hexagon:2611.00698v3

A two-chirality twistor correspondence for smooth Lorentzian Ricci-flat germs

LIOR ISTHMUS[1]

Abstract

We give a local twistor presentation of each smooth pointed, oriented and time-oriented Lorentzian Ricci-flat four-germ with spin, at its physical scale. The two inputs are ordinary cohomology classes of weights 2 and -6 for the flat hypersurface twistor CR structure on $\mathbb{R}^3 \times \mathbb{CP}^1$, whose projective line is the sphere of spin directions. In the selected weighted domain of the growth-controlled coefficient complexes, the coupled reality condition gives real transverse trace-free seeds for the metric and momentum. Starting with encoded test tensors, we take three successive normalized negative-chirality contact responses, recovering a tensor at each step. Contractions of the resulting tensors give the flat scalar and York inverses for the constraints. Initialized nonlinear transport yields maximal vacuum data and both spin gradient operators. Classical Cauchy evolution develops these data. Conversely, we extend normalized maximal-disk data to weighted vacuum data on $\mathbb{R}^3$, with exact agreement on a smaller disk, and decode the extension. The inverse gives a comparison class of presentations, with smooth local selections in finite-dimensional families and effective descent. A covariant wave kernel, determined modulo jointly smooth kernels, transports the degree-one evaluation section on each sky and its point connection between slices. Every pointed physical isometry with a specified actual spin lift has a unique full source lift modulo normalized presentation cells. On analytic germs we compare the wave source with the calibrated operator source of the analytic projective twistor correspondence. After real restriction and complete formal enrichment, the initialized comparison preserves coupled reality.

Date: 2026-11-19.

[1]Email: lioristhmus@proton.me. ORCID: 0009-0004-7908-9876.

Contents

Abstract 1
1. The local two-class correspondence 2
2. Intrinsic initial sources and local vacuum evolution 6
3. Ordinary classes and the initial-data construction 31
4. Actual source execution from ordinary classes 52
5. Local displays of arbitrary vacuum germs 84
6. The covariant wave enrichment 95
7. Geometric admission and full source comparisons 132
8. Local presentations, inverse functors and descent 146
9. Comparison on the analytic locus 151
Conclusion 165
Appendix A. Weighted construction and based gauge 166
Appendix B. Relative maximal extension on one smooth branch 189
AI and tool use disclosure 219
References 220

1. The local two-class correspondence

Penrose's nonlinear graviton construction encodes complex vacuum metrics with one Weyl chirality in the holomorphic geometry of a deformed twistor space (Penrose 1976). In his helicity convention the construction describes left-handed gravitational fields. The googly problem asks how the same twistor convention can also describe interacting right-handed fields. Penrose proposed adjoining twistor differentiation operators to holomorphic functions to form a noncommutative twistor algebra (Penrose 2015).

Twistor actions also describe interactions between the two chiralities. Adamo and Mason computed Einstein MHV amplitudes from the conformal gravity action and proposed an Einstein twistor action. For Euclidean gravity at $\Lambda = 0$, Sharma showed that his twistor action reduces to the Plebański action. Here the inputs are an ordinary pair in a specified weighted domain, and the outputs are smooth local Lorentzian Ricci-flat germs with comparisons over isometries carrying specified spin lifts.

A second approach starts on a spacelike hypersurface. LeBrun recovered the conformal geometry of a Riemannian three-manifold from the CR structure on its sphere bundle (LeBrun 1984). Mason explained how hypersurface twistors also encode the trace-free second fundamental form and distinguished these conformal data from the scale and mean curvature (Mason 1989). Mason and Penrose subsequently investigated a twistorial formulation of the full vacuum equations (Mason and Penrose 1989).

In the analytic projective twistor correspondence (Isthmus 2026a), two classes encode the two chiral curvature records, and analytic continuation

recovers an Einstein germ. A smooth germ can have a nonzero field whose derivatives vanish at a point. For a smooth local correspondence the data must therefore describe a field on a neighborhood. We work at $\Lambda = 0$ and retain the physical metric, both orientations and the actual spin structure.

Two localization issues arise. The decoder acts on whole weighted data, and a cutoff of vacuum data introduces a constraint error. Appendix B uses the maximal-data extension of Czimek and its explicit linear right inverses to obtain a vacuum extension that agrees with the prescribed data on a smaller disk. Two displays of the same germ may also use different slices. Radzikowski characterized the Hadamard condition by the wavefront set on a fixed spacetime, and Gérard and Wrochna gave a pseudodifferential factorization of the wave equation. Section 6 uses the positive-frequency branch and proves uniqueness of the kernel modulo joint smoothing in families. It transports the full degree-one evaluation between slices, canceling the direction-dependent Doppler factor already present for boosted flat hyperplanes.

On the fixed input space $N_0 = \mathbb{R}^3_x \times \mathbb{CP}^1$, each spin direction determines a complex null line with homogeneous representative $z \in \mathbb{C}^3$. Write $\mathscr{E}^\bullet_m$ for the weight-m CR complex with growth-controlled coefficients. Its first differential is $d^0 u = (\bar\partial u, Du)$, where $\bar\partial$ acts on the sphere and $D = z \cdot \partial_x$. The ordinary inputs lie in the first cohomology $H^1(\mathscr{E}^\bullet_m)$ at the two weights $m = +2, -6$ fixed by contour normalization. The coupled reality condition on the class pair makes both decoded seeds real, so the metric and momentum seeds enter a single vacuum development. In the diagram below the left category records local displays of these pairs, the middle category admits sources by independent geometric and operator equations, and the right category records physical germs and isometries with their specified spin lifts.

Theorem 1.1 (Smooth two-class local correspondence).
Let $\mathsf{Vac}_0^{\mathrm{loc}}$ have the physical germs above as objects and pointed physical isometries with their specified actual spin lifts as arrows. There are equivalences

$$\mathsf{WavePair}_0^{\mathrm{loc}} \xrightarrow{E} \mathsf{WaveSrc}_0^{\mathrm{loc}} \xrightarrow{R} \mathsf{Vac}_0^{\mathrm{loc}}.$$

The first domain consists of local displays of ordinary pairs as in Definition 8.1. All freely specified functional data of a chart are contained in the ordinary pair of Definition 3.11 up to its effective cochain arrows. The positive number r, retained neighborhood and markings specify the display. A local maximal initial slice is part of the presentation. The middle domain is defined by independent equations on actual geometric and spin fields and on complete operator and module data in Definition 7.7. The independent source uses a common causally convex, globally hyperbolic wave neighborhood for its retained carrier hits and connecting null segments, as specified in Definition 6.12. Every object of either of the last two categories is locally represented by a display. For each physical arrow between recognized germs there is a unique full source arrow over it, up to normalized presentation cells. The

inverse identifications are natural and obey the triangle identities. The physical target includes flat and algebraically special germs and imposes no asymptotic condition.

For smooth finite-dimensional families on common representatives with smooth markings, a local selection has one positive normalization radius and one buffered parent on retained compact parameter sets. This parent is a larger common domain in which the retained patches lie a positive distance from the boundary. For each stipulated finite derivative and coefficient depth the selected output has a bound in finitely many input seminorms on that parent and the positive margins. Supplied comparison families have the corresponding estimates with their map and inverse bounds. The local presentations and their full arrows admit effective atlas descent over open covers of Hausdorff representatives with effective actual spin descent.

On the analytic geometric locus, using the calibrated operator correspondence of Isthmus 2026a (v1.1, Definition H.1 and Theorem 9.10), there are structured comparisons in both directions after real restriction and complete formal enrichment. The analytic transport and descent results used for this clause are specified in Section 9. These comparisons carry the analytic chiral structures in their simultaneous presentations, the optical and numerical scales, both associators, actual root squares and base factors, the canonical finite-rank structures, actual spin and the original coupled square cell. The initialized framing compares the full degree-one evaluation section and its point connection in every direction. The finite-order bounds there apply to conditioned families on common analytic parents and their retained real restrictions.

Different whole-space extensions of the same local data can present the same pointed germ. Definition 8.1 therefore takes arrows between displays to be full comparisons of their executed sources on a common local restriction. This makes E fully faithful. Coverage of the physical germs follows from Theorem 5.7, and full faithfulness of R follows from the lifting theorem for physical arrows with specified spin lifts (Theorem 7.11). The proof of the combined statement concludes with Theorem 9.8.

Inputs and physical scale. The two functional data are classes in $H^1(\mathscr{E}_2^\bullet) \oplus H^1(\mathscr{E}_{-6}^\bullet)$. Membership uses actual primitives, the weighted linear root test and explicit polynomial obstructions. Coupled reality is imposed on the two inputs together. Smallness concerns the normalized whole input of the chart. In physical variables the same display has $h^{[r]} = r^2 h$, $K^{[r]} = rK$ and $\mathbb{D}_\pm^{[r]} = r^{-1}\mathbb{D}_\pm$. An isometry between outputs preserves the metric at this restored scale, even if the displays use different normalization radii. Classical Einstein evolution develops the smooth vacuum data supplied by the pair. The resulting metric fixes the wave kernel up to addition of a jointly smooth kernel. An exact quantum state requires further data.

Relation to the analytic correspondence. The original independently admitted algebroid contains an actual half-canonical square, a framed local-system factor and factorially bounded analytic operators (Isthmus 2026a, Definition H.1). The smooth evaluation has degree one and moves with the physical point. Theorem 9.7 transports this evaluation while retaining the original degree-zero calibration. Even when the metrics agree, a comparison must fix the relative scale of the evaluation sections, respect their sign under conjugation and compose coherently along physical arrows. In the analytic global correspondence, complete complex skies and their point spin maps glue, and null-path monodromy is recomputed under point restriction (Isthmus 2026b, Lemma 5.5). Here the bases are local smooth Lorentzian representatives at $\Lambda = 0$ and the descended objects are atlases of whole-input displays on local cores. All the comparisons are local to the retained real ray plaques.

Organization. In the forward construction the local transverse trace-free encoder has Fourier order four, and each negative contact response lowers that order by two, so three responses give $4 - 2 \cdot 3 = -2$, the order of the flat constraint inverses (Lemma 4.11 and Theorem 4.12).

Purpose	Construction	Proof
Construct a spacetime germ	Read the two seeds from the classes. Contract the three responses to obtain the flat inverses. Initialized nonlinear transport gives vacuum initial data and hence the two spin gradient operators. Classical Einstein Cauchy evolution develops the data.	§3.6, Theorem 4.18, Theorem 2.6
Present a given germ	Extend normalized maximal-disk data to the whole space with exact agreement on a smaller disk. Each such extension in the fixed gauge decodes to a unique pair of classes. The germ determines the comparison class of its displays.	Theorem B.1, Theorems 3.13 and 5.7
Compare distinct slices	The scalar wave kernel transports the full evaluation between carriers. A physical isometry with its specified spin lift then has a unique full source comparison up to normalized cells.	Theorems 6.10 and 7.11

The resulting local inverses and the descent of display atlases are proved in Theorem 8.3 and Theorem 8.5.

2. Intrinsic initial sources and local vacuum evolution

The source has two chiral gradient operators with a common real Clifford symbol. We first give independent local tests for these operators and prove that they recover vacuum initial data and their actual spin comparisons. This recognition applies without weighted smallness or a maximality assumption. The ordinary-pair construction of Sections 3–4 produces the maximal subclass used in the local coverage argument.

An actual finite Cartan collar then carries the source off its marked initial hypersurface. Its soldering, connection and Weyl field obey finite equations with a closed subsidiary system. Their geometric recognition is proved directly; smooth existence and geometric uniqueness use the precise classical Cauchy input stated below. Finite-dimensional families are constructed on one low-order-selected parent, with every higher estimate referring to that same solution.

The comparison results in this section retain the initial hypersurface. They identify all physical and actual spin comparisons in that category, and compare exterior, gauge and scale choices for a fixed physical initial germ. The wave enrichment and comparisons that change the initial hypersurface require their additional structures.

2.1. Intrinsic initial sources. The initial source constructed from an ordinary pair has a canonical weighted presentation. Its geometric recognition uses only local differential identities. We formulate those identities on an arbitrary oriented smooth three-manifold Σ, with no weighted coordinate system or prescribed mean curvature. This separates the construction of a display from the recognition and comparison of its actual initial fields.

All manifolds are Hausdorff and second countable, and initial open sets have no boundary as manifolds. A vacuum initial datum is $d = (\Sigma, h, K, P_\Sigma)$, where h is a positive metric, K is a real symmetric covariant two-tensor, and P_Σ is an actual Spin(3) lift of the oriented h-orthonormal frame bundle. Put $\tau = \operatorname{tr}_h K$ and require

$$R(h) + \tau^2 - |K|_h^2 = 0, \qquad \operatorname{div}_h K - \mathrm{d}\tau = 0.$$

A datum comparison over an orientation-preserving diffeomorphism $f : \Sigma \to \Sigma'$ consists of $f^*h' = h$, $f^*K' = K$ and an actual spin-bundle lift of its orthonormal frame map. Denote this groupoid by Data_0^∞. The definitions apply componentwise when the initial manifold is disconnected.

The algebraic marking of a gradient source is a real oriented and time-oriented Lorentz rank-four bundle $(\mathbb{E}, \eta)$ of signature $(+---)$, its actual $\mathrm{Spin}^+(1,3)$ lift, a future unit section N, and the associated complex Dirac module $\mathscr{S} = \mathscr{S}_+ \oplus \mathscr{S}_-$. It includes the real Clifford action c and the fixed invariant nondegenerate complex alternating form B. The normalization is

$$c(v)c(w) + c(w)c(v) = 2\eta(v,w)I, \qquad \operatorname{tr}_{\mathscr{S}}(c(v)c(w)) = 4\eta(v,w).$$

For clarity, in a local spin frame the model is

$$c(v)=\begin{pmatrix}0 & V\\ \widetilde{V} & 0\end{pmatrix},\qquad V=v^0I+\sum_a v^a\sigma_a,\qquad \widetilde{V}=v^0I-\sum_a v^a\sigma_a,$$

$$\mathsf{B}=\begin{pmatrix}\epsilon & 0\\ 0 & \epsilon\end{pmatrix},$$

$$\epsilon=\begin{pmatrix}0 & 1\\ -1 & 0\end{pmatrix},$$

$$\varrho(A)=\operatorname{diag}(A,(A^\dagger)^{-1}),\quad A\in SL(2,\mathbb{C}).$$

Here σ_a are the Pauli matrices. The oriented Clifford product fixes the chirality labels. The Clifford vectors generate the full complex endomorphism algebra of each Dirac fiber. These are algebraic bundle markings; a metric on $T\Sigma$, a soldering or a connection is not supplied.

Definition 2.1 (Intrinsic vacuum gradient source).
Supply two actual smooth complex-linear first-order operators

$$\mathbb{D}_\pm:\Gamma_{\mathrm{loc}}(\mathscr{S}_\pm)\longrightarrow\Gamma_{\mathrm{loc}}(\mathbb{E}_{\mathbb{C}}\otimes\mathscr{S}_\pm).$$

Their scalar commutators must have a common real orientation-preserving bundle isomorphism as symbol, and their complete outputs must be tangent to the internal spatial plane:

$$[\mathbb{D}_\pm,m_a]\psi=b(\mathrm{d}a)\otimes\psi,\qquad b:T^*\Sigma\xrightarrow{\sim}N^\perp,\qquad \eta(N,\mathbb{D}_\pm\psi)=0$$

for every real local scalar a. Define the algebraic and directional readouts

$$h^{-1}(\xi,\zeta)=-\eta(b\xi,b\zeta),\qquad \Theta X=b(h^\flat X),$$
$$P_X^\pm\psi=-\eta(\Theta X,\mathbb{D}_\pm\psi),\quad P_X=P_X^+\oplus P_X^-.$$

Require, for real X and real Clifford sections v,

$$X\mathsf{B}(\psi,\chi)=\mathsf{B}(P_X\psi,\chi)+\mathsf{B}(\psi,P_X\chi),$$
$$[P_X,c(v)]\in\Gamma_{\mathrm{loc}}(c(\mathbb{E})).$$

The connection D specified by $c(D_Xv)=[P_X,c(v)]$ is justified below. The remaining admission tests are

$$D\Theta=0,\qquad \Omega_D\in\operatorname{im}\mathsf{W}_\Theta,\qquad \mathsf{W}_\Theta(C)(X,Y)=\mathcal{C}(\Theta X,\Theta Y),$$

where C ranges over real algebraic Weyl tensors on $\mathbb{E}$ and $\eta(\mathcal{C}(u,v)z,w)=C(u,v,z,w)$. The curvature in this test is computed from the recovered connection. Its Weyl preimage is a readout, rather than an independently supplied curvature field.

The finite image test can be carried out explicitly. In an adapted orthonormal frame $v_0=N$, $v_a=\Theta(e_a)$, compute

$$E_{ab}=\sum_c(\Omega_D)^c{}_b(e_c,e_a),\qquad B_{ab}=\frac{1}{2}\epsilon_a{}^{cd}(\Omega_D)^0{}_b(e_c,e_d).$$

Require E, B to be symmetric and trace-free. Form the algebraic tensor $C(E,B)$ from curvature symmetries and

$$C_{0a0b} = E_{ab}, \qquad C_{cd0b} = -\epsilon_{cd}{}^{a} B_{ab}, \qquad C_{abcd} = -\epsilon_{ab}{}^{m}\epsilon_{cd}{}^{n} E_{mn},$$

and test the equality $\Omega_D = \mathsf{W}_\Theta(C(E,B))$. The proof below shows that this is exactly the image test and that the readout is independent of the adapted frame.

A physical source comparison over $f : \Sigma \to \Sigma'$ is an actual multiplication-semilinear isomorphism of the sheaves of chiral sections, $\widehat{U} = \widehat{U}_+ \oplus \widehat{U}_-$. It preserves B, the real Clifford-vector action and its orientation and time orientation, and N. Its induced vector map is denoted by L. Require the exact operator identities

$$\mathbb{D}'_\pm \widehat{U}_\pm = \widehat{L \otimes U_\pm}\, \mathbb{D}_\pm.$$

The passage from a semilinear sheaf map to the smooth bundle maps U, L is proved below. The map U itself is retained as a morphism. These objects and comparisons form Grad_0^∞.

Definition 2.1 tests the algebraic marking and the two operators directly on the initial source, before either developing the data or seeking an ordinary-pair presentation.

Lemma 2.2 (Intrinsic recovery of the directional connection). The symbol, tangential-output and normalization tests determine a positive metric h, an oriented spacelike soldering Θ, a unique real Lorentz connection D and its exact spin connection P. They recover the complete supplied operators by

$$\mathbb{D}_\pm \psi = \sum_{a=1}^{3} \Theta(e_a) \otimes P^\pm_{e_a} \psi$$

in every oriented h-orthonormal frame. If $D\Theta = 0$, then

$$\begin{aligned} K(X,Y) &= -\eta(D_X N, \Theta Y), \\ D_X N &= \Theta(K^\sharp X), \\ D_X(\Theta Y) &= \Theta(\nabla^h_X Y) + K(X,Y) N, \end{aligned}$$

where K is symmetric. Conversely, the last two formulas define the unique metric connection with the given h, K, N, Θ and zero torsion.

Proof. Positivity follows from the negative-definiteness of η on $N^\perp$. The readout satisfies $-\eta(\Theta X, \Theta Y) = h(X,Y)$ and preserves orientation. The directional operator is function-linear in X, and its scalar commutator is

$$[P_X, m_a]\psi = -\eta(\Theta X, b\,\mathrm{d}a)\psi = h^{-1}(h^\flat X, \mathrm{d}a)\psi = X(a)\psi.$$

It is therefore a connection. Its commutator action on real Clifford vectors obeys the real Leibniz rule and defines D. Commuting P_X with the Clifford anticommutator gives $X\eta(v,w) = \eta(D_X v, w) + \eta(v, D_X w)$, so D is Lorentz.

The connection D has a unique lift to the supplied spin bundle. Its induced Dirac connection and P act on Clifford vectors through the same

D, so their difference commutes with every Clifford vector. The full-matrix-algebra property makes that difference $\lambda_X I$. Both connections preserve B, and subtraction of their preservation identities gives $2\lambda_X \mathsf{B}(\psi,\chi)=0$. Thus $\lambda_X=0$. This removes the scalar connection ambiguity.

For $v_a=\Theta(e_a)$, every vector in $N_{\mathbb{C}}^{\perp}$ is $\sum_a v_a(-\eta(v_a,\cdot))$. Apply this identity to the entire output of $\mathbb{D}_{\pm}$. Its contractions are the $P^{\pm}_{e_a}$, and tangential output excludes a normal zero-order component. This proves the full recovery formula and its frame independence.

Metric compatibility makes $D_X N$ spatial and gives $\eta(D_X(\Theta Y),N)=K(X,Y)$. The normal component of $D\Theta=0$ is $K(X,Y)-K(Y,X)=0$. The tangential component defines a metric, torsion-free connection on $T\Sigma$, hence ∇^h. This proves the Gauss–Weingarten identities. Conversely they determine a connection on the splitting $\mathbb{E}=\mathbb{R}N\oplus\Theta(T\Sigma)$. Direct differentiation of the three pairings of N and $\Theta(T\Sigma)$ proves its metric compatibility; symmetry of K and torsion-freeness of ∇^h prove $D\Theta=0$. The same decomposition proves uniqueness. □

2.2. Vacuum recognition and the initial unit and counit. The curvature test has a unique algebraic output even at zero Weyl curvature. The initial recognition therefore uses no curvature-rank inverse.

Theorem 2.3 (Intrinsic vacuum recognition).
An independently admitted object of Grad_0^{∞} determines a datum of Data_0^{∞} by the metric and second-form readouts above and the actual spin reduction at N, transported by Θ. Its unique Weyl readout is

$$E=\mathrm{Ric}(h)+\tau K-Kh^{-1}K,\qquad B_{ab}=\epsilon_a{}^{cd}\nabla_c K_{db},\qquad \tau=\mathrm{tr}_h K.$$

Conversely, each datum d gives an independently admitted source. The two operations recover all initial tensors, the actual initial spin structure and the complete supplied operator pair under canonical identifications. No maximality, decay, smallness or curvature-rank condition is required for this recognition.

Proof. The preceding lemma supplies h,K,D with symmetric K. In an adapted frame put $\kappa_a=K_{ab}\theta^b$ and let ω be the Levi–Civita connection of h. The Lorentz connection has blocks $A^0{}_a=A^a{}_0=\kappa_a$ and $A^a{}_b=\omega^a{}_b$. Consequently

$$\Omega^a{}_b=(\Omega^h)^a{}_b+\kappa_a\wedge\kappa_b,$$
$$\Omega^0{}_b(e_c,e_d)=\nabla_c K_{db}-\nabla_d K_{cb}.$$

The positive sign in the quadratic term follows from the two equal mixed blocks in signature $(+---)$. Contracting gives exactly the displayed E,B. The matrix E is symmetric. The trace of B vanishes by symmetry of K, and

$$\mathrm{tr}\,E=R(h)+\tau^2-|K|_h^2,\qquad \epsilon_a{}^{bc}B_{bc}=\nabla^j K_{ja}-\nabla_a\tau.$$

The stated component formula for $C(E,B)$ gives a real algebraic Weyl tensor whenever E,B are symmetric and trace-free: alternation and pair symmetry define the other components. In an oriented orthonormal frame the cyclic sum

of $C_{0abc} = -\epsilon_{bc}{}^{d} B_{da}$ over a, b, c is $-\epsilon_{abc} \operatorname{tr} B$. The mixed Ricci contraction is, up to the contraction sign, $\sum_{b,d} \epsilon_{ab}{}^{d} B_{db}$, which vanishes exactly when B is symmetric. The purely spatial cyclic identity follows from the spatial component formula with symmetric E. The time-time Ricci contraction is a signed trace of E; in the spatial Ricci contraction the two E_{ab} terms cancel and only a multiple of $\delta_{ab} \operatorname{tr} E$ remains. Conversely, the contractions in the definition recover E, B from every algebraic Weyl tensor. These matrices determine all its components by the same formula. In particular the restriction W_Θ is injective and has the asserted finite inverse on its image. Equivalently, if a Weyl tensor vanishes on every spatial first pair, pair symmetry removes every component except possibly $C(N, v_a, N, v_b)$; its Ricci trace then removes those components as well. The uniqueness makes the reconstruction independent of the adapted frame.

Admission therefore gives symmetric trace-free E, B. The two identities above imply the vacuum constraints, with their stated signs. Reduction of the actual Lorentz spin bundle at N gives a Spin(3) lift of the spatial orthonormal frame bundle in $N^\perp$. Pulling it back through the oriented isometry $\Theta : (T\Sigma, h) \to (N^\perp, -\eta)$ gives P_Σ.

For the converse, given $d = (\Sigma, h, K, P_\Sigma)$, set

$$\begin{aligned} \mathbb{E}_d &= \mathbb{R} \oplus T\Sigma, \\ \eta_d((a, X), (b, Y)) &= ab - h(X, Y), \\ N_d &= (1, 0), \\ \Theta_d X &= (0, X), \\ D^d_X(a, Y) &= \big(Xa + K(X, Y),\ \nabla^h_X Y + aK^\sharp X\big). \end{aligned}$$

Extend P_Σ through $\mathrm{Spin}(3) \subset \mathrm{Spin}^+(1,3)$ and take its associated chiral modules and connections $\nabla^{d,\pm}$. Define

$$\mathbb{D}^d_\pm \psi = \sum_a \Theta_d(e_a) \otimes \nabla^{d,\pm}_{e_a} \psi.$$

The expression is invariant under an oriented orthonormal change of frame. Its symbol is $b_d(\xi) = \Theta_d(\xi^{\sharp_h})$, and its full output is spatial. Its directional readout is the actual spin connection of D^d, so it passes the alternating and real Clifford tests. The previous lemma gives torsion-freeness.

The constraints make the displayed E, B symmetric and trace-free. The mixed curvature block is the epsilon reconstruction of B. The spatial block is an algebraic three-dimensional curvature tensor with Ricci contraction E. Such a tensor is determined by that contraction: writing its components in the order $G_{cdba} = \epsilon_{cdm} \epsilon_{ban} S_{mn}$ identifies it with a symmetric matrix S, and the contraction is $S - (\operatorname{tr} S) I$, an invertible linear map since its trace is $-2 \operatorname{tr} S$. The spatial Weyl restriction of $C(E, B)$ has the same contraction, so it equals the Gauss block. Metric skewness determines the other mixed block. Thus $\Omega_{D^d} = \mathsf{W}_{\Theta_d}(C(E, B))$, establishing the independent image test.

Decoding this source recovers h, K exactly. Reduction of its extended spin bundle at N_d canonically recovers P_Σ. For a source followed by decoding and reconstruction, the vector comparison is

$$L_d : \mathbb{R} \oplus T\Sigma \longrightarrow \mathbb{E}, \qquad L_d(a, X) = aN + \Theta X.$$

It is an oriented and time-oriented isometry, and the Gauss–Weingarten formula identifies its transported connection with D. The source spin bundle is canonically the extension of its reduction at N: the map on associated principal bundles is $[p, g] \mapsto pg$. The corresponding Dirac counit is the evaluation map $[p, \psi] \mapsto p(\psi)$. It has its original actual spin label. The exact uncontracted recovery lemma now identifies the reconstructed pair with the supplied $\mathbb{D}_\pm$, including every zero-order coefficient. These are the inverse object identifications. □

The symbol metric retains physical scale. In particular a physical comparison must preserve h itself. On the maximal weighted chart, the source of Theorem 4.18 satisfies these intrinsic tests by its four commutator equations, and its intrinsic readout agrees with its canonical one. Restricting that constructed source to an initial open set restricts the fields and these local readouts. The whole weighted decoder of Section 3 is used before this restriction.

2.3. Physical comparisons on a fixed initial hypersurface. The comparison maps retain the actual spin action on spinor sections. We prove that the source intertwining equations determine precisely the physical comparisons of the decoded data.

Theorem 2.4 (Full physical Hom on the initial manifold). The intrinsic constructor and decoder extend to an equivalence

$$\mathsf{Grad}_0^\infty \simeq \mathsf{Data}_0^\infty$$

which is bijective on Hom sets over each specified initial diffeomorphism. Its unit and counit are natural exact comparisons and satisfy the triangle identities. The two central spin lifts on a connected initial domain remain distinct morphisms.

Proof. First, a multiplication-semilinear isomorphism of local section sheaves over f is a unique smooth bundle isomorphism over f. Indeed, identify the target with the pullback bundle on the source and write a section in a local frame. Semilinearity says that its image is the same linear combination of the images of those frame sections. These images are smooth, so they form the columns of a smooth matrix. The inverse sheaf map gives its inverse matrix. Frame changes give the bundle transformation law. This argument is local and requires no functional-analytic assumption on the section map.

The induced conjugation on real Clifford vectors defines a smooth bundle map L. The Clifford relation makes it a Lorentz isometry. The required algebraic orientations and preservation of N make it proper and orthochronous.

In local spin frames choose one spin lift of L, with Dirac action U_0. Any other Clifford intertwiner U differs by a matrix commuting with the entire complex Clifford algebra; hence $U_0^{-1}U = \lambda I$. Preservation of B gives $\lambda^2 = 1$. Thus the permitted intertwiners are precisely the two actual spin actions. Their local lifts glue because their Dirac actions are the given smooth U and the Dirac representation is faithful. The central element -1 acts as $-I$, so its two actions have not been identified.

Take scalar commutators of the exact gradient intertwining equation. At each x they give

$$b'_{f(x)}\big(((\mathrm{d}f_x)^{-1})^*\xi\big) = L_x b_x(\xi), \qquad \xi \in T^*_x\Sigma.$$

The metric readout then gives $f^*h' = h$ and $L\Theta X = \Theta' f_* X$. Contracting the full gradient identity with these soldered vectors gives

$$P'_{f_*X}\widehat{U} = \widehat{U}P_X, \qquad D'_{f_*X}(Lv) = L(D_X v).$$

The second equality follows by commuting the first with each Clifford vector. Preservation of N now gives $f^*K' = K$. The spin map restricts at N and, through Θ, becomes precisely a datum spin comparison. The unique Weyl readout is carried to its counterpart because the connection and its curvature are carried exactly.

Conversely, a datum comparison induces the oriented Lorentz isometry between the split bundles in the constructor and the actual extension of its supplied spin lift. The formula for D^d shows that it intertwines the full connections. Its uncontracted gradient formula then gives the exact source intertwining equations. Thus every datum comparison has a source comparison, and the two assignments on arrows are inverse.

The unit identifies the recovered h, K, P_Σ with the original datum. The counit is the vector map L_d and the associated principal-spin and spinor evaluation maps in the recognition theorem. Every physical comparison commutes with these evaluations, proving naturality. Reconstructing and then evaluating acts as the identity on each vector, spin frame and spinor; this gives the two triangle identities. It also verifies that the Hom bijection retains the supplied actual spin map on both passes. □

This Hom statement compares sources over a specified map of their initial hypersurfaces. Comparisons in which the initial hypersurface itself moves inside a spacetime require the additional local evolution and point-family structures.

2.4. Initial descent and finite local estimates.

Proposition 2.5 (Initial descent and local smooth dependence). Intrinsic vacuum gradient sources and their physical comparisons admit effective descent over open covers of a fixed initial manifold. The constructor, decoder, Weyl readout and their unit and counit commute with restriction and preserve smooth finite-dimensional families with supplied local algebraic and spin identifications.

More precisely, in retained coordinate and spin charts let $K \Subset \operatorname{int} K'$ be compact spatial sets and $\mathcal{K} \Subset \operatorname{int} \mathcal{K}'$ compact parameter sets. Impose a positive lower bound for the symbol metric and finite bounds for the inverse marking matrices on $K' \times \mathcal{K}'$. Every prescribed finite spatial and parameter norm of the readouts and comparison maps on $K \times \mathcal{K}$ is controlled by finitely many such norms of the supplied fields and comparison maps on $K' \times \mathcal{K}'$. For the object constructor and decoder, spatial input order $m + 2$ suffices for output order m, with constants depending on the stated finite coefficient bounds and inverse margins.

In particular, locally smooth convergence of admitted sources and their markings to a nondegenerate spacelike-symbol limit gives an admitted limit and locally smooth convergence of the decoded data and Weyl readout, including zero-curvature limits.

Proof. Let sources be supplied on $\{O_i\}$ with actual overlap module maps U_{ij} covering the identity and satisfying

$$U_{ij}U_{jk} = U_{ik}$$

on every triple overlap. The preceding theorem makes these smooth bundle maps and actual spin maps. The ordinary bundle cocycle glues the Dirac modules, Clifford-vector bundles, Lorentz and alternating forms, future section and actual spin reductions. The scalar commutator identity glues the symbol and soldering. Exact operator intertwining glues the two first-order operators as actions on the sheaves of sections.

All admission conditions are local. Their finite readouts agree under the same transitions, so the torsion and Weyl-image conditions hold for the glued object. Its Weyl field glues by the uniqueness of the algebraic preimage. Compatible morphisms glue as actual bundle maps and retain the exact intertwining equations. If two gluings have the prescribed local identifications, those identifications give compatible isomorphisms on every O_i and hence a unique global comparison. This is effective descent over the given manifold, including the exact spin signs on triples.

For the local estimates, work in finitely many retained spin frames. The inverse metric and Θ are finite matrix operations on the symbol and markings. The directional connection is obtained by contraction of the full supplied operator coefficients. Its vector connection follows by finite Clifford commutators, and the second form uses this connection and one derivative of N. Its curvature and the algebraic Weyl readout use one further spatial derivative. The positive inverse-matrix margins bound every denominator in these operations.

In the converse construction, an oriented local frame for h is obtained from a fixed coordinate frame by its positive square root. The connection uses one derivative of h and the value of K; its curvature uses two derivatives of h and one of K. The matrix square root and inverse are smooth on the retained positive set. Differentiating these finite formulas proves the stated spatial bound. Repeated passive-parameter differentiation obeys the finite product

and inverse-matrix rules and therefore uses finitely many input norms, with no additional symbol-smallness condition.

For a supplied comparison, semilinearity reads its local bundle matrix directly from its values on a frame. The vector map follows by Clifford conjugation; the actual spin map is its locally initialized lift through the smooth double covering. On a connected lift chart its sign is fixed. The unit and counit are the finite associated-bundle evaluations already described. The supplied comparison and the bounds on each lift chart give the same finite parameter estimates, with the sign fixed only locally.

Every local formula commutes with restriction. Their compositions with overlap maps agree by the Hom theorem, which proves compatibility of recognition and descent. For a convergent sequence, a retained compact set has a uniform positive symbol margin for all sufficiently late members. All finite operations then converge in each finite derivative norm. Torsion and curvature admission pass to the limit because they are finite differential identities and the Weyl-image test uses the explicit continuous algebraic reconstruction. It follows that the limit is admitted and that all the stated readouts converge, even where the curvature evaluation loses rank. □

The estimates in this proposition are estimates for finite initial differential constructions. The common spacetime domain and evolution estimates are supplied by the local Cauchy and collar results below.

2.5. The smooth Cauchy input and its local use. The geometric construction uses the smooth Einstein Cauchy theorem. The finite Cartan equations below identify the resulting source and propagate its constraints; existence is supplied by the hyperbolic Cauchy theory.

Theorem 2.6 (Local smooth Einstein Cauchy input).
Let (h, K) be smooth vacuum initial data on an open three-manifold, with h positive definite and a chosen future direction. There is a smooth Ricci-flat development germ along the initial hypersurface. Two developments of the same data have a unique germ of time-orientation-preserving isometry extending a fixed initial identification.

The external geometric input is the existence and common-development statement of Sbierski (2016, Theorem 2.6). His Lemma 3.1 and Corollary 3.2 give uniqueness of the comparison with fixed initial identification. We use the local germs of these conclusions.

For the reduced hyperbolic construction we use Hughes, Kato and Marsden (1977, Theorem III and its following regularity remark; Theorem V): the harmonic metric equation has a unique solution on one time interval for data in a fixed $H^{s+1} \times H^s$ neighborhood; the solution depends continuously on the data there, and additional smooth regularity persists on that interval. We choose an integer base index $s_0 \geq 5$ and retain a positive hyperbolicity margin.

Proof. We justify the local application. The vacuum constraints are exactly the smooth initial-data conditions in the cited geometric theorem.

Multiplying a metric of signature $(-+++)$ by -1 leaves its Levi–Civita connection and its Ricci tensor unchanged and gives our signature $(+---)$. With $K = \frac{1}{2}\partial_t h_t$ in Gaussian coordinates, the initial tensors have the stated convention. Restricting the common development gives local isometries. A time-oriented isometry fixing the initial hypersurface also fixes its normal, so the cited comparison uniqueness applies. On overlapping initial opens the unique germs agree. This proves the germ assertion; it does not require an open product collar itself to be a globally hyperbolic development.

Here is the localization used for estimates. In a coordinate ball choose $V \Subset W \Subset W_2 \Subset U$, where the initial data are supplied on U. Use a fixed cutoff to extend h from a neighborhood of $\overline{W_2}$ to a smooth positive metric on $\mathbb{R}^3$, equal to a fixed Euclidean metric outside a compact set. Convex interpolation of positive matrices gives such an extension with a positive lower bound. Extend K smoothly by the same fixed cutoff. For a family near one parameter, these extensions depend smoothly on that parameter and have finite estimates in every specified order.

For the reduced metric equation take, at coordinate time $t = 0$,

$$g_{00} = 1, \qquad g_{0i} = 0, \qquad g_{ij} = -h_{ij}, \qquad \partial_t g_{ij} = -2K_{ij}.$$

Choose the four remaining time derivatives to make $\mathcal{H}^\mu = g^{\alpha\beta}\Gamma^\mu_{\alpha\beta} = 0$ initially. Explicitly, with $\tau = \operatorname{tr}_h K$ and $\Gamma^k_h = h^{ij}\Gamma(h)^k{}_{ij}$, they are

$$\partial_t g_{00} = 2\tau, \qquad \partial_t g_{0i} = -h_{ik}\Gamma^k_h.$$

These are finite differential operations on (h, K). The reduced equations $\mathrm{Ric}(g) - \nabla_{(\mu}\mathcal{H}_{\nu)} = 0$ have principal part $-\frac{1}{2}g^{\alpha\beta}\partial_\alpha\partial_\beta g_{\mu\nu}$ and lower terms smooth in $g^{-1}, g, \partial g$. Subtract a fixed extended initial metric. The new unknown and its initial time derivative have the Sobolev behavior required by the hyperbolic input. The coefficients are smooth, their spatial derivatives are bounded, and the spatial principal matrices and time coefficient are uniformly definite after the sign convention is adjusted. At $t = 0$ this follows from the positive bound for h and the zero shift; it persists in one sufficiently small solution neighborhood. Thus the hypotheses of the cited reduced theorem hold.

The cutoff extension need not satisfy the vacuum constraints away from W_2. To see why this is sufficient, the contracted Bianchi identity for a reduced solution gives the homogeneous gauge equation

$$\nabla^\alpha\nabla_\alpha\mathcal{H}_\nu + R_\nu{}^\mu\mathcal{H}_\mu = 0.$$

Initially $\mathcal{H} = 0$. On W_2 its normal derivative is zero as well: the constraints give the normal components of the Einstein tensor equal to zero, while the reduced equation expresses these components, at $\mathcal{H} = 0$, as one half of the corresponding normal derivatives of $\mathcal{H}$. Tangential derivatives of its zero initial value vanish. Uniqueness for the linear gauge equation on the domain of dependence of W_2 therefore gives $\mathcal{H} = 0$ there, and the reduced metric is Ricci-flat on that domain.

Finite propagation follows from the wave energy identity on truncated cones whose lateral speed exceeds the characteristic speeds; their outward flux has the required sign. Uniform low-order metric bounds give one such speed bound. Hence a common smaller time interval and the fixed spatial buffers put the retained reduced-coordinate region, and subsequently the normal geodesics from W, inside this domain of dependence. This proves the required local applicability without an exterior extension of the vacuum constraints. □

2.6. Actual finite Cartan collars. The intrinsic initial source determines a soldering Θ, a real metric connection D and a unique algebraic Weyl tensor C_0. We now give its evolution an independent finite-field presentation. Its unknowns are a spacetime soldering, a Lorentz connection and an algebraic Weyl tensor. The frame formulation by these fields is classical (Friedrich and Nagy 1999, §2, equations (2.2)–(2.5)). The temporal gauge and subsidiary identities used here are proved below. Existence is supplied by the smooth local Cauchy theorem.

A collar is a neighborhood of the marked zero slice in a product with a time interval. We work with its germ along a retained initial patch, and may shrink that patch or the neighborhood. A product $(-T, T) \times V$ with spatial boundary is used as a coordinate domain; Cauchy uniqueness is applied on smaller domains of dependence inside a buffered initial patch. A common positive width is asserted on retained compact patches, not on an arbitrary noncompact initial manifold.

Let $\mathbb{G}$ be an independently admitted intrinsic vacuum gradient source on an initial open set U, as in the intrinsic definition. Write its readout as $(\mathbb{E}, \eta, N, \Theta, D, C_0)$, with its specified actual spin bundle. Extend these algebraic bundles by product in time and put $n(t, x) = N(x)$. In an adapted initial spin frame extended without time dependence, n has the constant components e_0. Changes between such frames are independent of time.

For an algebraic curvature tensor C, write $\mathcal{C}$ for its endomorphism form, with

$$\eta(\mathcal{C}(u, v)z, w) = C(u, v, z, w).$$

This uses the same last-slot convention as the intrinsic Weyl readout.

Definition 2.7 (An initialized finite Cartan collar).
An initialized finite Cartan collar over $\mathbb{G}$ consists of actual smooth fields

$$\theta = n\,\mathrm{d}t + \theta_i\,\mathrm{d}x^i, \qquad A = A_i\,\mathrm{d}x^i, \qquad C.$$

The spatial θ_i are an oriented basis of $n^\perp$, the A_i are the matrices of a real η-metric connection, and C is a real algebraic Weyl tensor. The connection on the product bundle is $\widehat{D}_0 = \partial_t$ and $\widehat{D}_i = \partial_i + A_i$, with its induced actual spin connections. Thus its temporal connection coefficient vanishes in the initialized product gauge. The following identities are required on the actual collar:

$$\partial_t \theta_i = A_i n, \qquad \partial_t A_i = \mathcal{C}(n, \theta_i),$$

$$\mathcal{B}_{0ij} := (\widehat{D}_0\mathcal{C})(\theta_i,\theta_j) + (\widehat{D}_i\mathcal{C})(\theta_j,n) + (\widehat{D}_j\mathcal{C})(n,\theta_i) = 0.$$

The connection acts on every tensor slot. In particular, for bundle sections u, v,

$$\widehat{D}_i^{\mathrm{End}}(\mathcal{C}(u,v)) = (\widehat{D}_i\mathcal{C})(u,v) + \mathcal{C}(\widehat{D}_i u, v) + \mathcal{C}(u, \widehat{D}_i v).$$

The initial values are prescribed by the source:

$$\theta_i(0) = \Theta_i, \qquad A_i(0) = D_i - \partial_i, \qquad C(0) = C_0.$$

All equations are finite first-order identities on supplied fields. Their definition uses the independent initial source tests and no spacetime metric. The Weyl field has ten real algebraic components, equivalently two Lorentz-related chiral Weyl spinors. Higher curvature derivatives may subsequently be read from the solution, but are not additional evolution variables in this definition.

The equations are invariant under time-independent initial coordinate and adapted spin-frame changes, so this definition patches on initial coordinate opens. The local collar germs are always taken with their initialized vector and actual spin identifications.

Proposition 2.8 (The exact finite subsidiary system).
Suppose fields of these algebraic types satisfy the two temporal equations. Define their spatial defects by

$$\begin{aligned}\mathcal{T}_{ij} &= \widehat{D}_i\theta_j - \widehat{D}_j\theta_i,\\ \Omega_{ij} &= \partial_i A_j - \partial_j A_i + [A_i, A_j],\\ \Delta_{ij} &= \Omega_{ij} - \mathcal{C}(\theta_i,\theta_j).\end{aligned}$$

Then

$$\partial_t\mathcal{T}_{ij} = \Delta_{ij}n, \qquad \partial_t\Delta_{ij} = -\mathcal{B}_{0ij} + \mathcal{C}(n,\mathcal{T}_{ij}).$$

At each fixed spatial point this is a linear system for thirty real defect components, with forcing $(0, -\mathcal{B})$. If L bounds its coefficient matrix in a fixed positive component norm on a retained compact cylinder, then

$$|(\mathcal{T},\Delta)(t,x)| \le e^{L|t|}\left(|(\mathcal{T},\Delta)(0,x)| + \int_{\min(0,t)}^{\max(0,t)} |\mathcal{B}(u,x)|\,\mathrm{d}u\right).$$

In particular, the actual mixed Bianchi identities and zero initial spatial defects imply zero spatial defects throughout the retained interval.

Proof. Differentiate the torsion defect and use the temporal equations:

$$\begin{aligned}\partial_t\mathcal{T}_{ij} = {}& \mathcal{C}(n,\theta_i)\theta_j - \mathcal{C}(n,\theta_j)\theta_i\\ &+ \widehat{D}_i(A_j n) - \widehat{D}_j(A_i n).\end{aligned}$$

The algebraic first Bianchi identity makes the first line $-\mathcal{C}(\theta_i,\theta_j)n$. Because the components of n are constant in the adapted product frame, the second line is $\Omega_{ij}n$. This proves the first identity.

For the connection curvature, direct differentiation gives

$$\partial_t\Omega_{ij} = \widehat{D}_i^{\mathrm{End}}(\mathcal{C}(n,\theta_j)) - \widehat{D}_j^{\mathrm{End}}(\mathcal{C}(n,\theta_i)).$$

Here $\widehat{D}_i n = A_i n = \partial_t \theta_i$. Expand each evaluated tensor using the derivative rule, and subtract

$$\partial_t(\mathcal{C}(\theta_i, \theta_j)) = (\widehat{D}_0\mathcal{C})(\theta_i, \theta_j) + \mathcal{C}(\partial_t\theta_i, \theta_j) + \mathcal{C}(\theta_i, \partial_t\theta_j).$$

The terms involving $\partial_t\theta$ cancel, by alternation of the curvature arguments. The remaining expression is

$$\begin{aligned}\partial_t\Delta_{ij} = (\widehat{D}_i\mathcal{C})(n, \theta_j) - (\widehat{D}_j\mathcal{C})(n, \theta_i) - (\widehat{D}_0\mathcal{C})(\theta_i, \theta_j)\\ + \mathcal{C}(n, \widehat{D}_i\theta_j - \widehat{D}_j\theta_i),\end{aligned}$$

which is the second asserted identity.

There are three spatial pairs. Each has a four-component torsion defect and a six-component Lorentz-skew curvature defect, giving thirty components in total. The coefficients of their homogeneous system are finite algebraic functions of the supplied C and fixed n. The integral equation and Gronwall's inequality give the displayed estimate for positive time. Reversing time gives its stated form for negative time. With zero forcing and zero initial value, uniqueness of this finite ODE gives the zero solution. □

The mixed Bianchi defect is the forcing term throughout the collar. The definition requires its vanishing there, a condition strictly stronger for smooth fields than vanishing of its normal jets on the initial slice.

2.7. Recognition and initialized comparison.

Theorem 2.9 (Recognition and existence of the initialized collar). An initialized finite Cartan collar reconstructs a smooth Lorentzian Ricci-flat metric with its specified initial data and actual spin marking. Under its soldering, the supplied connection is the full Levi–Civita connection and C is the actual Weyl curvature. The full differential Bianchi identity and its vacuum divergence equation hold.

Conversely, a smooth vacuum development of the initial data recognized from $\mathbb{G}$ has a Gaussian collar germ with exactly this initialized finite Cartan presentation. Its germ is unique after the initial vector and actual spin identifications are fixed. The existence and geometric uniqueness used here are those of the external smooth Cauchy input.

The initial value C_0 is determined by the initial source connection. It contributes no additional freely prescribed normal-curvature datum.

Proof. The intrinsic source recognition gives

$$D\Theta = 0, \qquad \Omega_D(X, Y) = \mathcal{C}_0(\Theta X, \Theta Y).$$

Thus the spatial defects in the preceding proposition vanish initially. The temporal equations already give

$$(\widehat{D}\theta)_{0i} = \partial_t\theta_i - A_i n = 0, \qquad \Omega_{0i} = \partial_t A_i = \mathcal{C}(n, \theta_i).$$

The actual mixed Bianchi equation and the subsidiary proposition now imply all the remaining components:

$$\widehat{D}\theta = 0, \qquad \Omega_{\widehat{D}} = \mathcal{C}(\theta, \theta).$$

The soldering is an isomorphism because $\theta(\partial_t) = n$ and the three θ_i are a basis of $n^\perp$. Define

$$g(X,Y) = \eta(\theta X, \theta Y), \qquad \nabla_X Y = \theta^{-1}\widehat{D}_X(\theta Y).$$

These are actual smooth fields. Metric compatibility of $\widehat{D}$ gives $\nabla g = 0$, and its torsion is $\theta^{-1}\widehat{D}\theta = 0$. Hence ∇ is the Levi–Civita connection. The full curvature identity identifies its curvature with C under the soldering. The algebraic Weyl conditions imply $\mathrm{Ric}(g) = 0$. Consequently C is the actual Weyl tensor, and the differential Bianchi identity for the Levi–Civita connection, followed by its Ricci-flat contraction, gives both remaining assertions about Bianchi.

The metric has the Gaussian form

$$g = \mathrm{d}t^2 - h_t, \qquad (h_t)_{ij} = -\eta(\theta_i, \theta_j).$$

Its future normal along the zero slice is ∂_t. Our second-form convention is $K(X,Y) = -g(\nabla_X \partial_t, Y)$ there. The temporal equation gives

$$\frac{1}{2}\partial_t (h_t)_{ij}\Big|_{t=0} = -\frac{1}{2}\left(\eta(D_i N, \Theta_j) + \eta(\Theta_i, D_j N)\right) = K_{ij}.$$

The last equality uses the symmetry of K already obtained from initial torsion. Thus the initial metric, full second form and orientations are exactly the source readouts. The soldering transports the product actual spin bundle to the oriented and time-oriented orthonormal frame bundle of g, with the prescribed initial reduction. On a smaller domain of dependence this is the required Cauchy development germ.

For the assertion about C_0, suppose an algebraic Weyl tensor has zero endomorphism restriction when its first two arguments are spatial. Pair symmetry then also annihilates components with spatial last pair. Only components with index pattern $0a0b$ could remain. Their Ricci trace-free equations express them as spatial contractions, which already vanish. Therefore the tensor is zero. The initial spatial curvature and Θ determine at most one Weyl tensor; the intrinsic admission test supplies its existence. This proves the claimed uniqueness of C_0 directly.

For existence, decode $\mathbb{G}$ by intrinsic initial recognition and apply the smooth Cauchy input to its vacuum data. Near a retained initial patch, the normal exponential map gives Gaussian coordinates after shrinking to avoid a focal point and to retain an injective coordinate map. The initial Lorentz identification is

$$J_0(aN + \Theta X) = a\nu + \mathrm{d}\iota(X),$$

where ι is the marked initial embedding and ν its future unit normal. Extend J_0 by Levi–Civita parallel transport along the normal geodesics. The pullback soldering and connection then have $\theta(\partial_t) = n$ and vanishing temporal connection coefficient. Transport the product actual spin bundle through this vector identification: its frame projection is the original product frame projection followed by J. This equips the collar with the required actual

spin structure and its initial reduction. If the development already carries an actual spin extension with the prescribed initial identification, path lifting along the normal geodesics gives the unique compatible spin-bundle identification. In either description the induced spin connection has vanishing temporal coefficient.

Put C equal to the actual Weyl tensor in this identification. Torsion-freeness gives the first temporal equation, and the curvature identity gives the second. The actual differential Bianchi identity gives the mixed equation. At zero, the full initial tangent connection agrees with D by the Gauss–Weingarten formula and equality of h, K. Its curvature therefore has the same spatial restriction, so the unique Weyl readout just proved gives $C(0) = C_0$. Every initial value is exactly the stipulated one.

For uniqueness, take two initialized Cartan germs over $\mathbb{G}$. Their recognized metrics induce the same vacuum data on a buffered initial open set. Restrict to common Cauchy neighborhoods of a smaller retained patch and apply geometric Cauchy uniqueness. The resulting isometry fixes the initial embedding and carries its future normal to the future normal. It therefore carries each normally parametrized geodesic to its counterpart, with the same proper-time parameter. In the Gaussian coordinates its point map is the identity on a smaller collar germ. The two metrics agree there.

Their initialized parallel identifications solve the same finite parallel-transport equation for this metric, so the solderings agree. Levi–Civita uniqueness then identifies the connections, and curvature identifies C. The two actual spin identifications agree as lifts of the same initialized frame identification: their difference is a locally constant central sign, fixed on each initial component. This proves uniqueness of the entire initialized source germ. □

The smooth Einstein Cauchy theorem supplies a collar realizing the initial data. Conversely, the Cartan identities recognize the vacuum metric of a supplied collar. The homogeneous subsidiary system propagates vanishing of the spatial torsion and curvature defects from the initial slice when the mixed Bianchi identity holds throughout the collar.

A comparison between initialized Cartan collars consists of a local point diffeomorphism F carrying the marked initial hypersurface to its counterpart, an oriented and time-oriented Lorentz bundle isometry L over F, and an actual compatible spin lift $\widetilde{L}$. Require the full identities

$$\theta' \circ \mathrm{d}F = L \circ \theta,$$
$$\widehat{D}'_{F_*X}(Lv) = L(\widehat{D}_X v),$$
$$\mathcal{C}'(Lu, Lv)Lw = L\big(\mathcal{C}(u, v)w\big).$$

The fields on the right and left are evaluated at corresponding points. These conditions are imposed directly on supplied fields. Their restriction is an intrinsic initial gradient comparison, with its actual spin map. Germs and comparisons are restricted to a common collar whenever they are composed.

Theorem 2.10 (Full comparison for a marked initial hypersurface).

Restriction gives a bijection between actual Cartan collar comparisons and intrinsic initial gradient comparisons over each specified initial diffeomorphism. Every initial comparison extends uniquely. These assignments preserve identities, inverse maps, composition and restriction.

Reconstruction also identifies these comparisons with all orientation-preserving and time-orientation-preserving spin isometries of the recognized vacuum germs that carry their marked initial hypersurfaces to one another. Thus initialized finite Cartan collar germs, intrinsic initial gradient sources and their marked vacuum development germs form equivalent groupoids.

Proof. A Cartan comparison preserves $g = \eta(\theta, \theta)$ by the soldering identity. It therefore induces an isometry of the recognized metrics. On the initial slice its connection and soldering identities recover the initial gradient intertwining relation by the exact intrinsic decoder. The stipulated spin lift restricts to the actual initial lift, so restriction is well defined.

Conversely, let an intrinsic source comparison cover $f : U \to U'$, with induced Lorentz map L_0 and actual spin map $\widetilde{L}_0$. By initial full Hom, it is exactly a physical comparison of the decoded h, K and initial spin structures. Regard the target metric as a development of the source data using the initial embedding $\iota' f$. Geometric Cauchy uniqueness, on the buffered domains just used, supplies an isometry germ F extending this identification. It preserves the future normal. Normal-geodesic uniqueness gives

$$F(t, x) = (t, f(x))$$

in the two Gaussian presentations. Parallel transport of L_0 and $\widetilde{L}_0$ along these geodesics gives the full vector and spin comparisons.

In the initialized product frames the transported maps are independent of time. Their complete coefficient identities are

$$\theta'_j(t, f(x))\, \partial_i f^j = L_0(x)\theta_i(t, x),$$
$$A'_j(t, f(x))\, \partial_i f^j = L_0 A_i(t, x) L_0^{-1} - (\partial_i L_0) L_0^{-1}.$$

The first follows from the soldering and normally parallel identifications; the second is the full connection transformation law, including the derivative of the comparison. Curvature gives the displayed $\mathcal{C}$ identity. The induced actual spin connection transforms by the corresponding principal spin law, with the same prescribed $\widetilde{L}_0$. These formulas prove that the extension satisfies every source comparison condition.

An isometry carrying the marked initial hypersurfaces and preserving time orientation must carry their future normal geodesics with their proper-time parametrizations. Hence its point germ is already determined by its initial map. The full soldering identity then determines its Lorentz bundle map. Two compatible spin lifts differ by a locally constant central sign, and their initial lift fixes that sign on every component. This proves uniqueness and

faithfulness. It also shows that every spin isometry of marked vacuum germs has precisely the source comparison just constructed.

The extension of an identity is the identity. The composite of two extensions has the prescribed composite initial comparison, so uniqueness makes it the extension of that composite; the same argument gives inverses and compatibility with restriction. The source-to-metric and metric-to-source object constructions are inverse up to the initialized identifications in the recognition theorem. The data-to-source unit is the intrinsic unit, and the collar counit is the initialized normal-geodesic and parallel-spin identification. They commute with every comparison by uniqueness; their composites are the identity on the corresponding vector and spinor bundles. This proves the equivalences and their triangle identities. Distinct initial central spin lifts remain distinct throughout. □

The marked hypersurface is part of this Hom statement. Comparison between displays based on different hypersurfaces will use the later unmarked source comparison theorem.

2.8. Ordinary displays and physical scale.

Proposition 2.11 (Physical scaling of the actual collar).
Let $r > 0$ be constant on a given collar and put $\psi_r(t', x) = (t'/r, x)$. On the corresponding rescaled time domain define

$$\theta^{(r)} = r\,\psi_r^*\theta, \qquad A^{(r)} = \psi_r^* A, \qquad C^{(r)} = r^{-2}\psi_r^* C,$$

with the pulled-back actual spin bundle and its initialized marking. These fields satisfy the same finite Cartan equations. Their recognized metric and initial data are

$$g^{(r)} = r^2\psi_r^* g, \qquad h^{(r)} = r^2 h, \qquad K^{(r)} = rK.$$

With the fixed internal Clifford metric and unit normal, the initial chiral gradients obey

$$\mathbb{D}_{\pm}^{(r)} = r^{-1}\mathbb{D}_{\pm}.$$

For an ordinary display, first restrict the already constructed complete initial contact presentation to its retained patch. In the fixed cotangent-ray coordinates of Proposition 4.15, its physical-scale coefficients satisfy

$$\beta^{(r)} = \beta, \qquad \eta_c^{(r)} = \eta_c, \qquad \omega^{(r)} = \omega + 3\log r, \qquad \nu^{(r)} = r^3\nu.$$

The local coefficient laws retain the maximal marking $\tau = 0$. The whole-space normalization relative to δ applies before this change of scale and need not hold for $r^2 h$.

The construction is compatible with restriction, initialized actual spin comparisons and successive positive scale changes.

Proof. In the new clock, $\theta^{(r)} = n\,\mathrm{d}t' + r\theta_i(t'/r, x)\,\mathrm{d}x^i$, and $A_i^{(r)} = A_i(t'/r, x)$. Therefore

$$\partial_{t'}\theta_i^{(r)} = A_i^{(r)} n, \qquad \partial_{t'} A_i^{(r)} = r^{-1}\mathcal{C}(n, \theta_i) = \mathcal{C}^{(r)}(n, \theta_i^{(r)}).$$

In the mixed Bianchi equation the time derivative of $C^{(r)}$ contributes r^{-3} and its two spatial soldering arguments contribute r^2. Each spatial derivative term has factor r^{-2} from C and factor r from its spatial argument. Thus the entire mixed defect is $r^{-1}\psi_r^*\mathcal{B}$ and vanishes. Weyl symmetries, reality and the initial spatial identities are preserved.

The metric formula follows by applying η to the rescaled soldering. Differentiating $h^{(r)}_{t'} = r^2 h_{t'/r}$ gives the displayed initial K law. At zero, the connection itself is unchanged in the fixed internal frame, while an orthonormal tangent vector changes from e_a to $r^{-1}e_a$ and the soldering changes from Θ to $r\Theta$. Consequently

$$\sum_a \Theta^{(r)}(r^{-1}e_a) \otimes P^{\pm}_{r^{-1}e_a} = r^{-1}\sum_a \Theta(e_a) \otimes P^{\pm}_{e_a},$$

which proves the gradient law.

For the local contact formulas, $A = h(z,z)$, $B = h(z,\bar{z})$ and Δ_c each acquire a factor r^2, so β and $\ell = z - \beta\bar{z}$ are unchanged. Also ρ_n acquires r^{-1}, whereas K acquires r and the Christoffel symbols are unchanged. The complete shear formula therefore gives $\eta_c^{(r)} = \eta_c$. Finally $\det(r^2 h) = r^6 \det h$ gives the stated ω and density laws. The initial mean curvature remains zero.

Pullback carries the supplied spin identification in the same internal frames. The even tensor coefficients alone cannot distinguish its central sign. The formulas compose under multiplication of positive scales and commute with the initialized comparisons, proving the remaining claims. □

Corollary 2.12 (The actual collar of an ordinary-pair display). Let c belong to the ordinary domain in Theorem 4.18, with its fixed physical and actual spin markings. Its constructed initial gradient source has an initialized finite Cartan collar germ. All initial collar values are read from that source, and the recognized metric is a smooth Ricci-flat development of its exactly recovered data.

Conversely, suppose a supplied initialized Cartan collar has an initial source satisfying the independent canonical tests and retained-domain conditions of Theorem 4.18. The ordinary class

$$c = \sigma_\gamma \mathscr{P}_\gamma^{-1}\pi\mathbb{G}$$

reconstructs that initial source exactly. Its constructed collar is identified with the supplied collar by the unique comparison inducing the identity on the initial source and its actual spin bundle.

Effective ordinary cochain comparisons give these initialized collar comparisons. The physical scale r is retained by the preceding proposition. On a smaller initial open set, the collar construction uses restriction of the already constructed source; it does not apply the whole weighted ordinary decoder to a spatially truncated field.

Proof. Theorem 4.18 gives an actual initial source with independent kinematic and vacuum admission, exact ordinary-class recovery, and the

specified spin marking. Its positive-frame presentation is a special case of the intrinsic source definition. The intrinsic recognition theorem provides Θ, D, C_0 and the vacuum data. The collar recognition theorem then gives an actual initialized Cartan germ. Its only initial values are those finite readouts; smooth Cauchy existence supplies the geometric evolution from them.

For the converse, Theorem 4.18 applies to the initial source and the displayed decoder. It reconstructs the same complete initial operator pair, or equivalently its complete canonical contact presentation with scale and spin. The uniqueness part of collar recognition identifies the entire initialized finite-field germ. Initial full Hom and its extension theorem give the same conclusion for the effective ordinary comparisons. In particular the initial ordinary-class unit persists after passing to the collar.

The scale formulas apply directly to these actual fields. Restriction of the finite Cartan equations and their initialized comparison is local, and uniqueness identifies the restricted construction. The inverse in the displayed ordinary formula remains the whole weighted inverse on its stated canonical domain. This proves all assertions. □

This construction supplies the actual geometric collar used by a local display. The full wave section and its point connection, and the comparison of displays with different initial hypersurfaces, retain their separate admission and comparison statements.

2.9. Higher estimates for the same reduced solution. Write $\mathcal{P} \subset \mathbb{R}^d$ for a parameter domain and use fixed local coefficient trivializations. A smooth family means joint smoothness of its supplied coefficients in parameter and spatial variables. All compact parameter estimates below use $\mathcal{K} \Subset \operatorname{int} \mathcal{K}' \Subset \mathcal{P}$. The input seminorm $\mathcal{N}_N$ is the maximum of a finite collection of spatial and parameter derivatives up to order N on the stated larger compact set, including the supplied coordinate and frame markings.

Lemma 2.13 (One reduced solution family at every finite order). For a smooth local family of initial data near any fixed parameter value, the localized reduced construction in the Cauchy input has one parameter neighborhood and one positive time interval chosen at the finite base order. On every fixed buffered compact part of that interval, the selected solutions are jointly smooth. Every finite spacetime and parameter seminorm is bounded by a finite locally bounded function of finitely many input seminorms and the inverses of the retained hyperbolicity and domain-buffer margins.

For two such families in the same fixed gauge and cutoff construction, write d, d' for their full reduced Cauchy data. On a common retained compact domain, for each m, q there are a finite $N = N(m, q)$ and a finite locally bounded function $G_{m,q}$ such that

$$\|g - g'\|_{C^q_\lambda C^m_{t,x}} \le G_{m,q}\big(\mathcal{N}_N(d), \mathcal{N}_N(d'), \mathfrak{m}^{-1}\big)\, \mathcal{N}_N(d - d').$$

Here all input norms use the larger compact parameter and spatial sets, and $\mathfrak{m}$ records the retained positive margins. At each output order the estimate uses a finite, possibly higher, input order.

Proof. Fix the cutoff and the reference extension. The initial reduced fields depend continuously on the parameter in the base $H^{s_0+1} \times H^{s_0}$ norm. Shrink its neighborhood once so that they lie in one existence neighborhood of the external theorem. Its common time interval and a smaller retained interval are then fixed. The base solution remains in a set with uniform positive hyperbolicity and bounded low-order coefficients.

For completeness, the estimates behind the higher and parameter assertions can be obtained on this same solution. For a linear equation with the reduced wave principal part, let

$$\mathcal{L}_g v = g^{\alpha\beta}\partial_\alpha\partial_\beta v + B^\alpha \partial_\alpha v + Dv = f,$$

where the matrices B^α, D have the required finite coefficient bounds. For an integer $n \geq 0$, a spatially differentiated energy is equivalent to $\|v(t)\|_{H^{n+1}} + \|\partial_t v(t)\|_{H^n}$. One may use the square root of

$$\sum_{|\beta|\leq n} \int_{\mathbb{R}^3} \left(g^{00}|\partial_t\partial_x^\beta v|^2 - g^{ij}\partial_i\partial_x^\beta v \cdot \partial_j\partial_x^\beta v + |\partial_x^\beta v|^2 \right) \mathrm{d}x.$$

In the retained coordinates both quadratic forms are positive. Multiplication of the commuted equation by $\partial_t\partial_x^\beta v$, followed by integration by parts, gives

$$E_n(v;t) \leq C_n \left(E_n(v;0) + \int_0^{|t|} \|f(\pm u)\|_{H^n}\, \mathrm{d}u \right).$$

The constant uses the fixed time interval, positive margins and finitely many coefficient norms. A spatial cutoff or cone version gives the same estimate on retained nested domains. The terms involving g^{0i} are flux terms and terms involving its first derivatives. The remaining commutators are bounded by the product rule, Sobolev multiplication and interpolation. Smooth approximation justifies the estimate for solutions at the stated finite regularity.

For the nonlinear reduced equation put $u = g - g_{\mathrm{ref}}$, where g_{ref} is the fixed extended initial reference. Its principal coefficients depend on g and not on ∂g. In the commuted equation the term with the largest derivative of a coefficient is a derivative of g multiplied by a bounded low-order derivative of g; the other products admit the same tame bound. Thus, on the fixed low-order neighborhood, the high energy satisfies an inequality of the form

$$E_n(u;t) \leq C_n \left(E_n(u;0) + \int_0^{|t|} (1 + E_n(u;\pm u'))\, \mathrm{d}u' \right),$$

with fixed-reference terms included in C_n. Applying the estimate to regularized initial data, taking the base limit by uniqueness, and using weak compactness and the estimate proves the bound for the already selected solution. The usual difference estimate then gives strong continuity at each

lower retained order. Equivalently, the persistence clause of the external input identifies the high-regularity solution with this same base solution. Gronwall's inequality supplies each finite bound on the same time interval. No high-order smallness is introduced.

To justify parameter differentiation, differentiate the reduced equation formally once and move all terms linear in the highest parameter derivative to the left. For $v = \partial_{\lambda^a} g$ its principal part is $\mathcal{L}_g$; the variation of $g^{\alpha\beta}$ times $\partial_{\alpha\beta} g$ is a zeroth-order matrix acting on v. The variation of the nonlinear first-derivative terms contributes first- and zeroth-order matrices. Their finite norms are controlled by higher spatial bounds just proved.

This formal equation identifies an actual derivative. Difference quotients $v_\varepsilon = (g_{\lambda+\varepsilon e_a} - g_\lambda)/\varepsilon$ satisfy linear equations with averaged coefficients and uniformly bounded finite data. The linear energy bound controls them at any fixed desired order after finitely many stronger input norms have been retained. Base continuity and uniform higher spatial bounds imply coefficient convergence in every required finite lower norm by Sobolev interpolation. Subtract the limiting linear equation from the quotient equation. The remaining initial data and forcing tend to zero, so the same energy estimate makes the difference tend to zero. Existence for the linear equation follows by regularization and this energy estimate, or by the linear case of the stated hyperbolic input. Thus v is the actual parameter derivative.

For a multi-index α, the equation for $\partial_\lambda^\alpha g$ has the same linear principal operator. Its forcing is a finite sum of products involving only lower parameter derivatives, each with at most two additional spacetime derivatives. Induction on $|\alpha|$, allowing these finite spatial derivative costs, proves existence, continuity and finite estimates for every specified mixed derivative. Its difference-quotient justification is the same. Compact containment of $\mathcal{K}$ in the interior of $\mathcal{K}'$ keeps all quotients inside the supplied parameter set.

Time derivatives are recovered successively from the equation because g^{00} has a positive lower bound. Sobolev embedding then gives every specified mixed C^m estimate on a retained compact region. For two families, subtraction of their equations and repeated parameter differentiation give linear difference equations with coefficients controlled by both families and forcing proportional to finitely many input differences. The same argument proves the final assertion. Every step consumes finitely many stronger norms; the base interval was fixed before their orders. □

2.10. Common Gaussian collars and their finite-family bounds.

Theorem 2.14 (Common collars for smooth finite-dimensional families).

Let a smooth finite-dimensional family of independently admitted intrinsic initial gradient sources be supplied on common coordinate and spin representatives. Near each parameter value and over each retained compact initial patch, there are one parameter neighborhood, a positive T, and buffered

spatial parents such that their recognized finite Cartan collars are defined on a common parent $(-2T, 2T) \times W$.

The metric, soldering, connection, actual Weyl field, initialized normal-geodesic map and initialized actual spin parallel transport are jointly smooth there. For every $m, q < \infty$, their coefficient seminorms on $\mathcal{K} \times [-T, T] \times \overline{V}$, with $V \Subset W$, satisfy

$$\begin{aligned} &\|\mathcal{Y}\|_{m,q} + \|\Phi\|_{m,q} + \|\Phi^{-1}\|_{m,q} \\ &\quad \le F_{m,q}\left(\mathcal{N}_N(\text{initial source and markings}; \mathcal{K}' \times \overline{U}),\, \mathfrak{m}^{-1}\right) \end{aligned}$$

for some finite N and a finite locally bounded function $F_{m,q}$. The map and inverse norms use their respective retained source and image coordinate buffers. Here $\mathcal{Y}$ is any fixed finite list of the preceding coefficient fields and transports. The margins $\mathfrak{m}$ include coordinate and symbol nondegeneracy, hyperbolicity, spacelikeness, domain buffers and the retained Gaussian inverse margin. The parameter neighborhood and parent are chosen before m, q.

A supplied smooth family of initial physical comparisons with actual spin lifts extends to a smooth family of the marked collar comparisons, after one common buffered restriction. It has the same kind of finite estimates, with the supplied comparison, inverse and spin-marking seminorms included. A supplied smooth positive scale is allowed, with its finite derivatives and a positive lower bound included.

Proof. Intrinsic recognition reads (h, K) and the actual initial spin structure by finite differential and matrix operations. On each retained compact set, symbol nondegeneracy gives one positive margin. These operations therefore preserve the stipulated family and supply the finite bounds for its local Cauchy data. Apply the preceding lemma to the fixed localized reduced gauge. On the smaller domain of dependence its metric is Ricci-flat by the gauge-propagation argument.

Solve the normal geodesic equation for this actual reduced metric, with initial position $(0, x)$ and the future unit initial normal. On the chosen initial gauge this velocity is ∂_t. Its coefficients use the first metric derivatives, and its derivatives obey finite variational ODEs. Low-order bounds and the spatial buffer give a common interval for all retained x and parameters. At time zero the spacetime differential of the map

$$\Phi_\lambda(t, x) = \exp^{g_\lambda}_{(0,x)}(t\, n_\lambda(0, x))$$

is invertible. Shrinking the time interval once gives a uniform inverse margin. In sufficiently small convex coordinate patches its differential remains uniformly close to the initialized one; the mean-value estimate then gives injectivity. For a general retained compact patch, take a finite cover by such smaller patches with buffered parents. By its Lebesgue number, sufficiently close initial base points belong to one injective parent. Pairs of base points separated by at least that number remain separated for small time, since $\Phi_\lambda(t, x)$ tends uniformly to $(0, x)$ as $t \to 0$. Compactness makes this choice uniform over the retained parameters. Thus the common restricted collar

map is injective. On overlaps the maps describe the same initialized normal geodesics, so they give the same collar germ.

Retain a larger parent than the final compact observation set before differentiating. The ODE variation equations give finite mixed bounds for Φ. The inverse margin and the spatial buffers also give common image buffers. Differentiating $\Phi^{-1} \circ \Phi = \mathrm{id}$ gives the inverse bounds from that margin. Pullback gives the Gaussian metric. Parallel transport along these normal geodesics identifies the tangent Lorentz bundle with its product extension and gives the temporal gauge $A_0 = 0$. Principal spin parallel transport with the supplied initial identification gives its actual spin lift. These finite-dimensional ODEs have unique solutions, and their variations obey linear ODEs with forcing involving finitely many earlier derivatives. Their finite estimates follow successively from Gronwall's inequality. The initialized spin lift retains its initial central sign.

In that identification, θ is the metric soldering, A the Levi–Civita connection and C its Weyl curvature. Reading them uses finitely many derivatives of the already constructed fields. Cartan recognition identifies this presentation with the independently admitted finite Cartan collar. This proves joint smoothness and the displayed bound. All higher estimates use the same low-order-selected reduced solution, geodesic domain and parallel-transport parent.

For a supplied initial comparison $(f_\lambda, \widetilde{f}_\lambda)$, marked collar comparison has Gaussian point map

$$(t, x) \longmapsto (t, f_\lambda(x)).$$

In the reduced-coordinate representatives it is $\Phi'_\lambda \circ (\mathrm{id}, f_\lambda) \circ \Phi_\lambda^{-1}$. If P_λ and P'_λ denote the initialized principal spin transports, its actual spin map is

$$\widetilde{F}_\lambda(t, x) = P'_\lambda(t, f_\lambda(x))\, \widetilde{f}_\lambda(x)\, P_\lambda(t, x)^{-1}.$$

The supplied initial images and inverse margins permit one common restriction of the parents. Composition, inversion on this retained nondegenerate domain and the finite ODE estimates give the asserted comparison bounds. Uniqueness gives their identity and composition laws. The comparison is chosen by the supplied initial physical and spin data.

Finally apply physical scale transport. For a parameter-dependent $r > 0$, its physical clock is $t' = rt$. On a retained compact parameter set, a positive lower bound for r gives one common physical time parent. The chain rule for t'/r, the factors r in the soldering and r^{-2} in the internal Weyl tensor use only finitely many derivatives of r and its reciprocal at each order. The connection and initialized spin transport obey the same pullback law. This proves the scaled family assertion with the stated finite dependence. □

The theorem applies in particular to the ordinary-pair initial sources of Theorem 4.18, by composition with that theorem's finite-family estimates. Its present output is the finite geometric collar and its marked comparisons. The full wave kernel, evaluation section and all-direction point connection have their own subsequent constructions. On a noncompact initial manifold

these conclusions are local over retained compact patches; their time widths need not have a positive global lower bound.

2.11. Exterior, gauge and scale choices on one initial germ. The ordinary input is a whole pair, while its physical display retains only a local source. We now compare the initial and finite Cartan parts of two such displays directly. Their comparison conditions are the source identities already defined, and do not require equality of the whole pairs.

Corollary 2.15 (Geometric choice independence on a fixed initial germ).

Let $d = (O, h, K, P_O)$ be a physical maximal vacuum initial germ. For $a = 0, 1$, let c_a be an ordinary pair in the selected domain of Theorem 4.18. Execute the pair on its whole domain and restrict its actual initial source to V_a. Supply a positive scale r_a, constant on that spacetime display, and a physical embedding $\iota_a : V_a \to O$, with an actual spin identification Ψ_a over it, such that

$$\iota_a^* h = r_a^2 h(c_a)|_{V_a}, \qquad \iota_a^* K = r_a K(c_a)|_{V_a}.$$

On a common germ in the two images, the scaled initial gradients and initialized Cartan collars have a unique physical comparison inducing the common initial identification and its specified actual spin map. Its initial point and spin maps are

$$f_{01} = \iota_1^{-1}\iota_0, \qquad \widetilde{U}_{01} = \Psi_1^{-1}\Psi_0.$$

The initial complete contact graph and physical density are also carried by the induced positive cotangent-ray map. These geometric comparisons have exact restriction, identity, inverse and cocycle laws. They include changes of exterior extension, based spatial gauge, coordinate or frame normalization, cutoff and scale that preserve the displayed physical initial germ.

For smoothly supplied finite-dimensional families of such displays and embeddings, the comparisons are smooth locally in the parameter on a common retained parent. Each prescribed finite derivative norm is bounded by finitely many norms of the supplied constructions and embeddings, with positive inverse and buffer margins. The positive scales may depend on the passive parameters. This assertion uses the initialized local family choices; it does not select isometries or spin signs over an arbitrary parameter space.

Proof. Put $\mathbb{G}_a^{(r_a)} = r_a^{-1}\mathbb{G}_a$ for the two uncontracted gradient operators, keeping their algebraic Clifford markings and unit normals. Physical scale transport gives their readouts as

$$h_a^{\mathrm{phys}} = r_a^2 h(c_a), \qquad K_a^{\mathrm{phys}} = r_a K(c_a).$$

The supplied embeddings imply $f_{01}^* h_1^{\mathrm{phys}} = h_0^{\mathrm{phys}}$ and $f_{01}^* K_1^{\mathrm{phys}} = K_0^{\mathrm{phys}}$. Together with $\widetilde{U}_{01}$ these are a physical initial-data comparison. Initial full Hom gives its exact source intertwiner. In particular, if $U_{01,\pm}$ are its chiral

actions and L_{01} its Lorentz vector map, then

$$(r_1^{-1}\mathbb{D}_{1,\pm})\widehat{U}_{01,\pm} = L_{01}\widehat{\otimes}\,\widehat{U_{01,\pm}}(r_0^{-1}\mathbb{D}_{0,\pm}).$$

This is precisely the independently defined operator comparison, including its zero-order coefficients. The marking determines the actual spin action; passage to even tensor coefficients has not identified its two central lifts.

Apply marked collar full Hom after restoring the physical clocks by the scale proposition. It extends the initial comparison uniquely to the finite Cartan fields. In their physical Gaussian coordinates its point map is

$$F_{01}(t,x) = (t, f_{01}(x)).$$

The vector and spin maps are the initialized normal parallel transports of the initial ones. Thus the extension preserves the physical clock, soldering, full connection and actual Weyl field. This applies on buffered Cauchy neighborhoods and hence to the retained collar germs.

For the initial contact assertion, the induced map is

$$\widehat{f}_{01}(x,[\xi]) = \left(f_{01}(x), [((\mathrm{d}f_{01,x})^{-1})^*\xi]\right).$$

The cotangent lift preserves the tautological form before positive projectivization, so it preserves the projective contact distribution. In the intrinsic description of Proposition 4.15, the graph is defined from the oriented metric screen, its cross product, the Levi–Civita horizontal splitting and the second form. An oriented h-isometry preserves the screen and cross product, and its derivative carries the Levi–Civita splitting to the target splitting. Preservation of K then carries the shear term as well. It therefore intertwines the whole graph, including the terms arising from the derivative of the coordinate transformation. The density comparison is $f_{01}^*\,\mathrm{vol}_{h_1^{\mathrm{phys}}} = \mathrm{vol}_{h_0^{\mathrm{phys}}}$. The contact coefficient formulas are used after restriction and physical scaling; no claim is made that the scaled local tensors remain in the original whole weighted chart.

For an explicit exterior example, let two whole extensions (H_a, L_a) coincide with the same physical data on O and let their based gauges be ϕ_a. On the common local images take $f_{01} = \phi_1^{-1}\phi_0$. Then

$$f_{01}^*(\phi_1^*H_1, \phi_1^*L_1) = (\phi_0^*H_0, \phi_0^*L_0).$$

Initialized actual spin lifts give the stipulated spin comparison. Additional local coordinate, frame and scale choices are already included by the embeddings and scale factors above. This proves the stated applicability without comparing whole Green kernels or recomputing a whole decoder after truncation.

For three displays, the actual formulas give $f_{12}f_{01} = f_{02}$ and $\widetilde{U}_{12}\widetilde{U}_{01} = \widetilde{U}_{02}$. Initial full Hom preserves these compositions, and uniqueness of the initialized Cartan extension gives the same equality on their common collars. The cotangent-ray maps compose by the ordinary chain rule. The same argument proves identities, inverses and compatibility with restriction.

For families, first retain compact spatial and parameter sets inside the supplied larger parents. The inverse-function bounds for ι_a, the positive scale bounds and the supplied spin identifications control each finite derivative of f_{01} and $\widetilde{U}_{01}$. The finite initial estimates and the common-collar family theorem then control the induced fields and comparisons on one retained parent. The contact comparison is the displayed cotangent formula and the preceding finite tensor construction, so it has the same finite-order dependence. This proves the family assertion. □

For the initial and finite Cartan parts, the local inverse retains the source-comparison class of an ordinary display. The two whole pairs can differ. Also, changing the radius used to select a maximal disk can change the initial hypersurface itself. Such a change is outside this fixed-hypersurface result unless its physical initial germs actually agree. Full wave comparisons and changes of initial hypersurface will be treated with their additional source structures.

3. Ordinary classes and the initial-data construction

Each local display starts from two ordinary classes in the smooth coefficient complexes on the fixed flat Cauchy source. Their root values give linear coordinates for the free transverse-traceless seeds. A selected nonlinear constraint correction then gives whole maximal vacuum data, with an exact inverse on an independently specified tensor domain.

Three distinctions determine the construction. The linear root complex is a polynomial quotient, whereas nonlinear coefficient operations use actual representatives. The actual cohomology has polynomial classes that the root map does not see; explicit obstruction tests select the ordinary domain. Finally, the smallness condition is imposed at one finite base order, and all higher weighted estimates concern that same selected solution.

The weighted analytic proofs are given in Appendix A. Theorem 3.13 supplies the ordinary input and tensor-decoder foundation. Section 4 constructs the actual initial source from those classes by normalized negative responses and current input transport. Section 5 uses this initial source in the local coverage construction, with its stated Cauchy and wave-completion inputs.

The classical linear Penrose transform relates twistor cohomology to massless fields (Eastwood, Penrose and Wells). The normalization of the two contour residues is that of the original analytic correspondence (Isthmus 2026a, Theorem 3.1 and Appendix C.2).

3.1. Root coefficients on the flat Cauchy source. The root calculation supplies linear coordinates for the weighted free seeds. We first compute it in a coefficient complex modulo spatial polynomials. This quotient is used only for the linear transform. Actual ordinary classes and their permitted primitives are defined in the separate actual coefficient complex; products of arbitrary coefficients do not descend through the polynomial quotient.

Fix the Euclidean orientation, scale and spin marking on $\mathbb{R}^3$. Fourier transformation uses

$$\widehat{f}(\xi) = (2\pi)^{-3/2} \int e^{-ix\cdot\xi} f(x)\,\mathrm{d}x, \qquad f(x) = (2\pi)^{-3/2} \int e^{ix\cdot\xi} \widehat{f}(\xi)\,\mathrm{d}\xi.$$

On $N_0 = \mathbb{R}^3_x \times \mathbb{CP}^1$, write the homogeneous conic as

$$z(s,t) = (s^2 - t^2,\ i(s^2+t^2),\ 2st), \qquad z \cdot z = 0.$$

In the affine chart $s = 1$, the antiholomorphic distribution is generated by $\partial_{\bar\zeta}$ and $W = z(1,\zeta)\cdot\partial_x$. The conic line is $\mathcal{O}(-2)$, with its retained spin square root $\mathcal{O}(-1)$. Evaluation of a weight-m one-form on W therefore has weight $m+2$.

For $\xi \neq 0$, put $r = |\xi|$, $n = \xi/r$ and

$$q_\xi(s,t) = \xi \cdot z(s,t), \qquad L_\xi = iq_\xi.$$

The two zeros of q_ξ are simple. Indeed, $q_\xi = 0$ means that ξ is perpendicular to both the real and imaginary parts of the null conic vector. The corresponding unit normal to this oriented two-plane is n or $-n$. Thus its divisor D_ξ consists of the antipodal roots $n_\varepsilon = \varepsilon n$, $\varepsilon = \pm 1$. Let

$$(\mathcal{R}_m)_\xi = \mathcal{O}(m+2)|_{D_\xi} = \mathcal{O}(m+2)|_{n_+} \oplus \mathcal{O}(m+2)|_{n_-}.$$

This is a radially pulled-back rank-two bundle over punctured frequency space. Its individual root lines are not trivialized by a polarization frame.

Definition 3.1 (Moderate coefficients and the polynomial quotient).

For a smooth bundle V on $\mathbb{CP}^1$, let $\mathscr{A}_M(V)$, $M \geq 0$, consist of smooth sections on $(\mathbb{R}^3 \setminus \{0\}) \times \mathbb{CP}^1$ satisfying, in finitely many bundle charts,

$$|\partial_\xi^\alpha \nabla_\zeta^\nu a(\xi,\zeta)| \leq C_{\alpha,\nu} |\xi|^{-M-|\alpha|}, \qquad 0 < |\xi| \leq 1,$$
$$|\partial_\xi^\alpha \nabla_\zeta^\nu a(\xi,\zeta)| \leq C_{\alpha,\nu,N} (1+|\xi|)^{-N}, \quad |\xi| \geq 1,$$

for every finite derivative order and every N. One exponent M works for all fiber derivatives. Write $\mathscr{A}(V) = \bigcup_{M\geq 0} \mathscr{A}_M(V)$. Each fixed-$M$ space has its displayed seminorms; continuity between these spaces is allowed to replace M by one larger finite exponent. For bundles pulled radially from the frequency sphere, the same definition uses finite angular trivializations.

For a bundle V over $\mathbb{CP}^1$, let $\widetilde{\mathscr{A}}(V)$ consist of actual smooth sections $f(x,\zeta)$ with polynomial growth in x together with all derivatives, whose distributional Fourier restriction away from zero belongs to $\mathscr{A}(V)$. At each fixed spatial differential order, a single polynomial growth exponent must work for all fiber derivatives. Let $\mathscr{P}(V)$ be the finite-degree polynomials in x with smooth V-valued coefficients, and set

$$\mathscr{C}(V) = \widetilde{\mathscr{A}}(V)/\mathscr{P}(V).$$

The weight-m complex has terms

$$\begin{aligned}\mathscr{C}^0_m &= \mathscr{C}(\mathcal{O}(m)),\\ \mathscr{C}^1_m &= \mathscr{C}(\Lambda^{0,1}\otimes\mathcal{O}(m))\oplus\mathscr{C}(\mathcal{O}(m+2)),\\ \mathscr{C}^2_m &= \mathscr{C}(\Lambda^{0,1}\otimes\mathcal{O}(m+2)),\end{aligned}$$

and differential

$$d^0_m u = (\bar\partial u, Wu), \qquad d^1_m(\alpha,\beta) = \bar\partial\beta - W\alpha.$$

Here $\bar\partial$ differentiates only in the projective fiber. The two generators commute because z is holomorphic. The operators preserve spatial polynomials, so the differential is defined on the quotient. A lift of a closed quotient cochain has polynomial differential. Closedness of the lift requires that polynomial to vanish.

Lemma 3.2 (Fourier restriction and the polynomial ambiguity). Fourier restriction induces an isomorphism

$$\mathscr{C}(V) \longrightarrow \mathscr{A}(V).$$

It identifies the differential with

$$d^0_m u = (\bar\partial u, L_\xi u), \qquad d^1_m(\alpha,\beta) = \bar\partial\beta - L_\xi\alpha.$$

Every punctured coefficient has an actual smooth spatial representative of polynomial growth. The ambiguity between its tempered extensions is precisely a polynomial spatial section. On a fixed amplitude order these constructions have bounds in finitely many seminorms, with a finite choice of distributional order.

Proof. Choose a scalar $\chi \in C^\infty_c(\mathbb{R}^3)$ equal to one near zero and an integer $L \geq 0$ with $L + 4 > M$. For $a \in \mathscr{A}_M(V)$ define, in a bundle chart,

$$\langle E_L a, \varphi\rangle = \int a(\xi)\left[\varphi(\xi) - \chi(\xi)\sum_{|\mu|\leq L}\frac{\partial^\mu\varphi(0)}{\mu!}\xi^\mu\right]\mathrm{d}\xi.$$

The bracket is $O(|\xi|^{L+1})$ near zero. Including the radial measure, its absolute integral is bounded there by a constant times $\int_0^1 r^{L+3-M}\,\mathrm{d}r$, which converges. At infinity rapid decrease suffices. This gives a tempered extension with order bounded in terms of M, L and finitely many amplitude seminorms. The same argument applies after every fiber derivative without increasing L, because the fixed exponent M controls them all. Since the cutoff and Taylor operation act only in ξ, the construction commutes with the bundle transition maps and gives a smooth family of distributions over the fiber.

Split this distribution with a second frequency cutoff. The part away from zero is Schwartz, with smooth fiber dependence. The part near zero has compact support and finite distributional order. Its inverse Fourier transform is an actual smooth function: differentiating in x inserts powers of ξ in the compactly supported distribution's action on the exponential. Each such derivative is bounded by a polynomial in x. The same order works for fiber

derivatives at each fixed spatial order. Thus the inverse transform belongs to $\widetilde{\mathscr{A}}(V)$.

Two tempered extensions differ by a distribution supported at zero. If its order is at most L', it depends only on the Taylor polynomial of a test function through degree L'. To verify this assertion, subtract that Taylor polynomial near zero. The remainder and all derivatives through order L' vanish there; multiplication by a cutoff with shrinking support and the distributional estimate make its value tend to zero. Consequently the distribution is a finite linear combination of $\partial_\xi^\mu \delta_0$. Smooth fiber dependence gives smooth coefficients, and the uniform distributional order gives one finite degree. Its inverse Fourier transform is a polynomial section.

Conversely, every polynomial has Fourier transform supported at zero. If a section of $\widetilde{\mathscr{A}}(V)$ has zero punctured Fourier restriction, the same point-support argument makes it a polynomial. This proves injectivity and surjectivity. Fourier transformation changes ∂_{x^j} to $i\xi_j$ and commutes with fiber differentiation. The differential and its polynomial ambiguity therefore have exactly the stated form. The estimates in this construction use finitely many seminorms at each requested order. □

Theorem 3.3 (The two-root cohomology calculation).
For $m = 2, -6$, evaluation of the horizontal component at D_ξ gives

$$\mathscr{H}^1_m := H^1(\mathscr{C}^\bullet_m) \xrightarrow{\sim} \mathscr{A}(\mathcal{R}_m), \qquad [(\alpha, \beta)] \longmapsto \beta|_{D_\xi}, \qquad \mathscr{H}^0_m = \mathscr{H}^2_m = 0.$$

There are linear normal representatives and primitives continuous between fixed coefficient orders, allowing a finite shift of M. They commute with smooth finite-dimensional parameter differentiation on such fixed orders. A Fourier-flat Schwartz amplitude is a Schwartz section all of whose frequency derivatives vanish at zero. The restriction to those amplitudes agrees with their root cohomology and embeds it injectively.

Proof. Work in punctured frequency space using the preceding lemma. Fix the Fubini–Study metric and invariant Hermitian metrics on $\mathcal{O}(k)$. Let H^j_k denote harmonic projection in Dolbeault degree j, and let $T_k = \bar\partial^* G_k$ on one-forms. The fiber identities are

$$\bar\partial T_k = 1 - H^1_k, \qquad T_k \bar\partial = 1 - H^0_k.$$

These are fixed continuous maps on smooth sections of the compact fiber. Their estimates use finitely many stronger fiber seminorms. They commute with spatial spin rotations for the chosen invariant metrics. The usual cohomology of $\mathcal{O}(k)$ on $\mathbb{CP}^1$ gives $H^1(\mathcal{O}(2)) = H^1(\mathcal{O}(4)) = 0$, $H^0(\mathcal{O}(-6)) = H^0(\mathcal{O}(-4)) = 0$, and dimensions $3, 5$ for $H^0(\mathcal{O}(2)), H^0(\mathcal{O}(4))$ and $5, 3$ for $H^1(\mathcal{O}(-6)), H^1(\mathcal{O}(-4))$.

For $m = 2$, subtract $d^0(T_2\alpha)$ from a closed pair. Its first component becomes zero and its second becomes the holomorphic section $\beta_0 = \beta - L_\xi T_2 \alpha$. Multiplication

$$L_\xi : H^0(\mathcal{O}(2)) \longrightarrow H^0(\mathcal{O}(4))$$

is injective because q_ξ is nonzero. Put $L_\xi^+ = (L_\xi^* L_\xi)^{-1} L_\xi^*$. Subtracting the remaining holomorphic coboundary gives the unique normal representative

$$(0, \beta_+), \qquad \beta_+ = (1 - L_\xi L_\xi^+)(\beta - L_\xi T_2 \alpha) \in \ker L_\xi^*.$$

If this expression vanishes, an explicit primitive is $T_2\alpha + L_\xi^+(\beta - L_\xi T_2\alpha)$.

For $m = -6$, put

$$U_6 = H^1(\mathcal{O}(-6)), \quad U_4 = H^1(\mathcal{O}(-4)), \quad A_\xi = H^1_{-4} L_\xi|_{U_6}.$$

After subtracting $d^0(T_{-6}\alpha)$, the first component is $\alpha_0 = H^1_{-6}\alpha$. Closedness gives $A_\xi \alpha_0 = 0$. Since $H^0(\mathcal{O}(-4)) = 0$, the second component is uniquely $T_{-4} L_\xi \alpha_0$. The normal representative is therefore

$$(\alpha_0, T_{-4} L_\xi \alpha_0), \qquad \alpha_0 \in \ker A_\xi.$$

It is exact only when $\alpha_0 = 0$, in which case $T_{-6}\alpha$ is a primitive of the original pair. The dual of A_ξ is, up to the fixed nonzero pairing factor, multiplication by q_ξ from $H^0(\mathcal{O}(2))$ to $H^0(\mathcal{O}(4))$. That map is injective, so A_ξ is surjective and has a two-dimensional kernel.

All finite-dimensional inverses used here have the required estimates. At $\xi = rn$, $L_\xi = rL_n$ and $A_\xi = rA_n$. Their nonzero singular values have a positive lower bound for $r = 1$, by constant rank and compactness of the unit sphere. Thus L_ξ^+ and $A_\xi^*(A_\xi A_\xi^*)^{-1}$ are smooth homogeneous multipliers of degree -1; the kernel and complement projections have degree zero. A smooth multiplier of any fixed degree changes a small-frequency exponent by only a finite amount, and each Cartesian derivative has the corresponding additional power of r^{-1}. At infinity it preserves rapid decrease. The normal forms and primitives consequently preserve $\mathscr{A}$, with the claimed fixed-order continuity.

The divisor sequence

$$0 \longrightarrow \mathcal{O}(m) \xrightarrow{L_\xi} \mathcal{O}(m+2) \longrightarrow \mathcal{O}(m+2)|_{D_\xi} \longrightarrow 0$$

identifies the positive cokernel and the negative kernel above with the two root fibers. For the negative weight the connecting class is represented by $L_\xi^{-1}\bar\partial\beta$, which is α for a closed pair. Thus in the total complex the common root map is exactly $\beta|_{D_\xi}$. The root-to-positive-normal-form map has degree zero. In the negative normal form the map to α_0 has degree -1, and the map to its horizontal component $T_{-4}L_\xi\alpha_0$ has degree zero. These maps are smooth in the direction and satisfy the same multiplier estimates.

One can see surjectivity without an implicit choice of a sheaf representative. Extend prescribed root values to a smooth section β_b of $\mathcal{O}(m+2)$, holomorphic near each root. Choose uniformly separated caps around the antipodal roots in finitely many angular charts, use holomorphic bundle frames within the caps, and patch by a partition of unity in the frequency direction. Put $\alpha_b = L_\xi^{-1}\bar\partial\beta_b$. The numerator is zero near the roots, so this division is smooth on the whole projective line. Division costs one power of r^{-1}; all direction derivatives have finite uniform bounds. The pair is closed and

has the stipulated root values. Its inverse Fourier transform is an actual smooth cochain representative; closedness at this stage is closedness in the polynomial quotient.

A degree-zero cocycle is fiber holomorphic and annihilated by the nonzero q_ξ, hence zero. For degree two, the positive-weight right inverse is $b \mapsto (0, T_4 b)$. For the negative weight, given a degree-two form b, choose

$$\alpha = -A_\xi^*(A_\xi A_\xi^*)^{-1} H^1_{-4} b, \qquad \beta = T_{-4}(b + L_\xi \alpha).$$

Then $b + L_\xi \alpha$ has zero harmonic part, and $d^1(\alpha, \beta) = b$. These operations have the already verified estimates, proving the remaining cohomology assertions.

All displayed operators act continuously on each specified coefficient order, with a finite order shift. They therefore commute with finite-dimensional parameter differentiation and give finite seminorm bounds for each specified derivative. On Fourier-flat amplitudes the same homogeneous operations preserve flatness at zero and Schwartz decay. A Fourier-flat class whose root values vanish has the primitive just constructed in that Fourier-flat space itself. This proves the final injectivity and agreement assertion. □

3.2. Decay and the zero-frequency normalization. For a real exponent $\beta < 0$, let $\mathscr{S}^\beta$ consist of actual smooth fields with seminorms

$$\sup_x \langle x \rangle^{-\beta + |\alpha|} |\partial_x^\alpha f(x)| < \infty$$

at every finite order. Let $\mathscr{B}_\beta$ consist of smooth functions on punctured frequency space, rapidly decreasing at infinity together with all derivatives, and satisfying

$$|\partial_\xi^\alpha a(\xi)| \le C_\alpha |\xi|^{-\beta - 3 - |\alpha|}, \qquad 0 < |\xi| \le 1.$$

These definitions apply componentwise to tensors. The amplitude is locally integrable at zero because $\beta < 0$. Its specified locally integrable extension is used whenever an inverse Fourier integral is taken.

Lemma 3.4 (Weighted Fourier transformation and its inverse estimate).

For every $\beta < 0$, the ordinary inverse Fourier integral defines a continuous map $\mathscr{B}_\beta \to \mathscr{S}^\beta$. For $-3 < \beta < 0$, Fourier transformation and this inverse give continuous inverse maps

$$\mathscr{S}^\beta \xrightarrow{\sim} \mathscr{B}_\beta.$$

In the two-sided assertion, the distributional Fourier transform is precisely the locally integrable extension of its punctured restriction, with no additional origin-supported distribution. Reality, tensor symmetries and the transverse-traceless conditions are preserved. For any fixed $0 < \alpha < 1$, $\mathscr{S}^\beta = \bigcap_{j \ge 0} C^{j,\alpha}_\beta$ with equivalent families of seminorms, using the weighted Hölder norms of Definition A.1. The one-way inverse estimate for $\beta \le -3$ does not assert this two-sided Fourier isomorphism outside $(-3, 0)$.

Proof. First take $-3 < \beta < 0$ and $f \in \mathscr{S}'^{\beta}$. A dyadic spatial partition decomposes its exterior part into $f_j(x) = 2^{j\beta} g_j(2^{-j}x)$, $j \geq 0$, where g_j is supported in one fixed annulus and has uniformly bounded derivatives of every prescribed finite order. Integration by parts gives

$$|\partial_\xi^\mu \widehat{f}_j(\xi)| \leq C_{\mu,N} 2^{j(\beta+3+|\mu|)} (1 + 2^j|\xi|)^{-N}.$$

For $0 < r = |\xi| \leq 1$, split at $2^j r = 1$. The sum with $2^j r \leq 1$ is bounded by $Cr^{-\beta-3-|\mu|}$ because $\beta + 3 + |\mu| > 0$. Choose $N > \beta + 3 + |\mu|$ for the other geometric tail; it has the same bound. If $r \geq 1$, this second estimate applies to every annulus and gives arbitrary rapid decrease. The central compactly supported term is Schwartz. The differentiated series converges uniformly on compact subsets away from zero, so it gives a smooth punctured coefficient in $\mathscr{B}_\beta$. The original dyadic series converges to f as a tempered distribution, identifying that coefficient with its distributional Fourier restriction.

For the inverse direction, now let any $\beta < 0$ and $a \in \mathscr{B}_\beta$. Separate a Schwartz high-frequency part and annular pieces supported where $|\xi| \asymp 2^{-j}$. On a fixed annulus in $\eta = 2^j \xi$, the rescaled amplitudes $2^{-j(\beta+3)} a(2^{-j}\eta)$, including their cutoffs, have bounded derivatives of every prescribed finite order. The inverse Fourier contribution of such a piece satisfies

$$|\partial_x^\mu f_j(x)| \leq C_{\mu,N} 2^{j(\beta-|\mu|)} (1 + 2^{-j}|x|)^{-N}.$$

For $|x| \geq 1$, split at $2^{-j}|x| = 1$. In the small-frequency tail, $\beta - |\mu| < 0$ gives a sum bounded by $C|x|^{\beta-|\mu|}$. For the other part choose $N > |\mu| - \beta$; its finite geometric sum has the same bound. On bounded sets the tail converges with every spatial derivative. The high-frequency part is Schwartz. This proves membership in $\mathscr{S}'^{\beta}$. Each target seminorm uses finitely many amplitude seminorms, and the same is true for the preceding forward estimate.

All $\xi^\mu a$ are integrable near zero and at infinity. Hence the inverse integral and its derivatives are actual functions and the inverse tends to zero at infinity. For an original $f \in \mathscr{S}'^{\beta}$ with $-3 < \beta < 0$, the difference between its distributional Fourier transform and the integrable extension just obtained is supported at zero. The Fourier restriction lemma makes its inverse a polynomial. Both f and the inverse integral tend to zero, so this polynomial is zero. This proves the two inverse identities and the absence of an added zero-frequency mode.

Conjugation and reversal of ξ give the real Fourier relation. The tensor symmetry conditions commute with Fourier transformation. If $\xi_i a_{ij} = 0$ off zero, the same identity holds for the locally integrable extension, hence its inverse satisfies $\partial_i f_{ij} = 0$. Trace is handled in the same way. Conversely, transforming those actual identities gives the punctured tensor conditions.

On each rescaled spatial annulus, the Hölder seminorm of a derivative is bounded by its supremum and the supremum of its next derivative on a slightly larger annulus. The symbol estimates thus give every weighted Hölder seminorm using finitely many stronger symbol seminorms. The pointwise

part of the weighted Hölder norm gives the converse. This proves the final seminorm equivalence. □

The lower bound $-3 < \beta$ was used in the forward spatial-annulus sum. It was not used in the inverse-frequency sum. In particular, the one-way estimate applies to a specifically constructed amplitude of order $\beta = \gamma - 3$; it does not assign that Fourier behavior to every field with the same faster spatial decay.

3.3. Calibrated root tensors and the weighted seed transform. For $\xi \neq 0$, let $P_n = 1 - n \otimes n$ and

$$\mathcal{T}_\xi = \{Q \in \mathrm{Sym}^2(\mathbb{C}^3) : Qn = 0,\ \mathrm{tr}\, Q = 0\},$$
$$\Pi_n Q = P_n Q P_n - \tfrac{1}{2}\,\mathrm{tr}(P_n Q P_n) P_n.$$

Put $j_n v = n \times v$ and $\mathcal{J}_n Q = j_n Q$ on $\mathcal{T}_\xi$. In an oriented orthonormal basis of $n^\perp$, a tracefree symmetric matrix anticommutes with j_n. Hence $j_n Q$ is again symmetric and tracefree, $\mathcal{J}_n^2 = -1$, and its two eigenspaces are lines. Define

$$\mathcal{L}_\sigma(\xi) = \ker(\mathcal{J}_n - \sigma i), \qquad \Pi_\sigma^{\mathrm{hel}}(\xi) = \tfrac{1}{2}(1 - \sigma i \mathcal{J}_n)\Pi_n, \quad \sigma = \pm 1.$$

These are smooth homogeneous degree-zero projectors. Their parity is

$$\Pi_\sigma^{\mathrm{hel}}(-\xi)\overline{Q} = \overline{\Pi_\sigma^{\mathrm{hel}}(\xi) Q},$$

because $j_{-n} = -j_n$. The projector formulas define the helicity lines intrinsically, without a choice of polarization vectors.

At n_ε, the conic vector has j_n eigenvalue $-\varepsilon i$. Its tensor square identifies $\mathcal{O}(-4)|_{n_\varepsilon}$ with $\mathcal{L}_{-\varepsilon}(\xi)$. The complex bilinear pairing $\mathrm{tr}(QU)$ is nondegenerate between opposite helicity lines and zero on each line with itself. It therefore identifies the dual root line $\mathcal{O}(4)|_{n_\varepsilon}$ with $\mathcal{L}_\varepsilon(\xi)$. Denote these degree-zero maps by $\iota_{-6,\varepsilon}$ and $\iota_{2,\varepsilon}$.

For a paired root class $b = (b_+, b_-) \in \mathscr{H}^1_2 \oplus \mathscr{H}^1_{-6}$, let

$$U_\varepsilon^+ = \iota_{2,\varepsilon}(b_+|_{n_\varepsilon}), \qquad U_\varepsilon^- = \iota_{-6,\varepsilon}(b_-|_{n_\varepsilon}).$$

The superscript labels the projective weight, the subscript labels the root branch, and σ labels helicity. These labels have different roles. The ordinary pair in the actual coefficient cohomology is denoted by c; its root readout will be $\rho(c) = b$.

We retain the original residue calibration. With these line identifications its metric-amplitude factors are

$$k_\varepsilon^+(r) = 4i\varepsilon r, \qquad k_\varepsilon^-(r) = \frac{i\varepsilon}{4r^3}.$$

Their normalization can be checked by one nonzero-frequency plane, using the original contour convention (Isthmus 2026a, §3, Theorem 3.1). Use the Hermitian soldering

$$X^{AA'} = \frac{1}{\sqrt{2}} \begin{pmatrix} t - x_3 & x_1 - i x_2 \\ x_1 + i x_2 & t + x_3 \end{pmatrix}, \qquad u = X^{00'},\ v = X^{01'},\ q = X^{10'}.$$

For $F_m = Ae^{\lambda w^0}/\zeta$, the ordered north-south Čech differential and a north cutoff ρ_N give the smooth CR representative

$$\alpha = Ae^{\lambda u}\zeta^{-1}\bar{\partial}\rho_N, \qquad \beta = -\sqrt{2}\lambda Ae^{\lambda u}\rho_N.$$

Indeed subtract the differential of $\rho_N G$, where $G = Ae^{\lambda u}(e^{\lambda v\zeta} - 1)/\zeta$, from the usual cutoff representative. The identities $Wu = -\sqrt{2}\zeta$, $Wv = \sqrt{2}$, and $WG = \sqrt{2}\lambda Ae^{\lambda u}$ give the displayed pair. At the north root its coefficient is $a = -\sqrt{2}\lambda A$. Put $\lambda = -i\sqrt{2}\varepsilon r$; the plane has spatial frequency $\xi = \varepsilon r e_3$ and time factor $e^{-i\varepsilon rt}$, with $a = 2i\varepsilon rA$.

The positive and negative contours have highest curvature components $A\lambda^4 e^{\lambda u}$ and $Ae^{\lambda u}$. In the original last-slot curvature convention they are realized by metric variations $-2A\lambda^2 e^{\lambda u}(\mathrm{d}v)^2$ and $-2A\lambda^{-2}e^{\lambda u}(\mathrm{d}q)^2$, since $\Psi^+_{0000} = -C_{uvuv}$ and $\Psi^-_{0'0'0'0'} = -C_{uquq}$ (Isthmus 2026a, §2 and Appendix C.2). At the north root $z = (1, i, 0)$, and

$$\iota_{-6}(a) = a\, z \otimes z, \qquad \iota_2(a) = \tfrac{a}{4}\,\overline{z} \otimes \overline{z},$$

because $(z\cdot\overline{z})^2 = 4$. The corresponding spatial metric amplitudes are $A\lambda^2\overline{z}\otimes\overline{z}$ and $A\lambda^{-2}z\otimes z$. Substitution gives precisely the two factors in the calibration formula. Spin rotation equivariance carries this computation to every root, with the same scalar factors. The plane fixes the local normalization; no plane wave is asserted to belong to the decaying weighted coefficient domain. The structured nonlinear analytic comparison has its own hypotheses and is separate from this flat normalization.

Definition 3.5 (The admitted weighted root pair).
Fix $-1 < \gamma < -1/2$. For the root tensors of b, put

$$A_\varepsilon(\xi) = 4i\varepsilon r\, U^+_\varepsilon(\xi) + \frac{i\varepsilon}{4r^3}U^-_\varepsilon(\xi),$$

$$Q_b = A_+ + A_-, \qquad P_b = -\frac{ir}{2}(A_+ - A_-).$$

The real space $\mathscr{W}_\gamma$ consists of paired root classes satisfying

$$U^+_-(\xi) = \frac{1}{16r^4}\overline{U^-_+(-\xi)}, \qquad U^-_-(\xi) = 16r^4\overline{U^+_+(-\xi)},$$

and

$$Q_b \in \mathscr{B}_\gamma, \qquad P_b \in \mathscr{B}_{\gamma-1}.$$

Its seminorms are those of these two amplitude spaces. The tests use only the linear root map, fixed multipliers, conjugation and amplitude estimates, before the nonlinear constraint chart is applied. Let $\mathcal{X}^\infty_\gamma$ be the real transverse-traceless pairs (q, p) with $q \in \mathscr{S}^\gamma$ and $p \in \mathscr{S}^{\gamma-1}$, equipped with all their weighted seminorms. In particular these fields belong to every finite weighted Hölder space required by the nonlinear chart.

Theorem 3.6 (The exact weighted root-to-seed transform).
The map

$$\mathscr{P}_\gamma : \mathscr{W}_\gamma \longrightarrow \mathcal{X}^\infty_\gamma, \qquad b \longmapsto (\mathcal{F}^{-1}Q_b, \mathcal{F}^{-1}P_b)$$

is a continuous real-linear isomorphism. The inverse transforms are convergent integrals with the specified zero-frequency normalization. The exact inverse is obtained by setting

$$A_\varepsilon = \tfrac{1}{2}\widehat{q} + \frac{i\varepsilon}{r}\widehat{p}$$

and then

$$U_\varepsilon^+ = \frac{1}{4i\varepsilon r}\Pi_\varepsilon^{\mathrm{hel}} A_\varepsilon, \qquad U_\varepsilon^- = \frac{4r^3}{i\varepsilon}\Pi_{-\varepsilon}^{\mathrm{hel}} A_\varepsilon,$$

followed by the inverse root-line maps and the root cohomology theorem. On the Fourier-flat subspace it is the calibrated residue Cauchy transform with the same constants. On smooth finite-dimensional families it commutes with every parameter derivative, and each specified finite output seminorm is bounded by finitely many input seminorms on the same retained parameter domain.

Proof. The projective root lines have precisely the helicity types displayed above, so Q_b and P_b are symmetric, tracefree and transverse away from zero. The coupled reality equations imply

$$A_-(\xi) = \overline{A_+(-\xi)}.$$

For example, the negative-helicity part of the left side is $-4irU_-^+(\xi) = -i\overline{U_+^-(-\xi)}/(4r^3)$, which is the corresponding component of $\overline{A_+(-\xi)}$. The other line is identical with the factors exchanged. It follows that $Q_b(-\xi) = \overline{Q_b(\xi)}$ and $P_b(-\xi) = \overline{P_b(\xi)}$. Both weights γ and $\gamma - 1$ lie in $(-3, 0)$. The weighted Fourier lemma therefore puts the inverse transforms in the stated real smooth weighted TT spaces, with no polynomial addition.

Conversely, take $(q, p) \in \mathcal{X}_\gamma^\infty$ and its unique Fourier amplitudes Q, P. The first inverse formula solves $Q = A_+ + A_-$ and $P = -ir(A_+ - A_-)/2$. The two helicity projectors split each A_ε into its two line components, and the second formula divides by the nonzero calibrated factors on those components. All line maps and projectors are smooth homogeneous degree-zero maps; the remaining factors have finite homogeneous orders. Thus the four raw root tensors belong to the moderate amplitude spaces. The root cohomology theorem gives a unique pair b with these values.

Reality of Q, P gives $A_-(\xi) = \overline{A_+(-\xi)}$. Apply the parity of the projectors and divide by the calibrated factors. This yields exactly the coupled residue equations, so $b \in \mathscr{W}_\gamma$. The recovered amplitudes are the original Q, P; hence the admission estimates are unchanged. The two compositions are identities because the root evaluation is an isomorphism and each displayed amplitude operation has the explicit inverse just used.

The amplitude seminorms and the weighted Fourier lemma give continuity in both directions. The inverse root formulas and the Hodge normal forms also bound each requested representative seminorm by finitely many seed seminorms, with a fixed finite small-frequency exponent shift. All these maps

are linear and continuous on the stated spaces. Difference quotients therefore show that they commute with finite-dimensional parameter derivatives; applying the same finite seminorm estimate to each derivative proves the family assertion. No smallness is needed for this linear theorem. On the Fourier-flat subspace the root evaluation, line maps and calibrated factors are the same formulas, proving exact agreement there. □

The coupled equations are also the fixed-point equations of the antilinear involution

$$(\mathfrak{c}_{\mathrm{res}}U)^+_\varepsilon(\xi) = \frac{1}{16r^4}\overline{U^-_{-\varepsilon}(-\xi)}, \qquad (\mathfrak{c}_{\mathrm{res}}U)^-_\varepsilon(\xi) = 16r^4\overline{U^+_{-\varepsilon}(-\xi)}.$$

Its square is the identity: reflection preserves r, the reciprocal factors cancel, and conjugation squares to one. The actual spin marking remains supplied. These even projective weights do not detect the two central spin lifts by their action alone.

The linear isomorphism covers all weighted smooth TT seeds in the specified orders, including permitted low-frequency singularities. It introduces no independent polynomial mode and imposes no Fourier-flatness or analytic-background condition on those seeds. To obtain the ordinary inputs used for source execution, one must next lift these root classes to actual closed cochains and impose the actual polynomial-obstruction tests. That next step is not a nonlinear operation on the polynomial quotient used here.

3.4. Actual coefficient complexes and ordinary classes. The root complex is a linear model for the two free fields. To perform nonlinear operations on smooth representatives, we retain the coefficient space before taking its polynomial quotient. We now determine the additional cohomology in that space.

For a smooth bundle V over the sphere, write $\mathscr{S}^\mu(V)$ for the space of actual smooth coefficients with seminorms

$$\sup_{x,\zeta}\langle x\rangle^{-\mu+|\alpha|}\left|\partial_x^\alpha\nabla_\zeta^\beta f(x,\zeta)\right|, \qquad \alpha,\beta\ge 0,$$

in the fixed finite bundle charts. One spatial order μ applies to every derivative. Put $\mathscr{E}(V)=\bigcup_{\mu\in\mathbb{R}}\mathscr{S}^\mu(V)$, and let $\mathscr{P}(V)$ denote the finite-degree polynomials in x with smooth fiber coefficients. Statements of continuity on this union will specify an input order, an output order and, when relevant, a polynomial-degree bound.

With $D = z\cdot\partial_x$, the actual weight-m complex is

$$\begin{aligned}
\mathscr{E}^0_m &= \mathscr{E}(\mathcal{O}(m)),\\
\mathscr{E}^1_m &= \mathscr{E}(\Lambda^{0,1}\otimes\mathcal{O}(m))\oplus\mathscr{E}(\mathcal{O}(m+2)),\\
\mathscr{E}^2_m &= \mathscr{E}(\Lambda^{0,1}\otimes\mathcal{O}(m+2)),\\
d^0u &= (\bar\partial u, Du), \qquad d^1(\alpha,\beta)=\bar\partial\beta - D\alpha.
\end{aligned}$$

The differential squares to zero because z is holomorphic in the sphere variable and independent of x. An ordinary class is represented by a closed

smooth cochain in this complex, modulo differentials of smooth zero-cochains in the same coefficient space.

Proposition 3.7 (Actual coefficients and their polynomial kernel). The actual moderate coefficient space used to define the root complex equals $\mathscr{E}(V)$. Punctured Fourier restriction gives an exact sequence

$$0 \longrightarrow \mathscr{P}(V) \longrightarrow \mathscr{E}(V) \xrightarrow{\mathcal{F}|_{\xi\neq 0}} \mathscr{A}(V) \longrightarrow 0.$$

For every fixed M, a continuous linear extension $\mathscr{A}_M(V) \to \mathscr{S}^\mu(V)$ exists when $\mu > 0$ is nonintegral and $\mu > M - 3$. Its choice may depend on the spatial anchor. Actual multiplication and differentiation satisfy

$$\mathscr{S}^\mu(V)\mathscr{S}^\nu(W) \subset \mathscr{S}^{\mu+\nu}(V \otimes W), \qquad \partial_{x^i}\mathscr{S}^\mu(V) \subset \mathscr{S}^{\mu-1}(V).$$

The collection of actual CR complexes over all integer weights is a differential graded algebra.

Proof. If $f \in \mathscr{S}^\mu(V)$, increase μ to a nonnegative number for the Fourier estimate. A spatial dyadic decomposition, rescaled to fixed annuli, gives

$$\left|\partial_\xi^\alpha \nabla_\zeta^\beta \widehat{f}_j(\xi,\zeta)\right| \le C_{\alpha,\beta,N} 2^{j(\mu+3+|\alpha|)}(1+2^j|\xi|)^{-N}.$$

Splitting the sum at $2^j|\xi| = 1$ gives the punctured estimate $C|\xi|^{-\mu-3-|\alpha|}$ near zero and rapid decrease, with all derivatives, at infinity. This is the same annular estimate as in the decay Fourier lemma; here it is used only for the restriction away from zero and does not require decay of f. Thus $\mathscr{E}(V)$ lies in the actual moderate space.

Conversely, let $a \in \mathscr{A}_M(V)$. Separate a Schwartz coefficient away from zero and frequency annuli a_j supported where $|\xi| \asymp 2^{-j}$, $j \ge 0$. Their inverse transforms can be written

$$f_j(x,\zeta) = 2^{j(M-3)} g_j(2^{-j}x, \zeta),$$

where the g_j are bounded in every Schwartz and fiber seminorm by finitely many seminorms of a. Let $L = \lfloor \mu \rfloor$. The strict inequality $\mu > M - 3$ implies $M - 4 - L < 0$. If $T_L f_j$ is the spatial Taylor polynomial at zero through degree L, define

$$Ea = f_{\text{high}} + \sum_{j\ge 0}(f_j - T_L f_j).$$

For $|\alpha| \le L$, Taylor's formula bounds a differentiated summand on a compact set by $C2^{j(M-4-L)}|x|^{L+1-|\alpha|}$. For $|\alpha| > L$, its polynomial part is absent and the bound is $C2^{j(M-3-|\alpha|)}$. Both tails are summable. The same statements hold after each specified fiber derivative, so Ea is actual smooth.

To estimate it at $R = |x| \ge 2$, split the sum at $2^j = R$. The tail is bounded by $CR^{M-3-|\alpha|}$. The nonpolynomial part of the remaining finite sum has the same bound by Schwartz decay. A Taylor term of degree d contributes at most

$$CR^{d-|\alpha|} \sum_{2^j \le R} 2^{j(M-3-d)}, \qquad |\alpha| \le d \le L.$$

According to the sign of $M-3-d$, this is bounded by a multiple of $R^{M-3-|\alpha|}$, $R^{d-|\alpha|}$, or $R^{d-|\alpha|}\log R$. Since $\mu > M-3$ and $\mu > L$, each expression is $O(R^{\mu-|\alpha|})$. These estimates prove $Ea \in \mathscr{S}^{\mu}(V)$ and continuity in finitely many source seminorms for each target seminorm. They also control partial sums as tempered distributions. The subtracted polynomials have Fourier support at zero, so the punctured Fourier restriction is a.

The point-support statement in the Fourier quotient lemma identifies the kernel with $\mathscr{P}(V)$. For any coefficient in the original actual moderate space, subtract the extension just constructed from its Fourier restriction. The difference is polynomial, hence belongs to $\mathscr{E}(V)$. This proves the equality and exactness.

The multiplication and derivative bounds are direct consequences of the Leibniz rule. Fiber derivatives preserve the order. The operators $\bar\partial$ and D commute and obey the graded Leibniz rule. In degree one, the product is

$$(\alpha,\beta)\wedge(\alpha',\beta') = \alpha\beta' - \beta\alpha',$$

with the line-bundle weights added. This proves the differential graded algebra assertion. □

The polynomial coefficients form a differential subcomplex, so the quotient supports the linear root calculation. Since the polynomials contain 1, the algebra ideal they generate is the entire coefficient algebra. We therefore form nonlinear products in the coefficient algebra $\mathscr{E}$.

Proposition 3.8 (The polynomial cohomology and actual root map).

Let $\mathcal{P}$ be the scalar polynomials in x. Normalize the fiber identifications by

$$v \longmapsto v_i z^i \in H^0(\mathcal{O}(2)), \qquad S \longmapsto S_{ij} z^i z^j \in H^0(\mathcal{O}(4)),$$

and use their Serre-dual pairings to identify $H^1(\mathcal{O}(-6))$ with constant complex symmetric tracefree tensors and $H^1(\mathcal{O}(-4))$ with complex vectors. Put

$$(\mathcal{K}v)_{ij} = \frac{1}{2}(\partial_i v_j + \partial_j v_i) - \frac{1}{3}\delta_{ij}\partial_k v_k, \qquad (\operatorname{div} S)_j = \partial_i S_{ij}.$$

Then

$$H^1(\mathscr{P}_2^\bullet) = (\mathcal{P}\otimes \operatorname{STF}^2\mathbb{C}^3)/\mathcal{K}(\mathcal{P}\otimes\mathbb{C}^3),$$
$$H^1(\mathscr{P}_{-6}^\bullet) = \ker\bigl(\operatorname{div} : \mathcal{P}\otimes\operatorname{STF}^2\mathbb{C}^3 \longrightarrow \mathcal{P}\otimes\mathbb{C}^3\bigr).$$

For both weights $H^2(\mathscr{P}_m^\bullet) = 0$. The negative degree-zero cohomology vanishes, while $H^0(\mathscr{P}_2^\bullet) = \ker\mathcal{K}$ is the ten-dimensional complex space of polynomial conformal Killing fields.

For $m = 2, -6$, the actual root map fits into the exact sequence

$$0 \longrightarrow H^1(\mathscr{P}_m^\bullet) \longrightarrow H^1(\mathscr{E}_m^\bullet) \xrightarrow{\rho_m} \mathscr{H}_m^1 \longrightarrow 0.$$

Moreover $H^0(\mathscr{E}_m^\bullet) = H^0(\mathscr{P}_m^\bullet)$ and $H^2(\mathscr{E}_m^\bullet) = 0$. Every root class has an actual closed lift, obtained by finite polynomial corrections after a coefficient extension. These constructions have finite estimates on fixed coefficient orders and bounded polynomial degrees.

Proof. The fixed fiber Hodge operators preserve polynomial spatial coefficients. In positive weight they reduce the polynomial complex to

$$\mathcal{P} \otimes H^0(\mathcal{O}(2)) \xrightarrow{D} \mathcal{P} \otimes H^0(\mathcal{O}(4)).$$

Since $z \cdot z = 0$, this is precisely $\mathcal{K}$ in the displayed tensor identifications. The relevant higher fiber cohomology vanishes. This proves the positive formulas and its degree-two vanishing.

In negative weight the reduced map is

$$\mathcal{P} \otimes H^1(\mathcal{O}(-6)) \xrightarrow{H^1_{-4}D} \mathcal{P} \otimes H^1(\mathcal{O}(-4)).$$

It is exactly div. Indeed, the dual of multiplication by z^i pairs S with the tracefree part of $\frac{1}{2}(e_i \otimes v + v \otimes e_i)$, giving $S_{ij}v_j$. Differentiation in x^i gives $\partial_i S_{ij}$ with no additional scalar.

This polynomial divergence map is onto. Set $\mathcal{L} = 2\mathcal{K}$ and

$$\mathcal{N}v = \operatorname{div} \mathcal{L}v = \Delta v + \frac{1}{3}\nabla \operatorname{div} v.$$

The coefficient of ∂_1^2 in $\mathcal{N}$ is $M = \operatorname{diag}(4/3, 1, 1)$. Write its other terms as $B^A\partial_1\partial_A + C^{AB}\partial_A\partial_B$, $A, B \in \{2, 3\}$. For a polynomial vector w of degree at most d, prescribe $v_0 = v_1 = 0$ and recursively define

$$v_{n+2} = M^{-1}\left(\partial_1^n w|_{x^1=0} - B^A\partial_A v_{n+1} - C^{AB}\partial_A\partial_B v_n\right).$$

Induction gives $\deg v_n \le d + 2 - n$. The recursion therefore terminates, and $v = \sum_{n=0}^{d+2}(x^1)^n v_n/n!$ satisfies $\mathcal{N}v = w$. Thus $\mathcal{L}v$ is a polynomial tracefree tensor with divergence w. This proves negative degree-two vanishing. Its degree-zero group vanishes because $H^0(\mathcal{O}(-6)) = 0$.

For the positive kernel, writing $\sigma = \operatorname{div} v/3$ in $\mathcal{K}v = 0$ and differentiating gives

$$\partial_i\partial_j v_k = \delta_{jk}\partial_i\sigma + \delta_{ik}\partial_j\sigma - \delta_{ij}\partial_k\sigma.$$

Commuting one more derivative and contracting gives $\partial_j\partial_l\sigma + \delta_{jl}\Delta\sigma = 0$. Taking the trace yields $\Delta\sigma = 0$, hence $\nabla^2\sigma = 0$. Substitution gives precisely

$$v(x) = a + \Omega x + \lambda x + 2(c \cdot x)x - |x|^2 c, \qquad \Omega^{\mathsf{T}} = -\Omega.$$

The dimensions of these summands are $3, 3, 1, 3$.

Apply the long exact cohomology sequence to the coefficient exact sequence above. The root-cohomology theorem gives $\mathscr{H}^0_m = \mathscr{H}^2_m = 0$. The polynomial degree-two vanishing just proved therefore gives the displayed exact sequence and the assertions in degrees zero and two. In particular its first map is injective; a nonzero polynomial first-cohomology class remains nonzero in actual cohomology.

Actual closed lifts can also be obtained directly. Extend a root cocycle coefficientwise to an actual cochain e. Its residual $r = d^1 e$ is a polynomial degree-two cochain. For positive weight, $p = (0, T_4 r)$ satisfies $d^1 p = r$. For negative weight, let $R = H^1_{-4} r$, solve $\mathcal{N}v = R$ by the finite recursion,

and let a be the harmonic representative corresponding to $-\mathcal{L}v$. Then $H^1_{-4}(r + Da) = 0$, and

$$p = \big(a, T_{-4}(r + Da)\big), \qquad d^1 p = r.$$

Thus $e - p$ is an actual closed lift. All operations are fixed fiber maps, coefficient extensions with their stated order bounds, and finite polynomial linear operations. Their estimates have the asserted finite dependence. □

The map ρ_m is therefore not injective on the full actual cohomology. We next specify the lift used for the admitted weighted pairs and give a direct test for its actual cohomology class.

Theorem 3.9 (The canonical actual lift of the weighted pair). Let $-1 < \gamma < -1/2$ and put $\vartheta = \gamma + 1 \in (0, 1/2)$. There is a continuous real-linear section

$$\sigma_\gamma : \mathscr{W}_\gamma \longrightarrow H^1(\mathscr{E}_2^\bullet) \oplus H^1(\mathscr{E}_{-6}^\bullet), \qquad \rho\sigma_\gamma = I.$$

For a paired root class $b \in \mathscr{W}_\gamma$, let $\beta_+(\xi)$ and $\alpha_-(\xi)$ be its positive harmonic normal coefficient and negative harmonic first coefficient in the root-cohomology theorem. A fixed actual representative of $\sigma_\gamma(b)$ is

$$\begin{aligned} e_+(b) &= (0, B_+(b)), \\ B_+(b)(x) &= (2\pi)^{-3/2} \int_{\mathbb{R}^3} (e^{ix\cdot\xi} - 1)\beta_+(\xi)\, \mathrm{d}\xi, \\ e_-(b) &= (A_-(b), B_-(b)), \\ A_-(b) &= \mathcal{F}^{-1}\alpha_-, \qquad B_-(b) = T_{-4} D A_-(b). \end{aligned}$$

These cochains are exactly closed and satisfy

$$\begin{aligned} &B_+ \in \mathscr{S}^\vartheta(\mathcal{O}(4)), \qquad B_+(0) = 0, \\ &A_- \in \mathscr{S}^{\vartheta-3}(\Lambda^{0,1} \otimes \mathcal{O}(-6)), \\ &B_- \in \mathscr{S}^{\vartheta-4}(\mathcal{O}(-4)). \end{aligned}$$

The positive representative is unique with its punctured harmonic normal coefficient, the displayed growth, and the anchor. The negative representative is unique in its harmonic normal form with the displayed decay. Changing the positive anchor changes the actual representative by an explicit ordinary coboundary, so its actual class is unchanged. The section is translation-equivariant in actual cohomology, preserves the calibrated coupled root reality, and preserves smooth passive families in the stated seminorms.

Proof. For $(q, p) = \mathscr{P}_\gamma b$, let $Q = \widehat{q}$ and $P = \widehat{p}$. The weighted root-domain bounds give $Q = O(r^{-\gamma-3})$ and $P = O(r^{-\gamma-2})$, $r = |\xi|$, with their differentiated versions. The inverse root-seed formulas give $U_\varepsilon^+ = O(r^{-\gamma-4})$ and $U_\varepsilon^- = O(r^{-\gamma})$. The positive root-to-normal map has degree zero and the negative map has degree -1. Consequently

$$\begin{aligned} &|\partial_\xi^\alpha \nabla_\zeta^\beta \beta_+| \le C_{\alpha,\beta} r^{-\vartheta-3-|\alpha|}, \\ &|\partial_\xi^\alpha \nabla_\zeta^\beta \alpha_-| \le C_{\alpha,\beta} r^{-\vartheta-|\alpha|}, \qquad 0 < r \le 1. \end{aligned}$$

Both coefficients decrease rapidly at infinity with all derivatives.

For B_+, the single bracket $e^{ix\cdot\xi}-1$ supplies one power of r near zero on compact spatial sets. Its two summands are not integrated separately. Including the three-dimensional volume factor leaves the integrable radial bound $Cr^{-\vartheta}\,\mathrm{d}r$. Every positive-order spatial derivative is also integrable. This proves actual smoothness, the fiber holomorphic property, and $B_+(0)=0$.

The one-way inverse assertion in the decay Fourier lemma applies to every prescribed punctured amplitude with a negative spatial exponent. Its two-sided characterization is used only in its stated interval $-3<s<0$. In particular the one-way estimate can be applied below to the constructed coefficient of B_-, whose spatial order is below -3.

Apply it to the first derivatives of B_+, whose Fourier coefficients have exponent $s=\vartheta-1<0$. Their higher spatial derivatives have the corresponding lower orders. Integrating their bound along the segment from 0 to x gives $|B_+(x)|\le C\langle x\rangle^{\vartheta}$, also after fiber differentiation. This proves the first order statement. The differentiated integral gives $\partial_j B_+=\mathcal{F}^{-1}(i\xi_j\beta_+)$. On punctured frequency space these three identities imply $\widehat{B}_+=\beta_+$, because at least one coordinate of ξ is nonzero. The anchor fixes the constant left undetermined by the spatial derivatives. Since B_+ is a holomorphic quartic in the fiber, $d^1e_+=0$ exactly.

The coefficient α_- and every spatial Fourier multiple are locally integrable since $\vartheta<3$. The reverse estimate with $s=\vartheta-3$ gives the stated order for A_-. The coefficient of B_- is $T_{-4}L_\xi\alpha_-$, with low-frequency order $r^{1-\vartheta}$; the same reverse estimate, now with $s=\vartheta-4$, gives its stated order. The identity $H^1_{-4}L_\xi\alpha_-=0$ holds for these locally integrable extensions as an actual distributional identity. Multiplication by ξ produces no point-supported term. The fiber homotopy formula therefore gives $\bar\partial B_-=DA_-$, proving exact closedness.

Two positive normal representatives with the same punctured coefficient differ by a polynomial. Their order $\vartheta<1$ makes it constant, and the anchor makes it zero. For the negative normal forms, the difference of first components is a decaying polynomial and hence zero. The second component is then fixed because $H^0(\mathcal{O}(-4))=0$.

An anchor a replaces $B_+(x)$ by $B_+(x)-B_+(a)$. Write $B_+(a)=S_{ij}z^iz^j$ with S tracefree, and set $u_a=S_{ij}x^iz^j$. Then

$$d^0u_a=(0,B_+(a)),$$

so the change is the coboundary $-d^0u_a$. These comparisons add exactly for successive anchor changes. The anchored representative of the translated Fourier coefficient $e^{ia\cdot\xi}\beta_+$ is $B_+(x+a)-B_+(a)$, proving translation equivariance of its ordinary class. The negative inverse transform already commutes with translation on representatives.

Every output seminorm has been estimated using finitely many input seminorms. Differentiation of a supplied smooth finite-dimensional family

under the integrals therefore gives the same estimates for all specified parameter derivatives. Root readout recovers the initial normal representatives. Thus $\rho\sigma_\gamma = I$ and the section is injective on $\mathscr{W}_\gamma$. It preserves the projective weights and the specified coupled root reality. This assertion does not impose an extra pointwise conjugation law on arbitrary off-root representatives. □

Theorem 3.10 (Independent normalization tests for actual cochains).

Let $e = (e_+, e_-)$ be an actual closed cochain pair with root readout $b \in \mathscr{W}_\gamma$. There are polynomial obstructions

$$\kappa_+(e) \in (\mathcal{P} \otimes \mathrm{STF}^2\,\mathbb{C}^3)/\mathcal{K}(\mathcal{P} \otimes \mathbb{C}^3),$$
$$\kappa_-(e) \in \ker\bigl(\mathrm{div} : \mathcal{P} \otimes \mathrm{STF}^2\,\mathbb{C}^3 \longrightarrow \mathcal{P} \otimes \mathbb{C}^3\bigr)$$

such that

$$e - e(b) = d^0 u \quad \Longleftrightarrow \quad \kappa_+(e) = \kappa_-(e) = 0.$$

The obstructions and a primitive when they vanish are computed by fixed fiber Hodge maps, punctured Fourier multipliers, coefficient extension and finite polynomial linear algebra. The negative primitive is unique. The positive primitive is unique modulo the ten-dimensional polynomial conformal Killing kernel. The tests are invariant under actual coboundaries and have finite estimates on fixed symbol orders and bounded polynomial degrees, allowing a finite increase of order. They preserve smooth passive families and use no spacetime metric or Einstein evolution.

Proof. Write $e_+ = (\alpha, \beta)$ and $e_+(b) = (0, B_b)$. Fiber normalization gives

$$\beta_0 = \beta - DT_2\alpha, \qquad \bar\partial\beta_0 = 0.$$

At nonzero frequency put $v = L_\xi^+ \widehat{\beta}_0$. Its homogeneous order is -1, so v belongs to a fixed moderate coefficient space. Extend it to an actual coefficient Ev by the coefficient proposition, keeping its values in the finite-dimensional space of holomorphic quadratic sections. Then

$$r = \beta_0 - B_b - D(Ev)$$

has zero punctured Fourier transform, because $(I - L_\xi L_\xi^+)\widehat{\beta}_0$ is exactly the positive normal coefficient defining B_b. Thus r is a polynomial holomorphic quartic. Define $\kappa_+(e) = [r]$ in the stated quotient. Another extension differs by a polynomial quadratic and changes r by an element of $\operatorname{im}\mathcal{K}$, so this class is independent of that choice.

If $[r] = 0$, solve $Dp = r$ among polynomial holomorphic quadratic sections. For a nonzero residual of degree D, a solution can be chosen of degree at most $D + 1$: every higher homogeneous component of any polynomial solution lies separately in $\ker\mathcal{K}$ and can be removed. For $r = 0$, take $p = 0$. On the resulting bounded-degree space this is finite-dimensional linear algebra on the image; choose a complement to the conformal Killing kernel to fix a solution. Then

$$u_+ = T_2\alpha + Ev + p$$

satisfies $d^0 u_+ = e_+ - e_+(b)$. Conversely, if that difference equals $d^0 u_+$, the identity $T_2 \bar{\partial} u_+ = u_+ - H_2^0 u_+$ gives $\beta_0 = B_b + D(H_2^0 u_+)$. Since the canonical positive normal coefficient lies in $\ker L_\xi^*$, v is the punctured Fourier restriction of $H_2^0 u_+$. Hence $H_2^0 u_+ - Ev$ is polynomial and r lies in $\operatorname{im} \mathcal{K}$. This proves the positive equivalence.

For $e_- = (\alpha, \beta)$, let $\alpha_0 = H_{-6}^1 \alpha$, and write $e_-(b) = (A_b, B_b)$. The two harmonic first coefficients have the same punctured Fourier restriction. Thus

$$S = \alpha_0 - A_b$$

is a polynomial harmonic section. Closedness gives $H_{-4}^1 DS = 0$, which is exactly $\operatorname{div} S = 0$ in the normalized tensor identification. Set $\kappa_-(e) = S$. If $S = 0$, subtract $d^0(T_{-6}\alpha)$ from $e_- - e_-(b)$. Its first component is zero and its second component is holomorphic in $\mathcal{O}(-4)$, hence zero. This gives $u_- = T_{-6}\alpha$. Necessity follows by taking H_{-6}^1 of the first component of an exact cochain.

Both obstructions are precisely the polynomial cohomology class of $e - e(b)$ in the injective map of the actual cohomology sequence. Adding a coboundary in the coefficient complex preserves both obstructions. Two primitives differ by a closed zero-cochain, so the degree-zero groups computed above give the stated ambiguity. The coefficient extensions and multipliers have finite seminorm estimates, and the remaining inverses are fixed linear maps on bounded-degree polynomial spaces. This proves continuity and the family assertions at the specified stages. □

3.5. The selected ordinary domain. Let $\mathcal{B}_\gamma$ be a fixed open finite-base neighborhood of zero in the smooth TT seed space, contained in the retained nonlinear chart domain and, when source execution is invoked, in its stated domain. Put

$$\mathscr{D}_{\gamma,\mathcal{B}} = \mathscr{P}_\gamma^{-1}(\mathcal{B}_\gamma), \qquad \mathscr{C}_{\gamma,\mathcal{B}}^{\mathrm{ord}} = \sigma_\gamma(\mathscr{D}_{\gamma,\mathcal{B}}).$$

For an ordinary pair $c \in \mathscr{C}_{\gamma,\mathcal{B}}^{\mathrm{ord}}$, use the distinct notations

$$b = \rho(c) \in \mathscr{D}_{\gamma,\mathcal{B}}, \qquad s(c) = \mathscr{P}_\gamma \rho(c) \in \mathcal{B}_\gamma.$$

The first is a paired root class, and the second is the actual whole TT seed. No quotient by spatial polynomials is applied to the ordinary pair itself.

Definition 3.11 (Admitted actual cochains and effective ordinary arrows).

An admitted cochain pair is an element $e \in \mathscr{E}_2^1 \oplus \mathscr{E}_{-6}^1$ satisfying

$$d^1 e_+ = d^1 e_- = 0, \qquad \rho(e) \in \mathscr{D}_{\gamma,\mathcal{B}}, \qquad \kappa_+(e) = \kappa_-(e) = 0.$$

The positive obstruction is tested in its polynomial quotient, rather than by setting a chosen residual polynomial equal to zero. An effective ordinary arrow $e \to e'$ is an actual zero-cochain pair u with $e' - e = d^0 u$, modulo

$$u \sim v \quad \Longleftrightarrow \quad d^0(u - v) = 0.$$

Composition is addition of primitives. Coefficient topology and family estimates are taken at fixed symbol orders and bounded polynomial degrees, together with the $\mathscr{W}_\gamma$ topology of the root readout.

Corollary 3.12 (The actual ordinary input represented by the tests).

The actual cohomology classes represented by the admitted cochains are precisely $\mathscr{C}^{\mathrm{ord}}_{\gamma,\mathcal{B}}$. On this selected ordinary domain, root readout and σ_γ are inverse. Two admitted cochains admit an effective ordinary arrow exactly when their roots agree, and that effective arrow is then unique. The representatives and comparison primitives have the finite seminorm and parameter estimates stated above.

Proof. If e is admitted, the normalization test gives $e - e(\rho(e)) = d^0 u$, so $[e] = \sigma_\gamma \rho(e)$. Conversely every $e(b)$ with $b \in \mathscr{D}_{\gamma,\mathcal{B}}$ is closed and has zero obstruction, as does every actual cohomologous representative. The identity $\rho\sigma_\gamma = I$ gives the inverse assertion on classes.

Every coboundary in the coefficient complex has zero root readout, as required for the arrow statement. If two roots agree, subtract their two explicit comparisons with the same canonical representative to obtain an actual primitive for their difference. Any two such primitives differ by a closed zero-cochain and hence define the same effective arrow. The displayed formulas give addition under composition and the asserted estimates. □

These effective arrows compare cochain representatives of the ordinary input. Comparisons of completed sources also carry a physical map and its actual spin lift, as specified in §7.

3.6. The ordinary-pair initial-data constructor. The root transform, actual lift and nonlinear chart now have compatible domains. We state their exact composition before using it to cover local spacetime germs. All fields in this construction remain defined on the whole marked $\mathbb{R}^3$; a local restriction is taken only after construction and decoding.

Let $\mathcal{U} \subset \mathcal{X}_{k_0}$ and $\mathcal{V} \subset \mathcal{Y}_{k_0}$ be the convex base neighborhoods of the constraint chart. Choose the finite-base neighborhood $\mathcal{B}_\gamma$ of the preceding subsection inside $\mathcal{U} \cap \mathcal{X}^\infty_\gamma$. In applications requiring the initial source execution, it is also contained in that operation's retained domain, as specified in §5.1. This extra use of a smaller seed ball does not change the coefficient complexes, transform or decoder.

For a whole tensor pair in the ambient gauge domain, write the already defined algebraic and differential readout as

$$\mathrm{Dec}(h,K) = \big(\mathrm{Dec}_s(h,K), \mathrm{Dec}_y(h,K)\big) = \big((q,p),(u,W)\big).$$

Define $\mathsf{Data}^{\mathrm{g}}_{\gamma,\mathcal{B}}$ by the following tests on the tensors themselves:

$$h - \delta \in \mathscr{S}^\gamma, \qquad K \in \mathscr{S}^{\gamma-1}, \qquad \tfrac{1}{2}\delta < h < \tfrac{3}{2}\delta,$$
$$\mathrm{div}_\delta \,\mathrm{TF}_\delta\, h = 0, \qquad \mathrm{tr}_h\, K = 0, \qquad R(h) = |K|^2_h, \qquad \mathrm{div}_h\, K = 0,$$
$$\mathrm{Dec}_s(h,K) \in \mathcal{B}_\gamma, \qquad \mathrm{Dec}_y(h,K) \in \mathcal{V}.$$

The fields are real covariant tensors, and the coordinate, scale and actual spin markings are retained. The decoder is defined on the ambient gauge domain before the two constraint equations are imposed. Thus these tests require no supplied generating pair and no membership in a constructor image. They select the part of the independent tensor target in the fixed uniqueness neighborhood.

Theorem 3.13 (Ordinary classes and the independent maximal-data domain).
The maps

$$\begin{aligned} \mathfrak{D}(c) &= \mathcal{T}\big(\mathscr{P}_\gamma \rho(c)\big), && c \in \mathscr{C}^{\mathrm{ord}}_{\gamma,\mathcal{B}}, \\ \mathfrak{C}(h,K) &= \sigma_\gamma \mathscr{P}_\gamma^{-1}\,\mathrm{Dec}_s(h,K), && (h,K) \in \mathsf{Data}^{\mathrm{g}}_{\gamma,\mathcal{B}}, \end{aligned}$$

are inverse bijections. In particular,

$$\mathfrak{C}\mathfrak{D}(c) = c, \qquad \mathfrak{D}\mathfrak{C}(h,K) = (h,K).$$

The forward operation uses the ordinary pair, the fixed coefficient normalizations and the retained flat markings. Its nonlinear correction is the selected solution of the constraint chart. Both maps are continuous in their full weighted seminorm topologies. They preserve every supplied smooth finite-dimensional family on its retained domain. For retained compact parameter sets $\mathcal{K} \Subset \operatorname{int}\mathcal{K}'$, each specified finite output estimate on $\mathcal{K}$ is controlled by finitely many input spatial and parameter seminorms on $\mathcal{K}'$ and the fixed base margins.

If the input is given by an admitted actual cochain pair, the result is independent of its effective ordinary presentation. Its comparison with the canonical representative uses a primitive in the coefficient complex satisfying the finite estimates in the obstruction theorem.

Proof. For an admitted ordinary pair, set $s = \mathscr{P}_\gamma \rho(c)$. The selected-domain definition gives $s \in \mathcal{B}_\gamma$. The base chart gives $\mathcal{T}(s) = \mathcal{A}(s, \mathcal{C}(s))$ satisfying the metric gauge, maximality and the two vacuum constraints. The all-order theorem puts its tensors in the required symbol spaces. The exact ambient decoder returns $(s, \mathcal{C}(s))$, so the two neighborhood tests also hold. Hence the forward map lands in the independently defined target.

Conversely, take an independently admitted tensor pair and decode it to (s, y). By definition $s \in \mathcal{B}_\gamma \subset \mathcal{U}$ and $y \in \mathcal{V}$. The ambient identities give $(h, K) = \mathcal{A}(s, y)$. Its vacuum constraints are exactly $\mathcal{F}(s, y) = 0$. Uniqueness of the base correction therefore forces $y = \mathcal{C}(s)$. This proves that its inverse-class formula constructs an ordinary input whose output is exactly the initial tensor pair.

On the ordinary domain, the actual normalization theorem and its corollary give $\sigma_\gamma \rho(c) = c$. Together with the two inverse identities for $\mathscr{P}_\gamma$ and the seed decoder, this proves $\mathfrak{C}\mathfrak{D} = I$. The identity $\rho\sigma_\gamma = I$ and the preceding constraint-uniqueness argument prove $\mathfrak{D}\mathfrak{C} = I$.

Continuity and the finite-family estimates follow in the order in which the maps are applied: the exact root transform and actual section, the same-branch nonlinear chart, and the exact ambient decoder. For any admitted cochain presentation, the polynomial-obstruction theorem constructs the comparison primitive at fixed coefficient orders and bounded polynomial degrees. Its effective-arrow assertion makes the resulting ordinary class independent of that representative. Every finite estimate is a finite composition of the displayed estimates. The one base neighborhood is fixed before choosing any higher output order. □

Corollary 3.14 (Whole maximal data with the based comparison retained).

Let (H, L) be whole smooth maximal vacuum data with $H - \delta \in \mathscr{S}^{\gamma}$ and $L \in \mathscr{S}^{\gamma-1}$, sufficiently small in the finite weighted norms of Corollary A.11. After one restriction of that threshold, there are the initialized based gauge ϕ_H and a unique ordinary pair $c_H \in \mathscr{C}^{\mathrm{ord}}_{\gamma,\mathcal{B}}$ such that

$$\phi_H^*(H, L) = \mathfrak{D}(c_H), \qquad c_H = \sigma_\gamma \mathscr{P}_\gamma^{-1} \operatorname{Dec}_s(\phi_H^*(H, L)).$$

The comparison retains the actual spin lift initialized along the based gauge isotopy. Its gauge, ordinary representatives and decoded data have all specified finite spatial and parameter estimates on the same retained family domain, with the compact-set and finite-input quantifiers of the preceding theorem.

Proof. Corollary A.11 supplies the actual gauge and an exact seed in the selected constraint chart. The gauge estimate and continuity of its decoder at the base order allow the input threshold to be reduced once so that this seed lies in the prescribed neighborhood $\mathcal{B}_\gamma$. Apply the preceding theorem. Its uniqueness is uniqueness of the ordinary pair for the normalized whole datum in the retained chart; the based gauge itself is unique in its displacement ball. The spin and family assertions are those of the gauge comparison followed by the finite estimates of the preceding theorem. □

Theorem 3.13 recovers the normalized whole datum from its ordinary pair. Section 4 constructs the chiral gradient and contact operators of the same pair. Theorem 4.18 proves that the operators recover the whole datum. The local coverage argument in Section 5 uses the following initial-data results.

Input to local coverage	Proof in this article
Weighted root pair and exact TT seed	Theorem 3.6
Actual ordinary lift and independent cochain tests	Theorems 3.9–3.10 and Corollary 3.12
Nonlinear constraints and exact whole decoder	Theorem 3.13, Proposition A.4 and Theorems A.6–A.7
Based coordinate gauge with actual spin	Theorem A.9, Proposition A.10 and Corollary 3.14

4. Actual source execution from ordinary classes

Starting from the fixed flat source, we construct the initial gradient and contact operators of a selected ordinary pair. The weighted negative-to-dual operation acts on closed cochain representatives and includes the primitives that compare them. Contact normalization over the complete sphere then gives the negative response. A finite composition of those responses realizes the scalar and York inverses required by the nonlinear current correction.

The source curvature supplies four independently defined defects. At each current source, its vertical response agrees with the transverse differential of these defects. The resulting horizontal field changes the prescribed ordinary seed and preserves the defects. Its endpoint is selected on one finite base neighborhood, and its weighted smoothness and family estimates concern that same endpoint. The argument in Appendix A supplies actual Picard convergence, path independence and a finite-execution error bound.

The canonical initial target is specified by tests on supplied operators or complete contact coefficients. The resulting unit and counit identify this target with the ordinary-pair domain. Full spacetime admission, wave enrichment and changes of initial carrier use the additional structures identified in Section 5.

4.1. The actual weighted negative-to-dual operation. The negative source determines a positive class for the conjugate projective sphere. This is a complex-linear operation on the original negative class. Its conjugate complex structure changes the angular differential; the operation itself preserves the real spatial Fourier frequency. We first fix the original scalar calibration, then prove the operation on actual weighted cochains and exhibit the primitives required to compare their representatives.

Retain the flat conic z, Fourier normalization and physical root-line identifications of the root transform. Write $D = z \cdot \partial_x$ and $D^\vee = \bar{z} \cdot \partial_x$. The conjugate positive line is $\overline{\mathcal{O}(2)}$: it has degree 2 for the conjugate sphere complex structure and Chern number -2 for the original sphere orientation. Its differential and horizontal frequency multiplier are

$$d^0_\vee v = (\partial v, D^\vee v), \quad d^1_\vee(\alpha, \beta) = \partial\beta - D^\vee\alpha,$$
$$L^\vee_\xi = i\xi \cdot \bar{z}, \qquad\qquad L_\xi = i\xi \cdot z.$$

Here ∂ acts only in the projective fiber. In particular $L^\vee_\xi$ is not $\overline{L_\xi}$ at fixed ξ. Conjugation of an actual spatial coefficient also reflects its frequency.

The roots remain $n_\varepsilon = \varepsilon\xi/r$, where $r = |\xi|$ and $\varepsilon = \pm 1$. Write U^-_ε for the original negative root tensor and $V^\vee_\varepsilon$ for the dual positive root tensor. Both take values in the line generated by $z(n_\varepsilon) \otimes z(n_\varepsilon)$. At the north root the scalar-to-tensor maps are

$$c \longmapsto c\, z \otimes z \quad \text{and} \quad c \longmapsto \frac{c}{4}\, z \otimes z$$

for the negative and dual positive weights, respectively. The second is the conjugate of the positive root-line normalization already used.

Proposition 4.1 (The original derived-dual multiplier).
On smooth Fourier amplitudes with compact support away from zero, the original analytic derived-dual class has root tensors

$$V_\varepsilon^\vee(\xi) = \frac{1}{16r^4} U_\varepsilon^-(\xi).$$

The multiplier is complex-linear, has positive sign, and uses the same frequency and root label.

Proof. The original dual coefficient formula (Isthmus 2026a, §5, equation for the derived dual class) sends the negative plane cocycle $F_- = Ae^{\lambda w^0}/\zeta$, $\lambda \neq 0$, to

$$G(\eta, \varrho) = \frac{A\lambda^{-4}}{\varrho}\left(e^{-\lambda\eta^0} - \sum_{j=0}^{3} \frac{(-\lambda\eta^0)^j}{j!}\right).$$

Indeed its original coefficients are $c^-_{j,0,0} = A\lambda^j/j!$, and the dual formula multiplies this coefficient by $(-1)^j j!/(j+4)!$. Summation gives the displayed expression. Each subtracted term is a positive-weight coboundary: $(\eta^0)^j/\varrho$ becomes $(\tilde\eta^0)^j\tilde\varrho^{3-j}$ in the weight-two south frame, with $\tilde\eta = \eta/\varrho$ and $\tilde\varrho = 1/\varrho$, and is holomorphic there when $j \leq 3$.

The derived class is consequently represented by $A\lambda^{-4}e^{-\lambda\eta^0}/\varrho$. Lorentz incidence is $\eta = -\bar w$, $\varrho = \bar\zeta$. Its restriction is $A\lambda^{-4}e^{\lambda\bar w^0}/\bar\zeta$. Apply the conjugate-angular version of the plane-to-root calculation preceding the weighted root definition, keeping λ and A unchanged. The relevant identities are $D^\vee u = -\sqrt{2}\bar\zeta$ and $D^\vee\bar v = \sqrt{2}$. Subtracting the same smooth cutoff primitive gives

$$c^\vee = -\sqrt{2}\lambda A\lambda^{-4} = \lambda^{-4}c^-, \qquad c^- = -\sqrt{2}\lambda A.$$

The pole exponent is still $e^{\lambda u}$, so the real spatial frequency is unchanged. At $\xi = \varepsilon r e_3$, the fixed soldering convention gives $\lambda = -i\sqrt{2}\varepsilon r$, hence $\lambda^4 = 4r^4$. Combining this with the additional tensor-line factor $1/4$ proves the multiplier and its sign.

The conic, incidence and original dual spin pairing commute with spatial spin rotations, so the axial calculation applies to both root branches. Let the frequency support lie in $r_0 \leq |\xi| \leq R_0$ with $r_0 > 0$. For each branch take a finite partition of unity in direction, supported in small caps with smooth spin frames. On each cap fix the angular chart at its central direction, so the corresponding pole has coordinate $a(\xi)$ and $a = 0$ at the cap centre. Fix a bounded annulus $\rho < |\zeta| < R$ in this chart. The rotation at the centre is the identity in the chosen frame. By shrinking the cap its Möbius denominator stays bounded away from zero on the closed annulus and $|a(\xi)| \leq \rho/2$. Thus $|\zeta - a(\xi)| \geq \rho/2$ throughout the overlap. Use the south coordinate at infinity and the conjugate coordinate for the dual cocycle.

Restrict the remaining complex coordinates to a common bounded open set with closure in the coordinate domains, taking the minimum of the available radii over the finite cover. For either the plane or its derived cocycle, write $F_\xi(Z)$ for its coefficient in these fixed complex coordinates. If S is the supported frequency set in the cap and K is relatively compact in the resulting overlap, $\sup_{(\xi,Z)\in S\times K}|\partial_\xi^\beta D_Z^\nu F_\xi(Z)|$ is finite for every specified pair of multi-indices. Indeed the frame factors have bounded derivatives there, $|\lambda|$ is bounded above and away from zero, and the exponential series and all its finite derivatives converge normally on a slightly larger compact set. These bounds justify integration against the smooth amplitude term by term, differentiation under the frequency integral and the contour computation of the root coefficient on a fixed circle enclosing the moving pole. In the real cochain calculation, subtract the cutoff primitive before reading the root, as in §3.3. On the real initial slice the corrected coefficients have the spatial phase $e^{ix\cdot\xi}$ times smooth compactly supported frequency amplitudes, also after fiber differentiation. Integration by parts in ξ gives rapid spatial decay of their superposition. □

Definition 4.2 (Weighted actual negative and dual complexes). Fix $-1<\gamma<-1/2$ and put $\vartheta=\gamma+1\in(0,1/2)$. Use the actual symbol spaces, including every fiber derivative, of the coefficient algebra. The negative complex is

$$\begin{aligned}\mathscr{C}^0_{-,\vartheta}&=\mathscr{S}^{\vartheta-3}(\mathcal{O}(-6)),\\ \mathscr{C}^1_{-,\vartheta}&=\mathscr{S}^{\vartheta-3}(\Lambda^{0,1}\otimes\mathcal{O}(-6))\oplus\mathscr{S}^{\vartheta-4}(\mathcal{O}(-4)),\\ \mathscr{C}^2_{-,\vartheta}&=\mathscr{S}^{\vartheta-4}(\Lambda^{0,1}\otimes\mathcal{O}(-4)),\end{aligned}$$

with $d^0_- u=(\bar\partial u, Du)$ and $d^1_-(\alpha,\beta)=\bar\partial\beta-D\alpha$. The dual positive complex is

$$\begin{aligned}\mathscr{C}^0_{\vee,\vartheta}&=\mathscr{S}^{\vartheta+1}(\overline{\mathcal{O}(2)}),\\ \mathscr{C}^1_{\vee,\vartheta}&=\mathscr{S}^{\vartheta+1}(\Lambda^{1,0}\otimes\overline{\mathcal{O}(2)})\oplus\mathscr{S}^{\vartheta}(\overline{\mathcal{O}(4)}),\\ \mathscr{C}^2_{\vee,\vartheta}&=\mathscr{S}^{\vartheta}(\Lambda^{1,0}\otimes\overline{\mathcal{O}(4)}),\end{aligned}$$

with the differential in the dual formula. These are actual complexes: a horizontal derivative lowers the spatial order by one and commutes with the fiber differential.

Their root spaces consist of tensor sections on the specified two root lines with rapid decrease at infinity, with all frequency derivatives, and bounds

$$\begin{aligned}|\partial_\xi^\mu U^-_\varepsilon(\xi)|&\le C_\mu r^{1-\vartheta-|\mu|},\\ |\partial_\xi^\mu V^\vee_\varepsilon(\xi)|&\le C_\mu r^{-\vartheta-3-|\mu|},\qquad 0<r\le 1.\end{aligned}$$

Derivatives are taken in the tensor realization or finitely many angular root frames. The complexes carry their full symbol seminorms, their cohomologies the quotient topologies, and the root spaces the displayed amplitude

seminorms. These are marginal weighted domains; an auxiliary negative input is not required by itself to be a Lorentz-fixed two-class pair.

Theorem 4.3 (Actual weighted dualization and its comparison primitives).

Root readout identifies the degree-one cohomology of each displayed actual complex with its root space. The two maps are continuous isomorphisms, with continuous choices of actual closed representatives and actual primitives for closed cochains of zero root. Consequently

$$\mathcal{D}_-^{\mathrm{dual}} : H^1(\mathscr{C}^\bullet_{-,\vartheta}) \xrightarrow{\sim} H^1(\mathscr{C}^\bullet_{\vee,\vartheta}), \qquad V_\varepsilon^\vee = \frac{1}{16r^4} U_\varepsilon^-$$

is a continuous complex-linear isomorphism, with inverse multiplier $16r^4$. It extends the original derived-dual operation in the preceding proposition.

For the fixed real anchor $x = 0$, its selected dual representative is

$$\begin{aligned} e_\vee &= (0, B_\vee), \\ T_\vee(x) &= (2\pi)^{-3/2} \int_{\mathbb{R}^3} (e^{ix\cdot\xi} - 1) \sum_{\varepsilon=\pm1} V_\varepsilon^\vee(\xi)\, \mathrm{d}\xi, \\ B_\vee(x,\zeta) &= T_\vee(x)(\bar z(\zeta), \bar z(\zeta)). \end{aligned}$$

The tensor $T_\vee$ is smooth, transverse-traceless, of order ϑ, and vanishes at the anchor. Every closed cochain in either weighted complex is recovered up to an actual coboundary within that same complex. The change of dual anchor is the explicit coboundary in that complex. Inclusion of either weighted complex into the corresponding unrestricted actual coefficient complex is injective on degree-one cohomology.

Each specified output seminorm, also after specified finite parameter differentiation, is bounded by finitely many input seminorms. The weight, anchor and coefficient domains are fixed before those orders are chosen.

Proof. We first classify the negative cochains. For $e_- = (\alpha, \beta)$, subtract $d_-^0(T_{-6}\alpha)$, using the fixed fiber Hodge operators. The first component becomes $\alpha_0 = H^1_{-6}\alpha \in \mathscr{S}^{\vartheta-3}$. Closedness and $H^0(\mathcal{O}(-4)) = 0$ give the actual identity

$$e_- - d_-^0(T_{-6}\alpha) = (\alpha_0, T_{-4}D\alpha_0), \qquad H^1_{-4}D\alpha_0 = 0.$$

Since $\vartheta - 3 \in (-3, 0)$, the two-sided Fourier estimate applies to α_0. Its punctured amplitude has order $r^{-\vartheta}$ near zero, with the corresponding derivative bounds. Root evaluation of $T_{-4}L_\xi\widehat{\alpha}_0$ has degree one and gives the asserted U^- bounds.

Conversely, the negative root inverse in the root-cohomology theorem has degree -1 for its harmonic first coefficient. Applied to the stipulated U^-, it gives $\alpha_-(\xi)$ with bounds $r^{-\vartheta-|\mu|}$. This coefficient and every spatial Fourier multiple are locally integrable. The inverse integral gives an actual $A_- \in \mathscr{S}^{\vartheta-3}$, and

$$e_-(U^-) = (A_-, T_{-4}DA_-) \in \mathscr{C}^1_{-,\vartheta}.$$

For its second component, the known amplitude is $T_{-4}L_\xi\alpha_-$, of low-frequency order $r^{1-\vartheta}$. The one-way inverse estimate, valid for every negative spatial exponent, gives order $\vartheta-4$. This does not use a two-sided Fourier characterization below -3. The harmonic identity holds for locally integrable amplitudes; multiplication by ξ introduces no point-supported term. The fiber homotopy identity consequently proves exact closedness.

For a supplied closed cochain with these roots, α_0-A_- has zero punctured Fourier transform. The actual coefficient exact sequence makes it polynomial, and its negative order makes it zero. The normalized second components then agree by the absence of holomorphic sections of $\mathcal{O}(-4)$. Thus the actual comparison primitive is $T_{-6}\alpha$ and has the required order $\vartheta-3$. This proves both directions on the negative side.

For the dual side, use the conjugate fiber Hodge operator $T_2^\vee=\overline{T_2\,\overline{(\,\cdot\,)}}$, which satisfies the analogous ∂ identities. Set

$$v_0=T_2^\vee\alpha,\qquad \beta_0=\beta-D^\vee v_0.$$

Here $v_0\in\mathscr{S}^{\vartheta+1}$, $\beta_0\in\mathscr{S}^{\vartheta}$, and closedness makes β_0 an anti-holomorphic quartic on the whole sphere. The positive-order annular Fourier estimate in the proof of Proposition 3.7 gives $\partial_\xi^\mu\widehat{\beta}_0=O(r^{-\vartheta-3-|\mu|})$ on punctured frequency space. Its root values therefore obey the required dual bounds.

Given such root values, the tensor sum $\widehat{T}_\vee=\sum_\varepsilon V_\varepsilon^\vee$ is transverse-traceless, and its quartic

$$\beta_\vee(\xi,\zeta)=\left(\sum_\varepsilon V_\varepsilon^\vee(\xi)\right)(\bar z(\zeta),\bar z(\zeta))$$

has exactly those roots. The complementary tensor line evaluates to zero at either root, while the retained line evaluates with the normalization in the root-line formula. Thus this is a right inverse to root readout.

The integral defining $T_\vee$ is one bracket integral; its two terms are not integrated separately. Near zero the bracket contributes one power of r, giving the radial majorant $Cr^{-\vartheta}\,\mathrm{d}r$ on each compact spatial set. Every positive spatial derivative is also integrable. The derivative amplitudes have spatial exponent $\vartheta-1<0$, so the one-way inverse estimate gives $\partial T_\vee\in\mathscr{S}^{\vartheta-1}$ with all higher derivatives. Integration from the anchor yields $T_\vee\in\mathscr{S}^{\vartheta}$. The differentiated integrals identify its punctured Fourier transform, and the anchor fixes the constant left by differentiation. These estimates hold also after fiber differentiation.

Its trace vanishes under the convergent bracket integral. Its divergence vanishes under the convergent first-derivative integral because $\xi_iV^\vee_{\varepsilon,ij}=0$. The coefficient $B_\vee$ is anti-holomorphic in the fiber, so $e_\vee$ is exactly closed. This proves surjectivity and the claimed representative estimates on the dual side.

It remains to compare a general β_0 to that representative. At nonzero frequency their difference has zero roots. Multiplication by $L_\xi^\vee$ maps anti-holomorphic quadratics injectively to anti-holomorphic quartics; the root

divisor sequence identifies its image with the zero-root quartics. Choose its invariant Hermitian left inverse $(L_\xi^\vee)^+$ and put

$$\widehat{w} = (L_\xi^\vee)^+(\widehat{\beta}_0 - \beta_\vee).$$

This is an anti-holomorphic quadratic coefficient with low-frequency bounds $|\partial_\xi^\mu \widehat{w}| \le C_\mu r^{-\vartheta-4-|\mu|}$ and rapid decrease at infinity. The left inverse is smooth homogeneous of degree -1: at $r = 1$ its nonzero singular values have a positive minimum on the compact direction sphere, as in the root-cohomology proof.

Use the twice-anchored integral

$$w(x) = (2\pi)^{-3/2} \int_{\mathbb{R}^3} (e^{ix\cdot\xi} - 1 - ix\cdot\xi)\widehat{w}(\xi)\,\mathrm{d}\xi.$$

Its bracket gains two powers of r near zero, again giving the integrable radial majorant $Cr^{-\vartheta}\,\mathrm{d}r$. Rapid decrease controls infinity, including the subtracted terms. The second derivatives are ordinary inverse integrals with spatial order $\vartheta - 1$. Consequently $D_x^2 w \in \mathscr{S}^{\vartheta-1}$ at all orders. Since $w(0) = 0$ and $D_x w(0) = 0$, integration along segments gives $D_x w \in \mathscr{S}^\vartheta$ and $w \in \mathscr{S}^{\vartheta+1}$. The same argument controls every fiber derivative. This proves the required actual primitive order directly.

Its punctured Fourier transform is $\widehat{w}$, as follows by taking two spatial derivatives. Hence

$$R = \beta_0 - B_\vee - D^\vee w$$

has zero punctured transform and spatial order $\vartheta < 1$. It is a constant anti-holomorphic quartic, specifically $R = \beta_0(0)$. Write $R = C_{ij}\bar{z}^i\bar{z}^j$ with C tracefree symmetric, and set $w_C = C_{ij}x^i\bar{z}^j$. Then $\partial w_C = 0$ and $D^\vee w_C = R$. The actual identity is

$$(\alpha, \beta) - (0, B_\vee) = d_\vee^0(v_0 + w + w_C).$$

Every term on the right has degree-zero order $\vartheta + 1$. This proves classification within the displayed complex, without discarding a polynomial first-cohomology class. A primitive is determined modulo a closed zero-cochain in that same complex; the formula provides one continuous choice.

All operations used are fixed compact-fiber Hodge maps, homogeneous multipliers with finite order shifts, convergent anchored integrals and finite tensor readouts. Each output seminorm therefore uses finitely many input seminorms. Root readout is continuous by the forward estimates; the representative maps are continuous by the inverse estimates. The explicit comparisons identify the induced inverse maps on the quotient topologies. Multiplication by $r^{-4}/16$ and its inverse interchange the displayed root bounds, proving the actual cohomology isomorphism and both inverse identities.

If a weighted closed cochain is exact in the unrestricted actual complex, its root readout is zero. The just-proved weighted classification then supplies a primitive within the weighted complex itself. This proves the two asserted

injections into unrestricted actual cohomology. It also identifies the negative component of the selected ordinary pair with a unique class in the present weighted negative group.

An anchor a changes $B_\vee$ to $B_\vee(x)-B_\vee(a)$. If $B_\vee(a)=C_a(\bar z,\bar z)$, the change is the actual coboundary of $-C_{a,ij}x^i\bar z^j$, of allowed order $\vartheta+1$. Successive anchor comparisons add exactly.

Finally the constructions are continuous linear maps between the stated spaces. They commute with differentiation of any supplied smooth finite-dimensional family in those seminorms. On retained compact parameter sets, applying the finite seminorm estimates to each prescribed derivative gives the asserted bounds. No choice of a primitive by an additional smooth-selection theorem is required. □

Proposition 4.4 (Coupled reality and the retained markings).
For an admitted ordinary pair $c=(c_+,c_-)$, the calibrated residue reality is equivalent to

$$\mathcal{D}_-^{\mathrm{dual}}c_-=\overline{c_+},\qquad V_\varepsilon^\vee(\xi)=\overline{U_{-\varepsilon}^+(-\xi)}.$$

Here $\overline{c_+}$ is formed by conjugating an actual positive representative and using the conjugate positive complex. With the common real anchor and the transverse-traceless quartic representatives, the equality is literal conjugation of actual closed cochains. Other retained representatives have the actual comparison primitives proved above.

The class transformation commutes with translations and the retained spatial spin rotations. A translation of its anchored representative is accompanied by the explicit anchor coboundary. All these operations preserve the stipulated smooth finite-dimensional families and their finite seminorm estimates.

Proof. Conjugation reflects the Fourier frequency, and $n_\varepsilon(\xi)=n_{-\varepsilon}(-\xi)$. The original coupled reality is

$$U_\varepsilon^-(\xi)=16r^4\,\overline{U_{-\varepsilon}^+(-\xi)}.$$

The multiplier proved above gives precisely the displayed dual root equality, with its converse obtained by multiplying by $16r^4$. Root classification and injectivity into actual cohomology turn this equality into the claimed class identity.

To see the representative assertion, choose the positive representative $e_+=(0,S_+(z,z))$, with

$$S_+(x)=(2\pi)^{-3/2}\int(e^{ix\cdot\xi}-1)\sum_\varepsilon U_\varepsilon^+(\xi)\,\mathrm{d}\xi.$$

It is the conjugate-angular counterpart of the dual representative construction, has order ϑ, and represents the original positive class by the conjugate of the theorem's actual comparison. This also compares it with any retained harmonic normal representative from Theorem 3.9. In the bracket integral, conjugation followed by $\xi\mapsto-\xi$ returns the same bracket at a real anchor.

The root identity therefore gives $T_\vee = \overline{S_+}$ and $B_\vee = \overline{S_+(z,z)}$. Both vertical components vanish, proving literal cochain equality for these representatives.

The multiplier is radial on geometrically identified root lines. It commutes with the translation factor $e^{ia\cdot\xi}$ and with spatial spin rotations. The prescribed representative after translation is $T_\vee(x+a) - T_\vee(a)$; its comparison with the literal translate is exactly the anchor coboundary in the theorem. Rotation equivariance follows from the invariant Hodge choices and fixed conic identifications. The continuous linear formulas give the asserted parameter statements. □

The dual class is thus derived from the negative ordinary input; it is not an additional independent class. We retain the spin marking separately because the even projective weights cannot distinguish the two central lifts. On unrestricted actual cohomology, the surviving polynomial kernel prevents defining this operation by root values alone. The construction above is a whole-space weighted operation with actual primitives; local source restrictions are taken after it has been performed.

4.2. Actual contact cochains and the normalized linear response. The fixed inverse needed by the nonlinear correction will be constructed from actual negative cochains. We first define the contact operation used in each auxiliary call. This identifies its complete linear initial graph before any Fourier multiplier is read from that graph.

On the marked flat carrier $N_0 = \mathbb{R}^3 \times \mathbb{CP}^1$, retain the conic and Fourier conventions of the root calculation. In a north spin chart put

$$d = 1 + |\zeta|^2, \qquad f = d^{-2}, \qquad a_0 = \frac{2\bar\zeta}{d}, \qquad b_0 = -\frac{2\zeta}{d},$$
$$n = \frac{i\, z \times \bar z}{z \cdot \bar z}, \qquad \alpha_0 = n \cdot \mathrm{d}x, \qquad R = n \cdot \partial_x,$$
$$V = \partial_{\bar\zeta}, \qquad U = \partial_\zeta, \qquad H = z \cdot \partial_x, \qquad \bar H = \bar z \cdot \partial_x.$$

Here $\mathcal{A}_0 = \mathrm{span}_{\mathbb{C}}(V, H)$ is involutive, $\overline{\mathcal{A}_0} = \mathrm{span}_{\mathbb{C}}(U, \bar H)$, and R is the Reeb field of α_0. For a complex contact Hamiltonian h, fix

$$\alpha_0(X_h) = h, \qquad \iota_{X_h}\mathrm{d}\alpha_0 = (Rh)\alpha_0 - \mathrm{d}h.$$

The Bott connection on $Q = T_{\mathbb{C}}N_0/\mathcal{A}_0$ is $\nabla^B_Y[Z] = [Y, Z] \bmod \mathcal{A}_0$ for $Y \in \mathcal{A}_0$. Let d_B be its differential and set $\mathcal{G}_B h = d_B[X_h]$. A horizontal Bott cocycle is contact-compatible when $\mathrm{d}\alpha_0(\Phi Y, Z) = \mathrm{d}\alpha_0(\Phi Z, Y)$ for $Y, Z \in \mathcal{A}_0$.

The positive and dual positive coefficient lifts are the global scalar operators

$$\mathcal{D}_+ u = -f\bar H u, \qquad \mathcal{D}_- v = -fHv.$$

Their domains are, respectively, $\mathcal{O}(2)$ and its conjugate line on the real sphere. On a north-south overlap, $u_S = \zeta^{-2}u_N$, $v_S = \bar\zeta^{-2}v_N$, $z_S = \zeta^{-2}z_N$, and $d_S = d_N/|\zeta|^2$. The factors cancel in both formulas, since the spatial derivative does not differentiate an angular transition. The lifts therefore act on actual smooth sections, with one spatial derivative and the corresponding symbol-order decrease.

These factors retain the original contact normalization. With the soldering X of §3.3, put $w = X(0,x)\pi$, $\pi = (1,\zeta)$, $\kappa = (1,\bar\zeta)$ and $\eta = -\kappa X(0,x)$. The original contact form $\theta = W\,\mathrm{d}Z$ restricts to $-(d/\sqrt{2})\alpha_0$, and direct differentiation gives

$$\bar H w^0 = -\sqrt{2}d\,\bar\zeta, \quad \bar H w^1 = \sqrt{2}d,$$
$$H\eta^{0'} = \sqrt{2}d\,\zeta, \qquad H\eta^{1'} = -\sqrt{2}d.$$

Consequently the original lifts $F_{w^1} - \bar\zeta F_{w^0}$ and $-G_{\eta^{1'}} + \zeta G_{\eta^{0'}}$, after contact-line conversion by $-\sqrt{2}/d$, give precisely $\mathcal{D}_+$ and $\mathcal{D}_-$ on actual holomorphic restrictions (Isthmus 2026a, §5 and Appendix F.1). The displayed differential formulas define their actual smooth extensions here. The negative lift acts after the derived dual-class transformation; its domain is not the original weight-minus-six line.

Proposition 4.5 (The actual cochain-to-Bott map).
For an actual closed positive cochain $e = (\alpha,\beta)$, whose local equation is $V\beta = H\alpha$, define

$$\Gamma_+(\alpha,\beta)(V) = \frac{f}{2}\big[\bar H(V\alpha - b_0\alpha) + 4R\alpha - 2f\beta\big]\bar H - \frac{f}{2}\big[\bar H(V\beta) + 2R\beta\big]U,$$
$$\Gamma_+(\alpha,\beta)(H) = \frac{f}{2}\big[\bar H(H\alpha) + 2R\beta\big]\bar H - \frac{f}{2}\bar H(H\beta)U.$$

This is a global linear differential operation of order at most two on closed cochains. Its output is a horizontal contact-compatible Bott cocycle, and

$$\Gamma_+ d^0 u = \mathcal{G}_B\mathcal{D}_+ u.$$

It commutes with spatial restriction and smooth passive parameter differentiation. On each finite collection of retained charts, every finite output norm is bounded by finitely many stronger input norms. The conjugate construction

$$\Gamma_\vee(e^\vee) = \overline{\Gamma_+(\overline{e^\vee})}$$

is complex-linear from actual closed dual positive cochains to the corresponding Bott cocycles over $\overline{\mathcal{A}}_0$, with $\Gamma_\vee \bar d^0 = \bar{\mathcal{G}}_B\mathcal{D}_-$, where $\bar{\mathcal{G}}_B h = \overline{\mathcal{G}_B\bar h}$.

Proof. The conic formulas give

$$[V,\bar H] = -b_0\bar H + 2R, \qquad [H,U] = -a_0 H - 2R, \qquad [V,R] = -fH,$$
$$[V,H] = [H,\bar H] = [H,R] = 0, \qquad Vf = b_0 f, \qquad Hf = 0,$$
$$\mathrm{d}\alpha_0(U,H) = -2, \qquad \mathrm{d}\alpha_0(\bar H,V) = 2.$$

The other pairings between these two horizontal frames vanish. Solving the contact-field equations modulo $\mathcal{A}_0$ gives $[X_h] = hR + \frac{1}{2}Hh\,U - \frac{1}{2}Vh\,\bar H$. Taking its Bott derivative yields

$$\mathcal{G}_B h(V) = -\frac{1}{2}(V^2 h - b_0 Vh)\bar H + \frac{1}{2}VHh\,U,$$
$$\mathcal{G}_B h(H) = -\frac{1}{2}HVh\,\bar H + \frac{1}{2}H^2 h\,U.$$

For $h = \mathcal{D}_+ u$, $\alpha = Vu$ and $\beta = Hu$, the commutators give $Vh = -f(\bar{H}\alpha + 2Ru)$ and $Hh = -f\bar{H}\beta$. Substitution, using $VRu = R\alpha - f\beta$, gives the formula for Γ_+ and the exact gauge identity. If $d^0 u = 0$, the displayed expression is zero.

For a general closed cochain, no actual local primitive is required. At a point choose j with $z_j \neq 0$. In the complexified real Taylor ring the analytic coordinates

$$s = \frac{x^j - x_0^j}{z_j(\zeta)}, \qquad y^k = x^k - x_0^k - \frac{z_k(\zeta)}{z_j(\zeta)}(x^j - x_0^j), \quad k \neq j,$$

together with $\zeta, \bar{\zeta}$, make $H = \partial_s$ and $V = \partial_{\bar{\zeta}}$. With $\sigma = \bar{\zeta} - \bar{\zeta}_0$, a formally closed pair has the formal primitive

$$\widehat{u}(s, \sigma) = \int_0^\sigma \widehat{\alpha}(0, r)\, \mathrm{d}r + \int_0^s \widehat{\beta}(t, \sigma)\, \mathrm{d}t.$$

The integrals are coefficientwise. The primitive through order $k+1$ depends linearly on the cochain through order k, with smooth coefficients in the center. A formally CR section is independent of s, σ; truncating in the remaining analytic coordinates realizes each of its finite jets by a local analytic CR section.

Define the comparison at the point by the finite jet of $d_B[X_{\mathcal{D}_+ \widehat{u}}]$. Two choices differ by a formally CR section and contribute zero by the preceding differential identity. This remains true at the second projective pole: the required finite CR jet is realized analytically, the global lift has zero Bott derivative on the north-chart overlap, and continuity gives zero at the pole. The construction is therefore independent of the formal primitive and angular frame. It agrees with the north formula and patches globally. Only the finite jet of the formal primitive enters the comparison at a point. Neither convergence of the series nor a smooth local solution of the primitive equation is needed.

The closedness equation gives

$$\begin{aligned} \mathrm{d}\alpha_0(\Gamma_+ e(V), H) &= f\big(\bar{H}(V\beta) + 2R\beta\big) \\ &= f\big(\bar{H}(H\alpha) + 2R\beta\big) = \mathrm{d}\alpha_0(\Gamma_+ e(H), V). \end{aligned}$$

Thus the cocycle is contact-compatible. Its horizontality is explicit. To verify $d_B \Gamma_+ e = 0$, choose the finite formal primitive one order farther and apply $d_B^2 = 0$ for the involutive flat distribution. This proves an identity of actual finite jets at every point.

The formula is differential with smooth coefficients on a finite angular cover, so the asserted restriction, parameter and finite-order estimates follow. Complex conjugation identifies the positive and dual complexes and interchanges the two distributions. Conjugating the proved identities gives the dual assertions. □

For symmetric complex tensors q, p, write their linear initial-sky graph as

$$\dot\Phi_{q,p}(V) = \beta_1 U, \qquad \dot\Phi_{q,p}(H) = -\beta_1 \bar H + \eta_1 U,$$
$$\beta_1 = \frac{f}{4} q(z,z),$$
$$\eta_1 = \frac{1}{2}(p(z,z) - n_a \Gamma(q)^a{}_{ij} z^i z^j),$$
$$\Gamma(q)^a{}_{ij} = \frac{1}{2}(\partial_i q_{aj} + \partial_j q_{ai} - \partial_a q_{ij}).$$

Traces of q, p do not change this graph, so we choose their trace-free representatives. The symbol η_1 is a graph coefficient, separate from the decay exponent γ. The graph over $\overline{\mathcal{A}}_0$ obtained by replacing (V, H, U, z) by $(U, \bar H, V, \bar z)$, while keeping q, p fixed, is denoted by $\dot{\bar\Phi}_{q,p}$.

Define

$$(\mathsf{C}S)_{ij} = \frac{1}{2}\left(\epsilon_i{}^{k\ell}\partial_k S_{\ell j} + \epsilon_j{}^{k\ell}\partial_k S_{\ell i}\right),$$
$$\mathcal{P}(S) = 2\Delta S - 4\,\mathrm{TF}_\delta(\mathrm{sym}\,\partial\,\mathrm{div}_\delta\, S).$$

Here sym means one half the sum with the two free indices exchanged. On TT tensors, $\mathsf{C}^2 = -\Delta$ and $\mathcal{P}(S) = 2\Delta S$. The first identity follows from the two helicity eigenvalues $\pm|\xi|$ of the Fourier curl symbol; equivalently it is the constant-coefficient curl identity after setting trace and divergence to zero.

Proposition 4.6 (Complete-sky normalization of the positive cochain).

Let $e = (\alpha, \beta)$ be an actual closed positive cochain on a complete sphere bundle. With the fixed fiber Hodge operator, put $u = T_2\alpha$ and write

$$e - d^0 u = (0, S(z,z)), \qquad S^{\mathsf{T}} = S, \quad \mathrm{tr}_\delta\, S = 0.$$

The tensor S is unique and actual smooth. Set $k = -S(n,n)$. Then

$$\Psi_e = \Gamma_+(e) + \mathcal{G}_B(k - \mathcal{D}_+ u) = \dot\Phi_{4i\mathsf{C}S,\,\mathcal{P}(S)}.$$

This is a global contact-compatible Bott cocycle preserving the vertical sky tangent bundle. All operations have finite estimates on the actual coefficient stages and commute with smooth parameters.

Proof. Since $H^1(\mathbb{CP}^1, \mathcal{O}(2)) = 0$, fiber Hodge decomposition gives $Vu = \alpha$. Closedness then makes $\beta - Hu$ a holomorphic quartic. The quadratic restriction map from trace-free symmetric tensors to $H^0(\mathcal{O}(4))$ is an isomorphism, giving the unique S. These are fixed continuous fiber maps, so they preserve actual smoothness and the indicated families.

The cochain-gauge identity reduces Ψ_e to $\Gamma_+(0, S(z,z)) + \mathcal{G}_B k$. Since $Vn = -fz$ and $Vf = b_0 f$,

$$Vk = 2fS(n,z), \qquad (V^2 - b_0 V)k = -2f^2 S(z,z).$$

The $\bar{H}$ component of $\mathcal{G}_B k(V)$ cancels that of $\Gamma_+(0, S(z,z))(V)$. The remaining graph coefficients are

$$\beta_1 = f\big[(\partial_z S)(n,z) - (\partial_n S)(z,z)\big],$$
$$\eta_1 = -\frac{f}{2}(\partial_{\bar{z}}\partial_z S)(z,z) - \frac{1}{2}(\partial_z^2 S)(n,n).$$

In spatial differentiation the angular vectors are held fixed. The identity $n \times z = -iz$ gives $\beta_1 = if(\mathsf{C}S)(z,z)$, the required conic coefficient.

To verify the momentum, rotate the constant frame at one angular point to $n = e_3$, $z = e_1 + ie_2$ and $f = 1$. Put $D = \partial_1 + i\partial_2$, $\bar{D} = \partial_1 - i\partial_2$, $N = \partial_3$, $A = S(z,z)$, $B = S(n,n)$ and $C = S(n,z)$. For $q = 4i\mathsf{C}S$, direct expansion gives

$$q(z,z) = 4(DC - NA), \qquad q(n,z) = \bar{D}A + 3DB - 2NC,$$
$$(\operatorname{div} S)(z) = \tfrac{1}{2}\bar{D}A - \tfrac{1}{2}DB + NC.$$

The coefficient formula gives $2\eta_1 = -\bar{D}DA - D^2B$. Hence

$$\begin{aligned} 2\eta_1 + n_a\Gamma(q)^a{}_{ij}z^i z^j &= 2\eta_1 + Dq(n,z) - \tfrac{1}{2}Nq(z,z) \\ &= 2N^2A + 2D^2B - 4DNC \\ &= 2\Delta A - 4D\big((\operatorname{div} S)(z)\big) = \mathcal{P}(S)(z,z). \end{aligned}$$

These tensor identities are rotation-covariant and respect the conic transitions, so the calculation proves the formula on the entire sphere. Globality and Bott closedness also follow from the preceding proposition and the global Hamiltonian k. Each operation is a fixed fiber map or a finite spatial differential operator. This proves the estimates and parameter assertions. □

Theorem 4.7 (The actual normalized negative graph).
For an actual closed dual positive cochain $e^\vee = (\alpha^\vee, \beta^\vee)$, define

$$v = \overline{T_2\overline{\alpha^\vee}}, \qquad e^\vee - \bar{d}^0 v = (0, T(\bar{z},\bar{z})), \qquad T^{\mathsf{T}} = T, \quad \operatorname{tr}_\delta T = 0.$$

Put

$$\Psi^\vee_{e^\vee} = \Gamma_\vee(e^\vee) + \bar{\mathcal{G}}_B\big(-T(n,n) - \mathcal{D}_- v\big).$$

There is a complex-linear conversion $\mathfrak{C}$ between the global vertical-preserving contact-compatible Bott cocycles over $\overline{\mathcal{A}}_0$ and those over $\mathcal{A}_0$, characterized by $\mathfrak{C}\dot{\bar{\Phi}}_{q,p} = \dot{\Phi}_{q,p}$. The calibrated negative graph is

$$\mathcal{B}_-(e^\vee) = -\mathfrak{C}\Psi^\vee_{e^\vee} = \dot{\Phi}_{q_-,p_-}, \qquad q_- = 4i\mathsf{C}T, \qquad p_- = -\mathcal{P}(T).$$

These are equalities of actual complete-sphere linear graph fields. If $T \in \mathscr{S}^\vartheta$, then $q_- \in \mathscr{S}^{\vartheta-1}$ and $p_- \in \mathscr{S}^{\vartheta-2}$, with finite seminorm estimates and the same statements for all supplied finite parameter derivatives.

Proof. A vertical-preserving contact-compatible cocycle over $\overline{\mathcal{A}}_0$ has the form

$$\Xi(U) = b_1 V, \qquad \Xi(\bar{H}) = -b_1 H + c_1 V.$$

Its Bott equations are $Ub_1 = -a_0 b_1$ and $Uc_1 = \bar{H}b_1$. Since $Uf = -a_0 f$, the first makes $4f^{-1}b_1$ an antiholomorphic quartic on the complete sphere.

Thus $b_1 = fq(\bar z, \bar z)/4$ for a unique trace-free tensor q. Using $Un = -f\bar z$ and $U\bar z = 0$ gives

$$U\left[-\tfrac{1}{2}n_a\Gamma(q)^a{}_{ij}\bar z^i\bar z^j\right] = \bar H b_1.$$

Consequently $2c_1 + n_a\Gamma(q)^a{}_{ij}\bar z^i\bar z^j$ is another antiholomorphic quartic, giving the unique p. Reversing these calculations proves the converse. The same reasoning in the opposite polarization proves the asserted conversion and its global finite angular readout.

It is useful to distinguish this conversion from the antipodal map. At $a\zeta = -1/\bar\zeta$, $n(a\zeta) = -n(\zeta)$, $\bar z(a\zeta) = -\zeta^{-2}z(\zeta)$ and $f(a\zeta) = |\zeta|^4 f(\zeta)$. The converted coefficients are

$$\beta_1 = \frac{\zeta^2}{\bar\zeta^2}b_1(a\zeta), \qquad \eta_1 = \zeta^4 c_1(a\zeta) - n_a\Gamma(q)^a{}_{ij}z^i z^j.$$

The connection term retains the same extrinsic tensor p. Literal antipodal pullback instead yields $\dot\Phi_{q,-p}$. The tensor construction proves extension across both chart poles.

Conjugating the positive normalization, with T held as the resulting actual coefficient, gives $\Psi^\vee_{e^\vee} = \dot{\bar\Phi}_{-4i\mathbf{C}T,\,\mathcal{P}(T)}$. The factor i changes sign; the real differential operator $\mathcal{P}$ does not. Conversion preserves the two tensors. The final minus sign is the original coordinate-operator convention $E = (\Phi^{-1})^*$: the point-map infinitesimal generator is the negative of the coordinate-operator generator. This proves the calibrated graph formula.

The calibration can be checked directly against §3.3. For the derived dual TT representative write $\widehat T = \sum_\varepsilon U_\varepsilon^-/(16r^4)$, using the actual anchored integral when necessary. At the root $n_\varepsilon = \varepsilon\xi/r$, the negative line has curl eigenvalue εr. Thus

$$\widehat{q}_- = \sum_\varepsilon \frac{i\varepsilon}{4r^3}U_\varepsilon^-, \qquad \widehat{p}_- = \sum_\varepsilon \frac{1}{8r^2}U_\varepsilon^-.$$

These are the original negative metric and momentum amplitudes, including their signs. For the selected positive TT quartic representative, with $\widehat S = \sum_\varepsilon U_\varepsilon^+$ up to its constant anchor, the calibrated graph $-\Psi_e$ likewise has $\widehat{q}_+ = \sum_\varepsilon 4i\varepsilon r\,U_\varepsilon^+$ and $\widehat{p}_+ = 2r^2\sum_\varepsilon U_\varepsilon^+$, because the positive root line has curl eigenvalue $-\varepsilon r$. Hence the sum for these selected representatives agrees with the exact two-class seed transform already proved. A retained longitudinal quartic is instead governed by the full tensor formula and its contact gauge; its graph is not determined by root values alone. Fourier reflection enters only when conjugating spatial fields for Lorentz reality; it is not part of the complex-linear conversion.

For the weighted dual representative the single anchored integral is differentiated before the conic tensor is read. Its constant anchor vanishes under these derivatives. The actual weighted dual theorem supplies the finite bounds, and the operators above use at most two additional spatial derivatives. Fixed fiber maps and finite angular coefficient readouts preserve all stated parameter estimates. □

For use with retained representatives, the contact gauge has the elementary tensor readout

$$\mathcal{G}_B(n \cdot W + F) = \dot{\Phi}_{-\mathcal{L}_W \delta,\, \nabla^2 F}.$$

Indeed $V(n \cdot W) = -fz \cdot W$ gives the conic coefficient in the formula for $\mathcal{G}_B$. Its H coefficient is $\frac{1}{2}\partial_z^2(n \cdot W + F)$, and $\Gamma(-\mathcal{L}_W \delta)^a{}_{ij} = -\partial_i \partial_j W^a$ supplies the same connection term. Traces are invisible to $\dot{\Phi}$. This verifies the full graph identity, not only its conic part.

4.3. Actual representative changes in the linear contact response. The actual primitives of the weighted dualization theorem also determine the comparison of retained linear contact responses. We use the fixed-polarization graph $\dot{\Phi}_{q,p}$, its actual cochain realization, and the contact gauge $\mathcal{G}_B$ defined in the linear contact construction. This paragraph concerns that linearized response.

Proposition 4.8 (The contact gauge of an actual dual primitive). Let $e^{\vee\prime} - e^{\vee} = d^0_{\vee} w$ be an actual dual cochain change in the weighted complex above. Write the fiber anti-holomorphic projection of w as

$$w_0 = w - T_2^{\vee} \partial w = Y_i(x) \bar{z}^i.$$

Then the calibrated negative graph satisfies

$$\mathcal{B}_-(e^{\vee\prime}) - \mathcal{B}_-(e^{\vee}) = \mathcal{G}_B h_-, \qquad h_- = -i\, n \cdot \operatorname{curl} Y + \frac{2}{3} \operatorname{div} Y.$$

For a positive representative change whose holomorphic projection is $Y_+ \cdot z$, the calibrated positive graph has

$$h_+ = i\, n \cdot \operatorname{curl} Y_+ + \frac{2}{3} \operatorname{div} Y_+.$$

Thus $Y = \overline{Y_+}$ gives a real paired Hamiltonian $h_+ + h_-$. Strict additive equations for supplied degree-zero cochain comparisons give strict additive equations for these Hamiltonians. All formulas are global on the complete sphere and obey finite symbol and parameter estimates.

Proof. Fiber normalization of the change gives

$$\delta\beta_0 = D^{\vee} w - D^{\vee} T_2^{\vee} \partial w = D^{\vee} w_0,$$
$$\delta T = \mathcal{K} Y = \mathrm{TF}_{\delta}(\operatorname{sym} \partial Y).$$

Here $\mathcal{K}$ is the tracefree symmetric gradient already defined in the polynomial cohomology calculation, with sym containing the factor $1/2$. To compute the calibrated graph, retain the operators $\mathsf{C} = \operatorname{curl}_{\mathrm{sym}}$ and $\mathcal{P}(T) = 2\Delta T - 4\,\mathrm{TF}_{\delta}(\operatorname{sym} \partial \operatorname{div} T)$ from the linear contact theorem. Direct differentiation gives

$$\mathsf{C}(\mathcal{K} Y) = \frac{1}{2} \operatorname{sym} \partial(\operatorname{curl} Y),$$
$$\operatorname{div}(\mathcal{K} Y) = \frac{1}{2} \Delta Y + \frac{1}{6} \operatorname{d}(\operatorname{div} Y),$$
$$\mathcal{P}(\mathcal{K} Y) = -\frac{2}{3}\, \mathrm{TF}_{\delta}\, \nabla^2(\operatorname{div} Y).$$

For the first identity, insert $(\mathcal{K}Y)_{ij} = (\partial_i Y_j + \partial_j Y_i)/2 - \delta_{ij}\partial_a Y_a/3$. Terms with two commuting derivatives contracted against the alternating tensor vanish, and the symmetrized trace terms cancel. The surviving terms are $\frac{1}{4}(\partial_i(\operatorname{curl} Y)_j + \partial_j(\operatorname{curl} Y)_i)$. The second identity follows by taking the divergence of the same formula. Substituting it into $\mathcal{P}$ cancels the two $\mathrm{TF}_\delta(\operatorname{sym}\partial\Delta Y)$ terms and gives the third identity.

The calibrated negative graph has $q_- = 4i\mathsf{C}T$ and $p_- = -\mathcal{P}(T)$, so its change is

$$\delta q_- = 2i \operatorname{sym}\partial(\operatorname{curl} Y), \qquad \delta p_- = \frac{2}{3}\,\mathrm{TF}_\delta\,\nabla^2(\operatorname{div} Y).$$

Use the scalar-vector contact identity

$$\mathcal{G}_B(n\cdot W + F) = \dot{\Phi}_{-\mathcal{L}_W\delta,\,\nabla^2 F}$$

with $W = -i\operatorname{curl} Y$ and $F = \frac{2}{3}\operatorname{div} Y$. The metric tensor is exactly δq_- because $\mathcal{L}_W\delta = 2\operatorname{sym}\partial W$. The momentum differs from δp_- only by a trace, which vanishes upon contraction with $z\otimes z$ in the graph. This proves the negative formula. The calibrated positive graph uses $q_+ = -4i\mathsf{C}S_+$ and the same $p_+ = -\mathcal{P}(S_+)$; the identical calculation uses $W = i\operatorname{curl} Y_+$ and gives the positive formula.

The vector n is globally defined on the marked sphere, and the remaining tensors are actual fields on the initial space. The Hamiltonians therefore patch globally. Their expressions are linear in the primitive through its fixed Hodge projection and tensor readout. This proves strict additivity. Complex conjugation gives $h_- = \overline{h_+}$ for the stated paired change, hence a real sum.

If $w \in \mathscr{S}^{\vartheta+1}$, then Y has that order and $h_- \in \mathscr{S}^{\vartheta}$. The metric and momentum changes have orders $\vartheta - 1$ and $\vartheta - 2$, respectively. Their finite seminorms use at most the displayed finite number of spatial derivatives and the compact-fiber projection. The same linear formulas commute with every supplied finite parameter derivative. □

When the constructor first takes its fixed canonical cohomology representative, an ordinary original coboundary leaves that canonical input and its response unchanged; the comparison is then the identity. The proposition instead applies to actual representatives that are retained, including a supplied dual cochain and the selected representative with the same class. It proves equality of actual linearized graph fields and their infinitesimal contact gauges. The nonlinear original point maps and their finite compositions require the separate nonlinear comparison laws.

4.4. A local encoder for actual negative cochains. Let $\mathscr{J}_{-6}$ be the harmonic Serre-dual identification from trace-free symmetric tensors to $H^1(\mathcal{O}(-6))$, normalized by

$$\langle \mathscr{J}_{-6}(B), Q(z,z)\rangle_{\mathrm{Serre}} = B : Q.$$

The Serre integral is $(2\pi i)^{-1}\int p(\zeta)\,\mathrm{d}\zeta\wedge\alpha$, with the contour orientation of §3.3. For an actual weighted TT field $B\in\mathscr{S}^{\gamma-2}_{\mathrm{TT}}$, with the fixed $-1<\gamma<-1/2$, define

$$e_-(B)=(\mathscr{J}_{-6}(B),T_{-4}(z\cdot\partial_x)\,\mathscr{J}_{-6}(B))\,.$$

Proposition 4.9 (The harmonic negative input and its exact roots).

The cochain $e_-(B)$ is exactly closed and preserves the spatial support of B. Its original negative root tensors obey

$$U^-_+(\xi)+U^-_-(\xi)=2i\mathsf{C}(\xi)\widehat{B}(\xi).$$

With $\vartheta=\gamma+1$, its coefficients have orders $\vartheta-3$ and $\vartheta-4$, so it lies in the actual weighted negative complex used by the dual transformation. The differential closedness and support assertions alone also hold for arbitrary smooth TT fields.

Proof. Under the normalized Serre identifications, the fiber harmonic projection $H^1_{-4}(z\cdot\partial_x)\,\mathscr{J}_{-6}(B)$ is exactly $\operatorname{div}B$, as in the polynomial cohomology calculation. This is an identity of constant-coefficient spatial differential operators and fixed fiber maps, so it applies to all actual smooth tensors. It vanishes for TT B. The fiber Hodge identity therefore gives $\bar\partial T_{-4}(z\cdot\partial_x)\,\mathscr{J}_{-6}(B)=(z\cdot\partial_x)\,\mathscr{J}_{-6}(B)$, which proves closedness. Both operations are local in the spatial variable and have the displayed derivative orders.

On the two-dimensional TT plane at $\xi\neq0$, the sum of roots defines a rotation-equivariant smooth homogeneous multiplier of degree one. The root sum is odd under $\xi\mapsto-\xi$: its horizontal coefficient contains $i\xi\cdot z$, and root reversal just interchanges the two summands. The stabilizer of a unit covector has distinct characters on the two helicity lines. Hence the multiplier is $a|\xi|I+b\mathsf{C}(\xi)$ for constants a,b; oddness forces $a=0$.

The remaining constant is fixed by the plane used in §3.3. For $Ae^{\lambda w^0}/\zeta$, $\xi=\varepsilon re_3$ and $\lambda=-i\sqrt{2}\,\varepsilon r$, the corrected first CR coefficient is $Ae^{\lambda u}\zeta^{-1}\bar\partial\rho_N$. The normalized contour pairing against $Q(z,z)$ gives $Ae^{\lambda u}Q(z_0,z_0)$, where $z_0=(1,i,0)$. Thus its harmonic tensor is $Ae^{\lambda u}z_0\otimes z_0$. Its original root tensor is $2i\varepsilon r$ times that tensor, while the curl eigenvalue on this line is εr. Therefore $b=2i$. This proves the root formula with its normalization. The actual symbol estimates follow directly from the finite differential and fiber operations. □

Proposition 4.10 (The actual normalized negative response).

For $B\in\mathscr{S}^{\gamma-2}_{\mathrm{TT}}$, form $e_-(B)$, apply the actual negative-to-dual operation with its prescribed anchored TT representative, apply the normalized contact graph above, and read its conic tensor. Denote this operation by $\mathcal{N}_-(B)$. It is a continuous complex-linear map

$$\mathcal{N}_-:\mathscr{S}^{\gamma-2}_{\mathrm{TT}}\longrightarrow\mathscr{S}^{\gamma}_{\mathrm{TT}},\qquad\widehat{\mathcal{N}_-(B)}(\xi)=-\frac{1}{2|\xi|^2}\widehat{B}(\xi),\quad\xi\neq0.$$

The equality is between actual decaying tensors. The response is unchanged by ordinary coboundaries before the canonical input projection, and commutes with translations, spin rotations, conjugation and every specified finite parameter derivative. Retained noncanonical dual representatives have the actual comparison primitives of the dual theorem and their linear contact gauges.

Proof. The preceding proposition proves that the actual input is closed and belongs to the stated negative domain. Its root sum and the calibrated multiplier $1/(16r^4)$ give the dual coefficient tensor

$$\widehat{T}(\xi) = \frac{i}{8r^4}\mathsf{C}(\xi)\widehat{B}(\xi).$$

The normalized actual graph has conic tensor $q_- = 4i\mathsf{C}T$. Since $\mathsf{C}(\xi)^2 = r^2 I$ on TT tensors, its multiplier is exactly $-1/(2r^2)$. Thus the multiplier is computed from the specified cochain and contact operations.

Write $\vartheta = \gamma + 1 \in (0, 1/2)$. Since $B \in \mathscr{S}^{\vartheta-3}$ with $\vartheta - 3 \in (-3, 0)$, the weighted Fourier lemma gives differentiated low-frequency bounds $r^{-\vartheta-|\alpha|}$ and rapid decrease at infinity. The conic amplitude has order $r^{-2-\vartheta}$, which is locally integrable in three dimensions. The curl of the anchored dual integral is precisely this inverse integral, because it annihilates the constant subtraction. The inverse Fourier estimate gives the actual order $\vartheta - 1 = \gamma$ with finite seminorm bounds. The TT identities hold for the locally integrable extensions. There is no added polynomial because the resulting field decays.

Canonical projection makes ordinary input coboundaries identical before the graph is constructed. For retained representatives, the dual theorem supplies actual primitives and the stated contact comparison. The final scalar multiplier is real, even and rotation-invariant, so translations, conjugation and spin rotations commute with the actual response. In particular the intermediate anchor has disappeared. Applying the same continuous linear maps to each finite parameter derivative proves the family assertion. □

The auxiliary calls are made in the complex source domain. A negative-only call is not required to be a Lorentz-fixed pair. Its conic response is nevertheless real on real tensor inputs by the proved multiplier, and the final contractions preserve the real residual fields. The actual spin marking is retained separately throughout.

4.5. The fixed inverse from finitely many negative responses. Define the local fourth-order operator on symmetric tensors by

$$D_{ij} = \delta_{ij}\Delta - \partial_i\partial_j, \qquad (\mathcal{R}_4 T)_{ij} = \left(D_{ik}D_{j\ell} - \frac{1}{2}D_{ij}D_{k\ell}\right)T_{k\ell}.$$

The identities $\partial_i D_{ij} = 0$, $\operatorname{tr} D = 2\Delta$ and $D^2 = \Delta D$ show directly that its output is symmetric, trace-free and divergence-free. At $\xi \neq 0$, put

$$P = I - \frac{\xi \otimes \xi}{r^2}, \qquad \Pi_{ij,k\ell} = \frac{1}{2}(P_{ik}P_{j\ell} + P_{i\ell}P_{jk}) - \frac{1}{2}P_{ij}P_{k\ell}.$$

Then $\mathcal{R}_4$ has polynomial symbol $r^4\Pi$. The displayed differential formula implements this symbol without first applying a nonlocal TT projection.

Choose a real orthonormal basis E_A, $1 \le A \le 5$, of trace-free symmetric tensors. For a scalar $g \in \mathscr{S}^{\gamma-2}$, form

$$F_A^{(0)}(g) = \mathcal{R}_4(E_A g), \qquad F_A^{(j+1)}(g) = \mathcal{N}_-(F_A^{(j)}(g)), \quad j = 0, 1, 2.$$

After each response its conic tensor is read before the next negative cochain is formed. The auxiliary graph also has a momentum tensor, which is not asserted to vanish.

Lemma 4.11 (Actual domains of the three successive responses). All four tensors are TT and obey

$$F_A^{(j)}(g) \in \mathscr{S}^{\gamma-6+2j}, \qquad \widehat{F}_A^{(j)}(\xi) = \left(-\frac{1}{2}\right)^j r^{4-2j}\Pi(\xi)E_A\widehat{g}(\xi), \quad 0 \le j \le 3.$$

Every prescribed finite output seminorm is bounded by finitely many input seminorms, uniformly in this finite basis and for specified finite parameter derivatives. The actual coefficient domains in the three calls are as follows. The last column gives the orders of both tensors in the auxiliary contact graph.

Call	TT input	Negative first and second coefficients	Conic and momentum tensors
0	$F_A^{(0)} \in \mathscr{S}^{\gamma-6}$	$\mathscr{S}^{\gamma-6} \oplus \mathscr{S}^{\gamma-7}$	$\mathscr{S}^{\gamma-4} \oplus \mathscr{S}^{\gamma-5}$
1	$F_A^{(1)} \in \mathscr{S}^{\gamma-4}$	$\mathscr{S}^{\gamma-4} \oplus \mathscr{S}^{\gamma-5}$	$\mathscr{S}^{\gamma-2} \oplus \mathscr{S}^{\gamma-3}$
2	$F_A^{(2)} \in \mathscr{S}^{\gamma-2}$	$\mathscr{S}^{\gamma-2} \oplus \mathscr{S}^{\gamma-3}$	$\mathscr{S}^{\gamma} \oplus \mathscr{S}^{\gamma-1}$

Each negative cochain lies in $\mathscr{C}^1_{-,\vartheta}$ for the same $\vartheta = \gamma + 1$. Each selected dual cochain lies in $\mathscr{C}^1_{\vee,\vartheta}$, with actual tensor coefficient of order ϑ, and the dual comparison primitives lie in $\mathscr{C}^0_{\vee,\vartheta}$ of order $\vartheta + 1$. The selected coefficients, their anchor-comparison primitives, and the displayed graph tensors have finite seminorms bounded by finitely many seminorms of g. For a separately supplied representative, the comparison bounds also use finitely many seminorms of that representative.

Proof. The differential encoder gives the first assertion and its symbol at $j = 0$. Its order is at most $\gamma - 2$, so the normalized response applies. Inductively the Fourier multiplier is the stated one. It remains to prove the stronger orders needed to apply the next response without losing the domain.

Put $\vartheta = \gamma + 1$. The scalar amplitude has low-frequency estimates $r^{-\vartheta-|\alpha|}$, since its input order $\vartheta - 3$ lies in $(-3, 0)$. With $\mu = 4 - 2j$, the explicit amplitude $r^\mu\Pi E_A\widehat{g}$ therefore has differentiated bounds $r^{\mu-\vartheta-|\alpha|}$ and rapid decrease at infinity. It is in the amplitude space of the one-way inverse Fourier estimate with exponent $\beta = \vartheta - 3 - \mu = \gamma - 6 + 2j < 0$. That

estimate was proved for every negative β, so it applies also to the first two exponents below -3. It gives the stated spatial orders and finite seminorm bounds. Equivalently, a low-frequency annulus $r \asymp 2^{-k}$ contributes at most

$$C_{\alpha,N}\, 2^{-k(3+\mu-\vartheta+|\alpha|)}(1+2^{-k}|x|)^{-N}$$

to the absolute value of its αth spatial derivative; splitting at $2^k \asymp \langle x\rangle$ gives the same bound. For all four values of j, $3+\mu-\vartheta>0$, so the amplitudes and their spatial Fourier multiples are integrable at zero.

The improved orders follow from the explicit factors in the encoded amplitude at low frequency. The gain thus applies to these encoded inputs. Faster decay alone would not justify it for a general Newton potential. At each stage the source response and the displayed inverse transform are decaying actual fields with the same punctured Fourier transform. Their difference is a polynomial and hence zero. This proves the exact identities and the zero-frequency normalization of every conic output.

For the other intermediate objects, the local harmonic encoder preserves the first coefficient's spatial order and lowers the second by one. All its inputs have order at most $\gamma-2$, so the fixed weighted dual theorem applies to every call. It gives the dual representatives and comparison primitives in the common stated spaces, with finite seminorm bounds. In calls 0 and 1, the dual tensor is more precisely $\tau-\tau(0)$, where the same inverse-frequency estimate gives $\tau\in\mathscr{S}^{\gamma-3+2j}$. The constant subtraction is retained; these two anchored tensors are not claimed to have negative spatial order. Their first and second derivatives do have the improved orders, and the graph formulas give the table. In call 2, the once-anchored integral has order ϑ and its differentiated integral gives the same conclusion. The degree-zero primitive of any anchor change is the explicit linear cochain of the dual theorem and has allowed order $\vartheta+1$.

All these operations are continuous linear maps on the displayed symbol stages. For each finite list of output, fiber and parameter derivatives, one first chooses the finite input seminorms required by the dual theorem and the finite differential operations, and then those required by the inverse-frequency estimate. For retained representative comparisons, that list includes the supplied representative. The composition proves the finite primitive and coefficient bounds for each of the three calls. The list is finite and the weight and anchor are unchanged. This proves the parameter assertion. □

Set $\mathcal{A}_A(g)=F_A^{(3)}(g)$. Define scalar and vector operations by

$$\mathcal{N}^{\mathrm{src}}g = 4\sum_A E_A : \mathcal{A}_A(g),$$

$$(\mathcal{M}F)_i = \sum_{A,j,k}(E_A)_{kj}[\mathcal{A}_A(F_k)]_{ij},$$

$$\mathcal{V}^{\mathrm{src}}F = \frac{3}{4}\mathcal{N}^{\mathrm{src}}F + 2\mathcal{M}F,$$

where the scalar operation in the last line acts componentwise.

Theorem 4.12 (The complete flat inverse from negative source responses).

On actual weighted smooth scalar and vector fields of order $\gamma - 2$,

$$\mathcal{N}^{\mathrm{src}} = \Delta^{-1}, \qquad \mathcal{V}^{\mathrm{src}} = \Delta_L^{-1},$$

with the decaying normalization of Lemma A.2. Thus the complete inverse of the flat constraint correction operator is

$$\mathsf{G}_0^{\mathrm{src}} = \operatorname{diag}\left(-\frac{1}{2}\mathcal{N}^{\mathrm{src}}, \mathcal{V}^{\mathrm{src}}\right) = \mathsf{A}_0^{-1}, \qquad \mathsf{A}_0 = \operatorname{diag}(-2\Delta, \Delta_L).$$

The operations are independent of the auxiliary orthonormal basis, preserve real inputs, and commute with Euclidean translations and rotations together with the supplied spatial spin lifts. Each smooth operation uses a local differential encoder, finitely many actual normalized negative responses, and finite coefficient contractions. A spin lift is retained as part of the marking; it is not recovered from the tensor contraction.

Proof. The preceding lemma gives the multiplier $-\Pi E_A/(8r^2)$ for $\mathcal{A}_A$. Since Π is the rank-two orthogonal projector on the TT plane in the five-dimensional trace-free symmetric space,

$$\sum_A E_A : \Pi E_A = \operatorname{tr}\Pi = 2.$$

Therefore $\mathcal{N}^{\mathrm{src}}$ has multiplier $-r^{-2}$. Its actual output has order γ and decays. Its difference from the decaying inverse of Δ has zero punctured Fourier transform, hence is a decaying polynomial and is zero.

Completeness of the basis gives $\sum_A [\Pi E_A]_{ij}(E_A)_{k\ell} = \Pi_{ij,k\ell}$. The partial trace is

$$\sum_j \Pi_{ij,kj} = \frac{1}{2}(P_{ik}\operatorname{tr}P + (P^2)_{ik}) - \frac{1}{2}(P^2)_{ik} = P_{ik}.$$

Hence $\mathcal{M}$ has multiplier $-P/(8r^2)$ and

$$\widehat{\mathcal{V}^{\mathrm{src}}}(\xi) = -\frac{1}{r^2}\left(\frac{3}{4}I + \frac{1}{4}P\right) = -\frac{1}{r^2}\left(I - \frac{1}{4}\frac{\xi\otimes\xi}{r^2}\right).$$

This is the inverse of the Fourier symbol $-r^2(I + \frac{1}{3}\xi\otimes\xi/r^2)$ of Δ_L. Decay and the same polynomial-kernel argument identify the actual output with its uniquely normalized inverse. The scalar factor $-1/2$ gives the inverse of -2Δ.

The basis sums are invariant traces and partial traces, so the result is basis-independent. The encoder and final contractions are real and rotation-equivariant, and the normalized response has the required conjugation and covariance laws. There are five chains of three responses for one scalar component, and finitely many such chains for the vector operation. The preceding lemma has verified the actual cochain domain and finite estimates of every call. No extra Green or biharmonic primitive is used by the execution. □

The encoder takes four derivatives, so the individual stages need not act boundedly on the base residual Hölder space. After contraction the exact

identity identifies the composite with the bounded inverse $\mathcal{Z}_m \to \mathcal{Y}_m$, with its two-derivative estimate. The identity specifies the realization at finite base order directly, without a density argument in the large weighted Hölder space.

At a smooth current state Z, the fixed inverse in the current Neumann correction can therefore be replaced by this equal source operation:

$$\mathsf{E}_Z = \mathsf{G}_0^{\mathrm{src}}(\mathsf{A}_Z - \mathsf{A}_0), \qquad \mathsf{G}_Z^{\mathrm{src}} = \sum_{j\geq 0}(-\mathsf{E}_Z)^j \mathsf{G}_0^{\mathrm{src}}.$$

The current inverse identity uses the single base margin in Proposition A.13. At smooth coefficients, the preceding formulas provide cochain representatives for every finite word. The series converges in the base operator norm. At higher orders one estimates the same inverse solution through its equation, without taking a limit of the partial sums in those norms.

On the analytic domain the comparison identifies the calibrated linear negative response. For smooth input the anchored cochains and the graph formulas on the complete sphere define each auxiliary response chain directly, without summing a holomorphic Hamiltonian series. The construction therefore needs no stagewise identification with the nonlinear analytic constructor.

4.6. Actual initial source coefficients. The ordinary classes will generate two uncontracted chiral gradient operators and a complete contact graph on the initial sphere bundle. We first specify these fields and their independent kinematic tests, before imposing the vacuum constraints. The physical scale is retained by the gradient symbol and, in the contact presentation, by the volume density. The initial spin bundle and its actual marking are part of both presentations.

Use the ambient weighted coordinates $Z = (s, y)$ of Proposition A.4, restricted to the fixed product D of the current inverse proposition. Thus

$$h = \delta + q + u\delta,$$
$$\kappa = p + L_0 W,$$
$$K = \mathrm{TF}_h\, \kappa,$$
$$s = (q, p), \quad y = (u, W).$$

These tensors are auxiliary coefficients of a source operation. Its output will consist of actual operators and a contact graph, with explicit inverse readouts. This distinction is needed to verify that tensor correction and source correction have the same meaning.

The smooth carriers below use the intersection of the stated weighted spaces. At a finite order, a carrier and its tangents mean their corresponding coefficient completions, with the same exact identities wherever defined. The base order is high enough for all displayed curvature tests. Thus a residual direction of finite regularity produces an actual tangent of that regularity; weighted smooth residuals produce weighted smooth tangents.

Fix the trivial real Lorentz bundle $\mathbb{E} = \mathbb{R}^3 \times \mathbb{R}^{1,3}$, with signature $(+---)$, future unit vector $N = E_0$ and oriented spatial orthonormal vectors E_a. Fix its actual spin reduction and complex Dirac module $\mathscr{S} = \mathscr{S}_+ \oplus \mathscr{S}_-$, with the invariant nondegenerate complex alternating form B. The Clifford conventions are

$$c(v)c(w) + c(w)c(v) = 2\eta(v,w)I, \qquad \operatorname{tr}_{\mathscr{S}}\big(c(v)c(w)\big) = 4\eta(v,w).$$

The Clifford vectors generate the full complex endomorphism algebra. For $\Omega \in \mathfrak{so}(1,3)$, let $\varrho_\pm(\Omega)$ be the chiral restrictions of $\frac{1}{4}\Omega^{\mu\nu}c(E_\mu)c(E_\nu)$, where $\Omega^{\mu\nu} = \Omega^\mu{}_\lambda \eta^{\lambda\nu}$. Thus $[\varrho(\Omega), c(v)] = c(\Omega v)$. The choice of actual spin bundle is retained even when its central sign acts trivially on a vector or contact coefficient.

Definition 4.13 (Canonical kinematic gradient carrier).
Supply two actual first-order complex-linear operators $\mathbb{D}_\pm : \Gamma(\mathscr{S}_\pm) \to \Gamma(\mathbb{E}_{\mathbb{C}} \otimes \mathscr{S}_\pm)$. Their commutators with real scalar multiplication must have a common symbol given by a real orientation-preserving bundle isomorphism $b : T^*\mathbb{R}^3 \xrightarrow{\sim} N^\perp$, and their complete outputs must be tangent to $N^\perp$:

$$[\mathbb{D}_\pm, m_f]\psi = b(\mathrm{d}f) \otimes \psi, \qquad \eta(N, \mathbb{D}_\pm\psi) = 0.$$

Define

$$h^{-1}(\xi,\zeta) = -\eta(b\xi, b\zeta), \qquad \Theta X = b(h^\flat X),$$
$$P_X^\pm\psi = -\eta(\Theta X, \mathbb{D}_\pm\psi), \quad P_X = P_X^+ \oplus P_X^-.$$

Require, for real X and real Clifford sections v,

$$X\mathsf{B}(\psi,\chi) = \mathsf{B}(P_X\psi,\chi) + \mathsf{B}(\psi, P_X\chi),$$
$$[P_X, c(v)] \in \Gamma(c(\mathbb{E})).$$

As proved below, $c(D_X v) = [P_X, c(v)]$ then defines a real Lorentz connection. Require the torsion test $D\Theta = 0$, and set

$$K(X,Y) = -\eta(D_X N, \Theta Y).$$

The canonical frame condition is that the matrix of b in the fixed coordinate covectors and the fixed spatial vectors is the positive matrix $h^{-1/2}$. Require the decoded (h, K) to satisfy the weighted ambient gauge and maximality tests of Definition A.3, and its exact ambient decoder to belong to D. No curvature-image test is imposed in this kinematic carrier. Its seed readout π is the s component of that ambient decoder.

These tests use only the supplied fields and therefore apply before any choice of a generating ordinary pair.

The ambient seed readout is evaluated on the whole weighted field. A later local restriction restricts the constructed operators and graph; it does not apply this whole inverse to a spatially truncated field.

Lemma 4.14 (Exact gradient construction and readout).
Write $F = h^{1/2}$, $e = F^{-1}$, $\Gamma_i = (\Gamma^a{}_{ib})$, and $K_i = (K_{ij})_j$ as a row. Define

$$\Theta X = (0, FX),$$
$$\Omega_i = \begin{pmatrix} 0 & K_i e \\ e K_i^{\mathsf{T}} & F\Gamma_i e - (\partial_i F)e \end{pmatrix},$$
$$\mathbb{D}_\pm(Z) = \sum_{a,i} E_a \otimes e_a^i(\partial_i + \varrho_\pm(\Omega_i)).$$

Here e_a^i are the components of the oriented h-orthonormal frame e_a. These formulas define a map $\mathrm{Enc}_{\mathrm{gr}}$ from D to the independently tested kinematic carrier. Its exact inverse is the source readout followed by Proposition A.4. At every finite weighted order the forward and inverse maps are smooth, with finite estimates on retained bounded sets.

More precisely, measure a coefficient difference by the $C^{m,\alpha}_\gamma$ norm of its symbol and the $C^{m-1,\alpha}_{\gamma-1}$ norms of both zero-order coefficients. Denote the resulting norm by $\|\cdot\|_{\mathrm{op},m}$. On such sets,

$$\|\delta \,\mathrm{Enc}_{\mathrm{gr}}\|_{\mathrm{op},m} \asymp \|\delta h\|_{C^{m,\alpha}_\gamma} + \|\delta K\|_{C^{m-1,\alpha}_{\gamma-1}} \asymp \|\delta s\|_{\mathcal{X}_m} + \|\delta y\|_{\mathcal{Y}_m}.$$

The assertion holds for tangents and finite differences, with constants depending on the fixed positive margin and finitely many coefficient bounds.

Proof. The scalar commutator in the definition gives the Leibniz rule for P_X, since $-\eta(\Theta X, b\,\mathrm{d}f) = Xf$. It also gives function-linearity in X, so P is a connection. Its commutators with real Clifford vectors define a real connection D. Differentiate the Clifford relation by P to see that $D\eta = 0$. The difference between P and the spin lift of D commutes with every Clifford vector, hence is a scalar one-form because those vectors generate the full matrix algebra. Preservation of B makes this scalar zero. Thus P is the exact spin connection, with no undetected scalar connection term.

For any oriented h-orthonormal frame, tangential output gives the full recovery identity

$$\mathbb{D}_\pm\psi = \sum_a \Theta(e_a) \otimes P^\pm_{e_a}\psi.$$

Indeed, contraction with each $\Theta(e_a)$ recovers its corresponding component, and the output has no component along N.

The connection is Lorentz, so $D_X N$ is spatial and defines K by the displayed readout. The normal component of $D\Theta = 0$ says that K is symmetric. Its tangential component says that the metric connection induced on $T\mathbb{R}^3$ is torsion-free, hence is ∇^h. Consequently

$$D_X N = \Theta(K^\sharp X), \qquad D_X \Theta Y = \Theta(\nabla^h_X Y) + K(X,Y)N.$$

In the positive frame these identities are exactly the displayed matrices Ω_i. Conversely those matrices obey the same identities; their spatial block is skew because $\nabla^h h = 0$. They therefore define a Lorentz connection whose spin lift preserves B, the real Clifford action and the torsion test. The

constructed symbol is $h^{-1/2}$, so every canonical-frame condition holds. The two constructions are inverse by the full recovery identity and the unique Gauss–Weingarten formula.

The forward formulas use finite matrix operations, one derivative of h, and the value of K. The inverse reads h algebraically from the symbol, and K from the directional connection and the fixed N. Their weighted product estimates give the first equivalence. Proposition A.4, including its bounded decaying York inverse, gives the second. The same argument applied to differences or derivatives gives the stated smoothness and estimates. □

Proposition 4.15 (Complete contact coefficients and their exact inverse).

On $N_0 = P_+T^*\mathbb{R}^3 = \mathbb{R}^3 \times S^2$, retain $\alpha_0 = n \cdot \mathrm{d}x$, its contact distribution and its projection to $\mathbb{R}^3$. In an affine angular chart put

$$z = (1-\zeta^2, i(1+\zeta^2), 2\zeta), \qquad d = 1 + |\zeta|^2,$$
$$n = d^{-1}(-2\operatorname{Re}\zeta, -2\operatorname{Im}\zeta, 1-|\zeta|^2), \qquad Z_0 = \partial_{\bar\zeta}, \quad W_0 = z \cdot \partial_x.$$

Thus $\partial_{\bar\zeta} n = -z/d^2$, $z \cdot \bar z = 2d^2$ and $n \times z = -iz$. For an ambient pair define

$$A = h(z,z), \qquad B = h(z,\bar z), \qquad \Delta_c = (B^2 - |A|^2)^{1/2},$$
$$\beta = \frac{A}{B+\Delta_c}, \qquad \ell = z - \beta\bar z, \qquad \rho_n = (n^{\mathsf T} h^{-1} n)^{1/2},$$
$$\eta_c = \frac{1}{2}\left(\rho_n K(\ell,\ell) - n_a \Gamma^a{}_{ij}\ell^i\ell^j\right), \quad \omega = \frac{1}{2}\log\det h.$$

The graph spanned by

$$Z_0 + \beta\bar Z_0, \qquad W_0 - \beta\bar W_0 + \eta_c \bar Z_0$$

is a real transverse, contact-compatible, integrable CR distribution. The formulas agree on a finite rotated angular cover of the whole sphere. Together with the density $e^{\omega}\mathrm{d}^3x$, the fixed value $\tau = 0$ of the mean curvature, and the actual spin marking, they define $\mathrm{Enc}_{\mathrm{CR}}(Z)$.

There is an independent test and exact decoder for these complete coefficients. Fix a sufficiently small uniform graph neighborhood of the flat graph, reducing D once at the base order if needed so that its encodings lie there. Require the supplied graph to lie in this neighborhood, to have $|\beta| < 1$ and the displayed triangular form on every chart, and to obey angular compatibility and real transversality. At three fixed sphere directions with flat null vectors proportional to $(1,i,0)$, $(1,0,i)$ and $(0,1,i)$, use the supplied β to form ℓ. Solve five independent real equations among $h^\sharp(\ell,\ell) = 0$, together with $\operatorname{tr}_\delta h^\sharp = 3$. Require the candidate $h^\sharp$ to be positive and to reproduce the supplied conic at every angular direction. Set

$$h = \left(\frac{e^{2\omega}}{\det h^\sharp}\right)^{1/3} h^\sharp.$$

Using this metric, solve five independent real null-square equations

$$K(\ell,\ell) = \rho_n^{-1}\big(2\eta_c + n_a\Gamma^a{}_{ij}\ell^i\ell^j\big)$$

together with $\operatorname{tr}_h K = 0$. Require this real symmetric candidate to reproduce the supplied shear at every angular direction. Finally apply the ambient gauge, decay and D tests to the decoded tensors. These are the complete kinematic contact tests. They require no vacuum equation and no prior membership in an ordinary-pair image. They are equivalent to exact reconstruction by $\mathrm{Enc}_{\mathrm{CR}}$.

For each fixed finite angular order l, give coefficient differences the balanced norm

$$\|\delta Q\|_{\mathcal{B}_{m,l}} = \|\delta\beta\|_{C^{m,\alpha}_{\gamma}(C^l_{\mathrm{ang}})} + \|\delta\eta_c\|_{C^{m-1,\alpha}_{\gamma-1}(C^l_{\mathrm{ang}})} + \|\delta\omega\|_{C^{m,\alpha}_{\gamma}}.$$

On retained bounded sets this norm is equivalent, in both finite-difference and tangent estimates, to $\|\delta h\|_{C^{m,\alpha}_{\gamma}} + \|\delta K\|_{C^{m-1,\alpha}_{\gamma-1}}$. The forward and inverse maps preserve smooth finite-dimensional families with finite seminorm estimates.

Proof. Positivity of h on the real plane spanned by $\operatorname{Re} z, \operatorname{Im} z$ gives $B > 0$ and $B^2 - |A|^2 > 0$. The chosen root has $|\beta| < 1$, and $h(\ell,\ell) = A - 2B\beta + \bar{A}\beta^2 = 0$. Its line is the oriented $-i$ screen eigenspace, since it equals the flat line at $h = \delta$ and positive metrics form a connected set.

For an intrinsic verification of the entire graph, let $u = h^{-1}n/\rho_n$ and $j_h X = u \times_h X$ on $E_n = \ker n$. Define

$$M_h w = \rho_n^{-1}h^{-1}w - \rho_n^{-3}h^{-1}n\,(n^{\mathsf{T}}h^{-1}w),$$
$$C_h X = P_n\big(\Gamma^a{}_{ji}X^j n_a\big)_i, \qquad P_n = I - n\otimes n.$$

Differentiation of the ray map $n \mapsto h^{-1}n/\rho_n$ gives M_h, and substitution gives $M_h^{-1}v = \rho_n P_n h v$. A parallel covector has coordinate derivative $\Gamma^a{}_{ji}X^j n_a$, so the fixed contact splitting converts to the metric splitting by $(X,w) \mapsto (X, M_h(w - C_h X))$.

Let B_K be the tracefree part of $h^{-1}K$ on the h-screen $u^\perp = E_n$. It anticommutes with j_h, and the intrinsic complex structure has matrix $\left(\begin{smallmatrix} j_h & 0 \\ 2j_h B_K & j_h \end{smallmatrix}\right)$. For screen vectors, $h(X, M_h w) = \rho_n^{-1} X\cdot w$. This shows that $Z_0 + \beta\bar{Z}_0$ has eigenvalue $-i$. For the second graph vector its base component is ℓ and its fixed vertical component is $-\eta_c\bar{z}/d^2$. Its required $+i$ vertical projection is $-B_K\ell$. Pairing this condition with ℓ gives

$$-\frac{2\eta_c}{\rho_n} - \frac{n_a\Gamma^a{}_{ij}\ell^i\ell^j}{\rho_n} + K(\ell,\ell) = 0,$$

which is the defining formula. Thus the graph is the entire intrinsic distribution; this proves its angular compatibility and real transversality. Its contact compatibility also follows directly from $(z+\beta\bar{z})\cdot(z-\beta\bar{z}) = 0$.

To prove integrability, fix a base point and choose an oriented h-orthonormal frame normal there. In its angular coordinate the antiholomorphic vectors are $\partial_{\bar{\zeta}_h}$ and $z_h^i H_i^h + \frac{1}{2}K(z_h,z_h)\partial_{\zeta_h}$, where H_i^h are

the Levi-Civita horizontal lifts. At the chosen point the vertical coefficients of the lifts vanish along the entire fiber; their vertical derivatives vanish there as well. The remaining coefficients are holomorphic in ζ_h. The bracket of these two vectors therefore vanishes on that fiber, in each angular chart. The obstruction to involutivity is tensorial and the point was arbitrary, so the distribution is integrable. No constraint or spacetime evolution has been used.

At the flat metric the three imaginary null-square equations read off the three off-diagonal entries of a real symmetric form, while two real equations read off its diagonal differences. Adding its trace gives an invertible real six-by-six system. The same selected equations remain invertible in a fixed small graph neighborhood. They therefore give the claimed candidate $h^\sharp$. The complete conic test makes that candidate conformally equal to the supplied conic, and the density gives its unique physical scale by the determinant formula. The corresponding five null-square equations determine a symmetric K modulo the metric; its prescribed zero trace determines the missing component. Their coefficient matrix is again invertible near flat data. The complete shear test is precisely the remaining equality with the supplied graph. This proves both inverse identities and the independent nature of the tests.

All these formulas are smooth finite-dimensional operations with uniformly nonzero denominators in the retained angular charts. The forward map uses one spatial derivative of h and no derivative of K. The inverse obtains h by fixed angular evaluations and uses one derivative of this recovered metric when obtaining K. Weighted multiplication gives the displayed balanced estimates, in both directions. Taking a fixed number of angular derivatives costs only further finite derivatives of these angular formulas, with no additional spatial derivative. Parameter differentiation obeys the same finite rules. □

4.7. The commutator residual and its normalized response. The kinematic source already determines its connection. Its curvature therefore gives an independent vacuum test before a corrected source is constructed. The test has four components; no inversion of a curvature record is involved.

Proposition 4.16 (The four source curvature defects).
For a kinematic gradient source form the actual multiplication operator

$$\mathcal{R}^{\mathscr{S}}(X,Y) = [P_X, P_Y] - P_{[X,Y]}.$$

Let e_a be an oriented h-orthonormal frame, $v_a = \Theta(e_a)$ and θ^a its dual coframe. The following are finite differential readouts of the supplied operators:

$$E_{ab} = -\frac{1}{4}\sum_c \operatorname{tr}_{\mathscr{S}}\left(c(v_c)[\mathcal{R}^{\mathscr{S}}(e_c,e_a), c(v_b)]\right),$$

$$B_{ab} = \frac{1}{8}\epsilon_a{}^{cd}\operatorname{tr}_{\mathscr{S}}\left(c(N)[\mathcal{R}^{\mathscr{S}}(e_c,e_d), c(v_b)]\right),$$

$$H = \operatorname{tr} E, \qquad M_a = \epsilon_a{}^{bc}B_{bc}, \qquad M = M_a\theta^a, \qquad \mathfrak{d} = (H, M).$$

The matrix E is symmetric, B is tracefree, and on the maximal weighted carrier

$$H = R(h) - |K|_h^2, \qquad M = \operatorname{div}_h K, \qquad \mathfrak{d}(\mathrm{Enc}_{\mathrm{gr}}(Z)) = \mathcal{F}(Z).$$

For any real symmetric tracefree matrices E_0, B_0, define $W(E_0, B_0)$ by the algebraic curvature symmetries and

$$C_{0a0b} = (E_0)_{ab}, \qquad C_{cd0b} = -\epsilon_{cd}{}^{a}(B_0)_{ab}, \qquad C_{abcd} = -\epsilon_{ab}{}^{m}\epsilon_{cd}{}^{n}(E_0)_{mn}.$$

These are all real algebraic Weyl tensors. Put $E^\circ = E - H\delta/3$ and $B^{\mathrm{s}} = (B + B^{\mathsf{T}})/2$. With $\eta(\mathcal{C}(u,v)z, w) = C(u,v,z,w)$,

$$\begin{aligned}
\Omega_D - W(E^\circ, B^{\mathrm{s}})(\Theta, \Theta) &= \mathcal{C}_\Theta(H, M),\\
\mathcal{C}_\Theta(H, M)(e_a, e_b)N &= \frac{1}{2}(M_b v_a - M_a v_b),\\
\mathcal{C}_\Theta(H, M)(e_a, e_b)v_c &= \frac{1}{2}(\delta_{ac}M_b - \delta_{bc}M_a)N + \frac{H}{6}(\delta_{bc}v_a - \delta_{ac}v_b).
\end{aligned}$$

Consequently $\mathfrak{d} = 0$ is equivalent to the finite Weyl-image test $\Omega_D \in \mathrm{im}(C \mapsto \mathcal{C}(\Theta, \Theta))$. The Weyl preimage is unique. In the adapted orthonormal frame, using the positive Frobenius norm of the endomorphism matrices,

$$\sum_{a<b} \|\mathcal{C}_\Theta(H, M)(e_a, e_b)\|_{\mathrm{Frob}}^2 = \frac{H^2}{6} + |M|^2.$$

The contact presentation has the same independent test by its exact decoder and the gradient construction. Its complete kinematic tests and $\mathfrak{d} = 0$ are therefore equivalent to the maximal vacuum constraints, before referring to any generating ordinary pair.

Proof. The connection Leibniz rule cancels the first-order part of the commutator, and $[\mathcal{R}^{\mathscr{S}}(X,Y), c(v)] = c(\Omega_D(X,Y)v)$. The Clifford trace normalization therefore makes E the spatial Ricci contraction of Ω_D and B its magnetic contraction.

The Gauss–Weingarten identities give, without a constraint hypothesis,

$$E = \mathrm{Ric}(h) + \tau K - K h^{-1} K, \qquad B_{ab} = \epsilon_a{}^{cd}\nabla_c K_{db}, \qquad \tau = \operatorname{tr}_h K.$$

The first matrix is symmetric. Symmetry of K gives $\operatorname{tr} B = 0$. Contraction yields $H = R(h) + \tau^2 - |K|_h^2$ and $\epsilon_a{}^{bc}B_{bc} = \nabla^j K_{ja} - \nabla_a \tau$. Since the carrier imposes $\tau = 0$, these are the two rows of $\mathcal{F}$. In particular the second row is the actual covariant one-form M, not just its moving frame components.

The component formula for W verifies the algebraic curvature symmetries, the cyclic identity and all vanishing traces. Conversely the electric and magnetic contractions of an algebraic Weyl tensor are symmetric tracefree matrices and recover all its components by that formula. Restriction to spatial first arguments retains these contractions and is injective.

The spatial curvature block of a torsion-free Gauss–Weingarten connection is a three-dimensional algebraic curvature tensor. It is determined by its symmetric Ricci contraction: in dimension three its components are the metric wedge with the Ricci tensor, minus half its scalar trace times the

metric wedge with itself. Subtracting the displayed Weyl restriction leaves Ricci contraction $H\delta/3$, hence the spatial term in the complement formula. The remaining magnetic matrix is $(B-B^{\mathsf{T}})/2=\epsilon_{ab}{}^{c}M_c/2$. Epsilon inversion gives the normal component in that formula, and Lorentz skewness gives the value on N. This proves the entire complement and its equivalence with $\mathfrak{d}=0$. Summing the squares of the displayed matrices gives the norm identity: the spatial terms contribute $H^2/6$ and the normal-spatial terms contribute $|M|^2$. All operations used are contractions, matrix operations and finitely many derivatives of the supplied coefficients. □

Write $\mathsf{A}_Z=D_y\mathcal{F}_Z$ and $\mathsf{B}_Z=D_s\mathcal{F}_Z$. Their explicit current formulas are the residual differential, including the variation of tracefree projection. The actual flat inverse is

$$\mathsf{G}_0^{\mathrm{src}}=\operatorname{diag}\left(-\frac{1}{2}\mathcal{N}^{\mathrm{src}},\mathcal{V}^{\mathrm{src}}\right)=A_0^{-1},\qquad A_0=(-2\Delta,\Delta_L),$$

with its finite negative-response realization in the source Green theorem. The current inverse is the bounded Neumann limit of Proposition A.13. The individual cochain words are evaluated on weighted smooth fields; the bounded composite gives their base-space extension.

For later use, the differential of the coefficient encodings is explicit. For $H_1=\dot h$, $L_1=\dot K$, put

$$F\dot F+\dot F F=H_1,$$
$$\dot e=-e\dot F e,$$
$$\dot\Gamma^{a}{}_{ij}=\frac{1}{2}h^{ab}(\nabla_i(H_1)_{jb}+\nabla_j(H_1)_{ib}-\nabla_b(H_1)_{ij}).$$

The Sylvester equation has a unique solution because every eigenvalue of F is positive. Direct differentiation gives

$$\dot\Omega_i=\begin{pmatrix}0 & (L_1)_ie+K_i\dot e\\ \dot eK_i^{\mathsf{T}}+e(L_1)_i^{\mathsf{T}} & \dot F\Gamma_ie+F\dot\Gamma_ie+F\Gamma_i\dot e-(\partial_i\dot F)e-(\partial_iF)\dot e\end{pmatrix},$$
$$\dot{\mathbb{D}}_{\pm}=\sum_{a,i}E_a\otimes\left[\dot e_a^{\,i}(\partial_i+\varrho_{\pm}(\Omega_i))+e_a^{\,i}\varrho_{\pm}(\dot\Omega_i)\right].$$

The complete contact variation is

$$\begin{aligned}\dot\beta&=\frac{H_1(\ell,\ell)}{2\Delta_c},\\ \dot\rho_n&=-\frac{\rho_n}{2}H_1(u,u),\\ u&=h^{-1}n/\rho_n,\\ \dot\eta_c&=\frac{1}{2}\left(\dot\rho_nK(\ell,\ell)+\rho_nL_1(\ell,\ell)-n_a\dot\Gamma^{a}{}_{ij}\ell^i\ell^j\right)\\ &\quad-\dot\beta\left(\rho_nK(\ell,\bar z)-n_a\Gamma^{a}{}_{ij}\ell^i\bar z^j\right),\\ \dot\omega&=\frac{1}{2}\operatorname{tr}_h H_1.\end{aligned}$$

For the first identity differentiate $h(\ell,\ell)=0$ and use $h(\ell,\bar z)=B-\beta\bar A=\Delta_c$. The remaining formulas differentiate ρ_n, $\ell=z-\beta\bar z$ and the defining shear. Thus the variation of the null direction itself is retained.

Proposition 4.17 (Normalized source response to the curvature defect).

At any weighted smooth state $Z\in D$, and for $r\in\mathcal{Z}_m$, set

$$\begin{aligned} v &= \mathsf{A}_Z^{-1}r=(v_u,v_W),\\ \mathscr{R}_r^{\mathrm{gr}} &= D\,\mathrm{Enc}_{\mathrm{gr},Z}[0,v],\\ \mathscr{R}_r^{\mathrm{CR}} &= D\,\mathrm{Enc}_{\mathrm{CR},Z}[0,v]. \end{aligned}$$

They are the unique canonical kinematic tangents with

$$D\pi\,\mathscr{R}_r=0,\qquad D\mathfrak{d}\,\mathscr{R}_r=r.$$

Their actual coefficients are the preceding variation formulas with

$$\begin{aligned} H_1 &= v_u\delta,\qquad \dot\kappa=L_0v_W,\qquad \vartheta_\kappa=\mathrm{tr}_h\,\kappa,\\ L_1 &= \dot\kappa-\frac{1}{3}\left(\mathrm{tr}_h\,\dot\kappa-\mathrm{tr}(h^{-1}H_1h^{-1}\kappa)\right)h-\frac{\vartheta_\kappa}{3}H_1. \end{aligned}$$

On retained bounded sets, $\|\mathscr{R}_r^{\mathrm{gr}}\|_{\mathrm{op},m}\asymp\|r\|_{\mathcal{Z}_m}$, and the corresponding estimate holds in $\mathcal{B}_{m,l}$ for every fixed finite l. Constants include the higher bounds required by the current inverse. The operations preserve smooth finite-dimensional families, with the same finite-order interpretation.

For a seed direction a, let $\mathscr{S}_a=D\,\mathrm{Enc}_Z[a,0]$ be the raw seed tangent. The horizontal source operation is

$$\mathscr{V}_a=\mathscr{S}_a-\mathscr{R}_{\mathsf{B}_Z a}=D\,\mathrm{Enc}_Z[a,-\mathsf{A}_Z^{-1}\mathsf{B}_Z a].$$

It satisfies $D\pi\,\mathscr{V}_a=a$ and $D\mathfrak{d}\,\mathscr{V}_a=0$. Every tangent T to either canonical carrier has the unique decomposition

$$T=\mathscr{V}_{D\pi(T)}+\mathscr{R}_{D\mathfrak{d}(T)}.$$

In particular the vertical response, the raw seed tangent and the corrected horizontal seed tangent have different readouts; they are not interchangeable operations.

Proof. The exact kinematic encodings and their inverses identify carrier tangents with (a,v). The formulas for H_1,L_1 are the variation of the actual ambient coefficients at fixed seed. The two readouts of $\mathscr{R}_r$ are consequently 0 and $\mathsf{A}_Zv=r$. Invertibility of A_Z proves uniqueness.

The derivative of the defect can also be read directly from the source commutators. In coordinates write $P_i=\partial_i+A_i$, $\Pi_i=\dot P_i$ and $\mathcal{R}_{ij}=[P_i,P_j]$. Then

$$\begin{aligned} \dot{\mathcal{R}}_{ij} &= [P_i,\Pi_j]-[P_j,\Pi_i],\\ \dot{\mathcal{R}}_{ab} &= (\dot e_a^{\,i}e_b^{\,j}+e_a^{\,i}\dot e_b^{\,j})\mathcal{R}_{ij}+e_a^{\,i}e_b^{\,j}\dot{\mathcal{R}}_{ij},\\ \dot M_i &= \dot F^a{}_iM_a+F^a{}_i\epsilon_a{}^{bc}\dot B_{bc}. \end{aligned}$$

The last equation uses $M_i = F^a{}_i M_a$ in the positive frame. Its first term is needed off the admitted set. Differentiating the Clifford contractions with these moving-frame terms gives precisely the covariant residual derivative displayed in the current differential lemma. Thus the operator computation agrees with the tensor differential throughout the kinematic carrier, not only after the defect vanishes.

The upper estimate follows from the higher current inverse and the balanced coefficient estimates. Conversely $D\mathfrak{d}$ is a bounded differential operation of order two on H_1 and order one on L_1, so the identity $D\mathfrak{d}\,\mathscr{R}_r = r$ gives the lower estimate. Differentiating the inverse identity and the finite coefficient formulas proves the family statements.

Finally $D\mathfrak{d}\,\mathscr{S}_a = \mathsf{B}_Z a$. Subtracting its normalized vertical response gives the displayed horizontal operation and both its readouts. For any tangent T, subtract the two terms in the asserted decomposition. The remainder has both readouts zero; the exact decoder gives zero seed component and a correction component in the kernel of A_Z, hence zero. This proves the decomposition and its uniqueness. □

4.8. Initial execution from the ordinary pair. The variable input is the ordinary pair $c \in \mathscr{C}^{\mathrm{ord}}_{\gamma,\mathcal{B}}$. Its actual whole seed is $s = \mathscr{P}_\gamma \rho(c)$. Fix the reference state $(0,0)$ and its flat physical and spin markings. Choose the transport tube with the preceding graph-neighborhood reduction included, and choose $\mathcal{B}_\gamma$ inside its input ball S in the current transport theorem, with its one fixed base smallness margin. The initial source operations are the ordinary normalization and root transform, the finite coefficient formulas above, the actual negative-response word for $\mathsf{G}_0^{\mathrm{src}}$, its bounded current inverse limit, and initialized parameter integration. The negative-response word may use complex auxiliary calls; the coupled input condition and the actual resulting coefficients retain Lorentz reality.

Theorem 4.18 (Actual initial source execution from an ordinary pair).

For each admitted ordinary pair in the selected base domain, these operations construct a weighted smooth gradient source $\mathbb{S}(c)$ and its complete initial contact presentation $Q(c)$ on the whole fixed initial manifold. They satisfy the independent kinematic and vacuum tests and the exact input identities

$$\pi\mathbb{S}(c) = \pi Q(c) = \mathscr{P}_\gamma \rho(c),$$

$$\sigma_\gamma \mathscr{P}_\gamma^{-1} \pi\mathbb{S}(c) = c, \qquad \mathfrak{d}\mathbb{S}(c) = \mathfrak{d}Q(c) = 0.$$

Conversely, every independently tested kinematic source in either canonical presentation, with zero defect and decoded state in the retained product $\mathcal{B}_\gamma \times B(0, r_y)$ of the transport theorem, is recovered exactly by executing the ordinary pair $\sigma_\gamma \mathscr{P}_\gamma^{-1}\pi$. Both assertions concern the marked canonical presentations; the actual spin marking is retained.

For any piecewise C^1 seed path in S, transport on the constructed flat branch is the actual source evolution

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{S}(t) = \mathscr{V}_{\dot{s}(t)}(\mathbb{S}(t)), \qquad \frac{\mathrm{d}}{\mathrm{d}t}Q(t) = \mathscr{V}^{\mathrm{CR}}_{\dot{s}(t)}(Q(t)).$$

Its endpoint depends only on the endpoint seed and it obeys the identity and composition laws. Along the radial path from zero it can be executed by the controlled finite Neumann–Picard stages of Corollary A.15. Every finite stage is an actual kinematic source with the exact prescribed endpoint seed. Its remaining defect tends to zero with the base error bound in the finite-execution estimate. The limit has every finite weighted order on the same base domain; this last assertion does not require convergence of the raw approximants in all higher norms.

On retained sets with the required finite higher bounds, for $m \geq k_0$ and each fixed finite angular order l,

$$\|\mathbb{S}(c) - \mathbb{S}(\widetilde{c})\|_{\mathrm{op},m} \asymp \|s(c) - s(\widetilde{c})\|_{\mathcal{X}_m},$$
$$\|Q(c) - Q(\widetilde{c})\|_{\mathcal{B}_{m,l}} \asymp \|s(c) - s(\widetilde{c})\|_{\mathcal{X}_m}.$$

The maps preserve smooth finite-dimensional families. For compact parameter sets $\mathcal{K} \Subset \operatorname{int}\mathcal{K}'$, every specified finite spatial and parameter estimate on $\mathcal{K}$ uses finitely many input seminorms on $\mathcal{K}'$ and the fixed base margins. The seed ball and the spatial domain are fixed before these higher orders. Effective ordinary cochain comparisons give the same canonical source, with the retained actual spin identification.

After this execution and the independent tests have been proved, the tensor readout of either source equals the earlier datum $\mathcal{T}(\mathscr{P}_\gamma\rho(c))$ of Theorem 3.13. That datum is not an additional generating input.

Proof. Normalize the ordinary pair using Theorems 3.9–3.10 and Corollary 3.12, and execute its root transform to obtain s. Apply the initialized transport of Theorem A.14 with $(s_*, y_*) = (0, 0)$. It constructs one actual base endpoint $Y(s)$, by the specified differential and negative-response operations, and puts it in every higher weighted space for smooth input. Form

$$\mathbb{S}(c) = \mathrm{Enc}_{\mathrm{gr}}(s, Y(s)), \qquad Q(c) = \mathrm{Enc}_{\mathrm{CR}}(s, Y(s)).$$

The exact coefficient results above prove that these are actual fields on the entire initial manifold, with both chiral operators, all angular directions, the density and the actual spin marking. They meet every independent kinematic test before the vacuum equation is used. By the chain rule their evolution is the displayed horizontal source equation. Its defect derivative is zero by the normalized response proposition. The flat initial defect is zero, so the resulting fields satisfy the independent Weyl-image test. This proves admission by a source differential identity.

The exact decoder returns $(s, Y(s))$. Thus its seed readout is $s = \mathscr{P}_\gamma\rho(c)$, and the identity $\sigma_\gamma\rho(c) = c$ on the selected ordinary domain gives the ordinary-class unit. In particular the unit concerns an actual ordinary

cohomology class, not only its polynomial quotient or its punctured root values.

Conversely, independently test a supplied source and use its exact kinematic decoder to obtain (s, y) in the stated product. The commutator proposition gives $\mathcal{F}(s, y) = 0$. The current transport theorem identifies the entire zero-residual set in that product with $(s, Y(s))$. Its exact coefficient inverse therefore returns the supplied source, including its complete symbol and directional coefficients, or its complete contact graph and density. Executing $\sigma_\gamma \mathscr{P}_\gamma^{-1} s$ gives that same endpoint. This proves both object-level inverse assertions on the independently specified initial domain.

Transport and its composition laws follow by conjugating the actual coefficient transport with the exact encodings. The construction along an arbitrary path in S stays on the same retained tube and uses no path-length smallness assumption. For the finite realization, let $Y_{J,N}(s)$ be the endpoint in Corollary A.15. The stage $\mathrm{Enc}(s, Y_{J,N}(s))$ is an actual weighted smooth kinematic source and has exact seed s, even when its defect is nonzero. The balanced coefficient estimate transfers the base endpoint error to the operator and contact norms. The identity $\mathfrak{d}\,\mathrm{Enc} = \mathcal{F}$ transfers its residual bound. This proves the claimed convergence and its precise finite-stage meaning. All higher regularity concerns the same limit obtained at the base order, by Theorem A.14.

The upper difference estimates follow from the higher endpoint difference estimate and the finite coefficient encodings. The lower estimates follow from their exact inverse readout and the bounded ambient seed decoder. Repeated parameter differentiation uses the same current inverse and finite coefficient formulas, so the family quantifiers and finite seminorm dependence are precisely those of the current transport theorem and the ordinary normalization. Two effectively equivalent admitted cochains have the same root and ordinary class, hence the same whole seed and the same constructed canonical field. Their identification retains the chosen actual spin marking; no quotient by a central spin sign has been introduced.

Finally $Y(s)$ solves $\mathcal{F}(s, Y(s)) = 0$. Uniqueness in the independently defined tensor chart identifies it with the earlier correction $\mathcal{C}(s)$, and its tensor readout is consequently $\mathcal{T}(s)$. At this point the source operations have already produced a solution, verified the vacuum test and recovered the input pair. The argument applies at every curvature rank in the smooth category. □

The construction so far takes an admitted ordinary pair to its marked initial source. Section 5 starts with a physical germ. We choose a maximal disk and construct a pair whose development presents that germ. Sections 6 and 7 compare the completed sources when their initial hypersurfaces differ.

5. Local displays of arbitrary vacuum germs

A small weighted construction can represent an arbitrary local vacuum geometry once its physical scale and initial embedding are retained. We first choose a local maximal disk with its point and tangent plane fixed. We then extend its normalized data exactly across a protected region and apply the whole-data decoder. Finally, local Cauchy uniqueness identifies the generated development with the original physical germ.

The exterior extension supplies the whole-space input to the constructor. Only its agreement with the disk on the protected region enters the identification of the physical germ.

5.1. Conventions and analytic inputs. A spacetime germ is four-dimensional, pointed, oriented and time oriented, and its actual local spin structure is retained. We use signature $(+---)$. If j is a spacelike embedding and n its future unit normal, its induced data are

$$h = -j^* g, \qquad K(X,Y) = -g(\nabla_X n, Y).$$

Thus a Gaussian presentation has $g = \mathrm{d}t^2 - h_t$ and $K = \frac{1}{2}\partial_t h_t$. A maximal vacuum datum satisfies

$$\mathrm{tr}_h K = 0, \qquad \mathrm{div}_h K = 0, \qquad \mathrm{Scal}(h) = |K|_h^2.$$

Write $B_a \subset \mathbb{R}^3$ for the Euclidean ball of radius a and $\langle x\rangle = (1+|x|^2)^{1/2}$. For a tensor in fixed Cartesian coordinates, the weighted Sobolev norm is

$$\|u\|_{H^s_\gamma}^2 = \sum_{|\beta|\le s} \int_{\mathbb{R}^3} \langle x\rangle^{-2\gamma-3+2|\beta|} |\partial^\beta u|^2 \,\mathrm{d}x.$$

The notation $\mathscr{S}^\mu$ means that $\sup_x \langle x\rangle^{-\mu+|\beta|}|\partial^\beta u(x)|$ is finite for every multi-index β. Weighted Hölder norms are those of Definition A.1, as used in the whole-data spatial-gauge theorem. We fix one weight $\gamma = -1/2 - \varepsilon_\gamma$, with $0 < \varepsilon_\gamma < 1/2$, in the interval furnished by the exterior inverse theorem. No choice of an inverse at every weight in $(-1, -1/2)$ is needed.

The following analytic inputs are used in their stated forms.

Relative extension. Fix the exterior normalization described in Appendix B. Theorem B.1 selects an operation $\mathfrak{E}$ on smooth whole prefields $(\widehat{h}, \widehat{K})$ satisfying $\mathrm{tr}_{\widehat{h}} \widehat{K} = 0$ on $\mathbb{R}^3$, the maximal constraints on B_1, finiteness of all weighted norms and one finite base-smallness condition. Its output (H, L) is maximal vacuum data on $\mathbb{R}^3$ and agrees exactly with the prefields on B_1. For the same selected output,

$$H - \delta \in H^M_\gamma, \qquad L \in H^{M-1}_{\gamma-1} \quad \text{for every integer } M \ge S_0.$$

Each norm and each specified finite parameter derivative is bounded by finitely many input seminorms and fixed positive margins. The base domain is chosen before those higher orders. The original extension result is Czimek (2018, Theorem 3.1); the protected relative normalization, nearby stronger

weight and same-branch assertion used here are the additional inputs specified in Appendix B.

Ordinary classes and their whole seeds. Let $\mathscr{W}_\gamma$ be the paired root-class space of Definition 3.5, with its coupled residue reality and weighted Fourier-amplitude tests. The isomorphism of Theorem 3.6 is

$$\mathscr{P}_\gamma : \mathscr{W}_\gamma \longrightarrow \mathcal{X}_\gamma^\infty .$$

Here $\mathcal{X}_\gamma^\infty$ consists of the smooth Euclidean transverse-traceless metric and momentum seeds, of symbol orders γ and $\gamma - 1$, in the indicated weighted chart spaces. All their weighted derivatives are finite; smallness is imposed at the chosen finite base order. The root-class complex used for this transform is linear. Nonlinear source operations use the actual coefficient complexes $\mathscr{E}_2^\bullet$ and $\mathscr{E}_{-6}^\bullet$ of Proposition 3.7.

Theorem 3.9 gives a continuous real-linear section

$$\sigma_\gamma : \mathscr{W}_\gamma \longrightarrow H^1(\mathscr{E}_2^\bullet) \oplus H^1(\mathscr{E}_{-6}^\bullet), \qquad \rho\sigma_\gamma = \mathrm{id}\,.$$

The first cohomology here is formed using actual closed cochains and actual coboundaries. The map ρ is the root readout. The section has selected closed representatives, is compatible with smooth passive parameters, and preserves the calibrated root reality. Its positive representative is anchored at the origin and has symbol order $\gamma + 1$; changing the anchor gives an actual coboundary. The negative representative has the decaying orders stated in that theorem.

Fix the flat source as reference and choose one rotation-invariant base neighborhood $\mathcal{B}_\gamma$ of zero in $\mathcal{X}_\gamma^\infty$ on which the initialized transport of Theorem A.14 and the source execution of Theorem 4.18 are defined. These neighborhoods contain a ball for the finitely many base norms; a rotation-invariant ball can be taken by using their suprema over the compact rotation group. All finite chart and inverse thresholds are included in this choice. Define

$$\mathscr{C}_{\gamma,\mathcal{B}}^{\mathrm{ord}} = \sigma_\gamma(\mathscr{P}_\gamma^{-1}(\mathcal{B}_\gamma)).$$

Throughout this section, $c = (c_+, c_-)$ denotes an actual ordinary pair in this domain; its whole seed is

$$s(c) = \mathscr{P}_\gamma \rho(c).$$

On the selected neighborhood, the canonical representatives supplied by σ_γ satisfy Definition 3.11. For other supplied representatives test actual closedness, root membership in $\mathscr{P}_\gamma^{-1}(\mathcal{B}_\gamma)$, and vanishing of the polynomial obstructions of Theorem 3.10. The positive obstruction is tested in its specified quotient. Effective arrows are actual primitives modulo closed zero-cochains that induce the same cochain change.

Whole-data gauge and decoding. For smooth whole maximal vacuum data (H, L), the hypotheses of Corollary A.11 are

$$H - \delta \in \mathscr{S}^\gamma, \qquad L \in \mathscr{S}^{\gamma-1}, \qquad \|H - \delta\|_{C_\gamma^{k+2,\alpha}} + \|L\|_{C_{\gamma-1}^{k+1,\alpha}} < \epsilon_{\mathrm{g}}.$$

The based gauge of Theorem A.9 has the form

$$\phi_H = \mathrm{id} + X_H, \qquad X_H(0) = 0, \qquad X_H \in \mathscr{S}^{\gamma+1}, \qquad \|DX_H\|_\infty < \frac{1}{2}.$$

Thus its displacement is sublinear. Its isotopy from the identity has an initialized actual spin lift. Apply the exact seed decoder of Proposition A.4 to $\phi_H^*(H, L)$ and then apply $\mathscr{P}_\gamma^{-1}$ and σ_γ. After reducing ϵ_{g} once so that the seed lies in the fixed $\mathcal{B}_\gamma$, the resulting ordinary pair satisfies

$$\phi_H^*(H, L) = \mathcal{T}(s(c)) =: (h(c), K(c)).$$

The map $\mathcal{T}$ is the nonlinear constraint chart of Theorem A.6. The ordinary-pair constructor of Theorem 3.13 is $c \mapsto \mathcal{T}(s(c))$. The decoder, actual class lift and based gauge have finite smooth-family estimates on the same selected base domain. Proper Euclidean rotations preserve the chart and gauge. These operations act on the whole extended fields.

Source execution, geometric completion and wave completion. For a selected actual representative of c, Theorem 4.18 constructs a canonical admitted initial gradient source and its complete initial contact presentation, with seed $s(c)$ and exact input recovery. Effective ordinary arrows give the same canonical output. The fixed reference and the supplied pair determine the source responses and hence the constructor. The spacetime metric is recovered when the resulting vacuum data are developed.

The initialized Cartan completion of Corollary 2.12 is the subsequent geometric development of the readout $(h(c), K(c))$. Its independent admission, subsidiary identities and geometric recognition are proved in Theorem 2.9, using the exact smooth Einstein Cauchy input of Theorem 2.6. Theorem 2.10 carries every specified initial physical and actual spin comparison. Theorem 2.14 supplies a common local parent and all finite family estimates. Corollary 2.15 compares exterior, gauge and scale choices for the initial and finite Cartan parts when the physical initial germ is fixed.

Finally, Section 6 constructs the wave completion from this source collar. Work on a smaller relatively compact causally convex globally hyperbolic collar, retaining every nonzero null direction over its point core by a finite angular cover of the full skies. Mixed carrier compositions use compactly contained ray domains whose rays meet each retained spacelike carrier once, with properly supported representatives and a positive buffer from the patch boundaries. Its scalar kernel class Λ_g^+ is characterized by the independent bisolution, commutator and positive-frequency wavefront tests of Definition 6.9 and Theorem 6.10, modulo kernels jointly smooth in the points and supplied parameters. Let C_{ab} denote the mixed trace of Definition 6.12 with its two slice half-density factors, and let C_{ap} be the one-slice point column of Definition 6.15, with only the carrier half-density factor. Set $W_a = C_{aa}^{-1/2}$ with the specified positive principal root, and put

$$T_{ab} = W_a C_{ab} W_b, \qquad e_a(p) = W_a C_{ap}.$$

Theorem 6.14 and Theorem 6.19 give the complete carrier comparison and the ambient point connection on all module sections. Proposition 6.16 proves the full evaluation-section identity, including every lower coefficient. These are derived structures. Chosen exact smoothing representatives are not additional generating inputs. The quantitative assertions used below concern finite complete coefficients and joint derivatives on the retained parents, as proved in Proposition 6.11, Proposition 6.17 and Proposition 6.22. Theorem 6.24 supplies independent admission and recovery for the normalized scalar enrichment over the constructed finite Cartan collar.

Choose S_0 large enough for these inputs. If the spatial gauge is measured in the fixed weighted $C^{k+2,\alpha}$ and $C^{k+1,\alpha}$ norms, require

$$S_0 > k + 2 + \alpha + \frac{3}{2},$$

as well as the finite lower requirements of the extension and source construction. Take one base threshold small enough for extension, gauge, ordinary decoding and the selected actual source-input ball. The radius below is chosen only after this order and threshold have been fixed.

5.2. Normal coordinates and maximal disks.

Lemma 5.1 (Normalization of a smooth metric family).
Let $(g_\lambda, p_\lambda, P_\lambda)$ be a smooth finite-dimensional family of Lorentzian metric germs with a specified spacelike tangent three-plane P_λ at p_λ. Locally in the parameter, choose a smooth oriented orthonormal frame $e_\lambda : \mathbb{R}^{1,3} \to T_{p_\lambda} M_\lambda$ whose time vector is future pointing and whose spatial image is P_λ. On common representatives and after restriction of the parameter neighborhood, there is $r_0 > 0$ such that

$$\phi_{\lambda,r}(z) = \exp_{p_\lambda}^{g_\lambda}(r e_\lambda z), \qquad g_{\lambda,r} = r^{-2}\phi_{\lambda,r}^* g_\lambda$$

are defined on $Q = (-2, 2) \times B_3$ for $0 < r \le r_0$. The coefficients extend smoothly to $r = 0$ and have the form

$$g_{\lambda,r} = \eta + r^2 G(\lambda, r, z), \qquad \eta = \mathrm{d}t^2 - \sum_{i=1}^{3} (\mathrm{d}x^i)^2.$$

The same r_0 works at every finite derivative order. On retained compact parameter subsets, each finite mixed derivative of G is bounded by finitely many derivatives of the marked family, frame and trivialization and by the retained coordinate margins. Ricci-flatness is preserved for $r > 0$.

Proof. Trivialize the family near its marked points. The geodesic equation depends smoothly on its metric, initial point and initial vector, so its normal-coordinate coefficients $\widetilde{g}_{\lambda,\mu\nu}(z)$ are smooth on one common neighborhood. The orthonormal frame and normal-coordinate identities give

$$\widetilde{g}_\lambda(0) = \eta, \qquad \partial_a \widetilde{g}_\lambda(0) = 0.$$

The normalized coefficients are $\widetilde{g}_{\lambda,\mu\nu}(rz)$. Taylor's integral formula therefore gives

$$G_{\mu\nu}(\lambda, r, z) = z^a z^b \int_0^1 (1-s)\partial_a\partial_b \widetilde{g}_{\lambda,\mu\nu}(srz)\,\mathrm{d}s.$$

This formula is smooth at $r = 0$ and gives every finite mixed estimate on the same parent. Multiplication of a metric by a positive constant leaves its Levi–Civita connection and covariant Ricci tensor unchanged. Naturality under pullback consequently gives $\mathrm{Ric}(g_{\lambda,r}) = \phi_{\lambda,r}^* \mathrm{Ric}(g_\lambda)$. □

Theorem 5.2 (A maximal disk with prescribed point and tangent plane).

Fix $0 < \alpha < 1$. Each smooth metric in a sufficiently small $C^{4,\alpha}$ neighborhood of η on Q has a selected small spacelike maximal graph

$$j_g(x) = (u_g(x), x), \qquad x \in B_2, \qquad u_g(0) = 0, \quad Du_g(0) = 0.$$

Its boundary value is $a_g + b_g \cdot x$ on ∂B_2. It is unique among small graphs with affine boundary value and the displayed point conditions, and is smooth near $\overline{B}_{7/4}$.

For a jointly smooth finite-dimensional family in this base neighborhood, the selection is jointly smooth on that retained interior ball. Each finite spatial and parameter derivative has a bound in finitely many derivatives of the metric family and the fixed base margins. The base neighborhood and retained ball do not depend on the requested higher order.

Proof. Let $E_i = \partial_i + u_i\partial_t$ be the graph tangent vectors, let N be its future unit normal, and put $\nu = g(N, \partial_t)$. Near the flat zero graph, ν is positive and the induced metric h is positive definite. Differentiation of $g(N, E_j) = 0$ yields

$$K_{ij} = \nu u_{ij} + g\big(N, \Gamma(E_i, E_j)\big), \qquad H_g(u) = \nu h^{ij} u_{ij} + h^{ij} g\big(N, \Gamma(E_i, E_j)\big).$$

The coefficient νh^{ij} is uniformly positive on one small graph neighborhood. At the flat metric,

$$H_\eta(u) = \operatorname{div}\left(\frac{Du}{\sqrt{1-|Du|^2}}\right), \qquad D_u H_\eta(0) = \Delta.$$

Write $u = v + a + b \cdot x$, with $v \in C_D^{2,\alpha}(\overline{B}_2)$ of zero Dirichlet trace. This condition imposes no boundary normal derivative. Consider

$$\mathcal{F}(g, v, a, b) = \big(H_g(v + a + b \cdot x),\, v(0) + a,\, Dv(0) + b\big).$$

It is continuously differentiable into $C^{0,\alpha} \times \mathbb{R}^4$ on the indicated metric and graph spaces; the extra metric derivatives give the composition margin. Its derivative in the unknowns at the flat zero graph is

$$(v, a, b) \longmapsto \big(\Delta v,\ v(0) + a,\ Dv(0) + b\big).$$

For a prescribed (f, c, d), solve the zero-Dirichlet problem $\Delta v = f$ and set $a = c - v(0)$ and $b = d - Dv(0)$. The Dirichlet estimate and bounded point evaluations give a bounded inverse. The implicit-function theorem selects

the graph and its affine boundary modes. Restrict the base neighborhood once so that the graph is spacelike, stays in Q, and the augmented derivative remains invertible. These four boundary modes impose the point and tangent conditions without adding conditions to a fixed zero-Dirichlet problem.

For smooth input, interior elliptic estimates improve this existing solution. At each order the highest derivative has the same elliptic principal part; the remaining terms contain finitely many previously controlled derivatives. For a requested order, use a finite chain of nested balls between $B_{7/4}$ and B_2. Every chain ends on the same retained ball. Its constants may depend on the order, but no higher norm must be small.

On a smooth finite-dimensional family, the augmented map restricted to that family is smooth to each required finite order because the family supplies the extra metric derivatives needed for composition. Its base implicit solution is smooth into $C^{2,\alpha}_D \times \mathbb{R}^4$. Parameter differentiation gives the same invertible augmented linearization with lower-order forcing. Interior elliptic estimates bound spatial derivatives successively in parameter order. Difference quotients in the base space identify those solutions with derivatives of the already selected graph. Higher spatial estimates and interpolation give continuity in each higher spatial space; the fundamental theorem of calculus gives differentiability there. Induction proves joint smoothness and the finite dependence asserted above. This is a family statement on smooth inputs, not a same-regularity smooth action on the full finite Hölder space. □

Corollary 5.3 (Normalized maximal data).
Apply Theorem 5.2 to Lemma 5.1, including $r = 0$. On $B_{7/4}$ the selected graph and induced data have the form

$$u_{\lambda,r} = r^2 U_{\lambda,r}, \qquad \bar{h}_{\lambda,r} = \delta + r^2 A_{\lambda,r}, \qquad \bar{K}_{\lambda,r} = r^2 B_{\lambda,r},$$

with jointly smooth U, A, B. For Ricci-flat input the data satisfy the maximal vacuum constraints. With $\Psi_{\lambda,r} = \phi_{\lambda,r} \circ j_{\lambda,r}$, their physical normalization is

$$-\Psi^*_{\lambda,r} g_\lambda = r^2 \bar{h}_{\lambda,r}, \qquad K_{\Psi_{\lambda,r}} = r \bar{K}_{\lambda,r}.$$

Proof. At $r = 0$ the selected graph vanishes, and the first r derivative of the normalized metric is zero. Differentiate the augmented equation and use its invertible derivative to obtain $\partial_r u_{\lambda,0} = 0$. Taylor's integral formula gives $u = r^2 U$. The induced metric and the graph equation give the remaining factorizations. Ricci-flatness and maximality give the vacuum constraints.

Under $g \mapsto r^2 g$, the future unit normal is divided by r, the induced metric is multiplied by r^2, and the covariant second fundamental form is multiplied by r. Pullback gives the stated physical identities. □

Each prescribed spatial or passive-parameter derivative of the three vanishing quantities is $O(r^2)$. With ℓ additional r derivatives the corresponding bound is $O(r^{\max(2-\ell,0)})$. Thus one positive radius can meet all finite base requirements while every higher requested norm remains finite at that radius.

5.3. Protected extension and ordinary decoding.

Lemma 5.4 (Extension of buffered maximal data).

Smooth maximal vacuum data $(\bar{h}, \bar{K})$ near $\overline{B}_{3/2}$ that are sufficiently small relative to $(\delta, 0)$ in $H^{S_0} \times H^{S_0-1}$ have a selected smooth maximal vacuum extension (H, L) on $\mathbb{R}^3$ with

$$(H, L)|_{B_1} = (\bar{h}, \bar{K}), \qquad H - \delta \in H^M_\gamma, \quad L \in H^{M-1}_{\gamma-1} \quad \text{for every finite } M.$$

The base norm is bounded by a fixed constant times the local input norm. Each higher norm is bounded using finitely many higher local norms and the fixed base margins. Supplied smooth finite-dimensional families have one such smooth selection, with every specified mixed estimate on the same base domain.

Proof. Choose $0 \le \chi \le 1$, supported in $B_{3/2}$ and equal to one near $\overline{B}_1$. Extend the cutoff products by zero and put

$$\widehat{h} = \delta + \chi(\bar{h} - \delta), \qquad \widehat{K} = \mathrm{TF}_{\widehat{h}}(\chi \bar{K}).$$

The metric is positive, and $\widehat{K}$ is globally tracefree for that metric. On B_1 these prefields agree exactly with the supplied constraint data. Their perturbations are compactly supported, so they have every required weighted norm. Multiplication, metric inversion and tracefree projection give the base estimate

$$\|\widehat{h} - \delta\|_{H^{S_0}_\gamma} + \|\widehat{K}\|_{H^{S_0-1}_{\gamma-1}} \le C \left(\|\bar{h} - \delta\|_{H^{S_0}(B_{3/2})} + \|\bar{K}\|_{H^{S_0-1}(B_{3/2})} \right).$$

Theorem B.1 therefore applies, with its normalization fixed as in §5.1. It gives the protected equality, the base bound and all higher bounds for the same initialized output. The cutoff and projection preserve smooth parameter dependence and finite mixed bounds, so the same input also gives the family assertion. Compact support already supplies any stronger decay required of the prefield; no additional strict-weight assumption on the local data is needed. □

By the fixed Sobolev-to-Hölder margin, the base output lies in the gauge domain after one reduction of the local threshold. Its higher weighted estimates imply the symbol bounds. Apply the whole-data gauge and decoder to obtain

$$(h(c), K(c)) = \phi^*_H(H, L) = \mathcal{T}(s(c)).$$

All inverse operations act on the whole extended field. The completed source is subsequently restricted to a local patch.

Theorem 5.5 (Ordinary pairs for marked maximal initial germs).

Every smooth pointed maximal vacuum initial germ (Σ, h, K, o) with its actual spin marking has an ordinary-pair presentation in the fixed small weighted domain: there are c, $r > 0$ and an orientation-preserving embedding $\iota : V \to \Sigma$, with $\iota(0) = o$, such that

$$\iota^* h = r^2 h(c)|_V, \qquad \iota^* K = r K(c)|_V.$$

The embedding has the specified actual spin comparison. The pair generates its canonical admitted initial source and complete initial contact presentation. The construction can be selected smoothly near every parameter value of a smooth finite-dimensional family, using one positive radius, one protected ball and one base domain.

Proof. Choose oriented h-orthonormal frames at the marked points and their actual spin lifts, smoothly near the chosen parameter value. Set

$$\psi_r(y) = \exp_o^h(rey), \qquad \bar{h}_r = r^{-2}\psi_r^* h, \qquad \bar{K}_r = r^{-1}\psi_r^* K.$$

For the unscaled normal-coordinate coefficients,

$$(\bar{h}_r)_{ij}(y) = \widetilde{h}_{ij}(ry) = \delta_{ij} + r^2 A_{ij}(r, y), \qquad (\bar{K}_r)_{ij}(y) = r\,\widetilde{K}_{ij}(ry).$$

All prescribed finite derivatives are bounded on one buffered ball. In this general initial-data normalization the momentum is $O(r)$; Corollary 5.3 obtains $O(r^2)$ for its particular spacetime-selected disk.

The constraints are preserved. For constant $a > 0$, under $h' = a^2 h$ and $K' = aK$,

$$\operatorname{Scal}(h') = a^{-2}\operatorname{Scal}(h), \qquad |K'|^2_{h'} = a^{-2}|K|^2_h,$$
$$\operatorname{tr}_{h'} K' = a^{-1}\operatorname{tr}_h K, \qquad \operatorname{div}_{h'} K' = a^{-1}\operatorname{div}_h K.$$

The spatial connection is unchanged. Apply these identities with $a = r^{-1}$ after pullback.

With the finite base order and all constructor thresholds already fixed, choose r so that Lemma 5.4 and the whole gauge and decoder apply. Let (H_r, L_r) be the extension. If a particular orthonormal point frame is retained, define

$$R_r = D\phi_{H_r}(0)\big((\phi_{H_r}^* H_r)(0)\big)^{-1/2} \in SO(3), \qquad O_r = R_r^{-1}.$$

Here $H_r(0) = \delta$. Decode the whole datum $O_r^*\phi_{H_r}^*(H_r, L_r)$, which remains in the spatial gauge, and call its ordinary pair c_r. If no particular orthonormal frame is marked, omit this rotation. On one smaller V such that $\phi_{H_r} O_r(V) \subset B_1$, set

$$\iota_r = \psi_r \circ \phi_{H_r} \circ O_r|_V.$$

Protected equality gives the two tensor identities in the statement. The polar formula carries the canonical orthonormal frame of the rotated datum to the selected normalized frame.

Lift the gauge isotopy and the small rotation with their initialized spin signs, and use the frame map $e_a \mapsto r^{-1}e_a$ for positive metric scaling. This produces the actual spin comparison with the prescribed central sign. On a connected local germ the lift is determined by its value at the marked frame.

The actual-class decoder supplies canonical closed representatives and the prescribed primitive comparisons. Theorem 4.18 constructs the admitted initial source from this pair and proves exact input recovery. The original metric is used to obtain a representing input in this proof; it is not a separate argument of that source constructor.

For a family, the initial smallness estimates permit a single positive r. Relative extension, based gauge, polar normalization, decoding and actual class transformation are smooth on the same selected domain. Their finite estimates, followed by those of source execution, give the claimed selection. Raising the requested order changes the required finite norms and their bounds, not the extension or radius. □

5.4. Physical scale on the generated source.

Proposition 5.6 (Physical scale and the wave completion).
Let an admitted source collar have metric g and the wave completion specified in §5.1. For $r > 0$ constant on spacetime, the scale restoration $g^{[r]} = r^2 g$ has

$$\mathbb{D}_{\pm}^{[r]} = r^{-1}\mathbb{D}_{\pm}, \qquad h^{[r]} = r^2 h, \qquad K^{[r]} = rK.$$

The internal Lorentz and spin pairings and the future unit section are kept fixed. The recovered Lorentz and spin connections are unchanged, while the soldering is multiplied by r. The actual spin identification follows the positive scaling path and retains the supplied central sign.

For the scalar wave structures in their fixed-coordinate realizations,

$$\Lambda^{+}_{r^2 g} = r^{-2}\Lambda^{+}_{g}, \qquad C^{[r]}_{ab} = rC_{ab},$$
$$W^{[r]}_{a} = r^{-1/2}W_a, \quad T^{[r]}_{ab} = T_{ab}, \qquad e^{[r]}_{a} = r^{-1}e_a.$$

The ambient point connection is unchanged. These identities are identities in the specified complete coefficient and joint-smoothing calculus. The numerical scale remains part of the physical display.

Proof. The scalar commutator of $r^{-1}\mathbb{D}_{\pm}$ multiplies the spatial symbol by r^{-1}. Its metric readout is therefore $r^2 h$ and its soldering is $r\Theta$. Contracting the whole directional operator against this soldering leaves the directional connection unchanged. Its normal readout is consequently rK. This is the full-gradient calculation of Proposition 2.11; it uses the uncontracted operators, retaining the full second fundamental form. The proposition also gives the actual finite Cartan transport with the physical clock $t' = rt$.

The positive metric rescaling identifies orthonormal frames by $v_a \mapsto r^{-1}v_a$. Initialization at $r = 1$ gives the spin identification along positive scales. A supplied opposite central lift remains distinct. Constant metric scaling leaves the unparametrized null rays, skies and contact distribution unchanged and multiplies the physical point density by r^4. Proposition 7.4 and Proposition 7.5 reconstruct their selected optical and canonical finite-rank fields from the rescaled soldering and the same connection. This restores their numerical physical scale; it does not declare an isometry between the two unscaled metrics.

For the wave part, $\Box_{r^2 g} = r^{-2}\Box_g$ and the four-dimensional volume form scales by r^4. The scalar causal kernel therefore has weight r^{-2}. The independent kernel characterization gives that same weight to Λ^{+}_{g}. Each slice

half-density has weight $r^{3/2}$, so the two-slice and one-slice kernels have weights

$$C^{[r]}_{ab} = rC_{ab}, \qquad C^{[r]}_{ap} = r^{-1/2}C_{ap}.$$

The positive-principal inverse square root gives the stated law for W_a. Substituting into the formulas for T_{ab} and e_a proves their weights. The conormal point families and their actual point derivatives are unchanged in the fixed-coordinate identification, giving the connection assertion. These are the complete-class scale laws of Proposition 6.17 and Proposition 6.22, and also give $\Box^{\nabla^{\mathrm{amb}}}_{r^2 g} e^{[r]} = r^{-3}\Box^{\nabla^{\mathrm{amb}}}_{g} e$. All comparisons of physical germs are made after this numerical scale has been restored. □

5.5. Coverage of spacetime germs.

Theorem 5.7 (Ordinary-pair coverage of smooth Ricci-flat germs). Every smooth pointed, oriented and time-oriented Lorentzian Ricci-flat germ, with a chosen spacelike tangent plane and actual local spin marking, has an ordinary-pair presentation of a local maximal Cauchy disk through that point and tangent plane. The pair generates an admitted initial source, its geometric collar and the wave completion specified in §5.1. The physically rescaled development is isometric to the original germ, with the retained future normal and actual spin comparison. Its generated complete structures have the physical scale laws of Proposition 5.6.

The construction is locally available for smooth finite-dimensional families on common representatives, with one positive normalization radius and one retained parent. Local geometric Cauchy uniqueness identifies the normalized development with the normalized germ on a common smaller domain of dependence. Restoring the initial data to (r^2h, rK) returns the development to the physical metric scale of the given germ.

Proof. Lemma 5.1, Theorem 5.2 and Corollary 5.3 give maximal vacuum data on a buffered ball. Choose one radius satisfying the finite base conditions. Apply Lemma 5.4, the whole gauge and decoder, and the polar frame correction when required. Compose their initial embedding with $\Psi_{\lambda,r}$. Its physical initial tensors are exactly $r^2h(c)$ and $rK(c)$.

Execute the pair by Theorem 4.18 to obtain the admitted initial source. Corollary 2.12 supplies its initialized Cartan completion from these derived data. Lemma 7.2 constructs its actual full null incidence on a smaller retained core; Proposition 7.4 and Proposition 7.5 construct its optical and canonical finite-rank fields. Shrink to a relatively compact causally convex globally hyperbolic collar and apply Theorem 6.24 to construct its independently admitted scalar enrichment, using the wave formulas of §5.1. The retained-carrier construction in Definition 6.12 uses spacelike transversality and compactness of the projective sky to give one smaller point neighborhood with full angular coverage, single carrier intersections and boundary buffers. No additional free initial field is supplied.

Before restoring scale, compare that development with $g_{\lambda,r}$. The initial tensors, future normal and actual spin identification agree. Theorem 2.10, using local geometric Cauchy uniqueness, gives the isometry and actual spin comparison on a common smaller domain of dependence. The physical graph identities of Corollary 5.3 then restore g_λ, including its metric scale. Proposition 5.6 restores the scale of the full generated source and its wave completion.

For families, every initial operation is made on one smooth base selection. Theorem 2.14 gives the smooth Cauchy construction on a common buffered time interval near the chosen parameter value; actual characteristic flow and initialized transport have common buffered domains after the same local restriction. The next proposition states the finite-order interpretation. □

5.6. Estimates on one selected family.

Proposition 5.8 (Finite estimates for one selected family). Let a smooth finite-dimensional family of pointed Ricci-flat germs, or of pointed maximal initial germs, be supplied on common local representatives with smooth markings. The selections in Theorems 5.5 and 5.7, with the scale restoration of Proposition 5.6, can be made on one parameter neighborhood with fixed positive radius, extension branch, initial patch and buffered spacetime and angular parents.

Fix compact parameter sets $\mathcal{K} \subset \operatorname{int} \mathcal{K}'$ in the selected parameter neighborhood, finitely many source charts, a finite complete-coefficient depth J, a point derivative order m and a parameter derivative order q. The selected ordinary pair, executed source, physical embedding and spin comparison satisfy a bound of the form

$$\|\mathcal{Y}\|_{m,q,J}+\|s(c)\|_{m,q}+\|\iota\|_{m,q} \leq F_{m,q,J}\left(\mathcal{N}_N(\text{input and markings}), r^{-1}, \mathfrak{m}^{-1}\right).$$

Here $\mathcal{Y}$ is the stipulated finite list of generated source and wave coefficients and selected carrier and spin-identification maps. Output seminorms are taken over $\mathcal{K}$; $\mathcal{N}_N$ includes the required finite spatial and parameter seminorms of the input and markings over $\mathcal{K}'$. The index N is finite, $F_{m,q,J}$ is finite and locally bounded, and $\mathfrak{m}$ records the positive coordinate, spacelike, elliptic, optical, inverse and domain-buffer margins. The seed seminorm is a finite weighted seminorm of the whole decoded seed. The selected family gives a full source with its physical embedding and spin identification. We choose the parent before specifying m, q, J and estimate the indicated finite list on that parent.

Proof. The normal-coordinate and augmented Dirichlet constructions give the graph and initial data on fixed buffered balls, or one uses the direct initial normalization of Theorem 5.5. Choose the finite base order before the radius. The cutoff prefields then lie in one relative-extension domain.

The extension theorem concerns the same base output at every spatial order. Its base smooth parameter dependence and uniform higher spatial bounds imply the stipulated mixed estimates by Lemma B.13. The gauge,

polar rotation, decoder and ordinary-class transform have finite estimates on that same domain. Include the finite norms of point, tangent-plane and frame markings in $\mathcal{N}_N$.

The source realization has the finite coefficient and parameter estimates of Theorem 4.18. Its Cauchy completion and initialized actual spin transport obey Theorem 2.14. The time interval is fixed by lower Sobolev norms and positive hyperbolic and spacelike margins. Lemma 2.13 proves the higher and actual parameter-derivative estimates for that same solution, allowing the stated finite spatial derivative slack.

Null flow, transverse hit maps, Jacobi equations and finite-rank transport are finite smooth systems on compact buffered domains. At a fixed complete-symbol or module depth, the transport equations involve finitely many earlier coefficients. The wave kernel, mixed carrier transports and all-direction point connection have the finite-coefficient estimates of Proposition 6.11, Proposition 6.17 and Proposition 6.22. Finite angular refinement and those initialized transport estimates therefore give each stipulated coefficient bound on the same parents. Compactness of the real sky supplies one finite angular refinement; increasing the depth does not require changing the parent.

Composition of these estimates yields the displayed bound. Its order and required input seminorms may depend on m, q, J. A smoothly varying positive numerical scale is also allowed when its finite derivatives and positive lower bound are included. □

6. The covariant wave enrichment

The Cartan collar determines the scalar wave kernel class and hence a complete distribution module on every retained spacelike carrier. We construct these objects from the recognized source metric, compare their full sections across carriers and define differentiation in every point direction. The construction applies to the collars generated from ordinary pairs in Section 5, without supplying a further Einstein solution.

To see the effect of changing carriers, take two spacelike hyperplanes through the origin in Minkowski space whose future unit normals have relative rapidity $\theta \neq 0$. Compare the two descriptions of the same null covector at the origin under the identity map of spacetime. If $\rho_i = |\xi_i|$ is its spatial frequency magnitude on hyperplane i, then in boost-adapted coordinates $J = \rho_1/\rho_0 = \cosh\theta - \sinh\theta\,(\xi_0)_1/\rho_0$. Fourier half-density transport multiplies the amplitude by $J^{-1/2}$. It carries the degree-zero metric amplitude to $J^{-1/2}\rho_0^{-3/2} = J\rho_1^{-3/2}$. The factor J varies with the direction at this single point, so a point scalar cannot remove it. For the degree-one wave evaluation the same transport gives $J^{-1/2}\rho_0^{-1/2} = \rho_1^{-1/2}$, with the common normalization constant suppressed. Thus the wave evaluation compares the two slices over the physical identity. Theorem 6.10 establishes uniqueness of the kernel class modulo joint smoothing from its own admission conditions. The relative line in §9 retains its degree-zero calibration.

We first specify the complete-symbol calculus, its internal adjoint and the shifted module conventions. A normal Riccati equation then constructs the positive wave branch; its Klein–Gordon normalization gives a kernel characterized by independent wavefront, bisolution and commutator tests. Mixed traces of that single kernel yield the carrier cocycle and the full degree-one evaluation section. Actual differentiation of the moving distribution families supplies the point connection.

All complete classes have actual classical representatives. The quotient removes kernels jointly smooth in the spatial, point and supplied parameter variables. Each finite estimate controls finitely many homogeneous coefficients and their specified derivatives on the same selected geometric parent. The final recognition theorem concerns this normalized scalar enrichment over an admitted finite Cartan collar. Its comparisons under supplied physical maps leave the source's actual spin lifts distinct.

6.1. Actual realizations and complete coefficient classes. The wave enrichment is constructed on the actual source collar of Section 2. Its metric and physical scale are therefore source readouts. Spatial operators act on complexified real half-densities. The underlying representatives are smooth classical pseudodifferential operators and Fourier integral operators, and the resulting source objects retain their complete classes modulo smoothing. We specify this quotient, its parameter dependence and its algebra and module conventions before using it.

Definition 6.1 (The retained complete calculus).
On an oriented spatial coordinate patch $V \subset \mathbb{R}^3$, a classical left symbol of upper order m has an expansion

$$a(x,\xi;\lambda) \sim \sum_{j\le m} a_j(x,\xi;\lambda), \qquad a_j(x,r\xi;\lambda) = r^j a_j(x,\xi;\lambda), \qquad r>0.$$

The source algebra uses integer orders. Auxiliary positive square roots may have half-integer orders, with the positive physical principal branch specified each time. A fixed cutoff removes zero frequency. Symbols, their asymptotic remainders and all prescribed spatial, angular and finite-dimensional parameter derivatives have locally uniform symbol bounds on larger retained compact parents. Operators are properly supported on the retained domain, with coordinate kernel convention

$$K_A(x,y) = (2\pi)^{-3} \int e^{i(x-y)\cdot\xi} a(x,\xi)\, \mathrm{d}\xi\, |\mathrm{d}x\, \mathrm{d}y|^{1/2}.$$

Two such operator families have the same complete class if their difference is smoothing jointly in the spatial and supplied parameter variables. A complete symbol is the sequence of all homogeneous coefficients in a local presentation. Its finite seminorms concern finitely many of those coefficients and their derivatives on a unit angular section; they do not include every seminorm of a chosen smoothing remainder.

For a conic Lagrangian, use actual local oscillatory distributions with real nondegenerate homogeneous phases and classical amplitudes. Phases and amplitudes depend smoothly on the supplied parameters on fixed buffered charts. Equivalent phase descriptions represent the same distribution class modulo a jointly smooth family. General Lagrangian distributions are allowed even where their projection to the spatial base is singular. A hypersurface-conormal graph description is used only in a chart where that description is valid.

We use the classical conormal and pseudodifferential calculus in Melrose (2006, Chapters 2–4 and 7), and the Fourier integral calculus in Duistermaat (1972, §1, Theorem 1.1). We use the coordinate invariance of oscillatory classes, restriction under the wavefront transversality condition, and the symbol formulas for composition and adjoints. For properly supported transverse canonical graphs, composition adds the orders and multiplies the principal symbols. The identity relation gives the pseudodifferential calculus. An elliptic graph operator has an inverse in this local quotient, and acts on the Lagrangian module carried by its graph. The parameter and finite-coefficient consequences used below are explained in the finite-calculus lemma.

All microlocal identities are asserted on open retained cones with a larger conic buffer. Finite cones cover the whole positive-ray sphere when a full-sky assertion is made. The full kernel characterization will impose its wavefront and smoothing tests on the entire retained point neighborhood squared. A statement on one angular cap alone does not imply that full assertion.

Proposition 6.2 (The complete spatial algebra and its internal adjoint).

The integer-order complete operators form a filtered sheaf of algebras on $P_+T^*\Sigma$, where P_+ means quotient by positive real scaling. In a cotangent chart, rephase each physical coefficient by

$$p_j = i^{-j}a_j.$$

The resulting composition and complex-linear internal adjoint are

$$P \circ Q = \sum_{\alpha} \frac{1}{\alpha!}(\partial_\xi^\alpha P)(\partial_x^\alpha Q), \qquad P^* = \sum_{j,\alpha} \frac{(-1)^j}{\alpha!}\partial_\xi^\alpha \partial_x^\alpha p_j.$$

Each output degree is a finite sum. The internal adjoint is an involutive anti-automorphism covering the same positive ray, and the actual half-density coordinate comparisons preserve it and satisfy strict cocycles.

We call an integer-order operator formally real when every p_j is real. On that real subalgebra the physical Hermitian adjoint $\dagger$ agrees with $*$. Physical bilinear transpose, denoted T, is a different operation: it reverses the cotangent covector before the lower-order corrections. These statements do not identify formal reality with reality of the physical Schwartz kernel.

Proof. The physical left-symbol product is

$$a \# b = \sum_{\alpha} \frac{i^{-|\alpha|}}{\alpha!} (\partial_{\xi}^{\alpha} a)(\partial_{x}^{\alpha} b).$$

A term from degrees j, k and $|\alpha|$ derivatives has output degree $j + k - |\alpha|$. Substituting $a_j = i^j p_j$ and $b_k = i^k q_k$ cancels the degree rephasing exactly and gives the stated product. In particular, the actual differential operator ∂_{x^i} has formal symbol ξ_i.

The physical Hermitian-adjoint symbol is $\sum_{\alpha} i^{-|\alpha|} \partial_{\xi}^{\alpha} \partial_{x}^{\alpha} \overline{a}/\alpha!$. Since $\overline{i^j p_j} = (-1)^j i^j \overline{p_j}$, its rephased form is the displayed expression for P^* with p_j replaced by $\overline{p_j}$. This proves the stated equality on formally real operators, while retaining the different linearities of the two operations.

To verify the internal involution and coordinate law without a reality assumption, work on a cone with one nonzero signed covector component, say ξ_3. Every finite jet of a homogeneous coefficient of integer degree j at a specified nonzero covector can be matched by

$$\xi_3^j p(x, \xi_1/\xi_3, \xi_2/\xi_3),$$

with p polynomial in the angular variables and smooth coefficient functions in x. Such homogeneous Laurent symbols are finite expressions in multiplication, ∂_{x^i} and the microlocal inverse of ∂_{x^3}. On these expressions the internal adjoint is product reversal with $f^* = f$, $(\partial_{x^i})^* = -\partial_{x^i}$, extended to inverses. It is therefore an involutive anti-automorphism. The bilinear differential transpose on half-densities commutes with real coordinate changes, and so does its extension to these microlocal inverses.

At any specified output degree, every composition, coordinate transformation and adjoint coefficient uses finitely many jets at the corresponding cotangent point. The preceding Laurent test proves the involution, anti-multiplication and covariance identities for arbitrary smooth homogeneous coefficients at that degree. It applies at all degrees, establishing the complete identities. This argument uses physical transpose only on the differential/Laurent test algebra; it does not replace the antipodal argument of an arbitrary angular symbol by degree parity.

Finally, a coordinate diffeomorphism f acts by actual half-density pushforward U_f, and the operator comparison is $A \mapsto U_f A U_f^{-1}$. Actual pushforward satisfies $U_{f \circ g} = U_f U_g$, so these comparisons have strict cocycles before and after the quotient. Local asymptotic summation realizes each complete sequence and changes none of these quotient identities. □

6.2. The whole-sky module convention.

Proposition 6.3 (Shifted sky modules and their local normalization).

At $b \in \Sigma$, use conormal half-densities with local unitary Fourier expression

$$(2\pi)^{-3/2} \int e^{i(x-b)\cdot\xi} \sum_{j\le m} i^j a_j(b,\xi)\, \mathrm{d}\xi\, |\mathrm{d}x|^{1/2}, \qquad a_j(b, r\xi) = r^{j-3/2} a_j(b,\xi).$$

The module filtration index is j. Complete classes, with smooth parameter dependence and joint smoothing as above, form modules over the complete algebra on the entire sky $P_+ T_b^*\Sigma$. Actual coordinate comparisons and specified invertible Fourier integral graph comparisons carry their full classes and satisfy the same composition laws as those operators.

In a normalized angular contact chart with sky $z = s(q)$, the physical Fourier amplitude in the positive radial variable τ has the form

$$e^{-i\tau s(q)} \sum_{j\le m} i^j \tau^{j-1/2} a_j(q).$$

Writing its formal class as $\sum_{j\le m} \tau^j a_j(q) u_s$, the exact generator actions are

$$D_z = \tau, \qquad D_{q^a} = \partial_{q^a} - s_{q^a}\tau, \qquad z = -\partial_\tau + \frac{1}{2\tau} + s(q).$$

The generator u_s has degree zero. Any section with a nowhere-zero highest coefficient is cyclic microlocally: every complete module section is obtained by a complete operator acting on it. This assertion holds for smooth families with smooth local choices of the expressing operators. It applies in particular to a degree-one elliptic evaluation section. Specializing the central smooth parameter functions at a fixed parameter gives the corresponding fixed-parameter module.

Proof. The actual pseudodifferential action preserves point conormality. Taylor expansion at b followed by integration by parts gives its rephased amplitude action

$$P \cdot a = \sum_{j,l,\alpha} \frac{(-1)^{|\alpha|}}{\alpha!} \partial_\xi^\alpha \big((\partial_x^\alpha p_j)(b,\xi) a_l(b,\xi) \big).$$

Every output degree uses finitely many coefficients. In this formula ∂_{x^i} acts as ξ_i and $x^i - b^i$ as $-\partial_{\xi_i}$. The principal amplitude times the fiber half-density has homogeneity j, so the convention is coordinate invariant. Explicitly, under $y = Ax$ the leading scalar amplitude becomes $|\det A|^{1/2} a(A^T\eta)$, the required fiber half-density transformation. The general coordinate statement follows from the conormal calculus with its successively lower integer terms.

Here is the angular normalization, including the shift. On $\xi_3 > 0$ set

$$q^a = \xi_a/\xi_3, \qquad z = x^3 + q^a x^a, \qquad \vartheta_a = -x^a, \qquad \tau = \xi_3, \qquad a = 1, 2.$$

The cotangent one-form is $\tau(\mathrm{d}z + \vartheta_a \mathrm{d}q^a)$. Define the local Radon graph operator

$$Rf(q,z) = \int_{\mathbb{R}^2} f(x', z - q \cdot x')\, \mathrm{d}x', \qquad \mathcal{U} = (2\pi)^{-1}(-i\partial_z)R.$$

With unitary Fourier transforms, $\widehat{\mathcal{U}f}(q,\tau) = \tau \widehat{f}(\tau q, \tau)$. The Jacobian is $\mathrm{d}\xi = \tau^2 \mathrm{d}q\mathrm{d}\tau$, so this is an invertible order-zero canonical-graph operator on the retained positive cone. Its inverse is the inverse Fourier-coordinate change. Applying it to the point-conormal expression gives the asserted power $\tau^{j-1/2}$ and phase with $s(q) = b^3 + q \cdot b'$. The same normalized conormal form for a general graph $z = s(q)$ follows by the actual graph translation.

The physical actions are differentiation in z, q and multiplication by z. On the displayed Fourier expression, differentiation of the phase gives the s_{q^a} terms. Physical multiplication by z is $i\partial_\tau$; its action on $\tau^{j-1/2}$ and the degree rephasing give exactly $-\partial_\tau + 1/(2\tau) + s$. Thus the one-half term is fixed by the half-density convention.

For completeness, this angular comparison also respects the internal adjoint. Its exact rephased images on coordinate generators are

$$\partial_{x^a} \longmapsto q^a D_z, \qquad \partial_{x^3} \longmapsto D_z,$$
$$x^a \longmapsto -D_z^{-1} D_{q^a}, \quad x^3 \longmapsto z + D_z^{-1}(E_q + 1), \qquad E_q = \sum_a q^a D_{q^a}.$$

They follow by integration by parts in R. The derivative images are internally skew and the multiplication images are internally self-adjoint. For the last image use $(E_q + 1)^* = -E_q - 1$ and $(D_z^{-1})^* = -D_z^{-1}$. The finite-jet Laurent argument of the preceding proposition extends these identities to all complete smooth angular symbols. No arbitrary contact quantization is asserted by this calculation.

Finitely many signed component cones cover the whole sphere. Actual coordinate maps and the specified Radon operators define the comparisons on overlaps. Their compositions cancel to give exact cocycles in the quotient. A supplied elliptic graph operator carries the same distribution module to the image Lagrangian; an inverse carries it back. This defines the corresponding shifted modules even at a singular projection of a Lagrangian. Local conormal graph presentations are normalized by those actual comparisons, with their half-density and phase factors retained.

To prove cyclicity, let e have upper degree k and nonzero leading amplitude on a retained chart. For a proposed section m of degree at most l, choose an operator of degree $l - k$ whose leading symbol on the Lagrangian is the ratio of their leading amplitudes. Extend that symbol smoothly to a neighboring cotangent cone. Subtract its action on e, and repeat one degree lower. Each step divides only by the same nonzero leading amplitude. Asymptotic summation gives a complete operator P with $m = Pe$. The extensions can be chosen in one fixed local normal chart, so all coefficients depend smoothly on any supplied parameters. This proves existence of the expressing operator;

it need not be unique away from the module. The same argument gives the assertion after an invertible graph comparison.

For central specialization, use local parameter coordinates λ near λ_0 and a normalized phase family. Restriction evaluates each complete amplitude coefficient at λ_0 and is surjective by extending that coefficient smoothly in the parameter. If every restricted coefficient vanishes, Hadamard's integral formula writes each one as

$$a_j(\lambda, q) = \sum_{\mu} (\lambda^{\mu} - \lambda_0^{\mu}) \int_0^1 \partial_{\lambda^{\mu}} a_j\big(\lambda_0 + t(\lambda - \lambda_0), q\big)\, \mathrm{d}t.$$

These coefficient sequences define complete smooth module families of the same upper degree. Thus the kernel is exactly the ideal of smooth parameter functions vanishing at λ_0 times the family module. The coefficient equality is in the joint quotient and is unchanged by phase comparisons. This proves the specialization assertion. □

6.3. Finite coefficient estimates and joint smoothing.

Lemma 6.4 (Finite calculus on one retained parent).
On fixed buffered coordinate, phase and angular parents, complete composition, elliptic inversion with a retained principal inverse margin, chosen simple scalar roots, canonical-graph composition and the preceding module operations preserve smooth finite-dimensional coefficient families. Every specified finite output depth and derivative order is controlled by finitely many coefficient and phase derivatives and the inverses of the retained nondegeneracy margins.

All coefficient sequences have actual smooth classical realizations on those same parents. Different realizations of the same complete family differ by a jointly smoothing operator or a jointly smooth distribution. The estimates control any prescribed finite list of homogeneous coefficients and their indicated derivatives on the retained parent. A choice of asymptotic sum also fixes a smooth remainder whose seminorms need not be controlled by a finite jet.

Proof. The composition formula uses finitely many terms at a specified output degree. To invert an elliptic symbol, choose its inverse principal symbol, cancel the next error by multiplication with that inverse, and proceed by decreasing degree. A simple chosen scalar root is obtained similarly by dividing at each stage by the nonzero derivative of its defining polynomial at the chosen principal root. Thus a positive square root uses the denominator twice its chosen positive principal coefficient. On retained compact angular sets all these denominators have fixed lower bounds. Differentiating each finite recursion proves the claimed finite estimates, with larger but finite input orders.

For transverse canonical-graph composition, the stationary-phase expansion at a given degree uses finitely many amplitude and phase derivatives and inverse powers of the nondegenerate transverse Hessian. The proper support

and graph hypotheses are precisely those of the cited composition theorem. The retained Hessian margin controls those inverses. An elliptic operator on an inverse graph can first be chosen with inverse principal symbol; its product is an elliptic pseudodifferential operator with principal symbol one. The preceding complete inverse corrects it to an actual quotient inverse. Left and right inverses agree by associativity. This also proves the local algebra-conjugation and module-action assertions used above, with the same finite parameter estimates.

To realize a sequence, take homogeneous extensions of its coefficients and multiply its successive terms by cutoffs $\chi(\xi/R_j)$, with $R_j \to \infty$. Choose R_j successively so that the jth term contributes less than 2^{-j} to the first j relevant symbol seminorms after lowering the tested order by a fixed fraction of j. Include the first j parameter derivatives on the first j compact sets in an exhaustion. For every fixed derivative and remainder order, all sufficiently late terms then converge in that seminorm. The resulting sum is smooth in the parameters and has the prescribed expansion. This is the parameter version of classical asymptotic summation (Melrose 2006, Chapter 2, Proposition 5); the cutoff choices do not change the spatial or conic parent.

If all homogeneous coefficients vanish, the symbol and every parameter derivative are rapidly decreasing. Its localized Fourier integral is therefore jointly smooth, by differentiating the integral at arbitrary finite orders. Conversely, the local complete coefficient of a smooth distribution is zero. The same statements hold in an equivalent nondegenerate phase by the invariant symbol calculus. They prove that changing the asymptotic sum or its phase description leaves exactly the stated joint quotient. They also show that a parameter derivative of a joint smoothing remainder is again joint smoothing.

The phase of a flow graph is fixed by its actual Hamiltonian flow on the retained tube. Its finite derivatives satisfy finite variational ODEs; amplitude transport at each degree is a linear ODE with forcing from finitely many earlier degrees. With a uniform flow and phase nondegeneracy margin, the chain rule and Gronwall estimates give the corresponding finite bounds on the same tube. No later coefficient changes the flow domain. A finite angular refinement handles all retained directions. These observations establish the last needed parameter calculus statements, without a uniform estimate as the requested depth tends to infinity. □

In particular, a pointwise choice of smooth remainder at every parameter value is not the equivalence relation used here. Kernel admission and comparison will use the joint calculus just defined. The actual finite Cartan and spin fields remain separate geometric structures; quotienting scalar wave representatives by joint smoothing does not quotient their physical scale or actual spin comparisons.

6.4. The wave operator and its positive normal branch. The metric in this construction is the metric recognized from the actual Cartan source

in Theorem 2.9. For an ordinary-pair display it has already been constructed from the pair. The wave construction introduces no additional Einstein datum. Its arguments apply to any smooth Lorentzian collar of the stated signature.

Retain a Gaussian coordinate neighborhood $g = \mathrm{d}t^2 - h_t$, a smaller point neighborhood and buffered initial coordinate disks. The neighborhoods are small enough that the retained null trajectories stay in the larger parent. A finite angular cover retains all nonzero null directions over the point neighborhood. When invoking a scalar Cauchy or Green operator, restrict to a smaller causally convex globally hyperbolic neighborhood inside this parent. An arbitrary product coordinate cylinder is not itself assumed to have that property.

We use the actual classical operator and joint-smoothing calculus of the symbol conventions. All equalities of complete symbols and parametrices below are in that quotient. Physical Hermitian adjoint is written $\dagger$, physical bilinear transpose is written T, and the complex-linear internal adjoint of the integer-order formal algebra is written $*$. These operations have different meanings. In particular, the frequency branch is fixed by the flat solution $\exp(-it|D|)$; its temporal covector is $-H$ in the signature $(+---)$.

Proposition 6.5 (Half-density wave reduction and the instantaneous defect).
Put

$$\rho_t = (\det h_t)^{1/4},$$
$$K_t = \tfrac{1}{2}\dot{h}_t,$$
$$\tau_t = \mathrm{tr}_{h_t} K_t,$$
$$D_{h_t} = -\rho_t \Delta_{h_t} \rho_t^{-1},$$
$$H_t(x,\xi) = |\xi|_{h_t}.$$

There is a unique complete classical square root A_{h_t} of D_{h_t} with physical principal symbol H_t. It is physically Hermitian. The operator $B_{h_t} = iA_{h_t}$ has real integer-degree formal coefficients and is internally skew-adjoint. In physical left symbols,

$$\sigma_0(A_{h_t} - \mathrm{Op}(H_t)) = \frac{1}{2i}\sum_j \partial_{x^j}\partial_{\xi_j} H_t.$$

Writing $V_t = \dot{\tau}_t/2 + \tau_t^2/4$, one has

$$\rho_t \Box_g \rho_t^{-1} = \partial_t^2 + D_{h_t} - V_t.$$

An instantaneous root parametrix satisfying $(\partial_t + iA_{h_t})G_{t,s} = 0$ consequently has wave residual

$$(\rho \Box_g \rho^{-1})G_{t,s} = (-i\dot{A}_{h_t} - V_t)G_{t,s}.$$

Its order-one coefficient is $-i\dot{H}_t = iK_t^{ij}\xi_i\xi_j/H_t$.

Proof. The positive principal root determines every lower coefficient of A_h by division by $2H$. At the step determining degree $1-j$, this coefficient occurs in the square only through $2Ha_{1-j}$ at degree $2-j$; differentiated products involve already determined coefficients. The complete recursion is therefore unique on the nonzero cone. The parameter-dependent asymptotic realization gives an actual classical representative. Since $D_h^\dagger = D_h$, taking the physical adjoint gives another square root with the same principal symbol. Uniqueness proves $A_h^\dagger = A_h$.

For B_h, the square equation is $B_h^2 = -D_h$. Its formal principal symbol is H, and the differential operator $-D_h$ has real formal coefficients. The same recursion in the rephased product therefore gives real coefficients at every integer degree. Physical Hermitian adjoint and internal adjoint agree on such operators. The equality $(iA_h)^\dagger = -iA_h$ proves $B_h^* = -B_h$.

In coordinate half-densities the degree-one physical symbol of D_h is $(1/i)(\partial_{x^i} h^{ij})\xi_j$. It equals $(2i)^{-1}\sum_j \partial_{x^j}\partial_{\xi_j}(H^2)$. The degree-one coefficient of the square of $H + a_0 + \cdots$ is

$$2Ha_0 + \frac{1}{i}\sum_j (\partial_{\xi_j} H)(\partial_{x^j} H).$$

Comparison gives the stated coefficient. Equivalently, the physical subprincipal symbol of A_h on half-densities is zero. The root is taken in the local complete-symbol calculus.

The scalar divergence formula gives $\Box_g = \partial_t^2 + \tau\partial_t - \Delta_h$, and $\dot\rho/\rho = \tau/2$. Direct conjugation gives the displayed reduction. Twice differentiating the instantaneous evolution gives its residual. Finally $\partial_t h^{ij} = -2K^{ij}$ implies $\dot H = -K^{ij}\xi_i\xi_j/H$, proving the last statement. □

For a scalar wave solution let $v = \rho u$ and $\pi = \rho\partial_t u$. The two field variables are half-density coefficients on the same coordinate disk. Their equations are

$$\dot v = \frac{\tau}{2}v + \pi, \qquad \dot\pi = -D_h v - \frac{\tau}{2}\pi.$$

The normal branch below corrects the instantaneous defect while retaining the same real null characteristics.

Theorem 6.6 (The complete positive normal symbol).
There is a unique complete classical symbol Q_t of order one with physical principal symbol iH_t whose graph $\pi = -Q_t v$ is invariant under the scalar wave equation modulo joint smoothing. It obeys

$$Q_t^2 - \dot Q_t - Q_t\frac{\tau_t}{2} - \frac{\tau_t}{2}Q_t + D_{h_t} = 0.$$

Its integer-degree formal coefficients are real, and $Q_t^T = Q_t$. Put $B_t = Q_t - \tau_t/2$. There are initialized order-zero Fourier integral parametrices $E_{t,s}$ with

$$(\partial_t + B_t)E_{t,s} = 0, \qquad E_{s,s} = I, \qquad E_{t,r}E_{r,s} = E_{t,s}.$$

They solve the wave equation modulo joint smoothing and have the null canonical graph selected by $\exp(-itH)$. In Gaussian coordinates,

$$B_t = B_{h_t} + \operatorname{Op}\left(\frac{\dot H_t}{2H_t}\right) \pmod{\Psi^{-1}}.$$

Proof. Substitute $\pi = -Qv$ into the canonical system. The equation for $\dot\pi$ gives $Q^2 - \dot Q - Q\tau/2 - \tau Q/2 + D_h = 0$, with the indicated order of the noncommuting factors.

In the rephased calculus the highest equation is $q_1^2 - H^2 = 0$. The prescribed branch is $q_1 = H$. Suppose the coefficients above degree $1-j$ have been determined. At degree $2-j$ the new coefficient occurs only as $2Hq_{1-j}$. Time differentiation does not raise symbol order, the terms containing τ have order one, and every differentiated symbol product lowers total order. All their contributions at the current degree are therefore already known. Division by $2H$ determines q_{1-j}. The argument also removes successively every degree of the difference of two possible solutions. The real coefficients of the equation in the formal product show that every q_{1-j} is real.

Physical transpose reverses products, commutes with time differentiation and interchanges the two τ terms. Also $D_h^T = D_h$. Its action on the principal coefficient gives $iH(x,-\xi) = iH(x,\xi)$. Consequently Q^T satisfies the same branch equation. On the retained full angular cover, uniqueness proves $Q^T = Q$.

The complete recursion is realized by a classical family using the finite-symbol family calculus. For the actual parametrix, in a noncaustic chart solve

$$\partial_t S + H_t(x, \partial_x S) = 0, \qquad S(s,x,\eta) = x\cdot\eta,$$

by the actual Hamiltonian flow. The initialized mixed projection is invertible for a common sufficiently short interval on each compact unit-direction chart. A finite cover supplies one interval for all retained directions. Seek a kernel

$$E_{t,s}(x,y) = (2\pi)^{-3}\int e^{i(S(t,x,\eta)-y\cdot\eta)} w(t,s,x,\eta)\,\mathrm{d}\eta\,|\mathrm{d}x\,\mathrm{d}y|^{1/2}.$$

Composition with $\partial_t + B_t$ first gives the Hamilton–Jacobi equation. At the next order it gives a linear transport equation for the leading amplitude along the same characteristics, with initial value one. Every successive homogeneous amplitude satisfies that same transport operator with a forcing involving finitely many derivatives of earlier amplitudes and already determined coefficients of B_t; its initial value is zero. Integration along the retained flow determines each coefficient as an actual smooth function. The leading solution is nonzero because it is an initialized exponential transport. Uniform asymptotic summation produces a classical amplitude, with a jointly smoothing residual and identity initial value modulo joint smoothing.

These are the local geometric-optics parametrices in the shared Fourier integral calculus. The wave factorization is also obtained in Gérard and Wrochna (2014, Theorem 6.3, Remark 6.5 and Appendix A.3); the negative

time phase fixes which of its two branches is used here. The required localization is explicit. Extend h_t from a larger buffered coordinate patch to a positive metric on $\mathbb{R}^3$ by convex interpolation with a fixed Euclidean metric, and extend it smoothly in time outside the retained interval. The coefficients of $D_h - V$ then have bounded spatial derivatives and a uniformly positive principal spatial matrix on each retained time interval. Thus they satisfy the model conditions in that paper's §5.1. No vacuum equation is required of this auxiliary extension. Proposition 4.6 and Corollary 6.6 there give the corresponding characteristic propagation; in the local construction it is the canonical relation of the displayed phase. Every retained trajectory stays inside the original metric patch.

The amplitude recursion proves uniqueness of the initialized parametrix class. Composition on common buffered canonical graphs gives another parametrix with the same initialization, so it proves the composition law and the inverse $E_{s,t}$. Substitution of its equation and the Riccati identity into the canonical system proves the wave equation. The smoothing ideal is preserved under these properly supported compositions.

For the first correction write $B = ia$. The conjugated wave equation gives $a^2 + i\dot{a} = A_h^2 - V$. At degree one, $2H\sigma_0(a - A_h) + i\dot{H} = 0$. Multiplication by i gives the stated correction. Its principal Hamiltonian is unchanged. □

Proposition 6.7 (Klein–Gordon normalization of the positive branch).

Define

$$R_t = i(Q_t^\dagger - Q_t), \qquad W_t = R_t^{1/2}, \qquad U_{t,s} = W_t E_{t,s} W_s^{-1},$$

where the complete square root has physical principal symbol $\sqrt{2H_t}$. The operator R_t is physically Hermitian and has real, transpose-symmetric kernel modulo joint smoothing. Its principal symbol is $2H_t$. The normalized propagator is physically unitary and composes in the quotient. Its integer-order generator satisfies

$$B_t^{\mathrm{cov}} = W_t B_t W_t^{-1} - \dot{W}_t W_t^{-1},$$
$$(B_t^{\mathrm{cov}})^* = -B_t^{\mathrm{cov}},$$
$$B_t^{\mathrm{cov}} = B_{h_t} \pmod{\Psi^{-1}}.$$

Proof. From $Q^T = Q$ one obtains $Q^\dagger = \overline{Q}$, where the bar is physical kernel conjugation. Hence $R = i(\overline{Q} - Q)$ is real and transpose symmetric, as well as Hermitian, with principal symbol $2H$. Inverse and square-root recursion give W and W^{-1} in the complete calculus. Adjointing $W^2 = R$ preserves the prescribed principal root, so uniqueness gives $W^\dagger = W$. A globally positive exact spectral representative is not required.

The Klein–Gordon form in canonical half-density variables is

$$i \int_\Sigma (\overline{v}\,\pi - \overline{\pi}\,v).$$

On the positive graph it equals $\langle v, Rv\rangle$. Direct subtraction of the Riccati equation and its Hermitian adjoint gives

$$\dot{R} = B^\dagger R + RB, \qquad E_{t,s}^\dagger R_t E_{t,s} = R_s.$$

The second equality follows by differentiating its left-hand side and using its initial value. It proves $U_{t,s}^\dagger U_{t,s} = I$ and physical anti-Hermiticity of B^{cov}. Composition follows from the corresponding law for E.

For the internal assertion, formal reality of Q gives $Q^\dagger = Q^*$. Put $F = Q - Q^*$ in the integer-order formal algebra. Its complete coefficients are real and its formal principal coefficient is $2H$. The formal half-order root $F^{1/2}$ has real coefficients by its recursion. With $i^{1/2} = e^{i\pi/4}$ in the physical degree conversion, the positive physical root has formal expression

$$W = e^{-i\pi/4} F^{1/2}, \qquad W^2 = -iF = R.$$

The scalar phase cancels from both WBW^{-1} and $\dot{W}W^{-1}$. Their difference has real coefficients at every integer degree. Physical Hermitian adjoint and internal adjoint therefore agree on this generator, giving the asserted internal skew-adjointness. This reasoning does not extend the internal involution to arbitrary half-integer degrees.

At degree zero, the correction $\dot{H}/(2H)$ in B cancels the principal coefficient of $\dot{W}W^{-1}$. The possible degree-zero term in $WB_hW^{-1} - B_h$ is proportional to the Poisson bracket $\{\sqrt{2H}, H\} = 0$. All other terms have degree at most -1. This proves the final formula. □

6.5. The independent scalar kernel class. We first fix the causal-kernel normalization used in the source test. The point slots of the kernel are scalar; its integral action uses the physical volume measure. Half-density factors are inserted only when taking traces on spacelike disks.

Lemma 6.8 (Local causal kernel and its double Cauchy traces). On a sufficiently small protected Cauchy neighborhood there are local retarded and advanced inverses for $\Box_g$ on compactly supported smooth sources. Their difference $E_g = E_g^{\mathrm{ret}} - E_g^{\mathrm{adv}}$ is a real antisymmetric scalar bisolution kernel. With field and future-normal half-density traces on a common initial disk, its four double traces, in field/normal order, are

$$\left((\gamma_i \otimes \gamma_j) E_g\right)_{i,j=0,1} = \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix}.$$

Here $\gamma_0 u = \rho u|_\Sigma$ and $\gamma_1 u = \rho n(u)|_\Sigma$. The row index i refers to the first kernel variable and the column index j to the second. For a kernel admitting these noncharacteristic Cauchy traces, smooth errors in both wave equations and all four double traces determine a jointly smooth error on a smaller product neighborhood. The required restrictions exist automatically for the characteristic wavefront classes used below. The same conclusions hold jointly with supplied smooth finite-dimensional metric parameters.

Proof. Use the scalar linear specialization of the hyperbolic input specified in Theorem 2.6, with the same protected coordinate extension. Its estimates are the linear wave energy estimates written in Lemma 2.13. One may obtain the smooth scalar evolution directly by spatial Friedrichs regularization of the first-order wave system: the differentiated energy inequality is uniform in the regularization, compactness gives a solution, and the zero-data energy inequality gives uniqueness. Applying this at every Sobolev order gives the canonical Cauchy evolution for smooth compactly supported pairs (v, π) on one common short interval. Duality, followed by localization, extends the solution operators to compactly supported distributions. Their kernels are distributions by the Schwartz kernel theorem.

For a compactly supported forcing, choose an initial time before its support and solve the inhomogeneous equation with zero data by Duhamel's formula. This gives the retarded inverse on the retained domain. Solving backwards from a time after the support gives the advanced inverse. Finite propagation makes both independent of the auxiliary endpoint and extension when restricted to the smaller protected neighborhood. The local energy estimate gives uniqueness with these support conditions. Integration by parts for the scalar wave operator against the physical volume gives $(E^{\mathrm{ret}})^T = E^{\mathrm{adv}}$. The coefficients are real, so both kernels are real. Their difference is therefore real, antisymmetric and a bisolution.

For the trace normalization, let $\mathcal{S}(t, s)$ be the canonical evolution matrix for the system, initialized by the identity. The field part of its response to initial normal data is its upper-right block. Duhamel's formula identifies the half-density field-field kernel of $E_g(t, s)$ with this block. It vanishes at $t = s$, and its first normal trace there is the identity. Antisymmetry gives the opposite field-normal trace. At equal times the canonical generator has blocks

$$\begin{pmatrix} \tau/2 & I \\ -D_h & -\tau/2 \end{pmatrix}.$$

The upper-right block of its square is zero. Equivalently, differentiate the initialized field-normal trace, including the $\tau/2$ corrections that convert $\partial_t v$ to π. This gives the zero normal-normal trace. These are the four stated identities, with the sign of $E^{\mathrm{ret}} - E^{\mathrm{adv}}$ fixed by the forward normal response.

For the smooth-error assertion, first restrict a proposed kernel error and its normal derivative to the first initial disk. Each is a distribution in the second spacetime variable, depending on the first spatial variable and any external parameters. Its second-variable wave equation and both initial traces are jointly smooth. The linear Cauchy estimate at all orders, with these variables retained as parameters, makes both restricted errors jointly smooth. The first-variable Cauchy equation, its now smooth initial data and its smooth forcing then make the entire error jointly smooth. Derivatives in the passive parameters satisfy linear inhomogeneous equations with finitely many earlier derivatives. The same energy induction proves the joint parameter assertion on the same smaller parent. □

Let $\mathcal{C}_g^+$ denote the primed null relation whose first temporal covector in a Gaussian chart is $-H$ and whose two covectors lie on the same null bicharacteristic. This notation fixes a phase convention; it is independent of the auxiliary Gaussian chart under time-oriented physical comparisons. Its transposed unprimed relation corresponds, after priming, to the opposite null component. For a quasifree Klein–Gordon field on a fixed globally hyperbolic spacetime, Radzikowski (1996) characterized the Hadamard condition by the wavefront set of the two-point distribution. Definition 6.9 specifies the positive-frequency condition in our phase convention and its extension to joint parameter variables.

Definition 6.9 (Independent admission of the wave kernel).
Over an actual source collar, supply a scalar distribution kernel class $[\Lambda]$ modulo jointly smooth kernels. Require actual classical oscillatory representatives on the retained null relation, in the uniform family calculus, and test

$$\Box_{g,p}\Lambda = 0, \qquad \Box_{g,p'}\Lambda = 0, \qquad \mathrm{WF}'(\Lambda) \subset \mathcal{C}_g^+, \qquad \Lambda - \Lambda^T = -iE_g,$$

where the three kernel equations are modulo joint smoothing. All nonzero null directions over the retained point neighborhood are included.

For a parameter domain $\mathcal{P}$, every representative and residual uses the joint variables (λ, p, p'). In particular, its joint wavefront set is required to lie in

$$\left\{(\lambda, p, p'; \sigma, \xi, -\xi') : (p, \xi; p', \xi') \in \mathcal{C}_{g_\lambda}^+\right\}.$$

The spacetime covectors in this set are nonzero. Pure parameter singularities are excluded. The parameter covector σ is the one furnished by the uniform oscillatory family; it is not set to zero when the metric or phase varies. These conditions concern the supplied kernel and recognized source fields. Membership in a constructed kernel image, a choice of quantum state and a chosen exact smoothing tail are not admission conditions.

Theorem 6.10 (Canonical admitted wave kernel).
The source metric determines exactly one independently admitted scalar kernel class Λ_g^+ on a sufficiently small retained neighborhood squared. It is physically Hermitian modulo joint smoothing, natural under restriction and time-oriented physical isometry, and independent of the auxiliary Gaussian disk, quantization, cutoff and asymptotic realization.

On a Gaussian collar with initial time s, the class is represented by

$$\widetilde{\Lambda}^+(t, t') = E_{t,s} R_s^{-1} E_{t',s}^\dagger, \qquad \Lambda^+(t, x; t', y) = \rho_t(x)^{-1} \widetilde{\Lambda}^+(t, x; t', y) \rho_{t'}(y)^{-1}.$$

The uniqueness, equations and naturality hold in the joint class for smooth finite-dimensional families.

Proof. The normal-branch and charge results give the displayed expression. The two first-order evolutions and the Riccati identity make it a bisolution modulo joint smoothing. Its real phase is the retained null canonical relation, with the fixed frequency component. Products and proper support are

taken on compactly contained ray domains; the construction covers a smaller neighborhood in all directions by the finite angular cover. Changes of low-frequency cutoff contribute a jointly smooth kernel. The construction and its Hermitian transpose agree because R_s^{-1} is Hermitian.

Put $C = R_s^{-1}$. It is real and transpose symmetric. We verify all four double Cauchy traces of the antisymmetric difference, using $\pi = -Qv$ and $Q^T = Q$. They are

$$\begin{aligned}
\Lambda_{00} - \Lambda_{00}^T &= C - C^T = 0,\\
(\Lambda - \Lambda^T)_{10} &= (-Q + Q^\dagger)C = -iI,\\
(\Lambda - \Lambda^T)_{01} &= C(Q - Q^\dagger) = iI,\\
(\Lambda - \Lambda^T)_{11} &= QCQ^\dagger - Q^\dagger CQ = 0.
\end{aligned}$$

The last identity keeps the noncommuting order of every factor: since $Q - Q^\dagger = iR$ and $C = R^{-1}$, its difference is $iRCQ^\dagger - iQ^\dagger CR = 0$. The preceding Green lemma identifies these traces with those of $-iE_g$. Its successive Cauchy argument in the two variables makes the remaining bisolution error jointly smooth. Thus the commutator equation holds as a complete kernel identity, not just at principal order.

For uniqueness, let D be the difference of two admitted kernels. Their commutator tests give $D - D^T$ jointly smooth. Their joint wavefront conditions put $\mathrm{WF}(D)$ in the positive relation of the definition. Transposition sends its nonzero spacetime covectors to the opposite component. Since equality modulo joint smoothing identifies the wavefront sets of D and D^T, that wavefront set lies in the intersection of the two components. The intersection is empty. The exclusion of pure parameter covectors makes this a statement about the full joint wavefront set. Hence D is jointly smooth.

All auxiliary constructions satisfy the same tests, so uniqueness removes their choices in the stated quotient. Restriction to a smaller causally convex neighborhood preserves the causal kernel by finite propagation and preserves the other tests. A physical isometry carries the volume-normalized wave equation, causal kernel and chosen null component to their counterparts. Pullback therefore satisfies the same tests and gives the unique class. For a smooth isometry family its parameter derivative changes the parameter covectors of the phase, but it preserves the nonzero spacetime component and the joint smooth ideal. This proves joint naturality as well. □

Proposition 6.11 (Finite coefficient estimates on one wave parent).

Let a smooth finite-dimensional family of actual source collars be given on the common parent of Theorem 2.14. After one further restriction near the chosen parameter value, the normal symbols, charge roots, propagator classes and admitted kernels above are constructed on one point parent and a finite angular phase cover, chosen before any prescribed coefficient depth or derivative order.

For each finite coefficient depth J, spacetime/angular derivative order m and parameter derivative order q, the corresponding finite list $\mathcal{Z}_J$ of symbols, phase functions and amplitudes satisfies a bound

$$\|\mathcal{Z}_J\|_{m,q} \le F_{J,m,q}(\mathcal{N}_N(g, \text{markings}; \mathcal{K}' \times \text{larger parent}), \mathfrak{m}^{-1}),$$

for a finite N and a finite locally bounded function $F_{J,m,q}$. Here $\mathcal{K} \Subset \operatorname{int} \mathcal{K}'$ are the retained parameter sets, angular norms use compact unit-covector sections, and $\mathfrak{m}$ includes spacelike, inverse-phase, positive-root and domain-buffer margins. Every finite coefficient of Q, R, W is a local expression in a finite jet of the metric. The wave kernel has its unique joint class; the estimate selects no exact smoothing remainder.

Proof. On the compact normalized cones, positivity of h gives a uniform lower bound for H. The recursive divisions defining A_h, Q, R^{-1} and W have denominators consisting of positive powers of H or $\sqrt{2H}$. At each fixed depth only finitely many coefficients and derivatives occur. The product, inverse and derivative estimates of the shared symbol calculus therefore give the indicated finite bounds for this part.

The Hamiltonian flow and its mixed generating map satisfy finite ODE systems. Their low-order bounds select the retained time and phase parents, including a positive inverse margin, before higher derivatives are requested. Smooth dependence and their differentiated ODEs give every finite mixed estimate on those same parents. Compactness of the unit-direction sphere gives a finite cover with a common positive buffer. Increasing coefficient depth requires higher input norms but does not change that cover or the chosen trajectories.

Each amplitude coefficient solves a linear transport equation along this already selected flow. Its forcing is a finite differential expression in earlier coefficients, the phase and the operator symbol. The initialized leading coefficient is a nonzero exponential transport. Induction over coefficient depth and differentiation in the finite parameters give the required bounds by the corresponding integral equations. Composition and stationary phase at fixed depth use finitely many further derivatives and the retained inverse-phase margins. Multiplication by the scalar density factors has the same finite dependence.

Lemma 6.4 realizes the amplitude families by asymptotic summation on the selected phase charts. To verify the joint wavefront condition, observe that the nonzero spacetime covectors control the homogeneous phase variable on each chart, while the parameter derivatives of the phase are bounded on its compact unit section. These bounds exclude a pure parameter wavefront direction. The kernel theorem then identifies every such realization with the same joint class. This proves the stated finite-order family conclusion. □

6.6. Spacelike carriers and the complete comparison operators. The scalar kernel is determined on the actual source collar. A spacelike disk gives a carrier on which its traces become operators on half-densities. We compare

different carriers by the same kernel, including carriers that intersect. All operator identities in this subsection are identities in the complete-symbol and joint-smoothing quotient of the calculus convention.

Definition 6.12 (Retained wave carriers).
Let B be a relatively compact point neighborhood in a source collar and let $F_a : \Sigma_a \to M$ be smooth spacelike embeddings with future unit normals. Retain open homogeneous ray domains on which every relevant null ray meets each carrier once, transversely and away from its boundary. Compositions use a smaller common ray domain with compact angular closure and properly supported representatives. To obtain a full-sky assertion over B, these domains must contain every nonzero null direction over B, using a finite angular cover.

For each comparison use a single open wave neighborhood W in the developed collar. It is globally hyperbolic and causally convex in the collar and contains the point core, the retained carrier hits and the ray segments joining them. The kernel tests on $W \times W$ use the retarded-minus-advanced propagator for its physical volume. Only the retained portions of each carrier enter the comparison.

For a smooth family the preceding conditions hold for (W_λ, g_λ) at each parameter. Use a smooth spacelike Cauchy hypersurface S_λ of each W_λ as an auxiliary initial surface. Near a reference parameter λ_0, require smooth identifications $\Phi_\lambda : N \to W_\lambda$ onto their images on neighborhoods $N \subset W_{\lambda_0}$ of the compact sets used, with Φ_{λ_0} the inclusion and $\Phi_\lambda^{-1}(S_\lambda) = N \cap S_{\lambda_0}$. Retain positive spacelike, transversality and interior margins there. At a single metric, the existence of a smooth spacelike Cauchy surface follows from Bernal and Sánchez (2005), Theorem 1.1.

For the local displays of §5 one can choose these neighborhoods about any marked point. Inside the causally convex development take a small coordinate cylinder around the point with a spacelike coordinate slice through it. On a smaller cylinder a constant Lorentz cone strictly contains all the metric causal cones, including those for nearby parameters. Use the development time as coordinate time. On a compact subcylinder the spatial speeds of causal curves in this time are uniformly bounded. Choose the time span of a central wider-cone diamond so small that this bound times its span is less than the spatial buffer to the subcylinder wall. Time is monotone on the development. A causal curve between points of the diamond must therefore stay in the subcylinder, by the speed bound up to a first exit. Cone inclusion then keeps it in the diamond. The central slice of the diamond is Cauchy for each metric: toward that slice the sections of the diamond expand faster than any causal curve moves in the spatial coordinates, so such a curve continues to the slice. Strict time monotonicity gives a unique crossing. A smaller point core and an interior disk of the slice retain all directions, by compactness of the unit sky and smooth transverse dependence of the hits.

The comparison with any further carrier uses the common rays for which the specified margins hold.

Write $h_a = -F_a^* g$ and, in local coordinates, $\rho_a = (\det h_a)^{1/4}$. With the coordinate half-density factors understood, define

$$C_{ab}(x,y) = \rho_a(x)\rho_b(y)\Lambda_g^+(F_a(x), F_b(y)),$$
$$C_a = C_{aa}, \qquad W_a = C_a^{-1/2}.$$

Thus C_{ab} acts from half-densities on Σ_b to half-densities on Σ_a. On the retained nonzero cone W_a is the complete elliptic inverse square root of C_a with positive physical principal symbol.

The restrictions are defined. A nonzero null covector cannot annihilate the tangent space of a spacelike hypersurface, so the kernel wavefront relation misses the conormal obstruction to both pullbacks. By constructing the canonical kernel in a Gaussian collar based at Σ_a, its diagonal trace is

$$C_a = R_a^{-1}, \qquad \sigma_{-1}^{\mathrm{phys}}(C_a) = (2H_a)^{-1}, \qquad H_a = |\xi|_{h_a}.$$

Here $R_a = i(Q_a^\dagger - Q_a)$ is the normal charge operator of the charge normalization. Uniqueness of the canonical kernel makes this the same C_a for every auxiliary Gaussian construction. In particular $C_a^\dagger = C_a$, and the square-root recursion gives $W_a^\dagger = W_a$ and $W_a C_a W_a = I$.

The common ray domains identify $T^*\Sigma_b \setminus 0$ with $T^*\Sigma_a \setminus 0$ by the null characteristic flow. This is a homogeneous canonical graph: restriction of the characteristic symplectic form to either spacelike section gives its cotangent symplectic form. The field trace of the chosen wave branch is an elliptic order-zero Fourier integral operator on this graph. In a short phase chart this follows from the real Hamilton–Jacobi phase and its nonvanishing initialized transport amplitude. Finite composition of such charts gives the retained graph operator and its inverse in the quotient. The graph composition and inverse statements use the stated finite-symbol calculus.

Mixed traces on a common wave neighborhood. Write (W, S) for one of these wave neighborhoods and its Cauchy surface. Construct the kernel on W^2 from its four scalar Cauchy traces. In Gaussian patches of S let $J_s = (I, -Q_s)^{\mathsf{t}}$ be the positive Cauchy column in the field and physical normal half-density traces (v, π). By Theorem 6.10 the four initial traces form the matrix

$$A_s = J_s R_s^{-1} J_s^\dagger.$$

Here the row is $(I, -Q_s^\dagger)$, with the physical adjoint. On overlapping initial patches the four kernels agree near the diagonal modulo joint smoothing. Indeed, restrict the two short spacetime kernels to a common causally convex neighborhood of each point of the overlap and use uniqueness in Theorem 6.10 before taking traces. On S^2 every possible singular pair lies on the diagonal since S is acausal. Glue the four properly supported classes there by a locally finite partition in the first variable and right cutoffs equal to one near each partition support, extending each term within its patch. The

resulting matrix A has the four specified complete symbols. Changing cutoffs or partitions changes it by a smooth matrix kernel, jointly in parameters.

Let $\mathcal{U}$ send the two initial half-density traces to the spacetime field by ordinary scalar Cauchy evolution. For a fixed metric, $\Box_g$ is normally hyperbolic, W is globally hyperbolic and S is a smooth spacelike Cauchy hypersurface. Bär, Ginoux and Pfäffle (2007), Theorems 3.2.11–3.2.12 give smooth well-posedness, continuous dependence on the data and propagation inside their causal future and past. Testing against compactly supported smooth sources and solving the adjoint inhomogeneous equation gives the distributional evolution by duality. Evolve the matrix A in its two variables to obtain $\mathcal{U}A\mathcal{U}^{\dagger}$ with the volume and surface pairings. The second evolution solves the conjugate equation. The evolutions commute. Smooth changes in all four traces have smooth double evolutions by Lemma 6.8. Applying this to the Hermitian difference and to the four antisymmetric traces of Theorem 6.10 gives physical Hermiticity and the commutator $-iE_g$ in the quotient.

Localization of the initial data. Take compact neighborhoods K_a, K_b of the endpoint patches in W, with room to restrict the patches, and set $K = K_a \cup K_b$. At a reference parameter the set swept out by causal segments from K toward S up to their first hit is compact. Indeed it is the union of $J^-(K) \cap J^+(S)$ and $J^+(K) \cap J^-(S)$, both compact by Bär, Ginoux and Pfäffle, Corollary A.5.4. Its intersection with S is the causal shadow of K. A spatial cutoff equal to one near this shadow leaves the solution on the endpoint patches unchanged by finite propagation. Apply it in each initial variable.

Here the same compact localization works near the reference parameter. Choose a relatively compact open neighborhood O of the swept set, with closure in the reference wave neighborhood. Use the local identifications on a neighborhood of its closure, straightening the varying surface there to S. At the reference metric a temporal function t with $S = t^{-1}(0)$ exists by Bernal and Sánchez (2006), Theorem 1.2. On the closure of O, $\mathrm{d}t$ remains timelike for nearby metrics; hence future causal vectors normalized by $\mathrm{d}t(v) = 1$ have uniformly bounded auxiliary Riemannian length. If segments from the endpoint compacts toward S could first exit O before hitting it for parameters tending to the reference value, parameterization by t would yield a uniformly Lipschitz subsequence up to first exit. Its limit joins K to the boundary on the same side of $t = 0$, allowing a boundary endpoint at $t = 0$. The limit is causal: on short coordinate subarcs the tangent vectors lie in arbitrarily small convex enlargements of the reference causal cones, and the increment bound passes to the limit. Extend the limiting segment toward the Cauchy surface if necessary. Its boundary endpoint belongs to the reference swept set, which lies in the interior of O. This contradiction gives the common localization. One can take the initial cutoff equal to one on $S \cap \overline{O}$ with slightly larger compact support. The argument uses the metric only on compact neighborhoods and works in either time direction.

Positive branch on the localized data. Cover the entire unit cosphere over that compact support by finitely many coordinate cones. On smaller cones use order-zero symbols forming a partition of unity at high frequency, quantized to sum to the identity modulo smoothing there. In each larger cone the initial matrix is $J_s R_s^{-1} J_s^\dagger$ to every order. Insert the partition *inside* this factorization, on the left of R_s^{-1}, with a right cutoff equal to one on a neighborhood of the microsupport of the left cutoff. Thus each term has initial columns J_s on both sides, with the second column physically adjointed. The nested cutoffs preserve the complete matrix on summation by pseudolocality; the low-frequency term is jointly smooth. The transition regions of the outer spatial cutoffs lie beyond the causal shadow, so commutators supported there have no effect on the endpoint patches.

On a short Gaussian neighborhood the ordinary evolution of J_s agrees with the positive parametrix of Theorem 6.6 by its field and normal traces. At fixed metric Duistermaat and Hörmander (1972), Theorem 6.1.1 propagates the wavefront of a solution with smooth forcing along the null bicharacteristics: the real principal symbol $g^{-1}(\xi,\xi)$ has nonzero fiber derivative on its nonzero characteristic set. In passing to another spacelike slice the two short branch columns split the full Cauchy data, since the difference of their normal roots is elliptic. An opposite-branch coefficient there would propagate elliptically into the opposite cone of the overlap, where the incoming solution is smooth. Successive short evolutions therefore carry the positive column to the endpoint, including across changes of phase coordinates.

For the family argument the phase representations can be taken on a finite cover *before* selecting an endpoint cone. Subdivide the bounded t interval in $\overline{O}$ into short slabs. With spacelike t-sections, the two local phase constructions apply in flow boxes by using Gaussian coordinates and taking traces on the sections. Cover the compact part of each intermediate cosphere by smaller boxes inside larger ones. The nonzero normal roots stay separated and the short flow projections remain invertible for a common sufficiently short step, by compactness and smooth dependence. Compose the two-trace Cauchy parametrices across these sections with a partition on **all** intermediate directions; equivalently the paths needed at an endpoint are obtained by tracing its two null lifts back toward S. Finite propagation permits spatial localization at the successive sections in neighborhoods of the swept set. With cutoffs equal to one near each required domain of dependence, the short parametrices have smooth equation and initial errors there to every order. The differentiated Cauchy estimates of Lemma 6.8 carry these errors through the finite number of steps, also for smooth parameter derivatives. In each normalized short graph chart the opposite-branch symbol has every homogeneous coefficient zero at each parameter by the fixed-metric argument above, so all their parameter derivatives vanish and the uniform classical remainder estimates to arbitrary depth give joint smoothing after every finite mixed derivative. The properly supported graph compositions

are transverse on the spacelike sections, as required in Duistermaat (1972), §1, Theorem 1.1. Lemma 6.4 gives the corresponding parameter-dependent symbol compositions and remainder estimates on these buffered charts, including the terms subsequently discarded at the endpoints.

Mixed field traces. Fix smaller conic cutoffs at the two endpoints. A stationary point of a localized product must lift each endpoint to its positive null covector, follow the two Hamiltonian trajectories to S, and give the same point and the same tangential covector there in the primed convention. The last condition is the diagonal canonical relation of R_s^{-1}. On a patch about a matched pair, choose a disk $D \subset S$ and an initial cone containing both returned covectors, with a cutoff equal to one on a smaller neighborhood of their possible values. Shrink the endpoint cutoffs inside that neighborhood under the return maps. In the terms of the finite partition outside this initial cutoff the matching equations have no solution. Overlapping terms inside it sum to the identity microlocally.

To see the regularity of the discarded terms, localize to dyadic frequencies of size L. For comparable frequencies the joined phases, rescaled to unit frequency, have gradients bounded away from zero in the integration variables on the compact supports separated from the matching equations. For noncomparable adjacent frequencies, differentiation in the common spatial variable separates their covectors since the short graph maps preserve frequency up to bounded positive factors. One can equivalently subdivide into dyadic shells for each phase variable, integrating first at the largest unmatched scale. Repeated integration by parts in these variables gives arbitrary inverse powers of the largest frequency; bounded frequencies give smooth kernels. The separation persists on smaller parameter neighborhoods. Each specified endpoint or parameter derivative costs at most a fixed power of that frequency by the phase and amplitude estimates of Proposition 6.11. Choose the number of integrations above this loss, the symbol orders and the integration dimensions. The shells then sum absolutely with every prescribed derivative. This proves joint smoothing off the matched graph. On it, both nonzero endpoint covectors control the phase frequency, and parameter covectors are bounded by that frequency on the compact unit sections. These are the joint wavefront conditions of Definition 6.9.

Write E_{as} for the field trace on a of the ordinary evolution of J_s from D on this cone, including the carrier half-density factor ρ_a. The localized product is $C_{ab} = E_{as} R_s^{-1} E_{bs}^{\dagger}$ in the joint quotient. Taking the spacetime field in the second slot instead gives the kernel family used in C_{ap}, with the point retained as a variable. The graph operators for the carrier traces are elliptic by short-time transport and transverse restriction. On intersections of reference disks and ray patches all the products are traces of the same double evolution. Applying the construction on compact neighborhoods in W^2 gives the positive wavefront test everywhere; the transposition argument

of Theorem 6.10 identifies this class with every independently admitted kernel on W^2.

To compare two choices W, W' in the same collar, restrict to a connected globally hyperbolic open set $U \subset W \cap W'$ causally convex in both. One may use the intersection component containing the retained ray segment: causal convexity in the collar keeps each causal diamond between its points inside that component, and the diamond is compact by global hyperbolicity of W. On U the scalar Green operators from either neighborhood restrict to the unique Green operators of U by zero extension of test sources and causal compatibility (Bär, Ginoux and Pfäffle, Proposition 3.5.1 and Corollary 3.4.3). The two kernel classes thus have the same commutator on U^2 and the same positive wavefront constraint there. The transposition argument of Theorem 6.10 identifies them on all of U^2 modulo joint smoothing. For displaced carriers retain a common ray patch whose two hits and joining segments lie in U, with the margins of Definition 6.12; on this patch the mixed traces, diagonal charge roots and hence the normalized transports agree. For a family these comparisons use the same local identifications and buffers on the retained compact sets.

Lemma 6.13 (Normalized transport through a spacelike isotopy). Let F_s be a smooth local isotopy of spacelike embeddings of a fixed disk. Decompose its velocity as $N_s n_s + Z_s$, with n_s the future unit normal and Z_s tangential. The lapse N_s may vanish or change sign. For a scalar wave solution define its half-density traces $v_s = \rho_s F_s^* u$ and $\pi_s = \rho_s F_s^*(n_s u)$. Their equations are

$$\begin{aligned} \dot v_s &= L_s v_s + N_s \pi_s, \\ \dot\pi_s &= -D_s v_s - L_s^\dagger \pi_s, \\ L_s &= \mathcal{L}^{(1/2)}_{Z_s} + \frac{N_s \tau_s}{2}, \\ D_s &= -\rho_s \operatorname{div}_{h_s}(N_s \operatorname{grad}_{h_s})\rho_s^{-1}, \qquad \tau_s = \operatorname{tr}_{h_s} K_s. \end{aligned}$$

The normal positive-branch symbol Q_s gives $\pi_s = -Q_s v_s$ and hence

$$\dot v_s = -B_s v_s, \qquad B_s = N_s Q_s - L_s.$$

Charge normalization by $W_s = R_s^{1/2}$ gives an integer-order generator

$$B_s^{\mathrm{cov}} = W_s B_s W_s^{-1} - \dot W_s W_s^{-1}$$

whose complete formal symbol is real and which satisfies $(B_s^{\mathrm{cov}})^* = -B_s^{\mathrm{cov}}$. Two sufficiently small spacelike disks can be joined by such an isotopy on smaller common domains, including when they intersect.

Proof. Differentiation of the induced metric gives $\dot h_s = 2N_s K_s + \mathcal{L}_{Z_s} h_s$. Consequently $\dot\rho_s/\rho_s = N_s\tau_s/2 + \operatorname{div}_{h_s} Z_s/2$. Differentiating the scalar trace therefore gives the first equation, with the usual half-density Lie derivative. The wave divergence formula, or differentiation of the normal momentum

density, gives

$$\frac{\mathrm{d}}{\mathrm{d}s}(\rho_s^2 F_s^*(n_s u)) = \rho_s^2 \operatorname{div}_{h_s}(N_s \operatorname{grad}_{h_s} F_s^* u) + \mathcal{L}_{Z_s}^{(1)}(\rho_s^2 F_s^*(n_s u)).$$

Dividing by ρ_s and using its displayed derivative gives the second equation. In particular $D_s^T = D_s$, and the half-density Lie derivative of a real vector field is skew for both physical adjoints. These formulas involve multiplication by N_s and its derivatives; they do not divide by N_s.

For each s, construct Q_s in a separate normal Gaussian collar by the normal Riccati theorem. Its denominators are $2H_s$. The kernel built in that collar agrees with the canonical kernel by its independent characterization. Thus the same positive wave branch has graph $\pi_s = -Q_s v_s$ on every embedding. Substitution in the Hamiltonian system gives

$$Q_s N_s Q_s - \dot{Q}_s - Q_s L_s - L_s^\dagger Q_s + D_s = 0.$$

This is a compatibility identity for the already constructed branch. It is not a recursion divided by the deformation lapse. The construction therefore applies also where $N_s = 0$.

The Hamiltonian system preserves the Klein–Gordon form. Equivalently, subtract the displayed equation from its physical adjoint and use $R_s = i(Q_s^\dagger - Q_s)$ to obtain

$$\dot{R}_s = B_s^\dagger R_s + R_s B_s.$$

If E_{s,s_0} is the field-trace evolution, this gives $E_{s,s_0}^\dagger R_s E_{s,s_0} = R_{s_0}$. Thus $W_s E_{s,s_0} W_{s_0}^{-1}$ is physically unitary and its generator B_s^{cov} is physically anti-Hermitian.

For the internal adjoint, N_s, Q_s and L_s have real complete formal coefficients, so the same is true of B_s. The half-order calculation of the charge normalization gives, with $\mathcal{F}_s = Q_s - Q_s^*$, $W_s = e^{-i\pi/4}\mathcal{F}_s^{1/2}$ in rephased symbols. The recursive coefficients of $\mathcal{F}_s^{1/2}$ are real; the constant phase cancels in both terms of B_s^{cov}. Its integer-degree formal coefficients are therefore real. On those coefficients the physical Hermitian adjoint agrees with the internal adjoint, proving the asserted skew identity. No internal involution on half-integer orders is required.

Finally put the two disks into graph form over a common spacelike coordinate disk. Choose this local metric neighborhood sufficiently small relative to the retained endpoint spacelike margins that both graph conormals belong to the future timelike cone of one fixed reference Lorentz form. Their convex interpolations then have a positive timelike margin on the compact graph and interpolation parameters. Shrinking the common neighborhood further preserves that margin for the varying metric along the interpolated graphs. Interpolating the graph functions therefore gives the required spacelike isotopy. Its lapse can vanish where the graphs meet; the preceding proof already permits this. □

Theorem 6.14 (Complete carrier comparison and cocycle). On common retained ray domains define

$$T_{ab} = W_a C_{ab} W_b.$$

Then T_{ab} is an elliptic order-zero Fourier integral operator on the null canonical graph, and

$$C_{ab}C_b^{-1}C_{bc} = C_{ac},$$
$$T_{aa} = I, \qquad T_{ab}^\dagger = T_{ba}, \qquad T_{ab}T_{bc} = T_{ac}.$$

Conjugation $\mathcal{A}_{ab}(P) = T_{ab}PT_{ba}$ is a filtered isomorphism of the complete carrier pseudodifferential algebras and preserves their internal adjoints. The map $\mathcal{M}_{ab}(m) = T_{ab}m$ is an invertible semilinear map between the complete modules of the same retained sky:

$$\mathcal{M}_{ab}(Pm) = \mathcal{A}_{ab}(P)\mathcal{M}_{ab}(m), \qquad \mathcal{A}_{ab}(P^*) = \mathcal{A}_{ab}(P)^*.$$

These maps commute with restriction, have exact identity, inverse and triple-composition laws in the quotient, and are natural under supplied physical isometries and their induced carrier maps.

Proof. On each of the ray patches constructed before Lemma 6.13, take a common reference disk in the Cauchy surface for all carriers in this composition. The double Cauchy evolution proved there gives

$$C_{ab} = E_{as}R_s^{-1}E_{bs}^\dagger.$$

The trace maps are invertible graph operators in the quotient. Inserting the complete inverse

$$C_b^{-1} = (E_{bs}^\dagger)^{-1}R_sE_{bs}^{-1}$$

between C_{ab} and C_{bc} cancels all four adjacent inverse factors and proves the first identity. In particular, if $E_{ab} = E_{as}E_{bs}^{-1}$, then

$$T_{ab} = W_aE_{ab}W_b^{-1}.$$

This is order zero and elliptic. Hermitian symmetry of the factorization gives $C_{ab}^\dagger = C_{ba}$. Since the W_a are selfadjoint, it gives $T_{ab}^\dagger = T_{ba}$. The identities $W_b^2 = C_b^{-1}$ and $W_aC_aW_a = I$, together with the mixed identity already proved, give the identity and cocycle formulas. Consequently T_{ba} is the exact quotient inverse. No undetermined scalar occurs in these cancellations.

The canonical-graph composition theorem makes conjugation by T_{ab} a filtered pseudodifferential algebra isomorphism. First consider a pair of carrier patches joined by a spacelike isotopy as in Lemma 6.13, with interior transverse hits throughout. The normalized trace formula above identifies the transport from its initial patch to each moving patch with the normalized evolution of that lemma. In fixed coordinates on its parameter disk, the normalized evolution has the generator B_s^{cov}. Conjugation obeys

$$\partial_s\mathcal{A}_s(P) = [-B_s^{\mathrm{cov}}, \mathcal{A}_s(P)].$$

Since $(B_s^{\mathrm{cov}})^* = -B_s^{\mathrm{cov}}$, applying the internal anti-involution to this equation gives the same equation for $\mathcal{A}_s(P)^*$ as for $\mathcal{A}_s(P^*)$. Their initial values agree. The complete transport recursion, degree by degree along the retained flow, has a unique initialized solution, so the two expressions agree at every degree.

For separated hits, fix a ray of Definition 6.12 and follow its compact segment from the hit on b to the hit on a inside the retained parent. Subdivide the segment into finitely many pieces, each in a Gaussian coordinate cylinder, with each subdivision point interior to both adjacent cylinders. Within each piece use the Gaussian time slices through its two ends. In the space coordinates on these slices take disks centered along the segment, transporting the disk coordinates by translation of the center. The time slices are spacelike and Gaussian time is strictly monotone along the ray. By compact containment one can choose a positive disk radius and a buffer throughout that piece. These moving disks form a spacelike isotopy with the given ray meeting their centers. At each subdivision point the incoming and outgoing disks pass through the same point. Join their restrictions by the local graph interpolation of Lemma 6.13 in the overlap. Use the same construction at the two endpoint hits to join the given carriers to the first and last Gaussian disks. The interpolation at a shared hit keeps that point fixed. Intersecting endpoint carriers with coincident hits require only this last local comparison.

On the unit section about the chosen ray, shrink once for this finite list of isotopies. Spacelikeness and transverse intersection give a positive normal-pairing margin, and the central hits have positive distance from the disk boundaries. Compactness of each closed parameter interval and continuous dependence of the hits preserve these margins on the smaller ray patch. Label the successive patches $D_0 = b, \dots, D_m = a$, with restrictions understood, and use $T_{D_j D_{j-1}}$ from D_{j-1} to D_j. The kernel cocycle proved above gives

$$T_{ab} = T_{D_m D_{m-1}} \cdots T_{D_1 D_0}.$$

Conjugate first by the rightmost factor and then by the successive factors. Each conjugation preserves the internal adjoint by the initialized transport equation, hence so does their composition. On intersections of ray patches each product equals the same mixed-kernel operator. A finite patch cover of the retained compact unit section therefore gives the identity throughout the common domain.

The graph operator maps the Lagrangian distribution class of a sky on Σ_b to that same sky on Σ_a. The action includes its complete amplitude, as prescribed by the shifted module convention. Properly supported graph composition sends jointly smooth remainders to jointly smooth remainders. The semilinear identity is now ordinary operator composition. Invertibility and the operator cocycle prove the module inverse and cocycle laws as well.

A supplied physical isometry preserves the scalar kernel by its independent characterization. Its induced carrier half-density pushforwards therefore

carry C_{ab} to the corresponding mixed kernel. Uniqueness of the positive-principal complete square root carries W_a, and the defining product carries T_{ab}. This proves naturality. All constructions and cancellations are valid after restriction to common buffered ray domains, proving the remaining laws. □

The scalar kernel and these scalar module maps do not detect the central spin sign. The actual spin maps of the geometric source remain separately retained comparisons.

6.7. Evaluation on the whole moving sky.

Definition 6.15 (The full wave evaluation section).
The point slot remains scalar. For $p \in B$ set

$$C_{ap}(x) = \rho_a(x)\Lambda_g^+(F_a(x), p), \qquad e_a(p) = W_a C_{ap}.$$

Only the carrier variable has a half-density factor. These are actual distribution families modulo kernels jointly smooth in the point and carrier variables, and in any supplied passive parameters. Let $L_{a,p}$ be the homogeneous Lagrangian obtained by following the chosen-branch null covectors at p to their carrier hits and restricting the transported covectors to $T\Sigma_a$. The section $e_a(p)$ belongs to its complete Lagrangian module on $T^*\Sigma_a \setminus 0$.

This definition is valid when p lies on the carrier. The pullback with fixed p has no wavefront obstruction because the first, carrier covector cannot vanish. The same fact gives smooth dependence on p as a distribution-valued family. A finite phase cover describes the moving Lagrangians uniformly. If $p = F_a(b)$, then $L_{a,p} = T_b^*\Sigma_a \setminus 0$; a spatial light-cone cut need not remain a smooth hypersurface there. In a short Gaussian phase chart the relevant phase is $S(t, x, \xi) - b \cdot \xi$, with $S(0, x, \xi) = x \cdot \xi$. At $t = 0$ its critical equations are $x - b = 0$ and have full differential rank in x. Thus it remains a nondegenerate three-frequency phase through the carrier. Propagating and composing such phase charts gives the stated moving Lagrangian family.

A contact graph $z = s(p, q)$ can be obtained near any fixed point p_0 by first transporting to a spacelike disk through p_0 with the already constructed T_{ab}, and then using the normalized Radon charts of the shifted module convention on its point fiber. The graph projection remains invertible for nearby points after restriction. This supplies the graph form through actual carrier transport and the specified Radon normalization, with their established algebra and internal-adjoint laws.

Proposition 6.16 (Full evaluation and carrier covariance).
For every retained point,

$$T_{ab}e_b(p) = e_a(p).$$

The comparing operator is independent of p. If $p = (t, b)$ in a Gaussian collar, and $U_{s,t} = W_s E_{s,t} W_t^{-1}$ denotes normalized trace transport, then

$$e_s(t, b) = U_{s,t} W_t^{-1}(\rho_t(b)^{-1}\delta_b).$$

Here δ_b denotes the coordinate delta half-density. At $p = F_a(b)$, the physical Fourier amplitude of $e_a(p)$, using the unitary Fourier normalization of the shifted module, has principal coefficient

$$(2\pi)^{-3/2}(2H_a(b,\xi))^{-1/2}\rho_a(b)^{-1}.$$

Its formal degree-one coefficient is i^{-1} times this expression. The full section is elliptic and locally cyclic, of module filtration degree one. These are equalities and properties of complete classes; all lower coefficients are determined by the canonical kernel prescription.

Proof. The common factorization also holds with its last variable fixed at p. The same inverse cancellations used for the carrier Gram identity give $C_{ab}C_b^{-1}C_{bp} = C_{ap}$. Since $W_b^2 = C_b^{-1}$, multiplication by W_a proves the evaluation identity at every point. The definition of T_{ab} involves only the source metric and the two carrier embeddings, so it is independent of the evaluated point.

On Σ_t, the point column is $C_t\delta_b/\rho_t(b)$: multiplying a scalar kernel by two slice half-densities and then fixing one variable must remove exactly the half-density factor in that fixed variable. Hence on Σ_s it is $E_{s,t}C_t\delta_b/\rho_t(b)$. Multiplication by W_s, followed by $C_t = W_t^{-2}$, gives the displayed normal form.

The principal symbol of W_t^{-1} is $(2H_t)^{-1/2}$. Its operator kernel uses the prefactor $(2\pi)^{-3}$, whereas the conormal section uses $(2\pi)^{-3/2}$. Thus its scalar conormal amplitude has the extra fixed factor $(2\pi)^{-3/2}$ stated above. The scalar amplitude order is $-1/2$, which is $j - 3/2$ for $j = 1$ in the shifted filtration. Dividing by the degree rephasing i^1 gives the formal coefficient. This amplitude is nonzero on every retained direction.

The lower coefficients arise from the full operator action, with its finite Taylor and integration-by-parts terms; no replacement by the principal amplitude has been made. An elliptic conormal amplitude is locally a cyclic generator by division and successive correction in the complete symbols. Transport by the invertible graph operators preserves that property and the filtration degree. This proves ellipticity and cyclicity on the entire retained moving sky, using its finite angular cover. The exact kernel and Gram identities prove the full-class assertions. □

6.8. Finite families and physical scale.

Proposition 6.17 (Carrier families, finite estimates and scale). Let source collars and retained carrier embeddings vary in a smooth finite-dimensional family, with the common parents and positive margins of the geometric family theorem. Include uniform spacelike, ray-intersection, phase-inverse and boundary-buffer margins. The kernels C_{ab}, roots W_a, comparisons T_{ab} and complete evaluations e_a have smooth complete-coefficient families on those parents.

For every finite coefficient depth J, base and angular derivative order m, and passive-parameter derivative order q, there are a finite input order N

and a finite locally bounded function $F_{J,m,q}$ such that

$$\|\mathcal{Y}\|_{J,m,q} \le F_{J,m,q}\left(\mathcal{N}_N(g, F_a; \mathcal{K}' \times \text{larger parents}), \mathfrak{m}^{-1}\right).$$

Here $\mathcal{Y}$ is any specified finite list of these coefficients and transport amplitudes, the output is measured over $\mathcal{K} \Subset \operatorname{int} \mathcal{K}'$, and $\mathfrak{m}$ denotes the retained positive margins. The parent is selected before J, m, q. The same statement holds after natural comparison by a supplied smooth family of physical maps, with its finite map and inverse bounds included.

In four spacetime dimensions, constant physical scaling has the exact complete-class laws

$$\Lambda^+_{r^2 g} = r^{-2}\Lambda^+_g, \qquad C^{(r)}_{ab} = rC_{ab}, \quad W^{(r)}_a = r^{-1/2}W_a,$$
$$C^{(r)}_{ap} = r^{-1/2}C_{ap}, \quad T^{(r)}_{ab} = T_{ab}, \qquad e^{(r)}_a = r^{-1}e_a.$$

The numerical scale is retained. It may vary smoothly with passive parameters, with its finite derivatives and positive lower bound included in the estimates, but is constant on each spacetime member.

Proof. The kernel family's finite-coefficient construction is provided by the kernel family estimate. Restriction to a spacelike carrier uses the retained transverse phase charts and the finite derivatives of its embedding. On a unit angular section, the metric bounds give the required upper and lower bounds for H_a. The elliptic inverse and positive square-root recursions divide only by their nonzero leading coefficients. Thus each fixed coefficient of W_a and each of its prescribed derivatives uses finitely many previously determined coefficients and derivatives.

The ray maps solve the actual null Hamilton equations, and the transverse hit times follow from the implicit function theorem with the stated intersection margin. Their finite variations satisfy finite ODE systems and inverse-function identities. Phase coefficients consequently have finite bounds from the source metric, embeddings and retained margins. The amplitudes satisfy linear transport recursions on this same ray tube; differentiating finitely many of their equations gives finite Gronwall bounds. Stationary phase in a retained nondegenerate composition uses finitely many derivatives at every output degree. These facts, together with the operator and evaluation products above, prove the displayed estimate. No successive reduction of the parent is needed when the coefficient depth increases.

Uniform asymptotic summation on compact parameter sets supplies actual classical representatives of the coefficient families. The cutoffs for a full sum may depend on the entire sequence of symbol bounds. Lemma 6.4 gives the quantitative estimates for the prescribed finite list of carrier coefficients. Differences of two such sums are uniformly rapidly decreasing with every retained derivative. Integration by parts, or direct integration after choosing a sufficiently large decay order, shows that their kernels are jointly smooth in all variables and passive parameters. Properly supported composition and restriction with the stated transverse margins preserve this joint-smoothing

quotient. Supplied physical comparison maps act by the same finite pullback, inverse and density operations, proving the natural family assertion.

For scale, $\Box_{r^2 g} = r^{-2}\Box_g$ and the physical four-volume is multiplied by r^4. The scalar causal kernel normalized against that measure is therefore multiplied by r^{-2}. The independent canonical-kernel characterization gives the same law for Λ^+. A carrier half-density has weight $r^{3/2}$, so two carrier factors give $C^{(r)}_{ab} = rC_{ab}$, while one carrier factor and the scalar point slot give $C^{(r)}_{ap} = r^{-1/2}C_{ap}$. Uniqueness of the positive-principal inverse square root gives the W_a law. Substitution into T_{ab} and e_a proves their two weights. The chain rule gives the finite passive-parameter estimates when r varies in such a family. □

The wave evaluation is a full section with its canonical kernel normalization. Its point derivatives and covariant wave equation are defined next on the actual moving distribution families. The source's finite geometric and actual spin structures retain the separate admission and comparison laws already proved.

6.9. Moving modules and actual point differentiation. Fix a retained carrier disk Σ_a and a buffered spacetime point neighborhood B as in the carrier construction. The carrier is fixed while the evaluated point p varies. Let $L_{a,p} \subset T^*\Sigma_a \setminus 0$ be the conic Lagrangian determined by the full sky of p. The point slot is scalar; the distribution on Σ_a carries its spatial half-density. All geometric coefficients are read from the admitted finite Cartan collar.

The family of Lagrangians is smooth even when $p = F_a(b)$ lies on the carrier. At such a point it is $T^*_b\Sigma_a \setminus 0$. The nondegenerate phase-family construction in the evaluation construction covers this case. A graph phase $z = s(p, q)$ below refers to a fixed angular contact presentation in which the Legendrian is a graph; it does not require a smooth hypersurface cut in the physical carrier coordinates. In these presentations $\tau > 0$ denotes the radial covariable dual to z.

Definition 6.18 (The moving complete module).
Let $\mathscr{M}_a$ be the sheaf on B of smooth families of classical Lagrangian distributions valued in spatial half-densities and associated with $L_{a,p}$, in the shifted filtration of the sky-module construction, modulo jointly smooth families. Locally these families have actual oscillatory representatives

$$m(p, x) = \int e^{i\phi(p,x,\vartheta)} a(p, x, \vartheta)\, \mathrm{d}\vartheta |\mathrm{d}x|^{1/2},$$

where ϕ is a nondegenerate homogeneous phase smoothly parametrizing $L_{a,p}$, and a is classical with a locally bounded upper order and symbol estimates uniform with every prescribed finite point derivative. Different phase presentations are identified by the actual distribution calculus. The quotient is by $C^\infty(B \times \Sigma_a; \Omega^{1/2}_{\Sigma_a})$ on the retained microlocal patches. With an additional finite-dimensional parameter $\lambda \in D$, the discarded remainder must be jointly smooth on $D \times B \times \Sigma_a$.

Write $F_m\mathscr{M}_a$ for the filtration. The complete carrier algebra $\widehat{\mathscr{E}}_a$ acts on this sheaf. A smooth operator family is *basic in the point variable* if its actual representative is independent of p. The point derivative of any smooth operator family is taken on this same fixed carrier. The actual distribution realization and its normalization are part of the definition.

The uniform symbol-family hypotheses provide actual representatives by the finite symbol-family calculus. They also ensure that differences with zero complete expansion have jointly smooth kernels after all prescribed finite parameter derivatives. Choosing unrelated fiberwise smooth remainders would not define this quotient.

Theorem 6.19 (The full ambient point connection).
Actual differentiation in p induces a connection ∇^{amb} on the union of the filtered sheaves $\mathscr{M}_a$. For a smooth point vector field X, a smooth point scalar f, a smooth complete carrier-operator family P and a module family m, it satisfies

$$\begin{aligned}
\nabla^{\mathrm{amb}}_X(fm) &= X(f)m + f\nabla^{\mathrm{amb}}_X m,\\
\nabla^{\mathrm{amb}}_X(Pm) &= (XP)m + P\nabla^{\mathrm{amb}}_X m,\\
\nabla^{\mathrm{amb}}_X F_\ell\mathscr{M}_a &\subset F_{\ell+1}\mathscr{M}_a,\\
[\nabla^{\mathrm{amb}}_X, \nabla^{\mathrm{amb}}_Y] &= \nabla^{\mathrm{amb}}_{[X,Y]}.
\end{aligned}$$

The corresponding algebra connection is $\nabla^{\mathrm{alg}}_X P = XP$. In a fixed angular contact chart with sky graph $z = s(p,q)$, write a complete section as au_s in the rephased convention. Then

$$\nabla^{\mathrm{amb}}_X(au_s) = \big(Xa - \tau(Xs)a\big)u_s.$$

These are identities on every complete module section. The connection is determined by the actual realization. Equivalently, its algebra Leibniz rule together with the actual derivative of the full elliptic evaluation family e_a determines it uniquely.

Proof. Differentiate an actual oscillatory representative. Its derivative has the same phase and amplitude $Xa + i(X\phi)a$. Homogeneity of ϕ raises the amplitude order by at most one. Nondegeneracy and the parameter-uniform symbol conditions remain valid on the same retained phase patch. A derivative of a jointly smooth remainder is jointly smooth. Actual differentiation is independent of the phase chosen to represent the distribution, so it induces a well-defined sheaf operator with the asserted filtration loss. Its linearity in X and scalar Leibniz identity follow before taking the quotient.

Properly supported carrier operators act on these distribution families by the actual parameter-dependent calculus. Differentiating their action gives $X(Pm) = (XP)m + P(Xm)$. Every term respects the joint smoothing quotient. The equality therefore descends to the complete algebras and modules and to their angular presentations. The algebra derivative obeys its product rule and commutes with the internal adjoint, because that adjoint is defined by finite coefficient operations independent of p.

In the graph presentation the physical Fourier amplitude is

$$e^{-i\tau s(p,q)} \sum_{j\le m} i^j \tau^{j-1/2} a_j(p,q).$$

Differentiating its phase produces $-i^{j+1}\tau^{j+1-1/2}(Xs)a_j$. After rephasing degree $j+1$, this is the second term in the graph formula. In particular the complete coefficient at degree j becomes

$$a_j \longmapsto Xa_j - (Xs)a_{j-1},$$

with coefficients above the upper order set to zero.

The exact module actions in this chart are

$$D_z = \tau, \qquad D_{q^j} = \partial_{q^j} - s_{q^j}\tau, \qquad z = -\partial_\tau + \frac{1}{2\tau} + s.$$

They give a direct check of the operator identity. For D_{q^j}, the point derivative of $-s_{q^j}\tau$ is canceled by the commutator of $-\tau Xs$ with ∂_{q^j}. For z, the point derivative Xs is canceled by $[-\tau Xs, -\partial_\tau] = -Xs$. The $1/(2\tau)$ term is retained throughout, and D_z commutes with the derivative. These checks use full module formulas, rather than their principal symbols.

Ordinary derivatives of actual distribution families satisfy $[X,Y]m = X(Ym) - Y(Xm)$. Passing to the joint quotient proves flatness. In graph coefficients the possible curvature multiplier is

$$-\tau\big(X(Ys) - Y(Xs) - [X,Y]s\big) = 0.$$

For the uniqueness statement let ∇' obey the same algebra Leibniz rule and have the same derivative on the full evaluation family. The difference $L_X = \nabla'_X - \nabla^{\mathrm{amb}}_X$ is linear over point functions and commutes with every smooth carrier-operator family. The evaluation amplitude has a nonzero leading coefficient of degree one by the full evaluation comparison. In a graph chart write $e_a = a_e u_s$ and $m = a_m u_s$. Complete Laurent division gives a smooth coefficient series $b = a_m/a_e$. Each coefficient is obtained by finitely many divisions by the nonvanishing leading coefficient and finite operations on the preceding coefficients. The normal-ordered operator $P = b(q, D_z; p)$ acts on $a_e u_s$ by multiplication by $b(q,\tau;p)$, so $m = Pe_a$ in that chart. A classical representative of P exists by asymptotic summation. Therefore $L_X m = L_X(Pe_a) = PL_X e_a = 0$. The local conclusion agrees on overlaps, proving uniqueness. □

Point differentiation raises the order by one, and flatness is the commutator identity on the resulting filtered union. Parallel transport of arbitrary module elements between distinct skies would require an additional construction. The finite Cartan and spin connections retain their respective curvatures. Realization by distributions fixes the ambiguity left by algebra compatibility and flatness, namely addition of a closed point one-form times the identity.

6.10. Carrier and stored-chart comparisons.

Proposition 6.20 (Covariance in all point directions).
The normalized carrier comparison of the carrier cocycle theorem obeys

$$\nabla_X^{\mathrm{amb},a}(T_{ab}m) = T_{ab}\nabla_X^{\mathrm{amb},b}m$$

on every smooth complete module family. It also intertwines the algebra point connections under $P \mapsto T_{ab}PT_{ba}$.

If a specified physical isometry $f : B_0 \to B_1$ compares two recognized collars, compose half-density pushforward to the image carrier with a normalized target-carrier comparison. Its induced map $\widehat{F}$ satisfies

$$\nabla_{f_*X}^{\mathrm{amb},1}\widehat{F} = \widehat{F}\nabla_X^{\mathrm{amb},0}.$$

These constructions have exact restriction, identity, inverse and composition laws on common retained parents. In point-dependent stored coordinates $m' = F_p m$, a connection written as $X + A_X$ transforms by

$$A'_X = F_p A_X F_p^{-1} - (XF_p)F_p^{-1}.$$

Proof. The kernel defining T_{ab} uses the two fixed carrier embeddings and the source collar. It is independent of the evaluated point p. Differentiation of its properly supported oscillatory realization therefore commutes with its action on a moving module family. Jointly smooth errors remain jointly smooth, which proves the first identity. Differentiating $T_{ab}PT_{ba}$ proves the algebra identity for the same reason.

Under f, naturality of the canonical scalar kernel and the actual carrier half-density pushforward give the full carrier comparison already established in the evaluation comparison. The fixed operator depends on f and the carrier embeddings, not on a separately variable evaluated point. Applying the ordinary chain rule in that point proves the displayed physical comparison. In particular, when f_*X has both normal and tangential components, all those components occur in $\nabla_{f_*X}^{\mathrm{amb}}$. The connection has been defined in every direction.

The carrier cocycle gives the identity, inverse and composition laws; the actual constructions commute with restriction. Finally, if $m' = F_p m$, write the transported derivative as $F_p(X + A_X)(F_p^{-1}m')$. The derivative of $F_pF_p^{-1} = I$ gives $F_p(XF_p^{-1}) = -(XF_p)F_p^{-1}$, and expansion yields the stated stored-chart formula. This also gives the full derivative term in compositions of varying stored presentations. □

The physical isometry in this proposition is supplied. Recognizing every physical map from the complete source-side arrow equations is a separate full-Hom assertion. Here the proposition constructs and verifies the point-connection part of the induced comparison.

6.11. The differential law and finite estimates.

Proposition 6.21 (The intrinsic wave equation of the full evaluation family).

The normalized evaluation family $e_a = W_a C_{ap}$ satisfies

$$\Box_g^{\nabla^{\mathrm{amb}}} e_a = 0, \qquad \Box_g^{\nabla^{\mathrm{amb}}} m = |\det g|^{-1/2} \nabla^{\mathrm{amb}}_{\partial_\mu} \left(|\det g|^{1/2} g^{\mu\nu} \nabla^{\mathrm{amb}}_{\partial_\nu} m \right)$$

in the joint complete-module quotient. The equation is coordinate independent and is preserved by the physical comparisons above.

For a positive constant physical scale $g^{[r]} = r^2 g$, the fixed-coordinate ambient connection is unchanged, and

$$e_a^{[r]} = r^{-1} e_a, \qquad \Box_{r^2 g}^{\nabla^{\mathrm{amb}}} e_a^{[r]} = r^{-3} \Box_g^{\nabla^{\mathrm{amb}}} e_a.$$

Proof. The carrier restriction of the canonical kernel has no wavefront covector with zero carrier component: a nonzero null covector cannot annihilate the tangent space of a spacelike carrier. Its local phase constructions therefore give the smooth distribution family C_{ap} used in the definition. The bisolution test gives $\Box_{g,p} C_{ap} = 0$ modulo a jointly smooth family. Here only the carrier factor $\rho_a(x)$ has been inserted; there is no density in the scalar point slot.

The normalizing operator W_a acts on $x \in \Sigma_a$ and is independent of p. It commutes both with point differentiation and with the coefficients $g^{\mu\nu}(p)$ and $|\det g(p)|^{1/2}$ in the displayed expression. Proper support makes its action preserve joint smoothness. Thus

$$\Box_{g,p}(W_a C_{ap}) = W_a \Box_{g,p} C_{ap} = 0.$$

The ambient-connection theorem identifies this actual derivative calculation with the stated module equation. The scalar divergence expression is invariant under point coordinate changes; the proposition on all-direction covariance and preservation of the physical metric then give its covariance.

For scaling, the carrier scale law gives $T_{ab}^{[r]} = T_{ab}$ and $e_a^{[r]} = r^{-1} e_a$. Constant scaling preserves the conic null relation and the underlying moving distribution realizations, so their actual point derivative is unchanged. Since $\Box_{r^2 g} = r^{-2} \Box_g$ and r is constant in the point variable, the second identity follows. The physical-clock pullback of Proposition 2.11 may subsequently be applied by the already proved coordinate covariance. □

The wave equation is a differential law of the fully normalized evaluation family. Its complete initial value is supplied by the canonical kernel and carrier-root prescription. It is not a horizontality condition. Its prolonged identities retain derivatives of the metric coefficients. For example, write $b^\nu = |\det g|^{-1/2} \partial_\mu (|\det g|^{1/2} g^{\mu\nu})$. Flatness in a point coordinate chart and differentiation of the wave identity give

$$\Box_g^{\nabla^{\mathrm{amb}}} \nabla^{\mathrm{amb}}_{\partial_\kappa} e_a = -(\partial_\kappa g^{\mu\nu}) \nabla^{\mathrm{amb}}_{\partial_\mu} \nabla^{\mathrm{amb}}_{\partial_\nu} e_a - (\partial_\kappa b^\nu) \nabla^{\mathrm{amb}}_{\partial_\nu} e_a.$$

The full comparison carries these identities and every finite point jet because it intertwines the connection on all sections.

Proposition 6.22 (Finite estimates on one point-family parent). Take a smooth finite-dimensional family of admitted collars and carrier embeddings on the common parents and with the positive margins of the carrier-family result. Fix a finite angular cover, a homogeneous depth J, a number k of iterated coordinate point derivatives, a coefficient derivative order m and an external parameter derivative order q. On retained compact subsets, the corresponding evaluation and comparison coefficients satisfy

$$\|\nabla^{\mathrm{amb},k} e_a\|_{m,q,J} \leq F_{m,q,J,k}\left(\mathcal{N}_N(\text{source coefficients and carrier embeddings}), \mathfrak{m}^{-1}\right)$$

for some finite N and finite locally bounded function $F_{m,q,J,k}$. The filtration origin on the left is the upper degree $1+k$. The same kind of estimate holds for any stipulated finite list of the connection formulas and their normalized comparisons, including derivatives of supplied stored-coordinate maps with their inverse margins. The common parent is chosen before m, q, J, k.

Proof. In a graph chart one derivative acts on coefficients by the exact rule in the coefficient formula. Iteration uses only finitely many coefficients of e_a, finitely many derivatives of the same phase graph and the coefficients of the selected vector fields. The top possible degree increases by one at each step. Finite product estimates therefore bound every requested coefficient and mixed derivative on the retained compact chart.

The carrier-family result bounds the phase, evaluation and transition coefficients by finite source and embedding seminorms with fixed ellipticity, spacelike and domain-buffer margins. Finite-depth stationary phase for an angular change uses finitely many earlier coefficients and finite derivatives. The finite symbol-family calculus consequently gives the same estimate on the finite angular cover. This argument applies to the complete conormal phase families through points on the carrier as well; a physical spatial cut is not needed.

Stored-coordinate formulas use composition, inversion on the supplied nondegenerate domains and the explicit derivative term in the preceding proposition. Their finite-depth estimates therefore use only the stated finite extra seminorms and inverse margins. Jointly smooth quotient equalities may be differentiated with all retained parameters, so comparison identities do not require a new representative branch or a smaller parent at every order.

A positive scale may depend on an external parameter. Its finite derivatives and a positive lower bound are included in the input seminorms and margins, and the chain rule gives the same conclusion. Here we differentiate the supplied family with respect to its external parameters. A canonical connection is defined only in the point variables in this construction. For the point connection, Lemma 6.4 gives the claimed bounds on the coefficients at each prescribed filtration depth. □

6.12. Independent scalar wave enrichment of a Cartan collar. The following admission isolates the scalar wave part over a geometric carrier

whose independent recognition was proved in Section 2. It is used before the final category with all geometric and algebroid comparison data is assembled.

Definition 6.23 (Independent wave enrichment on a retained carrier).

Supply an independently admitted finite Cartan collar of Definition 2.7, retained spacelike carrier disks and buffered ray and point domains in a common wave neighborhood W as in Definition 6.12. Form $g = \eta(\theta, \theta)$ from its supplied soldering. A normalized scalar wave enrichment consists of the following supplied classes, realizations and tests.

Supply a scalar kernel class Λ^+ satisfying the independent kernel tests: in the joint smoothing quotient,

$$\Box_{g,p}\Lambda^+ = \Box_{g,q}\Lambda^+ = 0, \qquad \Lambda^+ - (\Lambda^+)^T = -iE_g, \qquad \mathrm{WF}'(\Lambda^+) \subset \mathcal{C}_g^+.$$

Here E_g is the retarded-minus-advanced scalar kernel on W of the differential operator computed from θ, normalized by its physical volume. The chosen positive null relation and its joint-family condition are the ones in that definition. They are read from the source null flow. No additional metric or wave state is supplied.

Supply complete carrier algebra and whole-sky module presentations with specified invertible normalizations to the half-density algebra and moving distribution modules above. Test these maps and their inverses on the complete products, module actions, filtrations, internal adjoints and homogeneous characteristic maps; test their overlap and central-specialization identities in the stated quotient. These are specified normalizations to the known cotangent and distribution models. They are not an existential membership condition for an ordinary-pair constructor.

Use the actual carrier half-densities to form C_{ab} and C_{ap} from Λ^+. Require the supplied C_a^{-1} and W_a to satisfy their complete inverse and square-root identities, with the positive physical principal symbol for $W_a = C_a^{-1/2}$. Require the complete normalized carrier maps to be $T_{ab} = W_a C_{ab} W_b$. In each supplied normalization the full point section and connections must obey

$$e_a = W_a C_{ap}, \qquad \nabla_X^{\mathrm{alg}} P = XP, \qquad \nabla_X^{\mathrm{amb}}(au_s) = (Xa - \tau(Xs)a)u_s$$

for every point vector field and every complete module coefficient family. The graph test is imposed in fixed angular actual realizations; equivalent general phase realizations use actual point differentiation. Point-dependent stored normalizations carry the full derivative term of the stored-chart formula.

All statements with external parameters use the uniform classical family and joint smoothing conditions already specified. Smoothing representatives are not retained as additional data. The actual spin bundle, its initial sign and the Cartan fields remain the supplied geometric part of the object. The scalar normalization does not identify their distinct spin lifts.

Theorem 6.24 (Construction and recovery of the wave enrichment).

Every admitted finite Cartan collar has the wave enrichment of the preceding definition on sufficiently small retained parents. The kernel, complete carrier maps, full point section and all-direction point connection are determined by the supplied geometric collar and the specified actual normalizations.

Conversely, every independently admitted enrichment agrees, in its supplied normalizations, with this construction from the same collar. Constructing and then reading these normalized fields, or reading and then constructing them, gives the identity on their complete classes. The identifications commute with restriction, normalized carrier changes and supplied physical isometries, and preserve smooth finite-dimensional families with the finite estimates above.

In particular an ordinary-pair display followed by Corollary 2.12 has this derived enrichment. Its wave fields require no further generating class or variable Einstein input.

Proof. The finite Cartan recognition theorem gives an actual Lorentzian Ricci-flat metric from the supplied fields, with the retained physical scale and actual spin structure. Choose buffered local domains and carriers as in the carrier construction. The canonical-kernel theorem constructs a class satisfying the independent bisolution, commutator and characteristic tests. For displaced hits the double Cauchy construction preceding Lemma 6.13 extends this class to the common wave neighborhood. Its actual representative families have the required joint regularity.

The complete inverse and positive-principal square-root recursions give C_a^{-1} and W_a uniquely. The carrier cocycle theorem and the full evaluation comparison give T_{ab} and e_a, including their complete coefficients and all retained angular directions. The ambient-connection theorem gives the required operator on every smooth module family. Thus every admission condition is satisfied, with the actual distribution normalizations supplied by the construction.

Now start with a supplied admitted enrichment. Kernel uniqueness identifies its Λ^+ with the constructed class, including joint parameters. The trace-density formulas then identify every C_{ab} and C_{ap}. Uniqueness of the complete inverse and positive-principal root identifies C_a^{-1} and W_a, and hence the supplied T_{ab} and full e_a. The connection test on all module families identifies the supplied connection with actual point differentiation; the algebra test identifies the algebra connection. The supplied normalizations and their inverse equations carry these equalities back to each stored presentation. These arguments prove both normalized identity laws without assuming representability in the admission definition.

Restriction commutes with every construction. Carrier changes obey the exact cocycles and all-direction connection identities already proved. A supplied physical isometry carries the kernel by its independent uniqueness and carries the remaining structures by the exact density and mixed-kernel

formulas, followed by the point chain rule. The finite-family assertions follow by composition of the common geometric-collar estimates, the kernel and carrier finite-coefficient estimates, and the preceding point-connection estimates.

Finally, the ordinary construction supplies its finite Cartan collar from the executed initial source. The metric used in this proof is its readout $g = \eta(\theta, \theta)$. The kernel and every subsequent field are determined from that readout and its carrier markings. This proves the last assertion. □

This theorem concerns the normalized scalar enrichment over an admitted geometric collar. Recognizing all abstract source arrows, carrying the complete algebroid associators, and obtaining effective descent for the final unmarked atlas require their full comparison results. The analytic coupled Real operation also retains its separate structured comparison; it is not inferred from the scalar wave construction.

7. Geometric admission and full source comparisons

The scalar enrichment of Section 6 is covariant under a supplied physical isometry. For independently supplied sources, the labelled skies and the physical density recover the underlying physical map from a source comparison. We then use the full evaluation family over an open neighborhood of the point to show that the physical map and its specified spin lift determine the scalar and finite-rank comparison.

All geometric representatives are smooth, Hausdorff and second countable. Point cores are connected and have buffered parents; null rays mean their retained connected plaques. When the point core is restricted, these plaques and their local quotients are taken again. Scalar equalities use the complete calculus and joint smoothing quotient of Section 6. Actual spin maps and geometric connections are not taken modulo that quotient.

7.1. Independent real incidence over the Cartan collar. Low studied the geometry of the space of null geodesics, and Bautista, Ibort and Lafuente gave a reconstruction theorem using light rays and skies. Here the point manifold and the comparison map f on points are supplied. On the retained ray plaques the null cones determine the conformal class. If $f^* g_1 = c g_0$ with $c > 0$, the point density scales by c^2 in four dimensions. Its preservation gives $c^2 = 1$ and hence $c = 1$.

Definition 7.1 (Geometric incidence tests).
Let an admitted finite Cartan collar have point manifold B, soldering θ, connection A and metric $g = \eta(\theta, \theta)$. Supply a cooriented contact five-manifold $(Z, \mathcal{H})$ and a double fibration

$$B \xleftarrow{p} F \xrightarrow{q} Z.$$

Here F has dimension six, p is a proper surjective submersion with connected fibers S^2, and q is a surjective submersion with connected one-dimensional

fibers. Each $q|_{F_b}$ is an embedded Legendrian sphere. Supply a positive point density μ and require $\mu = \mathrm{vol}_g$.

For every positive local contact form α, require the semibasic covector ℓ defined by

$$q^*\alpha = \ell \circ \mathrm{d}p$$

to be nonzero, past directed and null for g on every point fiber. Changing α by a positive ray function changes ℓ by that function. These tests use the supplied incidence and the readout of its supplied Cartan fields. They require neither an ordinary-class presentation nor a weighted inverse.

A geometric comparison supplies an orientation-preserving point diffeomorphism $f : B_0 \to B_1$, a coorientation-preserving contact diffeomorphism $\Psi : Z_0 \to Z_1$, and their incidence comparison, with

$$f^*\mu_1 = \mu_0, \qquad \Psi\big(q_0(F_{0,b})\big) = q_1(F_{1,f(b)}).$$

An actual principal spin comparison is separately retained, with its compatibility with the supplied soldering and frame reductions. No metric-isometry equation is included as a substitute for the incidence and density tests.

Lemma 7.2 (Full null incidence).

The preceding tests identify F with the positive projectivization of the past null covectors of g. They identify $\ker \mathrm{d}q$ with the null-geodesic line distribution and Z with its quotient by the retained connected leaves. In particular each supplied sky contains every real null direction exactly once. Conversely every admitted Cartan collar supplies this incidence after restriction to a sufficiently small causally convex point core with a buffered full-sky ray parent.

Proof. Put $\lambda = q^*\alpha$. It is nonzero because q is a submersion. It vanishes on $\ker \mathrm{d}p$ by the Legendrian hypothesis, so the covector ℓ exists smoothly and is unique. The radical of $\mathrm{d}\lambda$ on $\ker \lambda$ is $\ker \mathrm{d}q$: the differential of q maps that hyperplane onto $\mathcal{H}$, where $\mathrm{d}\alpha$ is symplectic.

The map $u \mapsto [\ell_u]_+$ is an immersion on a point fiber. Indeed, if $v \in \ker \mathrm{d}p$ differentiates ℓ only by $v(\ell) = c\ell$, then for every $w \in \ker \lambda$, $\mathrm{d}\lambda(v, w) = c\lambda(w) = 0$. Thus $v \in \ker \mathrm{d}q$. The sky embedding has injective vertical derivative, so $v = 0$. Its image lies in the past projective null sphere of g. The fiber map is consequently a local diffeomorphism between two-dimensional manifolds. Compactness makes its image closed as well as open. It covers the whole sphere, and is a one-sheeted covering because that sphere is simply connected and the source fiber is connected. These fiberwise diffeomorphisms give a diffeomorphism over B.

On $T^*B \setminus 0$, let $\vartheta = \xi_a \,\mathrm{d}b^a$ and $G = \frac{1}{2} g^{ab}\xi_a\xi_b$, with $\iota_{X_G}\mathrm{d}\vartheta = -\mathrm{d}G$. On $G = 0$, the radical of $\mathrm{d}\vartheta$ on the common kernel of ϑ and $\mathrm{d}G$ is spanned by the radial Euler field and X_G. Positive projectivization removes the Euler field. The remaining characteristic line pulls back to the radical of $\mathrm{d}\lambda$ just computed. Hamilton's equations for G are the affinely parametrized null-geodesic equations. This proves the line-distribution assertion.

The connected fibers of q are precisely the maximal connected leaves within the retained domain. A surjective submersion is open, and its local sections identify its quotient topology and smooth structure with those of Z. Thus this is the asserted local null-ray quotient, with its contact distribution.

For existence, take the actual null Hamiltonian flow of the recognized metric on a buffered small Gaussian collar. A spacelike disk transverse to that flow supplies its ray carrier as an open subset of $P_+T^*\Sigma$. Choose a smaller causally convex core so that every direction over it has one transverse hit inside the buffered disk; compactness of the sky and smooth transverse dependence permit this choice. Put $F = P_+\mathcal{N}_g^-|_B$ and let q be the hit map. Projection to B is a proper sphere bundle. The hit map is a submersion, and its fibers are the connected null plaques; causal convexity and the short flow parent exclude separated returns. Variation of a ray through a fixed point gives an embedded sky: injectivity follows from flow uniqueness, and the vertical derivative is injective by the same uniqueness in phase space. The tautological form vanishes on variations at that point, so the sky is Legendrian. The ray contact structure is the quotient of the tautological hyperplane field; Hamiltonian transport preserves it. Its positive contact frames give precisely the past null covectors required in the definition. Taking $\mu = \mathrm{vol}_g$ completes all tests. □

Theorem 7.3 (Physical-map recognition from the source).
A geometric comparison in Definition 7.1 is a time-orientation-preserving local isometry of the recognized metrics. Its ray map is uniquely determined by its point map. Conversely, every orientation- and time-orientation-preserving local isometry, with specified actual spin lift, induces the unique comparison of these incidence and density fields on common retained domains.

Proof. The map $(p, q) : F \to B \times Z$ is an embedding. Its derivative is injective because its kernel is vertical for p, where the sky map is an immersion. It is injective because each sky is embedded, and is proper because p is proper. The sky equation therefore induces a smooth incidence diffeomorphism Φ covering f and Ψ.

Write $\Psi^*\alpha_1 = a\alpha_0$, with $a > 0$. Pulling back the contact-normal equation gives

$$\ell_{1,\Phi(u)} \circ \mathrm{d}f_b = a(q_0(u))\ell_{0,u}.$$

Lemma 7.2 shows that these are the complete past null cones. Their quadratic annihilator is one-dimensional. To see this directly in a Lorentz frame, a quadratic form vanishing on $(1, n)$ for every $|n| = 1$ has the form $a_0 + 2c \cdot n + n^{\mathsf{T}}Dn$. Comparison at n and $-n$ gives $c = 0$; evaluation at the coordinate vectors and their pairwise normalized sums gives $D = -a_0 I$. Thus it is a multiple of the Lorentz form. Consequently $f^*g_1 = c_0 g_0$ with $c_0 > 0$. In dimension four its density is $c_0^2\mu_0$. Preservation of μ forces $c_0 = 1$. Positivity in the covector law preserves the selected past component, hence time orientation.

The covector lift is now necessarily $(b, [\xi]_+) \mapsto (f(b), [(\mathrm{d}f_b^{-1})^*\xi]_+)$. It determines the incidence map and, by surjectivity of q_0, the ray map. Conversely this lift of an isometry preserves the tautological hyperplane and the null-geodesic line. It descends to the retained connected-leaf quotients. Local sections of q_0, q_1 make the descended maps and their inverses smooth. They are contact and obey the sky and density equations. The specified principal spin map is included unchanged. Composition and restriction are the corresponding actual covector and principal-bundle operations. □

The use of the entire sky is essential. Finitely many directions can recover a candidate quadratic form, but do not test its equality on a supplied arbitrary sphere family. The density also matters: the contact incidence alone determines the conformal class, rather than its numerical metric scale.

7.2. Optical scale and the canonical finite-rank fields.

Proposition 7.4 (Actual optical and affine completion).
The incidence of Definition 7.1 has a canonical physical optical scale with value $y = 1$. For a contact form α, there is a unique null generator D with $\mathrm{d}p(D) = g^{-1}\ell$. It is affine for g. If the tangent-sky curve along a ray is written in a constant symplectic graph chart as a symmetric matrix $S(t)$, then S' is invertible and

$$\kappa_\alpha = \frac{1}{2}\operatorname{tr}\left(\frac{1}{2}(S')^{-1}S''' - \frac{3}{4}\big((S')^{-1}S''\big)^2\right) = \frac{1}{2}\operatorname{Ric}_g(k,k) = 0, \qquad k = g^{-1}\ell.$$

The selected value is constant on the whole point sky and solves $D^2y + \kappa_\alpha y = 0$. Its value, slope and positive branch determine the selected optical states and roots. Initialized fundamental matrices and companion solutions give the full affine transition maps. They have exact inverse and cocycle laws and are natural for the comparisons of Theorem 7.3.

Proof. Lift the incidence to $T^*B \setminus 0$ by $u \mapsto (p(u), \ell_u)$. Its Liouville pullback is $q^*\alpha$. The lift of D and X_G have the same base projection. Their difference is vertical and pairs to zero under $\mathrm{d}\vartheta$ with every lifted incidence tangent. Surjectivity of $\mathrm{d}p$ makes that vertical difference zero. Thus D is the affine Hamiltonian generator, including its covector normalization.

The contact plane along a ray is the four-dimensional space of screen Jacobi solutions; the tangent sky at a point is the kernel of their evaluation there. In a parallel orthonormal screen frame, let Y, Z solve $j'' + \mathcal{R}j = 0$ with identity and zero initial position and velocity blocks. Their symplectic Wronskian identities give, where Y is invertible,

$$S = -Y^{-1}Z, \qquad S' = -Y^{-1}Y^{-\mathsf{T}}, \qquad A_0 := (S')^{-1}S'' = -2(Y^{-1}Y')^{\mathsf{T}},$$

$$\frac{1}{2}A_0' - \frac{1}{4}A_0^2 = (Y^{-1}\mathcal{R}Y)^{\mathsf{T}}.$$

The screen tidal operator is $\mathcal{R}[J] = [R(J,k)k]$. Its trace is $\operatorname{Ric}_g(k,k)$: the two extra null slots in the full curvature trace vanish. Taking the trace proves the formula. Any constant symplectic change gives another fundamental pair

with the same Wronskians and Jacobi equation, so the scalar is independent of the graph chart. Such charts cover the curve. Ricci-flatness was already recognized from the actual Cartan source in Section 2, and gives its last equality.

For $\alpha_i = c_{ij}\alpha_j$, the factor c_{ij} is pulled back from the ray. Hence $D_i = c_{ij}D_j$ and $D_j c_{ij} = 0$, so $\kappa_i = c_{ij}^2\kappa_j$. The scalar optical equations agree on overlaps. Let $F_j(t)$ be the identity-initialized fundamental matrix for the state $(y, D_j y)$. A change of hit and contact frame has matrix $B_{ij} = \mathrm{diag}(1, c_{ij})F_j(\tau_{ij})$, with the hit time and coefficients evaluated in their moving arguments. Uniqueness of that finite ODE proves its inverse and triple identities.

More generally write the selected nonzero solution in an auxiliary scale as σ, and put

$$G_i = \begin{pmatrix} \sigma_i & 0 \\ D_i\sigma_i & \sigma_i^{-1} \end{pmatrix}, \qquad (G_i \circ \Phi_{ij})^{-1} B_{ij} G_j = \begin{pmatrix} 1 & b_{ij} \\ 0 & c_{ij} \end{pmatrix}.$$

The first column follows from preservation of the selected solution; the second diagonal entry follows from the determinant. The companion-to-selected ratio τ_i satisfies $D_i\tau_i = \sigma_i^{-2}$. Its affine change has slope c_{ij}^{-1} and translation $-b_{ij}/c_{ij}$. These give all affine entries, not only their linear parts. In the physical splitting $\sigma = 1$ and the selected state is $(1, 0)$.

For a physical incidence comparison with contact multiplier d, the physical affine generators obey $\Phi_* D_0 = d^{-1}D_1$. Therefore $\tau_1 \circ \Phi - d^{-1}\tau_0$ is constant on the connected ray plaque. Evaluate it at a retained hit to obtain the affine translation. It is independent of the hit, and these translations compose exactly. The positive root is the specified physical branch, with its initialized sign. The same ODE uniqueness carries every optical state and its inverse.

If a stored chart uses $\widehat{g} = \Omega^2 g$, its selected value is $\widehat{y} = \Omega$. With $u = D\log\Omega$, direct differentiation gives $\widehat{D} = \Omega^{-2}D$ and $\widehat{\kappa} = \Omega^{-4}(\kappa - Du + u^2)$. Thus $(\widehat{D}^2 + \widehat{\kappa})(\Omega y) = \Omega^{-3}(D^2 + \kappa)y$. This transports the same physical optical scale, rather than supplying an independent direction-dependent solution. Every construction above uses actual flow, differentiation and finite linear ODEs on the retained tubes, which proves its naturality. □

Proposition 7.5 (Canonical tractor, local-twistor and screen data).

The supplied Cartan soldering and actual spin connection determine the following finite-rank models in the physical Ricci-flat splitting. On $\mathcal{T} = \mathbb{R} \oplus \mathbb{E} \oplus \mathbb{R}$ set

$$\begin{aligned} H((\sigma, v, \rho), (\sigma', v', \rho')) &= \eta(v, v') + \sigma\rho' + \sigma'\rho, \\ \nabla^{\mathcal{T}}_X(\sigma, v, \rho) &= (X\sigma - \eta(\theta_X, v),\ D_X v + \rho\theta_X,\ X\rho), \\ \mathsf{I} &= (1, 0, 0), \qquad \mathsf{X} = (0, 0, 1). \end{aligned}$$

It preserves H, makes I parallel, and has $H(\mathsf{I}, \mathsf{X}) = 1$. Here D is the Lorentz connection induced by the actual spin connection; it is not the affine ray generator of the preceding proposition.

These are the explicit Ricci-flat normalization formulas for the finite-rank diagram (Isthmus 2026a, Proposition I.3). In the ordered rank-four spinor sum $(\omega^A, \pi_{A'})$, put $\vartheta(v) = 2^{-1/2}(v^0 I + v^a \sigma_a)$, where σ_a are the Pauli matrices, and $J = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$. With this soldering matrix convention, set

$$\nabla^{\mathrm{LT}}_X = P_X + \begin{pmatrix} 0 & i\vartheta(\theta_X) \\ 0 & 0 \end{pmatrix}, \qquad \mathcal{F}(\sigma, v, \rho) = \begin{pmatrix} \rho J & -iJ\vartheta(v) \\ -i\vartheta(v)^{\mathsf{T}} J & \sigma J \end{pmatrix}.$$

The map $\mathcal{F}$ intertwines the tractor connection with the induced connection on covariant alternating two-forms. Its normalization is $2\operatorname{Pf}\mathcal{F} = 2\sigma\rho + \eta(v, v)$. The rank-two infinity pair is

$$I_{\mathrm{cov}} = \begin{pmatrix} 0 & 0 \\ 0 & J \end{pmatrix}, \qquad I_{\mathrm{contra}} = \begin{pmatrix} J & 0 \\ 0 & 0 \end{pmatrix}, \qquad I_{\mathrm{cov}} I_{\mathrm{contra}} = 0.$$

Both are parallel, as are the spin volume and the Hermitian pairing with matrix $\left(\begin{smallmatrix} 0 & I \\ I & 0 \end{smallmatrix}\right)$. For a nonzero null k, the optical plane is $\operatorname{span}\{\mathsf{X}, (0, k, 0)\}$, with affine slice $\mathsf{X} + \mathbb{R}(0, k, 0)$; its complexification has the corresponding complex affine slice and spinor pencil. Along an affine null ray these data obey $\nabla_k \mathsf{X} = (0, k, 0)$, $\nabla_k (0, k, 0) = 0$.

The screen is $k^{\perp}/\mathbb{R}k$ with positive metric $-\eta$, its induced orientation and quarter-turn, and its induced connection. Its complete Jacobi state satisfies $D_k Y = Q$, $D_k Q = -\mathcal{R} Y$. These constructions and their actual bundle maps are natural under physical comparisons with specified spin lifts.

Proof. Differentiating H cancels the two soldering terms against the two ρ terms; the remaining equality is $D\eta = 0$. Direct substitution gives the assertions about I, X. The torsion identity $D\theta = 0$ gives curvature $R^{\mathcal{T}}(X, Y)(\sigma, v, \rho) = (0, R^D(X, Y)v, 0)$. This is the normal tractor formula since the Schouten tensor is zero in the recognized Ricci-flat metric.

The matrix convention gives $2 \det \vartheta(v) = \eta(v, v)$ and $\operatorname{tr}(\vartheta(\theta_X) \operatorname{adj} \vartheta(v)) = \eta(\theta_X, v)$. For the upper triangular soldering term Ω_X, multiplication of the displayed matrices gives

$$-\Omega_X^{\mathsf{T}} \mathcal{F}(\sigma, v, \rho) - \mathcal{F}(\sigma, v, \rho) \Omega_X = \mathcal{F}(-\eta(\theta_X, v), \rho\theta_X, 0).$$

The spin part is equivariant because $A^{\mathsf{T}} J A = J$ for determinant-one spin frames. This proves the full connection intertwining. Expanding the four-by-four Pfaffian gives $\sigma\rho + \det \vartheta(v)$. The selected infinity form is $\mathcal{F}(\mathsf{I})$, and its volume dual is the displayed contravariant form. Their ranks, product and parallelism follow from their matrices and volume preservation. For real X, $\vartheta(\theta_X)$ is Hermitian, so $\Omega_X^{\dagger} H_{\mathrm{LT}} + H_{\mathrm{LT}} \Omega_X = 0$; the spin part preserves that pairing as well.

The two generators of the optical plane are independent and have all three mutual pairings zero. Evaluation by I gives its affine slice. Every nonzero alternating form in its image under $\mathcal{F}$ has rank two. In a spin frame where $\vartheta(k)$ has one nonzero diagonal entry, its kernels contain one common line and lie in one common three-plane; solving the remaining linear equation gives the stated projective pencil. The tractor derivative formulas and the

Jacobi system follow directly from affine geodesic transport and the induced screen curvature.

Every formula uses the supplied soldering, spin connection and fixed determinant normalization. A physical map with its actual spin lift preserves precisely these data, hence intertwines all the formulas. The central spin element acts by -1 on the rank-four bundle and by $+1$ on its alternating square. Retaining the rank-four map keeps those two lifts distinct. The rank-two optical plane and the four-dimensional screen Jacobi state have different roles and are not identified. □

Lemma 7.6 (Independent test for supplied finite-rank fields).
Let a retained null flow tube have generator K_0, transverse section S and finite-rank connections D^g, D^s on two supplied bundles. Fix an actual initial isomorphism J_0. Define J by

$$D^s_{K_0} J - J D^g_{K_0} = 0, \qquad J|_S = J_0.$$

It intertwines the full connections if and only if their tangential initial connections agree under J_0 and $R^s(K_0, W)J = JR^g(K_0, W)$ throughout the actual tube for every transverse direction W. Parallel tensors agree if they agree initially. Selected tractor positions and screen Jacobi states agree under the additional actual equations and initial values of Proposition 7.5, including equality of the transported tidal endomorphisms. These are independent tests for supplied copies of precisely those canonical models.

Proof. If $P^g_{t,0}$ and $P^s_{t,0}$ are parallel transports on the flow lines, the unique candidate is $J = P^s_{t,0} J_0 (P^g_{t,0})^{-1}$. For a flow-transported transverse frame $[K_0, W] = 0$, its transverse defect $E(W) = D^s_W J - J D^g_W$ satisfies

$$\nabla^{\mathrm{Hom}}_{K_0} E(W) = R^s(K_0, W)J - JR^g(K_0, W).$$

The two tests make this a finite homogeneous ODE with zero initial value. Thus the defect vanishes in every direction. Necessity follows by taking the curvature of a connection intertwiner. If a different frame is used, the curvature formula contains its full bracket term $-D_{[K_0,W]}$; omitting that term is not an equivalent test.

Tensor products, contractions, duality and the stipulated real structures commute with parallel transport. Their initially zero differences therefore remain zero. For positions, the difference system is $De_X = e_V$, $De_V = 0$, $De_I = 0$. For Jacobi states it is $De_Y = e_Q$, $De_Q = -\mathcal{R}^s e_Y$ after the tidal test. Both are finite initialized systems. Principal spin transport gives the same assertion for an actual initial principal spin map and preserves its central sign.

These arguments also prove restriction, inverse and composition laws. At a fixed derivative order their differentiated finite ODEs give finite estimates on the same buffered tube, with constants depending on that order, the finite coefficient norms and the retained flow margins. The curvature tests hold throughout the tube in every retained angular direction. □

7.3. The complete independently admitted source.

Definition 7.7 (Complete local wave source and its arrows).

An object consists of the following supplied fields and normalizations, subject to the indicated equations.

Its initial gradients satisfy Definition 2.1 with maximal readout, and its initialized finite Cartan collar satisfies Definition 2.7 with exactly those initial values. Retain its actual principal spin bundle, physical scale and future orientation. The Cartan fields determine g by Theorem 2.9. Supply the incidence of Definition 7.1, the physical optical states and affine transitions of Proposition 7.4, and the canonical finite-rank diagrams of Proposition 7.5. Independently supplied realizations of those diagrams have specified initial comparison maps and satisfy Lemma 7.6, including the selected tensor, position and tidal tests. No additional unnormalized finite-rank object is included.

The object supplies the normalized scalar enrichment of Definition 6.23. The normalization maps for its fixed-carrier algebra and whole-sky modules are supplied with their inverses, full homogeneous characteristic maps, native half-density roots and central-specialization maps. They satisfy the inverse equations. The kernel, full evaluation section and point connections satisfy the independent tests of that definition.

The parent collar contains the marked maximal initial patch and carries the initialized evolution. The pointed spacetime core is a connected open restriction of this development. Its marked point may lie away from the initial surface. Germ comparisons restrict the core at that point. The parent retains the portions of evolution and the transverse carrier hits needed to define the compared fields. The choice of maximal surface and Gaussian clock belongs to the presentation and may change under an arrow. An auxiliary scalar carrier may be displaced from the point core, provided its retained hits and segments lie in the common wave neighborhood required by Definition 6.12.

Stored presentations retain their complete algebra comparisons φ_{ij}, normalized associators C_{ijk} and semilinear module comparisons τ_{ij}. The supplied normalizations, including their overlap comparison cells, identify them with the corresponding cotangent and actual distribution presentations of Section 6. In these normalizations the supplied maps and overlap cells satisfy the complete identities for products, inverses, filtration, internal adjoints, roots, characteristic maps, sections, point connections and central specializations. Their descent equations include

$$\begin{aligned} \varphi_{ij}\varphi_{jk} &= \mathrm{Ad}(C_{ijk})\varphi_{ik}, \\ C_{ijk}C_{ikl} &= \varphi_{ij}(C_{jkl})C_{ijl}, \\ \tau_{ij}\tau_{jk} &= C_{ijk}\tau_{ik}. \end{aligned}$$

The normalization data include the native density-line squares and their chosen branches. The degree-one evaluation and its connection in every point

direction are those of Section 6. They have no independent scalar connection or degree-zero full metric vector as additional data. The object retains the complete classes modulo joint smoothing.

An arrow supplies the actual geometric comparison of Definition 7.1 and its principal spin map, together with filtered internal-adjoint-preserving algebra maps and semilinear whole-module maps over its homogeneous ray comparison. On every graded piece the algebra map is the full homogeneous cotangent lift of that ray comparison. A freely adjustable radial lift of a projective map is not allowed. The maps preserve the canonical optical and finite-rank data, native density roots, full evaluation sections and both algebra and module point connections on every section. Thus, with pullback by the point map f understood,

$$\Xi e_0 = e_1 \circ f, \qquad \nabla^{\mathrm{amb},1}_{f_*X}\Xi = \Xi\nabla^{\mathrm{amb},0}_{X}, \qquad \nabla^{\mathrm{alg},1}_{f_*X}k = k\nabla^{\mathrm{alg},0}_{X}.$$

In stored charts its vertices k_i, Ξ_i and normalized overlap cells T_{ij} obey

$$\begin{aligned} k_i\varphi^0_{ij} &= \mathrm{Ad}(T_{ij})\varphi^1_{ij}k_j, \\ T_{ij}\varphi^1_{ij}(T_{jk})C^1_{ijk} &= k_i(C^0_{ijk})T_{ik}, \\ \Xi_i\tau^0_{ij} &= T_{ij}\tau^1_{ij}\Xi_j. \end{aligned}$$

Both source and target associators occur here. The stored sections and connections are carried by the same vertices; for a point-dependent module normalization F, its connection matrix is $A'_X = FA_XF^{-1} - (XF)F^{-1}$. A representative sends the marked spacetime point to the marked spacetime point. Two representatives define the same arrow germ if on a common smaller pointed core their point and principal spin maps agree and, on a common refinement of the tested ray charts there, their remaining data agree up to the normalized presentation cells. The transverse transports are performed on the corresponding buffered parents. Write $\mathsf{WaveSrc}^{\mathrm{loc}}_0$ for this groupoid.

More explicitly, compare vertices on the source core B after pulling target coefficient sheaves back by f and reindexing passive parameters by the supplied base map (the identity for families over the same base). Algebra coefficients basic in the point remain basic. The module maps are over the point map: multiplication by a smooth point or passive function on the target acts on the source by its pullback. The tested conic domain on the source carrier is $\Omega_B = \bigcup_{p\in B} L_p$; its image is the tested domain on the target. Restrict all charts and overlap equations to these open sets. With the point and actual spin maps fixed, pass to one common refinement. A normalized change of target presentation by units $U_i \in 1 + F_{-1}$, $U_i^*U_i = 1$, acts on the vertices and overlap cells by

$$\begin{aligned} k'_i &= \mathrm{Ad}(U_i)k_i, \qquad \Xi'_i = U_i\Xi_i, \\ T'_{ij} &= U_iT_{ij}\varphi^1_{ij}(U_j^{-1}). \end{aligned}$$

The sections and connections are expressed in the changed presentations, including the derivative term above. This formula follows by substitution in the three overlap equations; a change of source presentation is pulled through the invertible vertices. Refinements and these compatible changes identify representatives. Once the presentations are fixed, Lemmas 7.8–7.10 test the remaining ambiguity on Ω_B. Equality also requires equality of the actual principal spin maps, including their central sign.

The complete algebra and module maps preserve their displayed filtrations and act on the coefficientwise completions. Algebra comparisons are morphisms of the smooth coefficient sheaves. Smooth dependence on additional parameters means that their maps on every fixed finite filtration quotient act smoothly on smooth coefficient families, as do the supplied inverse maps and module normalizations. Admission uses only equations on supplied source fields, actual maps and complete coefficient families. It includes no ordinary pair, weighted smallness, global inverse or membership in a constructor image. The original analytic calibrated Hom-line and its coupled Real operation are addressed by the later analytic comparison; they are not identified with the new evaluation by this definition.

7.4. Normalized algebra and module ambiguities.

Lemma 7.8 (Normalized inner comparison).
Let Φ be an automorphism of the complete algebra of Proposition 6.2, covering the identity, fixing passive functions, preserving the internal adjoint and acting identically on every graded piece. There is a unique $U \in 1 + F_{-1}$ with $\Phi = \mathrm{Ad}(U)$. It satisfies $U^*U = 1$. For smooth comparison families, its every finite coefficient is smooth and determined by finitely many coefficients and derivatives of Φ. The construction commutes with restriction.

Proof. Suppose the first nonzero difference of Φ lowers order by k. Its leading action δ is a derivation of the commutative homogeneous algebra, by multiplicativity. Preservation of commutators makes it a Poisson derivation. The leading internal adjoint on degree m is $(-1)^m$. Commutation with that adjoint gives $(-1)^m \delta a = (-1)^{m-k} \delta a$. Thus the difference vanishes at every odd k, and the first possible k is even and at least two.

The local derivation is a smooth homogeneous vector field V on the symplectic cone. Indeed, the Hadamard formula for a smooth function minus its affine first-order part makes a derivation depend only on the coordinate derivatives; apply this in degree-zero angular coordinates and an invertible radial coordinate. The Poisson equation says that V is symplectic, and $[E, V] = -kV$ for the Euler field. If $\lambda = \xi \, \mathrm{d}x$ and $\omega = \mathrm{d}\lambda$, then $\iota_V \omega$ is closed and has Euler weight $1 - k$. Cartan's formula gives the actual primitive

$$b = \frac{\lambda(V)}{1 - k}, \qquad \mathrm{d}b = -\iota_V \omega, \qquad \delta a = \{b, a\}.$$

The nonzero weight removes any topological primitive ambiguity. The degree of b is the negative odd integer $1 - k$. Take a symbol with this leading

coefficient and skew-symmetrize it to obtain $B^* = -B$ with $\sigma_{1-k}B = b$. Conjugating Φ by $\exp(-B)$ removes its leading difference. Repeat at each next nonzero even lowering degree. The skew generators have orders tending to minus infinity, so their ordered exponential product converges in every finite filtration quotient to the required U. Its finite products are internally unitary, proving $U^*U = 1$.

For uniqueness, the normal product shows that commutation with all coordinate multiplications gives $\partial_\xi P = 0$, and commutation with all coordinate derivatives gives $\partial_x P = 0$. The local center consists of degree-zero constants, or passive functions. A central unit of $1 + F_{-1}$ is therefore one. The ratio of two implementing units has this property. Each finite recursive coefficient uses finite differentiation and division by a nonzero integer; uniqueness makes the local constructions independent of the coordinate presentation and makes them glue. This proves all family and restriction assertions. □

Lemma 7.9 (The remaining module scalar).

On a connected sky, a module automorphism over the identity algebra map is multiplication by a nowhere-zero passive-parameter scalar. Consequently, if two algebra comparisons differ by $\mathrm{Ad}(U)$ as in Lemma 7.8 and F_0 is a module comparison over the first, every module comparison over the second is cUF_0 for such a scalar c.

Proof. Flatten the sky in the normalized angular model of Proposition 6.3. Write a module endomorphism on its generator as $u \mapsto a(q,\tau)u$. Commutation with multiplication by the q coordinates and with $D_z = \tau$, followed by coefficientwise completeness, makes its action on every complete coefficient multiplication by a. Commutation with $D_{q^a} = \partial_{q^a}$ gives $\partial_{q^a} a = 0$. The generator is annihilated by $zD_z + \frac{1}{2} = -\tau\partial_\tau$, so $\tau\partial_\tau a = 0$. Its complete homogeneous expansion therefore contains only a degree-zero coefficient independent of q. It may depend only on the passive parameters. Connected angular overlaps make it the same scalar on the whole sky. Invertibility is equivalent to its nonvanishing. Undoing the graph and normalized Fourier integral comparisons proves the general statement. Multiplication by U changes semilinearity to the second algebra map; the ratio with any other comparison is the automorphism just classified. □

Lemma 7.10 (Faithfulness on the tested ray domain).

Let B be a connected open point core, let Σ be a retained scalar carrier, and let $L_p \subset T^*\Sigma \setminus 0$ be the conic sky of p. The union

$$\Omega_B = \bigcup_{p \in B} L_p$$

is open and conic. If a complete carrier operator P, basic in the point variable, satisfies $Pe_\Sigma(p) = 0$ for every $p \in B$ as a joint complete-family identity, then $P = 0$ on Ω_B. This does not require B to meet Σ.

Proof. Follow a carrier covector to a point $p \in B$ at its retained flow time. At that same local flow time, every sufficiently nearby carrier covector

reaches a point still in the open set B. Those covectors therefore belong to the union. This proves openness. Positive radial dilation preserves the unparametrized ray, proving conicity. The argument is on connected retained plaques, rather than an unrestricted long-ray quotient.

The evaluation has nonzero principal module symbol on every L_p by Proposition 6.16 and the ellipticity of carrier transport. If P has upper order m, the leading symbol of its annihilation identity is $\sigma_m(P)|_{L_p}\,\sigma_1(e_\Sigma(p)) = 0$. Hence the same actual homogeneous function $\sigma_m(P)$ vanishes on every L_p, and thus on their open union. Basicness is what makes it the same function for all p. The order of P drops by one. Repeating gives $P \in F_{m-j}$ for every finite j, so the complete class is zero. Every degree is tested on the same domain. Joint smoothing makes these coefficient identities compatible in the point and passive parameters. A single sky tests the symbol only along its own conic Lagrangian. As the point varies, these tests cover Ω_B, which is the domain of the conclusion. □

7.5. Complete comparisons across surfaces and carriers.

Theorem 7.11 (Full source Hom).

Every geometric comparison of two independently admitted sources of Definition 7.7, with its specified actual spin lift, has a unique full source lift modulo normalized presentation cells. Conversely every full source arrow has the physical point map recognized in Theorem 7.3 and is that lift. The lift preserves the full evaluation, all point directions of both connections, the internal adjoint, native density roots, source and target associators, and canonical optical and finite-rank diagrams. It has exact identity, inverse, composition and restriction laws on common retained domains. Its two initial hypersurfaces need not agree or be disjoint, and its scalar carriers need not meet the point cores.

Proof. Theorem 7.3 first recognizes the point map f as an actual isometry. Propositions 7.4–7.5 and Lemma 7.6 give its canonical optical and finite-rank comparisons, with the stipulated actual spin map. This step uses source incidence and density tests and the independently admitted Cartan fields.

Work first in subneighborhoods on which the restriction of f identifies both causal propagators. To obtain them choose a causally convex target diamond in the range of f, then a causally convex source diamond in its preimage, and finally a smaller target diamond inside the image of the source diamond. Its preimage is causally convex in the source wave neighborhood: a causal segment between its points stays in the source diamond and its image stays in the smaller target diamond by causal convexity there. The two isometric final domains are globally hyperbolic. Choose a small spacelike disk D through the marked source point in this common isometry domain. After one common restriction of the pointed cores, the null segment for each direction at a source point p meets D transversely, and its image segment at $f(p)$ meets $f(D)$ transversely in the target. Compactness of the full sky gives uniform transversality on unit directions and a positive margin to both disk

boundaries. Choose the restriction so that each segment from p to its hit on D lies inside the domain of f. Let C_f be actual half-density pushforward from D to $f(D)$. In each source separately, use the carrier maps of Theorem 6.14 to put

$$F_f = T^1_{\Sigma_1, f(D)}\, C_f\, T^0_{D,\Sigma_0}, \qquad k_f(P) = F_f P F_f^{-1}.$$

The outside factors require only the individual source neighborhoods. No value of f on the possibly displaced Σ_0 is required.

Kernel uniqueness in Theorem 6.10 identifies the kernels on the small isometric subparents. Together with Proposition 6.16 and the carrier cocycle, it gives $F_f e_0(p) = e_1(f(p))$ in the complete joint quotient. Each factor is independent of the evaluation point. Thus actual point differentiation and its full chain rule give both point-connection identities on every section. The internal-adjoint and semilinearity identities follow from Theorem 6.14 and half-density pushforward. Its root statement is the native one: the initialized leading amplitude is the half-density Jacobian root. In the normalized spacelike evolution of Lemma 6.13, internal skewness gives the left-symbol relation $2b_0 = \partial_x \partial_\xi b_1$, so the physical subprincipal symbol is zero and no additional subprincipal phase changes that root. The kernel formula fixes its branch and its composition. This is not a comparison with the separately calibrated analytic root.

For a second disk D' through the marked point, work on a small causally convex subparent in the domain of f and its isometric image. Restrict the two disks and the point core so that the intervening ray segments stay there. Write C_f^D for pushforward on D. Naturality of the mixed kernel and its diagonal square roots gives

$$C_f^{D'} T^0_{D',D} = T^1_{f(D'),f(D)} C_f^D.$$

Precompose with T^0_{D,Σ_0} and postcompose with $T^1_{\Sigma_1,f(D')}$. The two carrier cocycles give exactly the lift above, computed on D' on the left and on D on the right. For a second isometry g, take D small enough that $f(D)$ is an auxiliary disk for g. In $F_g F_f$ the consecutive transports from $f(D)$ to Σ_1 and back cancel. Half-density pushforwards compose as $C_g^{f(D)} C_f^D = C_{gf}^D$, giving $F_g F_f = F_{gf}$ on the common domain. The principal spin maps compose in the same order.

Conjugate by the supplied stored normalizations to obtain the stored algebra and module vertices, carrying their section and connection by the same maps. Substitution in the supplied normalization overlap equations gives the first and third equations in the two-associator test. The two ways to compose three algebra vertices implement the same comparison; Lemma 7.8 makes their normalized units identical. This gives its middle equation with both supplied associators. The module ratio in Lemma 7.9 is fixed by the transported full nonzero evaluation, so its overlap equation holds with the same units. To see the composition in these presentations, use one

refinement and the same intermediate charts for f and g, suppressing point pullbacks. The composite has vertices $k_i^g k_i^f$ and $\Xi_i^g \Xi_i^f$. Its overlap unit is

$$T_{ij}^{gf} = k_i^g(T_{ij}^f) T_{ij}^g.$$

Indeed, moving k_i^g past φ_{ij}^1 in the algebra equation gives conjugation by T_{ij}^g, and moving Ξ_i^g past T_{ij}^f in the module equation gives multiplication by $k_i^g(T_{ij}^f)$. For the triple overlap, apply the middle equation of Definition 7.7 first to g and then to f:

$$\begin{aligned} T_{ij}^{gf} \varphi_{ij}^2(T_{jk}^{gf}) C_{ijk}^2 &= k_i^g \left(T_{ij}^f \varphi_{ij}^1(T_{jk}^f) C_{ijk}^1 \right) T_{ik}^g \\ &= (k_i^g k_i^f)(C_{ijk}^0) T_{ik}^{gf}. \end{aligned}$$

Here the intermediate associator C_{ijk}^1 is the one carried by the intermediate normalization. More explicitly, a normalized change U_i of a target vertex sends k_i to $\mathrm{Ad}(U_i)k_i$ and Ξ_i to $U_i \Xi_i$. The unit in the new charts is $U_i T_{ij} \varphi_{ij}^1(U_j^{-1})$, by the first and third equations of Definition 7.7. Substitution in the middle equation cancels the two U_j factors and moves the last U_k^{-1} past C_{ijk}^1 by the target descent law. The source descent law gives the inverse change upon precomposition. For two different presentations of the intermediate object, insert their supplied transition in one arrow and its inverse in the other on a common refinement. Their vertices cancel there, and the displayed unit law gives the same composite up to the normalized endpoint cells. These formulas carry the section and the point connection by the same change of vertex.

For uniqueness, compose another lift with the inverse of this one. Its actual geometric and spin labels are identities. In fixed-carrier normalizations the algebra point-connection equation makes its action on every basic operator independent of p. Coefficientwise asymptotic summation then supplies a point-independent representative of each resulting complete class, so basicness has the actual meaning of Definition 6.18. Its full graded map is the identity. Lemma 7.8 gives a basic normalized unit U implementing it. The module comparison has the form $c(p)U$ by Lemma 7.9. Intertwining the point connection on every module section gives $(Xc)m = 0$ for every X, m, because U is basic. Hence c is constant on the connected point core. The evaluation equation is now

$$(cU - I)e_1(p) = 0 \qquad (p \in B_1).$$

Lemma 7.10 makes $cU - I = 0$ on the tested open ray domain, even when the carrier is displaced. Its degree-zero symbol gives $c = 1$, then $U = I$ in every degree. Thus the normalized scalar maps agree. The supplied canonical finite-rank tests fix their remaining bundle maps, and the actual spin label has been retained throughout. This proves uniqueness including the stored normalized cells.

The disk and composition formulas hold modulo jointly smooth kernels on the common tested ray charts over the restricted core. For the identity

isometry the outside transports cancel, and the lift of f^{-1} composes on either side to the identity there. Restriction of the parents restricts the same kernel formulas. The uniqueness just proved identifies their stored representatives on Ω_{B_1} up to the normalized cells. The two central spin lifts are not identified: their scalar maps can coincide, but their actual principal and rank-four maps remain distinct. □

Proposition 7.12 (Finite families of supplied comparisons).
For smooth finite-dimensional families of admitted sources and supplied smooth geometric comparisons with actual spin lifts, Theorem 7.11 gives smooth full comparisons locally on one retained parent. Every specified finite coefficient depth, point derivative order and passive-parameter order is bounded by finitely many source, comparison and normalization seminorms on larger buffered parents and their positive inverse margins. The parent is selected before these orders.

Proof. Near a fixed parameter value choose D and $f(D)$ with positive spacelike margins on one compact buffered neighborhood. The supplied maps and inverse maps remain defined there after one parameter restriction. Compactness of the full sky gives one finite transverse-hit cover and one common smaller output core. No coefficient recursion changes those geometric flow domains.

The geometric comparisons use finite differentiation, initialized geodesic and bundle transport, matrix inverses and the optical state formulas. Their derivatives at any fixed order obey differentiated finite ODEs on the same tubes. The middle half-density map in the full-lift formula uses finitely many derivatives of f and a square root of a positive determinant. The outer carrier maps have the estimates of Proposition 6.17; the point connections have those of Proposition 6.22. Finite composition has the estimates of Lemma 6.4. The stored normalization maps and their inverses contribute their stated finite norms. Finally every fixed coefficient of a normalized comparison cell uses only finitely many recursion steps in Lemma 7.8. Combining these bounds proves the assertion, with order-dependent constants and finite input orders. These estimates start with a supplied smooth family of physical isometries and spin lifts. □

The full comparison identifies arrows even when the chosen surfaces differ. We now use it to construct inverses on objects and arrows.

8. Local presentations, inverse functors and descent

Different whole inputs can present the same physical germ. The inverse therefore returns the comparison class of its local displays. At a point, the displays are compared by full source arrows over physical isometries with specified spin lifts. On an open cover, these arrows give the identifications on overlaps.

8.1. Execution and the two inverses.

Definition 8.1 (The three local groupoids).

Write $\mathcal{V} = \mathsf{Vac}_0^{\mathrm{loc}}$ for the groupoid of pointed, oriented and time-oriented smooth Lorentzian four-dimensional Ricci-flat germs with their numerical metric and actual local spin structures. An arrow is the germ of a pointed orientation- and time-orientation-preserving isometry together with a specified lift to the principal spin bundles, commuting with their orthonormal frame projections. Restriction to a smaller neighborhood defines the same germ. Write $\mathcal{S} = \mathsf{WaveSrc}_0^{\mathrm{loc}}$ for the independent sources and comparisons of Definition 7.7, on the corresponding pointed cores.

An ordinary-pair display $P = (c, r, V, \mathfrak{m})$ consists of a whole ordinary pair $c = (c_+, c_-)$ in the small class domain of §§3–4, a positive number r, a pointed initial patch V and markings $\mathfrak{m}$. These include the spin marking and a spacetime point x_* in the retained core of the developed collar. The spacetime base point of EP is x_*; the marking of V specifies an initial-data point. For a disk chosen through x_* one may take its initial marking at x_*. Execute the whole pair, develop the resulting initial source, restore its scale and retain the local fields on a buffered parent of V. The wave marking retains a neighborhood and carriers satisfying Definition 6.12. Such a neighborhood exists for every physical germ in Theorem 5.7: take a small causally convex diamond in its developed collar, a spacelike disk through the marked point, and a smaller point core whose whole sky has interior hits. Retain the full wave enrichment there. The domain of a comparison can be reduced about the spacetime point. For these reductions one keeps representatives on buffered parents long enough to perform all transverse carrier transports. Denote the resulting source by EP. An arrow $P \to P'$ is a full local source comparison $EP \to EP'$ in $\mathcal{S}$. This defines $\mathcal{P} = \mathsf{WavePair}_0^{\mathrm{loc}}$ with composition inherited from $\mathcal{S}$. The objects retain their whole inputs, and the arrows compare their outputs on a local core.

Proposition 8.2 (Execution, recognition and actual local embeddings).

Execution defines a fully faithful functor $E : \mathcal{P} \to \mathcal{S}$. Recognition defines a fully faithful functor $R : \mathcal{S} \to \mathcal{V}$. Every independent source has a display and a full comparison from it on the pointed core. One may use a display of the source's own marked initial patch when the comparison point belongs to the common Cauchy development of the patch on which the initial data agree. For every $M \in \mathcal{V}$ there is a display P_M and an actual physical and spin identification $b_M : REP_M \to M$.

Proof. Theorem 4.18 constructs the two gradient operators and reads back the ordinary inputs on their whole domain. Corollary 2.12 constructs the collar from the derived data. Lemma 7.2 and Propositions 7.4–7.5 construct the geometric incidence, optical data and finite-rank fields. Section 6 constructs the kernel, the carrier module, its full evaluation and its point connection. Proposition 5.6 restores their numerical scale, with $h \mapsto r^2 h$,

$K \mapsto rK$ and $e \mapsto r^{-1}e$. These operations give an object of Definition 7.7. On arrows E uses precisely the specified full comparison, which proves the first Hom assertion.

For a supplied source S, first recover the initial data by Theorem 2.3. The initial maximality test and the initialized Cartan equations apply on its own collar. Theorem 2.9 recognizes the metric $g_S = \eta(\theta, \theta)$ there as Ricci-flat. Lemma 7.2 identifies its actual null rays and Theorem 6.24 identifies the supplied scalar class and associated fields on their normalized domains. This procedure defines RS on the entire independent domain. On an arrow use the point map recognized by Theorem 7.3 and the supplied principal spin map. Conversely, an arrow of $\mathcal{V}$ between RS and RS' supplies that geometric comparison, and Theorem 7.11 gives its unique full lift. Write this lift as

$$L_{S,S'}(f) : S \longrightarrow S', \qquad RL_{S,S'}(f) = f.$$

Here and below f includes its spin lift. The restriction, identity, inverse and composition assertions of that theorem make R a functor and give its Hom bijection.

First retain the marked initial patch of S. Theorem 5.5 supplies a whole input whose data agree on a smaller patch, including the future normal. The initialized spin identification and the Cauchy isometry of Theorem 2.10 identify a common development of that patch. When it contains the comparison point, restrict the isometry there and apply L to compare the executed display to S.

For an arbitrary pointed physical germ M, choose a spacelike tangent plane and a spin frame at its point. Theorem 5.7 supplies a new maximal disk through that point, a whole input and a physical identification $b_M : REP_M \to M$ with its continued spin lift. In particular, take $M = RS$ at any marked point of an independent source. The full comparison $L_{EP_M,S}(b_M)$ of Theorem 7.11 then compares the display with S on a smaller pointed core, including any displaced scalar carriers. This gives essential surjectivity at every point and supplies the local choices used in Theorems 8.3 and 8.5. □

Theorem 8.3 (The inverses and their natural comparisons).
The functors E and R are equivalences. A local selection of the embeddings b_M gives inverse functors for both, with natural units and counits satisfying both triangle identities. All comparisons lift the actual physical arrows, including comparisons between displays with different maximal surfaces, clocks and scalar carriers.

Proof. Put $Q = RE$. Proposition 8.2 makes Q fully faithful. Choose (P_M, b_M) locally as in that proposition. Define $K : \mathcal{V} \to \mathcal{P}$ by $KM = P_M$ and by the unique lift

$$Q(Kf) = b_N^{-1}\, f\, b_M \qquad (f : M \longrightarrow N).$$

The right side is a composition of germs on a common restriction and of their actual spin maps. Identity and composition follow by cancellation and faithfulness of Q. Thus $b : QK \Rightarrow 1_{\mathcal{V}}$ is natural. For $P \in \mathcal{P}$ let

$\eta_P : P \to KQP$ be the unique arrow over b_{QP}^{-1}. For $h : P \to P'$, applying Q to $(KQh)\eta_P$ gives $b_{QP'}^{-1}Qh$, the image of $\eta_{P'}h$. Faithfulness gives naturality of η. The triangles are

$$b_{QP}\, Q\eta_P = 1_{QP}, \qquad K(b_M)\, \eta_{KM} = 1_{KM}.$$

The first is the defining inverse identity. In the second, $QK(b_M) = b_M^{-1} b_M b_{QKM} = b_{QKM}$, so the image of the left side is the identity, including its spin action. This proves the two laws for the composite.

For the equivalence R take the inverse EK. Its counit is b and its unit at S is $u_S = L_{S,EKRS}(b_{RS}^{-1})$. Naturality follows from the arrow formula and the Hom bijection for R. The first triangle is $b_{RS}Ru_S = 1$. For the second triangle the image of $(EK)(b_M)u_{EKM}$ under R is $b_{QKM}b_{QKM}^{-1} = 1$. For E take the inverse KR. Its unit is η and its counit at S is $v_S = u_S^{-1} : EKRS \to S$. At EP the arrow $v_{EP}E\eta_P$ lies over $b_{QP}b_{QP}^{-1} = 1$. At KRS the arrow $(KR)(v_S)\eta_{KRS}$ has image $b_{QKRS}b_{QKRS}^{-1} = 1$ under Q. Faithfulness gives the remaining two triangles. Thus each unit and counit in a source category is an arrow of the complete source comparison, with all of Definition 7.7's structure.

For a second selection (P'_M, b'_M), lift $(b'_M)^{-1}b_M$ through Q. Functoriality gives the natural comparison $P_M \to P'_M$ and its uniqueness over the specified physical identification. A third selection gives the composite of these lifts. In particular exterior choices, gauges and distinct initial disks of a fixed physical germ compare through their embeddings. The intermediate disk used in Theorem 7.11 lies inside the domain of the physical arrow, even when neither retained carrier lies there. The inverse returns a local display up to this comparison. An automorphism with the nontrivial central spin lift remains distinct from the identity. □

8.2. Families and descent over open covers.

Proposition 8.4 (A smooth selection on one parent).

For a smooth finite-dimensional family on common representatives with smooth markings, selections and the inverse identifications above exist locally in parameter space. On compact subsets of a sufficiently small selected parameter neighborhood, one positive radius, extension branch and set of buffered parents works for all prescribed finite orders. The same conclusion holds for the lifts of supplied smooth comparison and spin families.

Proof. The plane and lifted frame used to select a maximal disk can be chosen in smooth local bundle charts. Fix the base norm and its threshold in Proposition 5.8 before choosing the radius and extension. That proposition controls the whole input, embedding, continued spin frame and the derived wave fields. Positivity of the spacelike and inverse margins and compactness of the full sky give one finite set of transverse-hit patches for the selected representatives. For two selections of the same family the physical map is $(b'_\lambda)^{-1}b_\lambda$. Its finite derivative bounds follow from the inverse function formula

and the two embedding estimates, on a compact subset with positive inverse and image margins. Their actual spin maps compose smoothly, with the chosen initial sign. For an independently supplied comparison f_λ the middle factor in the same formula is f_λ.

Proposition 7.12 estimates the complete lifts of these maps, including conjugation by the stored normalizations and every required inverse. Consequently each stipulated point order m, parameter order q and homogeneous depth J has an estimate of the form in Proposition 5.8, with finite norms of any supplied maps and inverse normalizations on those parents included. Increasing (m, q, J) uses stronger finite input norms on the selected parents. The equations for the triangles concern those same maps at every degree. For a smoothly varying scale one includes its derivatives and a positive lower bound. The construction requires a smooth comparison as input only when a comparison between two different physical families is requested. □

Theorem 8.5 (Effective descent of local atlases).
On open covers of Hausdorff spacetime representatives, with effective descent for the actual spin bundles, local ordinary-pair atlases with their full comparison maps have effective descent. Arrows descend uniquely. The assertion holds locally in the smooth families of Proposition 8.4.

Proof. Use charts centered at variable points of an open representative, marking a distinguished point on every chart that contains it. Restriction refines the spacetime cores, the connected null plaques and their buffered parents. An atlas assigns a display locally and full arrows on overlaps, over the specified base identifications, satisfying the triple identity in the groupoid of Definition 7.7. Presentation cells obey the two associator equations there. Morphisms are compatible families after common refinement. Passing to germs at a point recovers Definition 8.1.

Given descent data for such atlases, apply RE to their charts and overlap arrows. On the given base, the metric, orientations and point markings agree under pullback on intersections. They therefore define smooth global fields on that representative. Ricci flatness holds locally. Glue the actual principal spin bundles by the specified lifted transitions. The cocycle identity holds for principal maps and therefore retains the central sign on each overlap. This gives the physical object and identifications with the input charts. The same argument applies to an atlas presented over an already effective Hausdorff geometric gluing. For arrows the compatible local point and spin maps glue over the prescribed base arrow.

About each point of the descended representative select a display by Proposition 8.2. On small overlap cores lift the induced physical and spin identifications by Theorem 7.11. The lift has the required evaluation, connections, finite-rank fields and stored cells, and agrees on further restrictions by uniqueness. On triple intersections the two composites lie over the same actual physical and spin map and hence agree by full faithfulness. Equations

in Definition 7.7 give the corresponding local representatives with both associators. This constructs a descended atlas. Comparison of its restrictions with the originally given charts uses the same lift of their geometric identification. Applying the faithful recognition functor to the overlap diagram verifies compatibility with the original descent data. Similarly the lift of a glued physical arrow coincides on every member of the cover with its prescribed local arrow, which proves uniqueness and effectiveness for arrows.

On each small causally convex parent, Theorems 6.10 and 6.24 determine the kernel class and its normalized fields. The overlap maps identify these local data even when the union is not causally convex. The whole exterior of each generating pair belongs to its own chart. For a smooth family the same argument uses Proposition 8.4 on retained compact parameter subsets and finite refinements. On a noncompact representative each chart has its own positive width. □

9. Comparison on the analytic locus

The analytic correspondence uses a holomorphic operator algebra and a framed family of simple modules on complex skies. Its framing fixes the scalar of a module comparison. The wave source has a full evaluation section of degree one and a connection in every point direction. We first identify the two underlying module presentations. The analytic framing then fixes the scalar of the comparison map, which in turn determines the transport of the lower wave coefficients and the point connection.

The analytic input throughout is the original projective correspondence (Isthmus 2026a, Definition H.1 and Theorem 9.10), on its Lorentz locus at $\Lambda = 0$. Its independently admitted object has a separated contact germ along a compact ordered quadric, the first-negative-order unit torsor, calibrated affine and optical data, an analytic filtered half-density algebroid with complex-linear adjoint, and a locally standard simple module. The original projective correspondence recovers the complete support family from the compact characteristic. The analytic object retains the actual relative root square and a frame of the rank-one local-system factor over the point and passive-parameter family. In particular a frame-carrying module comparison can have a nontrivial scalar between different framed objects. Normalized *algebra* units have principal symbol one.

For the comparison we use v1.1, Theorem F.4 and Lemma F.8 for analytic quantization and normalized inner comparisons, Lemma G.2 and Theorem G.5 for module transport and the spin and Real lifts, and Lemma H.2, Proposition H.3 and Theorem H.4 for supplied comparisons, calibrated supports and descent. The actual counit of the original Theorem 9.10 identifies this independent domain with the analytic source data and carries these structures. Its use here requires the stated hypotheses on parents, finite refinements, calibrated root maps and full spin labels.

9.1. Real sectors and their module presentations.

Lemma 9.1 (Physical comparisons and analytic carriers).
A smooth isometry between analytic pseudo-Riemannian germs, together with either specified spin lift, is analytic on a smaller representative. Through a point of an analytic Lorentzian germ one can choose an analytic spacelike maximal disk. A smooth scalar carrier can subsequently be used for the wave realization.

Proof. If f is an isometry, preservation of the Levi-Civita connection gives $f = \exp_{f(p)} df_p \exp_p^{-1}$ on a normal neighborhood. Analytic geodesic ODEs and the analytic inverse function theorem make the expression analytic. In analytic orthonormal frame charts the lift through the spin covering is analytic on its specified sheet.

For the disk start with the smooth maximal graph from Theorem 5.2 in analytic ambient normal coordinates. Shrink and dilate so that the graph has small first and second derivatives and the ambient metric is close to the constant metric on a buffered ball. This is possible in coordinates adapted to its tangent plane, by the scaling in §5. The linearization in the graph variable with fixed boundary data is then an invertible small perturbation of the flat Dirichlet Laplacian from $C_0^{2,\alpha}$ (zero trace here) to $C^{0,\alpha}$. For an interior cutoff χ put $\psi_z(x) = x + \chi(x)z$. Regard the mean curvature of $x \mapsto (u(x), \psi_z(x))$ as an equation with the actual fixed boundary value of the graph u_0. Its coefficients use only the holomorphic ambient metric on a smaller complex neighborhood, u, Du, D^2u and the explicit first two derivatives of ψ_z. The positive spacelike margin preserves the initialized square-root branch. It is an analytic Banach-space equation in (z, u) also for complex z near zero. The implicit function theorem gives u_z. For real z, the reparametrized graph has $u_z = u_0 \circ \psi_z$ by the fixed-boundary uniqueness. Evaluating where $\chi = 1$ gives a holomorphic extension in the translation variable, which proves interior real analyticity of the graph. For a displaced smooth carrier Theorem 6.14 subsequently transports its wave fields through the real geometry. □

In real-adapted coordinates on a conic sector, restrict the analytic coefficient functions and their homogeneous variable to the chosen real cone. Its *complete formal algebra* allows smooth complex homogeneous coefficients at every successively lower integer degree, with the rephased product and internal adjoint of Proposition 6.2. Use the same coefficient convention for the shifted simple module in a flattened sky chart (Proposition 6.3), with its local half-form generator. For mixed coordinates make the analytic change to an adapted chart first. The overlap maps of these local models are the finite-order extensions of the analytic transitions constructed below.

Proposition 9.2 (Real sector presentations).
For a Lorentz analytic object of the original H.1 with $\Lambda = 0$, choose a local analytic Cauchy disk D as in Lemma 9.1. On a common finite refinement about the whole real sky, after shrinking the point neighborhood, there are filtered adjoint-preserving isomorphisms k_i from the algebra of each analytic

object chart to the algebra of the corresponding adapted chart on D, over the full homogeneous ray maps. Over each k_i there is a same-order semilinear isomorphism F_i from the simple module of the object chart to the shifted conormal module of the adapted chart. Their analytic sector maps and inverses extend to the complete formal models above and the corresponding modules on D. At each degree the output depends on finite jets of the input coefficients at the real preimage of the endpoint, smoothly in the point and passive parameters, and the maps restrict to the given analytic actions on analytic inputs. A supplied analytic quantization may be used when, on each strict conic compact of the common analytic coordinate and operator domains, the actual coordinate images under the comparison and its inverse have total-symbol terms of degree $m - j$ bounded by $MA^j j!$ for $j \geq 0$, with m the order of the image and M, A uniform over the retained family (v1.1, §§F.1 and H.1–H.2).

On overlaps the extended maps satisfy the two-associator and module comparison equations of Definition 7.7 with unique normalized algebra units. For fixed algebra maps and their overlap cells, the full module comparisons form a torsor for nowhere-zero smooth complex functions of the point and passive parameters: any two differ by exactly one such function before the framing of Proposition 9.5 is imposed. Composition with T_{aD} carries the comparison to a smooth carrier a on the common ray domain of D and a in Definition 6.12.

Proof. Use the geometric comparison of the original counit and the disk of Lemma 9.1. On a real Cauchy carrier the past covector over ξ is $(-H(\xi), \xi)$, $H = (h^{ij}\xi_i\xi_j)^{1/2}$. Its positive root, the null flow and the transverse hit map extend holomorphically on sectors about each retained direction. The past extensions of ξ and $-\xi$ describe different real rays. The sector choice retains both when the complex spatial projectivization identifies their spatial lines. Compactness gives a common finite refinement of the whole real sky, also over a small closed point core.

Use an analytic disk D through the center of the point core. The transverse hit map identifies the two real ray germs on this disk, including their homogeneous lifts. Choose real angular coordinates there and use their holomorphic extensions on the analytic side. In a mixed analytic chart first make this coordinate change within the analytic algebra. The real formal coefficients will be taken in the resulting adapted coordinates. In these coordinates the rephased product and complex-linear adjoint of Proposition 6.2 have the finite coefficient formulas of the original Appendix F.1.

The analytic simple presentation in a flattened contact chart has cyclic vector u with $D_{q^a}u = 0$ and $(zD_z + \frac{1}{2})u = 0$. On its coefficient functions the normal actions are $D_z = \tau$, $D_{q^a} = \partial_{q^a}$ and $z = -\partial_\tau + 1/(2\tau)$. They are exactly the actions of Proposition 6.3, with the translated equations for $z = s(q)$ on moving skies. The radial normalization can be seen on

$\xi = \tau(q^1, q^2, 1)$, $\tau > 0$. Then

$$\tau^{-3/2}|d^3\xi|^{1/2} = |dq^1 dq^2\, d\tau/\tau|^{1/2}.$$

It follows that the degree-zero analytic generator and the shifted conormal generator define the same local model. The normalized Radon map of Proposition 6.3 and the graph projection in Definition 6.15 put the skies of all nearby points in the form $z = s(q)$ on D. For the analytic disk their graph functions are analytic by the holomorphic hit map. The substitution $z' = z - s(q)$, with its cotangent lift and half-density action, flattens each graph. On the analytic side use the locally standard presentations and moving supports of the original H.1 and H.3. In the two local standard presentations choose the standard cyclic generators, each satisfying $D_{q^a} u = 0$ and $(z\, D_z + \frac{1}{2})u = 0$ in the flattened coordinates. The displayed normal actions then agree on coefficients, so sending one generator to the other with its coefficient fixed intertwines the actions over the identity and preserves order. Exchanging the generators gives the inverse.

We now lift the given algebra comparison in the flattened charts. Write the contact form as $\alpha = dz + p^t dq$ and the common reference Legendrian as $\Lambda_* = \{z = 0,\ p = 0\}$. Here p is an angular momentum coordinate; the physical point and the passive parameters are held as parameters in this construction. The flattened contact map h carries Λ_* onto itself. Write h_0 for h followed by the translation in q that removes $h(0)$ along Λ_*. This translation varies analytically with the family. For h_0 use $\delta_t(z, q, p) = (t^2 z, tq, tp)$. Its weighted leading map $S = \lim_{t\to 0} \delta_t^{-1} h_0 \delta_t$ is the contact lift of a conformal symplectic linear map with multiplier $\lambda \neq 0$, as in the original H.2. This map preserves $p = 0$. Its (q, p) matrix thus has blocks $(A, B; 0, \lambda A^{-t})$, with A invertible and $K = A^{-1}B$ symmetric, by the conformal symplectic equation. Explicitly S is the shear followed by the diagonal map

$$(z, q, p) \longmapsto (z - \tfrac{1}{2} p^t K p,\ q + Kp,\ p),$$
$$(z, q, p) \longmapsto (\lambda z,\ Aq,\ \lambda A^{-t} p).$$

Indeed the shear pulls α back to itself, the diagonal map multiplies it by λ, and the value of z at the origin fixes the integration constant in the lift. Replacing K by tK gives a shear path. For the diagonal path, at one reference parameter join the identity to the reference (A, λ) in $\mathrm{GL}(2, \mathbb{C}) \times \mathbb{C}^\times$ by matrix and scalar exponentials. On a smaller parameter chart append the exponential of t times the local logarithm of the remaining near-identity factor. Translation, shear and each diagonal path preserve Λ_* throughout.

The residual map $R = S^{-1} h_0$ has weighted leading identity. Both R and R^{-1} preserve Λ_*. For $t \neq 0$ the same is true of $R_t = \delta_t^{-1} R \delta_t$ and its inverse. Beyond the identity the terms in (q, p) have weight at least two and those in z have weight at least three. Hence substitution gives normally convergent power series with positive powers of t for the remainders of both maps, and $R_0 = R_0^{-1} = I$. More precisely, on a bounded complex neighborhood of $[0, 1]$ choose a smaller weighted polydisk whose dilates stay

inside the common convergence polydisks for R and R^{-1}. The scaled series give uniform holomorphic maps there. Their inverse identities extend across $t = 0$. After a further fixed restriction their images lie in the prescribed larger charts, and the contact multipliers and their inverses stay bounded away from zero. For the translations use charts about their moving centers; for the finitely many linear paths the matrices and inverses are bounded on smaller complex neighborhoods of their parameter intervals. These choices give flow and inverse-flow tubes about the retained cone sectors, using the contact multipliers to lift the radial variable. They can be made simultaneously on compact subsets of each conditioned parameter chart. A finite refinement covers the retained angular and parameter cores.

For any path piece let $V_t = \dot{h}_t h_t^{-1}$. Its restriction to Λ_* is tangent to Λ_*, so $\alpha(V_t) = 0$ there. The homogeneous Hamiltonian on the symplectization is $\tau\alpha(V_t)$ up to the pullback sign, and vanishes on the reference cone. Use the skew order-one lift $B(H) = (Q_0(H) - Q_0(H)^*)/2$ of the original F.1 for this Hamiltonian H, as in H.2. Its analytic symbol bounds on the smaller parents and the tubes just constructed meet the hypotheses of F.4 and G.2. Starting with the identity of the standard modules, G.2 yields invertible semilinear evolution for the algebra evolution of F.4. Compose the finitely many evolutions and undo the flattenings. Denote the resulting algebra and module maps by k_0, F_0. Reverse evolution gives their inverses. This gives existence with $k = k_0$ and $F = F_0$.

Now let k be a supplied filtered adjoint-preserving analytic comparison with the same full graded action. For a conditioned diagram H.1 requires uniform factorial bounds on the actual operator coordinate images and inverse comparisons on common parents. In particular $k(y_i)$ satisfies the source-coordinate image bounds used in the second part of H.2. For a single analytic germ the finitely many coordinate images are analytic symbols; restrict their parents together with those of the constructed inverse. For a family the bounds must hold on the stipulated retained parameter set. Form the source automorphism $\psi = k_0^{-1}k$. It fixes the full graded algebra and preserves the adjoint. On the smaller buffered parents its coordinate differences are

$$\Delta_i = \psi(y_i) - y_i = k_0^{-1}(k(y_i)) - y_i.$$

The controlled inverse k_0^{-1} acts on the bounded coordinate-image symbols, so the Δ_i have the uniform factorial estimates required by F.8. This is the left-quotient argument of H.2. Lemma F.8 gives a unique normalized source unit Q with $\psi = \mathrm{Ad}(Q)$. Taking $U = k_0(Q)$ gives $k = \mathrm{Ad}(U)k_0$. Thus $F = UF_0$ is a semilinear isomorphism over the supplied k with inverse $F_0^{-1}U^{-1}$. The estimates for Q and its inverse pass to U, U^{-1} through the constructed quantization.

To pass to smooth complete coefficients, factor out the leading pullback along the endpoint cone map and the transported half-form in G.2. In the transported coordinates the remaining generator strictly lowers degree.

Its term lowering degree by $j \geq 1$ differentiates an input at most $j+1$ times. At a prescribed total drop n the iterated integral therefore contains at most n factors, of total differential order at most $2n$. Time integration takes the coefficients of these differential operators; the input jet is taken at the single preimage under the endpoint map. The algebra evolution has the same finite-jet property by the positive-drop expansion in F.4. The normalized unit, its inverse and the graph substitutions also use finite jets in each degree. In real adapted endpoint coordinates these formulas act on arbitrary smooth homogeneous coefficients by real differentiation with complex coefficients. Intermediate complex paths enter only the coefficients of the formulas. Products, semilinearity and both inverse identities extend: at each fixed degree and point they involve finite jets, and Taylor polynomials in angular and base coordinates, multiplied by the prescribed radial powers, realize those jets by analytic germs. The operations depend smoothly on the physical point and passive parameters, with finite derivative bounds on the stated parents. This constructs the map of the complete modules on D. More generally, fix any smooth formal algebra comparison k' in these presentations with the same full graded map, adjoint and action on passive functions. By Lemma 7.8 its quotient with the extended k is conjugation by a unique normalized target unit V. Then VF is a module map over k', with inverse $F^{-1}V^{-1}$. The finite-coefficient construction of V in that lemma preserves smooth dependence in families.

On an overlap, apply the second part of H.2 and Lemma F.8 to the supplied transition φ^0_{ij} with the H.1 coordinate-image bounds. As for k above, this gives $\varphi^0_{ij} = \mathrm{Ad}(P_{ij}) \circ K_{ij}$, where K_{ij} is an F.4 quantization of the same graded map and P_{ij} is a normalized unit in the target chart. The quantization and conjugation, and their inverses, use only finite jets of the input coefficients at each degree after the endpoint pullback is factored out. Thus $k_i \varphi^0_{ij} k_j^{-1}$ and its inverse have this property at their real endpoints in the adapted ray charts and recover the analytic transitions on analytic inputs. Let T_{ij} be the normalized unit comparing $k_i \varphi^0_{ij}$ with $\varphi^1_{ij} k_j$ as in Definition 7.7. These maps have the same homogeneous action and preserve the adjoint, so Lemma 7.8 supplies T_{ij} uniquely. On a triple the two ways of composing the algebra maps have equal inner action, hence

$$T_{ij} \varphi^1_{ij}(T_{jk}) C^1_{ijk} = k_i(C^0_{ijk}) T_{ik}.$$

Consequently the module transports, corrected by these units, act on the lines of local intertwiners with an ordinary cocycle: in a triple composition the two module associators cancel by this equality. Over the fixed algebra comparison any two local intertwiners differ by a scalar constant in the sky directions by Lemma 7.9. The resulting flat Hom line on the positive cone has trivial monodromy since the cone retracts to S^2. Choose a nonzero value at one angular section over a small point-parameter chart and propagate it on a finite good cover. The transition scalars vary smoothly on the base, and the cocycle and trivial monodromy give path independence. The resulting

maps obey $F_i\tau^0_{ij} = T_{ij}\tau^1_{ij}F_j$ and retain precisely one nonzero scalar freedom on the base.

For a displaced smooth carrier a, compose F_D with T_{aD} and k_D with $\mathrm{Ad}(T_{aD})$. With both hits and the intervening segments retained in the wave neighborhood, Theorem 6.14 gives the inverse and carrier cocycle in the joint smoothing quotient. Proposition 6.16 gives $T_{aD}e_D(x) = e_a(x)$, and the transport between the two fixed carriers is independent of x, so it also transports point derivatives. Proposition 9.5 will select the remaining scalar using the relative root and the initialized unit of Propositions 9.3–9.4. The resulting family is the interface used to transfer the full section and connection in Theorem 9.7. □

9.2. Calibrating the root and its local-system factor.

Proposition 9.3 (The relative root and the wave frame).
Let $S \to B$ be the analytic point family over its passive base and $p : F \to S$ the compact complex quadric family. Write $\mathcal{L}$ for the contact line pulled to F. The physical metric gives $R_g = \mathcal{L}^{-1}$ with an actual square map $\kappa_g : R_g^2 \to K_{F/S}$. The analytic root (R, κ_R) factors as

$$R = R_g \otimes p^*A, \qquad q_A : A^2 \longrightarrow \mathcal{O}_S,$$

where the product square is κ_R. Relative to R_g, the principal evaluation supplies a frame n_{W} of the native wave local-system factor N_{W}.

Proof. The first-jet normal evaluation in the proof of the original Proposition H.3, with the null incidence identification of D.7, identifies the symplectization $(\mathcal{L}^{-1})^\times$ with the nonzero null cotangent cone in $E = T^*_{S/B}$. In particular this identification has the actual homogeneous radial coordinate. Put $q_g(k) = g^{-1}(k,k)$ and let $\mu_g^\vee$ be the oriented metric volume on the cotangent fibers, continued analytically from the chosen Lorentz phase. With $\mathscr{E}$ the fiber Euler vector field put

$$\eta_g = -\operatorname{Res}_{q_g=0}(\mu_g^\vee/q_g), \qquad \omega_g = \iota_{\mathscr{E}}\eta_g.$$

At a nonzero null covector $dq_g \neq 0$. Division by the derivative in an eliminated coordinate gives the residue independently of that coordinate by its wedge identity. In Gaussian covectors $q_g = k_0^2 - H^2$, $\mu_g^\vee = (\det h)^{-1/2}dk_0 \wedge d^3\xi$. On $k_0 = -H$ this gives $(\det h)^{-1/2}d^3\xi/(2H)$ with the indicated sign. The residue has weight two. In coordinates (τ, z^1, z^2) it is $\tau d\tau \wedge a_g\, dz^1 \wedge dz^2$ with a_g a unit, so ω_g is a nowhere-zero section of $K_{F/S} \otimes \mathcal{L}^2$. A change of radial coordinate adds a projective differential to $d\log\tau$ which wedges to zero against the projective top form. Multiplying by ω_g defines κ_g. The same residue, before choosing the metric volume, identifies $p^*\det E^*$ with $K_{F/S} \otimes \mathcal{L}^2$ and so retains the point-base canonical factor.

On each complete $\mathbb{P}^1 \times \mathbb{P}^1$ the line $\mathcal{L} \otimes R$ is of type $(0,0)$. Its fiberwise constant sections form $A = p_*(\mathcal{L} \otimes R)$ with evaluation isomorphism $p^*A \cong \mathcal{L} \otimes R$. Indeed, ω_g and κ_R trivialize the square on the ruled charts, and the two local roots differ by signs with trivial monodromy on each simply

connected quadric. Local base sections and their transitions give the stated evaluation. In a base frame b, the quotient

$$q_A(b^2) = (1 \otimes \kappa_R)(b^2)/\omega_g$$

is holomorphic and constant on every compact fiber. It is a nonzero function on the base. For example, if ℓ is the contact frame evaluated by τ and $\kappa_R(\rho^2) = s\, dz^1 \wedge dz^2$, the two factors in the square of $\ell^{-1} \otimes (\ell \otimes \rho)$ give a_g and s/a_g. Their product is the supplied s. This retains each base transition and any two-torsion factor.

Let $w_g = 1$ in $\mathcal{L} \otimes R_g$ and let $d_0 = (2\pi)^{-3/2}$. On the chosen real cone the half-form representing w_g is the physically normalized square root of the residue density. Indeed multiplying the degree-zero half-form $\sqrt{a_g\, dz^1 dz^2\, d\tau/\tau}$ by τ squares to η_g. Proposition 6.16 gives through the point the principal physical conormal half-density $d_0(\det h)^{-1/4}(2H)^{-1/2}|d^3\xi|^{1/2}$. The rephasing of degree one therefore defines the frame by

$$\sigma_1^{\mathrm{for}}(e) = -i d_0\, w_g \otimes n_{\mathrm{W}}.$$

In the standard module the quotient on the right has constant local-system coefficient on every angular chart. Transport by the complete carrier comparison gives the identical section on every other carrier. Both the residue and e are natural under physical arrows, hence so is n_{W}. Relative to R, its local-system factor is $A^{-1} \otimes N_{\mathrm{W}}$. This is a tensor factorization, retaining the square and the original root comparison even when $q_A(b^2)$ is negative in a Real frame. □

Functoriality of the calibrated root maps. In Proposition 9.3 the fiber of A_X consists of the fiberwise constant sections of $\mathrm{Hom}(R_{g_X}, R_X)$ in its quadratic trivializations. Use the forward direction for an arrow $f : X \to Y$, identifying its sky fibers by the induced map. Write $I_R(f)$ for the calibrated square-preserving root comparison of the original H.1 and H.3, and $I_g(f)$ for the native root comparison given by the contact line and metric residue. On these sections put

$$I_A(f)b = I_R(f) \circ b \circ I_g(f)^{-1}.$$

This is a map from A_X to A_Y over the point map and respects q_A by the two square identities. The composite of two calibrated root maps carries the same square and calibration as the root map of the composite, so they agree by calibrated uniqueness in H.3. Equivalently, a normalized comparison cell acts as one on the leading symbol and the separately framed local-system factor is carried by both maps. Cancelling that factor identifies the root maps. The native residue maps compose by pullback of the residue, hence $I_A(gf) = I_A(g)I_A(f)$; identities and restrictions follow from the same formula. For a coupled conjugate arrow use the conjugate source lines. This gives a map of quadratic lines before choosing any section of prescribed square.

Proposition 9.4 (An initialized unit on the independent analytic domain).

With the flat Lorentz calibration of the original source there is, locally on the retained point and parameter charts of every admitted analytic object X, a section a_X of A_X with $q_{A_X}(a_X^2) = \varepsilon_0$. Here $\varepsilon_0 = \pm 1$ is the actual sign of the normalized flat square. Physical comparisons and the original coupled operation carry these units. A simultaneous change of physical scale to $r^2 g$, with the original root square held in fixed cotangent coordinates, gives $\kappa_{r^2 g} = r^{-2}\kappa_g$, $q_A^{[r]} = r^2 q_A$, $a^{[r]} = r^{-1}a$.

Proof. In a frame b with $q_A(b^2) = \lambda$, write $a = sb$. The equation $s^2\lambda = \varepsilon_0$ defines an unramified double cover, with differential $ds = -\frac{1}{2} s\, d\log\lambda$. The change $b \mapsto ub$ gives $\lambda \mapsto u^2\lambda$, so this equation is intrinsic. On the product of an amplitude disk and a point-parameter parent, a choice on the entire zero-amplitude family lifts uniquely to the product by its retraction. Local square roots show that the lift is holomorphic. Coupled conjugation transports a compatible initialized value to the compatible extension by the same uniqueness.

We specify the zero-amplitude choice, since an automorphism of an uninitialized quadratic line may act by -1. Write a flat nonzero null covector as $\tau(1, z) \otimes (1, w)$ in the two ordered spin factors. If determinant-one spin transformations contribute multipliers k, l to the two spinors, then $dz' = k^{-2}dz$, $dw' = l^{-2}dw$, $\tau' = kl\,\tau$. In the product with $dz' \wedge dw'$ the extra term in $d\log\tau'$ vanishes. The actual spin lift on the degree-zero cone half-form ρ_0 thus has multiplier $(kl)^{-1}$, and the contact frame has multiplier kl. Their product is invariant under the two spin transformations and under translations. Choose the compatible half-form phase and initial calibrated sign of the original Theorem G.5. The product gives a Real flat section b_0 with constant $\lambda_0 = q_{A_0}(b_0^2) \neq 0$. Set $a_0 = b_0/\sqrt{|\lambda_0|}$ and $\varepsilon_0 = \operatorname{sgn}\lambda_0$. These are statements about the supplied phase and square.

For a constructed analytic source U, continue a_0 along its constructor amplitude and evaluate at one. To compare two physical presentations use precisely the original G.5 family $U_\epsilon = (\epsilon U; 0)$, $V_\epsilon = R_{d,H}(U_\epsilon)$, where (d, H) denotes the displacement and full spin transformation. Construct the unit for V_ϵ along its own internal family tV_ϵ, keeping ϵ passive, then evaluate $t = 1$. The conditioned two-parameter construction in G.5 allows these two paths on common parents. The square-preserving physical root comparison acts on the two units by a holomorphic sign. At $\epsilon = 0$ it is the translation and determinant-one spin comparison just computed. The sign is therefore $+1$ on the connected amplitude domain. In particular this argument uses the actual transformed V_ϵ, which generally differs from ϵV_1. The original coupled root map has the compatible initial Real sign and the same continuation proof. The normalized square cell acts by its principal symbol one on these factor maps.

For arbitrary independently admitted X let $C_X : X^{\mathrm{can}} \to X$ be the actual source counit of the original Theorem 9.10, where X^{can} is obtained by its constructor after its decoder. Define $a_X = I_A(C_X)a_{X^{\mathrm{can}}}$ using its actual base-line comparison. For $f : X \to Y$ the original counit naturality cell compares fC_X and $C_Y f^{\mathrm{can}}$ with the same spin label. The normalized unit acts by one on degree zero, and the local-system factor frame is carried separately (original Definition H.1 and Proposition H.3). It follows that the root maps agree. A sign discrepancy would require changing the separately framed factor to obtain equal total module maps. Dividing by the common native residue-root comparison gives equality of the two I_A maps. Applying them to the initialized unit proves $I_A(f)a_X = a_Y$. The same counit argument for coupled conjugation gives $I_A(J_X)\overline{a_X} = a_{\mathfrak{c}X}$, where $\mathfrak{c}X$ is the coupled conjugate. It also proves independence from the presentations used for the counit and compatibility on restrictions and finite diagrams. Thus the deformation used for a canonical presentation supplies the unit for an independent source by transport through its counit.

Finally $\mu^{\vee}_{r^2 g} = r^{-4}\mu^{\vee}_g$ and $q_{r^2 g} = r^{-2}q_g$. Their residue ratio gives the asserted square laws. The positive factor r^{-1} preserves the initialized branch. The two line comparisons multiply by r on R_g and by r^{-1} on A, with identity product on the fixed original root. □

Proposition 9.5 (The calibrated comparison and its scalar).
Let ν_X be the original nonzero frame of the local-system factor. Take the constructed interface of Proposition 9.2 or its comparison over a supplied quantization on the conditioned parents specified there. Over the resulting formal algebra map k_X there is a unique module interface F_X satisfying the relative framing condition

$$\nu_X \longmapsto a_X^{-1} \otimes n_{\mathrm{W},X}.$$

Here a_X^{-1} is the dual vector paired to one with a_X. If (α_f, Ξ_f) and $(\mathcal{A}_f, T_f)$ are the analytic and wave representatives of a physical arrow $f : X \to Y$ in the same direction, there is a unique normalized unit $U_f \in 1 + F_{-1}\mathscr{E}_Y$ with $U_f^* U_f = 1$ and

$$k_Y \alpha_f = \mathrm{Ad}(U_f)\mathcal{A}_f k_X, \qquad F_Y \Xi_f = U_f T_f F_X.$$

Proof. The analytic physical comparison of the original G.5 acts on the real ray cone by the lift of the physical isometry. Proposition 9.2 therefore extends it in adapted coordinates; for a supplied representative on a bounded diagram we use H.1–H.2 as in that proposition. Both sides of the calibration belong to local-system factors relative to the actual original root. The proposed image is nonzero, horizontally constant in the cone directions and defined for the full family. The Hom line in Proposition 9.2 therefore supplies the unique required scalar. The numerical phase of the native symbol in Proposition 9.3 and the original κ_R both remain in this rule. Reselecting an auxiliary radial section uses parallel transport of the original frame along the cone. Its value as a flat factor section is independent of the new radial coordinate.

A true change $\nu'_X = u\nu_X$ of the framed object scales its interface to $u^{-1}F_X$. Its induced connection in the reframed presentation is $u\nabla^F u^{-1}$, which adds the scalar one-form $-u^{-1}\mathrm{d}u$ to its coefficient.

The two algebra maps have the same full homogeneous map on the retained covectors. Lemma 7.8 gives U_f. After this correction both module maps have the same semilinearity, so their ratio is a point-parameter scalar by Lemma 7.9. The original map carries ν_X to ν_Y, Proposition 9.4 carries the dual of a_X to the dual of a_Y, and Proposition 9.3 carries $n_{\mathrm{W},X}$ to $n_{\mathrm{W},Y}$. The unit has principal symbol one. Both module maps thus have the same value on the prescribed nonzero factor frame, proving that their ratio is exactly one. □

9.3. Real and physical comparisons.

Proposition 9.6 (Native conjugation of the complete evaluation). On a compatible real presentation, coefficient conjugation defines maps $\kappa_X : \overline{\mathscr{E}_X} \to \mathscr{E}_{\mathfrak{c}X}$ and $j^0_X : \overline{\mathscr{M}_X} \to \mathscr{M}_{\mathfrak{c}X}$ with strict square. They commute with the point connections and carrier comparisons. The evaluation has the law $j^0_X(\overline{e}_X) = -e_{\mathfrak{c}X}$ at every homogeneous degree.

Proof. In Proposition 6.2 the rephased product and internal adjoint have real coefficients. Proposition 6.3 gives real formulas for the module action and radial Jacobian. In a real graph chart $\nabla^{\mathrm{amb}}_Y(au_s) = (Ya - \tau(Ys)a)u_s$ for real Y. These formulas prove the algebra, square and differential assertions. Real coordinate pushforwards obey these formulas by the real half-density substitution and its Taylor expansion. Along a spacelike isotopy the normalized generator of Lemma 6.13 is formally real. After flattening the real moving graph its recursive coefficient transports have real coefficients and real identity initial values. Uniqueness at each depth, followed by the carrier cocycle, proves commutation with every retained T_{ab}. Complex stored charts carry the conjugates of these maps through their normalizations.

On a carrier through the evaluated point b, the exact column is $e(b) = C^{1/2}(\rho(b)^{-1}\delta_b)$. With $\mathcal{F} = Q - Q^*$, §6 gives $C = i\mathcal{F}^{-1}$ and $C^{1/2} = e^{i\pi/4}\mathcal{F}^{-1/2}$ in the formal convention. The coefficients of the latter inverse root before its constant phase are real by recursion. Its physical left symbol has degree $-1/2-\ell$ term $i^{-\ell}b_\ell$ with b_ℓ real, by the half-integer rephasing. Taylor expand at b and move each $(x-b)^\beta$ to the frequency coefficient by integration by parts. At degree $-1/2-k$ the physical column has phase i^{-k} times the real sum

$$\sum_{\ell+|\beta|=k} \frac{(-1)^{|\beta|}}{\beta!}\partial_\xi^\beta \partial_x^\beta b_\ell(b,\xi)$$

times the real density and Fourier factors. Indeed integration by parts gives $i^{|\beta|-\ell} = i^{-k}(-1)^{|\beta|}$. This is module degree $1-k$, so its formal coefficient acquires the factor $i^{-(1-k)}i^{-k} = -i$. Every coefficient transforms with the odd sign. Transport from a disk through each point gives the same identity for the full joint family on a fixed carrier. The principal-symbol formula

makes n_{W} Real in its compatible root presentation. Formal conjugation acts on the coefficients over the same ray, whereas conjugation of the physical distribution reverses the frequency. □

Theorem 9.7 (Structured comparison on the analytic locus). The independently admitted analytic Lorentz operator sources at $\Lambda = 0$ and the independently admitted wave sources with analytic physical germ have a natural comparison after sectorwise real restriction and complete formal enrichment. It has both inverse comparisons and preserves the two chiral operations, the actual root and its square and base factors, both associators, physical and optical scale, the finite-rank diagrams and the actual spin label. It compares the full evaluation and all point derivatives on every complete module section, carrying the original coupled operation with its own square cell. The calibrated interfaces over a fixed physical identification, with their simultaneous presentation changes, form a contractible groupoid locally on point-parameter charts. Every stipulated finite coefficient and derivative has a finite bound on common conditioned parents.

Proof. Work in the direction of the actual geometric arrow, using the inverse of an original covariant source instruction where necessary as in original Theorem G.5. For a physical comparison let $P_f = k_Y^{-1}(U_f)$ in the real smooth formal algebra of Y from Proposition 9.2, and put

$$\begin{aligned} \alpha_f^{\mathrm{W}} &= \mathrm{Ad}(P_f^{-1})\alpha_f, \\ \Xi_f^{\mathrm{W}} &= P_f^{-1}\Xi_f, \\ e_X^F &= F_X^{-1} e_X, \\ \nabla_X^F &= F_X^{-1}\nabla_X^{\mathrm{amb}} F_X. \end{aligned}$$

Proposition 9.5 gives $k_Y \alpha_f^{\mathrm{W}} = \mathcal{A}_f k_X$ and $F_Y \Xi_f^{\mathrm{W}} = T_f F_X$. Applying the latter first to e_X^F and then to an arbitrary section and its covariant derivative proves both full intertwining equations. The corresponding algebra derivative is transported through k_X. An operator $D_Y = \alpha_f(D_X)$ in either original chirality now has representative $P_f^{-1} D_Y P_f$. Under a varying gauge G, a connection coefficient B_Z goes to $GB_ZG^{-1} - (ZG)G^{-1}$. These equations transport every original coefficient. The first-unit affine torsor is preserved because commutators of a unit in $1 + F_{-1}$ with the first negative filtration piece vanish at first negative order. Proposition 7.5 identifies the canonical finite-rank comparison over the same spin and metric. The actual optical comparison is the one determined by the same physically selected scale.

For composable original representatives let $D_{g,f}$ denote their normalized composition cell, so $\Xi_g \Xi_f = D_{g,f}\Xi_{gf}$ and the algebra composite uses $\mathrm{Ad}(D_{g,f})$. Expanding the two squares with the exact wave composite gives

$$U_g \mathcal{A}_g(U_f) = k_Z(D_{g,f})\, U_{gf}, \qquad \alpha_g(P_f) P_g = D_{g,f} P_{gf}$$

for target Z. Equality of implementers follows from normalized-unit uniqueness. Substitution gives exact composition for the corrected maps, including

inverses and identities with their normalized cells. In stored charts put $T'_{ij} = P_i^{-1} T_{ij} \varphi^Y_{ij}(P_j)$, where T_{ij} is the original overlap unit. Moving P_j across φ^Y_{ij} and using the original overlap equations gives

$$\begin{aligned} \alpha'_i \varphi^X_{ij} &= \mathrm{Ad}(T'_{ij})\, \varphi^Y_{ij} \alpha'_j, \\ T'_{ij} \varphi^Y_{ij}(T'_{jk}) C^Y_{ijk} &= \alpha'_i(C^X_{ijk}) T'_{ik}, \\ \Xi'_i \tau^X_{ij} &= T'_{ij} \tau^Y_{ij} \Xi'_j. \end{aligned}$$

For example, the middle equality uses $\varphi^Y_{ij}\varphi^Y_{jk}(P_k) C^Y_{ijk} = C^Y_{ijk} \varphi^Y_{ik}(P_k)$ before applying the original comparison identity. This retains both original associators on every common refinement. If $k'_X = \mathrm{Ad}(V) k_X$, the same calibrated Hom computation gives $F'_X = V F_X$. The unique change on the analytic presentation is $(F'_X)^{-1} F_X$. When the carriers differ, the comparison inserts the canonical wave transport between the two interfaces. The multiplication and cancellation of these comparisons give identity, inverses and all finite coherences, proving the asserted contractibility over the fixed physical identification.

For the coupled operation write (J_X, j_X) for its original representatives with normalized square Z_X. Transport them to the wave presentation as $\widehat{J}_X = k_{\mathfrak{c}X} J_X \overline{k_X}^{-1}$ and $\widehat{j}_X = F_{\mathfrak{c}X} j_X \overline{F_X}^{-1}$. The original contact-valued normal evaluation embeds $\mathcal{L}^{-1}$ in $T^*_{S/B}$ and commutes with the physical cotangent comparison. Spin-factor exchange and conjugation fix the Hermitian real null covector, so there is no additional radial action on its injective image. The two real algebra comparisons have the same full homogeneous map. By Lemma 7.8 a unique normalized unit V_X gives $\widehat{J}_X = \mathrm{Ad}(V_X)\kappa_X$. The original root operation carries the initialized base unit by Proposition 9.4, and the factor n_{W} is Real by Proposition 9.6. The original map also carries the separately specified frame. Evaluating on the calibrated factor as in Proposition 9.5 fixes the residual module scalar, hence $\widehat{j}_X = V_X j^0_X$. The two squares then yield

$$V_{\mathfrak{c}X}\, \kappa_{\mathfrak{c}X}(\overline{V_X}) = k_X(Z_X), \qquad \widehat{j}_X(\overline{e}_X) = -V_X e_{\mathfrak{c}X}.$$

The first equation follows from normalized uniqueness on the algebra, and on the module the same unit acts by left multiplication. In fixed-carrier coordinates the connection defect is $(Y V_X) V_X^{-1} \widehat{j}_X$ for a real point direction Y, by differentiating $V_X j^0_X$ using Proposition 9.6. Equivalently one transports the target connection by V_X along with the target vector. In other charts the derivative is the transported covariant derivative of the unit. The odd marking on the conjugate native package is $-\overline{e}$, or equivalently *ie* provides a Real full vector. Two applications give the original square cell on the original representatives. All these comparisons retain both chiralities of the coupled operation.

For the two inverses, start with the source counit C_X of original Theorem 9.10. The actual root and framed-module part of this counit is what defined

a_X in Proposition 9.4. Apply the corrected representatives to that very arrow, obtaining an arrow that carries the added full fields and the original operators in their simultaneous presentations. Conversely take any independent wave source whose recognized metric admits analytic coordinates. Apply the original local correspondence to this analytic geometry. Its original calibrated source gives the interfaces of Propositions 9.2–9.5. Theorem 7.11 then compares the output over the chosen physical identification to the given wave source. These arrows use the full Hom on both independent domains. A smooth physical map in this restricted target and its specified lift are analytic by Lemma 9.1, so neither inverse discards physical arrows. Over a physical identity the composites in the unit and counit triangles agree with the original triangles after inserting the cells in the physical composition law, and the wave composites coincide by Theorem 8.3. The uniqueness of the normalized algebra implementers and of the framed module comparisons gives the representative cells for those equalities. The same argument with a physical arrow between two objects proves naturality. The whole exterior ordinary inputs used to present either geometry may be smooth.

For estimates retain the compact point, angular and parameter cores of the stipulated finite analytic diagram, with the root, inverse, phase and image margins in the original H.1. Its finite buffered refinement covers those cores. The original quantizers, tangent transports, counit maps and inverses have their stated analytic estimates on those parents. Finite real coefficient restrictions therefore have finite bounds. Derivatives of $s^2\lambda = \varepsilon_0$ involve finitely many derivatives of λ and divisions controlled by the initialized root margins. The inverse-frame margin controls normalization of the Hom scalar. The wave transports on a smooth displaced carrier have the bounds of Proposition 7.12. At each negative degree the Euler recursion of Lemma 7.8, composition, inversion and the differentiated gauge equations use only finitely many earlier coefficients and finite jets. This proves every specified bound on the retained parents. The new comparisons and full evaluations live in the complete formal category after real restriction. The original analytic inputs and their analytic comparisons carry the factorial estimates. □

9.4. The complete local correspondence.

Theorem 9.8 (The two-class atlas with its analytic comparison). For smooth pointed oriented and time-oriented Lorentzian Ricci-flat four-germs with actual spin, the execution and recognition functors of Definition 8.1 give the two equivalences in Theorem 1.1. Every independent source on a wave neighborhood of Definition 6.12 and every physical germ has a display from the two ordinary classes on the selected whole-input domain, with numerical scale restored. Physical isometries and their specified spin lifts are in bijection with complete local comparisons, including changes of initial surface and scalar carrier. The inverse returns the class of local displays with its natural identifications. The common-parent finite-family bounds and effective descent hold under the hypotheses of Proposition 8.4

and Theorem 8.5. On the analytic geometric locus, using the v1.1 operator correspondence (Isthmus 2026a, Theorem 9.10) and the analytic inputs of Section 9, these identifications carry the calibrated analytic operator data, after real restriction and complete formal enrichment, through the two inverse comparisons of Theorem 9.7, including its root, associator and coupled Real maps.

Proof. First consider whole inputs and their physical restrictions. Theorem 3.13 identifies the ordinary classes with maximal data in a fixed spatial gauge. Theorem 4.18 constructs the actual gradients from the classes by the curvature response and recovers the same whole input. It uses the negative and positive class operations before comparing their output with that tensor chart. Theorem 2.9 develops and recognizes its finite Cartan fields. Theorem 5.7 constructs a display of a given physical germ by choosing a maximal disk through its point, normalizing at one finite order, applying the protected vacuum extension of Appendix B and decoding the resulting whole data. Physical Cauchy uniqueness gives an actual local embedding with the restored metric scale and spin lift.

On *local presentations*, a different inverse uses these embeddings. The independent wave tests determine the scalar kernel and the derived carrier fields by Theorem 6.24. The geometry of the skies and the density recognize the isometry by Theorem 7.3. The full lift in Theorem 7.11 uses an intermediate disk inside its domain and tests the whole open family of moving skies. Theorem 8.3 uses that Hom bijection to prove the two natural inverse identifications and their triangles. Hence the inverse of a physical germ recovers its local comparison class, even when its representatives have different whole exterior inputs. Proposition 8.4 gives the family quantifiers and Theorem 8.5 proves effective descent.

Finally, on the analytic locus Propositions 9.3 and 9.4 identify the actual relative quadratic line and its initialized unit through the original counit. They fix the scalar of the complete module comparison. Theorem 9.7 constructs the corrected physical and Real representatives, proves their two inverse comparisons and gives the stated finite estimates. □

Conclusion

In the chosen gauge the two ordinary classes determine a whole initial-data representative. Different extensions of the same disk can produce different pairs for one physical germ, so the inverse on germs takes values in a local comparison class. The full lift of a specified spin isometry makes this class independent of the initial disk and of the carrier used for the scalar wave kernel.

The smooth correspondence retains both chiralities at the physical metric scale. Its comparison with analytic twistor data uses real restriction and the complete formal symbol and module. For a smooth family, one common parent supports the estimates at every prescribed finite order. On the

analytic locus, the additional hypotheses on the analytic inputs supply the factorial estimates.

Computational supplement. The supplement guide describes the code, inputs and recorded results for the following checks. The archive smooth-twistor-germs-computations-v1.0.zip accompanies version 1.0 at https://github.com/lioristhmus/smooth-twistor-germs/releases/tag/v1.0.

- **Scalar and York inverses (§4).** At the nonaxial frequency $(1, 2, 2)$, both helicities give the same single-response multiplier. At that frequency the contractions after three responses in Theorem 4.12 agree with the direct scalar inverse and the full York inverse, including its longitudinal and transverse components.
- **Exterior momentum range (Lemma B.10).** For spherical harmonic degrees $l \geq 2$, the magnetic calculation derives the range identity of Lemma B.10 by differential elimination. Dropping the angular shift $+2$ changes the residual, and a divergence-free example lies outside the selected range. On the electric side the script substitutes one transcribed Hodge equation into the other. For the boundary profiles it prescribes division by r to pass to the physical frame and checks the resulting ratios at the same cutoff argument for two degrees. Appendix B proves the low-mode equations, $M_\delta J = I$, and the uniform estimates and higher regularity across all modes.
- **Wave transport (§6).** In flat spacetime, three boosted rays distinguish the degree-zero Doppler factor from the invariant degree-one normalization, and a Gaussian-coordinate mode distinguishes the two frequency branches by their opposite charges.
- **Analytic comparison (§9).** In a flat analytic example the analytic and wave frames transported by the same spin lift have ratio $+1$. In a nonlinear chart both the leading evaluation coefficient and its nonzero degree-zero term change sign under formal conjugation.

For the York inverse, the spectral reconstruction is compared with a separate free-space Green tensor quadrature. Initialized transport from the two chiral inputs gives a nonlinear far-field coefficient for comparison with a separate constraint iteration on the same grid. The specified right inverse in Appendix B extends a nonanalytic plane-wave slice while preserving its interior. In the local reconstruction, the direct band-96 tidal error is 0.20% of the rising-side signal. For a separate smooth band-32 representative, the straight and broken paths differ by 7.17×10^{-11} in K. The additional supplement guide accompanies paper3-smooth-twistor-additional-computations-2026-11-24-r2.zip.

Appendix A. Weighted construction and based gauge

This appendix proves the analytic operations used in the ordinary-pair initial-data constructor. It first supplies the whole decaying scalar and York inverses, then the nonlinear maximal constraint chart with its exact ambient

decoder. We obtain the higher weighted and parameter estimates for the solution selected at base order. A separate anchored inverse normalizes the spatial coordinates and retains their initialized actual spin comparison.

The residual estimates of LeFloch and Nguyen for their seed-to-solution map imply that the correction vanishes for an already exact whole-space seed (Theorem 2.7, (2.6)). Our chart uses transverse trace-free seeds and a based spatial gauge. Its exact decoder recovers the seeds that determine the ordinary classes, with finite-order dependence along the selected branch. Agreement with data prescribed only on an interior ball is the extension problem of Appendix B.

All small neighborhoods below are chosen using finitely many base norms. Every subsequent higher-order assertion has finite dependence on higher input seminorms, without a new smallness requirement at those orders.

A.1. Weighted spaces and the decaying flat inverses. The initial fields in this appendix are defined on the whole marked $\mathbb{R}^3$ in fixed Euclidean coordinates. Write δ for its Euclidean metric and $\langle x\rangle = (1+|x|^2)^{1/2}$. All tensor norms refer to Cartesian components. A negative weight denotes decay. The inverse used to decode the momentum is the decaying whole-space inverse below; the anchored sublinear inverse used for a change of spatial coordinates has a different domain. The longitudinal and transverse trace-free parts of the momentum decomposition follow York.

Definition A.1 (Weighted Hölder spaces).
For an integer $m \geq 0$, $0 < \alpha < 1$ and a real weight ω, set

$$\|u\|_{C^{m,\alpha}_{\omega}} = \|u\|_{C^{m,\alpha}(B_2)} + \sup_{R\geq 1} R^{-\omega}\|u(R\,\cdot)\|_{C^{m,\alpha}(\overline{A})},$$
$$A = \{x : 1 < |x| < 2\}.$$

The resulting Banach space is denoted by $C^{m,\alpha}_{\omega}$. Tensor-valued spaces are defined componentwise. Taking dyadic radii or using slightly larger fixed annuli gives equivalent norms. Differentiation maps $C^{m,\alpha}_{\omega}$ into $C^{m-1,\alpha}_{\omega-1}$, and

$$\|uv\|_{C^{m,\alpha}_{\omega+\nu}} \leq C_{m,\alpha}\|u\|_{C^{m,\alpha}_{\omega}}\|v\|_{C^{m,\alpha}_{\nu}}.$$

For different differentiability orders, the product estimate holds at the smaller order. If $\omega \leq \nu$, then $C^{m,\alpha}_{\omega} \subset C^{m,\alpha}_{\nu}$ continuously. These assertions follow on the inner ball and on each rescaled annulus from the ordinary Hölder estimates, followed by the indicated supremum.

We write $\mathscr{S}^{\omega}$ for smooth fields with

$$\sup_x \langle x\rangle^{-\omega+|\beta|}|\partial^{\beta}u(x)| < \infty \quad \text{for every multi-index } \beta.$$

For a fixed α, this is the intersection of the spaces $C^{m,\alpha}_{\omega}$ over all finite m, with equivalent seminorm systems. One direction is immediate from the weighted Hölder norm. For the other, one additional derivative bounds the Hölder seminorm on each fixed rescaled annulus. Thus passing between the

symbol notation and all weighted Hölder orders uses finitely many extra seminorms.

Define

$$(L_0W)_{ij} = \partial_i W_j + \partial_j W_i - \frac{2}{3}\delta_{ij}\partial_a W_a,$$

$$\Delta_L W = \operatorname{div}_\delta L_0 W = \Delta W + \frac{1}{3}\nabla \operatorname{div} W.$$

The interior estimate. We prove the interior Schauder estimate needed for our second-order systems. Douglis and Nirenberg describe the extension of interior Schauder estimates to elliptic systems (Douglis and Nirenberg 1955, p. 503). Here every row has weight zero and every unknown has weight two. Write $P = a^{ij}(x)\partial_i\partial_j + b^i(x)\partial_i + c(x)$ with matrix coefficients on $A' \subset \mathbb{R}^3$. Choose the sign of each row as in the applications below. A uniform inverse bound for the principal matrix $\sum a^{ij}(x)\xi_i\xi_j$ on $|\xi| = 1$ suffices. For N components and bounded coefficients this is equivalent to $|\det(\sum a^{ij}(x)\xi_i\xi_j)| \geq c|\xi|^{2N}$, since the other singular values have a uniform upper bound on the unit sphere. Strong ellipticity implies this bound. On fixed nested bounded domains $A \Subset A'$, for $0 < \alpha < 1$, j a nonnegative integer, and coefficients bounded in $C^{j,\alpha}(A')$, the estimate is

$$\|v\|_{C^{j+2,\alpha}(A)} \leq C\big(\|Pv\|_{C^{j,\alpha}(A')} + \|v\|_{C^0(A')}\big).$$

For variable coefficients we start with an existing $C^{2,\alpha}_{\mathrm{loc}}$ solution. A locally bounded distributional solution is sufficient for the constant-coefficient applications, with the estimate on nested precompact sets. The constant depends on the domains, α, j, the inverse-symbol margin and the indicated coefficient bounds. For the strongly elliptic variable systems below, the interior regularity conclusion also follows from Giaquinta and Martinazzi (2012, Theorem 5.22). At each bootstrap step the known solution belongs to $W^{2,2}_{\mathrm{loc}}$ and the terms $b^i\partial_i v + cv$ are included in the forcing at the order used. For Δ and Δ_L the coefficients are constant. The positive symbol of $-\Delta_L$ is $|\xi|^2 I + \frac{1}{3}\xi\otimes\xi$, with eigenvalues $|\xi|^2$ and $\frac{4}{3}|\xi|^2$.

Here is a derivation that also explains which regularity is used in passing to higher order on the same branch. Freeze the principal coefficients at x_0 and call the resulting homogeneous operator P_0. Take a smooth inhomogeneous dyadic Fourier partition, with low block S_0 and blocks Δ_l, $l \geq 0$, supported in fixed dilates of a compact annulus. For nonintegral $s > 0$ the norm on bounded functions on $\mathbb{R}^3$ given by

$$\|S_0 v\|_\infty + \sup_{l\geq 0} 2^{ls}\|\Delta_l v\|_\infty$$

is equivalent to $C^{\lfloor s\rfloor, s-\lfloor s\rfloor}$. Indeed, the annular kernel has vanishing moments. Subtracting the Taylor polynomial of degree $\lfloor s\rfloor$ bounds its convolution by $C2^{-ls}\|v\|_{C^s}$. Conversely derivatives of order k of a block cost $C_k 2^{lk}$ by convolution with a slightly larger frequency cutoff. Sum the blocks through $k = \lfloor s\rfloor$. For the top derivative split a difference at $2^l|h| = 1$, using one

more derivative below that scale and twice the supremum above it. The two geometric series give $C|h|^{s-k}$. This proof also constructs the indicated representative of a distribution from its blocks.

The symbol matrix of P_0 is invertible on the unit frequency annulus. There, its inverse times an enlarged smooth cutoff has bounded derivatives to any prescribed finite order, by matrix inversion and the stated margins. Its inverse Fourier transform is integrable with a uniform bound, by integration by parts. Scaling the annulus shows

$$\|\Delta_l v\|_\infty \le C 2^{-2l} \|\widetilde{\Delta}_l P_0 v\|_\infty$$

with $\widetilde{\Delta}_l$ an enlarged block equal to one on the smaller support. The low block is bounded by $\|v\|_\infty$. At $s = 2 + \alpha$ this proves the full-space estimate for P_0 on compactly supported fields.

For the local estimate choose $A \Subset D \Subset A'$ and put $d(x) = \min(r_0, \operatorname{dist}(x, \partial D))$, $t = d(x)/8$. In coordinates $x + ty$, $|y| < 4$, the $C^{0,\alpha}$ norm of $a(x+ty)-a(x)$ is bounded by $C r_0^\alpha [a]_\alpha$. Choose r_0 once to make this coefficient small compared with the frozen estimate. Apply that estimate to a cutoff of $v(x + t\,\cdot)$ equal to one on B_1 and supported in B_2. An extension of the coefficient difference by a second cutoff inside B_3 has the same small bound. Its product with the second derivatives of the cut-off unknown is absorbed on the left. All cutoff commutators and lower terms are bounded by a constant times $\|v(x + t\,\cdot)\|_{C^{1,\alpha}(B_3)}$, besides the rescaled forcing $t^2(Pv)(x + t\,\cdot)$. The lower coefficients receive factors t, t^2.

Let M be the supremum over $x \in D$ of the rescaled $C^{2,\alpha}(B_1)$ norm. It is finite for the existing solution since $\overline{D} \Subset A'$. On $B(x, 3t)$ the radii $d(y)/8$ are comparable to t. A bounded number of such balls, using smaller concentric subballs for the cover, bounds the rescaled $C^{2,\alpha}(B_3)$ norm by CM. Pairs of points farther apart than a fixed fraction of the unit radius use the supremum bound in the Hölder quotient. On B_3 interpolation bounds the preceding $C^{1,\alpha}$ norm by $\delta CM + C_\delta \|v\|_\infty$ for each $\delta > 0$, by extension from the fixed ball and a split into low and high frequency blocks at a sufficiently large index. Taking the supremum and choosing δ small absorbs M on the right. On A the chosen radii have a positive lower bound. This proves the displayed interior estimate for $j = 0$.

For $j \ge 1$ use nested domains and induction. If $Pv = f \in C^{j,\alpha}$ and v is already $C^{j+1,\alpha}$ locally, differentiate the equation $j - 1$ times and take a first spatial difference quotient. The resulting unknown is $D_h \partial^\beta v$, $|\beta| = j - 1$, with $D_h w = (w(\,\cdot + h e_i) - w)/h$. For each fixed nonzero h this unknown has the base regularity. Its principal matrix is a translate of a or a itself, according to the product rule used. Every commutator term of highest order has a differentiated coefficient and at most $j + 1$ derivatives of v. Difference quotients of differentiated coefficients and of f have uniform $C^{0,\alpha}$ bounds from their next derivatives. Applying the base estimate bounds the unknown uniformly in $C^{2,\alpha}$ on a smaller domain, including its bounded supremum. Compactness in C^2 on smaller sets and passage of the Hölder

inequality to the limit give the bound for $\partial_i \partial^\beta v$ in $C^{2,\alpha}$. The preceding induction bounds its lower-order terms by the right side of the displayed estimate on a larger domain. For each finite j choose the intermediate domains between the same two endpoints. This gives the displayed bound on the prescribed inner domain. For a bounded distributional solution of a constant-coefficient equation, local convolution on a slightly smaller domain commutes with P_0. Applying the estimates to the mollifications and passing to their distributional limit gives the same conclusion. This covers the flat potentials in Lemmas A.2 and A.8. The nonlinear applications in Theorems A.7, A.9 and A.14 and the intervening linearized or averaged systems all start above the base regularity. Their rescaled principal symbols and coefficient bounds are checked where those equations are formed.

Lemma A.2 (The decaying scalar and York inverses).
For $m \geq 2$, $0 < \alpha < 1$ and $-1 < \omega < 0$, the maps

$$\Delta : C^{m,\alpha}_{\omega} \longrightarrow C^{m-2,\alpha}_{\omega-2}, \qquad \Delta_L : C^{m,\alpha}_{\omega}(T\mathbb{R}^3) \longrightarrow C^{m-2,\alpha}_{\omega-2}(T\mathbb{R}^3)$$

are bounded isomorphisms. Their inverses are convolution with

$$\Phi(x) = -\frac{1}{4\pi|x|}, \qquad G_{ij}(x) = -\frac{1}{32\pi}\left(\frac{7\delta_{ij}}{|x|} + \frac{x_i x_j}{|x|^3}\right),$$

respectively. They are the same operators at every finite order and satisfy

$$\|\Delta^{-1} f\|_{C^{m,\alpha}_{\omega}} + \|\Delta_L^{-1} F\|_{C^{m,\alpha}_{\omega}} \leq C_{m,\alpha,\omega}\big(\|f\|_{C^{m-2,\alpha}_{\omega-2}} + \|F\|_{C^{m-2,\alpha}_{\omega-2}}\big).$$

Proof. The differential maps are bounded by the weighted differentiation rule. Distributionally, $\Delta\Phi = \delta_0$. Put $\Psi(x) = -|x|/(8\pi)$, so that $\Delta\Psi = \Phi$. Then

$$G_{ij} = \delta_{ij}\Phi - \frac{1}{4}\partial_i\partial_j\Psi, \qquad \partial_i G_{ij} = \frac{3}{4}\partial_j\Phi.$$

Applying Δ_L cancels the two terms involving $\partial_i\partial_j\Phi$ and gives $(\Delta_L G)_{ij} = \delta_{ij}\delta_0$. Differentiating $|x|$ gives the displayed matrix kernel. Both kernels are bounded by $C|x|^{-1}$.

Suppose $|f(y)| \leq B\langle y\rangle^{\omega-2}$. The convolution is locally integrable at $y = x$. At infinity its absolute integrand, including the volume element, is bounded by a multiple of $r^{\omega-1}\,\mathrm{d}r$, which is integrable because $\omega < 0$. This proves absolute convergence, also componentwise for $G * F$. The distributional kernel identity gives the equation; testing against a compactly supported smooth function and using the same absolute bounds justifies the interchange of integrals.

For the weighted zeroth-order estimate let $R = |x| \geq 2$. On $|y| < R/2$,

$$\frac{1}{R}\int_{|y|<R/2} \langle y\rangle^{\omega-2}\,\mathrm{d}y \leq C_\omega R^\omega,$$

where $\omega > -1$ is used. On $|x - y| < R/2$, one has $\langle y\rangle \asymp R$, and

$$R^{\omega-2}\int_{|z|<R/2} |z|^{-1}\,\mathrm{d}z \leq CR^\omega.$$

On the remaining region, both $|y|$ and $|x-y|$ are at least $R/2$. The substitution $y = Rz$ gives R^ω times an integral bounded uniformly for $x/R \in S^2$; its tail is integrable because $\omega < 0$. The bounded region $|x| \le 2$ follows from local integrability and the same tail estimate. Thus

$$\|\Phi * f\|_{C^0_\omega} + \|G * F\|_{C^0_\omega} \le C_\omega(\|f\|_{C^0_{\omega-2}} + \|F\|_{C^0_{\omega-2}}).$$

Splitting off a small neighborhood of the kernel singularity proves continuity of the convolutions. The interior estimate gives their local higher regularity. On nested fixed annuli use $v_R(z) = v(Rz)$ and $Pv_R = R^2 f(Rz)$. The zeroth-order bound and the source norm give

$$\|v_R\|_{C^{m,\alpha}(\overline{A})} \le CR^\omega \|f\|_{C^{m-2,\alpha}_{\omega-2}}.$$

An interior-ball estimate controls the remaining part of the weighted norm. The same calculation applies to the vector system, proving existence and the claimed bounds.

A harmonic scalar in $C^{m,\alpha}_\omega$ tends to zero at infinity, hence is zero by the maximum principle. If $\Delta_L W = 0$, then taking the divergence gives $\frac{4}{3}\Delta \operatorname{div} W = 0$ distributionally. Local harmonic regularity and its decay imply $\operatorname{div} W = 0$. Hence $\Delta W = 0$, and each component tends to zero. Thus $W = 0$. Uniqueness proves that the inverse formula and all its finite-order realizations agree. □

No integral or moment condition is imposed on the residual in this lemma. For $-1 < \omega < 0$ a $1/|x|$ tail has the allowed decay at infinity; the solution itself is smooth on the whole space. Constants and other nondecaying homogeneous solutions are excluded. The restriction $\gamma < -1/2$ used below is compatible with these tails.

A.2. The two seed coordinates and the independent target. Fix once and for all

$$k_0 \ge 4, \qquad 0 < \alpha < 1, \qquad -1 < \gamma < -\frac{1}{2}.$$

The base index can be increased before choosing the construction domain when a later finite-order operation requires it.

Definition A.3 (Seed spaces and the maximal target in spatial gauge).

For $m \ge k_0$, let

$$\begin{aligned}
\mathcal{Q}_m &= \{q \in C^{m,\alpha}_\gamma(\operatorname{Sym}^2 T^*\mathbb{R}^3) : \operatorname{tr}_\delta q = 0,\ \operatorname{div}_\delta q = 0\},\\
\mathcal{P}_{m-1} &= \{p \in C^{m-1,\alpha}_{\gamma-1}(\operatorname{Sym}^2 T^*\mathbb{R}^3) : \operatorname{tr}_\delta p = 0,\ \operatorname{div}_\delta p = 0\},\\
\mathcal{X}_m &= \mathcal{Q}_m \oplus \mathcal{P}_{m-1},\\
\mathcal{Y}_m &= C^{m,\alpha}_\gamma \oplus C^{m,\alpha}_\gamma(T\mathbb{R}^3),\\
\mathcal{Z}_m &= C^{m-2,\alpha}_{\gamma-2} \oplus C^{m-2,\alpha}_{\gamma-2}(T^*\mathbb{R}^3).
\end{aligned}$$

The variables are $s = (q, p) \in \mathcal{X}_m$ and $y = (u, W) \in \mathcal{Y}_m$. Give finite direct sums their sum norms.

The ambient data domain consists of positive whole metrics and covariant symmetric tensors satisfying

$$h-\delta\in C^{m,\alpha}_{\gamma},\qquad K\in C^{m-1,\alpha}_{\gamma-1},\qquad \frac{1}{2}\delta<h<\frac{3}{2}\delta,$$
$$\operatorname{div}_\delta(\mathrm{TF}_\delta\, h)=0,\qquad \operatorname{tr}_h K=0,$$

where $\mathrm{TF}_h T=T-\frac{1}{3}(\operatorname{tr}_h T)h$. Its maximal vacuum subset is defined independently by

$$\mathcal{H}(h,K)=R(h)-|K|_h^2=0,\qquad \mathcal{M}(h,K)=\operatorname{div}_h K=0.$$

These conventions agree with $g=\mathrm{d}t^2-h_t$ and $K=\frac{1}{2}\partial_t h_t$. The target conditions contain no requirement that a datum already be generated by a seed or a cohomology class.

Proposition A.4 (Ambient coordinates and exact decoder). On a neighborhood of zero the map

$$\mathcal{A}(q,p,u,W)=(h,K),\qquad h=\delta+q+u\delta,\qquad \kappa=p+L_0W,\qquad K=\mathrm{TF}_h\,\kappa$$

gives smooth local coordinates on the ambient gauge domain. Its inverse is

$$\begin{aligned} q&=\mathrm{TF}_\delta\, h, & u&=\frac{1}{3}\operatorname{tr}_\delta h-1,\\ \kappa&=K-\frac{\operatorname{tr}_\delta K}{\operatorname{tr}_\delta h}h, & W&=\Delta_L^{-1}\operatorname{div}_\delta\kappa,\\ p&=\kappa-L_0W. \end{aligned}$$

This inverse is defined before imposing the vacuum constraints. The second seed is the flat transverse-traceless part of the displayed κ.

Proof. The tensors p and L_0W are Euclidean tracefree, so $\operatorname{tr}_\delta\kappa=0$. The metric formula gives $\mathrm{TF}_\delta\, h=q$ and its required divergence condition. The definition of K gives its h-tracefreeness.

Conversely, the first two decoder formulas return a Euclidean transverse-traceless q of weight γ and a scalar u of that weight. Put $c=(\operatorname{tr}_\delta K)/(\operatorname{tr}_\delta h)$. The denominator is uniformly positive on the metric neighborhood. Then $\kappa=K-ch$ is Euclidean tracefree and, using $\operatorname{tr}_h K=0$, has $\operatorname{tr}_h\kappa=-3c$. Consequently $\mathrm{TF}_h\,\kappa=K$.

Weighted multiplication gives $\kappa\in C^{m-1,\alpha}_{\gamma-1}$, so its Euclidean divergence belongs to $C^{m-2,\alpha}_{\gamma-2}$. Lemma A.2 gives $W\in C^{m,\alpha}_{\gamma}$. Thus $p=\kappa-L_0W$ has the required weight and regularity, and

$$\operatorname{tr}_\delta p=0,\qquad \operatorname{div}_\delta p=\operatorname{div}_\delta\kappa-\Delta_L W=0.$$

Starting from Euclidean tracefree κ and forming $K=\mathrm{TF}_h\,\kappa$, one obtains

$$\frac{\operatorname{tr}_\delta K}{\operatorname{tr}_\delta h}=-\frac{1}{3}\operatorname{tr}_h\kappa.$$

The decoder therefore returns the same κ. It recovers W by the uniqueness of the decaying York inverse, then p. It plainly recovers q and u. The opposite composition is the identity by the preceding construction.

Matrix inversion near δ, the scalar division above and tracefree projection are smooth weighted algebraic operations. Lemma A.2 is bounded at the specified orders. Both coordinate maps are therefore smooth. The assertions hold at every finite order with the same formulas. □

A.3. The selected nonlinear constraint chart.

Lemma A.5 (Constraint map and quadratic remainder).
The map

$$\mathcal{F}(s,y) = \big(R(h) - |K|_h^2, \operatorname{div}_h K\big), \qquad (h,K) = \mathcal{A}(s,y),$$

is smooth from a neighborhood of zero in $\mathcal{X}_m \oplus \mathcal{Y}_m$ to $\mathcal{Z}_m$. With

$$A_0(u,W) = (-2\Delta u, \Delta_L W),$$

it has the decomposition

$$\mathcal{F}(s,y) = A_0 y + \mathcal{N}(s,y).$$

The two components of $\mathcal{N}$ take values in $C^{m-2,\alpha}_{2\gamma-2}$, vanish to second order at the origin, and obey a quadratic norm bound on a smaller neighborhood. In particular,

$$D_y\mathcal{F}(0,0) = A_0, \qquad D_s\mathcal{F}(0,0) = 0.$$

Proof. For a covariant perturbation r,

$$DR_\delta(r) = \partial_i\partial_j r_{ij} - \Delta \operatorname{tr}_\delta r.$$

The transverse-traceless conditions give $DR_\delta(q) = 0$, while $DR_\delta(u\delta) = -2\Delta u$. Writing $r = q + u\delta$, the nonlinear scalar-curvature remainder is a finite sum of smooth contractions of the forms

$$O(r)\,\partial^2 r, \qquad O(1)\,\partial r\,\partial r.$$

Here the coefficients are smooth matrix functions with uniformly bounded derivatives on a retained positive metric set. Both terms have weight $2\gamma - 2$ and regularity $C^{m-2,\alpha}$.

Since $\operatorname{tr}_\delta \kappa = 0$,

$$K - \kappa = -\frac{1}{3}\big((h^{-1} - \delta^{-1}) : \kappa\big)h.$$

It is quadratic and belongs to $C^{m-1,\alpha}_{2\gamma-1}$. The term $|K|_h^2$ belongs to $C^{m-2,\alpha}_{2\gamma-2}$. The linear momentum term is $\operatorname{div}_\delta \kappa = \Delta_L W$. The other momentum terms have the forms $O(r)\partial\kappa$ and $O(1)\partial r\,\kappa$, hence the same quadratic order, weight and required regularity.

The weighted product estimate and smooth matrix inversion prove smoothness into the stated faster-decay spaces and give

$$\|\mathcal{N}(s,y)\|_{C^{m-2,\alpha}_{2\gamma-2}\oplus C^{m-2,\alpha}_{2\gamma-2}} \le C_m(\|s\|_{\mathcal{X}_m} + \|y\|_{\mathcal{Y}_m})^2.$$

Since $\gamma < 0$, the faster target embeds into $\mathcal{Z}_m$. The derivative formulas follow from the displayed linear terms. □

Theorem A.6 (The base constraint chart).
At the fixed base order k_0, choose sufficiently small neighborhoods $\mathcal{U} \subset \mathcal{X}_{k_0}$ and $\mathcal{V} \subset \mathcal{Y}_{k_0}$ of zero. There is a unique smooth correction

$$\mathcal{C} : \mathcal{U} \longrightarrow \mathcal{V}, \qquad \mathcal{F}(s, \mathcal{C}(s)) = 0, \qquad \mathcal{C}(0) = 0,$$

such that the constraint solutions in $\mathcal{U} \times \mathcal{V}$ are precisely its graph. The map

$$\mathcal{T}(s) = \mathcal{A}(s, \mathcal{C}(s))$$

is a smooth bijection onto a relatively open neighborhood of $(\delta, 0)$ in the independent constrained target, with smooth inverse given by the seed part of Proposition A.4. Moreover,

$$D\mathcal{C}(0) = 0, \qquad \|\mathcal{C}(s)\|_{\mathcal{Y}_{k_0}} \leq C\|s\|^2_{\mathcal{X}_{k_0}}, \qquad D\mathcal{T}(0)[q, p] = (q, p).$$

The neighborhoods can be chosen as convex base neighborhoods with a fixed transverse inverse margin: throughout their product,

$$\|D_y(A_0^{-1}\mathcal{N})(s, y)\|_{\mathcal{L}(\mathcal{Y}_{k_0})} \leq \kappa < \frac{1}{2}.$$

Every resulting metric is complete and asymptotically flat in the stated coordinates.

Proof. Lemma A.2 makes $A_0 : \mathcal{Y}_{k_0} \to \mathcal{Z}_{k_0}$ a bounded isomorphism. Lemma A.5 and the Banach implicit function theorem give the selected graph. Since the derivative of $\mathcal{N}$ vanishes at zero, continuity permits a convex product neighborhood on which the last displayed operator bound holds, while preserving the graph and metric positivity. Equivalently, the selected correction is the zero-initialized fixed point of

$$y \longmapsto -A_0^{-1}\mathcal{N}(s, y)$$

on a sufficiently small invariant ball. The product neighborhood and ball are chosen once at the base order.

Differentiating $\mathcal{F}(s, \mathcal{C}(s)) = 0$ at zero gives $A_0 D\mathcal{C}(0) = 0$. Hence $D\mathcal{C}(0) = 0$, and a bounded second derivative on a smaller neighborhood gives the quadratic estimate. The ambient formula then gives the derivative of $\mathcal{T}$.

Proposition A.4 supplies independent ambient coordinates with their exact inverse. In those coordinates the constraint subset in the chosen open product is the graph of $\mathcal{C}$. Its projection to the seed coordinates is therefore a smooth bijection onto $\mathcal{U}$ with the stated inverse. This proves both directions using the independently imposed gauge, maximality and vacuum equations.

The fixed metric bounds make the Riemannian distance uniformly comparable with Euclidean distance, so completeness follows. The weighted estimates give $h \to \delta$ and $K \to 0$ with their specified derivatives. No condition on the higher norms has been used in choosing this base chart. □

A.4. All weighted orders on the same branch. Let $\mathcal{X}_\gamma^\infty = \bigcap_{m \geq k_0} \mathcal{X}_m$ and $\mathcal{Y}_\gamma^\infty = \bigcap_{m \geq k_0} \mathcal{Y}_m$, with their indicated seminorms. Their smooth elements are the corresponding symbol-decay fields. The selected domain is $\mathcal{U} \cap \mathcal{X}_\gamma^\infty$: smallness is imposed only through $\mathcal{U}$ at the base order.

A smooth finite-dimensional family means a map from a finite-dimensional parameter manifold into this intersection whose realization in every $\mathcal{X}_m$ is smooth. On each retained compact parameter set, every specified finite parameter derivative therefore has finite $\mathcal{X}_m$ norm. These bounds are required separately at each finite order; no single bound over all orders is assumed.

Theorem A.7 (Weighted smoothness, continuity and finite families).

For every $s \in \mathcal{U} \cap \mathcal{X}_\gamma^\infty$, the same correction $\mathcal{C}(s)$ of Theorem A.6 belongs to $\mathcal{Y}_\gamma^\infty$. Its metric and momentum satisfy

$$h - \delta \in \mathscr{S}^\gamma, \qquad K \in \mathscr{S}^{\gamma-1}.$$

At each finite output order the norms are bounded by finitely many input seminorms and the fixed base margins. Higher input norms need only be finite. The maps $\mathcal{C}, \mathcal{T}$ and the ambient seed decoder are continuous in these full seminorm topologies.

For every supplied smooth finite-dimensional family in the selected domain, these maps are smooth into each finite weighted space. On a compact parameter set contained in the interior of a larger retained compact set, each specified finite parameter derivative at a specified spatial order is bounded by finitely many input spatial and parameter seminorms on that larger set and the fixed margins. The selected base domain precedes the choice of these orders.

Proof. We first establish the spatial improvement of the existing base correction, then the estimates for its transverse linearization.

Write $\mathcal{C}(s) = (u, W)$ and $h = \delta + q + u\delta$. The coefficient of the second derivatives of u in the Hamiltonian equation is

$$a^{ij}(h) = (h^{-1}\delta h^{-1})^{ij} - (\mathrm{tr}_h\, \delta) h^{ij}.$$

Indeed these are the coefficients of $h^{ia}h^{jb}\partial_i\partial_j r_{ab} - h^{ij}h^{ab}\partial_i\partial_j r_{ab}$ when $r = u\delta$. At the flat metric $a^{ij} = -2\delta^{ij}$, so its negative stays uniformly positive under the base bounds.

For the momentum equation, retain the second derivatives of W in $\mathrm{div}_h\, \mathrm{TF}_h(L_0 W)$. Its coefficients depend algebraically on h and are uniformly close to those of Δ_L. The Hamiltonian equation contains W only through first derivatives, and the momentum equation contains u only through first derivatives. Thus, after changing the sign of the scalar row, the second-order principal part in (u, W) is block diagonal and uniformly strongly elliptic. The fixed sufficiently small base neighborhood of δ supplies its ellipticity margin.

Suppose the existing correction has $C^{j,\alpha}$ regularity on a bounded region, with $j \geq k_0$. The principal coefficients then have the required $C^{j-1,\alpha}$ regularity. After moving the prescribed $\partial^2 q$ and ∂p terms to the right, all remaining nonlinear terms use at most first derivatives of (u, W). A finite additional collection of seed derivatives therefore makes the right-hand side $C^{j-1,\alpha}$. The interior estimate gives $C^{j+1,\alpha}$ regularity on a smaller

region. Difference quotients, or the regularity part of that estimate, apply to the already existing solution; no higher-order contraction is being asserted. Induction proves local smoothness.

For the weighted improvement, apply the same argument on fixed nested rescaled annuli. With $x = Rz$, normalize the seed and correction components as

$$q_R(z) = R^{-\gamma} q(Rz), \qquad p_R(z) = R^{1-\gamma} p(Rz), \qquad y_R(z) = R^{-\gamma}(u, W)(Rz).$$

The metric on the rescaled annulus is $\delta + R^{\gamma}(q_R + u_R\delta)$, and the scaled momentum expression uses $R^{\gamma}(p_R + L_0 W_R)$. Multiplying the constraint equations by $R^{2-\gamma}$ yields the same principal elliptic system with uniformly controlled rescaled coefficients. The base norms bound y_R at order k_0, uniformly in $R \geq 1$. At each later step the required finite seed and already obtained correction norms are uniformly bounded by their weighted seminorms. Factors R^{γ} do not grow, since $\gamma < 0$. The interior estimate consequently gives the next-order bound uniformly in R. Taking the supremum over annuli and adding the interior-ball estimate proves all weighted orders. Each finite induction has finite input dependence and uses only the original ellipticity margin.

Now let

$$\mathsf{A}_{s,y} = D_y \mathcal{F}(s, y).$$

At every weighted smooth point in the base product neighborhood, its base inverse is

$$\mathsf{A}_{s,y}^{-1} = \big(I + D_y(A_0^{-1}\mathcal{N})(s, y)\big)^{-1} A_0^{-1},$$

with the first inverse bounded by $(1 - \kappa)^{-1}$. This follows directly from the base inverse margin.

After the scalar sign change, write this linearized operator as $a^{ij}\partial_i\partial_j + b^i\partial_i + c$. Its principal blocks have the ellipticity margin just verified; the off-diagonal blocks have order at most one. The coefficients $a - a_{\infty}$, b and c have weights γ, $\gamma - 1$ and $\gamma - 2$, respectively. This follows by differentiating the constraint expressions: the principal coefficients are smooth algebraic functions of h; a first-order coefficient contains one derivative of the metric or one momentum factor; a zeroth-order coefficient contains two metric derivatives, one momentum derivative, or products with at least the same decay.

Under $v_R(z) = R^{-\gamma} v(Rz)$ the coefficients become $a(Rz)$, $Rb(Rz)$ and $R^2 c(Rz)$. They have uniformly bounded finite $C^{m-2,\alpha}$ norms on fixed nested annuli. The interior estimate and the corresponding inner-ball estimate give

$$\|v\|_{\mathcal{Y}_m} \leq C_m\big(\|\mathsf{A}_{s,y} v\|_{\mathcal{Z}_m} + \|v\|_{C^0_{\gamma}}\big).$$

The constant depends on the ellipticity margin and finitely many higher background seminorms. The base inverse controls the last term for a solution of $\mathsf{A}_{s,y} v = f$. Local elliptic regularity first improves this base solution, and

the annular estimate makes the improved norm global. Thus

$$\|\mathsf{A}_{s,y}^{-1} f\|_{\mathcal{Y}_m} \le C'_m \|f\|_{\mathcal{Z}_m}, \qquad m \ge k_0.$$

Base uniqueness identifies these inverses at every order. The bound uses a constant C'_m for each fixed m.

For continuity, let $s_n \to s$ in every $\mathcal{X}_m$, and put $y_n = \mathcal{C}(s_n)$, $y = \mathcal{C}(s)$. Base smoothness gives $y_n \to y$ in $\mathcal{Y}_{k_0}$. The spatial induction just proved gives uniform bounds for y_n at every finite higher order. Set $d_n = s_n - s$ and $e_n = y_n - y$. On the convex product neighborhood use the segment $(s_\theta, y_\theta) = (s + \theta d_n, y + \theta e_n)$, and define

$$L_n = \int_0^1 D_y \mathcal{F}(s_\theta, y_\theta)\, \mathrm{d}\theta, \qquad M_n = \int_0^1 D_s \mathcal{F}(s_\theta, y_\theta)\, \mathrm{d}\theta.$$

The exact difference equation is

$$L_n e_n = -M_n d_n.$$

The segment is an ambient segment; it need not lie in the constrained subset. Averaging preserves the fixed ellipticity margin and the order of the off-diagonal terms. The finite rescaled coefficient bounds are uniform in n, by the preceding spatial bounds. Therefore the same annular estimate gives

$$\|e_n\|_{\mathcal{Y}_m} \le C_m(\|M_n d_n\|_{\mathcal{Z}_m} + \|e_n\|_{\mathcal{Y}_{k_0}}) \le C'_m(\|d_n\|_{\mathcal{X}_m} + \|e_n\|_{\mathcal{Y}_{k_0}}).$$

The last bound uses that $D_s\mathcal{F}$ has at most two derivatives of the metric seed and one of the momentum seed, together with the uniform finite background norms. The right-hand side tends to zero. This proves continuity of $\mathcal{C}$ at all weighted orders. The ambient algebraic formulas give continuity of $\mathcal{T}$.

For a supplied smooth parameter a, base differentiation gives

$$D_y\mathcal{F}(s, \mathcal{C}(s))\, \partial_a \mathcal{C}(s) = -D_s\mathcal{F}(s, \mathcal{C}(s))\, \partial_a s.$$

The all-order inverse estimate applies to this actual derivative. Repeated differentiation retains the same invertible left side; its right side is a finite sum of input derivatives and previously obtained correction derivatives. At a fixed parameter point the Banach implicit function theorem can also be applied in any higher space, because the just-proved higher inverse is an isomorphism there. Base uniqueness identifies its local solution with the already selected $\mathcal{C}$. This use of a higher-space neighborhood around an existing solution imposes no additional smallness on that solution's higher norms. It identifies the derivatives and their continuity in the high space. Induction gives smooth parameter dependence at each finite order.

On a retained compact parameter set, the higher estimates have uniform constants when the requisite finite input seminorms and the base margins are uniform. A slightly larger retained compact set supplies the parameter neighborhoods and derivative bounds used in this argument. A finite cover and the differentiated equations give the asserted finite estimates on the smaller set.

Finally, the decoder has only algebraic operations, one divergence and the fixed decaying inverse Δ_L^{-1}. The orders in Proposition A.4 and Lemma A.2 show its continuity and smooth finite-family dependence at every weighted level. Its exact identities hold before the high-order argument and consequently hold for this same smooth branch. □

A.5. Based spatial gauge. The inverse used for a coordinate displacement has a different weight from the decaying inverse used in the constraint decoder. Fix

$$k = k_0 \geq 4, \qquad 0 < \alpha < 1, \qquad -1 < \gamma < -\frac{1}{2}, \qquad d = \gamma + 1 \in (0, 1/2).$$

The displacement is permitted to grow with order d; its derivative decays with order γ. Its value at the marked point $o = 0$ fixes the translation ambiguity. We use the same base index as the constraint chart, with the weighted Hölder spaces and Cartesian tensor conventions specified above.

Lemma A.8 (The anchored sublinear York inverse).
For $0 < d < 1$, the operator

$$\Delta_L : \{X \in C_d^{k+1,\alpha}(T\mathbb{R}^3) : X(0) = 0\} \longrightarrow C_{d-2}^{k-1,\alpha}(T\mathbb{R}^3)$$

is a bounded isomorphism. With G the matrix Green kernel of the decaying York inverse, its inverse is the actual integral

$$(\mathcal{G}_o F)_i(x) = \int_{\mathbb{R}^3} \big(G_{ij}(x-y) - G_{ij}(-y)\big) F_j(y)\, \mathrm{d}y,$$

and

$$\|\mathcal{G}_o F\|_{C_d^{k+1,\alpha}} \leq C_{k,\alpha,d} \|F\|_{C_{d-2}^{k-1,\alpha}}.$$

Proof. The two possible singularities, at $y = x$ and $y = 0$, are locally integrable in three dimensions. For $|y| > 2|x| + 2$, the degree -1 homogeneity of G and the mean-value formula imply

$$|G(x-y) - G(-y)| \leq C|x|\,|y|^{-2}.$$

The absolute tail is bounded by $C\|F\|_{C^0_{d-2}}|x| \int_{2|x|+2}^{\infty} t^{d-2}\,\mathrm{d}t$, which is finite because $d < 1$. The integral vanishes at $x = 0$.

For the weighted zeroth-order estimate, let $R = |x| \geq 2$. On $|y| \leq 2R$, the subtracted kernel term has integral bounded by

$$C \int_0^{2R} t\langle t\rangle^{d-2}\,\mathrm{d}t \leq C_d R^d.$$

For $G(x-y)$ split this region into $|y| < R/2$, $|x-y| < R/2$ and the remaining comparable annulus. On the first part the kernel is bounded by C/R, giving $CR^{-1}\int_0^{R/2} t^2\langle t\rangle^{d-2}\,\mathrm{d}t \leq CR^d$. On the second part the source is bounded by CR^{d-2} and the kernel has integral $O(R^2)$. On the remaining part rescaling by R gives the same bound. The outer difference term is bounded by $CR\int_{2R}^{\infty} t^{d-2}\,\mathrm{d}t \leq C_d R^d$. The region $|x| \leq 2$ has a uniform bound by local

integrability and the tail estimate. Splitting small neighborhoods of the integrable singularities also proves continuity. Hence

$$\|\mathcal{G}_o F\|_{C^0_d} \le C_d \|F\|_{C^0_{d-2}}.$$

Applying Δ_L in the x variable kills the subtracted term and gives F. This holds distributionally: on a compact test support the same tail difference is integrable, so the kernel calculation can be tested and integrated. The local elliptic regularity and interior estimate stated above give the higher regularity. On nested annuli set $v_R(y) = X(Ry)$; the equation is $\Delta_{L,y} v_R = R^2 F(Ry)$. Multiplication of the local estimate by R^{-d} gives exactly the asserted higher weighted bound. A fixed ball handles the origin.

For uniqueness, suppose $\Delta_L X = 0$ in the stated domain. Taking divergence gives $(4/3)\Delta \operatorname{div} X = 0$. The divergence is a smooth harmonic scalar of weight $d-1 < 0$, so it is zero by the maximum principle and decay at infinity. Thus every component of X is harmonic. Each first derivative is harmonic and has the negative weight $d-1$, so every first derivative is zero. The field is constant and its value at the anchor is zero. This proves injectivity. □

For a based displacement X, put $\phi_X(x) = x + X(x)$. If $\|DX\|_{L^\infty} < 1$, then ϕ_X is a global diffeomorphism. For every y the map $x \mapsto y - X(x)$ is a contraction of complete Euclidean space and has one fixed point. The differential of ϕ_X is invertible, so its inverse is smooth to the available order. The path $x \mapsto x + tX(x)$ has invertible differential for $0 \le t \le 1$, hence the diffeomorphism preserves orientation. Since $X(0) = 0$,

$$(1 - \|DX\|_\infty)|x| \le |\phi_X(x)| \le (1 + \|DX\|_\infty)|x|.$$

This comparability and the chain rule on rescaled annuli give the weighted composition estimates. In particular, if $X \in \mathscr{S}^d$ and $d < 1$, the inverse displacement is also in $\mathscr{S}^d$. To see the weight explicitly, write $Y = -X \circ \phi_X^{-1}$. Its zeroth-order estimate follows from comparability. Differentiating the inverse equation first gives $D\phi_X^{-1} = (I + DX \circ \phi_X^{-1})^{-1}$; subsequent derivatives are finite sums of derivatives of X and lower derivatives of the inverse. Induction gives the claimed symbol order for Y. On an annulus the normalized displacement is $R^{-1}X(R\cdot) = O(R^{d-1})$, so the same estimates are uniform in its radius.

For a positive metric h, the actual gauge residual is

$$\mathcal{B}_h(X) = \operatorname{div}_\delta \operatorname{TF}_\delta(\phi_X^* h).$$

In Cartesian coordinates, with $A_i{}^a = \delta_i{}^a + \partial_i X^a$, its input tensor is

$$(\phi_X^* h)_{ij}(x) = h_{ab}(x + X(x)) A_i{}^a A_j{}^b.$$

Both Jacobian factors belong to this definition.

Theorem A.9 (Based normalization of the metric gauge).
There are positive constants ϵ and ρ, fixed at the finite order in the weight convention, such that every metric satisfying

$$\|h - \delta\|_{C^{k+2,\alpha}_\gamma} < \epsilon$$

has a unique based displacement $X_h \in C^{k+1,\alpha}_d$ in the ball $\|X_h\| < \rho$ with $\mathcal{B}_h(X_h) = 0$. Its derivative has norm less than $1/2$. Thus $\phi_h = x + X_h$ is a global orientation-preserving diffeomorphism fixing the origin. Moreover,

$$\|X_h\|_{C^{k+1,\alpha}_d} \le C \| \operatorname{div}_\delta \operatorname{TF}_\delta h\|_{C^{k-1,\alpha}_{\gamma-1}}.$$

It follows that ϕ_h is the identity whenever h already satisfies the gauge. If $h - \delta \in \mathscr{S}^\gamma$, then $X_h \in \mathscr{S}^d$. For a finite-dimensional family smooth in all these weighted seminorms and remaining in the one base domain, the same X_h is smooth in every weighted spatial order. Each specified finite spatial and parameter estimate uses finitely many higher input seminorms and the fixed base margins. Higher seminorms need not be small.

Proof. At the flat zero input,

$$D_X \mathcal{B}_\delta(0)V = \operatorname{div}_\delta \operatorname{TF}_\delta(\mathcal{L}_V \delta) = \Delta_L V.$$

Write $r = h - \delta$ and

$$\mathcal{B}_h(X) = b_h + \Delta_L X + \mathcal{N}_h(X), \qquad b_h = \operatorname{div}_\delta \operatorname{TF}_\delta h, \qquad \mathcal{N}_h(0) = 0.$$

On a sufficiently small based displacement ball, the actual pullback formula gives

$$\begin{aligned}&\|\mathcal{N}_h(X) - \mathcal{N}_h(Y)\|_{C^{k-1,\alpha}_{\gamma-1}} \\ &\quad \le C\big(\|r\|_{C^{k+2,\alpha}_\gamma} + \|X\|_{C^{k+1,\alpha}_d} + \|Y\|_{C^{k+1,\alpha}_d}\big)\|X - Y\|_{C^{k+1,\alpha}_d}.\end{aligned}$$

Here is why the stated orders suffice. Differentiate the pullback once in space, expand its two Jacobian factors, and subtract the flat terms linear in X. Every remaining displacement term contains a metric perturbation or at least two displacement factors. Subtracting two such expressions and using the mean-value formula for composed coefficients gives terms of the types

$$r D^2 X, \qquad Dr\, DX, \qquad D^2 r\, X, \qquad DX\, D^2 X,$$

with smooth contractions and composed versions of the coefficients. Their weights are at most $2\gamma - 1$, which embeds into $\gamma - 1$ because $\gamma < 0$. On an annulus of radius R, all arguments lie in one enlarged comparable annulus by radius comparability. The normalized displacement is $R^{-1}X(R\cdot)$, of order R^γ. Ordinary Hölder product and composition estimates on fixed annuli therefore have uniform radius-independent constants. Taking $k-1$ derivatives and their Hölder differences is covered by the displayed $k+2$ metric norm and $k+1$ displacement norm. The additional metric derivatives supply the mean-value bound in the highest Hölder norm retained here. This accounts for the derivative cost of the diffeomorphism action. The bounded region is treated on fixed balls. These observations prove the remainder estimate.

Use Lemma A.8 to define

$$\mathcal{T}^{\mathrm{g}}_h(X) = -\mathcal{G}_o(b_h + \mathcal{N}_h(X)).$$

Choose ρ and then ϵ so that this map has Lipschitz constant at most $1/2$ and sends the closed based ball of radius ρ into itself. This follows from the

remainder bound, $\mathcal{N}_h(0) = 0$, and the small bound on b_h. Its fixed point is unique in that ball. The inverse identity of Lemma A.8 makes its equation equivalent to $\mathcal{B}_h(X_h) = 0$. Applying the norm estimate to the fixed-point equation and absorbing its Lipschitz part proves the residual bound. If $b_h = 0$, the initialized zero displacement is a fixed point, so uniqueness makes $X_h = 0$. Decrease the one base radius so that the derivative bound is less than $1/2$. The preceding global inverse argument then applies.

We next improve this same solution. The coordinate formula writes its equation as a quasilinear second-order system for X. The second-order coefficients depend on $h(x + X)$ and DX; all terms involving derivatives of h use at most its first derivative before the equation is differentiated. The positive principal symbol of $-D_X\mathcal{B}_h$ at $(h, X) = (\delta, 0)$ is

$$|\xi|^2 I + \frac{1}{3}\xi \otimes \xi.$$

Its eigenvalues are $|\xi|^2$ and $(4/3)|\xi|^2$. The small base metric and first-derivative bounds make the actual symbol a uniformly strongly elliptic perturbation. This ellipticity margin is fixed before every higher derivative order.

Suppose the existing X has weighted order $C^{m,\alpha}_d$, with $m \geq k + 1$. On a fixed annulus the principal coefficients then have $C^{m-1,\alpha}$ regularity, and the remaining terms have that same regularity using finitely many derivatives of the supplied smooth metric. The stated interior system Schauder estimate, applied to the existing solution on nested annuli, gives one additional derivative. In rescaled coordinates the equation has the same form for

$$X_R(y) = R^{-1}X(Ry), \qquad h_R(y) = h(Ry).$$

Both $h_R - \delta$ and X_R have size $O(R^\gamma)$ in the corresponding already controlled norms. Their nonlinear products have at least this decay, and the normalized composition arguments stay in fixed comparable annuli. The interior estimates consequently give a finite bound for $R^{-\gamma}\|X_R\|_{C^{m+1,\alpha}}$ using finitely many input seminorms and the preceding bound. Taking the supremum over $R \geq 1$ proves the weighted gain. A fixed ball handles the origin. Induction gives every finite weighted order without decreasing the base neighborhood.

For the linearization at the selected solution, put $d = \gamma + 1$. In this proof write $E_m = \{v \in C^{m,\alpha}_d : v(0) = 0\}$ and $F_m = C^{m-2,\alpha}_{d-2}$ for spaces of vector fields, with $m \geq k + 1$. Then $L = D_X\mathcal{B}_h(X_h) = \Delta_L + D_X\mathcal{N}_h(X_h)$ maps E_m boundedly to F_m at the regularity just obtained. At $m = k + 1$, Lemma A.8 and the contraction estimate give the inverse $(I + \mathcal{G}_o D_X\mathcal{N}_h(X_h))^{-1}\mathcal{G}_o$. For $f \in F_m$ apply this base inverse to obtain v. Interior Schauder regularity for $Lv = f$ on a fixed ball and on rescaled annuli gives

$$\|v\|_{E_m} \leq C_m(\|f\|_{F_m} + \|v\|_{C^0_d}).$$

Indeed rescaling $v(R\cdot)$ by R^{-d} and $f(R\cdot)$ by R^{2-d} gives uniform estimates on nested unit annuli by the coefficient bounds above. The base inverse bounds

the last term by $C\|f\|_{F_{k+1}} \le C'\|f\|_{F_m}$. Hence the same based solution lies in E_m and satisfies $\|v\|_{E_m} \le C'_m\|f\|_{F_m}$. This proves surjectivity at order m, and injectivity follows by inclusion in the base space.

Fix a finite differentiation order $j \ge 1$. On the open set of positive metrics and based displacements satisfying the strict base bounds and the diffeomorphism condition, $(h, X) \mapsto \mathcal{B}_h(X)$ is C^j from $(\delta + C^{m+j,\alpha}_\gamma) \times E_m$ to F_m. Before spatial differentiation the highest metric derivative in the jth displacement derivative is $(D^{j+1}h) \circ \phi_X$. Taking the $m-2$ target derivatives uses at most $m + j - 1$ metric derivatives. The remaining derivative gives continuity of composition in the top Hölder norm by the mean-value estimate on enlarged annuli. The same argument for a metric variation, together with the rescaled product estimates above for the Jacobian factors and the displacement variations, gives the mixed Fréchet derivatives with the stated weights. The Banach implicit-function theorem at (h, X_h) now gives a C^j solution map into E_m locally in this higher metric norm. Its values solve the base equation in the uniqueness ball, hence coincide with X_h there. In particular the parameter derivatives through order j are continuous in E_m. Differentiating the equation expresses each highest parameter derivative through the bounded inverse of L and terms of lower parameter order.

These differentiability arguments are applied at points of the already fixed parameter domain where the base smallness holds. For any compact subset, a finite cover and a slightly larger retained compact set bound the finitely many input seminorms and derivatives used at a specified order. The output is therefore smooth at every order on the original retained domain; no new solution or higher-order smallness condition is selected. □

Proposition A.10 (Rotation and actual spin markings).
Choose equivalent rotation-invariant base norms and the preceding domains invariant under proper Euclidean rotations. For $A \in \mathrm{SO}(3)$,

$$X_{A^*h}(x) = A^{-1}X_h(Ax), \qquad \phi_{A^*h} = A^{-1} \circ \phi_h \circ A.$$

The path $\phi_{h,t}(x) = x + tX_h(x)$, $0 \le t \le 1$, gives the initialized actual spin lift of the normalization. It varies smoothly with every stipulated finite-dimensional smooth family. If a comparison is supplied with the opposite central lift, that lift remains a distinct marking.

Proof. The Euclidean trace, divergence and anchor are preserved by proper rotations. The matrix kernel obeys $G(Ax) = AG(x)A^{-1}$, so the anchored inverse commutes with the corresponding tensor pullbacks. Alternatively, direct substitution in the gauge equation shows that the displayed conjugated displacement solves the rotated equation. It has the same rotation-invariant norm and belongs to the same based uniqueness ball. Theorem A.9 gives the two identities.

Every $\phi_{h,t}$ is a based orientation-preserving diffeomorphism by the fixed derivative bound. Transport orthonormal frames by the differential and the positive polar identification between the corresponding metrics. This yields a smooth path in the oriented frame bundle starting at the stipulated

identity marking. The path-lifting property of its spin double cover gives one actual lift with that initial value. Smooth dependence follows from local trivializations of the covering and the same initialization. The two lifts differ by the central deck transformation; fixing the initial lift determines which one is transported. □

Corollary A.11 (Access to the exact maximal-data decoder). Let (H, L) be whole smooth maximal vacuum data on the marked $\mathbb{R}^3$, with

$$H - \delta \in \mathscr{S}^{\gamma}, \qquad L \in \mathscr{S}^{\gamma-1},$$

and sufficiently small in $C^{k+2,\alpha}_{\gamma} \oplus C^{k+1,\alpha}_{\gamma-1}$. After applying the based gauge ϕ_H, the normalized data

$$(h, K) = \phi_H^*(H, L)$$

belong to the independently defined gauged maximal constraint neighborhood. Their exact ambient decoder gives a unique smooth pair of transverse-traceless seeds (q, p) in its selected branch, and

$$(h, K) = \mathcal{T}(q, p).$$

All indicated higher spatial and finite-family estimates hold for this same gauge and decoded pair. The actual spin comparison is the one in Proposition A.10.

Proof. Pullback preserves the three maximal vacuum constraints. Theorem A.9 imposes the metric gauge. Weighted composition with the sublinear displacement preserves the symbol orders of both tensors; the Jacobian factors are the identity plus a decaying tensor. The gauge estimate and the smallness of the whole pair put the normalized fields in the ambient decoder neighborhood after one restriction of the finite base threshold.

The exact decoder gives

$$q = \mathrm{TF}_{\delta} h, \qquad u = \frac{1}{3} \operatorname{tr}_{\delta} h - 1,$$
$$\kappa = K - \frac{\operatorname{tr}_{\delta} K}{\operatorname{tr}_{\delta} h} h, \qquad W = \Delta_L^{-1} \operatorname{div}_{\delta} \kappa, \qquad p = \kappa - L_0 W.$$

Its algebraic identities and the decaying York inverse put q, p in the stated transverse-traceless spaces. The independently satisfied constraints force $(u, W) = \mathcal{C}(q, p)$ by uniqueness of the nonlinear correction in its base branch. This proves the tensor equality. The same-branch weighted estimates for the chart and gauge, and continuity of the exact decoder at every finite order, give all stated finite-family bounds. Proposition A.10 supplies the initialized actual spin marking. □

The gauge and the decoder operate on the whole fields before the resulting data or source are restricted to a local patch. The allowed displacement is based and sublinear. The source comparisons for different whole extensions and different initial hypersurfaces use the later local comparison theorem; they are not prerequisites for this normalization.

A.6. Current differential operations. Use the spaces $\mathcal{X}_m, \mathcal{Y}_m, \mathcal{Z}_m$ of Definition A.3, with the same k_0, α, γ. A current state is $Z = (s, y)$, where $s = (q, p)$ and $y = (u, W)$. Its coefficients are

$$h = \delta + q + u\delta, \qquad \kappa = p + L_0 W, \qquad K = \kappa - \frac{1}{3}\vartheta_\kappa h, \qquad \vartheta_\kappa = \operatorname{tr}_h \kappa.$$

The residual $\mathcal{F}(Z) = (R(h) - |K|_h^2, \operatorname{div}_h K)$ is defined throughout the positive ambient product. Its current derivatives are finite differential operations, including off the zero-residual set. The inverse used below is the completed operation

$$\mathsf{G}_0^{\mathrm{src}} = \operatorname{diag}\left(-\frac{1}{2}\Delta^{-1}, \Delta_L^{-1}\right) = A_0^{-1}$$

constructed from the finite negative-response word in the source Green theorem. Its bounded realization on $\mathcal{Z}_m \to \mathcal{Y}_m$ is the decaying inverse of Lemma A.2. Thus the following limits have a specified Banach-space meaning as well as an actual finite-word realization on weighted smooth fields.

Lemma A.12 (The current residual differential).
For a seed variation $a = (a_q, a_p)$ and a correction variation $v = (v_u, v_W)$ put

$$\begin{aligned} r &= a_q + v_u\delta, \\ z &= a_p + L_0 v_W, \\ \dot\vartheta_\kappa &= \operatorname{tr}_h z - r^{ij}\kappa_{ij}, \\ k &= z - \frac{1}{3}\dot\vartheta_\kappa h - \frac{1}{3}\vartheta_\kappa r, \end{aligned}$$

where $r^{ij} = h^{ia}h^{jb}r_{ab}$. All covariant derivatives below use h. Define

$$C^\ell{}_{ij}(r) = \frac{1}{2}h^{\ell a}(\nabla_i r_{ja} + \nabla_j r_{ia} - \nabla_a r_{ij}).$$

Then the scalar and one-form components of $D\mathcal{F}_Z[a, v]$ are

$$\begin{aligned} \mathfrak{H}_Z(r,k) &= \nabla^i\nabla^j r_{ij} - \Delta_h \operatorname{tr}_h r - \operatorname{Ric}(h)^{ij} r_{ij} \\ &\quad - 2K^{ij}k_{ij} + 2(K\circ K)^{ij} r_{ij}, \\ \mathfrak{M}_Z(r,k)_i &= \nabla^j k_{ji} - r^{jk}\nabla_j K_{ki} \\ &\quad - h^{jk}(C^\ell{}_{jk}(r)K_{\ell i} + C^\ell{}_{ji}(r)K_{k\ell}), \end{aligned}$$

where $(K\circ K)^{ij} = K^i{}_\ell K^{\ell j}$. In particular,

$$\mathsf{A}_Z = D_y\mathcal{F}_Z, \qquad \mathsf{B}_Z = D_s\mathcal{F}_Z$$

are obtained by setting $a = 0$ or $v = 0$ in these formulas. At the flat state, $\mathsf{A}_0 = A_0 = (-2\Delta, \Delta_L)$ and $\mathsf{B}_0 = 0$. The maps are continuous at each indicated weighted order and real analytic on the positive base neighborhood.

Proof. Differentiating $h^{-1}h = I$ gives $\dot h^{ij} = -r^{ij}$. Differentiate $\vartheta_\kappa = h^{ij}\kappa_{ij}$ and $K = \kappa - \vartheta_\kappa h/3$ to obtain the first displayed variation. In particular it retains the variation of the tracefree projection.

Differentiating the Koszul formula gives $C^{\ell}{}_{ij}$. Its Ricci contraction gives $DR_h[r] = \nabla^i\nabla^j r_{ij} - \Delta_h \operatorname{tr}_h r - \operatorname{Ric}^{ij} r_{ij}$. Differentiation of the two inverse metrics in $|K|_h^2$ gives $D|K|_h^2[r,k] = 2K^{ij}k_{ij} - 2(K\circ K)^{ij}r_{ij}$. Finally, differentiate $h^{jk}\nabla_j K_{ki}$. The inverse-metric term is the displayed $-r^{jk}\nabla_j K_{ki}$, and differentiating each covariant slot of K gives the two terms involving C. This proves the formulas for the covariant momentum one-form.

The flat identities follow from the TT conditions and the definition of L_0. Each formula contains at most two derivatives of r and one of k. Weighted multiplication therefore gives the asserted maps into $\mathcal{Z}_m$. Matrix inversion is real analytic in the fixed positive neighborhood; the other operations are finite differentiations, contractions and products. This proves the final assertion. □

A.7. The current inverse and horizontal fields. At a current state set

$$\begin{aligned} \mathsf{E}_Z &= \mathsf{G}_0^{\mathrm{src}}(\mathsf{A}_Z - A_0),\\ \mathsf{G}_Z &= \sum_{j=0}^{\infty}(-\mathsf{E}_Z)^j \mathsf{G}_0^{\mathrm{src}},\\ \mathsf{V}_Z a &= -\mathsf{G}_Z \mathsf{B}_Z a, \qquad \mathscr{H}_a(Z) = (a, \mathsf{V}_Z a). \end{aligned}$$

The parameter direction a changes the whole input seed. These are vector fields on the space of source coefficients, distinct from the physical point directions of a spacetime.

Proposition A.13 (Current inverse and flatness).
There is a convex base product $D \subset \mathcal{X}_{k_0}\times\mathcal{Y}_{k_0}$ around zero on which the metric has a fixed positive margin and

$$\|\mathsf{E}_Z\|_{\mathcal{L}(\mathcal{Y}_{k_0})} \le \varkappa < \frac{1}{2}.$$

The series defining G_Z converges in the base operator norm and equals A_Z^{-1}. The maps $\mathsf{G}_Z, \mathsf{V}_Z$ are real analytic there. For constant seed directions a, b,

$$D\mathcal{F}_Z\, \mathscr{H}_a(Z) = 0, \qquad [\mathscr{H}_a, \mathscr{H}_b] = 0.$$

At a weighted smooth current state, this same inverse maps $\mathcal{Z}_m$ boundedly to $\mathcal{Y}_m$ at every finite $m \ge k_0$. Each bound depends on finitely many higher state seminorms and the fixed base margins. The same higher inverse estimate holds for the average of A along any segment in D whose endpoints have the stated higher bounds.

Proof. The quadratic residual estimate in Lemma A.5 gives, on a smaller product,

$$\|\mathsf{E}_Z\| + \|\mathsf{G}_0^{\mathrm{src}}\mathsf{B}_Z\| \le C(\|s\|_{\mathcal{X}_{k_0}} + \|y\|_{\mathcal{Y}_{k_0}}).$$

Choose the product convex and small enough for the stated margin. Since $\mathsf{A}_Z = A_0(I + \mathsf{E}_Z)$, the Neumann series gives the actual inverse with norm at most $(1-\varkappa)^{-1}\|\mathsf{G}_0^{\mathrm{src}}\|$. The current derivatives in Lemma A.12 are analytic. The convergent inverse, or the inverse identity differentiated on this

neighborhood, therefore gives analytic dependence. In particular, for a state variation ζ,

$$D\mathsf{V}_Z[\zeta]a = -\mathsf{G}_Z\big(D\mathsf{A}_Z[\zeta]\mathsf{V}_Z a + D\mathsf{B}_Z[\zeta]a\big).$$

Residual conservation is the identity $\mathsf{B}_Z a+\mathsf{A}_Z\mathsf{V}_Z a = 0$. Differentiate it in the direction $\mathscr{H}_b$ and subtract the corresponding equation with a, b exchanged. The Hessian terms cancel by symmetry, giving $D\mathcal{F}_Z[\mathscr{H}_a, \mathscr{H}_b] = 0$. The seed component of the bracket vanishes because a, b are constant. Its remaining component is killed by the invertible A_Z, proving flatness. No zero-residual hypothesis is used here.

The proof of Theorem A.7 already gives the required transverse estimate at every smooth point in a sufficiently small ambient product, rather than only on its constraint graph. After changing the sign of the scalar row, the principal blocks are uniformly strongly elliptic and the off-diagonal terms have order at most one. Relative to the flat blocks, their second-, first- and zeroth-order coefficients have weights $\gamma, \gamma - 1, \gamma - 2$. On nested rescaled annuli the estimate is

$$\|v\|_{\mathcal{Y}_m} \le C_m\big(\|\mathsf{A}_Z v\|_{\mathcal{Z}_m} + \|v\|_{C^0_\gamma}\big).$$

Apply local regularity to the existing base solution, then this annular estimate and the interior-ball estimate. The base inverse controls the last term. This proves the higher bound for the same inverse, without requiring high-order convergence of the raw Neumann partial sums.

For an average $\overline{\mathsf{A}}$ along a segment, the normalized perturbation is the average of E, so it has the same base norm margin. Convexity also preserves the near-flat principal ellipticity margin, and the finite coefficient estimates are bounded by the stipulated finite endpoint bounds. The same proof applies. □

A.8. Transport on one base tube. Fix a weighted smooth reference state $(s_*, y_*) \in D$. Choose outer balls $B(s_*, r_s) \times B(y_*, r_y)$ contained in D and constants $\|\mathsf{V}\| \le M$, $\|D_y\mathsf{V}\| \le L$ there. Take $\rho > 0$ such that

$$\rho < r_s, \qquad M\rho < r_y/4, \qquad L\rho < 1/4,$$

omitting a condition with zero left coefficient. Set $S = B_{\mathcal{X}_{k_0}}(s_*, \rho)$ and $J = B_{\mathcal{Y}_{k_0}}(y_*, r_y/4)$. The ordinary-pair constructor uses the fixed flat choice $(s_*, y_*) = (0, 0)$. The general fiber labels below describe transport within the chosen coefficient neighborhood; they are not additional generating inputs to that constructor.

Theorem A.14 (Actual current transport and its endpoint). For $s \in S$, $y_0 \in J$ and $a = s - s_*$, the iteration

$$y^{[0]}(t) = y_0,$$

$$y^{[j+1]}(t) = y_0 + \int_0^t \mathsf{V}_{(s_*+\tau a, y^{[j]}(\tau))} a \,\mathrm{d}\tau$$

converges in $C^0([0,1];\mathcal{Y}_{k_0})$. Its endpoint $\Psi(s,y_0)$ satisfies

$$\begin{aligned}\mathcal{F}(s,\Psi(s,y_0)) &= \mathcal{F}(s_*,y_0),\\ D_s\Psi(s,y_0) &= \mathsf{V}_{(s,\Psi(s,y_0))},\\ D_{y_0}\Psi(s,y_0) &= \mathsf{G}_{(s,\Psi(s,y_0))}\mathsf{A}_{(s_*,y_0)}.\end{aligned}$$

The map $(s,y_0)\mapsto(s,\Psi(s,y_0))$ is a diffeomorphism onto an open tube in D. Every piecewise C^1 seed path remaining in S has a unique horizontal lift from any point of that tube, with its fiber label fixed. Its endpoint depends only on the endpoint seed, and the transports satisfy

$$\mathscr{T}_{s_2,s_1}\mathscr{T}_{s_1,s_0} = \mathscr{T}_{s_2,s_0}, \qquad \mathscr{T}_{s,s} = I.$$

These assertions apply to paths of arbitrary finite length in S.

For weighted smooth seeds and fiber labels, the same base endpoint belongs to every $\mathcal{Y}_m$. On sets with the required finite higher bounds,

$$\|\Psi(s,y_0)-\Psi(t,\widetilde{y}_0)\|_{\mathcal{Y}_m} \le C_m\big(\|s-t\|_{\mathcal{X}_m} + \|y_0-\widetilde{y}_0\|_{\mathcal{Y}_m}\big).$$

Smooth finite-dimensional families are preserved. For compact parameter sets $\mathcal{K}\Subset\operatorname{int}\mathcal{K}'$, every specified finite spatial and parameter estimate on $\mathcal{K}$ uses finitely many input seminorms on $\mathcal{K}'$ and the fixed base margins. The input ball and the spatial domain are fixed before these higher orders.

If $\mathcal{F}(s_*,y_*)=0$, put $Y(s)=\Psi(s,y_*)$. The entire independently tested zero-residual set in $S\times B(y_*,r_y)$ consists precisely of $(s,Y(s))$.

Proof. On the closed path ball $\sup_t\|y(t)-y_0\|_{\mathcal{Y}_{k_0}}\le r_y/4$, the integral map takes values in the same ball and has contraction constant $q_{\mathrm{P}}=L\rho<1/4$. All its states stay in the outer product. Its iterates converge with error at most $q_{\mathrm{P}}^j M\|a\|/(1-q_{\mathrm{P}})$. The limit is C^1 in t and solves the integral equation. The integral difference estimate gives uniqueness while a competing solution remains in the ambient domain.

The derivative in the path variable has norm at most q_{P}. The differentiated fixed-point equation therefore has the bounded inverse of I minus that derivative. Repeated differentiation proves smooth dependence in the base Banach spaces. Along the actual limit the chain rule gives $\partial_t\mathcal{F}=\mathsf{B}a+\mathsf{A}\mathsf{V}a=0$. Differentiate this conserved identity in s and y_0 and use the current inverse to obtain all three endpoint identities.

There is an independent injectivity estimate on every convex fiber segment in D:

$$\begin{aligned}\mathsf{G}_0^{\mathrm{src}}\big(\mathcal{F}(s,y)-\mathcal{F}(s,z)\big) &= \left(I+\int_0^1 \mathsf{E}_{(s,z+\tau(y-z))}\,\mathrm{d}\tau\right)(y-z),\\ (1-\varkappa)\|y-z\|_{\mathcal{Y}_{k_0}} &\le \|\mathsf{G}_0^{\mathrm{src}}(\mathcal{F}(s,y)-\mathcal{F}(s,z))\|_{\mathcal{Y}_{k_0}}.\end{aligned}$$

The derivative in the fiber label in the endpoint identities is invertible, so the product map is a local diffeomorphism. If two labels have the same output over s, residual conservation and this injectivity estimate at s_* show that the labels agree. Its image is consequently an open tube and its local inverses agree to give a global inverse there.

For an arbitrary seed path $s(t)$, the field $\Psi(s(t), y_0)$ stays in this tube and, by its seed derivative, solves the horizontal equation. Local uniqueness gives its unique lift along the whole path. Thus the label remains fixed, proving endpoint independence and the composition laws without a length restriction beyond the stated finite path domain. If the reference residual is zero, conservation gives $\mathcal{F}(s, Y(s)) = 0$. The fiber estimate identifies every other zero-residual state in the larger outer fiber ball with $Y(s)$.

The constructed base endpoint solves the actual equation

$$\mathcal{F}(s, y) = \mathcal{F}(s_*, y_0).$$

For weighted smooth input, its right side lies in every $\mathcal{Z}_m$. Apply the same spatial induction as in Theorem A.7, now with this prescribed right side. If the existing correction has order j, the principal coefficients have the required order $j - 1$; the prescribed metric seed is used through order $j + 1$, the momentum seed through order j, and the residual through order $j - 1$. All other terms contain only the already controlled correction derivatives. Interior system regularity improves the existing solution by one order. Rescale nested annuli exactly as in that theorem; the factors R^γ stay bounded and the fixed base ellipticity margin is unchanged. Taking the supremum and adding the inner-ball estimate gives the next weighted norm. Each finite induction uses finitely many input norms. This proves all spatial orders for the one base endpoint.

For two endpoints $y, \widetilde{y}$, integrate the derivative of $\mathcal{F}$ along the ambient segment joining (s, y) to $(t, \widetilde{y})$. The exact equation is

$$\overline{\mathsf{A}}(y - \widetilde{y}) = -\overline{\mathsf{B}}(s - t) + \mathcal{F}(s_*, y_0) - \mathcal{F}(s_*, \widetilde{y}_0).$$

The segment lies in D. Proposition A.13 applies to its averaged transverse operator; the preceding spatial bounds control the finite coefficients. The residual difference at fixed s_* has $\mathcal{Z}_m$ norm controlled by the $\mathcal{Y}_m$ label difference, using finitely many reference and label bounds. The seed term has the corresponding $\mathcal{X}_m \to \mathcal{Z}_m$ estimate. This proves the higher difference bound and continuity at every weighted order.

At any already obtained smooth endpoint, the same transverse operator is an isomorphism at each higher order. Apply the Banach implicit function theorem in that higher space to the endpoint equation near this endpoint; base uniqueness identifies its local solution with the constructed one. This identifies higher-space parameter derivatives without imposing smallness on high norms. Differentiating at each finite parameter order leaves the same invertible transverse operator on the left and finite expressions in input and lower output derivatives on the right. Induction proves smooth family dependence. The finite input dependence is uniform on $\mathcal{K}$ when measured on the larger retained set $\mathcal{K}'$, by a finite cover and the fixed margins. □

For the flat initial label, the proof constructs $Y(s)$ from current differential operations and the executed inverse. Uniqueness in Theorem A.6 then gives $Y(s) = \mathcal{C}(s)$. This is an independent check of the endpoint against the

earlier constraint chart. The transport iteration itself uses only the current operations.

Corollary A.15 (Controlled finite execution on the flat branch). Take $(s_*, y_*) = (0, 0)$. Truncate the current inverse through power N and set

$$\mathsf{V}_Z^{[N]} a = -\sum_{j=0}^{N} (-\mathsf{E}_Z)^j \mathsf{G}_0^{\mathrm{src}} \mathsf{B}_Z a.$$

Starting at the zero path, perform J Picard steps along the radial seed path ts with this field, and denote the endpoint by $Y_{J,N}(s)$. After one fixed reduction of the seed radius and for N sufficiently large, all these iterates remain in the retained product and

$$\|Y_{J,N}(s) - Y(s)\|_{\mathcal{Y}_{k_0}} \le C\|s\|_{\mathcal{X}_{k_0}} (q_{\mathrm{P}}^J + \varkappa^{N+1}).$$

Their endpoint residuals obey the same bound in $\mathcal{Z}_{k_0}$. Every finite endpoint has exactly the seed s. For weighted smooth input, each finite stage is an actual weighted smooth field obtained by finite current differential operations, finite negative-response words and parameter integration. The higher bounds may depend on J, N. The convergence estimate is in the base norm.

Proof. Uniformly on the fixed base product, the geometric tail gives

$$\|(\mathsf{V}_Z^{[N]} - \mathsf{V}_Z) a\|_{\mathcal{Y}_{k_0}} \le \frac{\varkappa^{N+1}}{1-\varkappa} \|\mathsf{G}_0^{\mathrm{src}} \mathsf{B}_Z a\|_{\mathcal{Y}_{k_0}} \le C\varkappa^{N+1} \|a\|_{\mathcal{X}_{k_0}}.$$

Choose a fixed margin $\epsilon_0 > 0$, reduce the seed radius once so that $(M+\epsilon_0)\rho < r_y/4$, and choose N with the last operator error at most ϵ_0. The approximate integral map preserves the same path ball. Comparing an approximate iterate to the exact path, with error e_J in the base path norm, gives

$$e_{J+1} \le q_{\mathrm{P}} e_J + C\varkappa^{N+1} \|s\|_{\mathcal{X}_{k_0}}, \qquad e_0 \le C\|s\|_{\mathcal{X}_{k_0}}.$$

Summation proves the error bound. The residual bound follows from the bounded derivative of $\mathcal{F}$ in the correction variable and the zero residual of the exact endpoint.

The iteration changes only the correction coordinate over the prescribed seed path. Its endpoint therefore has seed s independently of its residual. Each finite inverse word on smooth residuals is actual by the source Green theorem, including the domain of every auxiliary call. The coefficient derivatives preserve weighted smoothness, and integration on the finite parameter interval preserves every fixed seminorm. Induction proves actual smoothness of every finite stage. Theorem A.14 gives higher regularity directly for the exact endpoint. □

Appendix B. Relative maximal extension on one smooth branch

Chruściel and Cong extend smooth vacuum data across a boundary, preserving the original metric and second fundamental form (Theorem 6.4). For sufficiently small maximal vacuum data on a ball, Czimek already gives a

maximal asymptotically flat extension on $\mathbb{R}^3$ agreeing with the data on that ball. Mao, Oh and Tao construct constraint gluings using linear solution operators with prescribed support.

Here the correction is normalized relative to a supplied whole-space extension of the interior fields and vanishes when that extension already solves the constraints. We keep the interior data fixed and allow a decaying exterior tail, without prescribing exterior fields or ADM charges. Czimek's explicit exterior right inverses and the nonlinear construction below give existence. Their ranges supply the estimates at weights slightly below $-1/2$ and the higher spatial and parameter estimates on the same selected branch. Write $R(h) = \mathrm{Scal}(h)$.

B.1. Weighted spaces and fixed exterior ranges. Put $\Omega = \mathbb{R}^3 \setminus \overline{B}_1$ and write $\varrho = |x|$ for the exterior radial coordinate. All tensor coefficients and derivatives in the weighted norm use the fixed Euclidean coordinates. For an integer $s \geq 0$ set

$$\|u\|^2_{H^s_\omega} = \sum_{|\alpha| \leq s} \int_{\mathbb{R}^3} \langle x \rangle^{-2\omega - 3 + 2|\alpha|} |\partial^\alpha u(x)|^2 \, \mathrm{d}x.$$

On Ω, $\overline{H}^s_\omega$ is the closure of $C^\infty_c(\Omega)$ in this norm. Equivalently, the zero extension belongs to the whole-space weighted Sobolev space. The bar records the vanishing boundary traces; the resulting field can have a noncompact decaying tail.

The first flat operator is

$$P_1 v = \mathrm{div}_\delta \, v, \qquad \mathrm{tr}_\delta \, v = 0.$$

For the second, write a metric in polar form as

$$\widehat{h} = a^2 \, \mathrm{d}\varrho^2 + \widehat{\gamma}_{AB}(b^A \mathrm{d}\varrho + \mathrm{d}\theta^A)(b^B \mathrm{d}\varrho + \mathrm{d}\theta^B)$$

and define

$$\mathfrak{V}_{\widehat{h}}(\phi, \beta) = a^2 \, \mathrm{d}\varrho^2 + e^{2\phi} \widehat{\gamma}_{AB} \big((b + \beta)^A \mathrm{d}\varrho + \mathrm{d}\theta^A \big) \big((b + \beta)^B \mathrm{d}\varrho + \mathrm{d}\theta^B \big),$$

where β is tangent to the coordinate spheres. Set

$$P_2(\phi, \beta) = DR_\delta [D\mathfrak{V}_\delta(0, 0)(\phi, \beta)] \, .$$

Fix the explicit spherical-harmonic inverses B_1, B_2 of the flat exterior constructions in Czimek (2018, Lemmas 4.6 and 5.10), with the sign of the scalar residual chosen to agree with this P_2. The particular inverse formulas and their ranges are retained. They are not replaced by arbitrary right inverses of the same differential operators.

Choose a fixed smooth radial function ℓ that vanishes near $\overline{B}_1$ and equals $\log \varrho$ for $\varrho \geq 2$, and put $w_{\varepsilon_\gamma} = e^{-\varepsilon_\gamma \ell}$. Lemma B.2 supplies $\varepsilon_* > 0$ such that, after fixing

$$0 < \varepsilon_\gamma < \min(\varepsilon_*, 1/2), \qquad \gamma = -\frac{1}{2} - \varepsilon_\gamma,$$

the conjugated inverses

$$B_{j,\gamma} = w_{\varepsilon_\gamma} B_j (I + E_{j,\varepsilon_\gamma} B_j)^{-1} w_{\varepsilon_\gamma}^{-1}, \qquad E_{j,\varepsilon_\gamma} = w_{\varepsilon_\gamma}^{-1} P_j w_{\varepsilon_\gamma} - P_j, \quad j = 1, 2,$$

are bounded right inverses with the mappings

$$B_{1,\gamma} : \overline{H}^{s-2}_{\gamma-2}(T^*) \longrightarrow \overline{H}^{s-1}_{\gamma-1}(\mathrm{Sym}^2_0 T^*),$$
$$B_{2,\gamma} : \overline{H}^{s-2}_{\gamma-2} \longrightarrow \overline{H}^{s}_{\gamma} \oplus \overline{H}^{s}_{\gamma}(TS_\varrho), \qquad s \geq 2.$$

The same operators are used at every integer order. Their residual-space Neumann inverses are selected at one base order; the fact that E_{j,ε_γ} has one lower differential order gives all higher estimates without reducing ε_γ. On their base spaces the exact range identities are

$$\mathrm{Ran}\, B_{j,\gamma} = w_{\varepsilon_\gamma} \mathrm{Ran}\, B_j, \qquad j = 1, 2.$$

The inverses preserve real fields, zero-extension traces and proper Euclidean rotations. The rotation statement uses the fixed radial cutoffs and the same radial formulas throughout each harmonic multiplicity space. Lemmas B.2 and B.5 prove these properties for the chosen normalization. They are needed for the regularity of the selected nonlinear branch.

B.2. The relative-extension theorem.

Theorem B.1 (Relative maximal extension on one smooth branch).

Fix the weight and inverse ranges above and an integer $S_0 \geq 4$. Increase S_0 beforehand if subsequent finite-order chart estimates require it. There is a base threshold $\epsilon > 0$, chosen with fixed positivity and inverse margins, with the following property.

Let $(\widehat{h}, \widehat{K})$ be smooth whole fields on $\mathbb{R}^3$ such that

$$\widehat{h} - \delta \in H^M_\gamma, \qquad \widehat{K} \in H^{M-1}_{\gamma-1} \quad \text{for every integer } M \geq S_0,$$

and

$$\|\widehat{h} - \delta\|_{H^{S_0}_\gamma} + \|\widehat{K}\|_{H^{S_0-1}_{\gamma-1}} < \epsilon.$$

Assume that $\widehat{h}$ is positive, that $\mathrm{tr}_{\widehat{h}}\, \widehat{K} = 0$ on the whole space, and that on B_1

$$R(\widehat{h}) = |\widehat{K}|^2_{\widehat{h}}, \qquad \mathrm{div}_{\widehat{h}}\, \widehat{K} = 0.$$

The initialized relative operation selects a whole pair

$$(h, K) = \mathfrak{E}(\widehat{h}, \widehat{K})$$

with

$$(h, K)|_{B_1} = (\widehat{h}, \widehat{K})|_{B_1},$$
$$R(h) = |K|^2_h, \qquad \mathrm{div}_h\, K = 0, \qquad \mathrm{tr}_h\, K = 0,$$
$$h - \delta \in H^M_\gamma, \qquad K \in H^{M-1}_{\gamma-1} \quad \text{for every integer } M \geq S_0.$$

The corrections $h - \widehat{h}$ and $K - \widehat{K}$ have the corresponding zero-extension regularity. At the base order their whole output norm is bounded by a fixed

multiple of the whole input norm in the smallness condition. At each higher finite order, the output norm is bounded in terms of finitely many higher input norms and the fixed base margins. No higher norm is required to be small. The operation fixes every supplied whole pair in its domain that already satisfies the maximal vacuum constraints everywhere.

Let the prefields form a smooth family on a finite-dimensional parameter domain, smooth in all the indicated weighted seminorms and remaining in the selected base neighborhood. Then the same output is smooth at every finite weighted order. On a compact parameter set contained in the interior of a slightly larger retained compact set, every finite number of parameter derivatives in any specified output order is bounded by finitely many weighted input seminorms and parameter seminorms on that larger set, together with the fixed margins. The branch and base neighborhood are chosen before the requested derivative orders.

The proof is completed in §B.8. The intervening results specify the inverse ranges and establish the estimates for this selected operation.

B.3. Fixed inverses, nonlinear repairs and the base branch. The sphere-tangent variable β is a vector field on the Euclidean exterior, with its Cartesian tensor norm. Thus its norm is the norm of the vector field itself, rather than of its components in an unscaled angular coordinate basis. The fixed Euclidean metric on each sphere is $\gamma^\circ(\varrho) = \delta|_{TS_\varrho} = \varrho^2 g_{S^2}$; in particular its Laplacian includes the factor ϱ^{-2} relative to the unit-sphere Laplacian. With these conventions the second flat operator is

$$P_2(\phi,\beta) = -4\partial_\varrho^2\phi - \frac{12}{\varrho}\partial_\varrho\phi - \frac{4}{\varrho^2}\phi - 2\Delta_{\gamma^\circ}\phi + \frac{2}{\varrho^3}\partial_\varrho(\varrho^3\,\mathrm{div}_{\gamma^\circ}\,\beta).$$

The external linear input is the constructive part of Czimek (2018, Lemmas 4.6 and 5.10 and their proofs in §§4.3 and 5.3). Fix the transition function in those constructions, using the same radial formula throughout each spherical-harmonic multiplicity space. They give real linear right inverses

$$B_1 : \overline{H}^q_{-5/2}(T^*) \longrightarrow \overline{H}^{q+1}_{-3/2}(\mathrm{Sym}^2_0\,T^*),$$
$$B_2 : \overline{H}^q_{-5/2} \longrightarrow \overline{H}^{q+2}_{-1/2} \oplus \overline{H}^{q+2}_{-1/2}(TS_\varrho), \qquad q \geq 0,$$

bounded at every integer order and defined by the same formulas at those orders. Part (1) of each cited lemma gives the base estimate $q = 0$. For $q \geq 1$, set the regularity index in part (2) to $q+2$: the tensor output has order $q+1$ in Lemma 4.6, and both outputs have order $q+2$ in Lemma 5.10. These are estimates in the barred spaces for the fields constructed in the respective proofs. In particular, Proposition 5.17 gives the boundary regularity of the sphere-tangent field ξ constructed by the angular Hodge equations, including its membership in the barred space. The transition function and mode formulas are fixed before the regularity index is chosen. With these choices the normalization here is

$$P_j B_j = I, \qquad j = 1, 2.$$

The equation displayed in the cited Lemma 5.10 uses the negative of the curvature derivative displayed in its Lemma 5.7. Accordingly, its constructive solution is applied to the negative residual when defining B_2. This fixes a sign and leaves its range unchanged. The concrete ranges of these constructive maps are retained; the Hilbert orthogonal complements introduced separately in the cited implicit-function arguments are not used.

The equation numbers in the next two paragraphs refer to Czimek, arXiv:1609.08814v2. On the momentum side, (4.39)–(4.56) construct the radial scalar δ, the mixed field ϵ and the tangential tracefree field $\widehat{\eta}$ of the *same* tensor k. Lemma 4.11 verifies the equations. Lemma 4.16, (4.93), bounds this k at order $w-1$ by the residual at orders $w-2$ and zero, and puts it in the barred space of weight $-3/2$. The auxiliary curl tensor has order $w-2$ and weight $-5/2$. Taking $w = q+2$ gives the asserted estimate on the explicit image of B_1. The low degrees in δ and ϵ remain in this inverse. Equations (4.53) and (4.56) specify the exact component of the angular divergence of $\widehat{\eta}$ only in degrees $l \geq 2$. The half-residual terms cancel on that projection. In the physical sphere frame the angular derivatives acquire ϱ^{-1}, giving the third equation in Lemma B.10. There is no tracefree tensor harmonic of degree zero or one.

On the scalar side, multiplying (5.31) by four identifies its left side with $-P_2$ and its right side with the cited residual h. Equations (5.33) and (5.34) select the radial and higher-mode scalar solutions. Equations (5.35)–(5.36) fix the two cutoff moments, including the coefficient $-1/2$ multiplying the residual moment. Equation (5.37) fixes the sphere-tangent vector by the radial divergence equation, vanishing angular curl and zero boundary value. Lemma 5.13 checks the equation. Proposition 5.14 gives the two moment cancellations (5.39)–(5.40), and the barred estimate (5.42) for $u^{[\geq 1]}$. Proposition 5.17 gives (5.52) for $u^{[0]}$ and (5.53) for ξ, both in the barred space of weight $-1/2$ at order w. Thus all required zero-extension traces through order $w-1$ for the scalar outputs (through $w-2$ for the momentum output) follow at that finite order. The critical-weight estimates concern these fixed linear solutions. Lemma B.2 passes to the stronger weight by conjugation and identifies the image there. Lemmas B.8–B.11 use the resulting image equations and angular control to raise the regularity of the one nonlinear solution selected at base order.

Lemma B.2 (One stronger weight at every order).

There is $\varepsilon_* > 0$ such that each choice $0 < \varepsilon_\gamma < \min(\varepsilon_*, 1/2)$ gives the operators $B_{j,\gamma}$ in the stated exterior spaces, with $\gamma = -1/2 - \varepsilon_\gamma$. They are bounded right inverses at every integer order, use one fixed formula, and preserve real fields, zero-extension regularity and proper Euclidean rotations. Their base ranges satisfy

$$\operatorname{Ran} B_{j,\gamma} = w_{\varepsilon_\gamma} \operatorname{Ran} B_j.$$

The threshold ε_* is chosen at order zero for the residual and is independent of every subsequent regularity order.

Proof. Write $m_1 = 1$ and $m_2 = 2$. For this proof let Y_j^q be the corresponding residual space at order q and weight $-5/2$, and let Z_j^q be the output space at order $q + m_j$ and weight $-5/2 + m_j$. Thus $B_j : Y_j^q \to Z_j^q$. The scalar output has two components.

Choose the fixed radial ℓ to vanish on a neighborhood of $\overline{B}_1$ and equal $\log \varrho$ for $\varrho \geq 2$. It satisfies $|\ell| \leq C(1 + \log\langle x\rangle)$, and its positive-order derivatives satisfy $|\partial^\alpha \ell| \leq C_\alpha \langle x\rangle^{-|\alpha|}$. Consequently multiplication by $w = e^{-\varepsilon_\gamma \ell}$ is an isomorphism $H^k_\omega \to H^k_{\omega - \varepsilon_\gamma}$ at every finite k, with inverse multiplication by w^{-1}. This follows by differentiating the product and using $|\partial^\alpha w| \leq C_{\alpha,\varepsilon_\gamma} \langle x\rangle^{-\varepsilon_\gamma - |\alpha|}$ and the corresponding inverse estimate. Both multipliers preserve the zero-extension subspaces and all algebraic tensor types.

In Cartesian components, a coefficient of differential order d in P_j has size $O(\varrho^{d-m_j})$, with the corresponding bound after each derivative. For P_1 this is immediate; for P_2 it follows either from its displayed polar expression or from the degree-zero coefficients of the linearized polar metric variation. Leibniz' rule therefore gives

$$w^{-1} P_j w = P_j + E_j, \qquad E_j = \sum_{|\alpha| \leq m_j - 1} e_{j,\alpha}(x) \partial^\alpha,$$

where $E_j = 0$ near $\overline{B}_1$ and, for $0 < \varepsilon_\gamma \leq 1/2$,

$$|\partial^\nu e_{j,\alpha}(x)| \leq C_{\alpha,\nu} \varepsilon_\gamma \langle x\rangle^{|\alpha| - m_j - |\nu|}.$$

In particular,

$$\|E_j B_j\|_{Y_j^0 \to Y_j^0} \leq C \varepsilon_\gamma.$$

Choose ε_* so that the two norms are less than $1/2$. The inverse $Q_j = (I + E_j B_j)^{-1}$ is then defined on Y_j^0 by its convergent Neumann series. Define

$$\widetilde{B}_j = B_j Q_j, \qquad B_{j,\gamma} = w \widetilde{B}_j w^{-1}.$$

The identity $(P_j + E_j)\widetilde{B}_j = I$ proves the required right-inverse equation after conjugation.

It remains to prove higher bounds for this same operator. The coefficient estimate above has a useful precise consequence:

$$E_j : Z_j^q \longrightarrow Y_j^{q+1}$$

is bounded for each $q \geq 0$. Indeed, a derivative of order $|\alpha| \leq m_j - 1$ leaves at least $q + 1$ derivatives of the output available. Its weight becomes $-5/2 + m_j - |\alpha|$; the coefficient has the compensating weight $|\alpha| - m_j$. Differentiating the coefficient preserves this balance. Moreover, E_j vanishes near the inner sphere, so the extra residual derivative does not require an extra trace of the input there.

Let $u = \widetilde{B}_j f$ be the solution already constructed in Z_j^0. It satisfies

$$u = B_j f - B_j E_j u.$$

If $f \in Y_j^{q+1}$ and $u \in Z_j^q$, both terms on the right belong to Z_j^{q+1}, with

$$\|u\|_{Z_j^{q+1}} \le C_{q+1}\big(\|f\|_{Y_j^{q+1}} + \|u\|_{Z_j^q}\big).$$

Induction proves all higher estimates. Only their constants depend on the order; the Neumann inverse and the chosen ε_γ remain fixed. Multiplication by w gives precisely the claimed shifted weights.

The base range identity follows directly from the surjectivity of $Q_j : Y_j^0 \to Y_j^0$:

$$\operatorname{Ran} \widetilde{B}_j = \operatorname{Ran} B_j, \qquad \operatorname{Ran} B_{j,\gamma} = w \operatorname{Ran} B_j.$$

Every step preserves real fields and zero-extension regularity. Proper rotations commute with P_j, the fixed radial multipliers, and the constructive flat inverses: the radial formulas depend on harmonic degree, while the angular projections and oriented Hodge operators commute with those rotations. They therefore commute with E_j, its Neumann inverse and $B_{j,\gamma}$. Finally, at each fixed order $B_{j,\gamma}P_j$ is a bounded projection onto $\operatorname{Ran} B_{j,\gamma}$, because $P_j B_{j,\gamma} = I$. The selected range is therefore closed in the indicated output space; limits of correction iterates remain in that range. □

Exterior repairs at one fixed Sobolev order. Fix a finite $S \ge 4$ and the weight chosen in Lemma B.2. Write

$$\begin{aligned}
\mathcal{X}_S &= H^S_\gamma(\operatorname{Sym}^2 T^*),\\
\mathcal{V}_S &= \overline{H}^{S-1}_{\gamma-1}(\operatorname{Sym}^2_0 T^*),\\
\mathcal{Z}_S &= \overline{H}^S_\gamma \oplus \overline{H}^S_\gamma(TS_\varrho),\\
\mathcal{Y}^1_S &= \overline{H}^{S-2}_{\gamma-2}(T^*), \qquad \mathcal{Y}^0_S = \overline{H}^{S-2}_{\gamma-2}.
\end{aligned}$$

Metrics are written as $\delta + \mathcal{X}_S$. The bar is omitted when a tensor is defined on the whole space; norms of zero-extended tensors agree with the corresponding exterior norms.

Lemma B.3 (Normalized exterior repairs).
On a sufficiently small neighborhood at this fixed order there are smooth operations

$$\mathfrak{D}(h,\rho) = \mathrm{TF}_h\, v, \qquad v \in \operatorname{Ran} B_{1,\gamma},$$

and

$$\mathfrak{G}(\widehat{h}, s) = \mathfrak{V}_{\widehat{h}}(z), \qquad z \in \operatorname{Ran} B_{2,\gamma},$$

such that

$$\operatorname{div}_h \mathfrak{D}(h,\rho) = \rho, \qquad \operatorname{tr}_h \mathfrak{D}(h,\rho) = 0, \qquad \mathfrak{D}(h,0) = 0,$$
$$R(\mathfrak{G}(\widehat{h}, s)) = R(\widehat{h}) + s, \qquad \mathfrak{G}(\widehat{h}, 0) = \widehat{h}.$$

Both corrections have zero-extension regularity and preserve the supplied fields on B_1. On a smaller retained neighborhood they satisfy

$$\|\mathfrak{D}(h,\rho)-\mathfrak{D}(h',\rho')\|_{H^{S-1}_{\gamma-1}} \le C_S(\|\rho-\rho'\|_{\mathcal{Y}^1_S}+\|h-h'\|_{\mathcal{X}_S}\|\rho'\|_{\mathcal{Y}^1_S}),$$
$$\|\mathfrak{G}(\widehat{h},s)-\mathfrak{G}(\widehat{h},s')\|_{\mathcal{X}_S} \le C_S\|s-s'\|_{\mathcal{Y}^0_S},$$
$$\|\mathfrak{G}(\widehat{h},s)-\widehat{h}\|_{\mathcal{X}_S} \le C_S\|s\|_{\mathcal{Y}^0_S}.$$

The neighborhoods and constants here may depend on S. These operations agree at different finite orders whenever both are defined on the same input.

Proof. We first record the estimates that control the two nonlinear equations. On a dyadic annulus $\varrho \asymp 2^\nu$, rescale to a fixed annulus. The weighted norm is equivalent to the square sum of $2^{-\nu\omega}$ times the ordinary Sobolev norms of the rescaled Cartesian coefficients. The usual three-dimensional product estimate on the fixed annulus, followed by this square summation, gives

$$\|uv\|_{H^{S-2}_{\omega_1+\omega_2}} \le C_S\|u\|_{H^S_{\omega_1}}\|v\|_{H^{S-2}_{\omega_2}}.$$

The same argument applies when both factors have $S-1$ derivatives and the output has at most $S-1$ derivatives. Here $S-2>3/2$, so each lowest undifferentiated Sobolev space needed in these products is an algebra. Smooth composition with matrix inversion near the identity, with the exponential near zero, and with the tracefree projection follows by differentiating these finite-dimensional functions and applying the same product estimate. Differentiation lowers the weight by one. A factor of negative weight can also be included without weakening the desired decay.

For clarity, these estimates also apply to the polar variation. The Euclidean radial vector, radial one-form and tangential projection have bounded coefficients with derivatives of order k bounded by $C_k\varrho^{-k}$ on Ω. The polar variables of h are obtained algebraically from its radial and tangential blocks, with the tangential block inverted on the tangent planes. Its inverse is uniformly bounded after rescaling each sphere. In Cartesian tensor components the map $(h,z)\mapsto \mathfrak{V}_h(z)-h$ is therefore smooth $(\delta+\mathcal{X}_S)\times\mathcal{Z}_S\to\mathcal{X}_S$ near $(\delta,0)$, vanishes when $z=0$, and has the corresponding bounded derivatives. This avoids assigning Cartesian weights to bare angular coordinate components.

The Cartesian scalar curvature formula consists of terms $h^{-1}h^{-1}\partial^2 h$ and $h^{-1}h^{-1}h^{-1}(\partial h)(\partial h)$ with fixed index contractions. The preceding estimates make R a smooth map from the metric neighborhood to $H^{S-2}_{\gamma-2}$. They also make

$$L_h v := \operatorname{div}_h \operatorname{TF}_h v$$

a bounded map $\mathcal{V}_S\to\mathcal{Y}^1_S$, smooth in h. Since $L_\delta = P_1$ on Euclidean tracefree tensors, one has

$$\|(L_h-P_1)v\|_{\mathcal{Y}^1_S} \le C_S\|h-\delta\|_{\mathcal{X}_S}\|v\|_{\mathcal{V}_S},$$
$$\|(L_h-L_{h'})v\|_{\mathcal{Y}^1_S} \le C_S\|h-h'\|_{\mathcal{X}_S}\|v\|_{\mathcal{V}_S}.$$

For the scalar equation put

$$\mathcal{S}_h(z) = R(\mathfrak{V}_h(z)) - R(h), \qquad \mathcal{N}_h(z) = \mathcal{S}_h(z) - P_2 z.$$

The smoothness just proved, $\mathcal{S}_h(0) = 0$, and $D_z\mathcal{S}_\delta(0) = P_2$ give, on a fixed small ball,

$$\|\mathcal{N}_h(z) - \mathcal{N}_h(z')\|_{\mathcal{Y}^0_S} \le C_S(\|h - \delta\|_{\mathcal{X}_S} + \|z\|_{\mathcal{Z}_S} + \|z'\|_{\mathcal{Z}_S})\|z - z'\|_{\mathcal{Z}_S}.$$

Indeed, integrate $D_z\mathcal{S}_h - P_2$ along the segment from z' to z and use the bounded mixed and second derivatives. All these residuals have the indicated barred regularity. A zero-extended correction is a whole Sobolev tensor vanishing on B_1; the local differential expressions formed from it are whole Sobolev tensors vanishing there as well.

For the momentum equation, restrict the unknown by writing $v = B_{1,\gamma} f$. Set

$$T_h = (L_h - P_1)B_{1,\gamma} : \mathcal{Y}^1_S \longrightarrow \mathcal{Y}^1_S.$$

After shrinking the metric neighborhood, $\|T_h\| < 1/2$. The required solution is explicitly

$$v = B_{1,\gamma}(I + T_h)^{-1}\rho, \qquad \mathfrak{D}(h, \rho) = \mathrm{TF}_h\, v.$$

This proves its equations, range and zero normalization. The resolvent identity

$$(I + T_h)^{-1} - (I + T_{h'})^{-1} = (I + T_h)^{-1}(T_{h'} - T_h)(I + T_{h'})^{-1}$$

and the estimates for L_h and TF_h prove the first difference bound. They also prove smooth dependence on h and ρ.

For the scalar equation use the iteration

$$z_{n+1} = B_{2,\gamma}\big(s - \mathcal{N}_{\widehat{h}}(z_n)\big), \qquad z_0 = 0.$$

Choose the metric and residual neighborhoods so that, on $\|z\|_{\mathcal{Z}_S} \le 2\|B_{2,\gamma}\|\|s\|_{\mathcal{Y}^0_S}$, the displayed remainder estimate gives contraction constant less than $1/2$. Since $\mathcal{N}_{\widehat{h}}(0) = 0$, the same estimate proves that this ball is invariant. The limit has $\|z\| \le 2\|B_{2,\gamma}\|\|s\|$, remains in the fixed inverse range, and satisfies $\mathcal{S}_{\widehat{h}}(z) = s$ after applying P_2. For $s = 0$ every iterate is zero. Difference of two fixed-point equations, followed by absorption of the contraction term, gives the scalar residual difference estimate; the smooth metric variation then gives both stated estimates for $\mathfrak{G}$.

Smooth dependence follows in the residual variable as well. The equation $\mathcal{S}_{\widehat{h}}(B_{2,\gamma} f) = s$ has derivative I in f at the flat zero input. On the selected small ball its derivative remains invertible by a Neumann series, so the Banach implicit-function theorem supplies all finite parameter derivatives. The initialized solution above is that branch by uniqueness. The zero extensions of z make $\mathfrak{G} - \widehat{h}$ a whole tensor in H^S_γ vanishing on B_1.

For two finite orders the momentum solution is also obtained by the iteration

$$v_{n+1} = B_{1,\gamma}(\rho - (L_h - P_1)v_n), \qquad v_0 = 0.$$

Both this iteration and the scalar iteration use the same operators and initial values, independently of the norm in which convergence is proved. If their input lies in both selected neighborhoods, the higher-order limit and the lower-order limit coincide by lower-order uniqueness. This proves compatibility. □

Proposition B.4 (The initialized relative extension at the base order).

For the prefields of Theorem B.1, choose the threshold at the one fixed order $S = S_0$. The repairs of Lemma B.3 determine an initialized operation $\mathfrak{E}$ whose output is a whole maximal vacuum datum in $H^{S_0}_{\gamma} \oplus H^{S_0-1}_{\gamma-1}$, with the metric measured relative to δ. It agrees with the prefields on B_1, its corrections have the corresponding zero-extension regularity, and its base output norm is bounded by a fixed multiple of the base input norm. It fixes already exact whole data. On every supplied smooth finite-dimensional family in its domain it is smooth in the base spaces, with finite parameter and difference estimates. These assertions require only the fixed base threshold.

Proof. The construction is carried out at $S = S_0$. All later regularity arguments concern this initialized output.

Let $d = (\widehat{h}, \widehat{K})$ satisfy the hypotheses of Theorem B.1, and set

$$\widehat{a} = \widehat{K} - \frac{\operatorname{tr}_\delta \widehat{K}}{\operatorname{tr}_\delta \widehat{h}}\widehat{h}.$$

The denominator is bounded away from zero by the base metric margin. One has $\operatorname{tr}_\delta \widehat{a} = 0$ and $\mathrm{TF}_{\widehat{h}}\, \widehat{a} = \widehat{K}$: the second identity follows by taking the $\widehat{h}$ trace and using $\operatorname{tr}_{\widehat{h}} \widehat{K} = 0$. In particular,

$$\|\widehat{a}\|_{H^{S-1}_{\gamma-1}} \le C_S \|\widehat{K}\|_{H^{S-1}_{\gamma-1}}.$$

Consider candidate whole fields (h, K) agreeing with d on B_1, and define

$$\begin{aligned} A_d(h) &= \mathrm{TF}_h\, \widehat{a}, \\ M_d(h) &= A_d(h) + \mathfrak{D}\big(h, -\operatorname{div}_h A_d(h)\big), \\ H_d(h, K) &= \mathfrak{G}\big(\widehat{h}, |K|^2_h - R(\widehat{h})\big). \end{aligned}$$

The first momentum residual vanishes on B_1 because $A_d(\widehat{h}) = \widehat{K}$ there. The scalar residual vanishes there by the local Hamiltonian constraint. Both are whole tensors in $H^{S-2}_{\gamma-2}$: for the quadratic term the sharper product weight is $2\gamma - 2 < \gamma - 2$. Their restrictions therefore belong to $\overline{H}^{S-2}_{\gamma-2}(\Omega)$. Whole-space Sobolev regularity, together with vanishing throughout the ball, gives every trace required for zero extension. Boundary values alone would control only the zeroth trace.

Lemma B.3 gives

$$\operatorname{div}_h M_d(h) = 0, \qquad \operatorname{tr}_h M_d(h) = 0,$$

and both blocks preserve the prescribed interior. The same estimates imply, for candidate fields in a fixed small neighborhood,

$$\begin{aligned}\|M_d(h) - M_d(h')\|_{H^{S-1}_{\gamma-1}} &\le C_S\|\widehat{a}\|_{H^{S-1}_{\gamma-1}}\|h - h'\|_{H^S_\gamma},\\ \|H_d(h,K) - H_d(h',K')\|_{H^S_\gamma} &\le C_S\big(\|K\|_{H^{S-1}_{\gamma-1}} + \|K'\|_{H^{S-1}_{\gamma-1}}\big)\|K - K'\|_{H^{S-1}_{\gamma-1}}\\ &\quad + C_S\|K'\|^2_{H^{S-1}_{\gamma-1}}\|h - h'\|_{H^S_\gamma}.\end{aligned}$$

For the first estimate, $A_d(h)$ and each of its metric variations are linear in $\widehat{a}$. Thus $\|\operatorname{div}_h A_d(h)\| \le C_S\|\widehat{a}\|$ and the difference of these residuals is at most $C_S\|\widehat{a}\|\|h - h'\|$. Apply the momentum difference estimate in Lemma B.3, retaining its residual factor. For the second, expand the difference of the two quadratic norms, use smooth dependence of h^{-1}, and apply the scalar residual difference estimate with the same fixed premetric $\widehat{h}$.

Define

$$\mathcal{T}_d(h,K) = \big(H_d(h,K), M_d(H_d(h,K))\big).$$

Here is an explicit way to choose its invariant ball. Put $e = \|\widehat{h} - \delta\|_{H^S_\gamma} + \|\widehat{K}\|_{H^{S-1}_{\gamma-1}}$. The scalar repair estimate and the scalar-curvature product estimate give

$$\|H_d(h,K) - \delta\|_{H^S_\gamma} \le C_S e + C_S\|K\|^2_{H^{S-1}_{\gamma-1}},$$

whereas $\|M_d(h)\|_{H^{S-1}_{\gamma-1}} \le C_S e$ on the fixed metric neighborhood. Choose a fixed C_0 larger than twice the sum of these linear constants. After choosing the base threshold for e sufficiently small, $\mathcal{T}_d$ maps the closed set

$$\left\{(h,K) : \begin{array}{c}(h,K)|_{B_1} = d|_{B_1},\\ \|h - \delta\|_{H^S_\gamma} + \|K\|_{H^{S-1}_{\gamma-1}} \le C_0 e\end{array}\right\}$$

to itself, and all metrics in that set have the chosen positivity margin. The two difference estimates make its Lipschitz constant less than $1/2$ after a further base restriction. The restriction depends on S, the fixed inverses and the margins, with no reference to any higher order.

Starting at $(h_0, K_0) = (\widehat{h}, \widehat{K})$, the iteration

$$(h_{n+1}, K_{n+1}) = \big(H_d(h_n, K_n), M_d(H_d(h_n, K_n))\big)$$

converges in the base space to a unique fixed point in this ball. At the fixed point the defining equations of the repairs give all three maximal constraints. The invariant-ball estimate gives the claimed base norm bound. Every iterate agrees with d on B_1, so the limit does as well; its differences from d are whole Sobolev fields vanishing there and hence have the claimed barred regularity.

If d already satisfies all maximal constraints on the whole space, then at $(\widehat{h}, \widehat{K})$ both residuals are zero. The identities $\mathfrak{D}(h, 0) = 0$ and $\mathfrak{G}(h, 0) = h$ show that the first iterate equals the initial value. Every later iterate and the selected output are therefore equal to d.

For base parameter smoothness, use the fixed Banach space of corrections

$$\mathcal{C}_S = \overline{H}^S_\gamma(\Omega; \operatorname{Sym}^2 T^*) \oplus \overline{H}^{S-1}_{\gamma-1}(\Omega; \operatorname{Sym}^2 T^*).$$

If $d = d(p)$ is a supplied smooth family satisfying the local constraints, write the fixed-point equation in the form

$$\zeta = \mathcal{T}_{d(p)}(d(p) + \zeta) - d(p), \qquad \zeta \in \mathcal{C}_S.$$

The right side is smooth in p and ζ. It takes values in the same fixed correction space because the residuals vanish identically on B_1 for every p and every zero-extended correction; their parameter derivatives do so as well. Its derivative in ζ has norm less than $1/2$. Hence the implicit-function theorem gives a smooth base solution. The inverse derivative is bounded by 2. Differentiating the equation q times gives the same inverse applied to finite sums of derivatives of the block maps, input derivatives up to order q, and lower derivatives of ζ. Induction gives all specified finite parameter bounds at the base order on a retained compact parameter set.

Differences of two locally constrained prefields can be estimated without drawing a line segment inside the constrained input set. Write their fixed-point corrections as $\zeta_d, \zeta_{d'}$. The contraction estimate gives

$$\begin{aligned}\|\zeta_d - \zeta_{d'}\|_{\mathcal{C}_S} \le \frac{1}{2}\|\zeta_d - \zeta_{d'}\|_{\mathcal{C}_S} \\ + \|\mathcal{T}_d(d + \zeta_{d'}) - d - \mathcal{T}_{d'}(d' + \zeta_{d'}) + d'\|_{\mathcal{C}_S}.\end{aligned}$$

Both block evaluations are defined after restricting the prefields to a smaller retained neighborhood inside the common small correction neighborhood: the same correction vanishes on B_1, leaving each prefield's own exact interior unchanged. The estimates proved above bound the second term by the finite input difference. Absorbing the first term proves the base difference estimate. The construction uses the same initialized block maps at every finite order where the corresponding contraction is available, so those constructions agree by base uniqueness. The higher regularity on one base domain requires the range arguments that follow. □

B.4. Angular regularity and radial profiles. Fix the exterior inverses, the weight $\gamma = -1/2 - \varepsilon_\gamma$, and the initialized nonlinear branches described above. Take $S_0 \ge 4$ and choose the base domains using rotation-invariant weighted norms. One can use the Euclidean norms of the full derivative tensors in the weighted integrals; these are equivalent to the coordinate norms. The inverse formulas use fixed radial cutoffs and apply the same radial operation to every member of a spherical-harmonic multiplicity space.

For $A \in \mathrm{SO}(3)$, the notation A^* includes the pullback of the actual tensor indices. Let Z_1, Z_2, Z_3 be the standard infinitesimal rotations. For a finite ordered word $I = (i_1, \dots, i_q)$, put $\mathcal{L}^I_Z = \mathcal{L}_{Z_{i_1}} \cdots \mathcal{L}_{Z_{i_q}}$.

Lemma B.5 (Rotation equivariance).
The fixed inverses and initialized nonlinear operations satisfy

$$\begin{aligned} B_{j,\gamma}A^*f &= A^*B_{j,\gamma}f, \qquad j=1,2,\\ \mathfrak{D}(A^*h,A^*\rho) &= A^*\mathfrak{D}(h,\rho),\\ \mathfrak{G}(A^*\widehat{h},A^*s) &= A^*\mathfrak{G}(\widehat{h},s),\\ \mathfrak{E}(A^*\widehat{h},A^*\widehat{K}) &= A^*\mathfrak{E}(\widehat{h},\widehat{K}). \end{aligned}$$

These identities hold on the fixed base domains, before any improvement of spatial regularity.

Proof. The flat inverses are composed of radial differential equations, round-sphere differential operators, spherical Hodge inverses on specified harmonic subspaces, and scalar multipliers depending on the harmonic degree. Proper rotations preserve the round metric, its orientation, and the radial foliation. They therefore commute with the gradient, divergence, curl, and the uniquely normalized Hodge inverses. The excluded low modes are entire harmonic subspaces, so their projections also commute with rotations. On each remaining multiplicity space the radial equations and their boundary conditions are identical. A change of orthonormal harmonic basis consequently leaves the constructed operator unchanged. This proves the assertion for B_1 and B_2.

The multiplier w_{ε_γ} is radial. Both P_j and $E_{j,\varepsilon_\gamma} = w_{\varepsilon_\gamma}^{-1}P_j w_{\varepsilon_\gamma} - P_j$ commute with rotations. Every partial sum of the base Neumann inverse of $I + E_{j,\varepsilon_\gamma}B_j$ commutes with rotations, and the identity passes to its operator-norm limit. The formula for $B_{j,\gamma}$ gives its asserted equivariance. In particular its retained range is invariant.

Tracefree projection and divergence commute with metric pullback. The polar metric variation does also, because a rotation preserves the radial coordinate and its tangent spheres. Rotating either initialized block solution therefore gives a solution of the rotated block equation in the same retained range and small branch. The uniqueness supplied by the base implicit equation proves the two block identities. The premomentum

$$\widehat{a} = \widehat{K} - \frac{\operatorname{tr}_\delta \widehat{K}}{\operatorname{tr}_\delta \widehat{h}}\widehat{h}$$

is equivariant as well. Thus the initial pair, both iteration blocks, and the invariant contraction ball for $\mathfrak{E}$ are preserved by rotations. Every iterate obeys the last identity, which passes to the base-space limit. □

Theorem B.6 (Angular derivatives of the selected output).
Let the smooth prefields $(\widehat{h},\widehat{K})$ have every finite weighted norm and satisfy the local constraints and the fixed base-smallness condition. For the single initialized output $(h,K) = \mathfrak{E}(\widehat{h},\widehat{K})$ and every finite angular word I,

$$\mathcal{L}_Z^I(h-\delta) \in H^{S_0}_\gamma, \qquad \mathcal{L}_Z^I K \in H^{S_0-1}_{\gamma-1}.$$

The scalar correction $z = (\phi, \beta)$ and the Euclidean tracefree momentum correction v have the corresponding angular regularity in their fixed base spaces. Each specified norm is bounded using finitely many higher weighted input norms and the fixed base margins. The conclusions are uniform on supplied finite-dimensional smooth families with these uniform bounds. No higher input norm is required to be small.

Proof. For every tensor field u and finite word I, the coefficients of Z_i being linear and their index actions constant give

$$\|\mathcal{L}_Z^I u\|_{H^s_\omega} \le C_{s,I,\omega} \|u\|_{H^{s+|I|}_\omega}.$$

Indeed a factor of x compensates the weight change of one spatial derivative in the defining norm. Applying the product rule proves the estimate successively for each letter of I. Difference quotients, first on smooth compactly supported fields and then by density in the required weighted spaces, show that $A \mapsto A^* u$ is smooth into a fixed H^s_ω when all higher weighted norms of u are finite. Its derivatives are the corresponding tensor Lie derivatives. The same argument gives locally uniform estimates for supplied smooth families.

The rotated prefields still satisfy the local constraints, and their base norms stay in the same rotation-invariant domain. Apply the smooth base parameter dependence of the initialized construction to this rotation orbit and, when present, to the supplied passive parameters. Lemma B.5 identifies the output as the orbit of the original (h, K). Its derivatives in the base output spaces, identified distributionally, are precisely the angular Lie derivatives in the statement. Products of one-parameter rotations give every ordered word.

The parameter estimates of the base implicit equations supply finite bounds for these derivatives. More explicitly, at each order their highest variation is acted on by the same uniformly invertible base derivative. The other terms contain lower variations and finitely many derivatives of the supplied orbit. The fixed-order multiplication and composition estimates and induction on the variation order therefore give bounds in finitely many of the norms in the right-hand side of the orbit estimate. The branch and its base threshold are unchanged.

The scalar block has a uniquely selected correction pair in $\operatorname{Ran} B_{2,\gamma}$, and that correction pair itself satisfies the same equivariance identity. Its arguments in the outer fixed point are $\widehat{h}$ and $|K|_h^2 - R(\widehat{h})$, whose angular derivatives have the required base regularity by the result just proved and the fixed-order product estimates. The identical argument applies to the selected $v \in \operatorname{Ran} B_{1,\gamma}$ in the momentum block. These are the internal variables of the original output, so this argument does not introduce another spatial branch. □

The rotation fields span the tangent directions of each sphere. The additional normal regularity comes from the ranges of the chosen inverses. Their boundary profiles have a useful property that we record first.

Lemma B.7 (Radial profiles with rapidly decreasing angular coefficients).

Let $Y_{l\mu}$, $l \geq l_0 \geq 1$, be one of the fixed scalar, vector, or tensor spherical-harmonic families on the unit sphere, radially identified with the spheres $\varrho \geq 1$. Suppose

$$\sum_{l,\mu} (1+l)^{2N} |a_{l\mu}|^2 < \infty \quad \text{for every finite integer } N \geq 0.$$

Let smooth radial functions p_l have support in $[1 + c/l, 1 + C/l]$, for fixed $0 < c < C$, and suppose that for every finite j

$$\sup_{\varrho} |\partial_\varrho^j p_l(\varrho)| \leq C_j (1+l)^{n_j}$$

for some finite n_j. Then

$$u(\varrho, \theta) = \sum_{l,\mu} a_{l\mu} p_l(\varrho) Y_{l\mu}(\theta)$$

has a smooth compactly supported zero extension to $\mathbb{R}^3$. Its restriction to the exterior belongs to every $\overline{H}^m_\omega$. Each specified finite norm is bounded by one sufficiently high finite sum in the coefficient condition. The assertion holds for finite sums of such families and for their fixed radial primitives when the corresponding total moment vanishes. For the latter assertion the fixed radial factors and their derivatives must have polynomial bounds in l on these collars, as do the factors occurring in the exterior inverses.

Proof. All supports lie in the bounded annulus $1 \leq \varrho \leq 1 + C$, where Cartesian Sobolev norms are equivalent to finite sums of radial and spherical Sobolev norms. Round-sphere elliptic estimates and orthogonality give, for a finite harmonic sum,

$$\|u\|^2_{H^m} \leq C_m \sum_{j=0}^{m} \sum_{l,\mu} (1+l)^{2(m-j)} |a_{l\mu}|^2 \|\partial_\varrho^j p_l\|^2_{L^2(1,1+C)}.$$

The same estimate holds in each of the indicated tensor bundles; their fixed harmonic eigenvalues are comparable with $(1+l)^2$. Fixed radial identifications contribute only bounded smooth factors on this annulus. Choose an integer $N_m \geq \max_{0 \leq j \leq m}(m - j + n_j)$. The radial bounds control the displayed expression by $C_m \sum (1+l)^{2N_m} |a_{l\mu}|^2$. Applying the estimate to tails proves convergence at every prescribed Sobolev order.

Every finite harmonic sum vanishes in a neighborhood of the inner sphere, so its zero extension is smooth. The just-proved convergence of these zero extensions at every order gives a smooth limit on the whole space. It vanishes on B_1 and outside the fixed outer ball. This also proves all zero-extension trace conditions. Weighted norms are equivalent to unweighted ones on the common support.

For a radial primitive, vanishing of its total moment makes the primitive zero on both sides of the supporting collar. Its zeroth-order estimate follows by integrating over an interval of length at most C/l. Each positive derivative

is estimated by the product rule from the original profile and its fixed radial factors. All such estimates remain polynomial in l, with order-dependent constants. The preceding argument then applies to the primitive. At no point is a constant required to be uniform in the derivative order. □

B.5. The scalar range and its elliptic equation. The scalar right inverse has a sphere-tangent variable whose radial dependence is fixed by moments. This range condition first makes that variable smooth; a scalar elliptic estimate then improves the other variable. The premetric in this argument is the supplied smooth $\widehat{h}$.

Lemma B.8 (Smoothness of the sphere-tangent scalar correction). Let $z = (\phi, \beta) \in \operatorname{Ran} B_{2,\gamma}$ have every angular Lie derivative in its base $\overline{H}^{S_0}_{\gamma}$ space. Then the zero extension of β is smooth, with

$$\operatorname{supp}\beta \subset \{x : 1 \le |x| \le 2\}.$$

For every finite m its spatial norm is bounded by finitely many angular base norms of z. These estimates are uniform on families with the corresponding uniform bounds. The support assertion does not require a common positive distance between the support and the inner sphere.

Proof. The exact range identity gives $z = w_{\varepsilon_\gamma}(u, \xi)$ with $(u, \xi) \in \operatorname{Ran} B_2$. Use the curvature derivative defined above. Use $f = -P_2(u, \xi)$ in the formulas of Czimek (2018, (5.31) and (5.35)–(5.37)). This is the sign of the residual in the displayed construction there. The normalized right inverse in this paper satisfies $P_2 B_2 = I$; changing the sign of the supplied residual changes neither its range nor the following regularity argument. The flat differential operator and radial multiplier show that $f \in \overline{H}^{S_0-2}_{-5/2}$ and that every angular derivative of f belongs to that space, with bounds from the angular base norms of z.

Use harmonics $Y_{l\mu}(\varrho, \theta) = \varrho^{-1} Y_{l\mu}(1, \theta)$, orthonormal on S_ϱ, and let $f_{l\mu}(\varrho)$ be the corresponding coefficients. Put

$$\lambda_l = \sqrt{\frac{l(l+1)}{2}}, \qquad \chi_l(\varrho) = \chi(l(\varrho - 1)), \qquad l \ge 1,$$

where $0 \le \chi \le 1$, $\chi = 0$ on $(-\infty, 1/10]$, and $\chi = 1$ on $[1, \infty)$. The retained profiles are

$$\begin{aligned}
c_{l\mu} &= -\frac{1}{2}\int_1^\infty \varrho^{1-\lambda_l} f_{l\mu}(\varrho)\,\mathrm{d}\varrho,\\
A_l &= \int_1^\infty \varrho^{\lambda_l+1}\chi_l'(\varrho)\,\mathrm{d}\varrho, \qquad D_l = \int_1^\infty \varrho^{\lambda_l+1}\chi_l''(\varrho)\,\mathrm{d}\varrho,\\
\widetilde{c}_{l\mu} &= c_{l\mu}\frac{A_l}{D_l},\\
\zeta_{l\mu}(\varrho) &= c_{l\mu}\varrho^{\lambda_l-1}\chi_l'(\varrho) - \widetilde{c}_{l\mu}\varrho^{\lambda_l-1}\chi_l''(\varrho),\\
\zeta(\varrho, \theta) &= \sum_{l\ge1,\mu} \zeta_{l\mu}(\varrho) Y_{l\mu}(\varrho, \theta).
\end{aligned}$$

These formulas determine the tangent variable by

$$\operatorname{div}_{S_\varrho}\xi = \varrho^{-3}\int_1^\varrho s^3\zeta(s,\theta)\,\mathrm{d}s, \qquad \operatorname{curl}_{S_\varrho}\xi = 0.$$

We verify all properties of these expressions needed here. Since $\lambda_l \geq 1$, Cauchy–Schwarz gives

$$|c_{l\mu}| \leq \frac{1}{2\sqrt{2\lambda_l - 1}}\|\varrho f_{l\mu}\|_{L^2(1,\infty)}.$$

The $H^0_{-5/2}$ norm controls the square sum of the norms on the right: orthonormality on S_ϱ leaves the radial measure $\mathrm{d}\varrho$, and the weight is comparable with ϱ^2 for $\varrho \geq 1$. The rotation Casimir multiplies the lth scalar harmonic by $l(l+1)$. Applying the same bound after finitely many rotations therefore yields

$$\sum_{l,\mu}(1+l)^{2N}|c_{l\mu}|^2 < \infty \quad \text{for every finite } N,$$

with finite bounds in the corresponding angular norms of f.

The denominator in the profile definition cannot vanish. Integration by parts gives

$$D_l = -(\lambda_l + 1)\int_1^\infty \varrho^{\lambda_l}\chi_l'(\varrho)\,\mathrm{d}\varrho.$$

Writing $b_l = 1 + 1/l$, another integration by parts and $0 \leq \chi_l \leq 1$ give

$$\int_1^{b_l}\varrho^{\lambda_l}\chi_l'(\varrho)\,\mathrm{d}\varrho = b_l^{\lambda_l} - \lambda_l\int_1^{b_l}\varrho^{\lambda_l - 1}\chi_l(\varrho)\,\mathrm{d}\varrho \geq 1.$$

Thus $D_l \leq -(\lambda_l + 1)$. On the supporting collar, ϱ^{λ_l+1} is bounded independently of l, so $|A_l| \leq C\|\chi'\|_{L^1}$. Consequently

$$|\widetilde{c}_{l\mu}| \leq C(1+l)^{-1}|c_{l\mu}|.$$

Each profile is supported in $[1 + (10l)^{-1}, 1 + l^{-1}]$. Every radial derivative of $\varrho^{\lambda_l - 1}$ costs a finite power of $1 + l$ on this collar, as does each derivative of χ_l. For example, after factoring out $c_{l\mu}$, the jth derivative of the combined profile is bounded by $C_j(1+l)^{j+1}$. The radial factor ϱ^{-1} in the harmonic normalization preserves this property. Lemma B.7 and the rapid coefficient estimates show that ζ has a smooth compactly supported zero extension.

The two moment identities can be checked directly:

$$\int_1^\infty \varrho^{1-\lambda_l}\left(\frac{1}{2}f_{l\mu}(\varrho) + \zeta_{l\mu}(\varrho)\right)\mathrm{d}\varrho = 0,$$

$$\int_1^\infty \varrho^2\zeta_{l\mu}(\varrho)\,\mathrm{d}\varrho = 0.$$

For the first, $\int\chi_l' = 1$ and $\int\chi_l'' = 0$, so the contribution of $\zeta_{l\mu}$ is $c_{l\mu}$. The definition of $c_{l\mu}$ cancels the residual term. For the second, the integral

is $c_{l\mu}A_l - \widetilde{c}_{l\mu}D_l$. This is zero by definition. The second identity is the cancellation required for the tangent variable. In unit-sphere components,

$$\int_1^{\varrho} s^3\zeta(s,\theta)\,\mathrm{d}s = \sum_{l,\mu}\left(\int_1^{\varrho} s^2\zeta_{l\mu}(s)\,\mathrm{d}s\right)Y_{l\mu}(1,\theta).$$

Each coefficient on the right vanishes before and after its original collar, by the second moment identity. Its derivatives have polynomial harmonic-degree bounds. The scalar constant mode is absent. The unique curl-free solution of the spherical divergence equation is therefore the spherical gradient of the mean-zero Poisson solution. On a nonconstant mode its inverse involves the factor $1/[l(l+1)]$ and fixed powers of ϱ, so it preserves polynomial bounds on the fixed annulus. Lemma B.7 applied to these primitives proves the asserted smoothness and support for ξ, hence for $\beta = w_{\varepsilon_\gamma}\xi$.

For any requested spatial order, the profile estimate uses a finite coefficient order, that coefficient order uses finitely many rotation derivatives of f, and the differential formula for f uses finitely many angular base norms of z. This proves the stated finite dependence and its uniform family version. The collars shrink toward the inner sphere as l increases; convergence of their zero extensions is what proves smoothness at that sphere. □

Proposition B.9 (Scalar regularity of the same base solution).
Let (h, K) be the initialized base fixed point for smooth supplied prefields, and write $h = \mathfrak{V}_{\widehat{h}}(\phi,\beta)$. Retain the fixed base positivity margins and all angular estimates of Theorem B.6. If, for an integer $m \geq S_0$,

$$\phi \in \overline{H}^m_\gamma, \qquad K \in H^{m-1}_{\gamma-1},$$

then $\phi \in \overline{H}^{m+1}_\gamma$ and $h - \delta \in H^{m+1}_\gamma$. Their finite bounds depend on the displayed norms, finitely many supplied premetric and angular norms, and the fixed base margins. If the hypothesis holds after every finite angular word, the conclusion holds after every finite angular word. For each specified word only finitely many preceding norms are needed. These conclusions hold uniformly on families with the indicated uniform bounds, without another smallness condition.

Proof. Lemma B.8 makes β smooth at every spatial order for this very solution. We give the elliptic equation and its estimate, including the behavior at the inner sphere and at infinity.

Write

$$\widehat{h} = a^2\,\mathrm{d}\varrho^2 + \widehat{\gamma}_{AB}(b^A\mathrm{d}\varrho + \mathrm{d}\theta^A)(b^B\mathrm{d}\varrho + \mathrm{d}\theta^B), \qquad N' = a^{-1}(\partial_\varrho - b - \beta).$$

For the following calculation use the positive second fundamental form of the radial leaves, namely $\frac{1}{2}\mathcal{L}_{N'}$ of their metric. Set

$$T_0 = \frac{1}{2a}(\partial_\varrho\widehat{\gamma} - \mathcal{L}_{b+\beta}\widehat{\gamma}), \qquad H_0 = \mathrm{tr}_{\widehat{\gamma}}T_0, \qquad Q_0 = |T_0|^2_{\widehat{\gamma}}.$$

These fields depend only on $\widehat{h}$ and β. The leaf metric of h is $q = e^{2\phi}\widehat{\gamma}$, with second fundamental form $e^{2\phi}(T_0 + (N'\phi)\widehat{\gamma})$. Its mean curvature and squared

norm are

$$H = H_0 + 2N'\phi, \qquad |T|_q^2 = Q_0 + 2H_0N'\phi + 2(N'\phi)^2.$$

The contracted Gauss identity in this convention is

$$R(h) = R(q) - 2N'H - H^2 - |T|_q^2 - 2a^{-1}\Delta_q a.$$

Substituting the two-dimensional conformal identities $R(q) = e^{-2\phi}(R(\widehat{\gamma}) - 2\Delta_{\widehat{\gamma}}\phi)$ and $\Delta_q a = e^{-2\phi}\Delta_{\widehat{\gamma}} a$ gives

$$\begin{aligned} R(h) = & -4N'(N'\phi) - 2e^{-2\phi}\Delta_{\widehat{\gamma}}\phi + \mathcal{Q}(\phi, \mathrm{d}\phi; \widehat{h}, \beta), \\ \mathcal{Q} = & e^{-2\phi}\left(R(\widehat{\gamma}) - 2a^{-1}\Delta_{\widehat{\gamma}} a\right) - 2N'H_0 - H_0^2 - Q_0 \\ & - 6H_0N'\phi - 6(N'\phi)^2. \end{aligned}$$

In particular the second-order part is

$$-4N'(N'\phi) - 2e^{-2\phi}\Delta_{\widehat{\gamma}}\phi.$$

This agrees with the scalar variation formula in Czimek (2018, Lemma 5.5). The displayed expression for $\mathcal{Q}$ shows that every remaining term contains at most one derivative of ϕ.

After changing the sign of the second-order operator, its principal form is

$$4\,\xi(N')^2 + 2e^{-2\phi}|\xi|^2_{\widehat{\gamma}^{-1},\tan}.$$

It is uniformly positive relative to the Euclidean squared norm. Indeed the tangential term controls the restriction of ξ to a leaf, and the nonzero radial component a^{-1} of N' then controls its remaining component. The base bounds give uniform upper and lower bounds for a, e^{ϕ}, the leaf metric relative to the Euclidean sphere metric, and the shift. These bounds fix the ellipticity margin independently of m.

The actual fixed-point equation is $R(h) = |K|_h^2$. At the induction order the scalar inhomogeneity and all first-order terms have $H^{m-1}_{\gamma-2}$ regularity. We spell out the cancellation needed for this weighted assertion. At $(\widehat{h}, \beta, \phi) = (\delta, 0, 0)$,

$$R(\widehat{\gamma}) = \frac{2}{\varrho^2}, \quad H_0 = \frac{2}{\varrho}, \quad Q_0 = \frac{2}{\varrho^2}, \quad N'H_0 = -\frac{2}{\varrho^2}, \quad \Delta_{\widehat{\gamma}} a = 0.$$

Their contribution to $\mathcal{Q}$ is zero. Subtract these flat values before estimating the prescribed terms. Their differences contain $\widehat{h} - \delta$, its derivatives through order two, or the compactly supported smooth β and its derivatives. They consequently have the required weight $\gamma - 2$. The term $(e^{-2\phi} - 1)\varrho^{-2}$ has the same weight. Terms linear in $\mathrm{d}\phi$ have coefficients of size $O(\varrho^{-1})$, plus faster-decaying prescribed terms. Quadratic terms in $\mathrm{d}\phi$ and in K have weight $2\gamma - 2$, which is at least the required decay since $\gamma < 0$. Weighted product and composition estimates, with $m - 1 \geq 3$, give the asserted H^{m-1} regularity. The individual flat ϱ^{-2} terms are used only in their cancelling combination.

On a bounded annulus, write the equation in Cartesian coordinates as

$$A^{ij}(x,\phi)\partial_i\partial_j\phi = F(x,\phi,\mathrm{d}\phi,K),$$

where A has the positive form above; the prescribed arguments $\widehat{h}$ and β are suppressed. Its coefficients have the known H^m regularity and a fixed ellipticity margin. The right-hand side is in H^{m-1}. For a multi-index $|\alpha| = m-1$, $w = \partial^\alpha\phi$ initially belongs to H^1 and satisfies

$$A^{ij}\partial_i\partial_j w = \partial^\alpha F - [\partial^\alpha, A^{ij}]\partial_i\partial_j\phi.$$

The right-hand side is in L^2. In the commutator the highest derivative of ϕ is m, not $m+1$. The standard product estimate is bounded by a constant times

$$\|\mathrm{d}A\|_{L^\infty}\|\phi\|_{H^m} + \|A\|_{H^{m-1}}\|\mathrm{d}^2\phi\|_{L^\infty},$$

with the usual lower-order terms on a slightly larger annulus. These quantities are finite because $m \geq 4$ and H^4 in three dimensions controls the second derivatives in L^∞. The coefficients are Lipschitz. Converting the equation for w to divergence form puts the term $(\partial_i A^{ij})\partial_j w$ in L^2. The interior difference-quotient H^2 estimate therefore applies to this existing H^1 solution. It gives $w \in H^2$ on a smaller annulus, hence $\phi \in H^{m+1}$ there. The constants depend on norms already estimated and on the ellipticity margin fixed at base order. No smallness condition at the new order enters this step.

For an annulus crossing ∂B_1, extend ϕ and β by zero. The initial zero-extension regularity makes $h = \mathfrak{V}_{\widehat{h}}(\phi,\beta)$ a whole H^{S_0} metric there. Its curvature has no boundary distribution. On the inner side $h = \widehat{h}$ and $K = \widehat{K}$, so the same equation is exactly the supplied local constraint. Thus the scalar equation holds distributionally across the sphere. The polar coefficients are smooth away from the origin, and the preceding interior estimate applies across the sphere. The gain is a gain for the zero extension itself, so it preserves the required vanishing traces.

For the global weighted estimate, cover the remaining exterior by overlapping annuli $\{R < \varrho < 2R\}$ with $R = 2^j$ and larger parents $\{R/2 < \varrho < 4R\}$. Rescale $x = Ry$. On a fixed annulus in y, the known weighted H^m norms control the coefficient norms uniformly in R, because $\gamma < 0$; the rescaled Euclidean polar coefficients are fixed smooth functions. The ellipticity margin is the base margin. For $R > 2$ the tangent correction is zero. The local estimate just proved has the form

$$\|\phi(R\cdot)\|_{H^{m+1}(A)} \leq C_m\left(\|R^2F(R\cdot)\|_{H^{m-1}(A^*)} + \|\phi(R\cdot)\|_{H^m(A^*)}\right),$$

where C_m is bounded by the fixed margin and finitely many already controlled weighted norms. Multiplication by $R^{-\gamma}$, squaring, and summation over the finite-overlap annuli gives the H^{m+1}_γ estimate, using the weighted $H^{m-1}_{\gamma-2}$ norm of the inhomogeneity. A bounded annulus handles the inner sphere. The zero extension remains zero on the ball, so its new regularity is precisely $\overline{H}^{m+1}_\gamma$. The algebraic formula for $\mathfrak{V}_{\widehat{h}}$ then gives $h - \delta \in H^{m+1}_\gamma$.

Finally retain angular words. First perform the zero-word improvement above, and then induct on word length. Applying an ordered word I to the scalar equation gives a linear equation for $\mathcal{L}_Z^I\phi$ with the same second-order principal coefficients. All other second derivatives of angular variations have smaller word length and have already gained the spatial derivative. The highest angular variations of K have the assumed $H^{m-1}_{\gamma-1}$ regularity; the prescribed fields and β have every required order. The zeroth-order linearized coefficient may contain second derivatives of the zero-word ϕ, which are now in H^{m-1}. Products with the known H^m angular unknown therefore have the required H^{m-1} regularity. The same differentiated H^2 estimate, the equation across the sphere, and the annular summation give the improved norm for I. Each finite step uses finitely many lower words. All estimates remain uniform for families with the stated bounds, and every step concerns the original initialized solution. □

B.6. The normal equations in the momentum range. The momentum equation has three components, whereas a symmetric tracefree tensor in three dimensions has five. The selected exterior range supplies the two remaining equations. We derive those equations before estimating them, since a right inverse with the same mapping properties need not select this range.

Throughout this section, $\varrho = |x|$ is the exterior radius. Tensor components are taken in Euclidean orthonormal radial frames parallel along each radial ray. A tangent tensor in such a frame is identified with a tensor on the unit sphere. Write D, d and c for the unit-sphere gradient of a scalar, divergence of a symmetric tracefree tensor, and scalar curl of a one-form, respectively. Thus a physical angular derivative contains the factor ϱ^{-1}.

For a Euclidean tracefree symmetric tensor u, put

$$u_{NN} = a_u, \qquad u_{NA} = e_A, \qquad u_{AB} = T_{AB} - \tfrac{1}{2}a_u(g_{S^2})_{AB}, \qquad \operatorname{tr}_{S^2} T = 0.$$

The round-sphere Hodge decomposition splits one-forms into exact and coexact parts, denoted by the subscripts E and H. On harmonic degrees $l \geq 2$, the map d is an isomorphism from tracefree symmetric tensor harmonics to one-form harmonics. Let L_E and L_H be its inverse on the exact and coexact parts, extended by zero on degrees below two. Both have angular order -1, and $L_E D$ has order zero. Write $T = T_E + T_H$ for the corresponding decomposition of T. There are no tracefree symmetric tensor harmonics of degrees zero or one. The operators and projections just defined commute with proper rotations.

On coexact one-forms define

$$\Delta_H = c^{-1}\Delta_{S^2}c.$$

Here c is invertible from coexact one-forms to mean-zero scalars. In degree l, the eigenvalue of Δ_H is $-l(l+1)$. This definition fixes the convention; Δ_H is not the rough Laplacian on one-forms.

Lemma B.10 (Normal equations of the selected flat momentum range).

Let $u = B_1 f$ lie in the explicit flat momentum range used to define the exterior normalization. Suppose $u \in \overline{H}^{S_0-1}_{-3/2}$, with $S_0 \geq 4$, after every finite angular Lie derivative. There are smooth tangent fields ζ_E, ζ_H, whose zero extensions are supported in $1 \leq \varrho \leq 3/2$, such that

$$T_E = -L_E(\tfrac{1}{2} D a_u + e_E + \varrho \zeta_E)$$

and

$$\begin{aligned} \partial_\varrho T_H - \frac{2}{\varrho} T_H = & - L_H \partial_\varrho e_H + \frac{1}{\varrho} L_H (\Delta_H + 2) e_H \\ & - \varrho L_H \partial_\varrho \zeta_H + L_H \zeta_H. \end{aligned}$$

Projection to degrees $l \geq 2$ is understood wherever L_E or L_H occurs. Each finite spatial norm of the profiles is bounded by finitely many angular base norms of u. Both identities hold distributionally across $\varrho = 1$ after zero extension.

Proof. The inverse B_1 is the construction of Czimek (2018, §4.3), with the fixed degree-dependent radial cutoffs stipulated above. Its residual is $f = \operatorname{div}_\delta u$. Since Euclidean divergence commutes with rotations, f has every angular derivative in $\overline{H}^{S_0-2}_{-5/2}$.

First consider the boundary profiles. Use the harmonic normalization of that construction: the scalar and one-form harmonics on S_ϱ have $L^2(S_\varrho)$ norm one. Write $f_{N,l\mu}$, $f_{E,l\mu}$ and $f_{H,l\mu}$ for the resulting coefficient functions. Equations (4.43)–(4.44) and (4.48)–(4.49) of the cited construction define

$$\begin{aligned} b_{E,l\mu} &= \int_1^\infty \varrho^{-l+1} \left(\frac{l}{\sqrt{l(l+1)}} f_{N,l\mu} - f_{E,l\mu} \right) \mathrm{d}\varrho, \\ b_{H,l\mu} &= -\int_1^\infty \varrho^{-1-\nu_l} f_{H,l\mu} \, \mathrm{d}\varrho, \qquad \nu_l = \sqrt{l(l+1)+4}, \qquad l \geq 2. \end{aligned}$$

The corresponding normalized profile coefficients are

$$\begin{aligned} (\zeta_E)_{l\mu} &= b_{E,l\mu} \varrho^{l-1} \partial_\varrho \chi(l(\varrho - 1)), \\ (\zeta_H)_{l\mu} &= b_{H,l\mu} \varrho^{1+\nu_l} \partial_\varrho \chi(l(\varrho - 1)). \end{aligned}$$

The fixed smooth transition function χ is zero on $(-\infty, 1/10]$ and one on $[1, \infty)$. The normalized harmonics become ϱ^{-1} times the unit-sphere harmonics when expressed in the physical frames used here. Consequently the radial profiles in physical components contain the powers ϱ^{l-2} and ϱ^{ν_l}, respectively. This conversion does not change the moment coefficients.

Cauchy–Schwarz gives

$$\begin{aligned} |b_{E,l\mu}| &\leq (2l-1)^{-1/2} \left(\|\varrho f_{N,l\mu}\|_{L^2(1,\infty)} + \|\varrho f_{E,l\mu}\|_{L^2(1,\infty)} \right), \\ |b_{H,l\mu}| &\leq (2\nu_l + 3)^{-1/2} \|\varrho f_{H,l\mu}\|_{L^2(1,\infty)}. \end{aligned}$$

For example, in the first line write $\varrho^{-l+1} f_{N,l\mu} = \varrho^{-l}(\varrho f_{N,l\mu})$ and use $\int_1^\infty \varrho^{-2l} \, \mathrm{d}\varrho = (2l-1)^{-1}$. With the stated harmonic normalization, the

square sum of the source norms on the right is equivalent to the critical weighted $L^2_{-5/2}$ norm of f. Applying powers of the rotation Casimir operator supplies every polynomial harmonic-degree weight. Hence, for every finite N,

$$\sum_{l,\mu}(1+l)^{2N}\left(|b_{E,l\mu}|^2+|b_{H,l\mu}|^2\right)<\infty,$$

with a bound in finitely many angular base norms of u.

Each nonzero profile in the coefficient formulas is supported in

$$1+(10l)^{-1}\le\varrho\le 1+l^{-1}.$$

On this collar, ϱ^{l-2} and ϱ^{ν_l} are uniformly bounded. Every specified number of radial derivatives costs a finite polynomial in $1+l$; the constant can depend on that number. The rapidly decreasing moment coefficients therefore give convergence in every Sobolev norm of the zero-extended harmonic sums. Indeed, any finite sum vanishes on a neighborhood of the inner sphere, and the polynomial derivative bounds make the tails converge in each prescribed norm. Sobolev embedding then proves that the full zero extensions are smooth. Their supports lie in $1\le\varrho\le 3/2$, since $l\ge 2$. This also proves the asserted finite dependence of each profile norm.

We next derive the tensor equations. Put $\rho=\operatorname{div}_\delta u$ and let σ be the symmetric curl tensor in the explicit construction. The radial divergence formula and the electric Hodge equation in Czimek (2018, equations (4.38), (4.53) and (4.56)) give, in physical components,

$$\begin{aligned}\rho_N&=\partial_\varrho a_u+3\varrho^{-1}a_u+\varrho^{-1}\operatorname{div}_{S^2}e,\\ \rho_{\tan}&=\partial_\varrho e+3\varrho^{-1}e+\varrho^{-1}dT-\tfrac{1}{2\varrho}Da_u,\\ \varrho^{-1}dT_E&=-\zeta_E-\varrho^{-1}\Pi_{\ge2}\left(\tfrac12 Da_u+e_E\right).\end{aligned}$$

In the third equation, $\Pi_{\ge2}$ projects one-forms on the unit sphere onto harmonic degrees $l\ge2$. Since $L_E\Pi_{\ge2}=L_E$, applying L_E to that equation gives the electric identity.

On the coexact subspace of degree at least two, equations (4.47) and (4.50) of that construction become

$$\sigma_{NN}=\varrho^{-1}ce_H,\qquad\left(\partial_\varrho^2+\frac{3}{\varrho}\partial_\varrho+\frac{\Delta_{S^2}-3}{\varrho^2}\right)\sigma_{NN}=\partial_\varrho\left(\varrho^{-1}c(\rho_H+\zeta_H)\right).$$

Substitution of the first identity and cancellation of c/ϱ give

$$\partial_\varrho^2e_H+\varrho^{-1}\partial_\varrho e_H+\varrho^{-2}(\Delta_H-4)e_H=\partial_\varrho(\rho_H+\zeta_H)-\varrho^{-1}(\rho_H+\zeta_H).$$

The coexact part of the divergence formula is $\rho_H=\partial_\varrho e_H+3\varrho^{-1}e_H+\varrho^{-1}dT_H$. Insert it into the last equation. The second normal derivatives of e_H cancel, and multiplication by ϱ leaves

$$d\partial_\varrho T_H-2\varrho^{-1}dT_H=-\partial_\varrho e_H+\varrho^{-1}(\Delta_H+2)e_H-\varrho\partial_\varrho\zeta_H+\zeta_H.$$

Application of L_H proves the magnetic identity.

These calculations are valid for the selected base solution. The range formulas and their derivatives are continuous as distributional operations on its base spaces, so one may first use the defining smooth approximations of the flat inverse. The zero-extended field u belongs to H^{S_0-1} locally and its angular derivatives have the same property. Since $S_0 - 1 \geq 3$, the traces needed in the second-order calculation vanish at $\varrho = 1$. The zero-extended residual has the corresponding barred regularity and the profiles are smooth zero extensions. Differentiation therefore produces no boundary delta terms. The identities consequently hold on an annulus crossing the inner sphere. □

The stronger-weight normalization preserves the normal principal part. Write

$$v = wu \in \operatorname{Ran} B_{1,\gamma}, \qquad w = e^{-\varepsilon_\gamma \ell}.$$

The exact range equality established for the conjugated inverse identifies $u = w^{-1}v$ with an element of the original flat range. The electric equation is

$$(T_v)_E = -L_E\big(\tfrac{1}{2} D a_v + (e_v)_E + \varrho w \zeta_E\big).$$

For the magnetic part set

$$W_H = (T_v)_H + L_H (e_v)_H + \varrho w L_H \zeta_H.$$

Multiplication of the flat identity by w gives exactly

$$\partial_\varrho W_H - \left(\frac{2}{\varrho} + \frac{w'}{w}\right) W_H = \frac{1}{\varrho} L_H \Delta_H (e_v)_H.$$

Thus the weight changes only known coefficients of order zero. All derivatives of w'/w satisfy $|\partial_\varrho^j (w'/w)| \leq C_j \varrho^{-1-j}$; it also vanishes near the inner sphere.

Proposition B.11 (One normal regularity gain in the fixed momentum range).

Let $m \geq S_0 \geq 4$. Let h be a positive whole metric satisfying one fixed sufficiently small L^∞ bound on $h - \delta$. Suppose $v \in \operatorname{Ran} B_{1,\gamma}$ is a base solution of

$$\operatorname{div}_h \operatorname{TF}_h v = \rho.$$

Assume that after every finite angular word,

$$h - \delta \in H^m_\gamma, \qquad v \in \overline{H}^{m-1}_{\gamma-1}, \qquad \rho \in \overline{H}^{m-1}_{\gamma-2}.$$

Then $v \in \overline{H}^m_{\gamma-1}$ after every finite angular word. Each specified finite output norm is bounded by finitely many of the displayed norms and the fixed base inverse margin. The smallness threshold is independent of m and of the angular order. Uniform finite input bounds give uniform output bounds for families.

Proof. Lemma B.10 applies to $u = w^{-1}v$ and supplies smooth boundary profiles, with bounds from finitely many angular base norms of v. Differentiate the weighted electric equation and expand the weighted magnetic equation. They parameterize the complete normal derivative as

$$\partial_\varrho v = JU + F_\varrho, \qquad U = (\partial_\varrho a_v, \partial_\varrho e_v).$$

The fixed unit-sphere operator J has components

$$(JU)_{NN} = U_0, \qquad (JU)_{NA} = (U_1)_A,$$
$$(JU)_{\tan,0} = -L_E\big(\tfrac{1}{2}DU_0 + (U_1)_E\big) - L_H(U_1)_H,$$
$$(JU)_{AB} = (JU)_{\tan,0,AB} - \tfrac{1}{2}U_0(g_{S^2})_{AB}.$$

It is bounded from the L^2 scalar–one-form pair to L^2 tracefree tensors: $L_E D$ has order zero and L_E, L_H have order -1. It commutes with rotations. No low harmonic is discarded from U; the zero action of L_E, L_H below degree two says only that the tangential tracefree part has no such modes.

The NN and NA components of F_ϱ vanish. Its tangential tracefree component is explicitly

$$\begin{aligned}(F_\varrho)_{\tan,0} = \; & - L_E\partial_\varrho(\varrho w\zeta_E) - \partial_\varrho(\varrho w L_H\zeta_H) \\ & + \left(\frac{2}{\varrho} + \frac{w'}{w}\right) W_H + \frac{1}{\varrho}L_H\Delta_H(e_v)_H.\end{aligned}$$

Its nonsmooth-input terms have at most one angular derivative of v and no normal derivative of v. All other terms are the already smooth profiles. Hence, using finitely many extra angular words, F_ϱ has radial Sobolev order $m-1$ at weight $\gamma-2$.

Write the normal principal coefficient of $\operatorname{div}_h \mathrm{TF}_h$ as M_h. In the radial orthonormal frame its action on a Euclidean tracefree tensor t is

$$(M_h t)_j = h^{Ni}t_{ij} - \tfrac{1}{3}\delta_{jN}h^{ab}t_{ab}.$$

This follows by differentiating $\mathrm{TF}_h t = t - \frac{1}{3}(h^{ab}t_{ab})h$ and retaining the derivatives of t in the normal direction. Derivatives of h and derivatives of the radial frame belong to the remaining terms. At the Euclidean metric, $M_\delta t = (t_{NN}, t_{NA})$, because $\operatorname{tr}_\delta t = 0$. Therefore

$$M_\delta J = I, \qquad \|(M_h - M_\delta)J\|_{L^2(S^2)\to L^2(S^2)} \le C\|h-\delta\|_{L^\infty(S^2)}.$$

Choose the base metric neighborhood once so that the last expression is at most $1/2$ at every radius. The operator

$$\mathcal{A}_h = I + (M_h - M_\delta)J$$

then has an inverse on $L^2(S^2)$ of norm at most two. This is the only smallness used in the normal regularity argument.

Substitute the parameterization into the momentum equation. It becomes

$$\mathcal{A}_h U = G_\varrho, \qquad G_\varrho = \rho - M_h F_\varrho - Q(h,v),$$

where $Q(h,v)$ contains only the angular derivatives of v, the radial-frame terms of size $\varrho^{-1}v$, and products of first derivatives of h with v. We verify the required regularity of this equation on a fixed annulus before summing over annuli.

All-angular H^m regularity of h implies H^m regularity in the radial variable with values in any prescribed finite sphere Sobolev space. The analogous statement holds at radial order $m-1$ for v and ρ. Conversely, a finite number

of angular words controls each specified sphere Sobolev norm; the rotation fields span the sphere tangent bundle and their Casimir is elliptic. Choose a sphere Sobolev index large enough for the finite products under consideration. The one-dimensional radial product estimates and the sphere Sobolev algebra estimates then show that G_ϱ has radial order $m-1$ with every specified angular derivative. The highest radial derivative of h appearing here has order m, from the differentiated first-derivative coefficient in $Q(h,v)$. It is an L^2 factor multiplied by lower derivatives of v, which are bounded by the preceding radial regularity. No derivative of v of radial order m occurs in G_ϱ.

The coefficients of $M_h - M_\delta$, after any fixed angular word, have bounded radial derivatives through order $m-1$ on this annulus. Indeed, the additional radial derivative of h gives H^1 control of its $(m-1)$st radial derivative, and the one-dimensional Sobolev embedding gives the bound with values in a sufficiently high sphere Sobolev space. Smooth dependence of inverse metric coefficients preserves these bounds on the positive metric neighborhood.

The preceding regularity of v already gives $U \in L^2$ locally. Thus the normal equation identifies it almost everywhere with $\mathcal{A}_h^{-1} G_\varrho$. To justify differentiating this identity without assuming the desired derivatives, use difference quotients. For a radial increment t, put $\delta_t U(\varrho) = (U(\varrho+t) - U(\varrho))/t$, and use the same convention for the other fields. Subtraction of the equation at two radii gives

$$\mathcal{A}_h(\varrho+t)\,\delta_t U = \delta_t G_\varrho - (\delta_t \mathcal{A}_h)U.$$

The same inverse bound applies at the shifted radius. For a one-parameter rotation, pull back the equation and subtract it; J commutes with the pullback, so the resulting difference equation has the same form and the same inverse margin. Repeating this procedure gives, for every finite mixed radial and rotation derivative D^α, an equation whose highest derivative is

$$\mathcal{A}_h D^\alpha U = D^\alpha G_\varrho - \sum_{0<\beta\le\alpha} C_{\alpha,\beta}(D^\beta \mathcal{A}_h) D^{\alpha-\beta} U.$$

For ordered rotation words, the sum means the corresponding finite Leibniz sum, with their order retained. Every term on the right uses a strictly lower total derivative of U. Induction, the coefficient bounds just proved, and the L^2 inverse bound therefore control all angular derivatives of U with radial order at most $m-1$. Bounded difference quotients identify the weak derivatives at each step. The constants involve finitely many coefficient and input norms, while the highest derivative is always inverted at the same base margin.

The normal system now gives radial order $m-1$ for $\partial_\varrho v$. Together with the available angular derivatives, this yields H^m regularity of v on the annulus. The argument also applies to an annulus crossing $\varrho = 1$: the range equations hold there for the zero extension, ρ is a barred field, and h is the given

whole metric. Consequently it improves the zero extension itself, including its boundary traces.

For completeness, the weighted summation uses the same margin. Cover the exterior by overlapping dyadic annuli A_R with $R = 2^j$, together with one annulus crossing the inner sphere, and rescale $x = Ry$. Use physical tensor components for $h_R(y) = h(Ry)$ and $v_R(y) = v(Ry)$. The rescaled normal equation has right-hand side $R\rho(Ry)$; J is unchanged, and the factors ϱ^{-1} and w'/w have uniformly bounded rescaled derivatives. Weighted norms are equivalent to sums with factors $R^{-2(\gamma-1)}$ for v_R and $R^{-2(\gamma-2)}$ for $\rho(Ry)$. The additional factor R in the latter equation gives exactly these matching weights.

Every fixed rescaled coefficient norm of $h_R - \delta$ is bounded by the corresponding global weighted norm, since its local bound carries R^γ and $\gamma < 0$. Product terms thus have uniform finite annular constants. Summing the squared local estimates proves $v \in H^m_{\gamma-1}$ after each retained angular word. Since the improved field is still the zero extension, it belongs to $\overline{H}^m_{\gamma-1}$. The entire argument applies with uniform constants on a family having the stipulated finite bounds. □

Corollary B.12 (Momentum regularity in the relative extension). Let (h, K) be the initialized relative-extension fixed point for smooth supplied prefields. Suppose its all-angular regularity has reached $h - \delta \in H^m_\gamma$ and $K \in H^{m-1}_{\gamma-1}$, where $m \geq S_0$. Then $K \in H^m_{\gamma-1}$ after every finite angular word, on the same base domain and with finite bounds of the stated kind.

Proof. The actual momentum block has the form

$$K = \mathrm{TF}_h(\widehat{a} + v), \qquad v \in \operatorname{Ran} B_{1,\gamma}, \qquad \operatorname{div}_h \mathrm{TF}_h\, v = -\operatorname{div}_h \mathrm{TF}_h\, \widehat{a},$$

where $\widehat{a}$ is the fixed smooth Euclidean tracefree premomentum. On Euclidean tracefree tensors, the inverse of TF_h is explicit. If $\operatorname{tr}_h K = 0$, it is

$$t = K - \frac{\operatorname{tr}_\delta K}{\operatorname{tr}_\delta h}\, h, \qquad \operatorname{tr}_\delta t = 0, \qquad \mathrm{TF}_h\, t = K.$$

Thus $v = t - \widehat{a}$ has the preceding spatial and angular regularity of K. The denominator has a fixed positive lower bound in the metric neighborhood. The residual $\rho = -\operatorname{div}_h \mathrm{TF}_h\, \widehat{a}$ has all-angular $H^{m-1}_{\gamma-2}$ regularity: it contains at most one derivative of h, and $\widehat{a}$ has every required finite weighted norm. On the inner ball $h = \widehat{h}$ and $\mathrm{TF}_{\widehat{h}}\, \widehat{a} = \widehat{K}$, so the supplied local momentum constraint makes ρ zero there. Since it is already a whole $H^{m-1}_{\gamma-2}$ field, its exterior restriction belongs to the barred space at that order. Apply Proposition B.11 to this existing v. The formula for K then gives its additional derivative and finite bound. No new fixed point is selected. □

B.7. Parameter regularity on the same branch. The spatial estimates concern one selected extension at all finite orders. To obtain smooth dependence on passive parameters in each of those spaces, it is enough to combine their locally uniform bounds with the base parameter theorem. The

following interpolation argument supplies the precise implication and the finite dependence of its estimates.

Lemma B.13 (Parameter regularity from uniform higher spatial bounds).

Let P be a finite-dimensional open parameter domain, let s_0 be the fixed base order, and put

$$X_s = H^s_\gamma \oplus H^{s-1}_{\gamma-1}, \qquad s \geq s_0.$$

Suppose $F : P \to X_{s_0}$ is smooth. Suppose also that for every finite integer $M \geq s_0$ and every compact $\mathcal{K}' \Subset P$,

$$F(p) \in X_M, \qquad \sup_{p \in \mathcal{K}'} \|F(p)\|_{X_M} < \infty.$$

Then $F : P \to X_s$ is smooth for every finite $s \geq s_0$. On $\mathcal{K} \Subset \operatorname{int} \mathcal{K}'$, every specified finite number of parameter derivatives in X_s is bounded by one sufficiently high spatial bound in the hypothesis, finitely many base-space parameter derivative bounds on $\mathcal{K}'$, and the parameter-chart and cutoff margins. The assertion applies to fixed finite products and to tensor components.

Proof. We first record the interpolation inequality in the precise weighted scale. Choose a smooth dyadic partition of unity in the spatial variable, including one cutoff near the origin. Rescale the jth exterior annulus by $x = 2^j y$. For each fixed weight ω, the squared weighted norm of a field u is equivalent to

$$\sum_{j \geq 0} 2^{-2j\omega} \big\| (\chi_j u)(2^j \cdot) \big\|^2_{H^s}.$$

The same annular factor occurs at every derivative order because the derivative weight in the original norm compensates the rescaling. Extend each cut-off rescaled field by zero to a fixed bounded region and use ordinary Sobolev interpolation there. Hölder's inequality in the annular sum gives

$$\|u\|_{H^s_\omega} \leq C \|u\|^{1-\theta}_{H^{s_0}_\omega} \|u\|^{\theta}_{H^M_\omega}, \qquad \theta = \frac{s - s_0}{M - s_0}, \qquad s_0 < s < M.$$

Apply this to both components, with orders shifted by one in the second component, and apply the finite-dimensional version of Hölder's inequality to their sum. The result is

$$\|u\|_{X_s} \leq C \|u\|^{1-\theta}_{X_{s_0}} \|u\|^{\theta}_{X_M}.$$

All orders here are fixed and finite. The constants need not be uniform as M increases. Noninteger intermediate orders may equivalently be defined by the same rescaled Sobolev construction.

Work in one parameter chart, write its coordinates as $p \in \mathbb{R}^d$, and multiply F by a smooth parameter cutoff supported in $\operatorname{int} \mathcal{K}'$ and equal to one on a neighborhood of $\mathcal{K}$. Extend this product by zero outside the chart. We continue to write F for the product. A finite number of cutoff derivatives enters the bounds below.

The parameter convolutions used next are well defined in each X_M, even before high-space continuity has been proved. One direct verification uses the Hilbert structure. In the integer weighted Sobolev norm, choose a countable orthonormal basis by Gram–Schmidt from a dense family of smooth compactly supported fields. For a fixed such basis element e, the functional $u \mapsto \langle u, e\rangle_{X_M}$ is continuous in the base distributional topology, by integration by parts against the smooth compactly supported coefficients. Its value on $F(p)$ is therefore measurable. The finite basis expansions of $F(p)$ are measurable X_M-valued maps and converge pointwise in X_M. This proves strong measurability. The local high-space bound, together with the parameter cutoff, then gives Bochner integrability. These integrals agree with their base-space versions because the inclusion $X_M \hookrightarrow X_{s_0}$ is continuous and injective.

Fix the desired derivative order q. Choose a finite integer $M > s$ and a finite integer $L > q$ such that

$$L(1-\theta) > q, \qquad \theta = \frac{s - s_0}{M - s_0}.$$

For example, one may take an integer $M > 2s - s_0$ and $L = 2(q+1)$. Choose a compactly supported smooth signed kernel κ of integral one, with all moments of total degrees $1, \dots, L-1$ equal to zero. Such a kernel is obtained from a finite linear combination of distinct dilations of any fixed integral-one smooth kernel: solve the finite Vandermonde system for their coefficients so that the dilation factors cancel each positive degree below L. Put

$$\kappa_t(p) = t^{-d}\kappa(p/t), \qquad F_t = \kappa_t * F.$$

For $0 < t \le t_0$, where t_0 is chosen below the parameter cutoff margin, F_t is smooth with values in every fixed X_M. Differentiating the kernels gives, for $|\alpha| \le q$,

$$\|\partial_p^\alpha (F_{t/2} - F_t)\|_{X_M} \le C t^{-|\alpha|} \sup_{\mathcal{K}'} \|F\|_{X_M}.$$

Taylor's formula for the base-space-valued map, with remainder of order L, and the vanishing moments give

$$\|\partial_p^\alpha (F_{t/2} - F_t)\|_{X_{s_0}} \le C t^{L-|\alpha|} \sum_{j=0}^{L} \sup_{\mathcal{K}'} \|D_p^j F\|_{X_{s_0}}.$$

One can obtain this bound directly by differentiating the difference kernel and subtracting the degree-$(L-1)$ Taylor polynomial of F. Integration by parts shows that every polynomial term has zero integral; the differentiated kernel contributes $t^{-|\alpha|}$ and the Taylor remainder contributes t^L. Thus no high-space parameter derivative of F is used.

Interpolation of the high-space bound and the base-space bound gives

$$\|\partial_p^\alpha (F_{t/2} - F_t)\|_{X_s} \le C t^{L(1-\theta)-|\alpha|} A_L^{1-\theta} B_M^{\theta},$$

$$A_L = \sum_{j=0}^{L} \sup_{\mathcal{K}'} \|D_p^j F\|_{X_{s_0}}, \qquad B_M = \sup_{\mathcal{K}'} \|F\|_{X_M}.$$

The exponent is positive for every $|\alpha| \le q$. Summation at $t = 2^{-n}t_0$ therefore proves uniform convergence in X_s of F_t and all its parameter derivatives through order q on a neighborhood of $\mathcal{K}$ with compact closure in $\operatorname{int}\mathcal{K}'$. The undifferentiated limit is the given map F, since convergence also holds in X_{s_0}. The fundamental theorem of calculus along coordinate segments identifies the limits of the first derivatives with derivatives of F in X_s; iteration identifies all derivatives through order q. Their continuity follows from uniform convergence of continuous maps.

The first mollified term and the summable differences bound these derivatives using only A_L, B_M, the finitely chosen orders and the cutoff margin. For example, the bound can be written as $CA_L^{1-\theta}B_M^{\theta}$, with a larger fixed constant accounting for the initial scale, and hence also as $C(A_L + B_M)$. Since s and q were arbitrary finite orders, this proves smoothness in every X_s. A finite parameter-chart cover proves the assertion on a general finite-dimensional parameter domain. □

Applying Lemma B.13 to the initialized extension now combines its base-order smooth dependence with its locally uniform higher spatial bounds. The parameter neighborhood chosen at base order therefore serves at each finite derivative order.

B.8. Proof of the relative-extension theorem.

Proof of Theorem B.1. Choose the weight offset by Lemma B.2 and the finite base order S_0 before the higher orders to be estimated. Choose the base neighborhood in Proposition B.4 small enough to retain its contraction, the scalar positivity margin of Proposition B.9, and the normal inverse margin of Proposition B.11. These are finitely many conditions at the one base order. Sobolev embedding supplies the required uniform metric and scalar bounds. Rotation-invariant equivalent norms make the neighborhood invariant under proper rotations.

Proposition B.4 gives the initialized pair $(h, K) = \mathfrak{E}(\widehat{h}, \widehat{K})$, its exact protected agreement, its base bound and its fixed-point property on exact whole data. Let $z = (\phi, \beta)$ be its actual scalar correction and v its Euclidean tracefree momentum correction. They belong to the prescribed ranges. Theorem B.6 gives all their angular derivatives in the base spaces. Lemma B.8 therefore makes β smooth with a compactly supported zero extension, with each finite norm controlled by finitely many input norms.

We improve the same pair by induction on the spatial order. At $m = S_0$, and after every finite angular word,

$$\phi \in \overline{H}^{m}_{\gamma}, \qquad K \in H^{m-1}_{\gamma-1}.$$

Assume this at one finite $m \ge S_0$. Proposition B.9 gives $\phi \in \overline{H}^{m+1}_{\gamma}$ and $h - \delta \in H^{m+1}_{\gamma}$, with every angular word retained. The momentum block is

$$K = \mathrm{TF}_h(\widehat{a} + v), \qquad \operatorname{div}_h \mathrm{TF}_h\, v = -\operatorname{div}_h \mathrm{TF}_h\, \widehat{a}, \qquad v \in \operatorname{Ran} B_{1,\gamma}.$$

The algebraic inverse in Corollary B.12 gives the already known $H^{m-1}_{\gamma-1}$ regularity of v after every angular word. The fixed prefield $\widehat{a}$ is smooth at all weighted orders. Together with the improved metric, it gives the required $H^{m-1}_{\gamma-2}$ regularity of the residual. That residual vanishes on B_1 as a whole Sobolev field, hence belongs to the required barred space. Proposition B.11 and Corollary B.12 give $v, K \in H^{m}_{\gamma-1}$ with every angular word. This closes the induction at $m+1$.

Each finite induction uses only finitely many input norms and angular words. The scalar step uses a positive elliptic principal form. The momentum step inverts the fixed angular operator at its base L^2 margin and then differentiates. Both steps use the smallness fixed at base order. The zero-extension estimates hold across the inner sphere, so the induction gives all whole weighted orders and the stated zero-extension regularity of the corrections.

For a supplied finite-dimensional smooth parameter family, Proposition B.4 first gives smooth dependence in

$$X_{S_0} = H^{S_0}_{\gamma} \oplus H^{S_0-1}_{\gamma-1}$$

for $F = (h - \delta, K)$. On a retained compact parameter set, Theorem B.6 gives locally uniform angular bounds. The finite spatial induction then gives locally uniform bounds in every X_M. Apply Lemma B.13 to this same map F. Every specified output derivative on a compact $\mathcal{K} \Subset \operatorname{int} \mathcal{K}'$ is bounded by one sufficiently high spatial bound and finitely many base parameter bounds on $\mathcal{K}'$. The preceding construction expresses both in finitely many supplied input seminorms on that larger set, with the fixed margins. This proves the parameter assertion in Theorem B.1 with its stated order of quantifiers.

Spatial induction and parameter differentiation apply to the solution initialized in Proposition B.4. Hence its equality with the prescribed data on the protected ball and its fixed-point identity on exact whole data still hold. This proves all assertions of Theorem B.1. □

For the supplied extension the initialization and the chosen ranges determine the normalized correction uniquely. The prescribed interior data can admit other whole maximal vacuum extensions.

AI and tool use disclosure

AI systems were used extensively throughout this research and the preparation of this manuscript. The models used were GPT-6 Pro, gpt-6.2-astra (high), Claude Fable 5.2 (xhigh), Claude Opus 5.5 (xhigh), and Gemini 4 Carbon (high). We used SymPy for symbolic checks and NumPy and SciPy for numerical experiments. The computational supplements document the software versions, execution environments, and reproduction steps.

References

Tim Adamo, Lionel Mason (2014). Conformal and Einstein gravity from twistor actions. *Classical and Quantum Gravity* **31**, 045014. DOI.
Christian Bär, Nicolas Ginoux, Frank Pfäffle (2007). *Wave Equations on Lorentzian Manifolds and Quantization*. European Mathematical Society. Author text.
A. Bautista, A. Ibort, J. Lafuente (2014). On the space of light rays of a spacetime and a reconstruction theorem by Low. *Classical and Quantum Gravity* **31**, 075020. DOI.
Antonio N. Bernal, Miguel Sánchez (2005). Smoothness of time functions and the metric splitting of globally hyperbolic spacetimes. *Communications in Mathematical Physics* **257**, 43–50. Author text.
Antonio N. Bernal, Miguel Sánchez (2006). Further results on the smoothability of Cauchy hypersurfaces and Cauchy time functions. *Letters in Mathematical Physics* **77**, 183–197. Author text.
Piotr T. Chruściel, Wan Cong (2023). Gluing variations. *Classical and Quantum Gravity* **40**, 165009. Published article.
Stefan Czimek (2018). An extension procedure for the constraint equations. *Annals of PDE* **4**, article 2. DOI. arXiv version.
Avron Douglis, Louis Nirenberg (1955). Interior estimates for elliptic systems of partial differential equations. *Communications on Pure and Applied Mathematics* **8**(4), 503–538. DOI.
J. J. Duistermaat (1972). Applications of Fourier integral operators. *Séminaire Équations aux dérivées partielles (Polytechnique)* **1971–1972**, exposé 27, 1–24. Online text.
J. J. Duistermaat, L. Hörmander (1972). Fourier integral operators. II. *Acta Mathematica* **128**, 183–269. DOI.
Michael G. Eastwood, Roger Penrose, R. O. Wells, Jr. (1981). Cohomology and massless fields. *Communications in Mathematical Physics* **78**, 305–351. DOI.
Helmut Friedrich, Gabriel Nagy (1999). The Initial Boundary Value Problem for Einstein's Vacuum Field Equation. *Communications in Mathematical Physics* **201**, 619–655. Online text.
Mariano Giaquinta, Luca Martinazzi (2012). An Introduction to the Regularity Theory for Elliptic Systems, Harmonic Maps and Minimal Graphs. Edizioni della Normale. Second edition. DOI.
Christian Gérard, Michał Wrochna (2014). Construction of Hadamard states by pseudo-differential calculus. *Communications in Mathematical Physics* **325**, 713–755. DOI. arXiv version.
Thomas J. R. Hughes, Tosio Kato, Jerrold E. Marsden (1977). Well-posed quasi-linear second-order hyperbolic systems with applications to nonlinear elastodynamics and general relativity. *Archive for Rational Mechanics and Analysis* **63**(3), 273–294. DOI. Online text.

Lior Isthmus (2026a). A projective twistor correspondence for local analytic Einstein germs. Version 1.1. DOI. Online text.
Lior Isthmus (2026b). Global twistor sources for analytic Einstein spacetimes. Version 1.0.1. DOI. Online text.
Claude R. LeBrun (1984). Twistor CR manifolds and three-dimensional conformal geometry. *Transactions of the American Mathematical Society* **284**(2), 601–616. DOI.
Philippe G. LeFloch, The-Cang Nguyen (2023). The seed-to-solution method for the Einstein constraints and the asymptotic localization problem. *Journal of Functional Analysis* **285**, 110106. DOI.
Robert J. Low (2001). The space of null geodesics. *Nonlinear Analysis* **47**, 3005–3017. Publisher.
Yuchen Mao, Sung-Jin Oh, Zhongkai Tao (2023). Initial data gluing in the asymptotically flat regime via solution operators with prescribed support properties. arXiv:2308.13031.
Lionel J. Mason (1989). Hypersurface twistors. *Twistor Newsletter* **28**, 20–23. Online text.
Lionel J. Mason, Roger Penrose (1989). A twistorial approach to the full vacuum equations. *Twistor Newsletter* **29**, 1–5. Online text.
Richard Melrose (2006). Lectures on Pseudodifferential operators. Massachusetts Institute of Technology. Online text. MIT 18.157, version 0.7E.
Roger Penrose (1976). Nonlinear gravitons and curved twistor theory. *General Relativity and Gravitation* **7**(1), 31–52. DOI.
Roger Penrose (2015). Palatial twistor theory and the twistor googly problem. *Philosophical Transactions of the Royal Society A* **373**, 20140237. DOI.
M. J. Radzikowski (1996). Micro-local approach to the Hadamard condition in quantum field theory on curved space-time. *Communications in Mathematical Physics* **179**, 529–553. DOI.
Atul Sharma (2021). Twistor action for general relativity. arXiv:2104.07031v2.
Jan Sbierski (2016). On the existence of a maximal Cauchy development for the Einstein equations: a dezornification. *Annales Henri Poincaré* **17**(2), 301–329. DOI. arXiv version.
James W. York, Jr. (1974). Covariant decompositions of symmetric tensors in the theory of gravitation. *Annales de l'Institut Henri Poincaré, Section A* **21**, 319–332. Original article.