{x_i\}$ and cover M by the balls of radius $2K^{-1}$ about its points. Choose a squared partition of unity $\sum_i \chi_i^2 = 1$ with $\operatorname{supp}\chi_i$ in these balls and $\|\nabla^j \chi_i\|_\infty \le C_j K^j$; such functions are obtained by normalizing the sum of squares of rescaled bumps. Packing of disjoint smaller balls bounds the overlap multiplicity independently of K. We colour the intersection graph of the closed supports with a fixed number h of colours, increasing h if necessary so that S_h is available. At most one summand of a given colour is nonzero at any point. This map S_h is selected before the large parameters K and T. For $T \ge K$ and $1 \le \ell \le a$ we define

$$z_{\ell,c}(x) = \sum_{i \text{ of colour } c} \chi_i(x) \exp\Big(2\pi \mathrm{i} T (R_{x_i}\phi(x))_\ell\Big), \qquad c = 1, \dots, h.$$

Disjointness within each colour gives $|z_\ell(x)|^2 = \sum_i \chi_i(x)^2 = 1$, so z_ℓ takes values in S^{2h-1} and $|F_{K,T}| \le \sup|S_h|$.

On $\operatorname{supp}\chi_i$ we have $R_{x_i}\mathrm{d}\phi_x - L_x = O(K^{-1})$. Hence, for a local vector field v of bounded length,

$$T^{-1} v(z_{\ell,c}) = 2\pi\mathrm{i}\,(L_x v)_\ell\, z_{\ell,c} + O(K^{-1} + K/T).$$

More precisely, the error is a sum with coefficients $T^{-1} v\chi_i$ or $2\pi\mathrm{i}\chi_i(R_{x_i}\mathrm{d}\phi_x - L_x)v$. In charts scaled by T^{-1} their derivatives at each fixed order satisfy the same bound: a derivative falling on a cutoff costs K/T, the derivatives of positive order of the exponent are bounded, and derivatives falling on slow data gain T^{-1}. Since $2\pi\mathrm{i} z_\ell = Y_\ell(z_\ell)$, the scaled first derivatives of $F_{K,T}$ are, up to $O(K^{-1} + K/T)$, the Hopf derivatives $Y_{L_x v} S_h$. Differentiating once more in the scaled frame gives the same error bound for the second derivative columns, whose leading terms are the successive Hopf derivatives $Y_{L_x u} Y_{L_x v} S_h$. The terms in which L_x, the source frame or $f(x)$ are differentiated carry a factor T^{-1}. These estimates hold at each fixed scaled derivative order, since the number of overlapping supports is bounded. Since S_h induces g_W on D_W and L_x is an isometry, the leading Gram matrix of the first columns with respect to a g-orthonormal frame is the identity. The leading normal projection and second fundamental form are exactly those of S_h along D_W at the image point, composed with L_x.

We use the charts $x = \exp_p(y/T)$ with centres p in a maximal T^{-1}-separated net. Their scaled transitions $y \mapsto T \exp_q^{-1}(\exp_p(y/T))$ have uniformly bounded derivatives at every fixed order, by smoothness on a compact manifold and rescaling, and packing again bounds the overlap multiplicity. In these charts $T^2 g$ equals $\mathrm{Id} + O(T^{-1})$ with uniform derivative bounds. This proves (i). The previous estimates give (ii) and (5.1). The best phase margin of S_h has a positive minimum on the compact domain of the model, and the normal projector depends smoothly on the first columns. Hence the phase margin of $F_{K,T}$ is bounded below for $K^{-1} + K/T$ small. □

Proof of Theorem 2.2. Apply Lemmas 5.1 and 5.2 with $K = T^{1/2}$, so that $K^{-1} + K/T = 2T^{-1/2}$. In the scaled charts all the data required by Proposition 3.1, namely chart radii, overlap multiplicity, transition bounds, ellipticity, reference derivative bounds and phase margin, are uniform in T. We apply Proposition 3.1 with a fixed tolerance $\varepsilon_0 \le 1$ in the scaled charts. Its threshold η_0 does not depend on the number of charts, hence not on T. By (5.1), for T large the proposition yields a phase-regular map F_T^* with

$\gamma(F_T^*) = T^2 g$ and $|F_T^* - F_{T^{1/2},T}| \leq 1$. The map F_T^*/T induces g, is phase regular, since scaling by a constant preserves tangent and normal spaces, and satisfies

$$\sup_M |F_T^*/T| \leq \frac{\sup|S_h| + 1}{T}.$$

We choose T so that its image lies in the ball of half the prescribed radius centred at the origin. We then apply Proposition 4.4 with a C^0 displacement smaller than the remaining radius, and translate to the prescribed centre. □

6. The four-section construction at dimension $s_n + 5$

In this section we describe the geometry that a compact model for a general tangent bundle TM has to accommodate. It accounts for the normal ranks used later, and for the difference between a construction of formal columns and the partial-metric hypothesis of Theorem 2.2.

6.1. **The three-dimensional singular set.** Let M be closed with $n \geq 5$. A generic bundle map $\Psi : \underline{\mathbb{R}}^4 \to TM$ has rank at least three everywhere, and it drops rank exactly on its rank-three locus P, a closed embedded three-manifold (possibly empty) along which Ψ is transverse to the rank stratification. Indeed, in a matrix chart in which a j-by-j minor is invertible, block elimination writes the matrix as an invertible lower triangular factor, times the block diagonal matrix formed by the minor and its Schur complement, times an invertible upper triangular factor. Hence rank j is equivalent to the vanishing of the $(n-j)$-by-$(4-j)$ Schur complement, and the rank-j stratum has codimension $(n-j)(4-j)$. For $j \leq 2$ this is at least $2(n-2) > n$, while for $j = 3$ it is $n - 3$. Transversality follows from local matrix perturbations, cut off on a finite cover, and the parametric transversality theorem. On P we write $K = \ker \Psi|_P$ for the kernel line and $Q = (\operatorname{im} \Psi|_P)^\perp$ for the cokernel, of rank $n - 3$. The normal derivative of the Schur complement identifies νP with $\operatorname{Hom}(K, Q)$.

On a sufficiently small tube around P we use the Schur complement itself as the normal coordinate $v \in \operatorname{Hom}(K, Q)$. In the splittings $\underline{\mathbb{R}}^4 = K^\perp \oplus K$ and $TM = \operatorname{im}(\Psi|_P) \oplus Q$, transported radially from P, block elimination factors Ψ into the diagonal blocks given by the rank-three isomorphism and by v, with invertible triangular factors on the two sides. Multiplying the off-diagonal entries of these factors by a radial cutoff switches them off on a smaller tube, keeps both factors invertible and creates no new rank changes. Let A_0 be the invertible block at the zero section and $U_0 = A_0(A_0^* A_0)^{-1/2}$ its polar isometry. After multiplication by U_0^{-1}, the block equals $(A_0^* A_0)^{1/2}$ at the zero section, so its symmetric part is positive definite on a smaller tube, and convex interpolation with the identity preserves invertibility. Altogether we obtain a prepared form

$$\Psi_{(p,v)} = \begin{pmatrix} U_0(p) & 0 \\ 0 & v \end{pmatrix} : K_p^\perp \oplus K_p \longrightarrow \operatorname{im}(\Psi_p) \oplus Q_p$$

on a neighbourhood of P, with $U_0 : K^\perp \to \operatorname{im}(\Psi|_P)$ a bundle isometry over P.

Fix an orthogonal complex structure J of $\underline{\mathbb{R}}^4$ and set $L = JK$ and $D_2 = (K \oplus L)^\perp$. Write

$$A = K \oplus D_2 = L^\perp, \qquad H = Q \oplus D_2, \qquad m = n - 4.$$

Via U_0, the pulled-back tangent bundle is $H \oplus L$ on the tube. Let $\Sigma = S(\mathrm{Hom}(K, Q))$ be the unit normal sphere bundle, of dimension $n - 1 = m + 3$. A point $v \in \Sigma$ is an isometric embedding of the line K_p into Q_p, and we put $E_v = Q \ominus vK$, a bundle of rank m over Σ. The link identification

$$E_v \oplus A \longrightarrow H, \qquad (e, k, b) \longmapsto (e + v(k), b),$$

uses v on the summand K of A. It is a metric identification and requires no choice of a unit vector in K, so it is defined even when K is nonorientable.

For a Euclidean bundle U write $J_2(U) = U \oplus \mathrm{Sym}^2 U$, and put $j_k = \mathrm{rank}\, J_2(\mathbb{R}^k) = k(k+3)/2$. The link identification splits $J_2(H)$ orthogonally into $J_2(E_v)$ and the small summand

$$A \oplus (E_v \otimes A) \oplus \mathrm{Sym}^2 A,$$

of rank $3m + 9$. The same small summand is the orthogonal complement of $v(K) \otimes A$ in the bundle

$$\mathcal{S} = A \oplus (Q \otimes A) \oplus \mathrm{Sym}^2 A$$

over P, of rank $3m + 12$. The second description allows us to compare the full columns over the whole link. We shall use the following standard facts.

Lemma 6.1. *Let X be a compact manifold of dimension d, possibly with boundary, and let U_1, U_2 be Euclidean bundles over X of ranks $k \le \ell$, with rank gap $g = \ell - k$.*

(i) *If $g + 1 > d$, every isometric monomorphism $U_1 \to U_2$ given near a closed subset $X_0 \subset X$ extends to an isometric monomorphism over X which agrees with it near X_0. In particular, if $g > d$, any two isometric monomorphisms $U_1 \to U_2$ are homotopic through isometric monomorphisms, with stationary endpoints.*

(ii) *Let $Z_t : U_1 \to U_2$, $t \in [0, 1]$, be a smooth path of isometric monomorphisms. The skew endomorphism*

$$\Omega_t = \dot{Z}_t Z_t^* - Z_t \dot{Z}_t^* - Z_t (Z_t^* \dot{Z}_t) Z_t^*$$

of U_2 satisfies $\Omega_t Z_t = \dot{Z}_t$. Its linear flow O_t, with $O_0 = \mathrm{Id}$, is orthogonal, satisfies $O_t Z_0 = Z_t$, maps $(\mathrm{im}\, Z_0)^\perp$ onto $(\mathrm{im}\, Z_t)^\perp$, and is the identity wherever Z_t is independent of t. Similarly, for a smooth path of orthogonal projectors Π_t, the skew generator $[\dot{\Pi}_t, \Pi_t]$ has an orthogonal flow V_t with $\Pi_t V_t = V_t \Pi_0$.

Proof. In $\mathrm{Hom}(\mathbb{R}^k, \mathbb{R}^\ell)$ the matrices of rank $j < k$ form a submanifold of codimension $(k - j)(\ell - j) \ge g + 1$. For (i), multiply the given monomorphism by a cutoff equal to one near X_0, perturb the resulting bundle map away from X_0 by relative transversality so that it avoids all rank-loss strata, and replace the injective map Z so obtained by its polar normalization $Z(Z^*Z)^{-1/2}$; this does not change Z where it was already isometric. The homotopy statement is (i) on $X \times [0, 1]$, relative to $X \times \{0, 1\}$. For (ii), differentiating $Z_t^* Z_t = \mathrm{Id}$ shows that $Z_t^* \dot{Z}_t$ is skew, whence $\Omega_t^* = -\Omega_t$ and $\Omega_t Z_t = \dot{Z}_t$. Thus O_t is

orthogonal, and $O_t Z_0$ and Z_t solve the same linear equation with the same initial value. For projectors, $\Pi_t \dot{\Pi}_t \Pi_t = 0$ gives $[[\dot{\Pi}_t, \Pi_t], \Pi_t] = \dot{\Pi}_t$, and the same argument applies. $\square$

Put $R = \operatorname{rank} \mathcal{S} - 3 = 3m + 9$. If $R > \dim \Sigma$ and $j_m > \dim \Sigma$, there is an isometric monomorphism $K \otimes A \to \mathcal{S}$ over P, and its pullback to Σ and the moving inclusion $k \otimes a \mapsto v(k) \otimes a$ are homotopic through isometric monomorphisms. Indeed, rank failure for an inclusion with rank gap R has codimension $R + 1$, which exceeds the dimension of the cylinder $\Sigma \times [0, 1]$; one extends a path of bundle maps from the endpoints, uses relative transversality and applies polar normalization (Lemma 6.1(i)). If Z_t is the resulting isometric path, the skew matrix

$$\dot{Z}_t Z_t^* - Z_t \dot{Z}_t^* - Z_t (Z_t^* \dot{Z}_t) Z_t^*$$

takes Z_t to $\dot{Z}_t$. Its linear orthogonal flow R_t transports the complement of the moving inclusion, that is, the small summand above, to a fixed small bundle S_0 pulled back from P. The two inclusions of S_0 into $J_2(H)$ obtained in this way, namely an inclusion i_0 over P and the small summand composed with R_1^{-1}, can be compared by the same procedure, now with rank gap j_m. The resulting orthogonal flow T_t of $J_2(H)$ over Σ gives an isometry T_1 of $J_2(E_v)$ over the whole link onto a bundle J_0 pulled back from P, and it can be used to clutch the old column bundle on the exterior to one on the disk tube. The two rank inequalities hold for $n \geq 7$: one has $R - \dim \Sigma = 2m + 6$ and $j_m - \dim \Sigma = (m - 2)(m + 3)/2$.

The corresponding extension of formal columns is as follows. We write $J_2(H) = J_0 \oplus i_0(S_0)$ for the fixed splitting over P obtained after the second transport. Let $I_0 : J_0 \to W_0$ and $I_{\mathcal{S}} : \mathcal{S} \to W_0$ be isometric inclusions, over P, with orthogonal ranges in a target bundle W_0. On the link we prescribe the inclusion

$$x + y \longmapsto I_0(T_1 x) + I_{\mathcal{S}}(y), \qquad x \in J_2(E_v), \quad y \in A \oplus (E_v \otimes A) \oplus \operatorname{Sym}^2 A.$$

First apply R_t in the second summand, keeping the first one unchanged. At $t = 1$ the result equals $I_* T_1$, where

$$I_*(x_0 + i_0 s) = I_0(x_0) + I_{\mathcal{S}}(s), \qquad x_0 \in J_0, \ s \in S_0,$$

is independent of v. Then follow $I_* T_t$ backwards to I_*. Placing these two paths on radial annuli with stationary endpoints, and using I_* near the zero section, we obtain a smooth monomorphism on the disk bundle which coincides with the prescribed one on a boundary collar. This is a statement about formal column data only.

The coefficient inclusion itself can also be extended across the tube, without changing the number of original Hopf factors. On the boundary collar, polar normalization of the prepared map sends $D_2 \oplus L$ identically to itself and $v(K)$ to K by v^*. Accordingly, if $I_Q : Q \to \underline{\mathbb{R}}^a$ is an isometric inclusion over P, the link inclusion sends E_v through I_Q, vK through $v^* : vK \to K$, and $D_2 \oplus L$ identically into $\underline{\mathbb{R}}^4$. For $q \in Q$, $b \in D_2$ and $\ell \in L$, the map

$$(q, b, \ell) \longmapsto \Big(I_Q(\pi_{E_v} q + \cos(\theta)\, \pi_{vK} q), \ \sin(\theta)\, v^* q + b + \ell \Big)$$

is isometric for $0 \leq \theta \leq \pi/2$, since $|v^* q| = |\pi_{vK} q|$. Choose θ as a function of the radius, equal to $\pi/2$ on a boundary collar, where the map is the link inclusion above, and equal

to zero near the zero section, where the map becomes $(q, b, \ell) \mapsto (I_Q(q), b + \ell)$ and is independent of v. This gives a coefficient inclusion of TM across the whole tube.

6.2. **The dimension count.** In the four-direction construction the old and active Euclidean blocks have dimensions $s_n - 152$ and 157. The full old columns $J_2(E)$ use j_m dimensions, and the four active first columns and ten symmetric second columns use $4 + 10 = 14$. Thus, if the corresponding partial maps have these ranks, the remaining dimensions in the two blocks are

old block complement	$s_n - 152 - j_m = 3m - 142$
active block complement	$157 - 14 = 143$
total complement	$3m + 1.$

Here we used $s_n = s_m + 4m + 10$, and the count involves $(s_n - 152) + 157 = s_n + 5$ coordinates.

On the core, the bundle $J_2(H) \oplus L$ has rank $j_{n-1} + 1 = s_n$, so any isometric column inclusion of it into $\mathbb{R}^{s_n+5}$ has a complement of rank five. For any such inclusion over the zero section, two perpendicular copies of L can be chosen in this complement: rank failure in $\mathrm{Hom}(2L, \mathbb{R}^5)$ has codimension at least four, which exceeds $\dim P$. Such a choice extends along any given homotopy of column inclusions on the disk bundle and a product collar, by orthogonal transport of their complements. For instance, if Π_t denotes the projector onto the complement, one may use the skew generator $[\dot{\Pi}_t, \Pi_t]$ of Lemma 6.1(ii).

6.3. **The characteristic-class condition.** Write

$$r = n - 3, \qquad a = w_1(Q), \qquad b = w_2(Q), \qquad c = w_3(Q), \qquad t = w_1(L)$$

on P. Let N_5 be the complement of an isometric monomorphism $J_2(H) \oplus L \to P \times \mathbb{R}^{s_n+5}$ over P. Such a monomorphism exists, and is unique up to isometric homotopy, by Lemma 6.1, since the rank gap is five and $\dim P = 3$.

Proposition 6.2. *On every component of P,*

$$w_3(N_5) = \begin{cases} 0, & n \text{ odd or } n \equiv 0 \pmod 4, \\ w_1(\det \mathrm{Hom}(K, Q))^3, & n \equiv 2 \pmod 4. \end{cases}$$

Three linearly independent sections of N_5 exist on a component precisely when this class vanishes there. If $n \equiv 2 \pmod 4$, put $\mathcal{D} = \det \mathrm{Hom}(K, Q)$; then $N_5 \simeq 3\mathcal{D} \oplus \underline{\mathbb{R}}^2$.

Proof. All characteristic classes are taken over $\mathbb{F}_2$, and we work on one component of P. Since J identifies K with L, we have $H \oplus 2L \simeq Q \oplus \underline{\mathbb{R}}^4$. Expanding Sym^2 of both sides, and using $\mathrm{Sym}^2(2L) \simeq \underline{\mathbb{R}}^3$ and $[H \otimes L] = [Q \otimes L] + 4[L] - 2[\underline{\mathbb{R}}]$, we obtain in stable bundle notation

$$[\mathrm{Sym}^2 H] = [\mathrm{Sym}^2 Q] + 4[Q] - 2[Q \otimes L] - 8[L] + 11[\underline{\mathbb{R}}].$$

The first columns have bundle $H \oplus L$, with $[H \oplus L] = [Q] + 4[\underline{\mathbb{R}}] - [L]$. Since $[J_2(H) \oplus L] + [N_5]$ is trivial, it follows that

$$w(N_5) = w(\mathrm{Sym}^2 Q)^{-1} w(Q)^{-5} w(Q \otimes L)^2 (1 + t)^9.$$

Put $r_2 = r \bmod 2$. Since $\dim P = 3$, we have $w(Q)^4 = (1+t)^8 = 1$ and $w(Q \otimes L)^2 = 1 + (a + r_2 t)^2$. Hence, before multiplying by the first factor, the terms of degrees one through three are

$$a + t, \qquad b + at + r_2 t^2, \qquad c + bt + r_2 a t^2 + r_2 t^3.$$

The positive-degree terms of $w(\mathrm{Sym}^2 Q)$ through degree three are

$$(r-1)a, \qquad \tbinom{r-1}{2} a^2 + r_2 b, \qquad \tbinom{r-1}{3} a^3 + r_2(ab + c).$$

These formulas follow by expanding $\prod_{i<j}(1 + x_i + x_j)$ in the formal roots x_i of Q. In degree three, the coefficients of x_1^3, $x_1^2 x_2$ and $x_1 x_2 x_3$ determine the coefficients in the basis a^3, ab, c.

The relations $TM|_P = TP \oplus (Q \otimes L)$ and $[TM|_P] = [Q] + 4[\underline{\mathbb{R}}] - [L]$ give

$$w_1(TP) = (r+1)t, \qquad at = \Big(\tbinom{r}{2} + r\Big) t^2,$$

the first in degree one and the second in degree two. For the latter we used $w_2(TP) = w_1(TP)^2$, which follows from the Wu formula and $v_2 = 0$ in dimension three. Inverting and multiplying the truncated polynomials above gives

$r \bmod 4$	$w_3(N_5)$
0	$c + (a+t)b + a^2 t$
1	$(a+t)t^2$
2	$c + (a+t)b$
3	$(a+t)^3$.

For $r \equiv 0 \pmod 4$ we have $at = 0$. Hence in both cases with r even, where $w_1(TP) = t$, the class is $Sq^1 b + w_1(TP) b = 0$, by $Sq^1 b = ab + c$ and the first Wu formula $Sq^1 u = w_1(TP) u$ for $u \in H^2$ of a closed three-manifold. For $r \equiv 1 \pmod 4$ we have $at = t^2$, so $(a+t)t^2 = 0$. For $r \equiv 3 \pmod 4$, $\mathcal{D} \simeq L^{\otimes r} \otimes \det Q$ with r odd, so $a + t = w_1(\mathcal{D})$. This proves the formula for $w_3(N_5)$.

If N_5 has three independent sections, then clearly $w_3(N_5) = 0$. Conversely, two independent sections exist by rank avoidance. The remaining rank-three bundle E_3 has a unit section exactly when its Euler class with orientation-twisted integer coefficients vanishes. This is the only obstruction, on the three-skeleton, since the unit sphere S^2 is simply connected. The Euler class is killed by two, by fibrewise negation, and reduces modulo two to $w_3(E_3) = w_3(N_5)$. By Poincaré duality the degree-three group with these coefficients is $\mathbb{Z}$ or $\mathbb{Z}/2$ on each component. In the first case it has no two-torsion; in the second, reduction modulo two is injective. This proves the equivalence.

If $r \equiv 3 \pmod 4$, then $at = 0$ and the same computation gives $w_1(N_5) = a + t$ and $w_2(N_5) = a^2 + at + t^2 = (a+t)^2$. Hence $w(N_5 \oplus \mathcal{D}) = (1 + w_1(\mathcal{D}))^4 = 1$, so $N_5 \oplus \mathcal{D}$ is orientable with $w_2 = 0$. An oriented bundle of rank at least four with $w_2 = 0$ over a three-complex is trivial: it lifts to a spin bundle, and the classifying space of the spin group has vanishing homotopy groups through degree three. The same applies to $4\mathcal{D} \oplus \underline{\mathbb{R}}^2$. We cancel the common line $\mathcal{D}$ by homotoping its two inclusions into the trivial six-plane and transporting the complements; the rank-failure codimension six exceeds the dimension four of the homotopy cylinder. □

Remark 6.3 (Occurrence of the exceptional case). Let $M = \mathbb{RP}^n$ with $n \equiv 6 \pmod 8$, $n \geq 14$, and write $x = w_1(\lambda)$ for the tautological line λ. The exact sequence $0 \to K \to \underline{\mathbb{R}}^4 \to TM|_P \to Q \to 0$ and the oddness of r give $w_1(\mathcal{D}) = w_1(Q) + w_1(K) = w_1(TM)|_P = x|_P$. The rank-three locus of any transverse preparation represents

$$\mathrm{PD}_M[P] = w_{n-3}(TM).$$

Indeed, on $M \times \mathbb{RP}^3$ the bundle map restricts to a section of $U^* \otimes TM$, with U the tautological line of the second factor. At a zero, the three projective directions map onto the image of the rank-three map and the normal directions of P onto its cokernel. Hence the section is transverse, and its zero set is the graph of the kernel line over P. If $u = w_1(U)$, the top Stiefel–Whitney class of $U^* \otimes TM$ is $\sum_j u^{n-j} w_j(TM)$ by the splitting principle. Integration over the projective fibre selects u^3 and gives the asserted dual class. Since $TM \oplus \underline{\mathbb{R}} = (n+1)\lambda$, it follows that

$$\langle x^3|_P, [P] \rangle = \binom{n+1}{4} \bmod 2 = 1.$$

The last parity follows from $(1+z)^{n+1} = (1+z)^7 (1+z^8)^{(n-6)/8}$ over $\mathbb{F}_2$. In particular, by Proposition 6.2, some component of P has $w_3(N_5) \neq 0$.

7. The reference map and the boundary class

Write

$$m = n - 4, \qquad N = s_n + 5, \qquad N_{\mathrm{o}} = s_n - 152,$$

so that $\mathbb{R}^N = \mathbb{R}^{N_{\mathrm{o}}} \oplus \mathbb{R}^{157}$. On the exterior of the small prepared tube the four sections have rank four, and we polar-normalize them. This gives an orthogonal splitting $TM = E \oplus A_4$, with $A_4 = \underline{\mathbb{R}}^4$. Since polar normalization divides the K-component by the radius, on the link collar E is the radial pullback of the original $E_v = (vK)^\perp \subset Q$. Fix an isometric inclusion I_Q of Q into a finite trivial bundle over P, and an isometric coefficient inclusion of E into the same trivial bundle which agrees on this collar with the restriction of I_Q. The latter exists by relative rank avoidance (Lemma 6.1(i)) in a sufficiently large trivial bundle: one first extends a bundle map off a smaller collar, then avoids its rank-loss strata and polar-normalizes. For A_4 we use the four original active slots. On the disk tube, the rotation of Section 6.1 extends these inclusions to an inclusion of TM. The line L is unchanged by the rotation. Thus the coefficient inclusion is fixed before the common Hopf size is chosen. Using the product collar, the coefficient inclusion and the exterior bundle extend to the doubled exterior, a closed n-manifold which we call the double.

We next specify the slow full column data, including their class along the attaching sphere bundle. Write $J_2(S) = S \oplus \mathrm{Sym}^2 S$. Since $\operatorname{rank} \mathrm{Hom}(K, Q) = m + 1 > 3$, rank avoidance gives a section $v_0 : K \to Q$ of the unit link bundle over P. The argument of Section 6.1, with v_0 as the fixed inclusion, gives an orthogonal homotopy T_t of $J_2(H)$ over the link, with $T_0 = \mathrm{Id}$, carrying $J_2(E_v)$ at its endpoint onto $J_0 = J_2(E_{v_0})$, pulled back from P. In detail, we first homotope the three-column inclusion $vK \otimes A$ in $A \oplus (Q \otimes A) \oplus \mathrm{Sym}^2 A$ to the one for v_0; its complement is transported to the fixed

small summand $A \oplus (E_{v_0} \otimes A) \oplus \operatorname{Sym}^2 A$ of rank $3m+9$. The two inclusions of that summand in $J_2(H)$ are then homotopic, since the rank gap is j_m, and extending this homotopy by orthogonal transport gives T_t. Both rank gaps exceed the dimension $m+4$ of the link cylinder.

Choose an isometric old monomorphism $I_{\mathrm{o}}^P : J_0 \to \underline{\mathbb{R}}^{N_{\mathrm{o}}}$ over P. The rank gap, over P and on the double, is $3m-142$. We extend the composite $I_{\mathrm{o}}^P \circ T_1$ on $J_2(E_v)$, taken radially constant on a collar, to an isometric formal old monomorphism

$$I_{\mathrm{o}} : J_2(E) \longrightarrow \underline{\mathbb{R}}^{N_{\mathrm{o}}}$$

on the double. Independently, we prescribe a constant isometric inclusion of the fourteen full active columns $J_2(A_4)$ into the 157-dimensional target block, with values in $\mathbb{R}^{143} \oplus \{0\}$; its complement is a fixed 143-plane containing the fourteen unused axes $\{0\} \oplus \mathbb{R}^{14}$. Over the link, since $A \oplus L = A_4$, we have the orthogonal decomposition

$$J_2(H) \oplus L = J_2(E_v) \oplus (A_4 \oplus \operatorname{Sym}^2 A) \oplus (E_v \otimes A)$$

into summands of ranks j_m, 10 and $3m$.

On the core, we extend the inclusion I_{o}^P of J_0 to an isometric inclusion of $J_2(H) \oplus L$ into $\underline{\mathbb{R}}^N$: the complement of J_0 has rank $3m+10$, the complement of the old columns in $\mathbb{R}^N$ has rank $3m+15$, so the remaining rank gap is five, and the base has dimension three. On the link we precompose this inclusion with $T_1 \oplus \mathrm{Id}_L$, so that its old columns agree with the slow old inclusion I_{o} specified above. In the orthogonal complement of the old columns, which has rank $3m+15$, its ten active columns $A_4 \oplus \operatorname{Sym}^2 A$ can be moved by an isometric homotopy to the corresponding columns in the fixed active block. The rank gap here is $3m+5$, larger than $m+4$, so this relative homotopy of the ten columns follows by rank avoidance on the link cylinder, and we extend it by orthogonal transport of the complement. We have thus obtained the slow augmented class over the link, and it extends over the disk: on a radial collar we run these two homotopies backwards to the original inclusion over P. At the attaching end, the comparison keeps the old columns and the ten active columns at their prescribed placements. We denote by J_Σ the resulting mixed reference inclusion of $E_v \otimes A$ into the complement of the old columns and the ten active base columns. Its normal bundle is the pullback of the rank-five bundle N_5 of Proposition 6.2.

The slow retained complement W on the double is the sum of the complement of the formal old columns and the fixed 143-plane in the active block. It has rank $(3m-142)+143 = 3m+1$, and on the collar it is pulled back from P, since there the old columns $I_{\mathrm{o}}(J_2(E_v)) = I_{\mathrm{o}}^P(J_0)$ are independent of v. Adding the four cofactor axes, namely the images of the columns $(L \otimes A) \oplus \operatorname{Sym}^2 L$ of the fixed active inclusion, gives the complement of the old columns and the ten selected active columns. Restricted to the marked section v_0, this complement is identified with $(E_{v_0} \otimes A) \oplus N_5$ by the class of the full injection. This gives the two bundle identifications for the ordinary and the exceptional moment models used in Sections 10 and 8.

On the original Hopf product of any fixed size at least two, we realize the formal old inclusion on the closed double by Proposition 14.1, with $K = \emptyset$ (the double is closed) and the unit old metric; this applies since $N_{\mathrm{o}} \geq j_m + 5$. We obtain an exact unit old

block F_{o} together with a homotopy of its full columns to the formal old ones. After polar normalization of the columns, orthogonal complement transport along this homotopy gives a smooth isometry of the slow old normal complement onto the complement of the full old columns of F_{o}. In the active block we use the map B of Section 8.6. The homotopy of its full active columns that gives the normal identification can be chosen to fix the fourteen unused axes throughout. For the rank-nine bank over a core we use the map η_κ of that section, with a positive constant κ on an annulus around the link and with a slow cutoff to zero before the end of a larger collar. We transport the normal complement from $\eta_0 = \nu$ to η_κ within the complement of the full active columns of B. This modifies the active normal isometry by an orthogonal path which is the identity off the collar. In particular it realizes the bank η_κ without changing the old part of the normal isometry. With these choices,

$$F_0 = (F_{\mathrm{o}}, B) : \text{original Hopf product over the double} \longrightarrow \mathbb{R}^{N_{\mathrm{o}}} \oplus \mathbb{R}^{157}$$

is the unit free product reference of Section 11.5. The old part of its full normal isometry I_x is independent of the active variables, and the active part is independent of the old variables. On every annulus used below, these frames and the slow augmented class determine the same data.

8. The active map on an exceptional link

Throughout this section we assume

$$n \equiv 2 \pmod 4, \qquad r = n - 3 \geq 10^{62}.$$

In this residue class Proposition 6.2 gives $N_5 \simeq 3\mathcal{D} \oplus \underline{\mathbb{R}}^2$, and three independent normal sections need not exist over P. We therefore work directly over the unit sphere bundle $\Sigma = S(Q \otimes L)$, which has dimension $r + 2$.

Write $A_4 = \underline{\mathbb{R}}^4 = L \oplus A$ with $A = L^\perp$ of rank three. An isometric identification of $\mathbb{R}^4$ with the quaternions $\mathbb{H}$ trivializes $U = L \otimes A$: for a unit vector $\ell \in L$ the vectors $\ell\mathbf{i}, \ell\mathbf{j}, \ell\mathbf{k}$ form an orthonormal basis of $A = \ell^\perp$, and the sections

$$e_1 = \ell \otimes \ell\mathbf{i}, \qquad e_2 = \ell \otimes \ell\mathbf{j}, \qquad e_3 = \ell \otimes \ell\mathbf{k}$$

of U do not depend on the sign of ℓ. Via the metric contraction $L \otimes L \simeq \underline{\mathbb{R}}$ we also identify $A \simeq L \otimes U$. Put $V = A_4 \oplus \mathbb{R}z$, with the inclusion $I : A_4 \hookrightarrow V$, and define over Σ

$$\mathcal{H}_v = E_v \oplus L\mathcal{D}, \qquad L\mathcal{D} = L \otimes \mathcal{D}.$$

The line $L\mathcal{D}$ belongs to the old generating bundle $\mathcal{H}$, which has rank r.

By rank avoidance choose two orthogonal copies of L in Q, so that $Q = D \oplus L \oplus L$ with $\operatorname{rank} D = r - 2$. To compare the old full columns we use the section of Σ given by the first copy of L. Along this section the transported bundle J_0 of Section 7 restricts to $J_2(Q - L)$, and the four first and ten second active columns have bundle

$$A_4 \oplus \operatorname{Sym}^2 A \oplus (L \otimes A) \oplus \operatorname{Sym}^2 L.$$

The last two summands are trivial, of total rank four. Hence the retained complement W of the old full columns and the ten active columns $A_4 \oplus \mathrm{Sym}^2 A$ has stable class

$$W + \underline{\mathbb{R}}^4 = \Big((Q - L) \otimes A\Big) + N_5.$$

Since $L \otimes A \simeq \underline{\mathbb{R}}^3$ and $N_5 \simeq 3\mathcal{D} \oplus \underline{\mathbb{R}}^2$, stable cancellation gives

$$W \simeq (D \otimes A) \oplus 3\mathcal{D} \oplus \underline{\mathbb{R}}, \qquad w = \operatorname{rank} W = 3r - 2.$$

For the cancellation we regard both sides as orthogonal complements of two isometric inclusions of the common trivial summand $\underline{\mathbb{R}}^4$ into one bundle of rank $w+4$; by the rank stratification argument these inclusions are homotopic through isometric inclusions, and we transport their complements along the homotopy. This applies because the rank gap w exceeds the dimension of the three-dimensional base and of its homotopy cylinder. The bundle W over Σ is pulled back from P.

8.1. **The angular seed and its four cofactor rows.** The splitting $Q = D \oplus L \oplus L$ and the contraction $L \otimes L \simeq \underline{\mathbb{R}}$ give $Q \otimes L \simeq (D \otimes L) \oplus \underline{\mathbb{R}} \oplus \underline{\mathbb{R}}$, and we write accordingly

$$v = (v_D, p, q), \qquad |v_D|^2 + p^2 + q^2 = 1.$$

Using $A \simeq L \otimes U$, a tensor in $\mathcal{H}_v \otimes A$ has coordinates (Z, x, y, ζ) with

$$Z \in (D \otimes L) \otimes U, \qquad x, y \in U, \qquad \zeta \in (L\mathcal{D}) \otimes A \simeq \mathcal{D} \otimes U = 3\mathcal{D},$$

subject to the condition that (Z, x, y) be orthogonal to the image of v, namely

$$v_D^* Z + px + qy = 0, \tag{8.1}$$

where v_D^* contracts the factor $D \otimes L$ with v_D.

In the basis (e_1, e_2, e_3) of U let

$$S_0 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \qquad P_0 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix};$$

thus S_0 is tridiagonal with zero diagonal, and P_0 interchanges e_1, e_3 and kills e_2. Identifying W with $((D \otimes L) \otimes U) \oplus 3\mathcal{D} \oplus \underline{\mathbb{R}}$, we define the retained seed T_A and the four base cofactor rows C_{base}, with values in $U \oplus \underline{\mathbb{R}} \simeq \underline{\mathbb{R}}^4$, by

$$T_A(Z, x, y, \zeta) = (Z + v_D \otimes P_0 y,\ \zeta,\ y_2), \qquad C_{\mathrm{base}}(Z, x, y, \zeta) = (x + S_0 y,\ y_3).$$

Lemma 8.1. *The augmented map (T_A, C_{base}) is injective on $\mathcal{H}_v \otimes A$ for every $v \in \Sigma$. Its restriction to $E_v \otimes A$ has normal bundle in $W \oplus \underline{\mathbb{R}}^4$ isomorphic to $\underline{\mathbb{R}}^2 \oplus 3\mathcal{D} \simeq \pi_\Sigma^* N_5$.*

Proof. If $(T_A, C_{\mathrm{base}})(Z, x, y, \zeta) = 0$, then $\zeta = 0$ and $y_2 = y_3 = 0$, and since $S_0 e_1 = e_2$, $P_0 e_1 = e_3$,

$$y = d_1 e_1, \qquad x = -d_1 e_2, \qquad Z = -d_1 v_D \otimes e_3, \qquad \zeta = 0.$$

Then (8.1) gives

$$d_1\Big(-|v_D|^2 e_3 - p e_2 + q e_1\Big) = 0.$$

The vector $\mathbf{w}(v) = -|v_D|^2 e_3 - p e_2 + q e_1$ satisfies $|\mathbf{w}(v)|^2 = (|v_D|^2 - \frac{1}{2})^2 + \frac{3}{4} > 0$ for every unit v, so $d_1 = 0$.

For the normal bundle we note that $\zeta = 0$ on $E_v \otimes A$, and we temporarily omit both (8.1) and ζ. The map of (Z, x, y) to the other retained outputs and the cofactor outputs goes from a space of dimension $3r$ to one of dimension $3r - 1$, and by the computation above its kernel is precisely the line spanned by $(-v_D \otimes e_3, -e_2, e_1)$; hence it is surjective. The constraint map $(Z, x, y) \mapsto v_D^* Z + px + qy$ is onto U, has kernel $E_v \otimes A$, and sends that line onto $\mathbb{R}\mathbf{w}(v)$. Imposing (8.1) therefore leaves the cokernel $U/\mathbb{R}\mathbf{w}(v)$. The unit vector $\mathbf{w}/|\mathbf{w}|$ has nonpositive e_3-component, so its normalized straight-line homotopy to $-e_3$ is defined, and the quotient is trivial of rank two. The unused summand $3\mathcal{D}$ of W accounts for the other three normals, and $\underline{\mathbb{R}}^2 \oplus 3\mathcal{D} \simeq \pi_\Sigma^* N_5$ by Proposition 6.2. □

This normal identification allows us to move the seed into the prescribed formal class. After the independent base columns are identified, let $J_{\mathrm{ref}} = J_\Sigma$ be the prescribed mixed injection of $E_v \otimes A$ into $W \oplus \underline{\mathbb{R}}^4$; its normal bundle is N_5. We polar-normalize both injections. By Lemma 8.1 and Proposition 6.2 their normal bundles are isometric, so there is an orthogonal bundle automorphism of $W \oplus \underline{\mathbb{R}}^4$ carrying the normalized seed injection to J_{ref}. The restriction of this automorphism to the four extra axes and the original inclusion of these axes are homotopic through isometric four-frames, since the rank-failure codimension in a path is

$$w + 1 = 3r - 1 > \dim \Sigma + 1.$$

We extend this homotopy by orthogonal transport. At its endpoint the resulting automorphism fixes the four axes, and so acts by an orthogonal change of W alone. Applying this change to the retained seed places its augmented mixed injection in the prescribed class. Along the comparison the four extra axes may move, but at the retained endpoint they are unchanged; the positive polar normalizations are themselves homotopies through injections. If the retained tensor is presented in moment form, the same change conjugates the selfadjoint moment operator and postcomposes the injector.

The rows C_{base} arise by contracting A against the reader

$$Q \longrightarrow \mathcal{B} = L \otimes \big(\mathrm{SymEnd}(U) \oplus U^*\big), \qquad (h_D, \ell_1, \ell_2) \longmapsto \big(\ell_1 \otimes \mathrm{Id}_U + \ell_2 \otimes S_0,\ \ell_2 \otimes e_3^*\big),$$

extended by zero on $L\mathcal{D}$; here $\mathrm{SymEnd}(U)$ denotes the symmetric endomorphisms of U, and $\mathcal{B} \simeq 9L$. By rank avoidance we split a copy of $9L$ off D and identify it with $\mathcal{B}$. Adding to the reader above a small positive multiple of the orthogonal projection onto these coordinates, we obtain a reader $R_{\mathrm{amb}} : Q \to \mathcal{B}$ which is onto. Let $R_v = (R_{v,\mathrm{sym}}, R_{v,*})$ be its restriction to $\mathcal{H}_v$, extended by zero on the generating line $L\mathcal{D}$. Since E_v is a hyperplane in Q, R_v has rank nine except on the sphere Σ_8 of $\mathrm{im}\, R_{\mathrm{amb}}^*$ in $Q \otimes L$, where it has rank eight. This angular stratum has codimension $(r-1)-8 = r-9$ in Σ. Define the new cofactor rows by

$$C_0(h \otimes d) = \big(R_{v,\mathrm{sym}}(h)d,\ R_{v,*}(h)d\big), \qquad h \in \mathcal{H}_v,\ d \in A,$$

using the contraction with L; the formula does not depend on the choice of a local unit of L. If the added multiple is small, the compact margin of the starting augmented injection persists under the change from C_{base} to C_0.

Lemma 8.2. *For every $v \in \Sigma$, including $v \in \Sigma_8$, and every complex two-plane $B \subset A_{\mathbb{C}}$, the restriction $C_0|_{\mathcal{H}_{v,\mathbb{C}} \otimes B}$ is onto $\mathbb{C}^4$.*

Proof. Identify $\mathrm{SymEnd}(U_{\mathbb{C}})$ with $\mathrm{Sym}^2 U_{\mathbb{C}}$ by the bilinear trace pairing. A dual vector $(u, s) \neq 0$ killing the image would force the map $d \mapsto (u \odot d,\, sd)$, $d \in B$, to take values in the annihilator of $\operatorname{im} R_v$, which has dimension at most one. But this map is injective on the two-plane B: if $s \neq 0$ this is clear, and if $s = 0$ then $u \neq 0$ and multiplication by u is injective in the symmetric algebra. This contradiction proves surjectivity. □

The retained seed extends to the five labels as a moment tensor. Diagonalize P_0 in U, keeping e_2 as one eigenvector; this splits W orthogonally into three target blocks, and on each block the seed uses only the corresponding single L-valued reader on A. Assign distinct real scalar nodes to the three blocks, and interpolate the three readers, extended by zero to the two other labels. Thus the map on each block is a fixed old injector composed with the evaluation, at the node of the block, of one common L-valued polynomial system of degree at most seven. It agrees with T_A on A. A slow orthogonal change of the target carries this presentation along by conjugation.

In each block we split off trivial lines, by rank avoidance over P, until a summand of rank at most three remains. Let X be the sum of these residual summands; it has rank at most nine. We assign distinct simple real nodes to the trivial lines, arbitrarily close to the original nodes, and let the normalized rows on X vary freely in $\mathrm{Hom}(\mathcal{H} \otimes V, X)$. The simple rows keep the evaluation form, and each old injector covector varies freely. We call these data the *simple moment and residual-row model*. By compactness, the augmented injection at $I : A_4 \hookrightarrow V$ persists under sufficiently small changes of the nodes, the polynomial system, the injector covectors and the rows on X.

8.2. **Polynomial labels.** The five L-valued label polynomials can be chosen homogeneous of degree 70, with no common projective zero and with linearly independent fifteen pair products. To the initial polynomials, of degree at most seven, we add

$$\varepsilon\, t^{40+b_i} \beta(t), \qquad (b_i) = (0, 1, 3, 9, 27), \qquad \beta(t) = \sum_{j=0}^{3} t^j \ell_j,$$

where ℓ_j are the projections to L of the four ambient coordinate vectors. The polynomial β is nonzero in every fibre. In a linear relation among the pair products, the terms that involve an initial polynomial have degree at most 77. The remaining terms have degree at least 80 and equal $\varepsilon^2 t^{80} \beta^2$ times a linear combination of the powers $t^{b_i + b_j}$, $i \leq j$, which are distinct. This proves independence.

We homogenize and perturb the coefficients over P. A common real projective zero imposes five independent conditions, and one condition is returned by its location; a nonreal zero imposes ten real conditions and returns its two real coordinates, since 1 and t interpolate an arbitrary complex value with real coefficients in an affine chart. Likewise, proportional evaluation rows at two distinct simple nodes, or the failure of the rows at five such nodes to span at least four dimensions, imposes codimension at least four. This follows from surjectivity of polynomial interpolation at up to five fixed real nodes together with the codimension of the determinantal strata. Since P has

dimension three, finite local parameter sprays over P avoid all these incidences. The perturbations preserve the positive compact margin at I.

Write $\mathcal{P}$ for the resulting five-dimensional space of scalar polynomials over $\mathbb{C}$, in a local unit of L.

Lemma 8.3. *Let $f, g \in \mathcal{P}$ be linearly independent. If u is a rational function with $uf, ug \in \mathcal{P}$, then u is constant. The same holds after a subsystem of $\mathcal{P}$ is divided by a common scalar factor.*

Proof. The identity $f(ug) = g(uf)$ holds in $\operatorname{Sym}^2 \mathcal{P}$, by independence of the pair products. Comparing coefficients in a basis of $\mathcal{P}$ extending f, g gives $uf = \alpha f$ and $ug = \alpha g$ for a constant α. After division by a common factor h, multiplication by h^2 reduces the claim to the previous case. □

8.3. **A path of active derivatives.**

Theorem 8.4. *There is a normalized tensor $T : \mathcal{H} \otimes V \to W$ in the simple moment and residual-row model with the following property. Let Z_{a} be any fixed product of four original Hopf spheres. On $\Sigma \times Z_{\mathrm{a}}$ there are a smooth map $a_{\mathrm{a}} : \Sigma \times Z_{\mathrm{a}} \to V$ and a smooth path $D_s : A_4 \to V$, $s \in [0,1]$, of linear maps of rank four, from $D_0 = I$ to the active derivative $D_1 = \mathrm{d}_{\mathrm{a}} a_{\mathrm{a}}$, such that throughout the path*

$$T(\mathrm{Id} \otimes D_s)|_{\ker C_0} \quad \text{is injective}, \qquad \ker C_0 \subset \mathcal{H} \otimes A.$$

The four rows C_0 remain fixed, and the injection has a positive uniform margin on the compact path.

Proof. We use the complexified incidence spaces with fixed simple nodes and with the polynomial system on its open locus. The simple injector rows and all residual rows on X vary independently. In what follows every generic preparation is chosen within the starting compact injection margin at I.

Step 1: small kernel tensors. We exclude rank-one kernel tensors and pairs of distinct projective rank-two kernel tensors. A projective rank-one tensor has $r + 3$ parameters. Its scalar entry is a nonzero polynomial of degree 70, so it vanishes at no more than 70 simple nodes, and the kernel condition imposes at least $w - 70$ equations. A projective rank-two tensor has $2r + 5$ parameters. Two distinct such tensors are linearly independent over the rational functions, by Lemma 8.3. Hence their wedge has a nonzero minor of degree at most 140, which bounds the sum of their rank deficits over the simple nodes; on the free rows over X, their constant independence gives two conditions on each row. Thus, after the projective tensor parameters are returned, the codimensions are at least

$$w - 70 - (r + 3), \qquad 2w - 140 - 2(2r + 5).$$

Both exceed $r + 2 = \dim \Sigma$.

Step 2: the pencil on a principal slice. Fix a principal covector on A_4 which is nonzero on A, and let $B \subset A$ be its kernel in A, the frozen two-plane. Write $F : B \to V$ for the frozen injection, $F_i = F(b_i)$ in a basis of B, and put $K_B = \ker(C_0|_{\mathcal{H} \otimes B})$. By Lemma 8.2, $\dim K_B = 2r - 4$. At an allowed jet, $T(\mathrm{Id} \otimes F)$ embeds K_B in W. Choose

$e \in A$ on which the principal covector equals one and a linear lift $k_e(h) \in \mathcal{H} \otimes B$ with $C_0(k_e(h)) = C_0(h \otimes e)$. The remaining quotient of $\ker C_0$ is then tested by the affine pencil

$$h \longmapsto T(h, D(e)) - T(\mathrm{Id} \otimes F)k_e(h) \pmod{T(\mathrm{Id} \otimes F)K_B}.$$

It maps an r-dimensional space to a space of dimension $w - (2r - 4) = r + 2$.

Pass to the highest homogeneous part in $z_* = D(e)$. For a label outside $\mathbb{P}\langle F_1, F_2\rangle$, a rank defect of the homogeneous pencil is witnessed by a tensor

$$q_* = X_1 \otimes F_1 + X_2 \otimes F_2 + h_* \otimes z_*, \qquad h_* \neq 0, \tag{8.2}$$

in $\ker T$, with $C_0(X_1 \otimes b_1 + X_2 \otimes b_2) = 0$. By Step 1 there is at most one projective rank-two kernel tensor, and it contributes at most a projective two-plane of labels: its row plane must meet the frozen plane in a line and determines the three-plane containing both, and if it equals the frozen plane it contributes no label away from $\langle F_1, F_2\rangle$. Consequently a hypersurface of rank defects of the homogeneous pencil requires an irreducible family of rank-three witnesses whose marked moving labels have image of dimension at least three.

Step 3: rank-three families. For a generic member of such a family the vector polynomial of old vectors has a fixed common scalar divisor g, up to scale. It divides the greatest common divisor of the two frozen polynomials, so there are finitely many possibilities. If g is nonconstant, the subsystem divisible by g has dimension exactly four: dominance of the moving label forces at least four, and five would contradict the absence of common projective zeros. We divide this subsystem by g. This deletes at most one simple node, since two distinct evaluation hyperplanes cannot both contain the four-system. Let $q_0 \in \{0, 1\}$ be the number of deleted nodes, with $q_0 = 0$ in the undivided case. After the division, generic vector polynomials of the family are primitive, that is, they have no common scalar zero.

Take ℓ generic members of the family that are linearly independent over the rational functions, where ℓ is the smaller of the rational dimension of the family and $\lfloor 13r/20 \rfloor + 1$. Let d be the dimension of their combined old span and $\delta = 3\ell - d$. Let R_{rel} be the δ-plane of linear relations among the 3ℓ separate old columns of the tensors (8.2). At each projective parameter the ℓ evaluations in the separate blocks span an ℓ-plane; let $\nu \leq \delta$ be the dimension of the span of its intersections with R_{rel} as the parameter varies. Let d_1, d_2, d_3 be the rank deficiencies of the arrays of first, second and third old columns, and put $j = \max(\nu, d_1, d_2, d_3) \leq \delta$.

On the span of dimension ν, the blockwise wedge with the separate evaluations is generically injective by rational independence. A nonzero ν-minor has degree at most 70ν, and its order of vanishing at a node is at least the rank deficiency there, by local row elimination. Hence the total number of equations lost at the distinct simple nondeleted nodes is at most 140ν. On the free rows over X there are exactly $(\operatorname{rank} X)\ell$ equations, by constant independence. Together the retained rows impose at least $(w - q_0)\ell - 140\nu$ equations; the deleted node has no effect on the free rows.

For the count of source parameters we require R_{rel} to contain the span of ν independent lines in the ℓ-planes of separate evaluations. Each line costs at most ℓ parameters

(one for the node and $\ell - 1$ for ratios), while the containment imposes $d\nu$ conditions. Alternatively, we mark a d_i-dimensional subspace in the fixed ℓ-space of one block slot and require the relation plane to contain it; this saves $d_i(d - \ell + d_i)$ parameters. The surjective combined column matrix with fixed old span projects submersively to its relation Grassmannian, so we may use the largest of these savings, which is at least $j(d - \ell)$.

We temporarily treat the nine-dimensional reader as independently variable. The four cofactor equations on each member of the tuple impose at least $4\ell - 4(d_1 + d_2)$ equations on the reader: we use two diagonal symmetric coordinates for X_1 and X_2, one off-diagonal coordinate involving the third basis direction for X_1, and one scalar coordinate for X_1; in a complex basis we use dual coordinates for the symmetric form. These variations are independent of the retained row variables. For each projective tensor with fixed combined old d-space, the tensor itself takes $3d - 1$ parameters, the extension of the frozen two-plane to a three-plane takes two, and its marked label takes two. The combined d-space takes $d(r - d)$ parameters. In the divided four-system we save one extension parameter per tensor, which covers $q_0 \leq 1$. The fibre of generic tuples in the same witness family returns at least 3ℓ parameters. Thus the codimension at fixed F, and hence also jointly with F as a free parameter, is at least

$$\delta(r - d) + 2\ell + j(d - \ell - 148). \tag{8.3}$$

For a whole rational basis we have $j \geq 1$. Indeed, if $\nu = 0$, the evaluated vectors are independent everywhere, so the rational coefficients of a polynomial of the same degree in their span are regular at every point of $\mathbb{P}^1$, hence constant. If moreover $d_3 = 0$, the moving old columns are independent. Modulo the frozen pair, a generic combination with all coefficients nonzero then has rank equal to the dimension of the span of the moving labels, and for a generic basis tuple this span has dimension at least two by dominance; this contradicts the form (8.2). A whole basis of length one cannot occur: by Lemma 8.3 it gives a single projective tensor, whose marked image has dimension at most two.

For $\ell = \lfloor 13r/20 \rfloor + 1$ the bound (8.3) exceeds $r + 300$: if its last bracket is negative we use $j \leq \delta$ and obtain $2\ell + \delta(r - \ell - 148)$, and otherwise the bound is at least 2ℓ. For a smaller whole basis of length greater than 300, a nonnegative bracket gives at least $r + \ell - 148 > r + 150$, since $\delta, j \geq 1$; a negative bracket gives the same strict bound, by $\delta > 2\ell - 148$ and $\ell \leq 13r/20 + 1$. For length at most 300 and $\delta \geq 2$ we use $d \leq 900$ and $j \leq \delta$ to obtain a bound greater than $r + 150$.

Step 4: a whole basis with one relation. It remains to treat a whole rational basis of length T_0, $2 \leq T_0 \leq 300$, with $\delta = 1$, so that $d = 3T_0 - 1$. Project each separate three-block away from the projection of the one-dimensional relation. This projection factors through the combined old space and leaves at least two independent scalar entries of the original polynomial. In the projection of another member of the family, its rational coefficient in that block multiplies both entries, so it is constant by Lemma 8.3. The family therefore lies in the constant span of the chosen basis. Its unique old relation involves only the moving columns, since independence of these columns would lead to the contradiction with the label quotient obtained above.

Write b_i for the coefficients of the relation, g_i for the moving labels in the quotient by the frozen pair, and g_* for the rank-one row direction of a generic combination with coefficients $c_i \neq 0$. Every covector α annihilating g_* has $(c_i\alpha(g_i))_i$ proportional to $(b_i)_i$. Hence the g_i span exactly two dimensions, by dominance and because the evaluation of functionals restricted to the annihilator of g_* has rank one. Generically g_* differs projectively from each g_i. Consequently all $b_i \neq 0$, and the projective combination is determined by $[c_i] = [b_i/\alpha(g_i)]$ for a nonzero quotient covector α killing g_* in their two-dimensional span. The unmarked tensor family is thus one-dimensional and uses one four-system over the frozen pair.

For a generic unmarked tensor, the marked lifts sweep a dense subset of its projective two-plane, since the marked image has dimension at least three. Fix a canonical lift by a linear complement to the frozen pair. Changing the lift by the two frozen labels changes the first two old columns by multiples of the moving old column. Hence, for a generic unmarked tensor, the reader annihilates its moving column tensored with both frozen directions, in addition to the four equations on the canonical frozen columns. The cost for the reader is

$$7(T_0 - 1) + 4T_0 :$$

the two symmetric matrix columns have five independent entries and the scalar reader has two, and the canonical frozen columns are independent modulo the span of the moving columns.

At a simple nondeleted node at most one equation is lost, and a loss forces both frozen polynomials to vanish there, because every coefficient of the unique relation is nonzero. Let d_0 be the number of such nodes. Then $d_0 \leq 4$, since the rows at five nodes span at least four dimensions in the original five-system. If $d_0 \geq 3$, requiring F to vanish under two of these evaluations imposes four independent additional conditions, because the evaluations are projectively distinct. We combine the retained cost $wT_0 - d_0 - q_0T_0$, the reader cost above and the source saving $d - T_0$. Canonical unmarked tensors take $3dT_0$ parameters and their fibre returns T_0; the common old span takes $d(r - d)$; and the four-system returns two parameters in the primitive case and none in the divided case. Substituting $d = 3T_0 - 1$ gives the codimensions

$$r + 9T_0 - 9 - d_0, \qquad r + (9 - q_0)T_0 - 7 - d_0,$$

respectively. Both exceed $r + 6$ when $d_0 \leq 2$; when $d_0 \geq 3$ we add the four extra conditions.

Step 5: the fixed reader and the principal slices. Apply a small isomorphism B_{o} of the source to T. In the incidence this replaces the effective reader by $R_vB_{\mathrm{o}}^{-1}$, and it preserves both types of retained row parameters. Where R_v has rank nine, this variation is onto all readers. On the rank-eight stratum it varies freely into the given eight-plane, and restricting readers to this plane loses at most r equations in the long counts. For the count with one relation we freeze the combined old span; the reader equations use at most its 900 directions and so lose at most 900 equations. Adding the codimension $r - 9$ of the rank-eight stratum, we obtain a bound of at least $r + 6$ in the combined

space in both cases, even if 72 frame parameters for the eight-plane are returned. In the original variables the reader, and hence C_0, is fixed.

The incidence counts involve finitely many tuple lengths and rank charts. The constraints and the changes of scalar system are constructible, since the polynomial degrees are bounded. We take the loci where the fibres have dimension at least 3ℓ, or at least T_0 in the canonical unmarked count, and then their closures; the bounds above apply to them without change.

Take four principal covectors on $A_4 = L \otimes \mathbb{R}^4$, represented locally by $(1, x_i, x_i^2, x_i^3)$ for four distinct real numbers x_i, and identify A with the first three coordinates. Any three of them restrict to independent covectors on A. Let B_i be their frozen two-planes and $B_{ij} = B_i \cap B_j$. Denote by Z_i the closed exceptional locus for the map on B_i. Project it to the map on B_{ij} and take the closure of the locus where the fibre has dimension at least four, in both orders; the union Z_{ij} has codimension at least $r + 5$, since the space of affine extensions has dimension five. Forgetting that line map and taking the closures where the fibres have dimension at least four for the three pairs of any triple, we obtain loci of codimension at least $r + 4$.

Small perturbations of the coefficient data and of an initial formal derivative D_0 close to I avoid all these loci over Σ, by stratified parametric transversality. The source sprays span the normalized row coordinates and the coordinates of the source isomorphism, and the sprays of the initial derivative span the coordinates of the frozen maps. The angular variable is the sphere variable itself, including the rank-eight stratum; after pullback, polynomial variations over P span their coefficients at each angular point. The dimension is at most $r + 2$, below each of the codimension bounds. The strict injection is retained along the straight path from I to D_0.

Step 6: the h-principle. In the open relation we require rank four, injectivity on $\ker C_0$, and avoidance of Z_i and Z_{ij} by the corresponding restrictions. In principal slice i the homogeneous pencil has no hypersurface of rank defects once its frozen data are fixed. Hence the affine pencil has a rank-defect set of complex codimension at least two: if its maximal minors had a common nonconstant factor, the highest homogeneous part of that factor would divide every maximal minor of top degree r, and these are precisely the minors of the homogeneous pencil. Failure of rank four itself has codimension two in the five free coordinates of the slice, since the other three active columns are independent. The varying exclusions Z_j, Z_{jk} have dimension at most three in that slice, by the fibre bounds for pairs and triples: for Z_j the map on B_{ij} is fixed, and for Z_{jk} with i distinct from j, k the covector i is nonzero on B_{jk}, so the bound for triples applies. The real trace of each excluded set therefore has codimension at least two. The complement of all excluded sets in the slice is connected and has full convex hull: a generic two-segment path joins any two points, and generic opposite displacements exhibit every point as a midpoint. Thus each principal slice is ample.

In Hopf flow coordinates, the four chosen covectors together with transverse coordinate covectors form integrable principal frames; one can use the functions $\lambda_i(s) \cdot \theta$ and the transverse coordinates, where θ denotes the flow coordinates and s the transverse variables. The transverse slices do not change the active test. We apply the open

principal-frame h-principle [MAdP21, Theorem 3.11] to a formal first jet with active part D_0. It gives a homotopy within the relation from this formal jet to a holonomic one. Appending it to the straight path from I and using stationary collars at the endpoints proves the theorem. The positive margin follows by compactness. □

8.4. **Removing the semidefinite target weight.** For a basis $v_1, \dots, v_5$ of V define the alternating Gram tensor

$$\Omega(T)_{ij}(e,f) = \langle T_i e, T_j f\rangle - \langle T_j e, T_i f\rangle, \qquad i<j, \qquad T_i = T(\,\cdot\,, v_i).$$

Let $S = \mathrm{Id} - \pi_X$. Each simple row is an old source functional times a label functional, so $\Omega(ST) = 0$. We may prepare the same normalized T so that the vertical derivative of $T' \mapsto \Omega(ST')$ is onto everywhere on Σ, as follows.

For $F = ST$ consider the map on label-indexed columns $Z_i \in \operatorname{im} S$ given by

$$Z \longmapsto \Big(e \mapsto \langle F_i e, Z_j\rangle - \langle F_j e, Z_i\rangle\Big)_{i<j}.$$

Its dual obstruction is an old-source array of skew labels. If the span of its coefficients in the ten-dimensional space of skew labels has dimension k, the projective parameter count is $k(r+10-k)-1$. The evaluation curve of $\mathcal{P}$ spans the label space, since $\dim \mathcal{P} = 5$; let u be the generic rank of the contraction of this span with a point of the curve.

Lemma 8.5. *We have* $1 \le u \le 4$ *and* $k \le 3u-2$.

Proof. The bounds on u follow from $k \ge 1$, the spanning property and skew symmetry. Choose a basis of five generic curve values, at each of which the contraction has rank u, and choose k independent edge coordinates of the complete graph on these five vertices that restrict to a basis of the dual of the skew span. At each vertex at most u selected edges can be incident, so $2k \le 5u$. This suffices for $u = 3, 4$.

For $u = 1$, the contractions of any two matrices A, B of the span along the curve are proportional, $Ax(t) = f(t)Bx(t)$. Pairing with $x(s)$ and using skew symmetry gives

$$\Big(f(t) - f(s)\Big)\, x(s)^T B x(t) = 0.$$

For $B \neq 0$ the last factor is a nonzero polynomial, by the spanning property, so f is constant and $k = 1$.

For $u = 2$ and $k = 5$, the selected edges would form a five-cycle. Every unselected edge coordinate lies in the spans of the two sets of selected coordinates incident to its endpoints, and these spans intersect in zero; hence all unselected entries vanish. At a curve point with all five basis coordinates nonzero, the contraction on the five-cycle has rank four, as one sees by rescaling its matrix to the incidence matrix of the odd cycle. This contradicts $u = 2$, so $k \le 4$. □

At most $70u$ simple nodes lose rank, because a nonzero u-minor of the contraction has degree at most $70u$. Even after the deletion of the at most nine residual rows and of seventy further chosen simple rows, the equations for each stress impose at least $u(w-79) - 70u^2$ conditions. After the projective parameters of the stress are returned, the codimension exceeds $r+2$, by Lemma 8.5 and $w = 3r-2$. The possible sets of

deleted rows are finite. The same perturbations of the simple injectors as in Theorem 8.4 therefore give surjectivity on the link, and this can be imposed together with the other incidence preparations.

For the derivative of Ω, prescribe a desired alternating tensor in the pairs (e, f) of old arguments. At fixed f we apply the linear right inverse of the surjective skew map just obtained, with right-hand side one half of the desired tensor at (e, f). Taking the skew difference in e, f yields the desired derivative of the curvature. All variations take values in $\operatorname{im} S$, so they lift to variations of T. This proves vertical surjectivity.

Now set $S_\eta = (S + \eta \operatorname{Id})/(1 + \eta)$. The finite-dimensional implicit function theorem, applied on the smooth complement of the vertical kernel given by the right inverse, yields $T_\eta \to T$ with $\Omega(S_\eta T_\eta) = 0$. The complement is obtained from the adjoint composed with the inverse Gram operator of the surjective derivative, so the construction is smooth over the compact link, and one positive η works at all points. Strict injectivity along the compact marked path persists for T_η and is preserved by the invertible row scaling S_η. The positive homotopy from Id to S_η retains the augmented comparison. Thus $G = S_\eta T_\eta$ is an exactly flat tensor of the original width with the same path of active derivatives. Its reader $J_z = G(\,\cdot\,, z)$ is a monomorphism, since the absence of projective rank-one kernel tensors is an open condition, by compactness.

There is also a radial version of the correction. Pull back the split rows and all angular data across a product collar. First keep the weight equal to zero on X while cutting the arbitrary part on X back to its original moment value. On a later subband increase the weight on X from zero to one, with only the pure moment remaining. The weighted curvature stays zero, and the simple rows still give a surjective derivative throughout. The implicit solution on this enlarged compact collar can fix the final pure moment germ, where the weight is already the identity. This matches the flat angular block to pure moment data on the exterior side of the collar, and the construction controls the marked active path and the curvature throughout the collar.

8.5. **Scaling and comparison with the bank.** Suppose $W = W_{\mathrm{old}} \oplus W_{\mathrm{a}}$ with $\operatorname{rank} W_{\mathrm{a}} = 143$, and that $J_z = G(\,\cdot\,, z)$ is a monomorphism. Let $R_{\mathrm{a}} : \mathcal{H} \to W_{\mathrm{a}}$ be a prescribed smooth map. We realize it as the full bank projection of a rotation of ΛJ_z. For sufficiently large Λ the operator

$$\Lambda^2 J_z^* J_z - R_{\mathrm{a}}^* R_{\mathrm{a}}$$

is positive. The rank gap of W_{old} over $\mathcal{H}$ exceeds $\dim \Sigma$, so there is a monomorphism $\mathcal{H} \to W_{\mathrm{old}}$, which can be normalized to have precisely this Gram operator. Adjoining R_{a} we obtain a map into W whose Gram operator is that of ΛJ_z. The two normalized inclusions into W are homotopic through isometric inclusions, by relative rank avoidance on $\Sigma \times [0, 1]$, and the orthogonal transport of the previous section implements the required placement by a slow isotopy. These inequalities still hold if the source is enlarged by at most $\lceil 1000\sqrt{n-4} \rceil$ auxiliary generating directions.

Suppose that the four cofactor rows of the derivative of this fixed bank projection, computed for the reference map, coincide with C_0 up to an invertible change of the four rows, and that the old retained block is independent of the active variables. Then the

active connection correction at the distinguished label depends only on the prescribed bank part, and it is bounded independently of Λ. On $\ker C_0$ the main retained tensor along the marked path has a lower singular-value bound proportional to Λ. Hence, for one sufficiently large Λ, the augmented comparison at z is injective at all marked trials and all independent samples of the active variables of the reference map. By compactness, a sufficiently small ball around z retains this comparison.

To place the values of the active map in this ball, we precompose on the four original spheres with their normalized k-th power maps $P_k(w) = \big(w_j^k/(\sum_l |w_l|^{2k})^{1/2}\big)_j$, and scale the active map minus z by $1/k$:

$$a_{\mathrm{a}}^{(k)}(y,w) = z + \tfrac{1}{k}\big(a_{\mathrm{a}}(y, P_k(w)) - z\big).$$

The Hopf derivatives gain a factor k, since P_k intertwines the Hopf flow with its k-fold speed-up, and this cancels the scaling of the values. We precompose the formal derivative path with the same power maps. Its initial derivative is still I, and its endpoint is the derivative of $a_{\mathrm{a}}^{(k)}$. These steps use a bank response of the reference map with the stated four rows; such a response is constructed in the next subsection.

8.6. **The 157-dimensional active block.** On each of the four input spheres there is a planar map whose Hopf derivative vanishes nowhere. Indeed, the partial immersion relation for the Hopf field is an open first-order relation. On a principal slice that changes the Hopf derivative, its excluded set is one point in a plane. Thus the slice is ample, and the open ample theorem [EM02, Theorem 18.4.1] applies to the formal derivative given by any nonzero constant vector. The same argument works for any complex sphere size at least two.

The speed can be normalized intrinsically. If the initial speed is $v > 0$, let L_{orb} be its mean over the orbit and put $q = v/L_{\mathrm{orb}} - 1$, which has zero mean on each orbit. For the Hopf action ρ_t the function

$$b(z) = \int_0^1 t\, q(\rho_t z)\, \mathrm{d}t$$

satisfies $Yb = q$, by integration by parts. The variable-time map $z \mapsto \rho_{b(z)} z$ preserves each orbit, has derivative $1 + Yb = v/L_{\mathrm{orb}} > 0$ along it, and is homotopic to the identity through $\rho_{sb(z)} z$; it is therefore a diffeomorphism. We compose the initial map with its inverse and divide by L_{orb}; the resulting speed is one. The product b_0 of four such maps has partial Gram operator the identity on the four generators. We place it in $\mathbb{R}^{15}$, using eight coordinates and padding with zeros.

Choose 64 covectors λ_ν on $\mathbb{R}^{15}$ which distinguish every nonzero quadratic and every nonzero cubic form on every real four-plane, that is, no such form vanishes at all the restrictions $\lambda_\nu|_E$. For cubics the projective marker space of pairs $(E, [p])$ has dimension $4(15-4) + (20-1) = 63$, and for quadratics the corresponding dimension is $4(15-4) + (10-1) = 53$. For a fixed nonzero form each equation $p(\lambda_\nu) = 0$ costs at least one dimension. Consequently the projected exceptional sets in $(\mathbb{R}^{15})^{64}$ have dimension at most $63 + 64 \cdot 14 = 959 < 960$, and we choose a tuple in their complement. For a small

constant $\varepsilon_0 > 0$ put

$$A_0 = \Big(\mathrm{Id} - \varepsilon_0^2 \sum_\nu \lambda_\nu \lambda_\nu^*\Big)^{1/2}, \qquad \Phi(x) = \Big(A_0 x,\, \big(\varepsilon_0 \cos(\lambda_\nu \cdot x),\, \varepsilon_0 \sin(\lambda_\nu \cdot x)\big)_{\nu=1}^{64}\Big).$$

The positive square root is invertible and $(\mathrm{d}\Phi)^* \mathrm{d}\Phi = \mathrm{Id}$. Thus $B_{143} = \Phi \circ b_0$ has unit partial metric in $\mathbb{R}^{143}$.

The four first, ten second and twenty third active columns of B_{143} are linearly independent. In a linear relation among these columns, the A_0 block first kills the corresponding linear combination of the first, second and third columns of b_0. The tangential component of the circle at frequency λ_ν then gives the homogeneous cubic in $\lambda_\nu \cdot Y_i b_0$ with the coefficients of the third columns, and by the choice of frequencies all these coefficients vanish. The radial components then give the analogous quadratic, so the coefficients of the second columns vanish as well. Finally the four first columns of b_0 are independent. This argument uses only the injectivity of the first array of b_0, and so it applies to any formal first, second and third arrays with this property. The four first columns lie in four fixed disjoint planes, and so they are homotopic through independent columns to four constant unit vectors: lift each into an unused independent constant coordinate and then rotate the endpoint. At the same time we interpolate the formal value and the second and third arrays to constants. Through the chain rule this gives a homotopy of the injection of all 34 columns, and in particular of the first 14 columns, to a constant injection. Orthogonal transport gives a normal identification along this homotopy.

Put $B = (B_{143}, 0)$ in $\mathbb{R}^{157}$, with fourteen unused coordinates. Let N be the orthogonal complement of all first and second active columns of B, and define

$$T_3 = \pi_N D_{\mathrm{a}}^3 B, \qquad R_3 = T_3 (T_3^* T_3)^{-1}.$$

By the independence just proved, R_3 is a smooth right inverse of T_3^*, with zero components in the unused block. The rank-nine bank bundle $\mathcal{B} = L \otimes (\mathrm{SymEnd}(U) \oplus U^*)$ admits an isometric inclusion ν into the unused $\mathbb{R}^{14}$ over P, by rank avoidance. For $\xi = \ell \otimes (M, \rho) \in \mathcal{B}$ in a local unit ℓ of L, we prescribe a symmetric cubic form $\mathcal{K}(\xi)$ on $A_4 = A \oplus L$ by its values on the four types of triples:

$$\begin{aligned} \mathcal{K}(\xi)|_{A^3} &= 0, \\ \mathcal{K}(\xi)(a, b, \ell) &= \langle Mu, v\rangle, \quad a = \ell \otimes u, \quad b = \ell \otimes v, \\ \mathcal{K}(\xi)(a, \ell, \ell) &= \rho(u), \\ \mathcal{K}(\xi)(\ell, \ell, \ell) &= 0. \end{aligned}$$

These formulas use $A = L \otimes U$ and are compatible with a change of sign of the unit ℓ. For sufficiently small $\kappa > 0$ put

$$\eta_\kappa = -\kappa R_3 \mathcal{K} + \nu\Big(\mathrm{Id} - \kappa^2 (R_3 \mathcal{K})^* R_3 \mathcal{K}\Big)^{1/2}.$$

Its two terms lie in orthogonal factors of the target, so η_κ is an isometric bank in N, and $\eta_\kappa^* D_{\mathrm{a}}^3 B = -\kappa \mathcal{K}^*$. Differentiating the orthogonality of η_κ to the first and second columns along $a \in A$ gives

$$\langle D_a \eta_\kappa(\xi), D_i B\rangle = 0, \qquad \langle D_a \eta_\kappa(\xi), D_{\mathrm{a}}^2 B(b, c)\rangle = \kappa\, \mathcal{K}(\xi)(a, b, c).$$

Thus $D_a\eta_\kappa$ is orthogonal to the four first columns and to the six second columns $D^2_{\mathrm{a}}B(b,c)$, $b,c\in A$. On the remaining four directions of the span of the active full columns we use the coordinates given by the pairings with the columns in $L\otimes A$ and $\operatorname{Sym}^2 L$, after projecting off the ten selected columns; this is an isomorphism because all fourteen columns are independent. In these coordinates the derivative of the bank has precisely the four rows $\kappa(Mu,\rho(u))$. The other component, normal to all full columns, is the derivative of η_κ itself and is retained.

At $\kappa=0$ the bank is the fixed slow inclusion ν into the fourteen unused axes. The smooth path in κ is a path of isometric banks in N, so its orthogonal transport matches the final bank to this slow inclusion. Together with the normal identification of B_{143}, this identifies the 143-dimensional active normal bundle, including the specified bank, with a fixed block. We compose its slow bank input with $\kappa^{-1}R_v$, in the notation of the exceptional link. The cofactor response of the reference map is then C_0 in the coordinates above. This realizes the bank hypothesis of the preceding subsection in 157 active coordinates. The construction does not depend on a trivialization of the line L.

9. Flat matching on the doubled exterior

Let $m=n-4$ and let B be the doubled slow exterior; it is a closed manifold with $\dim B\le m+4$. Put

$$k=\lceil 1000\sqrt{m}\,\rceil,\qquad R=m+k,\qquad w=3m+1.$$

The generating bundle is $\mathcal{E}=E\oplus\underline{\mathbb{R}}^k$, of rank R, while the retained target W keeps rank w; note that $w=3R+1-3k$. On an exceptional collar the generating line $L\mathcal{D}$ embeds in $\underline{\mathbb{R}}^k$ over the three-dimensional core P by rank avoidance, since $k>4$. The angular data of Section 8 therefore extend to $\mathcal{E}$ on the collar, by zero on the orthogonal complement of $L\mathcal{D}$ in $\underline{\mathbb{R}}^k$. In this section we extend these data to an exactly flat tensor on all of B, and we arrange at the same time that its old planes satisfy the following goodness conditions, both for generic full planes and for all reduced planes.

Definition 9.1. Let $G\colon \mathcal{E}\otimes V\to W$ be a five-label tensor at a point of B, and let $S\subset\mathcal{E}_{\mathbb{C}}$ be an old plane. A normal two-plane $N\subset V^*_{\mathbb{C}}$ is *bad* for S if the restriction of $G_{\mathbb{C}}$ to $S\otimes\operatorname{ann}N$ is not injective; here $\operatorname{ann}N\subset V_{\mathbb{C}}$ is three-dimensional. A *star* is the three-dimensional family of all normal two-planes containing a fixed normal line, its *center*. We call S *good* if the ordinary skew map of G on S is onto and the complex variety of bad normal two-planes for S has dimension at most four and contains no star. *Reduced old planes* are the complex $(m-1)$-planes in $E_{\mathbb{C}}\oplus\ell_{\mathrm{aux}}$, where $\ell_{\mathrm{aux}}\subset\mathbb{C}^k$ is any auxiliary line.

9.1. **Source estimates for polynomial moments.** Throughout this subsection all dimensions and incidences are over $\mathbb{C}$. Let P_d be a basepoint-free five-dimensional linear subsystem of $H^0(\mathbb{P}^1,\mathcal{O}(d))$ with independent pair products, that is, the fifteen products of pairs of elements of a basis are linearly independent in $H^0(\mathbb{P}^1,\mathcal{O}(2d))$. We also use four-dimensional subsystems, after division by their common factor; the pair-product property persists for a subsystem and after the division. We write $q\le 70$ for the degree after any such division.

Consider an irreducible projective family of rank-three old tensors with labels in such a system, satisfying one of the following two conditions on its normal image. In five labels, the unique annihilator two-planes of its members form a subvariety of $\mathrm{Gr}_2(V_{\mathbb{C}}^*)$ of dimension at least five. In four labels, the annihilator lines of its members are dense in $\mathbb{P}^3$. In both cases we assume that the general member is *primitive* as a vector polynomial, meaning that its scalar coordinates have no common zero on $\mathbb{P}^1$. Let T be the rank of the family over the field of rational functions on $\mathbb{P}^1$. Choose T generic members $u_1, \dots, u_T$ that are linearly independent over this field, and write the three old columns of u_i in the i-th summand of

$$\mathcal{V} = (\mathbb{C}^3)^{\oplus T}.$$

Let h be the dimension of the combined old span of $u_1, \dots, u_T$. The column map from $\mathcal{V}$ onto this span has a kernel R_{rel}, the *relation plane*, of dimension

$$\delta = 3T - h.$$

All old coefficients of members of the family lie in the combined span: a member is a rational combination of $u_1, \dots, u_T$, and its projection to the constant quotient by the span vanishes.

Lemma 9.2. *For such a tuple $(u_1, \dots, u_T)$,*

$$\delta \geq 2, \qquad 211\,\delta \geq T^{2/3}.$$

More precisely, there are integers $d_0, e_0 \geq 0$ with $T^{2/3} \leq 211(d_0 + e_0) \leq 211\,\delta$ such that, at fixed scalar row spaces of the u_i, the source incidence in $\mathrm{Gr}_\delta(\mathcal{V})$ admits the following codimension savings:

(i) *if $d_0 \geq 100$, a saving of at least $100(h - T - 7070)$;*
(ii) *if $d_0 < 100$ and $e_0 \geq 100$, a saving of at least*

$$d_0(h - T - 69d_0 - 70) + 100(h - T - 70d_0 - 69).$$

These savings are alternatives to the saving obtained from the span of evaluated relations alone.

Proof. Step 1: a function-space inequality. Let X be an irreducible variety, let U_b be a b-dimensional space of rational functions on X, and let x_1, x_2, x_3 be algebraically independent rational functions on X. We claim that

$$\dim(U_b + x_1U_b + x_2U_b + x_3U_b) \geq b + b^{2/3}. \tag{9.1}$$

Choose a smooth point p_0 at which all elements of U_b and the x_i are regular after clearing a common denominator, and at which $z_i = x_i - x_i(p_0)$, $i = 1, 2, 3$, are part of a system of local parameters; this is possible since the x_i are algebraically independent. By elimination in power series, choose a basis of U_b whose leading exponents, for a fixed monomial order, are distinct; they form a set A with $|A| = b$. Let π_i be the projection dropping the i-th coordinate. On each line parallel to the i-th axis that meets A, multiply the basis element with the topmost exponent on that line by z_i; the

leading exponent of the product is the successor of that point, which lies in $(A+e_i)\setminus A$. Functions with distinct leading exponents are independent, so

$$\dim(U_b + x_i U_b) = \dim(U_b + z_i U_b) \geq b + |\pi_i A|.$$

The discrete projection inequality $|A|^2 \leq \prod_{i=1}^{3} |\pi_i A|$ gives $\max_i |\pi_i A| \geq b^{2/3}$, which proves (9.1). For $A \subset \mathbb{Z}^3$ the projection inequality follows by applying Cauchy–Schwarz to the sum over A of the product of the three indicator functions of the pairwise projections. In general apply it on each slice with fixed remaining coordinates and use $\sum_j \sqrt{a_j b_j c_j} \leq \sqrt{\sum_j a_j \sum_j b_j \sum_j c_j}$.

We apply (9.1) to a constant basis $w_1, \ldots, w_b$ of any family with the specified normal image, writing its rank-three members again in b separate three-blocks; the kernel of the resulting column map is the *constant-basis relation plane*. On an affine chart of the family write a member as $\sum_j c_j w_j$, with linearly independent coefficient functions $c_1, \ldots, c_b$. In five labels choose a chart of $\mathrm{Gr}_2(V_{\mathbb{C}}^*)$ in which the normal plane has basis rows $(1, 0, x_1, x_2, x_3)$ and $(0, 1, y_1, y_2, y_3)$. Since the normal image has dimension at least five, one of the two triples, say (x_1, x_2, x_3), is algebraically independent. The contraction of a member by either row vanishes, so the separate contractions $(c_j\, \alpha(w_j))_j$ by a row α lie in the constant-basis relation plane. Before the separate five-to-three coefficient projections are applied, the contractions by the first row have coordinates $(c_j, 0, x_1 c_j, x_2 c_j, x_3 c_j)$ and span at least $b + b^{2/3}$ dimensions by (9.1). The contractions by the second row contribute at least b further dimensions, seen in the second normal slot of each block. The separate projections lose at most $2b$ dimensions. In four labels the single varying normal is $(1, x_1, x_2, x_3)$, with x_1, x_2, x_3 algebraically independent because the normal image is dense; the same argument uses only this row and loses at most b under the four-to-three projections. Consequently, in either case,

$$\dim(\text{constant-basis relation plane}) \geq b^{2/3}. \tag{9.2}$$

Step 2: torsion on $\mathbb{P}^1$. Return to the rational basis. The primitive columns u_i give subbundles $\mathcal{O}(-q) \to \mathbb{C}^3 \otimes \mathcal{O}$ on $\mathbb{P}^1$. Let Q_{pol} be their combined quotient; the map given blockwise by wedge product with u_i embeds Q_{pol} as a subbundle of a sum of copies of $\mathcal{O}(q)$. The map

$$R_{\mathrm{rel}} \otimes \mathcal{O} \longrightarrow Q_{\mathrm{pol}}$$

is generically injective by rational independence, so its cokernel has a torsion part of finite length L_{pol}. At each point the torsion classes are represented by the principal parts of meromorphic R_{rel}-valued sections with regular image in Q_{pol}, modulo regular relation sections. Let $D_{\mathrm{pol}} \subset R_{\mathrm{rel}}$ be the constant span of all coefficients of these principal parts, and put $d_0 = \dim D_{\mathrm{pol}}$. By the principal-part description, the torsion length of the map from $D_{\mathrm{pol}} \otimes \mathcal{O}$ is the same number L_{pol}. A nonzero d_0-minor of this map, composed with the wedge embedding, is a nonzero section of $\mathcal{O}(qd_0)$; its zero divisor bounds the torsion, so

$$L_{\mathrm{pol}} \leq q d_0 \leq 70\, d_0.$$

Let K' be a constant subspace with $D_{\mathrm{pol}} \subset K' \subset R_{\mathrm{rel}}$. The projected basis columns form a subsheaf of $(\mathcal{V}/K') \otimes \mathcal{O}$ isomorphic to $\mathcal{O}(-q)^T$, and its saturation has quotient

of length L_{pol} over this subsheaf. To see this, lift a meromorphic combination of the columns that is regular modulo K'; its principal part is a meromorphic K'-section with regular image in Q_{pol}. Conversely, every such relation principal part cancels the poles of a meromorphic combination of the columns, and holomorphic coefficients give back the original subsheaf. The natural maps between the saturated bundles for different K' are therefore generic isomorphisms between bundles of the same rank T and degree $-qT + L_{\mathrm{pol}}$; their determinant has no zeros, so they are isomorphisms. After twisting by $\mathcal{O}(q)$ the space of sections has dimension $T + L_{\mathrm{pol}}$, by the torsion exact sequence and $H^1(\mathbb{P}^1, \mathcal{O}) = 0$. In particular, taking $K' = D_{\mathrm{pol}}$, every member of the rational family lifts to this space of sections.

In $R_{\mathrm{rel}}/D_{\mathrm{pol}}$ take the span of the contractions of all lifted tensors by the annihilators of their three-dimensional scalar row spaces, using the full degree-q scalar space. All these contractions project to zero in $\mathcal{V}/R_{\mathrm{rel}}$. Let e_0 be the dimension of their span; thus $d_0 + e_0 \le \delta$. After passing to the quotient by this span, the lifted rank-three family has exactly its original scalar row spaces, including membership in the designated label subsystem. It contains the projected basis tensors in its closure, so its old span has dimension $3T - d_0 - e_0$, and its constant dimension b satisfies $T \le b \le T + L_{\mathrm{pol}}$. The estimate (9.2) for this family, together with $b - T \le L_{\mathrm{pol}} \le 70 d_0$, gives

$$T^{2/3} \le b^{2/3} \le 3(b - T) + d_0 + e_0 \le 211(d_0 + e_0) \le 211\,\delta.$$

Suppose $\delta \le 1$. Project each initial three-block away from the projection of R_{rel} into it. The resulting projection is defined on the old span and retains at least two independent scalar entries of that basis member. For any member of the family, its rational coefficient in that block multiplies both entries into the label subsystem. Such a multiplier is constant: if a, b are independent and fa, fb belong to the subsystem, then $a(fb) = b(fa)$ is an identity in its symmetric square, so by the pair-product property $fa = ca$ and $fb = cb$ for a constant c. Hence the constant dimension of the original family is T, and (9.2) gives $\delta \ge T^{2/3}$. Since $T \ge 1$, this forces $\delta = T = 1$; but a single projective tensor has a single point as normal image, which is incompatible with either condition on the normal image. This proves $\delta \ge 2$.

Step 3: the source savings. We count the constrained relation planes at fixed basis scalar row spaces. The surjective combined coefficient matrices onto a fixed h-space map to $\mathrm{Gr}_\delta(\mathcal{V})$, by taking the kernel, with equidimensional fibres. Requiring a given j-plane to be contained in the relation plane costs jh conditions.

If $d_0 \ge 100$, list the principal-part coefficients obtained from local Smith generators chain by chain, from each highest pole order down to the lowest, and stop at the hundredth independent one. The segments already processed and their shorter shifts give independent torsion classes for the current coefficient span: a truncated segment is represented by multiplying its full Smith generator by a power of a local parameter, and the shorter shifts form the basis of its cyclic submodule. Their independence holds in the full torsion module, hence in the module for the current span. Thus, by the same wedge-minor bound, at most qj positions have been processed when the span has dimension $j = 100$. If $d_0 < 100$, the argument applies to the full list with $j = d_0$.

The root locations cost at most qj parameters. With the preceding negative coefficients fixed, the next one is a known translate of a vector in the T-dimensional span of the separate block values at the root. Indeed, a meromorphic principal part mapping to a regular section of the quotient can be written locally as a combination of the separate columns plus a regular section, and its new coefficients enter through these independent block values. A step that increases the span therefore costs at most T parameters, and a step that does not increase it costs at most j. The possible j-planes thus form a family of dimension at most $jT + (q-1)j^2 + qj$, and requiring their containment in the relation plane saves at least

$$j\Big(h - T - (q-1)j - q\Big) \geq j(h - T - 69j - 70).$$

For $j = 100$ this gives the first saving. In the second case, after D_{pol} has been chosen (with $j = d_0$), the projected contractions of any lifted section are parametrized projectively by at most $T + L_{\mathrm{pol}} - 1 + q \leq T + 70d_0 + 69$ parameters: a projective lifted section and a scalar functional. Mark 100 independent such lines in the quotient by D_{pol} and require their containment in the relation plane as well. With $q \leq 70$ this gives the second saving.

These counts use only finite data: degrees, root multiplicities, finite arrays of Laurent coefficients, and finite rank charts. For a fixed proposed span the space of possible lifts in degree q is defined by the identical vanishing of the $(T+1)$-minors formed with the basis columns, which is linear in the extra column. All dimension estimates therefore apply to constructible source incidences, before or after projectivizing the members of the tuple. □

9.2. **Transversality and global flat matching.**

Proposition 9.3. *Suppose $m \geq 10^{62}$. Given the angular preparation of Section 8 on finitely many disjoint collars, there is a smooth tensor $G\colon \mathcal{E} \otimes V \to W$ over B with $\Omega(G) = 0$ which retains the marked path tests on the prescribed angular source in an inner band of each collar, and such that at every slow point:*

(i) *a nonempty Zariski-open set of full complex m-planes in $\mathcal{E}_{\mathbb{C}}$ is good;*
(ii) *every reduced old plane is good.*

Proof. The outer polynomial system. We first work in normalized coordinates, allowing a semidefinite target weight as in Section 8.4. After the residual-row collar we have a pure twisted moment, described by a slow selfadjoint operator K of W, an injector on $\mathcal{E} \otimes L$, and the L-valued five polynomials. We extend these data a little further across a product collar. On the untwisted side we choose a plain five-polynomial system, for example the span of

$$1 + \varepsilon, \quad t + \varepsilon t, \quad t^2 + \varepsilon t^3, \quad t^3 + \varepsilon t^9, \quad t^4 + \varepsilon t^{27}$$

for a small generic $\varepsilon > 0$, homogenized to degree 27. It is basepoint-free, since the first polynomial vanishes only at $t = \infty$, where the last one does not. It spans five labels, and its pair products are independent: the highest power of ε in a suitable maximal minor comes from the pair products of the monomials $1, t, t^3, t^9, t^{27}$, which are pairwise

distinct because the fifteen sums of two exponents from $\{0, 1, 3, 9, 27\}$ are distinct. The functional calculus of K realizes either moment by evaluating its polynomial row on the corresponding eigenspace.

The spectral bridge. Extend K to a smooth selfadjoint endomorphism of W over B. Perturb it away from the prescribed simple-moment collar so that all eigenvalue multiplicities are at most $M = \lfloor 2\sqrt{R} \rfloor$. By the dimension count of orthogonal orbits, a repeated eigenvalue of prescribed multiplicity j, with the eigenvalue free, has real codimension $j(j+1)/2 - 1$. For $j > M$ this exceeds $2R - 1$, hence $\dim B$. The residual eigenblocks on the prepared inner band have rank at most three, so this is a relative transversality condition. By compactness there is a window size such that every window of four times that size contains at most M eigenvalues, counted with multiplicity; otherwise a limit of windows of vanishing size would give an eigenvalue of multiplicity larger than M.

On the transition collar sweep a threshold from below the spectrum to above it, with room at both ends. Let u be the signed spectral displacement from the threshold divided by the window size. We use a weight $g(u)$ which vanishes for $|u| \le 2$, is positive for $|u| > 2$, and equals one for $|u| \ge 3$. The twisted moment is used on the positive side, multiplied by a cutoff supported in $u > 1/2$ and equal to one for $u \ge 3/2$; the plain moment is used with the reflected cutoff. In $|u| < 1.65$ we allow general normalized rows, by adding an arbitrary tensor multiplied by a scalar function which is positive there and vanishes for $|u| \ge 1.75$. Every normalized row is then either freely variable or an independent injector row evaluated against exactly one curve. Where the weight is nonzero the weighted tensor is a pure moment, so its alternating curvature vanishes: $\Omega(gT) = 0$.

The spectral formulas are smooth through multiple eigenvalues. For instance, write the cutoffs by Fourier inversion, after a smooth compactly supported modification outside a bounded interval containing the spectra. The exponential $e^{\mathrm{i}\xi K}$ and its derivatives of fixed order have polynomial bounds in the Fourier parameter ξ, which permits differentiation under the rapidly decaying integral. At the two ends of the collar the construction agrees with the respective pure moments.

Ordinaryness of the live rows. We first check that the rows with nonzero weight are ordinary, that is, that their skew map is onto. For a nonzero skew stress let u_0 be the smaller of its generic contraction ranks on the two evaluation curves that occur. The spanning-curve argument of Section 8.4 bounds its coefficient rank by $3u_0 - 2$, so the projective stress parameters number at most $(3u_0 - 2)(R + 10) - 1$. A nonzero u_0-minor on each curve has degree at most $70u_0$, so rank is lost at no more than that many nodes per curve. The zero-weight window costs at most M further rows. Even if 140 additional whole eigenblocks are deleted, the equations cost at least

$$u_0(w - 141M) - 140u_0^2 M.$$

Since $u_0 \le 4$, $k \le 1001\sqrt{R}$ and $M \le 2\sqrt{R}$, the net codimension exceeds $2R - 20000\sqrt{R} - 100 > m + 4$. Thus ordinaryness holds after the small simultaneous preparations described below. Near the angular band the same construction includes the

residual-row bridge of Section 8.4. At most nine residual rows are lost there. Their normalized variables are free while their weight vanishes, and their coefficient cutoff can be taken positive until the single-curve term has become nonzero. All counts in that band use one curve.

Low-rank exclusions and a kernel cap. The following dimension counts give the bound on the bad family for generic full old planes. We exclude rank-one kernel tensors and pairs of distinct projective rank-two kernel tensors in the ambient R-space. On either curve, rational proportionality of two rank-two tensors would force constant proportionality, by independence of pair products. The bounds on root losses, with two curves and multiplicities at most M, give the codimensions

$$w - 140M - (R + 3), \qquad 2w - 280M - 2(2R + 5),$$

where $R + 3$ and $2R + 5$ are the dimensions of the projective rank-one and rank-two tensors. Free rows impose the full one or two equations. Both codimensions exceed $m+4$. When these exclusions hold, a single rank-two kernel tensor with two-dimensional label image V_2 makes N bad only if $N \subset \operatorname{ann} V_2$; these N form a projective plane, which has dimension two and contains no star.

We also cap the kernel on each three-label plane: its dimension is at most $a - 1$, where

$$a = \lfloor 7000\sqrt{R}\,\rfloor + 1.$$

At each node where the three-plane does not vanish identically, an element with zero evaluation has coefficient rank at most two. Hence in an a-dimensional kernel of tensors the evaluation kernel has dimension at most one, and it is zero except at the roots of the single possible rank-two ray. On each curve there are at most 70 such nodes and at most 70 basepoint nodes of the three-plane. The corresponding loss over both groups is at most $140M(a+1)$. Marking the kernel plane and its three labels costs $a(3R-a)+6$. Consequently the forbidden incidence has codimension at least

$$a(a + 1 - 3k) - 140M(a + 1) - 6 > m + 4.$$

A freely variable normalized block incurs no loss in this count or in the other kernel-tuple counts.

General full old planes. Suppose that a generic m-plane in the old source has a bad family of dimension at least five. Then there is an irreducible projective rank-three kernel family Γ in the ambient space, of dimension at least $3k + 5$, whose normal image has dimension at least five. Indeed, the normal two-plane of a rank-three tensor is its unique annihilator. Marking an m-plane containing its three-dimensional old span takes $k(m - 3)$ parameters, while the m-planes vary in a family of dimension km. The single rank-two ray and the excluded rank-one tensors cannot account for a five-dimensional normal image. Thus a component that gives at least five dimensions in the generic marked m-plane has the asserted dimension before marking. For each label curve, the vector polynomials in Γ are generically primitive: a common root would put their normal line through a point of that evaluation curve, and the lines in $\mathbb{P}^4$ meeting a curve form a family of dimension at most four.

Choose

$$\ell = \min\Big(T_1, T_2, \lfloor R - 300\sqrt{R}\,\rfloor\Big)$$

generic members, independent over the rational function field on both curves, where T_1, T_2 are the two rational span dimensions; each condition is open on the same irreducible family of tuples. With one curve, use just its dimension. Let d be the combined old dimension, put $\delta = 3\ell - d$, and let ν be the larger of the dimensions of the spans of evaluated relations on the two curves, as in the proof of Theorem 8.4. The relation plane itself is the same for both curves, since both scalar systems are identifications of the same label space. The wedge-minor argument bounds the total spectral loss by $140M\nu$. The free rows give every equation for each constant-independent tuple. We use the larger relation-span saving $\nu(d - \ell)$, once. Each rank-three flag in the common old d-space takes $3d + 5$ parameters, including its projective three-label space, and the common old span takes $d(R - d)$. Returning the $(3k + 5)\ell$ parameters of the fibre of tuples in Γ and using $w = 3R + 1 - 3k$, we obtain the net codimension

$$w\ell - 140M\nu - \Big(d(R-d) + (3d+5)\ell - \nu(d-\ell) - (3k+5)\ell\Big) = \delta(R-d) + \ell + \nu(d-\ell-140M).$$

If ℓ is the truncation $\lfloor R - 300\sqrt{R}\rfloor$, this is at least $\ell > m + 4$: when the bracket is negative, use $\nu \le \delta$ and $R - \ell - 140M > 0$.

If instead ℓ is a shorter whole basis on one curve, then $\nu, \delta \ge 1$. Otherwise all evaluated vectors would be independent on that curve, and regularity on $\mathbb{P}^1$ would force all rational coefficients of any member of the family to be constant. The contraction-drop locus of this constant basis in the dual $\mathbb{P}^4$ would then be closed and disjoint from the evaluation curve, and so of dimension at most two by the projective dimension theorem. Every normal line of the family lies in that locus, and such a family of lines has dimension at most two, since two distinct marked points determine a line. This contradicts the dimension of the normal image. Using $\delta, \nu \ge 1$, a nonnegative bracket gives a bound of at least $R - 140M > m + 4$; for a negative bracket the same follows from $\nu \le \delta$. Thus the projected bad incidences have codimension greater than the slow dimension. At a tensor outside them there is, by the dimension of projective fibres, a Zariski-open set of full old planes with the claimed property. It contains real planes, because the real Grassmannian is Zariski-dense in the complex one.

All reduced old planes. Fix ℓ_{aux} and work in the old $(m + 1)$-space $E_{\mathbb{C}} \oplus \ell_{\mathrm{aux}}$. The restrictions of the injector and of the arbitrary rows to it are freely variable. Write $r_0 = m + 1$ and $X = \sqrt{R}$; then $w = 3r_0 - 2$. If some reduced plane has a bad family of dimension five, the low-rank exclusions leave an irreducible rank-three family with normal image of that dimension, generically primitive on both curves as before. Its members lie in the reduced plane, so all combined old spans counted below have $h \le r_0 - 2$, and no separate mark for the reduced plane is needed. For ℓ entries that are rationally independent on both curves we use the same relation-plane saving and spectral loss. The fibre of tuples has dimension at least 5ℓ. The resulting count in the r_0-space is at least

$$B_{\mathrm{red}} = \delta(r_0 - h) - 2\ell + \nu(h - \ell - 140M), \qquad \delta = 3\ell - h.$$

Let ℓ be the smaller rational rank or $\lfloor 4r_0/5 \rfloor$, whichever is less. In the truncated case with nonnegative last bracket, $r_0 - h \geq 2$ and $\delta \geq 3\ell - r_0 + 2$ give $B_{\mathrm{red}} \geq 4\ell - 2r_0 \geq 6r_0/5 - 4$. With negative bracket, $\nu \leq \delta$ and $\delta \geq 3\ell - r_0$ give $B_{\mathrm{red}} \geq \delta(r_0 - \ell - 140M) - 2\ell > 2r_0$.

For a shorter whole basis we apply Lemma 9.2 on the controlling curve. If $h \leq 0.9r_0$, then $B_{\mathrm{red}} \geq \delta r_0/10 - 2\ell$, using $\nu \leq \delta$ when the bracket is negative; this is greater than $3r_0$ for $\delta \geq 50$. For $\delta < 50$ the lemma gives $\ell < 2 \cdot 10^6$ and $h < 6 \cdot 10^6$, and $\delta \geq 2$ yields the sharper inequality

$$B_{\mathrm{red}} \geq 2(r_0 - 6 \cdot 10^6) - 4 \cdot 10^6 - 7000M > 1.1r_0.$$

If $h > 0.9r_0$, then $\ell > 0.3r_0$ and $h - \ell > 0.1r_0$. For $\nu \geq 50$ the bound exceeds $-1.6r_0 + 50(0.1r_0 - 140M) > 1.1r_0$. For $\nu < 50$, the lemma gives $d_0 + e_0 > 201$, so that one of its two source savings is greater than $9r_0$. Using that saving in place of $\nu(h - \ell)$ in B_{red} gives a bound greater than $1.1r_0$; this saving uses only the controlling curve. All these bounds exceed $m+4+k-1$. Letting the auxiliary line vary costs $k-1$ parameters, so the bound holds simultaneously for all reduced planes. The resulting condition is open, by properness of the projective flag space of auxiliary lines and reduced planes.

Regular tuples for stars. At the weighted tensor prepared so far, which we call the *seed*, consider a rank-three kernel family dominating the star through some center. Its old coefficients may lie anywhere in $\mathcal{E}_{\mathbb{C}}$, while all its labels lie in the four-dimensional annihilator of the center. On each curve divide these labels by their common factor; the removed nodes carry in total $q_\star \leq 140M$ rows. After the division a generic tensor of the family is primitive on both curves, since a common root would force its normal point in $\mathbb{P}^3$ to lie on the divided evaluation curve. Rank-one tensors are already absent, and a single rank-two tensor cannot dominate a star. Let b be the constant span dimension of the rank-three family, and put $c = w - 3R = 1 - 3k$.

Suppose $b \leq R/16$. Choose a whole rational basis of length $\ell \leq b$ on the curve of smaller rank, independent on both curves. Outside the common basepoint blocks the loss of evaluation equations is at most $140M\ell$, by the zero divisor of a nonzero ℓ-minor on each curve. A generic projective flag of a rank-three tensor in four labels takes $3h+2$ parameters in the combined h-space. Returning 3ℓ in the fibre of tuples and allowing four parameters for the shared center, we obtain a net count of at least

$$\delta(R - h) + (c + 1 - q_\star - 140M)\ell - 4.$$

By Lemma 9.2, $\delta \geq \max(2, \ell^{2/3}/211)$, and $h \leq 3\ell \leq 3R/16$. The negative linear term has magnitude at most $3700X\ell$, which is less than $0.01R\delta$. Hence this count exceeds $m + 4$.

Suppose $R/16 < b \leq 3R/4$. Choose a constant basis of b generic members and put $\delta = 3b - h$. At every moment node that is not a basepoint, the evaluated basis loses at most $a - 1$ in rank, by the cap on three-label kernels. With the center fixed, Step 1 of the proof of Lemma 9.2 in four labels gives $d_1 = \lceil b^{2/3} \rceil$ independent normal-contraction vectors in the relation plane. With the basis scalar spaces fixed, each such projective vector can be marked in a family of dimension at most $b+3$, since it is the contraction of a linear combination by one projective normal. Requiring the relation plane to contain

their span saves at least $d_1(h-b-3)$, and $\delta \geq d_1$. The retained equations cost at least $(w-q_\star)(b-a)$. Counting the combined source and the flags with the same center and fibre dimensions as before gives the lower bound

$$d_1(R-b-3)+(c+1-q_\star)b-wa-4,$$

which exceeds $0.03R^{5/3}-27000R^{3/2} > m+4$. Therefore, outside a set of codimension greater than $m+4$, every covering rank-three family at the prepared seed has $b > 3R/4$.

Let $l_* = \lfloor 13R/20 \rfloor + 1$. Since $b-a > l_*$, each covering family has l_* constant-independent members whose evaluations are independent at every moment node that is not a basepoint; the relevant nonvanishing conditions are finitely many open conditions on the irreducible power of the family. On the zero-weight blocks and on the basepoint blocks we first prescribe arbitrary outputs of the tuple by an unrestricted variation of the normalized tensor, and compensate its linearized curvature on the remaining live blocks by ordinaryness there. This uses the deletion of at most 140 whole basepoint blocks besides the zero-weight window, which is within the ordinaryness count above. On each remaining moment block we freely vary its injector to prescribe the remaining outputs; this variation preserves the weighted curvature to first order. Near the angular residual-row bridge the live residual blocks, if used, have pure moment form by construction, so the same argument applies. This proves that the l_* evaluations are jointly surjective on the tangent space of the normalized weighted-flat relation, with the weight held fixed. We call such a tuple *regular*.

Star avoidance. Star avoidance is a condition in the unrestricted weighted-flat relation. The moment estimates by themselves give only regular tuples, and a star may still be present, so we pass from the prepared seed as follows. Near the seed the derivative of the weighted curvature is onto. Regularity is the open full-rank condition on the stacked vertical evaluation and curvature matrices; it is an algebraic condition with the weight matrix among the parameters. Every nearby star failure admits at least $3l_*$ dimensions of regular tuples with its fixed center and source. Otherwise the constructible locus of failures with smaller fibres would specialize to the seed along an algebraic trait, the source and center marks being compact and projective. Take the projective closure over the trait of a generic covering component of the family of kernel flags, after a finite extension if needed. This closure is flat over the trait, since its local rings have no torsion over the discrete valuation ring. By properness, a component of the special fibre still covers the whole specialized star, because the generic star is covered and its closure is the star family. At the seed such a component has general rank three, by the low-rank exclusions, and it contains a regular tuple, by the seed calculation. The l_*-fold fibre product of the flat closure is flat, so its generic fibre is dense, and the open regular minor lifts to the power of the same generic component. An open subset there has dimension at least $3l_*$, a contradiction.

For a marked old s-plane and a fixed center there are $3s+2$ parameters per projective flag, and each entry of a regular tuple imposes w independent equations in the weighted-flat relation. Returning $3l_*$ fibre dimensions gives a codimension of at least $l_*(w-3s+1)$ before the marks are forgotten. For $s = m-1$ we have $w-3s+1 = 5$, and allowing all

reduced planes and the center costs $(k-1)+2(m-1)+4$; the resulting codimension is at least

$$5l_* - (k-1) - 2(m-1) - 4 > m + 4.$$

For $s = m$, use first the dimension bound in coefficient–plane pairs, allowing four parameters for the center. If all full planes contain a star, the projection to the coefficients returns the whole Grassmann fibre, of dimension km, and the resulting codimension is at least $2l_* - 4 > m + 4$. For a fixed tensor the full planes containing a star form a closed set, by properness of the incidence of normal flags. Thus, after simultaneous avoidance, there are nonempty open sets of full planes without a star and with the dimension bound, and their intersection is nonempty.

Stratification and transversality. The counts used in preparing the weighted tensor are finite constructible incidences on its spectral strata. On a stratum of fixed eigenvalue multiplicities we use separate local orthonormal frames of the eigenspaces and the eigenvalues as coordinates; the row-type strata also record the inequalities between these eigenvalues and the threshold. On a row with nonzero moment coefficient we absorb that scalar into the free injector, and on a freely variable row we absorb its moment part into the independent map of that row. These are invertible changes of parameters on the strata, and the counts above are the algebraic counts in the resulting injector and coefficient variables. We first choose the selfadjoint section transverse to the spectral and row-type strata, which is possible by local sprays of symmetric matrices relative to the pure collars. The same local sprays of the injector and free-row variables, applied to a stratification of the constructible incidences and their closures, then give the stated avoidances. Here we use that the real trace of a complex constructible set of complex codimension c lies in a finite union of submanifolds of real codimension at least c: near a smooth point, the real and imaginary parts of c local defining functions with independent differentials contain c functions with independent real differentials. On the pure inner band the counts use only the single twisted curve and the freely variable injector and residual rows. All these perturbations can be taken small there, so the compact marked path tests are preserved. In the two-curve region the twisted data are used only inside the product collar; the line L is not extended over B.

For the subsequent unrestricted perturbation we use local implicit charts of the normalized weighted-flat relation to spray its whole vertical tangent space. On each spectral and row-type stratum the weight has constant rank and inertia, and smooth local congruences identify it with a fixed weight of that rank for the dimension counts, without changing the kernel tests. The smooth spectral section may first be chosen transverse to these strata, as above. The real traces of the closures of the complex constructible failure sets then have codimension greater than the dimension of each slow stratum, and relative stratified transversality gives one arbitrarily small perturbation avoiding all of them. The implicit charts are smooth also where the rank of the weight changes, since the vertical curvature derivative is onto there. Thus the perturbations can be made on the whole product collar and the double, preserving the dimension bounds and the strict angular comparisons. On the inner band one may first prepare the single-curve

data on the link and pull them back radially; openness provides a buffer for all the relative perturbations.

Removal of the weight. On the resulting normalized solution T_* of the weighted equation $\Omega(gT_*) = 0$, ordinaryness of the live rows gives a surjective vertical curvature derivative, by the same half-pairing formula as in Section 8.4. Replace the weight by $g_\eta = (g + \eta)/(1 + \eta)$ with a single small constant $\eta > 0$ on the compact doubled exterior, and solve the curvature equation $\Omega(g_\eta T) = 0$ on the smooth normal complement given by the right inverse, obtaining $T_\eta \to T_*$; put $G = g_\eta T_\eta$. By openness this retains the simultaneous tests on full and reduced planes. Since g_η is positive, row scaling by it preserves each kernel test, and ordinaryness transfers by the transpose on variation columns. The marked angular path persists in its inner band throughout the small normalized correction and the invertible scaling, as in Section 8.4. This proves the proposition. □

10. Boundary data when the rank-five normal splits

In this section we treat the components of the three-dimensional core P on which the rank-five normal bundle N_5 of the formal disk columns splits as

$$N_5 \simeq N_2 \oplus \underline{\mathbb{R}}^3, \qquad \operatorname{rank} N_2 = 2.$$

By Proposition 6.2 this holds on every component of P when n is odd or $n \equiv 0 \pmod 4$, and on the components where $w_1(\det \operatorname{Hom}(K, Q))^3 = 0$ when $n \equiv 2 \pmod 4$. We use the core P, the kernel line K, the cokernel Q of rank $m + 1$, the line $L = JK$ and its orthogonal complement $A = L^\perp \subset \mathbb{R}^4$ from the four-section construction of Section 6.1. Choose an isometric monomorphism $v_0 \colon K \to Q$ over P, that is, a section of the unit link bundle, and identify its image with L. Put $D = Q \ominus L$, a bundle of rank $m = n - 4$. The old and active reference blocks of Section 7 give a slow bundle W over P of rank $w = 3m + 1$. Along the marked section of the link the complete columns decompose orthogonally as

$$J_2(D) \oplus (A_4 \oplus \operatorname{Sym}^2 A) \oplus (D \otimes A).$$

The four active columns omitted from the first two summands span $L \otimes A \oplus \operatorname{Sym}^2 L \simeq \mathbb{R}^4$, so the orthogonal complement of the first two summands is $W \oplus \underline{\mathbb{R}}^4$, while the complement of all columns is $N_5 \simeq N_2 \oplus \underline{\mathbb{R}}^3$. Hence $(D \otimes A) \oplus N_2 \oplus \underline{\mathbb{R}}^3 \simeq W \oplus \underline{\mathbb{R}}^4$, and cancelling three trivial lines over P gives

$$W \oplus \underline{\mathbb{R}} \simeq (D \otimes A) \oplus N_2.$$

Here and below, stable cancellation over P of bundles of rank greater than three is implemented by a homotopy of a common unit-line inclusion followed by orthogonal transport of its complement.

10.1. **A moment with one scalar row.** Let $p_b(t)$, $b \in V = \mathbb{R}^5$, be the fixed untwisted degree-27 label system introduced in the proof of Proposition 9.3 (Section 9). It is basepoint-free and has linearly independent pair products. Its evaluation curve $\mathbb{P}^1 \to \mathbb{P}^4$ is injective: on the affine part the ratio of the first two coordinates recovers t, and the point at infinity has a distinct image.

Proposition 10.1. *Over P there are a unit section e_* of D, a covector field ℓ_* on A, and a pure moment $G_D(h \otimes b) = p_{gb}(B_*)J_D(h)$ on D, with B_* a smooth selfadjoint endomorphism of W, g a smooth automorphism of V arbitrarily close to the identity and $J_D \colon D \to W$ an injector, such that, with*

$$\sigma(h, a) = \langle h, e_* \rangle \, \ell_*(a), \qquad h \in D, \quad a \in A,$$

the map $(G_D|_{D\otimes A}, \sigma) \colon D \otimes A \to W \oplus \underline{\mathbb{R}}$ is injective with normal bundle isomorphic to N_2. Moreover, on $Q = D \oplus L$ adjoin the three rows $h \otimes a \mapsto (\pi_L h) \otimes a$ and use G_D on $\pi_D h$. At the included derivative $I \colon A_4 \hookrightarrow V$ the resulting map $Q \otimes A \to W \oplus (L \otimes A) \oplus \underline{\mathbb{R}}$ is injective with normal N_2, and its restriction to $E_v \otimes A$ over the link is injective with normal N_5. Both statements persist under sufficiently small variations of G_D with σ fixed.

Proof. The spectrum of B_*. Split off a unit section e_*, so that $D = D' \oplus \mathbb{R}e_*$. Take a generic smooth real selfadjoint endomorphism B_* of W and a smooth label automorphism g close to the identity. A triple eigenvalue has codimension five and two distinct double eigenvalues have codimension four; since $\dim P = 3$, the eigenvalues of B_* have multiplicity at most two, with at most one double eigenvalue in each fiber. Write $x(t) \in \mathbb{R}^3$ for the three evaluations of $p_{gb}(t)$ on a local basis of A. At a simple eigenvalue the condition $x(t) = 0$ imposes three conditions on the label three-plane $g(A)$. The evaluations at two distinct eigenvalues are independent, so simultaneous vanishing costs six conditions, and vanishing at a double eigenvalue costs at least five, including the codimension of the multiplicity. Hence only simple eigenlines can be *inactive* on A, that is, have $x(t) = 0$; there is at most one in each fiber, and the points of P carrying one form a finite transverse set P_{inact}.

Injectivity on $D' \otimes A$. For a generic injector $J' \colon D' \to W$ the moment $G'(h' \otimes a) = p_{ga}(B_*)J'(h')$ is injective on $D' \otimes A$. Indeed, a projective rank-three kernel witness has $3(m-1) - 1$ parameters and imposes one equation on each spectral row except the inactive one, that is, at least $3m$ equations; this leaves codimension at least four. A witness of rank one or two is a nonzero vector polynomial of degree at most 27, which vanishes on at most 54 spectral rows; its projective parameter count is at most $2m - 1$, so the codimension is at least $m - 52 > 3$. Let C_4 be the rank-four orthogonal complement of $G'(D' \otimes A)$ in W. At a point of P_{inact} the inactive eigenline U_0 is orthogonal to $G'(D' \otimes A)$, hence $U_0 \subset C_4$.

The projected evaluation. We claim that the map

$$\Psi \colon W \to \mathrm{Hom}(A, C_4), \qquad \Psi(u)(a) = \pi_{C_4} \, p_{ga}(B_*)u,$$

is onto off P_{inact}, and onto $\mathrm{Hom}(A, C_4 \cap U_0^{\perp})$ at each point of P_{inact}. A cokernel witness is a triple of vectors in C_4; in an eigenbasis of B_* it has entries $y_i \in \mathbb{R}^3$ at the spectral rows t_i, and it satisfies $x(t_i) \cdot y_i = 0$. Thus the matrices $x(t_i)y_i^{\mathsf{t}}$ are trace-free. Let S be their span, of dimension $b \le 8$. Orthogonality of the witness to $G'(D' \otimes A)$ imposes $(m-1)b$ linear conditions on J', and marking S costs at most 16 parameters. With S fixed, let $d \le 2$ be the generic dimension of the space of y with $x(t) \cdot y = 0$ and $x(t)y^{\mathsf{t}} \in S$. A nonzero maximal minor of these constraints has degree at most 81, so the witness dimension is at most $wd + 486$. We bound b in terms of d.

- If $d = 2$, every trace-free matrix orthogonal to S preserves every evaluation line $\mathbb{R}x(t)$. Such a matrix is scalar: a finite union of its eigenspaces cannot contain the spanning polynomial curve unless one eigenspace contains the curve. Hence it vanishes and $b = 8$.
- If $d = 1$, a rational generator $y(t)$ cannot have constant projective direction, since $x(t)$ spans $\mathbb{R}^3$ and $x(t) \cdot y(t) = 0$. The entries of $y(t)$ therefore span at least two rational functions, and those of $x(t)$ span three. Their pair products span at least four: choose bases with distinct vanishing orders at a regular point and use that the sumset of two increasing integer lists of lengths a_1, a_2 has at least $a_1 + a_2 - 1$ elements. Hence $b \geq 4$.
- If $d = 0$ and $b > 0$, we use $b \geq 1$.

In each case $(m-1)b - wd - 502 > 3$. These fiberwise incidence counts give simultaneous real transversality over P for a generic J'. The remaining witnesses, with $b = 0$, have $y_i = 0$ whenever $x(t_i) \neq 0$; they are exactly the triples supported on an inactive row. This proves the claim; no global eigenbasis is needed.

The last column. Since $W = (D' \otimes A) \oplus C_4$, the identity $W \oplus \underline{\mathbb{R}} \simeq (D \otimes A) \oplus N_2$ shows that $C_4 \oplus \underline{\mathbb{R}}$ and $A \oplus N_2$ are stably isomorphic; as these bundles have rank five over the three-dimensional base P, they are isometric. Choose an isometric inclusion $i = (i_{C_4}, i_{\underline{\mathbb{R}}}) \colon A \to C_4 \oplus \underline{\mathbb{R}}$ with normal N_2, and put $\ell_* = i_{\underline{\mathbb{R}}}$. At each point of P_{inact} a local orthogonal isotopy moves i so that $i_{C_4}(A) \subset C_4 \cap U_0^{\perp}$ at that point; the isotopy can be supported in a small ball, since the special orthogonal group of a five-dimensional space acts transitively on its orthonormal three-frames. In that ball choose a local section u of W with $\Psi(u) = i_{C_4}$ at the center, and blend i_{C_4} to $\Psi(u)$ in a smaller ball while keeping $i_{\underline{\mathbb{R}}} = \ell_*$. The modified map remains injective by closeness to its value at the center. Away from the smaller balls solve $\Psi(u) = i_{C_4}$ using the surjectivity of Ψ, and patch the local solutions by a partition of unity; since the equation is linear in u, the patched section still satisfies it. Put $J_D|_{D'} = J'$ and $J_D(e_*) = u$. After eliminating the component of the last three columns lying in $G'(D' \otimes A)$, the last three columns together with the scalar row give the inclusion (i_{C_4}, ℓ_*). Hence $(G_D|_{D \otimes A}, \sigma)$ is injective with normal N_2.

Extension to Q and to the link. Adjoining the rows $(\pi_L h) \otimes a$ adds the trivial three-plane $L \otimes A$ to both source and target, so the map on $Q \otimes A$ is injective with normal N_2. Over the link, $Q = E_v \oplus vK$, and restriction to $E_v \otimes A$ removes the three dimensions $vK \otimes A \simeq L \otimes A$, a trivial three-plane. The normal of the restricted inclusion is therefore $N_2 \oplus \underline{\mathbb{R}}^3 \simeq N_5$. Both statements are open conditions over the compact base and persist under small variations of G_D with σ fixed. □

10.2. **Simultaneous rank and incidence preparation.** Embed K in the trivial bundle $\underline{\mathbb{R}}^k$ over P by rank avoidance, and let $K^{\perp}$ be its orthogonal complement. Over the link, the restriction of $\mathcal{E} = E \oplus \underline{\mathbb{R}}^k$ is identified isometrically with the pull-back of the slow bundle $Q \oplus K^{\perp}$ of rank R, by sending the auxiliary copy of K to vK, that is, $(e_v, k_0, k_{\perp}) \mapsto (e_v + v(k_0), k_{\perp})$. Extend the moment G_D of Proposition 10.1 by zero on $L \oplus K^{\perp}$, and vary its injector and spectrum on this rank-R bundle.

Lemma 10.2. *Arbitrarily close to the moment of Proposition 10.1 there is a prepared pure moment, together with an arbitrarily short path in the ordinary flat locus over P to a flat tensor G on $(Q \oplus K^{\perp}) \otimes V$, such that:*

(i) *the full and reduced goodness of Proposition 9.3 hold at every point of the link, with the same tensors over P;*
(ii) *$G(\cdot, z)$ is injective on $Q \oplus K^{\perp}$;*
(iii) *on D, the bad normal family $\mathcal{B}$ has dimension at most four, contains no full star, and contains no $\mathrm{Gr}_2(\operatorname{ann} b)$ for a projective label b.*

These properties persist under subsequent small variations in the ordinary flat locus, and the preparation retains the augmented injection of Proposition 10.1 and its normal class.

Proof. We carry out the pure-moment preparation of Proposition 9.3 over P. For reduced sources we mark both the angular point v and the auxiliary line; the complex angular quadric adds m parameters and the auxiliary line $k-1$. The seed bound $1.1(m+1)$ for reduced incidences exceeds $m+k-1+3$, and the reduced star bound becomes

$$5l_* - (k-1) - 2(m-1) - 4 - m > 3.$$

The generic full-plane condition is invariant under the orthogonal change of source at v. Thus the prepared seed, and its subsequent unrestricted flat perturbation, satisfy the exterior tests at every point of the link. This gives (i). Part (ii) follows from the rank-one exclusion, and is an open condition.

For (iii), avoidance of full stars on the fixed bundle D has codimension $2l_* - 4 > 3$ in the same unrestricted flat perturbation as in Proposition 9.3. The other two conditions are imposed at the pure moment seed, by the following additional counts. For a bad family of dimension at least five, the rank-three witnesses are generically primitive on the evaluation curve, because projective lines meeting that curve form a family of dimension at most four. If the bad family contains all normals of three-label planes through one label b, a common root of a rank-three witness must be one of the finitely many roots of p_{gb}, and the three-planes through b in such an evaluation hyperplane form only a two-dimensional family; so the witnesses are generically primitive in this case as well. The possible single rank-two ray cannot dominate either condition.

We now use the relation-plane count of Proposition 9.3, with source of dimension m. For ℓ rationally independent witnesses let d be their common old source dimension, put $\delta = 3\ell - d$, and let ν be the dimension of the span of evaluation relations. The two codimensions are at least

$$\delta(m-d) + \ell + \nu(d - \ell - 140M), \qquad \delta(m-d) + 2\ell - 4 + \nu(d - \ell - 140M).$$

In the first count the parameter count per flag is $3d+5$ and the tuple fiber returns 5ℓ. In the second these numbers are $3d+3$ and 4ℓ, with four parameters for the marked label. The scalar degree and the source relation fee are as before. Take ℓ to be the smaller of the rational rank and $\lfloor m/2 \rfloor + 1$. In the truncated case both bounds exceed three, using $\nu \le \delta$ when the bracket is negative. If ℓ is the rational rank, then $\nu, \delta \ge 1$. Indeed, if the evaluations of the basis were independent everywhere, all rational coefficients of

the relations would be regular on the projective line, hence constant. Their contraction-drop locus would then be disjoint from the evaluation curve, so of dimension at most two, and so would its family of contained lines; this contradicts either prescribed normal image. With $\nu, \delta \geq 1$ both bounds again exceed three, since $m - \ell - 140M > 0$. Both avoidances are open by properness of the incidences.

Finally, we keep the prepared moment together with the short flat path to the resulting tensor. Such a path in the ordinary part of the curvature equation is obtained in its smooth local implicit charts, using the smoothly varying right-inverse complement over the compact base. All preparations can be taken arbitrarily small, so they retain the augmented injection and the source normal class. □

10.3. **The active map on a split-normal core.** All projective incidences in this subsection are over $\mathbb{C}$; a normal two-plane in $V_{\mathbb{C}}^*$ is viewed as a line in $\mathbb{P}(V_{\mathbb{C}}^*) \simeq \mathbb{P}^4$. For the bad normal family $\mathcal{B}$ of G on D put

$$\mathcal{E}_{\mathrm{bad}} = \{N : \dim\{M \in \mathcal{B} : M \cap N \neq 0\} \geq 3\}, \qquad \mathcal{C}_0 = \mathcal{B} \cup \mathcal{E}_{\mathrm{bad}}.$$

For a closed family $\mathcal{C}'$ of normal planes let $\mathcal{L}(\mathcal{C}')$ be the set of projective labels b satisfying one of

$$\dim\{M \in \mathcal{C}' : M \subset \operatorname{ann} b\} \geq 3,$$
$$\{M : c \subset M \subset \operatorname{ann} b\} \subset \mathcal{C}' \quad \text{for some line } c \subset \operatorname{ann} b.$$

The family in the second condition is a *restricted star*. Set $\mathcal{L}_0 = \mathcal{L}(\mathcal{C}_0)$, $\mathcal{F} = \{N : \mathbb{P}(\operatorname{ann} N) \subset \mathcal{L}_0\}$ and $\mathcal{C} = \mathcal{C}_0 \cup \mathcal{F}$.

Lemma 10.3. *Suppose $\mathcal{B}$ satisfies Lemma 10.2(iii). Then the sets above are closed, also in coefficient families,* $\dim \mathcal{E}_{\mathrm{bad}} \leq 2$, $\dim \mathcal{L}_0 \leq 2$, *and $\mathcal{F}$ is finite. For every N outside $\mathcal{E}_{\mathrm{bad}} \cup \mathcal{F}$,*

$$\dim\{M \in \mathcal{C} : M \cap N \neq 0\} \leq 2, \qquad \mathbb{P}(\operatorname{ann} N) \not\subset \mathcal{L}_0.$$

Proof. We first bound $\mathcal{E}_{\mathrm{bad}}$. Let $\mathcal{B}_i$ be an irreducible component of $\mathcal{B}$ of dimension $d \leq 4$. Its point-line incidence has dimension $d + 1$, and by the star exclusion all its point fibers have dimension at most two. Let $S_2 \subset \mathbb{P}^4$ be the locus of points whose fiber has dimension two. If $d = 4$ and the projection to $\mathbb{P}^4$ is dominant, the generic fiber has dimension one and $\dim S_2 \leq 2$ by upper semicontinuity. If $d = 4$ and the image Y has dimension three, the lines through a general point p of Y form a two-dimensional closed subset of the projective plane of tangent directions $\mathbb{P}(T_pY)$, hence fill it. Then Y is a hyperplane and $\mathcal{B}_i$ its whole line Grassmannian, which is excluded. An image of smaller dimension would violate the fiber bound. If $d \leq 3$, the generic fiber cannot have dimension two: the image would then have dimension at most two, and at a general smooth point it admits at most a one-dimensional family of tangent line directions. Hence $\dim S_2 \leq d - 2$. In all cases $\dim S_2 \leq 2$.

Suppose a normal line N meets a family of lines of $\mathcal{B}_i$ of dimension at least three. Discarding N itself, every remaining line meets N in one point, so by the fiber bound N is contained in the closure of S_2. Lines in a projective surface form a family of dimension at most two, by the same incidence and tangent argument. Hence $\dim \mathcal{E}_{\mathrm{bad}} \leq 2$ on the

whole Grassmannian of normal planes. In particular $\mathcal{C}_0$ has dimension at most four and still contains no full star.

Next let $\mathcal{C}'$ be any closed family of normal planes of dimension at most four without a full star; we show $\dim \mathcal{L}(\mathcal{C}') \le 2$. For the first condition, a component of dimension d has a label incidence $\{(b, M) : M \subset \operatorname{ann} b\}$ of dimension $d+2$, since $\mathbb{P}(\operatorname{ann} M)$ is a projective plane. If $d \le 3$, subtracting three for the fiber gives the bound. If $d = 4$, the incidence is dominant over the label space $\mathbb{P}^4$. Otherwise its image, of dimension at most three, would contain a four-dimensional family of projective planes; through a general smooth point the planes in the image give at most two dimensions of tangent planes, whereas the incidence requires at least three. Thus the generic fiber has dimension two, and the proper closed locus of larger fibers has image of dimension at most two. For the second condition, fix a center c. The normals through c in $\mathcal{C}'$ form a proper closed subset of $\mathbb{P}^3$, with finitely many plane components. A restricted star is such a plane and determines the label; its center lies in the locus S_2 of two-dimensional fibers estimated above. The only possible four-dimensional component with a larger center locus is the line Grassmannian of a hyperplane, which contributes only the label of that hyperplane. This proves the bound for both conditions.

The incidences are closed in coefficient families by properness and the fiber-dimension theorem. Containment of a restricted star and containment of a plane are closed conditions, because the complement of the incidence has open projection in the corresponding projective-space bundles. A projective plane in a set of dimension two is one of its irreducible components, so $\mathcal{F}$ is finite. Adding the finite closed set $\mathcal{F}$ changes neither condition in the definition of $\mathcal{L}$, so $\mathcal{L}(\mathcal{C}) = \mathcal{L}_0$, and the two displayed properties follow from the definitions of $\mathcal{E}_{\mathrm{bad}}$ and $\mathcal{F}$. □

For a three-plane $T(A) \subset V$, the map $G_D(\mathrm{Id} \otimes T|_A)$ is injective on $D \otimes A$ precisely when $\operatorname{ann}(T(A)_{\mathbb{C}}) \notin \mathcal{B}$. Consider the open first-order relation on active derivatives $T\colon A_4 \to V$ given by rank four and injectivity of

$$\Big(G_D(\mathrm{Id} \otimes T|_A), \sigma\Big) \quad \text{on } D \otimes A,$$

where the scalar row uses the original input $a \in A$. We add the avoidance condition $\operatorname{ann}(T(A)_{\mathbb{C}}) \notin \mathcal{E}_{\mathrm{bad}} \cup \mathcal{F}$, and call the resulting open relation $\mathcal{R}_{\mathrm{ord}}$.

Proposition 10.4. *After an arbitrarily small change of G in the ordinary flat locus over P, there are an active map on the product of P with the four original Hopf factors, with values arbitrarily close to z, and a compact path of rank-four maps $A_4 \to V$ from the included derivative I to its Hopf derivative, along which the augmented test on $D \otimes A$ is injective. Along the same path the test on $Q \otimes A$, with the three rows $(\pi_L h) \otimes a$, and its restriction to $E_v \otimes A$ over the link are injective.*

Proof. Admission of the included derivative. The included derivative can be brought into the avoidance condition by a small change of G in the ordinary flat locus and a small formal change of T over P. Indeed, $T \mapsto \operatorname{ann}(T(A)_{\mathbb{C}})$ maps rank-four maps submersively onto the six-dimensional Grassmannian of normal planes, and the avoided family has dimension at most two for every coefficient by Lemma 10.3. Stratifying the

algebraic incidence in these joint parameters and using the regular flat coordinate charts gives codimension at least four, which exceeds $\dim P = 3$. The augmented comparison at I, the ordinary and exterior rank tests, and the absence of the full label and normal families of Lemma 10.2 are open, so they persist under sufficiently small such changes. With the coefficient now fixed, the short path from I to the admitted T stays in the augmented comparison.

Avoidance templates. We verify the hypotheses of the avoidance template [MAdP21, Definition 4.5] for $\mathcal{R}_{\mathrm{ord}}$. For a finite list of principal covectors $\tau_j \in A_4^*$ we require that any at most three of their restrictions to A be linearly independent. For J of size one or two put

$$G_J = T\Big(A \cap \bigcap_{j \in J} \ker \tau_j\Big).$$

For a single index we require that only finitely many members of $\mathcal{C}$ annihilate G_J; for two indices we require that the label point G_J lie outside $\mathcal{L}_0$. These open conditions are coherent under deletion of indices. For density, fix $N = \operatorname{ann}(T(A)_{\mathbb{C}})$ outside $\mathcal{E}_{\mathrm{bad}} \cup \mathcal{F}$. By Lemma 10.3 the members of $\mathcal{C}$ meeting N form a family of dimension at most two. Each such member $M \neq N$ annihilates exactly one two-plane of $T(A)$, namely $\operatorname{ann}(M + N)$, while N annihilates all of them and contributes one point to every fiber. Thus a generic two-plane in $T(A)$ is annihilated by only finitely many members of $\mathcal{C}$. Since $\mathbb{P}(T(A)_{\mathbb{C}}) = \mathbb{P}(\operatorname{ann} N) \notin \mathcal{L}_0$, a generic line in $T(A)$ avoids $\mathcal{L}_0$. The maps $T \mapsto G_J$ are submersive for independent covector restrictions, so any finite list of these conditions is dense.

Ampleness. Fix i and a vector e_i with $\tau_i(e_i) = 1$, and consider the principal slice in which $T(e_i) \in V$ varies with T fixed on $\ker \tau_i$. Rank failure has codimension two in V. The frozen two-plane $G_{\{i\}}$ is annihilated by finitely many members of $\mathcal{C}$, and each possible failing three-plane $T(A)$ is annihilated by one of them, which is a codimension-two condition on $T(e_i)$. This also bounds the failure of the augmented test, since injectivity of $G_D(\mathrm{Id} \otimes T|_A)$ alone implies it. The conditions with $i \in J$ are fixed in the slice. For a one-index condition $\{j\}$, $j \neq i$, the moving two-plane $G_{\{j\}}$ extends the fixed label line $b = G_{\{i,j\}}$, which lies outside $\mathcal{L}_0$; the extension is a point of $\mathbb{P}(V/b) \simeq \mathbb{P}^3$. The members of $\mathcal{C}$ annihilating b give a family of projective lines there of dimension at most two, containing no restricted star. The locus of points lying on a positive-dimensional family of these lines has dimension at most one. For a two-dimensional component whose lines sweep out $\mathbb{P}^3$, this follows from its three-dimensional incidence by dropping one dimension and subtracting one. If the lines of a two-dimensional component sweep out only a surface, then at a general point their directions fill the tangent plane, so the surface is a plane and the component consists of all its lines; under annihilator duality this is a restricted star, which is excluded. A family of smaller dimension gives the bound directly. For a two-index condition $\{j, l\}$ with $i \notin \{j, l\}$, the moving label line ranges freely over $\mathbb{P}^4$ and the excluded labels have dimension at most two. The maps from the moving vector to these projective spaces are submersive away from rank failure. Hence all failures lie in a finite union of sets of real codimension at least two. The allowed open slice is therefore connected, as any two points are joined by a path of

two generic segments, and its convex hull is all of V, by opposite displacements. This verifies ampleness for every finite configuration.

Conclusion. We apply the foliated lifting and the relative C^0-close h-principle for avoidance [MAdP21, Lemma 4.7 and Theorem 5.1]. It gives an active map on the four original Hopf factors and a compact formal path in $\mathcal{R}_{\mathrm{ord}}$ from the admitted derivative to its Hopf derivative, with the augmented test on $D \otimes A$ throughout. We prepend the short admission path. The test on $Q \otimes A$ is injective along the path, with the three rows $(\pi_L h) \otimes a$ on the original inputs $a \in A$: their vanishing forces a kernel tensor to lie in $D \otimes A$. The same holds on $E_v \otimes A$. As in the exceptional case (Section 8), precomposition with normalized powers on the active factors and the inverse scaling of the values place the value path arbitrarily close to z without changing the derivative tests. □

11. Integration of the mixed coefficients and exact metric repair

11.1. **The old primitive on the exceptional link.** On the exceptional link the old Hopf coefficients have rank $r-1$ and source E_v, while the slow bundle $\mathcal{H}$ has rank r. We first produce a map u with values in $\mathcal{H}$ whose old derivative $d_E u\colon E_v \to \mathcal{H}$ is a monomorphism homotopic, through monomorphisms, to the inclusion $E_v \to \mathcal{H}$. Consider the first-order open relation on sections requiring $d_E u$ to have rank $r-1$. In an effective principal slice, transverse to an old hyperplane $H_0 \subset E_v$, the restriction $d_E u|_{H_0}$ is fixed and the remaining column must avoid the affine $(r-2)$-plane $d_E u(H_0)$ in an r-dimensional space; this has codimension two. On the other slices the relation is constant. Hence the relation is ample, and the relative open ample theorem [EM02], applied to the formal inclusion, gives a section u together with a homotopy of monomorphisms from the inclusion to $d_E u$. Using a finite chart partition, we approximate u in C^1 by polynomials in the ambient sphere coordinates with smooth slow coefficients. Injectivity and the augmented boundary comparison persist along the short linear path. No global frame of E_v or of $\mathcal{H}$ is needed.

Let G be the exactly flat tensor of Section 8.4, with the slow placement of the bank, let $v_0\colon V \to W$ be a fixed slow frame of rank five, and put $v = v_0 + \lambda G(u, \cdot)$ with $\lambda > 0$ small. Since v_0 and G are slow, $d_E v = \lambda G(d_E u, \cdot)$ and

$$v^* d_E v = \lambda\, v_0^* G(d_E u, \cdot) + \lambda^2\, G(u, \cdot)^* G(d_E u, \cdot).$$

The last matrix is symmetric by the alternating Gram identity expressing the flatness of G. Consequently the skew part of $v^* d_E v$ has the primitive

$$P = \operatorname{skew}\Big(v_0^*(v - v_0)\Big), \qquad d_E P = \operatorname{skew}(v^* d_E v).$$

All coefficient derivatives here keep the slow quantities fixed. The mixed tensor is $d_E v$, and the four bank rows in the formulas of Section 8.6 are evaluated on $d_E u$. The homotopy from the inclusion to $d_E u$ preserves the augmented comparison: the comparison is precomposed with a monomorphism on the first factor, and the paired comparison is injective on all of $\mathcal{H} \otimes A$. This gives the old variable on the exceptional link, in the original target and with zero passive term.

11.2. **Derivatives of the old map on the exterior.** Let G be the exterior tensor of Proposition 9.3, with the relative ordinary preparation of Section 10.2 where it applies. On the collar band we use the generator and the placement of the bank constructed in Section 15.1.

For the generating section u of $\mathcal{E}$ on $B \times \mathcal{O}_{\mathrm{old}}$ we impose two conditions: $d_E u$ is injective, and its projection to E has rank at least $m-1$. In an effective principal slice the first condition excludes an affine subspace of codimension $(m+k)-(m-1) = k+1$. The second imposes no condition if the other projected columns have rank $m-1$, and excludes an affine subspace of codimension two if they have rank $m-2$. In either case the complement is connected and has full convex hull, and the formal inclusion satisfies the relation. On the inner band the prescribed generator and its derivative homotopy take values in E_v plus a single generating line, and there both rank conditions follow from injectivity. The relative open ample theorem therefore realizes the formal extension by a section u.

On the distant exterior we multiply the auxiliary coordinates of u by a small positive constant. Let the multiplier tend to zero. Any limit of the complex old image planes contains an $(m-1)$-dimensional subspace of $E_{\mathbb{C}}$: choose $m-1$ domain vectors whose projected images are independent at the limiting old point. Hence the limiting plane lies in $E_{\mathbb{C}}$ plus at most one auxiliary line, and all of its hyperplanes are good by Proposition 9.3. By compactness of the flag bundle and openness of goodness, one small positive multiplier works on the whole compact product. We interpolate to the multiplier one inside the product band, where the generator still takes values in E_v plus its generating line; for every positive multiplier all hyperplanes there are reduced planes of the required type. The multipliers are slow, so they enter the old derivative without a derivative term.

Write $A_u = d_E u$. The condition that $G(A\,\cdot\,,\cdot\,)$ is good is Zariski open in the matrix coordinates (G, A). Choose finitely many polynomials defining its complement. At each slow value of G at least one of them, written $g(G, A)$, is a nonzero polynomial in A, since a good plane exists. On a slow chart carrying such a polynomial we arrange that $f = g(G, A_u)$ has a nonzero old directional derivative of some finite order, possibly zero, at every point. This rests on the following jet estimate.

In a commuting old frame with first field Y, the truncated Taylor curve of A_u along Y through order J can be prescribed freely by the section jets of order $J+1$: the first old column uses $Y^{j+1}u$ and the other columns use $Y^j Y_i u$, $0 \le j \le J$, and these are distinct symmetric multi-indices. Let D_g bound the degrees of the defining polynomials. For fixed G and degree $d > 0$, select a scalar direction of A on which the highest homogeneous part of $g(G, \cdot)$ is nonzero; finitely many directions suffice for all degrees at most D_g. After division by a constant, the polynomial in this variable is monic of degree d. Fix all other coefficients of the curve jet and write the resulting polynomial as $q(s, v)$. Suppose that $d+1$ scalar jets $v_0(s), \dots, v_d(s)$, with pairwise distinct truncations modulo $s^{\lfloor J/d \rfloor + 1}$, satisfy $\operatorname{ord}_s q(s, v_i(s)) > J$. Lagrange interpolation over $\mathbb{C}((s))$ recovers

the leading coefficient:

$$1 = \sum_{i=0}^{d} \frac{q(s, v_i(s))}{\prod_{j \neq i}(v_i(s) - v_j(s))}.$$

Each difference $v_i - v_j$ has order at most $\lfloor J/d \rfloor$, so every denominator has order at most $d\lfloor J/d \rfloor \leq J$, and the right side has positive order, a contradiction. Thus the admissible truncations form at most d points, and vanishing through order J costs at least $\lfloor J/D_g \rfloor + 1$ jet conditions. When the polynomial is a nonzero constant there is no bad jet. We take J so large that this codimension exceeds the source dimension even after all independent matrix parameters G are allowed. The finitely stratified algebraic jet incidences can then be avoided by C^1-small jet perturbations on buffered charts. Since we project over the whole matrix parameter space before taking closures, no transversality of the particular slow section G is required.

These perturbations are first made on the exceptional link, where all auxiliary coordinates are allowed and the compact boundary comparison is kept by C^1-smallness. The result is carried radially on the inner band and joined to the unperturbed generator before the auxiliary multiplier changes. Finite contact is an open condition on a compact buffer of the fixed band, witnessed by finitely many derivatives, and the remaining jet perturbations are made relative to this buffer. All hyperplane tests continue to hold by compactness and openness. With one sufficiently small $\lambda > 0$, the map $v = v_0 + \lambda G(u, \cdot)$ is an injective old frame and has the exact skew primitive P of the previous subsection on the double.

The angular active map extends to the double as an activity of rank four. On the collar we extend the formal jets, using the given derivative path to the constant inclusion on the far side, and we extend the values and transverse formal derivatives freely. In an effective principal slice the rank-four relation excludes an affine three-plane in $\mathbb{R}^5$, so the relative open ample theorem realizes the extension by a map a. Later C^1-small perturbations on the collar band keep its boundary comparison by a short interpolation.

For every such a, the mixed tensor $(d_E v)(\mathrm{Id}_E \otimes d_A a)$ satisfies the tilted phase test everywhere. Choose any real old hyperplane. Its complexified image under A_u is good for G, hence for v when λ is small, so its bad family does not contain the complex star centred at the annihilator line of the active four-plane $d_A a(A)$. The real normal planes of that star are Zariski dense in it; hence some real normal plane containing the annihilator line lies outside the bad family, and its annihilator is a real three-plane passing the injection test. This three-plane is the image of an active hyperplane under $d_A a$. Taking the maximum of the least singular value over the two real hyperplane choices gives a positive margin, uniform by compactness.

11.3. **Differential identities across an old rank change.** The next lemma concerns operators in the old and active variables. Write M for the active part of the mixed metric operator and $\mathcal{Q} = M + d_E$ for the operator that also allows an active potential differentiated in the old variables. The old fields commute with the active fields; in particular multiplication by a function of the old variables alone commutes with M.

Lemma 11.1. *Suppose that on a chart there is a finite-order differential operator R with $\mathcal{Q}R = f^p \,\mathrm{Id}$, where f is a smooth function of the old variables alone and p is a positive integer. Let Y be a fixed old coordinate field, and suppose that for the old covector $\nu = d_E f$ there is a conormal identity*

$$MS_\nu + \nu \otimes V_\nu = \nu(Y)^b \,\mathrm{Id}, \qquad b \geq 0,$$

where S_ν, V_ν are finite-order operators differentiating only in the active variables, with coefficients smooth through $\nu = 0$. Then $(Yf)^{(b+1)p} \,\mathrm{Id} = \mathcal{Q}R'$ for a finite-order differential operator R'. If such conormal identities are also available for the successive derivatives $Yf, Y^2 f, \ldots$, and $Y^j f \neq 0$ at a point, then $\mathcal{Q}$ has a finite-order differential right inverse near that point.

Proof. Put $t = Yf = \nu(Y)$; it depends only on the old variables. For every $k \geq 0$ and every datum h, the conormal identity and the product rule give

$$f^k t^b h = M(f^k S_\nu h) + d_E\Big(\frac{f^{k+1}}{k+1} V_\nu h\Big) - \frac{f^{k+1}}{k+1} d_E(V_\nu h).$$

The first two terms lie in the image of $\mathcal{Q}$. We show by induction on $k \leq p$ that

$$t^{(b+1)p} h = \mathcal{Q}(W_k h) + f^k t^{(b+1)(p-k)} L_k h$$

for finite-order operators W_k, L_k; for $k = 0$ take $W_0 = 0$, $L_0 = \mathrm{Id}$. Given step $k < p$, write the residual as $f^k t^b \tilde{h}$ with $\tilde{h} = t^{(b+1)(p-k-1)+1} L_k h$ and apply the displayed identity to $\tilde{h}$. The active operator S_ν commutes with the powers of t, and d_E in the new residual removes at most one of them. The new residual therefore has the factor $f^{k+1} t^{(b+1)(p-k-1)}$, and no division by t occurs. At $k = p$ the residual $f^p L_p h = \mathcal{Q}(R L_p h)$ is absorbed by the given multiple, so $(Yf)^{(b+1)p} \,\mathrm{Id} = \mathcal{Q}(W_p + R L_p)$ on the whole chart.

Iterating with $Yf, Y^2 f, \ldots$ in place of f gives $(Y^j f)^N \,\mathrm{Id} = \mathcal{Q}\widetilde{R}$, and $h \mapsto \widetilde{R}\Big((Y^j f)^{-N} h\Big)$ is a right inverse where $Y^j f \neq 0$. On a compact source, a finite partition of unity $\{\psi_\alpha\}$ patches these local inverses: apply each local inverse to $\psi_\alpha h$ and extend by zero. Since $\mathcal{Q}$ is local, the sum is a right inverse. □

The argument requires conormal identities. A pointwise good hyperplane is only a rank condition on a tensor; it supplies neither finite differential identities nor their localization at parameters where the old rank changes.

11.4. **The active parameter input.** The required active input is stated on joint algebraic parameter charts. For the old ranks $s = m, m-1$, let $\mathcal{P}_s$ be the open locus of good tensors in the matrix space. We use the complexified spheres $\{(z, \zeta) : \zeta^{\mathsf{t}} z = 1\}$ and their active torus generators. For a polynomial activity with smooth slow coefficients, we require the following: on algebraic charts of its coefficient and source parameters that contain all real points, the intrinsic adjoint rows of the active operator generate the whole module of differential rows over $\mathcal{P}_s$, at both ranks. The adjoint rows are normal contraction and active divergence. The operator ring is generated by the algebraic functions and the four orbit derivations, which annihilate the matrix and coefficient parameters.

This input yields the joint inverse on the exterior. Take G, A to be independent algebraic variables and put $U = G(A\cdot, \cdot)$, including the nonzero factor λ if needed. Choose a principal matrix chart on which all complex old hyperplanes of U are good; since goodness is invariant under conjugation, the defining polynomial g may be taken with real coefficients. Over $g(G, A) \neq 0$ the row quotient of the active adjoint vanishes by the input at rank m. These row quotients localize as coefficient modules; normal ordering after localization uses the Leibniz rule, and the orbit derivations annihilate g. A coefficient module on an affine chart vanishes if all its localizations at maximal ideals vanish, by considering the annihilator of a nonzero element. Thus, after inverting g, each order-zero flux row lies in the span of the adjoint rows. Since $d_\alpha g = 0$, all powers of g^{-1} can be moved to the left, and a common finite power clears them. Substituting the slow parameters and the old coefficients, and then taking formal adjoints, we obtain

$$MR = f^p \,\mathrm{Id}, \qquad f = g(G, A_u),$$

where M is the active operator. The substitutions commute with all products of active operators, since the substituted quantities are annihilated by the orbit derivations. Formal adjoints are taken with respect to the invariant Hopf volume and an active-invariant old density, and real parts give real identities.

For the conormal test we adjoin an independent algebraic covector variable ν and add $\nu \otimes \mathrm{Id}$ to M. Where $\nu(Y) \neq 0$, the kernel hyperplane of ν has the frame $Y_r - \frac{\nu(Y_r)}{\nu(Y)} Y$, $r \geq 2$, obtained by eliminating the Y-coordinate in an old frame. We solve the restricted problem by the identity at rank $m - 1$ and extend its old potential by zero on Y. The residual of the original datum is then ν times an active potential, namely its Y-component divided by $\nu(Y)$. Hence the row quotient of the augmented adjoint vanishes after inverting $\nu(Y)$. Clearing a common finite power on the whole affine chart and taking formal adjoints gives

$$MS_\nu + \nu \otimes V_\nu = \nu(Y)^b \,\mathrm{Id},$$

with coefficients polynomial in ν, hence valid through $\nu = 0$. We use this identity at the covectors $d_E f, d_E(Yf), \ldots$. Lemma 11.1 and the finite contact of f give a local joint inverse everywhere, and a finite partition of unity subordinate to the input charts gives the exact joint identity on the product.

11.5. **Exact metric production.** We now carry out the analytic completion in the fixed target dimension, assuming the mixed inverse on the whole compact exterior. Let $F_0 = (F_\mathrm{o}, B)$ be the reference product; its old and active complete columns are independently free on the original old factors and on the four active Hopf factors, respectively. The selected distribution $D = E \oplus A_4$ has rank n, its two blocks are orthogonal, and its partial Gram matrix is denoted by h_0. Let

$$I_x \colon W \to (\text{span of the first and pure second columns})^\perp$$

be the normal isometry. The partial fields commute in slow frames. Let v, f, a be the old frame, the old passive reader, and the active map, and write $w = va$, $c = fa$ and

$C_{i\alpha} = d_i d_\alpha w$. Assume that for an old-basic function d the skew identity

$$d_E P = \mathrm{skew}(v^* d_E v + d^2 f^* d_E f)$$

holds. Assume further that on the full product the operator

$$(*) \qquad (\eta, u, b) \longmapsto \Big(\langle C_{i\alpha}, \eta\rangle + d_\alpha u_i + d_i b_\alpha\Big)_{i,\alpha}$$

has a smooth finite-order differential right inverse. Finally assume that at each point there are old and active covectors $\xi_{\mathrm{o}}, \xi_{\mathrm{a}}$ such that C is injective on $\ker \xi_{\mathrm{o}} \otimes \ker \xi_{\mathrm{a}}$. The covectors range over unit spheres, so this pointwise condition yields a positive best margin by compactness.

We use independent sphere variables x (for the map) and y (for the coefficients) over the same slow base. For an integer L_0 put $\tau = L_0^{-1}$; eventually y is evaluated at the normalized coordinatewise L_0-th powers $P_{L_0}(x)$ of x. These powers intertwine the Hopf actions, $P_{L_0}(e^{2\pi \mathrm{i}\theta} z) = e^{2\pi \mathrm{i} L_0 \theta} P_{L_0}(z)$, so the differential of an evaluated profile is the evaluation of $D_i = d_i^x + \tau^{-1} d_i^y$. For an integer $k \geq 4$ we start the retained part at $F_0 + \tau^k I_x w(y)$. The passive metric term is included as the separate tensor $\tau^{2k} d^2 (\mathrm{d}_D c_{\mathrm{ev}})^* \mathrm{d}_D c_{\mathrm{ev}}$; it uses only separate sources and no passive output coordinate.

The exact linearized strain of a profile Z is

$$D_j \langle (F_0)_i, Z\rangle + D_i \langle (F_0)_j, Z\rangle - 2\langle (F_0)_{ij}, Z\rangle.$$

For $Z = I_x Z'$ it vanishes identically, including the derivatives of the normal isometry: the pairings of $I_x Z'$ with the first and pure second columns vanish, and the mixed reference columns are zero. Inverting the Gram matrix of the complete selected columns gives two pointwise lifts. The first-pairing lift has prescribed first pairings and zero pure pairings; the pure-pairing lift has prescribed pure pairings and zero first pairings.

Since $d_i^y a = 0$ and $d_\alpha^y v = d_\alpha^y f = 0$, the mixed quadratic coefficient at order $2k-2$ is

$$d_i^y \Big\langle d_\alpha^y a, \Big(\tfrac{1}{2}(v^* v + d^2 f^* f) + P\Big) a \Big\rangle.$$

This follows from the skew identity, which gives $v^* d_i^y v + d^2 f^* d_i^y f = d_i^y \Big(\tfrac{1}{2}(v^* v + d^2 f^* f) + P\Big)$, together with $d_i^y d_\alpha^y a = 0$. We cancel it by the first-pairing lift at order $2k-1$, whose $\tau^{-1} d_i^y$-derivative contributes at order $2k-2$, and we cancel the remaining pure coefficients by the pure-pairing lift at order $2k-2$.

Suppose that all orders less than $j > 2k-2$ have been cancelled. We add a profile $\tau^{j-k+2} I_x Z'$ in the complete normal direction. In the mixed rows, its quadratic cross term with the leading profile at order j is

$$d_\alpha^y \langle d_i^y w, Z'\rangle + d_i^y \langle d_\alpha^y w, Z'\rangle - 2\langle C_{i\alpha}, Z'\rangle.$$

All derivatives falling on I_x or on the slow part of the profile carry at least one extra power of τ at this step. We apply the inverse of $(*)$, with x treated as a parameter, to the negative of the remaining mixed coefficient, and for its solution (η, u, b) we choose $Z' = -\eta/2$. The first-pairing lift at order $j+1$ supplies $u_i - \langle d_i^y w, Z'\rangle$ and $b_\alpha - \langle d_\alpha^y w, Z'\rangle$, and the pure-pairing lift at order j cancels the remaining pure coefficients. Earlier orders are unchanged, because $j - k + 2 \geq k + 1$ and the new self-interaction and the cross

terms with previous corrections have larger order. By induction we obtain finite profiles that solve the metric equation to any prescribed order J_{met} before evaluation.

As $\tau \to 0$, the first columns and the pure second columns converge uniformly to those of F_0, and the mixed columns divided by τ^{k-2} converge to $I_x C(y)$. For a permitted unit pair of old and active covectors put $\xi = (\xi_{\mathrm{o}} + \xi_{\mathrm{a}})/\sqrt{2}$. The kernel of the pure projection on $\mathrm{Sym}^2 \ker \xi$ is exactly $\ker \xi_{\mathrm{o}} \otimes \ker \xi_{\mathrm{a}}$, because the contraction of a mixed tensor with ξ splits into its contractions with the two covectors. We use the independent reference first and pure columns on the orthogonal complement of this kernel and the scaled mixed columns on the kernel. Their limiting images are independent with a uniform lower bound, by the compact best margin. After restoring the scaling and projecting off the first span, this gives a phase margin of at least $c_1 \tau^{k-2}$.

The normalized coordinatewise powers of a fixed complex sphere cost $O(L_0^j)$ in every fixed C^j norm. Indeed their components are

$$P_{L_0}(z)_a = \frac{z_a^{L_0}}{\left(\sum_b |z_b|^{2L_0}\right)^{1/2}},$$

and on the sphere some coordinate has modulus at least a fixed $c_0 > 0$. Each derivative of a numerator divided by the largest modulus to the power L_0 is bounded by $C_j L_0^j c_0^{-j}$, using $|z_a|^{L_0 - j}$ for $L_0 \geq j$. Likewise the derivatives of the sum in the denominator, relative to the sum itself, are $O(L_0^j)$, and the reciprocal is handled by the chain rule. These estimates patch over a fixed finite atlas. Consequently the evaluated metric error is $O(\tau^{J_{\mathrm{met}}+1-q})$ in every fixed C^q norm, and all finite derivatives of the reference map and of the metric are polynomially bounded in τ^{-1}.

To remove this error we need the following quantitative form of Proposition 3.1. On a fixed compact foliated source, suppose the initial data have ellipticity in a fixed interval, reference derivatives bounded by $O(\tau^{-a_0})$ up to one fixed order, and best phase margin at least $c_1 \tau^{b_0}$. Then an error $O(\tau^K)$ in full C^{440} can be corrected by a displacement of at most τ^q in full C^2, where K is finite and depends only on a_0, b_0, q and the fixed dimensions. We indicate the uniformity. For the transverse-minor polynomial used in the proof of Proposition 3.1, the lower bound on the best phase gives a polynomial lower bound on its norm in the finite-dimensional space of polynomials. On any fixed small ball of directions, the values on a finite interpolation grid bound this norm from below by a uniform constant times the global norm. Hence each direction of the positive-square template can be replaced by a nearby good direction whose restricted singular value has a polynomial lower bound; the positive-square decomposition itself is uniformly stable by ellipticity. We freeze the phases on charts of radius τ^A, with A chosen from the finite reference jet bounds and the phase margin. A cover on this mesh has uniformly bounded enlarged overlap, and its cutoff derivatives are bounded by powers of τ^{-1}. The fixed number of phase labels is separated by successive polynomial weights in τ^{-1}, so that each paired phase keeps the margin of its dominant phase after normalization. Inverse Gram matrices, phase-coordinate changes and square roots in the finite formulas then obey polynomial bounds, by the polynomial bounds on their denominators and the product rule. The initial thresholds of stage ten of the iteration are finitely many

inequalities with positive powers of its starting scale; we choose this scale to be τ^{μ} with μ large, which meets the geometric margins and the prescribed displacement τ^{q}, and we choose K large enough for the normalized initial C^{440} defect. The subsequent finite waits are carried out for each fixed τ as in the proof of the proposition and require no new uniform initial error order. The derivative order of the reference map required at the start is determined by the same fixed stage, independently of a_0, b_0, q.

We choose k larger than this finite reference order plus three. Then, for every truncation order, the reference maps have bounded required jets, since all new profile exponents are at least k. The target $h_0 - \tau^{2k} d^2 |\mathrm{d}_D c_{\mathrm{ev}}|^2$ differs from h_0 by $O(\tau^{2k-2})$, so it has fixed ellipticity for small τ. With $b_0 = k - 2$ we prescribe $q = 2k$ and choose J_{met} so large that $J_{\mathrm{met}} + 1 - 440 > K$. This gives an exact smooth retained map with the stated target, within τ^{2k} in full C^2 of the computed map; the phase margin and the target dimension are preserved. The completion uses a finite-order differential inverse of (∗) on the compact host; this input cannot be replaced by pointwise surjectivity of a symbol.

11.6. **The active cofactor complement.** On the attaching collar $H = E_v \oplus A$. In the reference product we keep all complete old columns, the four active first columns, and the six pure second columns in $\mathrm{Sym}^2 A$. Their orthogonal complement is $I_x W \oplus C_A(x)$, where $C_A(x)$ is the rank-four complement of these columns inside the complete active span. The projections of the LA and LL second columns identify $C_A(x)$ with $L \otimes A \oplus \mathrm{Sym}^2 L$. The normal isometry splits into an old block and an active block. In a slow coefficient frame the old derivative of I_x therefore has zero projection to C_A, while the active derivative has the cofactor response of Section 8.6. These are the reference columns for the boundary comparison of Section 17.

12. Characteristic directions of the active operator

In this section we carry out the algebraic computation on graphs that is needed for the active parameter problem of Section 13. All spaces and dimensions in this section are complex. Let $\mathcal{B}$ be a closed algebraic family of projective lines in $\mathbb{P}^4$, of dimension at most four and containing no full star. Let $x \mapsto (x, f(x))$ be a local holomorphic graph over $\mathbb{C}^4$, with affine normal $\eta(x) = (-\mathrm{d}f(x), 1)$, represented in the affine chart $\{p_0 = 1\} \subset \mathbb{P}^4$ by $p(x) = \mathrm{d}f(x)$. Its *raw cotangent cone* is the set of $(x, \xi) \in T^*\mathbb{C}^4$ with $\xi \neq 0$ such that the projective line spanned by $(1, p(x))$ and $(0, \xi)$ belongs to $\mathcal{B}$; thus it consists of the directions ξ of the lines of $\mathcal{B}$ through the normal $p = \mathrm{d}f(x)$. We say that the Gauss map $x \mapsto \mathrm{d}f(x)$ is *quasifinite* if its fibres are zero-dimensional.

A complex submanifold Λ of $T^*\mathbb{C}^4 \setminus \{0\}$ is *coisotropic* if $(T_q\Lambda)^{\omega} \subset T_q\Lambda$ for every $q \in \Lambda$, where ω is the canonical symplectic form; equivalently, the Poisson bracket of two local functions vanishing on Λ again vanishes on Λ. A germ invariant under the scalings $(x, \xi) \mapsto (x, \lambda\xi)$, $\lambda \in \mathbb{C}^*$, is called *conic*. A conic coisotropic germ has dimension at least four, so its fibrewise projectivization has dimension at least three.

Lemma 12.1. *Let J be a positive integer. If the raw cone of a graph with quasifinite Gauss map contains a nonzero conic coisotropic germ, then the graph satisfies one of*

finitely many algebraic J-point incidence conditions on jets of order at most five at J distinct source points. After the source points and the normal and line marks are allowed to vary, and the fibre dimension supplied by the germ is taken into account, each of these incidences has codimension at least J in any coefficient space that interpolates such jets independently at J distinct source points. These estimates hold fibrewise, at fixed parameters of the line family.

Proof. We work with an irreducible component $\mathcal{B}_0$ of the line family. Its marked incidence

$$I = \{(p, \ell) \in \mathbb{P}^4 \times \mathcal{B}_0 : \ p \in \ell\}$$

has dimension $\dim \mathcal{B}_0 + 1$. Since $\mathcal{B}$ contains no full star, the lines of $\mathcal{B}$ through a point p form a proper closed subset of the $\mathbb{P}^3$ of all lines through p; hence every projective direction fibre has dimension at most two.

Dominant components: the common image. Suppose first that $I \to \mathbb{P}^4$ is dominant. Then $\mathcal{B}_0$ has dimension $d \in \{3, 4\}$. At a general smooth line ℓ, the conormal space of $\mathcal{B}_0$ in the Grassmannian is a $(6-d)$-dimensional subspace of $\mathrm{Hom}(\mathbb{C}^5/\ell, \ell)$. We claim that it cannot, for generic ℓ, consist of maps with one common image line $c_\ell \subset \ell$. Indeed, suppose it does. Since $6 - d \in \{2, 3\}$ and $\dim \mathrm{Hom}(\mathbb{C}^5/\ell, c_\ell) = 3$, this defines a rational focal map $\ell \mapsto [c_\ell]$, and the conormal space consists exactly of the maps with image in c_ℓ that vanish on a common kernel $\Pi_\ell/\ell \subset \mathbb{C}^5/\ell$ of dimension $3 - (6 - d) = d - 3$. The preimage $\Pi_\ell \subset \mathbb{C}^5$ is the vector space of a projective $(d-2)$-plane containing ℓ. For a tangent variation $\delta\colon \ell \to \mathbb{C}^5/\ell$ of $\mathcal{B}_0$, the conormal equations require $\delta(c_\ell) \in \Pi_\ell/\ell$. Differentiating $c_\ell \subset \ell$ therefore shows that, at a general point, the projective tangent space of the focal image lies in $\mathbb{P}(\Pi_\ell)$. Let $s \leq d - 2$ be the dimension of the focal image.

- If $s = 0$, all lines of $\mathcal{B}_0$ pass through one point and lie in its star. This is impossible when $d = 4$, and when $d = 3$ it makes $\mathcal{B}_0$ a full star, which is excluded.
- If $s = d-2 > 0$, then $\mathbb{P}(\Pi_\ell)$ is the tangent plane of the focal image at $[c_\ell]$, and all lines with this focal point lie in the $(s-1)$-dimensional family of lines through $[c_\ell]$ in that plane. Hence $\dim \mathcal{B}_0 \leq s + (s-1) = 2s - 1 < d$, a contradiction.
- In the remaining case $d = 4$, $s = 1$, the family consists of four-dimensionally many lines meeting one irreducible curve. The family of all lines meeting that curve is irreducible of dimension four, so it equals $\mathcal{B}_0$ and contains a full star. This is also excluded.

Consequently the common-image locus, together with the singular locus of $\mathcal{B}_0$, is a closed line subfamily of smaller dimension, which still has no full star. It is handled by descending induction on the dimension.

Dominant components: the regular part. Off the closure C_{crit} of the critical normal values, the projection $I \to \mathbb{P}^4$ is smooth. The closure C_{crit} is a proper subvariety, as one sees by stratifying the critical set by the rank of the projection, which is at most three there. On the smooth locus we use local conic equations $P_i(p, \xi)$, $1 \leq i \leq 6 - d$, for the incidence, homogeneous of degree one in ξ, with linearly independent $a_i = \partial_\xi P_i$; such

equations are obtained from local Plücker equations divided by powers of a nonzero coordinate. Write $b_i = \partial_p P_i$. For the line spanned by $(1, p)$ and $(0, \xi)$, the conormal maps are the combinations of $(1, p) \otimes b_i + (0, \xi) \otimes a_i$; in particular, both coefficient covectors annihilate the line after homogeneous extension. Suppose that for every pair i, j the functional

$$Q_{ij}(H) = b_i^{\mathrm{t}} H a_j - b_j^{\mathrm{t}} H a_i$$

vanishes on all symmetric matrices H. Then $b_i \odot a_j = b_j \odot a_i$ as symmetric tensors, and comparison in a basis containing the independent covectors a_i, a_j gives $b_i = \lambda a_i$ with a common scalar λ. All conormal maps then have image in the line spanned by $\lambda(1, p) + (0, \xi)$, which is the excluded common image. Thus on the regular part there is a pair i, j for which $Q = Q_{ij}$ is a nonzero functional on $\operatorname{Sym}^2 \mathbb{C}^4$.

Put $e_i(x, \xi) = P_i(\mathrm{d}f(x), \xi)$. On a smooth part of a coisotropic germ the bracket of two functions vanishing there again vanishes, because their Hamiltonian vector fields are tangent to the germ. For the selected pair this gives

$$\{e_i, e_j\} = -Q\Big(D^2 f\Big) = 0,$$

and also $\{e_i, \{e_i, e_j\}\} = \{e_j, \{e_i, e_j\}\} = 0$. With the normal and the Hessian fixed, the last two conditions on the third derivative $T_3 = D^3 f$ have linear parts $-Q(T_3(a_i, \cdot, \cdot))$ and $-Q(T_3(a_j, \cdot, \cdot))$. They are independent, since the corresponding cubic polynomials are the same nonzero quadratic polynomial multiplied by the two independent linear polynomials a_i and a_j. We thus obtain three independent conditions on the jets of order two and three.

Now we count. The marked incidence has dimension at most five, and the source point contributes four parameters. The equation $\mathrm{d}f(x) = p$ imposes four conditions and the bracket identities three more, so each marked point increases the dimension by at most $5 + 4 - 4 - 3 = 2$. On the other hand, the projectivization of the conic coisotropic germ has dimension at least three, which gives fibres of dimension at least three for each mark. Distinct source points exist, since the projective direction fibres have dimension at most two, so the source projection of the germ is positive-dimensional. Returning the $3J$ dimensions of the fibres against the increase of at most $2J$ gives the desired codimension J.

Exceptional normals and lower-dimensional components. For the exceptional normal values of a dominant component, the part of the marked incidence over the proper subvariety C_{crit} has dimension at most four. For a nondominant component of dimension at most three, the marked incidence also has dimension at most four. Stratify the relevant proper normal image into smooth strata S of dimension $s \le 3$ over which the projective fibre dimension f_S is constant; thus $s + f_S \le 4$. Write F_p for the affine direction cone at p, including zero; it has dimension $f_S + 1$. By quasifiniteness, the projectivized raw cone over $(\mathrm{d}f)^{-1}(S)$ has dimension at most $s + f_S \le 4$, so a conic coisotropic germ there has dimension four or five.

Let Λ be a conic Lagrangian germ, that is, of dimension four, and let Y be its source projection. At a smooth point of the projection, Λ is an open part of the conormal bundle of Y: the Liouville form vanishes on Λ, being the contraction of the symplectic

form with the radial vector, and the dimensions agree. Consequently

$$D^2 f(N_p^* S) \subset N_x^* Y \subset F_p, \qquad p = \mathrm{d}f(x).$$

The first inclusion follows by differentiating $\mathrm{d}f(Y) \subset S$ and using the symmetry of the Hessian; singular Hessians are allowed. We show that this cone containment imposes at least two conditions on a symmetric Hessian. If $\dim N_p^* S \geq 2$, choose two independent conormals n_1, n_2. As H ranges over symmetric matrices, the pairs (Hn_1, Hn_2) range over the seven-dimensional hyperplane

$$\mathcal{H}_{12} = \{(y_1, y_2) : \ n_2^{\mathsf{t}} y_1 = n_1^{\mathsf{t}} y_2\},$$

and the span of Hn_1, Hn_2 must lie in one irreducible component C_0 of F_p. If $\dim C_0 \leq 2$, the condition costs at least $7-4 = 3$ in $\mathcal{H}_{12}$. If $\dim C_0 = 3$, then $C_0 \times C_0$ has dimension six and $\mathcal{H}_{12}$ cuts it properly: otherwise, taking $y_1 = 0$ and then $y_2 = 0$, the cone C_0 would lie in both independent hyperplanes $\ker n_1$ and $\ker n_2$, which is impossible for a three-dimensional cone. This costs at least $7 - 5 = 2$. If $\dim N_p^* S = 1$, then $s = 3$, hence $f_S \leq 1$ and $\dim F_p \leq 2$, while the single Hessian value Hn_1 ranges over the full four-space; this also costs two conditions. Thus the marked count gains at most $4 + 4 - 4 - 2 = 2$ dimensions per point, while the germ returns at least three per point.

Let now Λ have dimension five on these exceptional strata. By quasifiniteness it is a top component of the raw cone there: $s + f_S = 4$, it has full source dimension s over the normal stratum, and it fills a component C_0 of the direction fibre. By coisotropy its vertical tangent contains the conormal of its source projection. Hence, for a nonzero $n \in N_p^* S$, the vertical tangent $T_\xi C_0$ contains $D^2 f\, n$. At a generic smooth point of C_0 this tangent is a proper linear subspace, so at least one nonzero linear Hessian equation is imposed. The marked count gains at most three dimensions per point and returns at least four. Germs of larger dimension do not fit into this part of the raw cone.

Nondominant four-dimensional components. Finally, let $\mathcal{B}_0$ be a nondominant component of dimension four, and let s be the dimension of its normal image. At a generic smooth point of the image, each line of $\mathcal{B}_0$ through the point lies in the image, hence in its projective tangent space, while these lines form a family of dimension at least $5 - s$. Since lines through a point of $\mathbb{P}^s$ form a $\mathbb{P}^{s-1}$, we get $5 - s \leq s - 1$. Hence $s = 3$, the tangent directions at the point fill the whole $\mathbb{P}^2$, and the projective tangent hyperplane lies in the image. By irreducibility, $\mathcal{B}_0$ is the family of all lines in a fixed projective hyperplane. In an affine chart meeting the hyperplane, write it as $n^{\mathsf{t}} p = c$. The raw equations are $n^{\mathsf{t}} \mathrm{d}f = c$ and $n^{\mathsf{t}} \xi = 0$. Their repeated brackets force the five conditions

$$\partial_n f = c, \qquad \partial_n^2 f = \partial_n^3 f = \partial_n^4 f = \partial_n^5 f = 0,$$

which are independent as scalar jet conditions. There are only four parameters for the source point, and no additional normal marks. By the bound on projective fibres, a coisotropic germ has positive-dimensional source projection, so it supplies the required number of distinct points. This yields the same saving of J.

Parameters. All these counts use finitely many algebraic charts and their fibre-dimension loci. The parameters held fixed may also be stratified, including the components of the fibres after finite algebraic covers. To obtain these covers, one defines

the finitely many components of the geometric generic fibre over a finite field extension, spreads the equations over a smaller open set, and repeats the construction on its closed complement. The dimension bounds at fixed parameters are unchanged. The dimension tests on fibres for the analytic germs use their relatively open sets of regular points, whose algebraic closures have at least the indicated dimensions. This proves the lemma. □

13. A simultaneous active differential identity

Theorem 13.1. *On the original four Hopf factors, every rank-four activity over a compact slow base admits arbitrarily C^1-close activities, polynomial in the sphere coordinates with smooth slow coefficients, which satisfy the active parameter input of Section 11.4 simultaneously at the old ranks $s = m$ and $s = m - 1$. Every prescribed compact open paired comparison persists under a sufficiently small choice.*

Proof. We work at a fixed old rank s until the last step. Let $\mathcal{P}_s$ be the complex open set of matrices of Section 11.4, of dimension p_s. It parametrizes the whole tensor U, including the changes of its kernel rank. Each complex sphere is the affine variety

$$X_c = \{(z, \zeta) \in \mathbb{C}^h \times \mathbb{C}^h : \ \zeta^{\mathsf{t}} z = 1\}.$$

Its quotient by the action $(z, \zeta) \mapsto (tz, t^{-1}\zeta)$ is the variety of rank-one idempotents $z\zeta^{\mathsf{t}}$. This is an algebraically locally trivial $\mathbb{C}^*$-bundle: on the chart where a given diagonal entry is nonzero, we may normalize the corresponding component of z to one and recover both factors. The product of the four quotient varieties has dimension $t_{\mathrm{orb}} = 4(2h - 2)$. We include its parameters and the independent matrix U as slow algebraic marks. Let $\mathcal{A}_D$ be the affine space of all five-component ambient polynomial maps of degree at most D.

Interpolation. At any J distinct points of one complex orbit, the coefficients interpolate arbitrary jets of order d along the orbit, provided $D \geq (d+1)J - 1$. Indeed, choose an affine ambient linear function ℓ taking distinct values at the J points. For each point, the product of the differences of ℓ from its values at the other points, each raised to the power $d + 1$, is nonzero at the selected point and vanishes to order $d + 1$ at the others. Multiply it by the Taylor polynomial of the desired jet divided by this product. Since the orbit is immersed, ambient affine linear functions supply local coordinates for its jets, and the degree bound follows. In an immersive graph chart we keep the jets of the first four coordinates fixed and vary the fifth, so the same independence holds for the jets of the graph height.

The relative characteristic support. For each $U \in \mathcal{P}_s$ the bad normal family satisfies the hypotheses of Lemma 12.1. We apply its finitely many algebraic jet incidences with the matrix and orbit parameters held fixed, together with the interpolation above for $d = 5$. For every J this yields a closed exceptional set of coefficients, of codimension at least $J - p_s - t_{\mathrm{orb}}$, outside of which no nonzero conic coisotropic germ lies in the raw cone at an immersive point with quasifinite Gauss map, provided $D \geq 6J - 1$. Indeed, before the matrix and orbit marks are forgotten, the saving in dimension is J, and this holds on each of the finitely many fibre strata of that lemma. All jet

expressions are algebraic in local graph coordinates and in their invertible coordinate changes. Quasifiniteness means that the fibre of the Gauss map is zero-dimensional at the point; this is an algebraic open condition on the immersive source. Let Y° be this open subset of the source over $\mathcal{P}_s$ and the coefficients outside the exceptional set. Its projection to the parameter image is an open smooth morphism of relative dimension four. Each nonempty fibre is an open subset of the irreducible orbit torus, and is therefore analytically connected.

We now describe the intrinsic flux presentation. For an adjoint flux, push its four active components forward by the derivative of the activity and include the invariant density factor, that is, the Jacobian determinant of the graph coordinates. The resulting vector X takes values in the old space tensored with the five labels. The equations are

$$UX = 0, \qquad e_{\eta} X = 0, \qquad \operatorname{div} X = 0, \qquad \eta = (-\mathrm{d}f, 1).$$

Here e_{η} is contraction in the labels and the divergence differentiates the first four components in the graph coordinates. The first equation is the adjoint algebraic contraction, the second expresses tangency to the graph, and the third is the integration-by-parts identity for the active potential. To verify the latter, change variables in the pairing with the active differential of a test potential. Its derivative is the graph covector applied to the pushed-forward flux, and integration by parts in the four graph coordinates gives this divergence after absorption of the nonzero Jacobian density. Thus these equations represent the relative adjoint module, up to invertible changes of the flux and the adjunction of the graph tangency equation.

For this left relative differential module, the symbol equations at a nonzero cotangent vector ξ are

$$UX = e_{\eta} X = e_{\xi} X = 0,$$

with ξ extended by zero on the fifth component. Outside the raw cone the matrix (U, e_{η}, e_{ξ}) has full column rank, whatever the rank of U, so the characteristic support lies in the raw cone together with the zero section: the associated graded presentation contains the leading rows of these generators. We use the relative involutivity of the characteristic support [Wu22, Theorem 2.1]. At a general smooth point of a characteristic component not contained in the zero section, where its parameter projection has constant rank, the vertical tangent space is coisotropic. Indeed, a covector annihilating the vertical tangent space extends to an ambient covector annihilating the tangent space of the component. By involutivity, the relative Poisson tensor sends it into that tangent space, and its image is vertical. Thus the vertical smooth germ of the component in the cotangent directions is a conic coisotropic germ in the raw cone of the same orbit. This argument applies also to a component whose parameter image is a proper subvariety. Since the exceptional set of coefficients excludes all such germs, the characteristic support lies in the zero section. For a good order filtration, powers of the ideal of the fibre coordinates ξ then annihilate the associated graded module locally. Only finitely many graded degrees remain, which proves coherence over the ring of functions. The commuting relative derivatives give an integrable connection along each fibre.

Detection by compatibility rows. Fix U and write $K_U = \ker U$. For a general normal η, the projective directions of bad lines through η form a set of dimension at most one: the marked incidence has dimension at most five and the space of normals has dimension four, while nondominant components miss a general normal. Thus the bad affine direction cone has codimension at least two in the four-dimensional cotangent space, and (e_η, e_ξ) is injective on K_U for good ξ. For $X \in K_U \cap \ker e_\eta$ consider the forcing classes

$$[(e_\xi X, 0)] \quad \text{modulo} \quad \{(e_\eta Z, e_\xi Z) : \ Z \in K_U\}.$$

Suppose they vanish for every good ξ. Then the corresponding $Z = Z_\xi$ is unique and, by Cramer's rule on injective minors, holomorphic in ξ; it is homogeneous of degree one in ξ. Hartogs extension across the analytic subset of codimension two makes Z_ξ a linear function of ξ on the whole four-space. From $e_\eta Z_\xi = e_\xi X$ and $e_\xi Z_\xi = 0$, polarization gives an old-valued skew tensor in the labels whose η-column is X and whose horizontal columns are $-Z_\xi$, all in K_U. Since U is ordinary, this tensor vanishes, so $X = 0$. Thus at a general normal the forcing classes detect every nonzero tangent flux.

These symbolic tests come from exact differential compatibility rows of uniformly bounded order; we give the elimination. At a fixed parameter, eliminate the rows UX and split e_η on a meromorphic chart of constant rank. Write $X = B_1 x + B_2 g$, where the columns of B_1 span the tangent kernel $K_U \cap \ker e_\eta$ and g lies in the range of e_η. The divergence equation becomes

$$Px = h - Qg, \qquad P = \operatorname{div} \circ B_1, \qquad Q = \operatorname{div} \circ B_2.$$

The principal symbol $e_\xi B_1$ of P is generically injective, since $\ker e_\eta \cap \ker e_\xi \cap K_U = 0$ for good ξ. Select a square minor P_J of rows of P with invertible principal symbol. The ring of differential operators with meromorphic coefficients has no zero divisors, by multiplicativity of principal symbols. It admits nonzero common left multiples: for two nonzero operators of order at most b, choose multipliers of order a with

$$2\binom{a+4}{4} > \binom{a+b+4}{4},$$

for instance $a = 6b$. The linear system for the left coefficients making the two products agree then has a nonzero solution, and both multipliers are nonzero. This gives the division ring of left fractions. Gaussian elimination on P_J with nonzero principal pivots produces an inverse in this ring whose leading symbol is the commutative inverse, by the Schur complement formula at each step. The exact relations obtained by subtracting $P_{\mathrm{rem}} P_J^{-1}$ times the selected rows from the remaining rows annihilate P; we clear common left denominators. Together with the algebraic range constraints for e_η, the resulting compatibility rows on (g, h) have weighted principal symbols spanning the left cokernel of (e_η, e_ξ) on K_U, on a dense open set of directions. Indeed, the commutative row elimination gives a basis of the cokernel there, and each denominator multiplies its symbol by a nonzero scalar symbol. After substitution of $h - Qg$, the first block of each row has order at most one more than the divergence block. The eliminated rows UX can be reinserted to obtain identities in the original ambient flux. For fixed matrix dimensions there is a bounded number of fraction operations and denominator clearings,

so the argument for common multiples gives a bound on all orders independent of the polynomial degree of the activity.

Elimination of the coherent part. The support of the relative module, which is coherent over the functions, is saturated along the connected fibres. Indeed, its restriction to each fibre is a coherent module with integrable connection, hence locally free there by horizontal transport, and Nakayama's lemma identifies the points of the support with the points where this restriction is nonzero. Thus the parameter projection of the complement of the support is the open complement of the projected support. We stratify that algebraic support into finitely many smooth pieces over which the module, restricted to algebraic local sections of the orbit projection, is locally free. Such sections exist by the local triviality of the torus bundle, after shrinking around a point of a constant torus. Generic freeness and induction on closed complements give the stratification. On each nonzero stratum, flat relative transport over a small source ball gives a nonzero holomorphic family of solutions X_λ of the adjoint equations.

At a fixed parameter of a nonzero stratum, choose J' distinct source points at which the solution is nonzero and the normal is general, away from the meromorphic poles used in the elimination. This is possible because the quasifinite Gauss map is locally open and the nonzero solution is nonzero on a dense open subset of the source. At each of these points, forcing detection and the compatibility rows give an exact row whose first principal block detects $e_\xi X$ for some good direction ξ. Differentiate the equations along a tangent vector to the parameter stratum and apply this exact row. It annihilates the term involving the differentiated solution, and so gives a linear scalar constraint on the tangent vector. To see that these constraints are independent, vary only the jets of the graph height at the chosen points: keep all derivatives through the order b of the first block equal to zero, and use the homogeneous jet $\xi^{\otimes(b+1)}$ at the next order. In the differentiated equations the divergence row is unchanged in graph coordinates, and U is held fixed. The product rule in the first block leaves exactly the nonzero symbol applied to $e_\xi X$, while the derivatives of the solution and the lower derivatives of the graph drop out of this test. Independent interpolation of the ambient coefficients gives J' independent tangent constraints. The required order of interpolation is bounded before D is chosen, by the fraction elimination above. Hence the projected exceptional set of coefficients for these nonzero strata has codimension at least $J' - p_s - t_{\mathrm{orb}}$, and this remains true for its algebraic closure.

Choice of the activity. Finite complex Gauss germs at all real points can be arranged by a preparation of finite jets. For a centred scalar height jet, the loci where the Jacobian algebra modulo the $(j+1)$-st power of the maximal ideal has dimension at least $j+1$ are closed determinantal loci in the spaces of $(j+1)$-jets. They are nested under inverse images, because each nonzero successive quotient of powers of the maximal ideal has positive dimension, while a zero successive quotient terminates the filtration by Nakayama's lemma. Their codimensions tend to infinity. Otherwise the components of maximal dimension would eventually be full jet cylinders over components at the preceding level; since there are finitely many components at the starting level, one such cylinder would persist at all higher orders. However, every fixed Taylor jet extends

to a polynomial with an isolated critical point by adding $\sum_i x_i^q$ in strictly higher degree, because the homogeneous gradient terms have no common projective zero. This contradiction gives arbitrarily large codimension. For a real nonzero height covector, varying the activity in that scalar height direction surjects onto the scalar jets. Finite-jet transversality, with the projective height marks allowed to vary, therefore gives isolated critical Gauss germs at every real point. By compactness we obtain a fixed jet order and an open margin. We approximate in the corresponding norm by ambient polynomials with smooth slow coefficients, preserving the prescribed C^1 comparison. This gives quasifiniteness where it is needed on the complexified orbits.

Choose J and J' so that both codimension bounds exceed the dimension of the slow base, at both old ranks, and choose a single degree D large enough for both interpolations and for the approximation. By real stratified transversality, the coefficient map of the compact slow base can be moved by an arbitrarily small constant real translation of the coefficients so as to avoid the finite union of the closed exceptional sets for the two ranks. The complex conjugates may be added to this finite union, and the real traces have real codimension at least the stated complex codimensions. All rank-four and finite-Gauss conditions persist, as do the compact paired comparisons. On the joint good open sets of parameters and orbits there is now no support of the coherent module at either rank. Hence the order-zero flux rows are finite left differential combinations of the intrinsic adjoint equations on algebraic charts. Since the good order filtration and all coefficient incidences are algebraic, these identities hold on those charts before the smooth slow coefficients are substituted. They are the required active parameter identities at both ranks. □

14. Relative filling of the hyperplane metric

For a distribution D of rank d we write

$$j_d = d + \frac{d(d+1)}{2} = \frac{d(d+3)}{2}$$

for the rank of the leafwise second-jet bundle $J_2(D) = D \oplus \operatorname{Sym}^2 D$. A map, or a map defined on an open subset, is called *D-free* if its first and symmetric second leaf derivatives (its first and second leaf columns) define a monomorphism $J_2(D) \to \mathbb{R}^N$ at every point.

Proposition 14.1. *Let $\pi\colon X \to B_0$ be a compact smooth submersion, where B_0 may have boundary and $\partial X = \pi^{-1}(\partial B_0)$. Let $D \subset \ker \mathrm{d}\pi$ be an integrable vertical distribution of rank d with a smooth positive metric h, and let $N \geq j_d + 5$. Let $K \subset B_0$ be a closed subset containing ∂B_0. Suppose that F_{in} is a D-free h-isometric map into $\mathbb{R}^N$, defined on $\pi^{-1}(U)$ for an open neighbourhood U of K, and that its columns $J_2(D)F_{\mathrm{in}}$ extend to a formal monomorphism $\Phi_0\colon J_2(D) \to X \times \mathbb{R}^N$ over all of X. Then there is a smooth D-free map $F_{\mathrm{out}}\colon X \to \mathbb{R}^N$ with $\gamma_D(F_{\mathrm{out}}) = h$, equal to F_{in} on $\pi^{-1}(U')$ for a smaller neighbourhood U' of K, such that $J_2(D)F_{\mathrm{out}}$ is homotopic to Φ_0 through monomorphisms relative to $\pi^{-1}(U')$.*

Proof. Step 1: a free map in the given class. Using a connection and a complement of D, extend Φ_0 to a formal second jet over X which equals the jet of F_{in} near $\pi^{-1}(K)$. The open D-free relation and its first-jet projection are ample. Indeed, let τ be a principal covector. If $\tau|_D \neq 0$, choose a leaf basis $e_1, \dots, e_d$ with $e_1, \dots, e_{d-1} \in \ker \tau|_D$. The principal slice varies only the last pure second column, the other $j_d - 1$ columns being fixed and linearly independent, and freeness excludes their linear span, which has codimension $N - j_d + 1 \geq 6$. If $\tau|_D = 0$, the relation is constant on the slice. The first-jet projection is the partial immersion relation, and the same slice argument with the d first columns applies; the excluded span has codimension $N - d + 1 \geq 2$. Thus every slice is open and connected with full convex hull. The relative higher-order ample theorem [MAdP21, Theorem 3.5] gives a smooth D-free map $F\colon X \to \mathbb{R}^N$, equal to F_{in} near $\pi^{-1}(K)$, together with a homotopy of monomorphisms from Φ_0 to $J_2(D)F$ relative to a neighbourhood of $\pi^{-1}(K)$.

Step 2: supported fourth-power primitives. We describe the device used to adjust the metric, recording the local analytic results it relies on. Work in a contractible foliation box $\Omega \times Q$, with Ω a leaf ball, Q a compact transverse parameter ball, D spanned by the leaf coordinate fields and π constant on the plaques. A smooth compactly supported amplitude $a \geq 0$ and a linear leaf coordinate ϕ determine the increment $a^4(\mathrm{d}_D\phi)^2$. We use the free-map primitive construction [Z16, Sections 4.1–4.3, Theorem 4.1] and the supported perturbation theorem [Z17, Theorem 2.12]. In leaf coordinates and in target dimension at least $j_d + 5$, the first construction gives supported free maps whose leaf metric approximates the desired increment to arbitrarily high order in the wavelength ε, measured in leaf $C^{2,\alpha}$. More precisely, the finite correction through order r has metric error $O(\varepsilon^{r-2-\alpha})$ in leaf $C^{2,\alpha}$. Its determinant estimate uses normalized free columns and gives a lower bound independent of ε for fixed r, while the leaf derivatives of the columns of each fixed order are bounded by powers of ε^{-1}. Consequently, by the Gram inverse formula and the product rule, the free pairing inverse has $C^{2,\alpha}$ norm $O(\varepsilon^{-b})$ for some b independent of r.

These local results hold with the transverse parameters of the box. To obtain the relative and homotopy assertions, choose five orthonormal sections of the normal bundle $(J_2(D)G)^{\perp}$ of the current free map G over a buffered product box, by orthogonal transport along radial segments. The finite primitive formula uses only these sections, linear pairings, and zero-mean integration in an auxiliary periodic variable θ. For a fixed finite order, the normalized column determinants and all cutoff and inverse estimates are uniform on the compact parameter box. At a zero of the amplitude, the only apparent quotient by a^2 occurs in the first step and has numerator $|\partial_\theta(aR)|^2 = a^2|\partial_\theta R|^2$, where R is the fixed periodic curve; it is therefore extended as $|\partial_\theta R|^2$. The later pairing systems divide by a only on the large-amplitude side of the smooth partition in a^2 used in that construction. Inductively from the initial oscillator a^2R, all displacement terms vanish at a transverse parameter where a vanishes identically on the plaque.

Choose one outer leaf cutoff equal to one on the support of the remainder, and apply the supported perturbation theorem with this cutoff throughout the compact transverse family. If e is the prescribed metric remainder, the smallness required for the contraction

is the product of the $C^{2,\alpha}$ norm of the pairing inverse with the $C^{2,\alpha}$ norm of its value on $(0, e)$, which is $O(\varepsilon^{r-2-\alpha-2b})$; the displacement is $O(\varepsilon^{r-2-\alpha-b})$. We first fix the finite order r so that, for small ε, this displacement also lies within the freeness tolerance given by the normalized column bounds, and then fix one small ε, uniformly on the compact box. The contraction uses the fixed leaf Dirichlet inverse and has a common contraction constant less than one. Its solution depends smoothly on the parameters, by the implicit function theorem in leaf $C^{2,\alpha}$ and the higher leaf regularity of the same supported equation [Z17, Sections 3.1–3.3]. For mixed derivatives, a derivative in a further parameter satisfies an equation with the same contractive linearization acting on the highest parameter derivative; the higher leaf estimates absorb the highest leaf derivative in that equation and bound the lower derivatives inductively, and the fixed leaf Dirichlet inverse commutes with parameter derivatives. This gives joint smoothness for one fixed choice of ε. On the leaf and transverse boundary buffers, where the amplitude vanishes on whole plaques, the zero correction is the unique small solution; hence the map is unchanged off the buffered chart. Adding a compact parameter $t \in [0, 1]$ with amplitude ta gives a smooth path of D-free maps from the given free map to the one realizing the increment.

Step 3: shortening and decomposition of the defect. Before applying this device, shorten F by a positive basic function λ (that is, a function pulled back from B_0) equal to one near a smaller closed neighbourhood of $\pi^{-1}(K)$. Take $\lambda^2 = 1 - \beta^4$, where the basic function $\beta \geq 0$ vanishes near that set and equals a positive constant outside a region on which F is already h-isometric. Choose the constant so that $\lambda^2 \gamma_D(F) < h/2$ where β is constant. Then

$$h - \gamma_D(\lambda F) = \beta^4 \Delta$$

with a smooth positive definite Δ, equal to h wherever $\beta = 0$. Since $\mathrm{d}_D \lambda = 0$, multiplication by λ preserves freeness and the relative homotopy class of the columns.

In each small foliation chart, decompose Δ as a sum of smooth strictly positive coefficients times constant covector squares. For instance, take a Δ-orthonormal leaf coframe $e^1, \dots, e^d$ at the centre of the chart and the squares of the e^i and of $e^i \pm e^j$, $i < j$. This finite family spans the symmetric tensors and, because $(e^i + e^j)^2 + (e^i - e^j)^2 = 2(e^i)^2 + 2(e^j)^2$, represents the centre value $\sum_i (e^i)^2$ with strictly positive weights; a fixed linear right inverse keeps the weights positive nearby. Multiply their positive fourth roots by a fourth-power partition of unity and by β. This writes the defect as a finite sum

$$h - \gamma_D(\lambda F) = \sum_{\nu=1}^{M} a_\nu^4 (\mathrm{d}_D \phi_\nu)^2$$

of supported fourth-power increments, whose amplitudes vanish on a common neighbourhood of $\pi^{-1}(K)$ and, since β is basic, on every plaque meeting it. Apply Step 2 successively, each time to the current free map. The metric increments add exactly, and the concatenated relative paths of free maps give the required homotopy class. □

Corollary 14.2. *Over the normal disk bundle of the three-dimensional core, suppose that the H-isometric H-free collar map and its surviving unit line admit a relative*

formal monomorphism extension of $J_2(H) \oplus L$ into $\mathbb{R}^{s_n+5}$, pulled back from the core on the zero section. Then there is a relative H-isometric H-free filling into $\mathbb{R}^{s_n+5}$, together with a unit augmented line perpendicular to the first H-columns, agreeing with the given one on a smaller collar and independent of the $J_2(H)$-columns, and two orthonormal copies of L perpendicular to the span of the $J_2(H)$-columns and the augmented line.

Proof. Here $j_{n-1} = (n-1)+\frac{1}{2}(n-1)n = s_n - 1$, so Proposition 14.1 applies with six free normal dimensions. First make the formal L-column perpendicular to the formal first H-columns by subtracting its projection and rescaling positively; its component normal to the formal $J_2(H)$-columns is nonzero. Over the zero section the formal orthogonal complement of $J_2(H) \oplus L$ has rank $(s_n + 5) - s_n = 5$. The maps $2L \to \mathbb{R}^5$ of rank less than two form a subset of codimension four, which exceeds the dimension three of the core; rank avoidance and polar normalization therefore give two orthonormal copies of L in that complement. Radial contraction of the disk and orthogonal transport of the formal complement extend this choice over the whole product, including the fast variables.

Along the relative homotopy of $J_2(H)$-monomorphisms provided by Proposition 14.1, transport the initial normal component of the L-column and the two chosen copies by orthogonal complement transport. Write the initial tangential component as a combination of the $J_2(H)$-columns, using their metric left inverse. Evaluate the same combination on the final columns and subtract its projection to the first H-columns. Add the transported normal component, which is nowhere zero, and normalize by the length. This gives the unit augmented line; it is independent of the $J_2(H)$-columns because its normal component is nonzero. The two transported copies are orthogonal to the $J_2(H)$-columns and to the transported normal component, hence to the augmented line. The support and collar conditions follow from the relative homotopy, which is stationary on the collar. □

15. A single set of coefficients on the doubled exterior

When $n \equiv 2 \pmod 4$ we use the exceptional angular construction on every core component. In that residue the description $N_5 \simeq 3\mathcal{D} \oplus \mathbb{R}^2$ of Proposition 6.2 holds whatever the value of the orientation cube $w_1(\mathcal{D})^3$. In the other dimensions Proposition 6.2 gives $w_3(N_5) = 0$, and we use the split-normal construction on every component. Thus in each case the chosen boundary model is available on every component. There are finitely many components, with pairwise disjoint buffered collars.

We now describe the matching and state the boundary test precisely. The reference is the product of Section 7. On an exceptional collar we use $\mathcal{H}_v$ and its cofactor C_0. The adjustment of the retained part to the class J_Σ is the orthogonal comparison of Section 8. It can be made on the slow seed, since it conjugates the selfadjoint operator of the seed and postcomposes its injector. We carry out the normalized preparations and the two-curve switch of Proposition 9.3, retaining all derivatives along the path constructed in Section efsec:angular. The residual rows are normalized before the switch as in Section 8.4. The open tests concern the kernel of the fixed cofactor, so the final positive row scaling preserves them and their augmented class.

On a split-normal collar we use Section 10.3 and the prepared ordinary moment of Section 10.2. The slow bundle $\mathcal{E}$ is identified with $Q \oplus K^{\perp}$ through the original angular map. On $E_v \otimes A$, matching its normal N_5 to the normal of J_Σ gives the same adjustment of the retained part. Namely, match the two isometric injections into $W \oplus \underline{\mathbb{R}}^4$, including their normal bundles, and then homotope the four cofactor frame columns of the resulting isometry back to their original inclusion; the rank w of the complement allows this homotopy over the link cylinder. Orthogonal transport gives an endpoint isometry acting on W. We apply it to the seed and to its short flat path, on all old source columns. This gauge preserves the algebraic properties of the seed and all kernel comparisons, and the moment remains the fixed ordinary polynomial system, with the slow label chart g close to the identity.

In this split-normal case the short path in the ordinary part of the flat equation provides the transition to the prepared moment. Pull it back radially, keep the final coefficient on an inner band, and then use a smooth radial path parameter decreasing to its seed value. The seed remains valid on a larger product buffer. Extend the selfadjoint and injector parameters as sections over the double, and extend g by cutting off its difference from the identity in the larger collar, which preserves invertibility. Along the path the full and reduced weak-four conditions persist by smallness, and the path stays in the neighbourhood, used in the proof of Proposition 9.3, in which every star failure has a regular tuple. On the fixed inner band the full and reduced star exclusions already hold. Hence the unrestricted flat avoidance of that proposition can be applied relative to a smaller band, also on the interpolated part of the collar: the estimates there are the same regular-tuple codimension counts for the smooth flat relation, now with the angular source identifications fixed. On the rest of the double we first prepare the ordinary moment parameters and then apply the same flat avoidance, relative to the product buffer. These relative constructions give one smooth flat tensor on the double which satisfies the full-plane and reduced-plane tests and the prescribed small comparisons on the core collars. In either residue, the core and exterior perturbations are performed with this single target W and source $\mathcal{E}$.

15.1. **The normal bank and the boundary condition.** For $b \in A$, let $C_A(x)$ be the span of the four projected active columns of Section 11.6. For $l \in L$ put

$$M_x(l \otimes b) = \pi_{C_A(x)} D^2 B(l, b), \qquad J_x(b) = \pi_{C_A(x)} D_b^x I_x.$$

Let Z_x be the unique vector in $C_A(x)$ perpendicular to $\operatorname{im} M_x$ whose inner product with the projected LL-column is one. It exists, and depends smoothly on x, because the first and second active columns have full rank. The isomorphism $M_x \oplus Z_x$ from the four slow axes to $C_A(x)$ is homotopic to the isomorphism given directly by the four projected columns: subtracting the LA-component of the LL-column and changing its positive length are homotopies through isomorphisms.

On the boundary band we apply the placement of the whole bank from Section 8.4. Its Gram completion works on all of $\mathcal{E}$, since $G(\cdot, z)$ is a monomorphism there by the open rank-one exclusion and the rank gap of the old normal exceeds the dimension of the link. It sends the bank projection of $\Lambda G(\cdot, z)$ to the assigned row in the rank-nine

bank, extended by zero where necessary; the positive constant Λ may be taken as large as needed. The corresponding orthogonal transport of the target can be cut off in the larger collar. Equivalently, one applies the same transport and scalar to the seed and to the flat path before the relative selection on the exterior. The original class comparison is preserved, by the isotopy of this bank placement and the positive scaling.

In the exceptional case we assign the bank row of C_0 divided by κ. Then at z the value $J_x(b)G(h,z)$ is the prescribed cofactor, in the invertible four-row coordinates of Section 8.6. In the split-normal case we assign only the $L\otimes U^*$-part of the bank, namely $\kappa^{-1}\sigma(\pi_D h,\cdot)\in A^*$, and zero on the other old directions. The cubic of Section 8.6 then gives

$$J_x(b)G(h,z)=\sigma(\pi_D h,b)\,Z_x.$$

Indeed, the left-hand side has inner product zero with each of the three projected LA-columns and inner product $\sigma(\pi_D h,b)$ with the LL-column. This uses the derivative of I_x on the bank constructed in Section 8.6. The other normal component of $J_x(b)G(h,z)$ is a bounded correction at this distinguished value, since the bank projection there is fixed independently of Λ.

On a split-normal collar we first take an old generator u with monomorphic derivative $d_E u\colon E_v\to Q$, homotopic through monomorphisms to the old inclusion. The open ample argument of Section 11.2 applies with values in Q, because the missing dimension is one and the effective bad slice has codimension two. Write $u_D=\pi_D u$ and $u_L=\pi_L u$. On this collar we use the generator $u_D+\delta_0 u_L$ with a small constant $\delta_0>0$, and the passive profile

$$f=-\lambda\,u_L\otimes z^*,\qquad z^*\in V^*,\quad z^*(z)=1,$$

which takes values in the original active line L. At the comparison point $\delta_0=0$, $a=z$, and for any derivative T on the path of active derivatives, the paired mixed map on $Q\otimes A$, up to the common positive factor λ, is

$$h\otimes b\longmapsto\Big(G(\pi_D h,Tb),\ \sigma(\pi_D h,b)\,Z_x+M_x\big((\pi_L h)\otimes b\big)\Big).$$

It is injective by the path of active derivatives and the three LA-rows. This comparison at $\delta_0=0$ concerns columns only. We choose $\delta_0>0$ sufficiently small after the compact comparison and the bank placement have been fixed. The generator then has a monomorphic old derivative $E_v\to Q$, and by compactness the paired map keeps its rank and class, also along the derivative path. Restriction along the old derivative $E_v\to Q$ preserves rank and is homotopic to restriction along the original inclusion. In this way we avoid having to realize an old generator whose line component vanishes identically.

On exceptional collars we use instead an old generator into $\mathcal{H}_v$, as in Section 11.2, and set the passive profile identically zero. Both generators take values in E_v plus their single generating line in $\mathcal{E}$, and all reduced tests hold there. The classes of these derivatives are preserved by these homotopies of monomorphic old derivatives. After the compact tests have been fixed, we shrink the active values toward z by normalized powers and the inverse scaling of values. The paired comparisons for all independent samples of the points and of the coefficients persist at the endpoints.

We then use Section 11.2 to extend and prepare a single old generator on the double. The required rank of its projection to E holds on each link, since both generators take values in E_v plus one line, and it holds along the required monomorphic formal homotopies to the inclusion. The slow contraction of the auxiliary coordinates away from these bands and the small finite-contact preparations preserve every required reduced test. On the band the changes of coefficients can be taken arbitrarily small in C^1, so that the four cofactor comparisons are retained. The passive profiles may be kept fixed during these small changes, since they already satisfy all the compact open comparisons at the initial generators. We extend them off the split-normal collars by slow cutoffs in their original L-slots and by zero elsewhere. They have the same fixed rank-one label reader z^* everywhere, so $f^* d_E f = \lambda^2 \langle u_L, d_E u_L \rangle\, z \otimes z^*$ is symmetric and $\operatorname{skew}(f^* d_E f) = 0$. Choose an independent slow reference frame v_0 on the double by rank avoidance, and one sufficiently small $\lambda > 0$. Then $v = v_0 + \lambda G(u, \cdot)$ is an independent old frame, and $\operatorname{skew}\big(v_0^*(v - v_0)\big)$ is its exact old skew primitive, valid for every basic passive weight.

Extend the active maps relative to the collars over the double, using their formal rank-four paths and the open ample argument of Section 11.2. Choose the simultaneous activity of Theorem 13.1 so close on this compact double that all the boundary comparisons on the boundary band persist. The finite old contact and Lemma 11.1 give the differential right inverse $(*)$ on the same product, and the reduced tests give its compact tilted phase margin. On the boundary band, with

$$U(h, b') = (d_h^y v) b', \qquad S(h, b') = (d_h^y f) b',$$

the paired map on $E_v \otimes A$

$$h \otimes b \longmapsto \Big(U(h, d_b^y a),\ J_x(b) U(h, a) - M_x\big(S(h, a) \otimes b\big) \Big) \tag{P}$$

is uniformly injective. At $a \approx z$ one has $-M_x(S(h, a) \otimes b) \approx \lambda M_x((d_h^y u_L) \otimes b)$, which carries the same factor λ as $U(h, d_b^y a) = \lambda G(d_h^y u, d_b^y a)$. After adjoining the old columns and the ten active columns, the class of (P) agrees with J_Σ. The argument gives this class while the independent old and active samples are still separate, so it is preserved by the subsequent power substitution.

16. Completing the surviving line

We first construct the periodic family used to complete the metric along the surviving line. The construction is intrinsic for a metric real line bundle L, including a nonorientable one: all formulas are written in a local unit ℓ of L and are checked to be independent of the choice of ℓ.

Throughout this section, for a function f on a sphere bundle with Hopf flow ρ_t we write

$$\overline{f}(z) = \int_0^1 f(\rho_t z)\, \mathrm{d}t$$

for its orbit mean, and J denotes a fixed quarter turn of $\mathbb{R}^2$.

Lemma 16.1. *Let L be a metric line bundle over a compact slow base, and let $d_* \geq 1$ and $b_* \geq 2$ be integers. In a local unit of L, let X denote the period-one Hopf generator on the unit sphere bundle $S(\mathbb{R}^{b_*} \oplus L^{b_*})$. There is a smooth family $a(q, p, \theta)$ of sections of $2L = L \oplus L$ over this sphere bundle, defined for q in a small ball of $\mathbb{R}^{d_*}$, p in a small ball of L^{d_*}, and θ in a small interval of L, with the following properties.*

(i) *Each member of the family has unit speed $|Xa| = 1$, zero orbit mean $\overline{a} = 0$, and zero signed orbit area $\overline{\langle a, JXa\rangle} = 0$.*

(ii) *The orbit-mean parameter one-form of the family is*

$$\overline{\langle \mathrm{d}a, Xa\rangle} = \mathrm{d}\theta - \sum_{j=1}^{d_*} p_j \,\mathrm{d}q_j,$$

where the differentials are taken in the family parameters only.

The form in (ii) has the parity of L. The family and its parameter derivatives are smooth also in the sphere and slow variables.

Proof. In a local unit of L write the sphere coordinates as $z = x + \mathrm{i}y \in \mathbb{C}^{b_*}$, so that the Hopf flow is $\rho_t(z) = e^{2\pi \mathrm{i} t} z$. Changing the unit conjugates z and reverses the sign of X. The antipodal map $\sigma(z) = -z = \rho_{1/2}(z)$ is a globally defined free involution. On the quotient by σ the Hopf tangent line is the pullback of L; let γ be the sign line of the double cover. We first find a section of $L \oplus (L \otimes \gamma)$ over the quotient whose derivative along the Hopf line vanishes nowhere. A formal solution is given by the derivative that maps the Hopf line isometrically onto the first summand, and the relation is open and ample: each principal slice is the complement of one point in a plane. By the relative parametric h-principle for open ample relations [EM02, Theorem 18.4.1] such a section exists. Lifted to the sphere bundle, it becomes a $2L$-valued section a satisfying $a(\sigma z) = \mathcal{R} a(z)$ with $\mathcal{R} = \mathrm{diag}(1, -1)$.

For a function f with $\overline{f} = 0$ define

$$\mathcal{I}_X f(z) = \int_0^1 t\, f(\rho_t z)\, \mathrm{d}t.$$

Integration by parts gives $X\mathcal{I}_X f = f$, and Fubini's theorem gives $\overline{\mathcal{I}_X f} = 0$. Reversing the unit of L reverses this primitive. Let $v_0 = |Xa| > 0$ be the speed of the section and put $\lambda_0 = \overline{v_0}$. The orbit map with variable time $\mathcal{I}_X(v_0/\lambda_0 - 1)$ has positive derivative v_0/λ_0 along the orbits and degree one on every circle, so it is a diffeomorphism of each orbit. Composing the section with its inverse and dividing by λ_0 makes the speed equal to one. Since v_0 is σ-invariant, this operation commutes with the reflection symmetry. Finally we subtract the orbit mean. The resulting section is centered, and its signed area has mean zero, because $\mathcal{R}^T J \mathcal{R} = -J$ implies that the area integrand $\langle a, JXa\rangle$ changes sign under $\sigma = \rho_{1/2}$.

Write $v = Xa$ and $j = Jv$ in a unit of L. Since a and X both change sign under unit reversal, v and j are unchanged, and $Xv = \kappa j$ for an L-valued curvature κ. The mean curvature $\overline{\kappa}$ vanishes: the sphere fibre has dimension $2b_* - 1 \geq 3$ and is simply connected, so the unit tangent map v sends each Hopf circle to a null-homotopic loop

in S^1. Moreover $b_2 = \overline{\kappa^2} > 0$, since zero curvature on an orbit would give a constant unit tangent v, contradicting $\overline{v} = 0$. Both statements hold on every orbit.

We next produce enough unit-speed tangent pairs for the parameter family. On every orbit $\max v_1 > 0$, since otherwise $\overline{v} = 0$ and $|v| = 1$ would be incompatible. Choose $\eta_0 > 0$ such that this maximum exceeds $2\eta_0$ on every orbit, and a smooth function $\eta(v_1) \geq 0$ supported in $\{v_1 > \eta_0\}$ and positive on $\{v_1 > 2\eta_0\}$. Then

$$b_0 = \overline{a_2^2 \eta^2 v_1} > 0.$$

Indeed, if $a_2 = 0$ throughout $\{v_1 > 2\eta_0\}$, then $v_2 = Xa_2 = 0$ and hence $v_1 = 1$ there. This set would then be a nonempty open and closed subset of the orbit circle, contradicting $\overline{v_1} = 0$.

Choose distinct positive integers $c_1, \dots, c_{d_*}$ and, for a large positive integer K_*, set

$$h_{q,i} = a_2 \eta \cos(K_* c_i a_1), \qquad h_{p,i} = a_2 \eta \sin(K_* c_i a_1).$$

For either type of input h put

$$\beta_h = h - \overline{h} - \kappa \overline{\kappa h}/b_2, \qquad \alpha_h = \mathcal{I}_X(\kappa \beta_h), \qquad u_h = \alpha_h v + \beta_h j - \overline{\alpha_h v + \beta_h j}.$$

The primitive is defined since $\overline{\kappa \beta_h} = 0$; also $\overline{\beta_h} = 0$. From $X u_h = (\kappa \alpha_h + X \beta_h) j$ we get $\langle v, X u_h \rangle = 0$, and $\overline{\langle u_h, v \rangle} = \overline{\alpha_h} = 0$. The q-tangents have the parity of $2L$, and the p-tangents are ordinary two-vectors. Both have the same reflection symmetry as a: under σ, v transforms by $\mathcal{R}$ and j by $-\mathcal{R}$, while κ, β_h are odd and α_h is even.

The skew pairing of these tangent pairs is nondegenerate for K_* large. On the support of the oscillations $Xa_1 = v_1 > \eta_0$, so repeated integration by parts against $(\mathrm{i} K_* c\, v_1)^{-1} X$ shows that every mean of nonzero frequency tends to zero faster than any inverse power of K_*. One such antiderivative, together with a bounded zero-mean primitive of its remainder, gives $\alpha_h = O(K_*^{-1})$. Consequently

$$\overline{\langle u_{q,i}, X u_{p,i} \rangle} = \tfrac{1}{2} K_* c_i b_0 + o(K_*).$$

Indeed, the term of order K_* is $\beta_{q,i} X \beta_{p,i}$, whose nonoscillatory part is $\frac{1}{2} K_* c_i a_2^2 \eta^2 v_1$. All pairs of distinct frequencies, and all pairs of the same type, have zero mean at order K_*. These estimates hold uniformly on the compact orbit quotient. For the parameter form $\lambda_a = \overline{\langle \mathrm{d}a, Xa \rangle}$, the exterior derivative on two variations u, u' is $-2\overline{\langle u, Xu' \rangle}$. After rescaling each pair by $(K_* c_i b_0)^{-1/2}$ and reversing the sign of the p-member, the resulting form is close to

$$\omega_0 = \sum_{i=1}^{d_*} \mathrm{d}q_i \wedge \mathrm{d}p_i.$$

A fibrewise linear change makes it equal to ω_0 at the origin: interpolate linearly between the two nondegenerate forms, ω_t, and use the linear transport generated by $-\frac{1}{2} \omega_t^{-1} \dot{\omega}_t$.

Form the affine family with these tangents, and normalize its speed and center it exactly as above. At the origin the tangents are unchanged by this normalization, since their linearized speed and mean already vanish. The reflection symmetry gives zero area for all nearby parameters. Including the time shift θ along each Hopf orbit, the parameter form becomes $\mathrm{d}\theta + \lambda(q, p)$ with $\lambda(0) = 0$ and $\mathrm{d}\lambda(0) = \omega_0$. On a common smaller ball, the radial homotopy on the parameter fibres writes $\mathrm{d}\lambda - \omega_0 = \mathrm{d}\mu$ with

$\mu = O(|(q,p)|^2)$. The vector field Z_t defined by $\iota_{Z_t}(\omega_0 + t\,\mathrm{d}\mu) = -\mu$ vanishes to second order at the origin, so it has a smooth time-one flow on a smaller ball, and pullback by this flow converts $\mathrm{d}\lambda$ into ω_0. Now $\lambda + \sum_i p_i\,\mathrm{d}q_i$ is closed, hence exact by the radial homotopy, and changing θ by its negative primitive gives the required formula. All orbit means and all these parameter changes are smooth over the orbit quotient. The forms and times have the parity of L, so their defining vector fields and flows agree under unit reversal. This proves the lemma. □

16.1. **Line completion on a compact model.** Let B be a closed manifold and let V_0 be a rank-n metric bundle on B with an injective coefficient map $A_0\colon V_0 \to \underline{\mathbb{R}}^a$. On an open tube $U \subset B$ fix an orthogonal splitting $V_0|_U = H \oplus L$. The old host is $B \times (S^{2h-1})^a$, the product of B with the original a Hopf spheres. Denote by Y_0 the old coefficient field of a local unit ℓ of L, and by $Y_1, \dots, Y_{n-1}$ those of a local slow orthonormal frame of H. Fix smooth basic functions $c, \chi \geq 0$ with $c^2 + \chi^2 = 1$ and $\operatorname{supp}\chi$ compact in U, and fix a slow open neighbourhood of $\operatorname{supp}\chi$ with closure in U.

Proposition 16.2. *Let $N \geq s_n + 1$. Suppose there are an old map F into $\mathbb{R}^N$ and an L-valued function s on the old host whose slow support is compact in U, with $|\mathrm{d}_{V_0} s| < 1/10$. Suppose that on a neighbourhood of $\operatorname{supp}\chi$ the map F is H-free, and there is a line monomorphism $e\colon L \to \mathbb{R}^N$, independent of the columns $J_2(H)F$, satisfying $cY_0F = c\,e(\ell)$. Suppose that on the same neighbourhood there is an isometric inclusion $\nu\colon 2L \to \mathbb{R}^N$, slow or fast varying, orthogonal to $J_2(H)F \oplus e$. In a local unit of L write $\alpha = \mathrm{d}_H s$ and $p = Y_0 s$. Assume that the old metric has the blocks*

$$\text{(L)} \qquad \gamma_F|_{HH} = 1_H + c^2\alpha^2, \qquad \gamma_F|_{HL} = (1 + c^2 p)\alpha, \qquad \gamma_F|_{LL} = 1 + 2p + c^2p^2$$

where $c > 0$ in U, that $\gamma_F|_{HH} = 1_H$ where $c = 0$, and that outside U the old metric is the specified unit metric on V_0. Finally, suppose that F is phase regular on V_0 outside a relatively compact neighbourhood of $\operatorname{supp}\chi$ on which e and ν are defined.

Fix a slow isometric inclusion $\iota_L\colon L \to \underline{\mathbb{R}}^6$ on U and a basic cutoff β, equal to one near $\operatorname{supp} s \cup \operatorname{supp}\chi$ and supported in U. Define the coefficient map $A_1\colon V_0 \to \underline{\mathbb{R}}^{a+6}$ by

$$A_1(v_H + v_L) = \big(A_0(v_H + cv_L),\ \beta\,\iota_L(v_L)\big)$$

on U, and by $A_1 = (A_0, 0)$ outside U. Then on the old product with six further ordinary Hopf factors of the same sphere size there is a smooth map into the same $\mathbb{R}^N$ which induces the unit metric of V_0 along the distribution of A_1 and is phase regular along it.

Proof. All estimates in this proof are for one fixed original sphere size h. Choose two nested slow buffers around $\operatorname{supp}\chi$ on which e and ν are defined; they have compact closures in the region of definition. Choose finitely many smooth functions $q = (q_1, \dots, q_{d_*})$ of the old variables whose differential $\mathrm{d}_H q$ is injective on the larger buffer. For example, use cutoffs times the coordinates of a finite atlas of the compact old host, together with the cutoffs themselves. Scale q into the small parameter ball of Lemma 16.1. For the moment we work on the auxiliary twisted sphere bundle of that lemma; its size b_* will be fixed below, independently of all frequencies. In a local unit of L its generator X

commutes with the old fields. Put

$$a = a(q, 0, 0), \qquad v = Xa, \qquad Z = \nu a.$$

Then

$$\overline{Z} = 0, \qquad |XZ| = 1, \qquad \overline{\langle YZ, XZ\rangle} = 0$$

for every old field Y, including the field of the old line. To see the last identity, write $YZ = (Y\nu)a + \nu \sum_j \partial_{q_j} a\,(Yq_j)$. Since $\nu^*\nu = \mathrm{Id}$, the connection matrix $\nu^* Y \nu$ is skew, hence a multiple $\omega_Y J$ of the quarter turn, and its contribution is $-\omega_Y \overline{\langle a, JXa\rangle} = 0$ by the zero signed area. The remaining terms vanish because the parameter form of Lemma 16.1 restricted to $p = \theta = 0$ is zero.

Periodic coefficients. We use the formal independent new derivative $cY_0 + \varepsilon^{-1}X$ along the line and consider the formal series

$$F + \sum_{r\geq 1} \varepsilon^r F_r, \qquad F_1 = \chi Z.$$

The leading line column is

$$w = c\, e(\ell) + \chi\, \nu v,$$

where the local unit is understood in the last term. Its leading pairings with the old columns, together with its squared length, give the metric with blocks

$$k_{HH} = 1_H + c^2\alpha^2, \qquad k_{HL} = c(1 + c^2 p)\alpha, \qquad k_{LL} = (1 + c^2 p)^2. \tag{K}$$

Where $c = 0$ this uses only the H-metric and the unit speed. Elsewhere it follows from (L) and the orthogonality of the normal plane; for instance $|w|^2 = c^2(1+2p+c^2p^2)+\chi^2 = (1 + c^2 p)^2$. The first variation of the H-metric in F_1 vanishes, by orthogonality of F_1 to the first and second H-columns.

Suppose that $F_1, \dots, F_r$ have been constructed so that the (H, L) and (L, L) residuals vanish below order ε^r and the (H, H) residuals vanish below order ε^{r+1}. We keep the factorizations

$$F_j \in \chi\, C^\infty, \qquad \overline{F_j} \in \chi^2 C^\infty,$$

that is, $F_j = \chi \widetilde{F}_j$ and $\overline{F_j} = \chi^2 \widehat{F}_j$ with smooth $\widetilde{F}_j$ and smooth X-invariant $\widehat{F}_j$. At the next orders the mixed, line and H equations are, respectively,

$$\begin{aligned} X\langle Y_i F, F_{r+1}\rangle + C_{r,i} &= 0, \\ 2\langle w, XF_{r+1}\rangle + D_r &= 0, \\ \mathcal{L}_H F_{r+1} + Q_{r+1} &= 0, \end{aligned} \tag{R}$$

where the first two equations are of degree r and the third of degree $r + 1$. Here $\mathcal{L}_H$ is the linearized metric pairing of the old map,

$$(\mathcal{L}_H W)_{ij} = Y_i\langle Y_j F, W\rangle + Y_j\langle Y_i F, W\rangle - 2\langle Y_i Y_j F, W\rangle,$$

and $Q_{r+1,ij} = \sum_{m=1}^{r} \langle Y_i F_m, Y_j F_{r+1-m}\rangle$. These equations follow from the product rule for the quadratic metric expression, in which only the new derivative costs a factor ε^{-1}. The known terms C_r, D_r, which depend only on $F_1, \dots, F_r$, have a factor χ and their means have a factor χ^2, while Q_{r+1} has a factor χ^2. Indeed, each known quadratic term contains two factors of χ, and the means of the terms that are linear in old derivatives

of F_r acquire the second factor from the induction hypothesis, since old derivatives commute with orbit averaging.

At $r = 1$ the mixed mean is zero. Explicitly,

$$C_{1,i} = -2c\chi\langle Y_i e(\ell), Z\rangle + \chi^2\langle Y_i Z, XZ\rangle;$$

this follows by differentiating $\langle Y_i F, Z\rangle = \langle e, Z\rangle = 0$ and using $cY_0F = c\,e(\ell)$. Both terms have zero mean, by $\overline{Z} = 0$ and by the third identity for Z above. For $r \geq 2$ we replace F_r by

$$F_r + V_r, \qquad V_r = \chi\,\nu \sum_{j=1}^{d_*} \partial_{p_j} a(q, 0, 0)\,\zeta_j,$$

where the L-valued functions ζ_j are X-invariant. This preserves the earlier equations, since V_r is centered and normal to $J_2(H)F \oplus e$, and $\langle w, XV_r\rangle = 0$ by the differentiated unit-speed identity $\langle v, X\partial_{p_j} a\rangle = 0$. At degree r the mixed mean changes by $-\chi^2(Y_i q)\cdot\zeta$. Indeed, the linear old-derivative terms in V_r have zero mean because V_r is centered. The two remaining cross terms with F_1 form the variation of $\chi^2\,\overline{\langle Y_i Z, XZ\rangle}$ as p moves to $t\zeta$; by the contact formula of Lemma 16.1 this variation is $-\chi^2(Y_i q)\cdot\zeta$, and the area remains zero. A self-product of V_r would have degree $2r-1 > r$. Since $\mathrm{d}_H q$ is injective, the adjoint $(\mathrm{d}_H q)^*$ has a smooth right inverse; we use it to choose ζ making the mean of C_r zero. For this we use the smooth coefficient of $\overline{C_r}$ after the factor χ^2, which is available also where χ vanishes, and we then recompute the other known terms.

Put $a_i = -\mathcal{I}_X C_{r,i}$ and $b = -D_r/2$, and write $\overline{b} = \chi^2 b_m$ with the smooth mean factor b_m. Let e_N be the component of $e(\ell)$ orthogonal to the first H-columns, and set

$$w_N = c\,e_N + \chi\,\nu v, \qquad Q_N = |w_N|^2 = c^2|e_N|^2 + \chi^2.$$

The function Q_N is X-invariant and has a positive minimum on the compact buffer, since e is independent of the H-columns. With $Xv = \kappa j$ as in Lemma 16.1, the line correction is

$$b_1 = -\kappa\,b_m/\overline{\kappa^2}, \qquad u = Q_N^{-1}\mathcal{I}_X\big(b + \chi^2\kappa b_1\big), \qquad W_{\text{line}} = u\,w_N + \chi b_1\,\nu j.$$

Its pairings with the first H-columns are zero. The primitive has an input of zero mean, and differentiation gives

$$\langle w, XW_{\text{line}}\rangle = Q_N\,Xu - \chi^2\kappa b_1 = b.$$

The correction W_{line} has the factor χ and its mean has the factor χ^2: the mean of $u\,c\,e_N$ vanishes, the term containing $u\chi$ has both factors, and $\overline{b_1 j} = -(b_m/\overline{\kappa^2})\,\overline{Xv} = 0$.

Using the pairing inverse of the independent columns $J_2(H)F \oplus e$, choose W_H in their span with pairings

$$\langle Y_i F, W_H\rangle = a_i, \qquad \langle e(\ell), W_H\rangle = 0,$$
$$\langle Y_i Y_j F, W_H\rangle = \tfrac{1}{2}\big(Y_i a_j + Y_j a_i + Q_{r+1,ij}\big) - \langle Y_i Y_j F, W_{\text{line}}\rangle.$$

Then $F_{r+1} = W_{\text{line}} + W_H$ solves all three equations (R). The line pairing is not affected by W_H, since the pairings of W_H with e and ν vanish. The pairing inverse does not depend on the new orbit variable, so the prescribed pairings give W_H the same factor and mean factor. This completes the induction, to any finite depth. All coefficients

extend by zero off the buffer of the wave. Under a change of unit of L, the coordinates of a in that unit are odd while $Z = \nu a$ is unchanged, the coordinate vectors v, j are even, κ, u, b_1 have the parity of L, and w, w_N have this parity as columns evaluated in that unit. Hence the recurrence and all resulting ambient displacements F_r are well defined independently of the unit.

For truncation at depth R_*, the error in the formal first Gram matrix is $O(\varepsilon^{R_*})$. After substitution of a twisted normalized power of integer frequency $K_0 = \varepsilon^{-1}$, the bounds in derivative order J are $O(\varepsilon^{R_*-J})$ for the error and $O(1+\varepsilon^{1-J})$ for the map; this uses the power derivative estimate of the metric production argument. The limiting augmented columns for the good hyperplane H are $J_2(H)F$ together with w. They are independent, because the normal projection of $c\,e$ and the component $\chi\,\nu v$ are orthogonal and cannot vanish simultaneously. They give a positive lower bound, uniform on the compact buffer, for the normal second form on $\operatorname{Sym}^2 H$. In this test the first and pure H-columns and the first line column have errors $O(\varepsilon)$. Away from the support the approximation is the old map. Where the old augmented columns are available and $c = 1$, they themselves give a good hyperplane; on the rest of the complement the phase hypothesis on F supplies the bound.

Six ordinary Hopf factors and exact correction. This step uses only six new input factors. The coefficient map A_1 is injective: where $c < 1$ we are in $\operatorname{supp}\chi$, where $\beta = 1$ and the new line block ι_L is isometric; where $\beta \neq 1$ we have $c = 1$ and the old block is A_0. Fix a finite atlas of slow charts trivializing L over the compact tube containing the supports. For an integer T, subdivide these charts in every slow coordinate at scale $1/T$, and take a squared partition of unity $\sum_\lambda \rho_\lambda^2 = 1$ on a neighbourhood of the wave support, with all cutoffs in the trivializing charts. The enlarged boxes have bounded overlap and their intersection graph has bounded degree, so a fixed number d_c of colours suffices for every T. At a box centre b_λ, round the six coordinates of $T\iota_L(\ell_\lambda)(b_\lambda)$ to an integer row $Q_\lambda \in \mathbb{Z}^6$, where ℓ_λ is the local unit of L on that box.

On the six new spheres of size h, let Ψ_{Q_λ} be the tensor product of the normalized coordinate powers of degrees $|Q_{\lambda j}|$, conjugating negative powers and using one fixed unit vector for zero powers. It is a unit vector in $\mathbb{C}^{h^6}$, and the generator of factor j acts on it by $2\pi\mathrm{i}\,Q_{\lambda j}\Psi_{Q_\lambda}$. On each colour take the sum of

$$\rho_\lambda\Big(\operatorname{Re}\Psi_{Q_\lambda},\ \ell_\lambda \operatorname{Im}\Psi_{Q_\lambda}\Big)$$

over its boxes. Since the supports within one colour are disjoint, the sum over the orthogonal colour blocks is a unit section Z_T of the twisted sphere bundle of size $b_* = d_c h^6$. The cutoff supports allow an intrinsic extension of each summand by zero. This fixes the auxiliary sphere size used in the periodic coefficients before T varies. In a local unit of L, let Y_1 be the new coefficient field of that unit given by ι_L. Then Z_T satisfies the Hopf equation with speed T up to an error

$$T^{-1}Y_1 Z_T - X(Z_T) = O(T^{-1}),$$

where $X(Z_T)$ denotes the twisted generator at the image point. The estimate holds in all fixed scaled C^J norms. Indeed, on one summand the difference is

$$2\pi\mathrm{i}\,\rho_\lambda\left\langle \iota_L(\ell_\lambda),\, Q_\lambda/T-\iota_L(\ell_\lambda)\right\rangle \Psi_{Q_\lambda},$$

up to conjugation, and $Q_\lambda/T-\iota_L(\ell_\lambda)=O(T^{-1})$ on the box by the rounding and the variation of $\iota_L(\ell_\lambda)$ over a slow box of size $O(T^{-1})$. Scaled slow derivatives preserve this bound, as do scaled derivatives of the cutoffs and of the normalized powers, by their $O(T^j)$ bounds in each fixed derivative order j. The sums have bounded overlap. The sign of the unit gives the same intrinsic estimate on each overlap.

Use the ordinary T-th powers on *all* old input factors, and denote the powered sample by x_T. Use the twisted K_0-th power P_{K_0} on Z_T, with $K_0=\varepsilon^{-1}$. Evaluating every coefficient at these samples gives the finite map

$$F(x_T)+\sum_{r=1}^{R_*}\varepsilon^r F_r\Big(x_T,P_{K_0}Z_T\Big).$$

It is smooth across the boundary of the wave support, since all added coefficients vanish outside its compact slow buffer. Along the normalized coefficient fields $Y_{A_1(\cdot)}/T$, the old derivative is exactly the derivative of the old profile, and where the coefficients are supported we have $\beta=1$, so the new line part obeys the approximate transport above. Thus the formal first Gram matrix (K), with α,p evaluated at x_T, approximates the first Gram matrix of this map. Outside the wave support the map depends only on the old inputs, $c=1$, and the old metric is given exactly by the same formula. On the transition region of β the potential s also vanishes.

To quantify the estimates we use local torus angle charts, projective charts and slow charts, all scaled by $1/T$ on compact buffered nets. At each slow point, a local frame of $A_1(V_0)$ and a complementary frame in the angle coordinates straighten the coefficient fields; these frames are chosen on finitely many slow charts before rescaling. The coefficient map has a positive lower bound for its least singular value on the compact base, so these changes of coordinates and their inverses have bounded jets in the scaled charts. The chart transitions and their inverses are uniformly bounded at every fixed order, with uniformly bounded overlap of the enlarged charts. In these charts the derivatives of order J of the reference map are $O(1+\varepsilon^{1-J})$, and the defect of the first Gram matrix is

$$O(\varepsilon^{R_*-J})+O(\varepsilon^{-B_J}/T)$$

in scaled order J, for a finite exponent B_J. This follows by the chain rule from the finite recurrence and from the scaled transport estimate together with its differentiated versions. The same comparison of first and second columns retains a positive phase margin for small ε and large T, since the compact limiting test on $\mathrm{Sym}^2 H$ is independent of ε and T. On the compact complement we use the positive phase margin of the old map.

The target metric (K) has an exact description by a change of source. In a unit of L, translate the angle torus along the line generator $Y_c=cY_0+\beta Y_1$ by the variable time $c\,s(x_T)/T$, and extend this map by the identity off the support of the potential.

The map is intrinsic under reversal of the unit, since s and Y_c change sign together. Since the old power multiplies old derivatives by T and the new fields annihilate the potential, its differential along H and L adds $c\alpha(x_T)$ and $c^2 p(x_T)$ times the unit line, respectively. In particular, the rank-one determinant is $1 + c^2 p(x_T) > 9/10$, because $|p| \le |\mathrm{d}_{V_0} s| < 1/10$. On each compact angle torus the map is a local diffeomorphism of degree one, since scaling the variable time gives a homotopy to the identity. Hence it is a diffeomorphism Θ_T, and it preserves the coefficient distribution. The pullback by Θ_T of the unit metric of V_0 is exactly (K). This target metric is uniformly elliptic and has uniformly bounded scaled derivatives of every fixed order, by the bound on $\mathrm{d}_{V_0} s$ and the estimates for the old powers.

We now apply the polynomial-size entrance estimate of Proposition 3.1, in the form detailed in Section 11.5, in these uniformly controlled charts. The fixed reference derivative order is chosen first. The exponent of the reference bound in ε^{-1} is independent of the truncation depth, and the phase lower bound has exponent zero. Thus an error $O(\varepsilon^K)$ in order 440, for some finite K, suffices for a phase-preserving exact correction. Take $R_* > K + 440$, then ε a sufficiently small reciprocal integer, and then T sufficiently large for all of the finitely many entrance bounds. In the unscaled coefficient fields the corrected map has metric $T^2 \Theta_T^* h_0$, where h_0 is the specified unit metric. Divide it by T and precompose with Θ_T^{-1}. The resulting map is exactly isometric along A_1. Since Θ_T^{-1} preserves the leaves, the normal second form is transported by its first leaf derivative: the chain-rule term involving its second derivative lies in the span of the first columns and is removed by the normal projection. Phase regularity is therefore retained, and the target remains $\mathbb{R}^N$ throughout. □

17. Global assembly

We choose the slow radii on each tube before choosing the common Hopf size. The attaching sphere lies in an inner plateau. Farther out, choose a smooth slow angle ω which decreases from $\pi/2$ to zero within the band of Section 15.1. Set $c = \cos\omega$ and $\chi = \sin\omega$ on that band, and continue them by $c = 0$, $\chi = 1$ on the filled side and by $c = 1$, $\chi = 0$ on the far side. We keep an outer buffer of the band beyond $\operatorname{supp}\chi$. Choose a slow cutoff ρ equal to one on $\operatorname{supp}\chi$ and on its buffer of fixed H, and turning off on a farther collar on which $\chi = 0$ and L is still defined. The components of the core are disjoint, so all such cutoffs can be chosen simultaneously.

For each fixed Hopf size we apply the exact metric production of Section 11.5 on the doubled exterior, with passive weight $d = \chi$. On the annuli the passive value fa takes values in the old line L. We use the same exponent $k \ge 4$ and parameter $\tau = L_0^{-1}$ as in that production, and write b_τ for the evaluation of fa after the old and active powers of frequency L_0. The resulting exact map F_τ has partial metric

$$h_0 - \tau^{2k} \chi^2 |\mathrm{d}_{V_0} b_\tau|^2$$

on the exterior, where $V_0 = TM$, and it has a positive phase margin by the exterior construction.

In a unit ℓ of L, let Y_0 be its old line generator, and define the forward orbit translation along Y_0 by the variable time $t_\tau = \rho\,\tau^k b_\tau$, extended by the identity off the collar of the potential. Its first angle derivatives are $O(\tau^{k-1})$. The derivative of the translation on the angle torus is the identity plus a rank-one map, with determinant $1 + Y_0 t_\tau > 0$ for small τ. By degree one, it is a diffeomorphism of each torus. Denote its inverse by Ψ_τ and put $s = -t_\tau \circ \Psi_\tau$. This is a well-defined L-valued time, since reversing the unit reverses t_τ and Y_0 together. The leaf differential of Ψ_τ is $v \mapsto v + (\mathrm{d}_{V_0} s(v))\ell$, so

$$\Psi_\tau^* h_0 = h_0 + 2\ell^* \odot \mathrm{d}_{V_0} s + (\mathrm{d}_{V_0} s)^2.$$

Throughout the support of the passive weight we have $\rho = 1$, so $\tau^{2k}\Psi_\tau^* |\mathrm{d}_{V_0} b_\tau|^2 = |\mathrm{d}_{V_0} s|^2$, and $1 - \chi^2 = c^2$. Thus the old exterior map $\widetilde{F}_\tau = F_\tau \circ \Psi_\tau$ has exactly the metric

$$h_0 + 2\ell^* \odot \mathrm{d}_{V_0} s + c^2 (\mathrm{d}_{V_0} s)^2.$$

This is (L). In the inner plateau its H-metric is exactly the unit metric. Its phase regularity on the exterior is preserved by the line translation, since Ψ_τ preserves the leaves and the chain-rule term of its second derivative is tangential. On the collar where only the potential is present, $c = 1$ and the same formula holds, including where the potential is turned off.

17.1. **Boundary comparison with the passive term.** We verify the augmented rank throughout the wave buffer, including the transition region $0 < \chi < 1$.

Lemma 17.1. *Let $k \geq 4$. For all sufficiently small τ, the columns $J_2(H)\widetilde{F}_\tau$ and the surviving first column $\mathrm{d}_L \widetilde{F}_\tau$ are linearly independent throughout the wave buffer. Their augmented columns lie in the collar class of (P), and on the attaching sphere in the class J_Σ.*

Proof. We work at the inverse-translated point and its independently powered sample. Modulo the first columns, the Hessian of the composition is the Hessian of F_τ evaluated on the two moved inputs; the term from the second derivative of the translation lies in the first span. For $h \in E_v$,

$$\mathrm{d}\Psi_\tau(h) = h - \frac{\mathrm{d}_h t_\tau}{1 + Y_0 t_\tau}\,\ell = h - \tau^{k-1} S(h, a) + O(\tau^{2k-2}),$$

where we have used that $\rho = 1$ on this buffer whenever the passive term is nonzero. For $b \in A$ the change is $O(\tau^{k-1})\ell$. Therefore the first-order change in the mixed column of the reference map is $-\tau^{k-1} D^2 B(S(h, a), b)$. The mixed old–active columns of the reference map vanish, so the change coming from the moved b alone vanishes. The cross term from both substitutions is $O(\tau^{2k-2})$, and changing the inputs of the perturbed second columns costs $O(\tau^{2k-3})$.

In the unmoved mixed Hessian the leading retained term is $\tau^{k-2} I_x U(h, d_b^y a)$. The first term visible in the four cofactor rows is $\tau^{k-1} J_x(b) U(h, a)$: the other term with one fast derivative has zero cofactor projection by the separate reference blocks of Section 11.6. Every higher normal profile starts at order $k + 1$; its term with two fast derivatives lies in $I_x W$, so after projection to the four rows it has order at least k. The lifts of the pure pairings start at order $2k - 2$, so after two derivatives their error in the four

rows is $O(\tau^{2k-4})$. The lifts of the first pairings and the exact metric correction are of higher order still. The retained projection of the translation term of the reference map vanishes, since that term is an active second column; its cofactor projection is $-\tau^{k-1} M_x(S(h,a) \otimes b)$. Replacing the reference projector by the projector onto the complement of the first and pure columns of $\widetilde{F}_\tau$ costs $O(\tau^{2k-4})$: by the formula for the Gram inverse the projector moves by $O(\tau^{k-2})$, and the mixed columns being projected are $O(\tau^{k-2})$. These estimates are uniform over the buffer, since the slow weight is annihilated by all leaf derivatives.

Scale the projected retained rows by $\tau^{-(k-2)}$ and the four cofactor rows by $\tau^{-(k-1)}$. Since $k \geq 4$ gives $2k - 4 \geq k$, the resulting mixed columns are within $O(\tau)$ of the uniform paired injection (P). Moreover, $\mathrm{d}_L \widetilde{F}_\tau$ is close to the fourth active first column of the reference map, which is independent of the remaining reference columns. This proves the simultaneous independence of $J_2(H)\widetilde{F}_\tau$ and $\mathrm{d}_L \widetilde{F}_\tau$ through the whole buffer. The same estimate gives the collar class of (P), by linear interpolation over the $O(\tau)$ error after scaling, by undoing the positive row scalings, and by undoing the short transport of the base projector. On the attaching sphere the resulting comparison is with J_Σ. The inverse translation can be removed from the formal endpoint by its time isotopy. The powers and shifts of the independent samples are allowed throughout the comparison, so no new identification of normal bundles is needed after substitution. □

17.2. **Filling the core and completing the line.** On the attaching plateau, where the H-metric is the unit metric, orthogonalize the surviving first column against the first H-columns and normalize it positively. This does not change the class of the injection. Extend the augmented columns of $\widetilde{F}_\tau$ formally over an inward collar: make the radial coordinate stationary at a comparison sphere, run the class homotopy of Lemma 17.1 to J_Σ, and then use the fixed disk extension of Section 7. With smooth stationary end parameters this gives a smooth extension which is pulled back from the core P near the zero section and agrees with the columns of $\widetilde{F}_\tau$ near the outer collar. Corollary 14.2 now gives an H-free filling with unit H-metric which equals the exterior map on a smaller open collar, together with an isometric augmented line and two normal copies of L orthogonal to all augmented columns. In particular, the old map obtained by gluing is a single smooth map on the original Hopf product over M.

On the inner collar write the surviving derivative as

$$\mathrm{d}_L \widetilde{F}_\tau = \mathrm{d}_H \widetilde{F}_\tau \, T + \mu \, e_\perp,$$

where $e_\perp$ is the positively orthogonalized unit line and $\mu > 0$. After a radial retraction, cut off T and $\log \mu$ to zero inwards by a slow cutoff, and use the same formula with the filled H-columns and their augmented unit line. The augmented span is unchanged. This gives a column e which agrees with the old line column on the exterior side, as required where $c > 0$; inside, $c = 0$. The two normal copies extend outwards through the wave buffer by orthogonal transport in the rank-five complements of the augmented spans. To match them smoothly on an open overlap, we first trivialize those complements by radial transport, and reparametrize the given family radially so that

it is stationary before the end of the overlap. Thus both copies are smooth isometric inclusions on a neighbourhood of the wave support.

Extend s inwards by a slow radial retraction and a cutoff in the region $c = 0$, where only the H-metric is prescribed. The first original angle derivatives of its incoming germ are $O(\tau^{k-1})$ on the compact old host, by the bounds for the powers and the translation. The radial extension has derivatives of the same size and all fixed coefficients are bounded, so the hypothesis $|\mathrm{d}_{V_0} s| < 1/10$ holds for the extended potential when τ is small. Its support lies in a fixed slow tube. Away from the wave neighbourhood the old map is phase regular, since it is the corrected, possibly translated, exterior map. Its metric there is exactly the prescribed formula, and it is the unit metric outside the tube of the potential. All hypotheses of Proposition 16.2 therefore hold on this single old host, with $N = s_n + 5 \geq s_n + 1$.

Take $\iota_L \colon L \to \mathbb{R}^6$ to be the original inclusion of L into the four active coefficient slots followed by two zeros, and choose the slow cutoff β once on the tubes. Since there are finitely many components, all constants, frequencies and slow cutoffs can be chosen simultaneously. Proposition 16.2 gives a single exact phase-regular unit partial isometry over M into $\mathbb{R}^{s_n+5}$, defined on a product of ordinary Hopf spheres. The final injective coefficient map is fixed independently of their common size, and its coefficient bundle is isomorphic to TM. If the core is empty, the exterior production already supplies this compact model, with zero passive term and no additional inputs. All these arguments apply for every sufficiently large common sphere size.

Finally we apply Theorem 2.2 with slow base M, with the metric on the coefficient bundle transported from the chosen auxiliary Riemannian metric, and with the identity as classifying map; the dimension condition $N \geq s_n + 2$ holds for $N = s_n + 5$. (In its proof, the passage from the auxiliary metric to g uses Lemma 5.1 and Lemma 5.2, followed by Propositions 3.1 and 4.4.) The theorem gives the asserted smooth isometric embedding for the independently prescribed metric g, in the same $s_n + 5$ coordinates. The prescribed image ball and the phase regularity of the embedding also follow from that theorem. This proves Theorem 1.1.

Acknowledgements

Generative AI was used in developing the proofs.

References

[C27] É. Cartan, *Sur la possibilité de plonger un espace riemannien donné dans un espace euclidien*, Ann. Soc. Polon. Math. **6** (1927), 1–7. Original volume, pp. 1–7.

[EM02] Y. Eliashberg and N. Mishachev, *Introduction to the h-Principle*, Graduate Studies in Mathematics, vol. 48, American Mathematical Society, Providence, RI, 2002. doi:10.1090/gsm/048.

[G70] R. E. Greene, *Isometric embeddings of Riemannian and pseudo-Riemannian manifolds*, Mem. Amer. Math. Soc., no. 97 (1970). doi:10.1090/memo/0097.

[G86] M. Gromov, *Partial Differential Relations*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3), vol. 9, Springer-Verlag, Berlin, 1986. doi:10.1007/978-3-662-02267-2.

[G17] M. Gromov, *Geometric, algebraic, and analytic descendants of Nash isometric embedding theorems*, Bull. Amer. Math. Soc. (N.S.) **54** (2017), no. 2, 173–245. doi:10.1090/bull/1570.

[GR70] M. L. Gromov and V. A. Rokhlin, *Embeddings and immersions in Riemannian geometry*, Russian Math. Surveys **25** (1970), no. 5, 1–57. doi:10.1070/RM1970v025n05ABEH003801.

[Gü89] M. Günther, *Zum Einbettungssatz von J. Nash*, Math. Nachr. **144** (1989), 165–187. doi:10.1002/mana.19891440113.

[Gü91] M. Günther, *Isometric embeddings of Riemannian manifolds*, Proceedings of the International Congress of Mathematicians, Vol. I, II (Kyoto, 1990), Math. Soc. Japan, Tokyo, 1991, pp. 1137–1143.

[J26] M. Janet, *Sur la possibilité de plonger un espace riemannien donné dans un espace euclidien*, Ann. Soc. Polon. Math. **5** (1926), 38–43. Original article.

[MAdP21] J. Martínez-Aguinaga and Á. del Pino, *Convex integration with avoidance and hyperbolic* $(4, 6)$ *distributions*, preprint (2021), arXiv:2112.14632.

[MS74] J. W. Milnor and J. D. Stasheff, *Characteristic Classes*, Annals of Mathematics Studies, no. 76, Princeton University Press, Princeton, NJ, 1974. doi:10.1515/9781400881826.

[N56] J. Nash, *The imbedding problem for Riemannian manifolds*, Ann. of Math. (2) **63** (1956), no. 1, 20–63. doi:10.2307/1969989.

[WY26] Y. Wang and W. Yang, *Every closed smooth n-manifold embeds isometrically in* $\mathbb{R}^{n(n+1)/2+28}$, Hexagon Math (2026), hexagon:2611.00153v1.

[Wu22] L. Wu, *Characteristic cycles associated to holonomic $\mathcal{D}$-modules*, Math. Z. **301** (2022), no. 2, 2059–2098. doi:10.1007/s00209-022-02974-0. arXiv:2011.14527.

[YW26a] W. Yang and Y. Wang, *Smooth local isometric and conformal embeddings in optimal dimension*, Hexagon Math (2026), hexagon:2610.00378v1.

[YW26b] W. Yang and Y. Wang, *Isometric embeddings of parallelizable and stably parallelizable Riemannian manifolds*, Hexagon Math (2026), hexagon:2610.00690v1.

[Z16] N. Zergänge, *Isometrische Einbettung von geschlossenen Riemannschen Mannigfaltigkeiten*, preprint (2016), arXiv:1607.04259.

[Z17] N. Zergänge, *Stability analysis of isometric embeddings*, preprint (2017), arXiv:1712.02766.

Tsinghua University, Beijing, China

Tsinghua University, Beijing, China