Journal Name

Crossmark

RECEIVED
dd Month yyyy

REVISED
dd Month yyyy

PAPER

Spin-2 Sector of the Dyality-Ordered Electromagnetic Vacuum: An Effective Field Theory

Chris Scott[1*]

[1]Royal Australian Air Force, Department of Defence, Canberra, Australia
[*]Author to whom any correspondence should be addressed.

E-mail: christopher.scott7@defence.gov.au

Keywords: emergent gravity, effective field theory, spin-2 collective mode, two-potential electrodynamics, toroidal multipoles, Lorentz violation

Abstract

We develop a medium effective field theory for a transverse-traceless spin-2 collective response in the unreduced electromagnetic near-field sector. The starting point is the observation that the usual Ward-projected radiative description removes source-bound longitudinal and toroidal multipole structure from the asymptotic photon Hilbert space. Retaining this local near-field structure leaves a compact dyality phase which defines a local, constrained $U(1)_\phi$ connection over the ordinary electromagnetic $U(1)_\Theta$ bundle. The mixed curvature invariant $\int F_\Theta \wedge F_\phi$ then supports codimension-two defect structure and a natural symmetric transverse-traceless bilinear. We identify the corresponding excitation as the dyality plasmon: a massive $\ell = 2$ collective mode of the dyality-ordered near-field medium. This mode is not a fundamental graviton and does not define a dynamical metric. It carries only the two projected tensor polarisations of the medium response, avoiding the consistency problems of nonlinear massive gravity. The long-wavelength gravitational interpretation is of emergent $f(R)$-type vacuum elasticity, with the scalar and tensor branches belonging to distinct collective channels. We derive the medium-frame EFT, its projected source, propagator and admissible foliation structure. A schematic ultraviolet completion by electroweak-scale Witten dyons motivates the charge weights of the mixed-curvature sector; the gauge model itself is not constructed. Dirac quantisation then gives a conditional phase-velocity estimate $c_h/c = 1/\left(2\sqrt{\alpha(m_Z^2)\,\alpha(q^2)}\right)$, yielding $c_h \simeq 64.5c$ at the Z pole. The scalar and vector channels are derived rather than projected away: the vector channel is empty and the surviving scalar is the condensate amplitude mode. A conditional pole location for the tensor branch and a cosmological resonant near-field search in the 250–470 GHz band follow.

1 Introduction

The Ward identity in quantum electrodynamics enforces the transversality of the photon propagator. Non-transverse components do not propagate as independent external photons. In practical terms it enacts what Siegert [1] identified in 1937 for multipole radiation: the source-bound near-field components of a current, non-radiating, longitudinal, and toroidal, are traded for the charge distribution by parts, leaving no trace in asymptotic S-matrix observables. The reduction is exact for the radiative far field [2, 3]. Applied routinely to the vacuum, it removes the toroidal sector from the standard Hilbert space altogether. Annex C of [4] makes the point explicit:

> "in general, as a functional relation, the Siegert theorem is simply invalid. This approximation is valid only when the electromagnetic system has no toroid component."

The transverse-photon Fock vacuum is therefore the vacuum of the reduced radiative sector. It is not a completeness statement about all local near-field response.

The near-field content reappears as higher-order propagator corrections, fixed by the transverse data through the Ward identity, so it may seem already accounted for. This conflates two things. Dressing the reduced propagator reproduces every asymptotic observable, which is why the Siegert

reduction is exact for the S-matrix. It does not follow that the near field has no independent content. The toroidal configurations with $E \cdot B \neq 0$ are genuine local degrees of freedom of the oscillator cell, not artifacts recoverable within the reduced sector. The two descriptions agree on every asymptotic observable and diverge on the local response the reduction projects out. Structured metamaterials already show this content is physical: unreduced sub-wavelength toroidal dynamics carry signatures missed entirely by far-field probes [5, 6]. A covariant field theory of the same sector in the vacuum has been absent, a gap this paper attempts to fill.

Retaining the near-field response changes how the vacuum Hilbert space analogue is built. Conventional field modes map to non-compact amplitude quadratures on an open phase space. The unreduced cell instead carries a Hamiltonian structure with compact phase variables, and the resulting system is near-integrable. Section 2 defines this pre-reduced vacuum: the cell Hamiltonian, its surviving invariant tori, the defect remnants they leave, and the regime of validity. It supplies the compact T^3 on which the rest of the construction rests.

This medium is driven by electromagnetic dyality, a notion usually obscured in standard treatments. Passive duality rotates electric and magnetic fields into one another while the sources are held fixed, $E \to E\cos\phi + B\sin\phi$. Dyality is the active counterpart: it rotates the fields and the electric-magnetic source doublet together, as a dyonic configuration requires [7]. In a source-free radiative vacuum the distinction carries no local content. In the near field, where currents and fields are mutually enclosed, a dyality transformation must act on the sources and the fields simultaneously [4]. Field and source are locked together in the sub-wavelength regime, and no global gauge choice separates them. The dyality angle $\phi \in S^1$ is thus a genuine internal degree of freedom, a compact $\mathrm{U}(1)_\phi$ connection over spacetime.

The construction rests on identifying two structures of codimension-2. The first is topological: the phase defects of the $\mathrm{U}(1)_\phi$ bundle, where the dyality angle is undefined, are codimension-2 loci carrying mixed Chern charge in the sense of Chern and Ishida [8]. The second is dynamical: the resonances of the near-field flow, which break as the perturbation grows and release the web that carries transport. We propose that the first is what survives the breaking of the second. The dyality angle ϕ is the angle conjugate to the toroidal action of a cycle $\mathbf{w}$, so the dyality defects of that cycle, the amplitude zeros of its toroid doublet, carry winding $\mathbf{w}$ and are moved by that cycle's dynamics. While the cycle sits on an invariant torus these defects are confined to compact filaments. When the resonance $\mathbf{w} \cdot \omega(I) = 0$ breaks, the confinement goes but the winding cannot, so the defect persists and is released into the web. It is the topologically protected residue of a destroyed resonance; its charge is unchanged by the release and only the geometry of its locus is, as the computation of Appendix B illustrates. The topological reading then supplies the charge and the spin-2 composite; the dynamical reading supplies the substrate and its transport. This is the central structural hypothesis of the paper and of the proposed near-field oscillator model. It is not a theorem. The identification of ϕ with that conjugate angle is derived at leading order in the coherent sector (Section 4.2); the surviving hypothesis is the projection of the defect web onto the broken-resonance web.

Localising the dyality angle yields a mixed Chern-2 invariant $\int F_\Theta \wedge F_\phi$, a topological backbone invisible to Ward-projected far-field amplitudes. A bilinear of the electromagnetic and dyality curvatures, projected to the transverse-traceless sector, gives a spin-2 composite operator with no analogue in the on-shell photon sector. We term this excitation the dyality plasmon: the zero-sound elastic wave of the phase-ordered near-field medium.

This paper develops the medium effective theory of that excitation. It establishes the Lagrangian, the propagator, and a proposed ultraviolet completion. The central claim is narrow: the unreduced toroidal near-field sector admits a transverse-traceless $\ell = 2$ collective response described by a medium EFT. The bolder elements, the dyon completion and the scale at which c_h is fixed, are conditional and testable, and we mark them where they appear. The collective mode is an excitation of the medium, not a fundamental field. Its natural gravitational target is the emergent $f(R)$-type vacuum elasticity of Volovik [9, 10], the 4-form q-theory of Klinkhamer and Volovik [11, 12], and the emergent electroweak gravity of McElrath [13]. The dyality-ordered vacuum supplies the trans-Planckian substrate for these constructions.

The two-potential structure at the heart of the construction is not new physics but a revived formalism. Zwanziger's local Lagrangian for electric and magnetic charges [14] introduced a dual potential and a fixed timelike vector, and remained a formal niche for decades, for want of monopoles and because the fixed vector broke manifest Lorentz invariance [15]. The present construction inverts that liability: the medium possesses a physical rest frame u^μ, so the fixed vector, a liability for the formalism as fundamental physics, is here its content. The lineage through its modern covariantisation, and its structural match to this theory's foliation apparatus, are set

out in Section 3.1.

The main results separate into established core claims and conditional propositions.

(C1) The unreduced toroidal near-field sector admits a transverse-traceless $\ell = 2$ collective response, the dyality plasmon, described by a medium effective field theory. The collective pole is derived in Section 6.2 within a specified collective-response model, under the stated kernel and stability assumptions, rather than posited at the level of the action. This is the central claim of the paper.

(C2) The mode is an emergent collective tensor response of the dyality-ordered medium, not a fundamental graviton. It carries two polarisations and defines no metric; the massive-gravity counting, five healthy polarisations with the Boulware–Deser mode a sixth [16], never applies, and the long-wavelength gravitational context is of $f(R)$ type, whose only extra propagating mode is a scalar. The $\ell = 2$ branch is selected by the TT projector relative to the condensate rest frame u^μ.

(C3) The composite arises from the mixed Chern-2 defect structure of the $\mathrm{U}(1)_\Theta \times \mathrm{U}(1)_\phi$ bundle. The defect count and the elementary charge unit $|C_2| = \frac{1}{6}$ follow, given the comoving inertial match of Section 8.3: the count is $b_1(T^3) = 3$ doubled by orientation, the balance $\sum_\alpha C_\alpha = 0$ is the flatness of the comoving inertial match under the Regge reading of Section 8.3, and the unit $\frac{1}{6}$ follows from the isotropy of that match together with integrality of the total Chern number (Sections 4 and 8.3).

(C4) The medium-frame action is the transverse-traceless projection of a covariant field-strength polysupermetric, with $c_h^2 = Z_s/Z_t$ its eigenvalue ratio; the Lorentz-invariant (general-relativity-equivalent) case is the degenerate eigenvalue $Z_s = Z_t$. The associated many-fingered (Tomonaga–Schwinger) evolution is consistent on the foliation-preserving (projectable) class selected by the comoving frame, the obstruction to full refoliation invariance being a single term $\propto (c_h^2/c^2 - 1)$ that is removed by either of two independent conditions: projectability, which holds for any c_h, or the acoustic limit (Section 8).

The following carry explicit, testable hypotheses.

(P1) The phase velocity is $c_h/c = q_\phi/q_\Theta$, the stiffness-to-inertia ratio of the tensor channel, with the stiffness assignment selected by the transformation law (10) under the charged-rotor constitutive assignment of Section 7 (the dyality phase carries q_ϕ alone). Conditional on the emergent-gravity stiffness normalisation and the scale-running prescription of Section 7.3, Dirac quantisation fixes it at the Z pole to $c_h = c/(2\alpha(m_Z^2)) \approx 64.5\,c$, running as $c_h(q)/c = 1/(2\sqrt{\alpha(m_Z^2)\,\alpha(q^2)})$ to $\approx 66.5\,c$ in the infrared (Section 7). This infrared running of the coupling is distinct from the acoustic limit $Z_s \to Z_t$ invoked elsewhere: the former is the low-q value of the charge ratio, the latter an independent degeneracy of the kinetic coefficients that this construction does not reach. Decoupling of the gapped branch at low frequency does not imply either.

(P2) A Witten-dyon ultraviolet completion path places the defect-core scale at $m_\chi \geq m_W/\alpha_W^{(2)} \approx 2.3$ TeV via the 't Hooft–Polyakov BPS bound. This is the model's natural core scale, not a collider exclusion; assigning direct limits would require a dedicated analysis of highly ionising or magnetically charged states.

The argument of the paper runs as follows. The premise is that the Siegert reduction, exact for the S-matrix, removes source-bound toroidal structure that is nonetheless a local degree of freedom; Section 2 retains it and obtains a near-integrable vacuum with a surviving T^3. The construction localises the dyality phase of that sector, producing a two-$U(1)$ bundle whose mixed Chern-2 structure supports codimension-two defects with quantised fractional charge and a natural spin-2 bilinear (Sections 3 and 4). The derivation step obtains the collective spectrum of the ordered medium channel by channel, with the tensor branch the only propagating one of the retained near-field sector on the admissible frame class (Section 5), and writes its effective theory (Section 6). The inputs are then confined to one place: the charged-rotor constitutive assignment and the candidate dyon completion, which together set the conditional velocity benchmark (Section 7). Consistency is established canonically on the admissible foliations and causally through the front-velocity analysis (Sections 8 and 9). The chain terminates in its context and its consequences: the emergent-gravity reading (Section 10), a conditional pole location, and a specified experiment (Section 11), with the epistemic status of principal statements inventoried in

Section 12. The named hypotheses enter at four points: the defect-resonance identification (Section 4), the constitutive assignment and the dyon completion (Section 7), the acoustic-limit inertial match (Section 8.3), and the condensate identification (Section 11).

2 The pre-reduced near-field vacuum

We now construct the pre-reduced vacuum as a near-integrable Hamiltonian system. In conventional field theory, field modes are mapped onto non-compact amplitude quadratures spanning an open phase space ($\mathbb{R}^2$). Retaining the local near-field response instead assigns each oscillator cell a Hamiltonian structure with compact phase variables. The electromagnetic field on a toroidal near-field cell admits $H = H_0 + \varepsilon H_1$, where

$$H_0 = \frac{1}{8\pi}\int \left(\mathbf{E}^2 + \mathbf{B}^2\right) d^3\mathbf{r} \tag{1}$$

is the free Maxwell part, with $\mathbf{E}$ and $\mathbf{B}$ the spatial components of F^{Θ}, and the perturbation

$$\varepsilon H_1 = -\frac{1}{4\pi}\int \mathbf{E}\cdot\mathbf{P}\, d^3\mathbf{r} \;+\; 2\pi\int \mathbf{P}^2\, d^3\mathbf{r} \;-\; \frac{1}{4\pi}\int \mathbf{M}\cdot\mathbf{B}\, d^3\mathbf{r} \tag{2}$$

encodes the dyality coupling, with $\mathbf{P}$ and $\mathbf{M}$ the electric and magnetic toroidal polarisation fields [17], identified in Section 3 with the components of the residual dyality curvature ΔF^{ϕ} of (13). This follows from Dubovik's dual-potential formalism [17], in which the field equations contain four unknown vector fields and are underdetermined until constitutive relations are supplied.[1] The dyality decomposition $F^{\phi} = \star F^{\Theta} + \Delta F^{\phi}$, introduced in Section 3, closes the system.

The retained cell variable is fixed by the multipole structure of the near-field current, not chosen by hand. After the Helmholtz split removes the longitudinal charge sector, the transverse current carries two independent scalar families, poloidal (magnetic) and toroidal. The leading toroidal member is the toroid dipole

$$\mathbf{T} = \tfrac{1}{10}\int \left[(\mathbf{r}\cdot\mathbf{j})\mathbf{r} - 2r^2\mathbf{j}\right] d^3r, \tag{3}$$

the $\ell = 1$ toroidal harmonic, a polar vector with $2\ell + 1 = 3$ components [4, Eqs. (2.5), (2.11) and App. E]. Retaining it as the cell coordinate gives the phase space $\mathcal{P}_x = T^*\mathbb{R}^3$ with free part $H_0 = \sum_{k=1}^{3} \frac{1}{2}(p_k^2 + \omega_k^2 q_k^2)$, where $\mathbf{q} = (q_1, q_2, q_3)$ are the Cartesian amplitudes of $\mathbf{T}$ and the ω_k are the cell frequencies. This truncation to $n = 3$ degrees of freedom, one per cycle of the cell, is the leading term of the sub-wavelength expansion: higher toroidal multipoles enter at $O(ka)$ for a cell of size a and probe wavenumber k, so it holds when $ka \ll 1$ (Section 6). In action-angle variables $q_k = \sqrt{2I_k/\omega_k}\cos\theta_k$, the integrable part of H written as a function of the actions alone is the averaged (Birkhoff) normal form $h(I)$. The system is near-integrable: H_0 is integrable, the coupling εH_1 breaks it, and the surviving invariant tori prevent full ergodicity.

2.1 Nondegeneracy of the coarse cell

At leading order the toroidal spring $2\pi\int \mathbf{P}^2$ gives the quadratic stiffness $\frac{1}{2}\sum_k \omega_k^2 q_k^2$, anisotropic through the cell geometry, and the harmonic normal form $h_0(I) = \sum_k \omega_k I_k$ is isochronous, $D_I^2 h_0 = 0$. A nondegenerate twist must therefore come from the next term in the same stiffness. Rotational invariance fixes it: a polar vector admits no cubic invariant and a unique quartic one, so the stiffness expands as

$$2\pi\int \mathbf{P}^2 \;\longrightarrow\; \tfrac{1}{2}\sum_k \omega_k^2 q_k^2 + \frac{\kappa(a)}{4}\left(\mathbf{q}^2\right)^2 + O(q^6) \tag{4}$$

with $\kappa(a)$ the effective quartic coefficient at the coarse-graining scale a, generated by the unresolved subcell structure. We assign it no universal microscopic value; the only input is the hardening condition $\kappa(a) > 0$, equivalently boundedness of the cell energy from below. Averaging the quartic over the angles gives

$$h_{\mathrm{eff}}(I) = \sum_{k=1}^{3}\omega_k I_k + \kappa(a)\left[\tfrac{3}{8}\sum_k \frac{I_k^2}{\omega_k^2} + \tfrac{1}{2}\sum_{i<j}\frac{I_i I_j}{\omega_i\omega_j}\right] + O(I^3) \tag{5}$$

[1] Equation (2) reproduces the toroidal-polarisation sector of the Hamiltonian of Ref. [17], Eq. (2.10): the field energy, the $2\pi\int \mathbf{P}^2$ toroidal spring, and the $\mathbf{M}\cdot\mathbf{B}$ term, with $\mathbf{P} \to \mathbf{P} + \mathrm{curl}\,\mathbf{T}_e$. The prefactor convention here differs from that reference, which writes $\mathbf{D}$ where we write $\mathbf{E}$ and carries unit coefficients; the $2\pi\int \mathbf{P}^2$ term matches directly. The oscillator frequencies ω_k introduced below depend on this normalisation and should be fixed against Ref. [17] for a given cell.

with action Hessian

$$D_I^2 h_{\rm eff} = \kappa(a)\, DAD + O(I), \qquad D = \mathrm{diag}(\omega_k^{-1}), \quad A = \tfrac{1}{4}\mathbb{1}_3 + \tfrac{1}{2}J_3 \tag{6}$$

The eigenvalues of A are $\{7/4, 1/4, 1/4\}$, and $\det(D_I^2 h_{\rm eff}) = \frac{7}{64}\,\kappa(a)^3/(\omega_1\omega_2\omega_3)^2 > 0$, so $h_{\rm eff}$ is strictly convex. The anharmonic stiffness lifts the isochronous degeneracy, and supplies convexity, a property required to for a medium capable of diffusion.

2.2 Integrable backbone

The convex near-integrable structure organises the phase space into two coexisting sectors. The first is the integrable backbone, the surviving invariant tori. Here ε measures the defect-sourced perturbation; it is largest in the ultraviolet, where a cell of size a at a defect core, $\xi_{\rm core} \sim 1/m_\chi$, sees the strongest perturbation, so the system sits near the KAM transition ε_c rather than at small ε. By Kolmogorov's theorem the tori of H_0 with sufficiently irrational frequency ratios survive the perturbation [18, 19], and the most irrational, noble torus survives to ε_c. For $n = 3$ these tori have codimension two and do not separate the energy surface, but each is a stable T^3. This surviving T^3 is the rigid geometric backbone of the construction: it is the ambient manifold on which the Chern–Ishida topology of Section 4 is built, and Berry's adiabatic theorem [20] governs the transport over its cycles.

2.3 Diffuse medium

The second sector is the diffuse medium, the Arnold web threading the gaps between surviving tori. The resonant tori break first. A broken resonance $\mathbf{w}\cdot\omega(I) = 0$ leaves an island chain with a chaotic separatrix layer, and for $n \geq 3$ degrees of freedom, those layers connect into the web. The defects of Section 4 are the zeros of the toroidal doublet of each cycle, which carry integer winding and are moved by that cycle's dynamics: on an invariant torus a zero's worldline is confined, and when the governing resonance breaks the confinement is lost while the winding, being topological, is not. A broken resonance therefore releases the zeros of its cycle into the web, and it is these released defects that the web transports.[2] Under the convexity established above this web is not inert. Cheng and Xue prove that a nearly integrable convex system with $n \geq 3$ admits genuine Arnold diffusion for a typical perturbation. We assume, and do not verify here, that the present cell model satisfies the genericity and near-integrability hypotheses of that theorem; convexity alone does not establish them, and the existence of connecting orbits is not by itself a derivation of coherent transport. There exist orbits whose actions drift along the resonant web from one region to another, and, as a corollary, orbits whose closure is dense in the phase space [25]. The surviving Diophantine tori are never crossed because for $n \geq 3$ these tori no longer divide the energy surface. The web between tori is connected, and the diffusing orbits thread it without entering the tori. The region orbits reach is dense in action space but not of full measure. The web therefore carries genuine action transport between the defects, making the diffuse sector a true transport medium rather than a set of isolated stochastic layers. It is the dynamical substrate on which the proposed collective excitation, identified in Section 6 as the dyality plasmon, propagates. The convexity required to ensure the existence of this medium is exactly the condition $\kappa(a) > 0$ supplies.

The two sectors act on widely separated timescales, and should not be confused. The collective excitation propagates fast, at the stiffness-to-inertia velocity c_h of Section 7.3. Read on the stiffness side, this same split is the origin of that velocity: H_0 supplies the temporal inertia of the collective mode through the electromagnetic field energy, while εH_1 supplies its spatial stiffness through the dyality coupling, a correspondence made precise in Section 7.3. Arnold diffusion is exponentially slow, the action drift confined by Nekhoroshev's theorem $t \lesssim \exp(c\,\varepsilon^{-b})$ [26]; the stickiness of island boundaries [27, 28] adds to this but is not the mechanism, since for $n \geq 3$ nothing divides the energy surface and the suppression is the slowness of the drift, not a barrier. The relation is that of fast lattice waves to slow defect creep on the same medium: the excitation is not transported by the diffusion. The diffusion's role is only to keep the substrate dynamically connected rather than frozen, so that the surviving T^3 and its defect web remain coherent over times far exceeding any process in the EFT window. As the condensate develops below m_W the perturbation dilutes with scale, and the infrared inherits this structure as a coarse-grained defect sector. This picture is consistent with the bulk constraint $\Delta F^\phi = 0$ of (14).

[2] In the two-dimensional reduction the destroyed irrational circles leave Aubry–Mather cantori [21, 22], whose breaking is fixed by the Greene residue criterion [23]; these are remnants of irrational circles, not of resonances, and are not the defects. For the $n = 3$ cell weak KAM theory guarantees action-minimising invariant measures in general [24]; the release mechanism stated in the text is what the identification of Section 4.2 asserts, and Appendix B illustrates its dynamical half in a computable model at the same number of degrees of freedom.

The surviving torus is not only a dynamical invariant; its relative angle variables are compact phases, and the toroidal angle conjugate to the retained action is, by the cell map of Section 4, the dyality angle itself. Identifying the ambient manifold of the ordered vacuum with this T^3 and localising its compact phase over spacetime is what promotes the Hamiltonian cell to a principal $\mathrm{U}(1)_\Theta \times \mathrm{U}(1)_\phi$ bundle. The next section conducts the geometric lift of the structure just described.

3 Vacuum bundle and mixed Chern-2 structure

The vacuum sector considered here is not the source-erased radiative Fock sector. It is the pre-reduced near-field sector in which the electromagnetic field and the current or constraint data responsible for its reactive multipole structure are retained. This distinction is essential.

We describe this pre-reduced sector by a principal bundle

$$P \to M^4, \qquad \mathcal{G} = \mathrm{U}(1)_\Theta \times \mathrm{U}(1)_\phi \tag{7}$$

The connection is

$$A = (A_\Theta, A_\phi) \tag{8}$$

with curvature two-forms

$$F_\Theta = \mathrm{d}A_\Theta, \qquad F_\phi = \mathrm{d}A_\phi \tag{9}$$

Here A_Θ is the ordinary electromagnetic connection. The second connection A_ϕ is not an independently propagating gauge field. It is a constrained relative-phase connection which records the source-field phase information retained in the pre-reduced near-field sector.

The phase $\phi \in S^1$ is the dyality coordinate. Unlike a passive electromagnetic duality rotation, which relabels (E, B) while holding the source sector fixed, dyality acts simultaneously on field and source [7]:

$$\begin{pmatrix} E \\ B \end{pmatrix} \to R(\phi) \begin{pmatrix} E \\ B \end{pmatrix}, \qquad \begin{pmatrix} \rho_e \\ \rho_m \end{pmatrix} \to R(\phi) \begin{pmatrix} \rho_e \\ \rho_m \end{pmatrix} \tag{10}$$

where $R(\phi)$ is the SO(2) rotation matrix. This definition should be read literally. If the source sector is absent, dyality reduces to passive electromagnetic duality and carries no independent local content beyond the usual rotation of (E, B). Dyality becomes a physical local degree of freedom only when the source information has not been discarded.

In the vacuum sector studied here, this source information is not an externally inserted Maxwell current. We set

$$j^\mu_{\mathrm{ext}} = 0 \tag{11}$$

The relevant source data are instead boundary, edge, or topological defect data of the near-field configuration. Equivalently, one may regard the local vacuum response as a phase-space decomposition into real electromagnetic oscillator cells before their reactive multipole content is mapped into the transverse radiative basis. Within such a cell the simultaneous source-field rotation is meaningful because the local current or constraint data responsible for the near field are still present. After the far-field reduction, this information is no longer represented as an independent degree of freedom.

Because source and field co-rotate under (10), no external frame can distinguish the two within a self-contained toroidal configuration. Such configurations are characterised by mutually enclosed source currents and field lines [4]. The condition

$$E \cdot B \neq 0 \tag{12}$$

in the interior of the configuration is the local signature of this self-enclosure. It is the condition under which the dyality fibre becomes physical rather than a passive relabelling of the electromagnetic field.

Localising $\phi \to \phi(x)$ therefore does not introduce a new massless spin-1 field. It introduces a constrained relative connection A^ϕ_μ whose curvature decomposes as

$$F^\phi = \star F^\Theta + \Delta F^\phi \tag{13}$$

The term $\star F^\Theta$ is the Hodge-dual electromagnetic contribution associated with source-field co-rotation. The residual curvature ΔF^ϕ is supported only by boundary, edge, or topological defect sectors. It measures the part of the local source-field phase structure not captured by the Hodge-dual electromagnetic component. Thus A^ϕ_μ is not a dyality photon and does not carry an independent Maxwell kinetic term in the bulk.

The magnetic component revealed by ΔF^{ϕ} should not be interpreted as a free Dirac monopole observable from outside the configuration. It is accessible only in the dyonic sense, as the magnetic component exposed by the simultaneous source-field rotation (10). The situation is analogous to nested vortex structures in a superfluid [29]: the circulation of an inner vortex is meaningful only relative to the co-moving phase frame of the outer configuration. In the same way, the dyality curvature is meaningful only relative to the electromagnetic near-field structure whose source information has not yet been quotiented out.

Since A^{ϕ} carries no bulk kinetic term, varying H_1 with respect to A^{ϕ} yields a constraint rather than a wave equation:

$$d\star\Delta F^{\phi} = J_{\phi} \tag{14}$$

where J_{ϕ} is the toroidal defect 3-current localised on the worldsheets $I^{(\alpha)}$. In the bulk, $J_{\phi} = 0$ and Bianchi for F^{Θ} give $d\star\Delta F^{\phi} = 0$; with the constitutive relation of Section 7 and decay of the residual curvature at spatial infinity, the only solution is $\Delta F^{\phi} = 0$, and $F^{\phi} = \star F^{\Theta}$ recovers there. Harmonic residuals supported by nontrivial topology or by boundary data are excluded by that fall-off condition, which is part of the specification of the bulk sector. The independent dynamical content of the theory is therefore carried entirely by the defect sector, and the mixed Chern-2 form separates as

$$\int_{M^4} F^{\Theta} \wedge F^{\phi} = \int_{M^4} F^{\Theta} \wedge \star F^{\Theta} + \int_{M^4} F^{\Theta} \wedge \Delta F^{\phi} \tag{15}$$

where the first term is the standard Maxwell self-energy and the second is the defect-supported contribution absent from any far-field description.

The corresponding mixed topological invariant is built from the dimensionless normalized curvatures

$$\mathcal{F}_{\Theta} = \frac{q_{\Theta}}{\hbar c} F_{\Theta}, \qquad \mathcal{F}_{\phi} = \frac{q_{\phi}}{\hbar c} F_{\phi} \tag{16}$$

as the Chern-2 form [8]

$$C_2 = \frac{1}{8\pi^2} \int_{M^4} \mathcal{F}_{\Theta} \wedge \mathcal{F}_{\phi} \tag{17}$$

Prequantizability requires the periods of the mixed curvature to be integral,

$$\frac{1}{2\pi} \int_{\Sigma} \mathcal{F}_{\Theta} \wedge \mathcal{F}_{\phi} = n, \qquad n \in \mathbb{Z} \tag{18}$$

In terms of the physical field strengths this reads

$$\frac{q_{\Theta} q_{\phi}}{2\pi(\hbar c)^2} \int_{\Sigma} F_{\Theta} \wedge F_{\phi} = n \tag{19}$$

so the period condition fixes the charge product $q_{\Theta} q_{\phi}$ once the geometric flux is specified. In Section 7 the unit period reduces to the Dirac quantisation condition $q_{\Theta} q_{\phi} = \hbar c/2$, rather than that condition being inserted by hand. This structure is the topological backbone of the effective theory from which the spin-2 sector is extracted.

It is natural to ask whether the dyality-ordered vacuum follows from Maxwell electrodynamics, is an effective phenomenological description, or requires degrees of freedom beyond Maxwell theory. The answer is the third option, in a specific sense. The ordered vacuum does not follow from Maxwell theory supplemented by the standard radiative reduction, because that reduction removes the toroidal sector in which the order lives. Retaining the near-field sector is necessary but not sufficient: the ordering itself is carried by the defect sector, whose elementary charge carrier is the dual-charged state of Section 7. The description of the ordered medium given here is effective below the condensation scale m_W, with the microscopic origin of the order residing in the defect sector and its candidate completion. The stability of the ordered state is not established in this paper. The consistency conditions supplied here are the recovery of conventional electrodynamics at $k \to 0$ and the constraint structure of (14). A third condition applies not to the ordering itself but to the relic pair-correlated fraction that fixes the gap: the two stability results of Section 11 fix a double well in ϕ_c with a barrier between its minima, so that state is at least locally stable against both formation and decay. The energetics of the ordering condensate at m_W, and the coefficient magnitudes of that double well, remain among the open items of Section 12.

3.1 Gauge structure

The bundle carries two gauge actions and they play different roles. Under $U(1)_\Theta$, $A^\Theta \to A^\Theta + \mathrm{d}\lambda$ with the usual Maxwell orbit; the toroid polarisations $\mathbf{T}^e, \mathbf{T}^m$ are built from source distributions and are invariant, so the doublet section $\Psi_{\mathbf{w}}$ and its phase are Θ-gauge invariant. Under $U(1)_\phi$, $A^\phi \to A^\phi + \mathrm{d}\chi$ with $\phi \to \phi + \chi$: this relabels the Seifert leaves, the level sets of ϕ, and is the redundancy of the implicit representation of Section 4. The physical content of the ϕ sector is exhausted by the winding (38), the defect set (37), and the singular curvature (40), all invariant under both actions. Because A^ϕ carries no bulk kinetic term, its field equation is the constraint (14) rather than a wave equation, so localising ϕ introduces no new propagating gauge orbit: the constraint removes the would-be bulk degree of freedom, and $F^\phi = \star F^\Theta$ is recovered identically away from the defects. The one residual redundancy is the framing of open defect strings ending on a slice boundary; for closed defect loops the linking data of Section 4.3 are the complete gauge-invariant content.

This structure is not new to this construction; it is two-potential electrodynamics. In Zwanziger's local formulation [14] a single compact $U(1)$ is described by two potentials, one coupling to electric and one to magnetic charge, sharing one photon through a constraint that involves a fixed timelike vector, and the dual potential carries no independent propagating mode. The constrained A^ϕ above is that dual potential, and the fixed vector, historically the formalism's liability, is here the physical medium frame u^μ.

The correspondence extends to the modern covariant form: the Pasti–Sorokin–Tonin covariantisation replaces the fixed vector by $v_\mu = \partial_\mu a / \sqrt{-(\partial a)^2}$ with $a(x)$ an auxiliary scalar [30], a unit timelike hypersurface-orthogonal frame built from a time function, which is precisely the foliation-defining (khronon) structure of the admissible class of Section 8. The requirement that a admit no constant gauge fixing is the global-hyperbolicity condition that appears independently in Section 9 as the global time function [31]. Gauge fixing the covariant form in four dimensions returns the Zwanziger construction [32]. The medium formulation is therefore the natural habitat of the two-potential formalism: what was auxiliary structure for fundamental fields is physical structure here. No further constraints or redundancies arise.

4 Spin-2 composite, defect charge, and linking

Which local operator carries the tensor content of the mixed sector is fixed by symmetry. A rank-2 bilinear built from one factor of each curvature, invariant under both gauge actions, and vanishing whenever either sector is absent is necessarily a contraction of F^Θ with F^ϕ; its antisymmetric part carries the $\ell = 1$ content and its trace the $\ell = 0$ content, so the symmetric-traceless contraction is the minimal carrier of the $\ell = 2$ channel. The local spin-2 operator is accordingly defined by

$$\mathcal{O}^{(2)}_{ij}(x) = F^\Theta_{i\ell}(x)\, F^\phi_{j\ell}(x) + F^\phi_{i\ell}(x)\, F^\Theta_{j\ell}(x) - \tfrac{2}{3}\delta_{ij}\, F^\Theta_{k\ell}(x)\, F^\phi_{k\ell}(x) \tag{20}$$

Equation (20) is bilinear in curvature, symmetric in i and j, and traceless by construction; it vanishes if either curvature sector is absent. Using the decomposition (13), the operator separates into a locked co-rotation part and a relative-acceleration part:

$$\mathcal{O}^{(2)}_{ij} = \underbrace{E_iB_j + B_iE_j - \tfrac{2}{3}\delta_{ij}(\mathbf{E}\cdot\mathbf{B})}_{\Delta F^\phi = 0 \text{ limit}} + F^\Theta_{i\ell}\,\Delta F^\phi_{j\ell} + \Delta F^\phi_{i\ell}\, F^\Theta_{j\ell} - \tfrac{2}{3}\delta_{ij}\, F^\Theta_{k\ell}\,\Delta F^\phi_{k\ell} \tag{21}$$

The first term is the symmetric-traceless cross-field stress. It is nonzero whenever both fields are, since $Q_{ij}Q_{ij} = 2E^2B^2 + \frac{2}{3}(\mathbf{E}\cdot\mathbf{B})^2$ for $Q_{ij} = E_iB_j + B_iE_j - \frac{2}{3}\delta_{ij}\,\mathbf{E}\cdot\mathbf{B}$; the pseudoscalar $\mathbf{E}\cdot\mathbf{B} \neq 0$ marks the unreduced, dyality-active sector but does not by itself characterise the tensor. What decides whether the smooth term survives is the configuration, not the field strengths. Wherever the fields separate into independent source and radiated contributions, the Siegert reduction removes it. The toroidal cell does not separate, because its near field closes back on its own source (Section 2), so the term survives there.

The topological origin of this object lies in the defect spectrum of the bi-phasic $\mathrm{U}(1)_\Theta \times \mathrm{U}(1)_\phi$ bundle. Codimension-2 defects carry fractional mixed Chern-2 charge,

$$C_2\Big(I^{(\alpha)}\Big) = \pm\tfrac{1}{6}, \qquad \alpha = 1, \ldots, 6 \tag{22}$$

The total mixed Chern number is integer,

$$\sum_{\alpha=1}^{6} C_2\Big(I^{(\alpha)}\Big) = n, \qquad n \in \mathbb{Z} \tag{23}$$

Two separate facts need to be established: that the coherent channel is spin-2 rather than spin-0 or spin-1, and that balanced defect configurations support nonzero variance in that channel. First, however, we explain where the charge unit $\frac{1}{6}$ comes from.

4.1 Geometric origin of the fractional charge

Each defect worldsheet $I^{(\alpha)}$ is an *exact* codimension-2 submanifold of M^4: it bounds a codimension-1 Seifert hypersurface, which in our context is a level set of the dyality angle ϕ. This is precisely the class of submanifolds studied by Chern and Ishida [8], who represent each such $I^{(\alpha)}$ implicitly by a complex-valued function $\psi_\alpha \in C^\infty(M^4;\mathbb{C})$ with zero set $\psi_\alpha^{-1}(0) = I^{(\alpha)}$. The S^1-family of phase level sets $\{\psi_\alpha^{-1}(s)\}_{s\in S^1}$ is the family of Seifert surfaces bounded by $I^{(\alpha)}$ and parametrised by the dyality phase $s = e^{i\phi} \in S^1$.

Chern and Ishida show (their Theorem 5.8) that the space $\mathcal{F}_O$ of all such implicit representations is a prequantum G-bundle over the space O of defect configurations, with structure group $G = S^1 \times H^1(M^4;\mathbb{Z})$ and the Marsden–Weinstein symplectic form [33] as its curvature. The connection 1-form on this bundle is

$$\Theta_\psi(\dot\psi) = \frac{1}{2\pi|M|}\int_{M^4} \mathrm{Im}\left(\frac{\dot\psi\,\bar\psi}{|\psi|^2}\right)\mu \tag{24}$$

which measures the average swept volume of the S^1-family of Seifert surfaces under a deformation of ψ (their Remark 4.7). In our context, the relevant volume form is $\mu = F_\Theta \wedge F_\phi$ (the mixed curvature), so $|M| = \int_{M^4} F_\Theta \wedge F_\phi$ and the MW form is $\omega_{\mathrm{MW}}|_\gamma(X^u, X^v) = |M|^{-1}\int_\gamma \iota_v\iota_u(F_\Theta \wedge F_\phi)$.

The prequantizability condition, $[\omega_{\mathrm{MW}}] \in H^2(O;\mathbb{Z})$, is precisely the flux quantisation condition (18). By their Corollary 5.10, the mixed Chern charge of each defect worldsheet equals the holonomy of Θ around the corresponding loop in O, equal to the integral of ω_{MW} over the enclosed disk.

We now motivate the minimal symmetric normalisation $|C_2| = \frac{1}{6}$ from the topology of the invariant torus of Section 2. For $n = 3$ degrees of freedom the surviving torus is the compact T^3 established there, which serves as the smooth ambient manifold $M = T^3$ for the Chern–Ishida construction; Berry's adiabatic theorem [20] governs the geometric transport over its cycles.

The count of defects and the charge unit then follow from the homology of this ambient torus. The loops in O around which the holonomy of Θ is computed are labelled by $H_1(T^3;\mathbb{Z}) = \mathbb{Z}^3$: a defect is a phase singularity of the doublet section $\Psi_{\mathbf{w}}$ threading one primitive 1-cycle (38), and each primitive cycle carries the two orientations of (36), the reversed cycle giving the opposite-sign charge. The three primitive cycles of T^3 doubled by orientation are six elementary defect classes $I^{(\alpha)}$, $\alpha = 1,\ldots,6$; this is the origin of the count in (22). Their charges are the holonomies $\int_{\mathrm{disk}}\omega_{\mathrm{MW}}$ over the disks bounded by the six loops. Two facts then fix these charges, both supplied by the comoving constraint of Section 8.2. We read a codimension-2 defect as a Regge hinge of M^4 and its linking holonomy as the deficit angle concentrated there. This is a geometrical interpretation supplying a normalisation for the unit of charge. The comoving projection matches the emergent rest-frame geometry to the flat inertial metric $\eta_{\mu\nu}$ (Section 8.2), and flatness of the matched cell is the vanishing of the total deficit, which is exactly the balance condition $\sum_\alpha C_\alpha = 0$ used in Section 4. Isotropy of that same inertial match leaves the three T^3 cycle directions geometrically indistinguishable, so the six hinges carry a common $|C_2|$; the elementary classes are here taken to be the six signed basis classes of $H_1(T^3) = \mathbb{Z}^3$, a truncation that is part of the defect-sector definition, since the homology itself contains infinitely many primitive classes. The integer-total condition (23) then fixes that common value: six co-oriented elementary defects must carry one unit of mixed Chern number, $n = 1$, whence

$$\left|C_2\left(I^{(\alpha)}\right)\right| = \frac{1}{6} \tag{25}$$

with the sign set by the cycle orientation. The count 6 is $b_1(T^3) = 3$ doubled by orientation; the balance $\sum_\alpha C_\alpha = 0$ is the flatness of the inertial match; and the unit $\frac{1}{6}$ is the isotropy of that match shared with the integrality of (23).

4.2 The defect-cell map

We now state the map between the toroidal near-field cell and the codimension-two dyality defects. The construction avoids treating the dyality phase as an abstract Hamiltonian angle. The dyality angle is defined first as the source-field readout phase of the toroidal polarization doublet, where it

is fixed by the dyality transformation law (10); only after projection onto a coherent cell mode does it coincide with the phase angle of the corresponding toroidal oscillator.

Let $Q^a(x)$, $a = 1, 2, 3$, denote the retained toroidal amplitudes of the leading near-field cell. Before a time slicing is chosen, the cell is described by its directional gradients

$$V^a{}_\mu = \partial_\mu Q^a \tag{26}$$

and the conjugate momentum densities

$$\Pi^\mu{}_a = \frac{\partial \mathcal{L}_{\text{cell}}}{\partial(\partial_\mu Q^a)} \tag{27}$$

We call these polyvelocities and polymomenta only in this elementary sense: they are the velocity and momentum data of the cell before a preferred time direction is selected.[3] The ordinary Hamiltonian cell is recovered by projecting onto the comoving frame of the medium. With u^μ the medium four-velocity,

$$P_a^{(u)} = u_\mu \, \Pi^\mu{}_a \tag{28}$$

the pair $(Q^a, P_a^{(u)})$ is the phase-space coordinate of the cell on the comoving hypersurface. In the near-integrable regime it admits action–angle variables

$$(Q^a, P_a^{(u)}) \longmapsto (I_a, \theta^a), \qquad \theta^a \sim \theta^a + 2\pi \tag{29}$$

The bound source-field content of the cell is carried by the electric and magnetic toroid polarizations $\mathbf{T}^e, \mathbf{T}^m$, the space parts of the four-vectors $\Pi^e = (\varphi^e, \mathbf{T}^e)$ and $\Pi^m = (\varphi^m, \mathbf{T}^m)$ of Dubovik, Saha and Rubin [34], with the same Lorentz transformation law as the electromagnetic potentials. These are precisely the toroidal pieces lost in the far-field reduction. Because $\mathbf{T}^e$ and $\mathbf{T}^m$ are built from the electric and magnetic source distributions respectively, the dyality rotation (10), which rotates the electric–magnetic source doublet, co-rotates them as a unit,

$$\begin{pmatrix} \mathbf{T}^e \\ \mathbf{T}^m \end{pmatrix} \longrightarrow R(\phi) \begin{pmatrix} \mathbf{T}^e \\ \mathbf{T}^m \end{pmatrix} \tag{30}$$

For a primitive toroidal mode $\mathbf{w} = (w_1, w_2, w_3) \in H_1(T^3, \mathbb{Z})$ with cell profile $\mathbf{e}_\mathbf{w}$, define the projected amplitudes

$$T^e_\mathbf{w} = \int_{\text{cell}} \mathbf{T}^e \cdot \mathbf{e}_\mathbf{w} \, d^3x, \qquad T^m_\mathbf{w} = \int_{\text{cell}} \mathbf{T}^m \cdot \mathbf{e}_\mathbf{w} \, d^3x \tag{31}$$

The dyality angle of the primitive sector is the phase of the complex toroid-polarization section

$$\Psi_\mathbf{w}(x) = T^e_\mathbf{w}(x) + i \, T^m_\mathbf{w}(x), \qquad \phi_\mathbf{w}(x) = \arg \Psi_\mathbf{w}(x) \tag{32}$$

Under the dyality rotation (30) the doublet rotates as a unit, so $\Psi_\mathbf{w} \to e^{-i\phi}\Psi_\mathbf{w}$ and $\arg \Psi_\mathbf{w} \to \arg \Psi_\mathbf{w} - \phi$: the phase of $\Psi_\mathbf{w}$ shifts by exactly the dyality angle, which is the defining property (10) requires. At one value of $\phi_\mathbf{w}$ the cell reads as electric-toroidal field response, at another as magnetic-toroidal source-bound response, and at intermediate values as the mixed source-field configuration.

The relation to the Hamiltonian toroidal angle follows in the coherent single-mode sector, with no free normalisation. Take the scalar parts $\varphi^{e,m}$ fixed inside the cell and the primitive mode to be a Beltrami (curl-eigen) mode,

$$\nabla \times \mathbf{e}_\mathbf{w} = \kappa_\mathbf{w} \, \mathbf{e}_\mathbf{w} \tag{33}$$

The electric–magnetic toroid relations of [34], $\nabla \times \mathbf{T}^{e,m} = \pm\big(\frac{1}{c}\partial_t \mathbf{T}^{m,e} + \nabla\varphi^{m,e}\big)$, then reduce on this mode to the conjugate pair

$$\dot{T}^m_\mathbf{w} = -\, c\kappa_\mathbf{w} \, T^e_\mathbf{w}, \qquad \dot{T}^e_\mathbf{w} = c\kappa_\mathbf{w} \, T^m_\mathbf{w} \tag{34}$$

Two consequences follow at once. First, eliminating one amplitude gives $\ddot{T}^e_\mathbf{w} = -(c\kappa_\mathbf{w})^2 \, T^e_\mathbf{w}$, so the doublet oscillates at the free-Maxwell frequency

$$\omega_\mathbf{w} = c \, \kappa_\mathbf{w} \tag{35}$$

[3] no independent De Donder–Weyl structure is assumed here.

which is a consequence of the toroid relations, not an input. Second, at that frequency $T^e_{\mathbf{w}}$ and $T^m_{\mathbf{w}}$ are equal-amplitude quadratures, so the relative normalisation that would otherwise appear in (32) is fixed to unity. Identifying the canonical pair $Q_{\mathbf{w}} = T^e_{\mathbf{w}}$, $P_{\mathbf{w}} = \dot{Q}_{\mathbf{w}}$, the quadrature (34) reads $T^m_{\mathbf{w}} = P_{\mathbf{w}}/\omega_{\mathbf{w}}$, and on the harmonic torus $Q_{\mathbf{w}} = \sqrt{2I_{\mathbf{w}}/\omega_{\mathbf{w}}}\cos\theta_{\mathbf{w}}$, $P_{\mathbf{w}} = -\sqrt{2I_{\mathbf{w}}\omega_{\mathbf{w}}}\sin\theta_{\mathbf{w}}$, so that, [4]

$$\Psi_{\mathbf{w}} = Q_{\mathbf{w}} + i\,\frac{P_{\mathbf{w}}}{\omega_{\mathbf{w}}} = \sqrt{\frac{2I_{\mathbf{w}}}{\omega_{\mathbf{w}}}}\,e^{-i\theta_{\mathbf{w}}}, \qquad \phi_{\mathbf{w}} = -\theta_{\mathbf{w}} \pmod{2\pi} \tag{36}$$

The sign is the orientation of the primitive cycle: the reversed cycle gives $e^{+i\theta_{\mathbf{w}}}$ and the opposite-sign charge in (22). The equality of dyality phase and toroidal oscillator phase is therefore derived in the coherent projected sector, not imposed globally.

Away from the harmonic order the anharmonic stiffness $\kappa(a)$ and the coupling εH_1 shift the effective frequency from $c\kappa_{\mathbf{w}}$, so the pointwise equality $\phi_{\mathbf{w}} = -\theta_{\mathbf{w}}$ holds only at leading (Beltrami) order. The topological content is insensitive to this. A positive rescaling of either toroid amplitude is an orientation-preserving homeomorphism of the punctured plane and leaves the winding (38), the zero set, and the mixed Chern charge unchanged. The defect spectrum and the unit $|C_2| = \frac{1}{6}$ therefore depend only on the doublet structure and the primitive homology class, not on the normalisation or on the harmonic approximation; only the explicit phase–angle identification does.

The argument $\arg\Psi_{\mathbf{w}}$ is well defined only after the comoving projection. Under a boost the toroid moment is not a clean tensor: $\mathbf{T}^{e,m}$ mix with the scalar parts $\varphi^{e,m}$ and, at higher multipole order, with the mean-square-radius and quadrupole moments [34]. In a boosted frame the doublet is contaminated by these other moments and its phase is no longer the dyality angle. The polyvelocity and polymomentum data are exactly what render $\Psi_{\mathbf{w}}$ frame-clean: the comoving projection $P^{(u)}_a = u_\mu \Pi^\mu{}_a$ selects the rest-frame doublet $(\mathbf{T}^e, \mathbf{T}^m)$ in which $\arg\Psi_{\mathbf{w}}$ is the dyality coordinate. This is the same toroid-moment non-covariance invoked in Section 10 to fix the medium frame from the source side; the cell map and the comoving condition there select one and the same u^μ.

The dyality defect is the zero set of the section,

$$I^{(\alpha)} = \Psi^{-1}_{\mathbf{w},\alpha}(0) = \left\{\, T^e_{\mathbf{w},\alpha} = 0, \quad T^m_{\mathbf{w},\alpha} = 0 \,\right\} \tag{37}$$

Since $\Psi_{\mathbf{w}}$ is complex, $\Psi_{\mathbf{w}} = 0$ imposes two real conditions, so the generic zero set in four-dimensional spacetime is two-dimensional: a codimension-two worldsheet. The winding around a small loop ℓ_α linking it is

$$\nu_\alpha = \frac{1}{2\pi}\oint_{\ell_\alpha} d\phi_{\mathbf{w}} = \frac{1}{2\pi}\oint_{\ell_\alpha} d\arg\left(T^e_{\mathbf{w}} + i\,T^m_{\mathbf{w}}\right) \in \mathbb{Z} \tag{38}$$

Away from the zero set the induced relative connection is the doublet phase connection a_ϕ, the local realisation on the cell section of the bundle connection A^ϕ of Section 3,

$$a_\phi = \mathrm{Im}\left(\frac{\overline{\Psi}_{\mathbf{w}}\, d\Psi_{\mathbf{w}}}{|\Psi_{\mathbf{w}}|^2}\right) = d\phi_{\mathbf{w}} \tag{39}$$

with curvature vanishing in the smooth bulk and singular on the defect set,

$$da_\phi = 2\pi \sum_\alpha \nu_\alpha\, \delta^{(2)}\big(I^{(\alpha)}\big) \tag{40}$$

This supplies the defect-supported part of the dyality curvature in the decomposition $F^\phi = \star F^\Theta + \Delta F^\phi$ of (13), with $\Delta F^\phi = da_\phi$: the first term the locked source-field co-rotation of the smooth cell, the second the residual curvature carried by phase singularities of the toroid doublet. This is the local reason that A^ϕ is a constrained relative-phase connection rather than an independently propagating gauge field, consistent with the bulk constraint (14). The mixed Chern contribution of a defect is then

$$C_2\big(I^{(\alpha)}\big) = \frac{1}{8\pi^2}\int_{N(I^{(\alpha)})} \mathcal{F}_\Theta \wedge \mathcal{F}_\phi \tag{41}$$

[4] The harmonic form follows from the cell Lagrangian of Ref. [17], Eqs. (2.6) and (2.9): the Power–Zienau–Woolley reorganisation (2.9) differs from (2.6) by a total time derivative and a total divergence, leaving the dynamics unchanged while promoting the toroid polarization to the dynamical cell variable. The toroidal sector is then a polarization oscillator, with the $2\pi\int\mathbf{P}^2$ self-energy as the spring and the field energy $(1/8\pi)\int\mathbf{E}^2$ as the inertia, the $\mathbf{E}\cdot\mathbf{P}$ term renormalising the latter. Its action–angle form is $|\Psi_{\mathbf{w}}| = \sqrt{2I_{\mathbf{w}}/(Z_{\mathbf{w}}\omega_{\mathbf{w}})}$, with $Z_{\mathbf{w}}$ the (profile-dependent) inertia normalisation; writing the mode in canonical, unit-inertia variables ($Z_{\mathbf{w}} = 1$) gives $|\Psi_{\mathbf{w}}| = \sqrt{2I_{\mathbf{w}}/\omega_{\mathbf{w}}}$. The relation $H_{\mathbf{w}} = \omega_{\mathbf{w}} I_{\mathbf{w}}$ is the oscillator virial identity, and the inertia normalisation is a rescaling of $Q_{\mathbf{w}}$ that leaves the winding and the mixed charge unchanged.

with $N(I^{(\alpha)})$ a small tubular neighbourhood of the worldsheet, so the fractional charge (22) is a statement about the primitive flux and winding sectors of the toroidal cell.

Finally, intersecting the spacetime defect with a comoving hypersurface Σ_t gives the dynamical reading. In the coherent sector $\phi_{\mathbf{w}} = -\theta_{\mathbf{w}}$, so the same primitive label $\mathbf{w}$ labels the Hamiltonian resonance $\mathbf{w}\cdot\omega(I) = 0$ and the doublet $\Psi_{\mathbf{w}}$ whose zeros it moves. The proposed defect-cell map is that the defect is what a broken resonance leaves behind:

$$I^{(\alpha)} \cap \Sigma_t \;=\; \{\Psi_{\mathbf{w}} = 0\} \text{ released by the breaking of } \mathbf{w}\cdot\omega(I) = 0 \tag{42}$$

Three statements fix the content of this map. First, the mechanism: the dyality defect (37) carries the integer winding (38), and a winding cannot be removed by any continuous deformation; the destruction of an invariant torus is a dynamical event and has no power over a topological index. While the cycle is quasi-periodic the defect worldline is confined and the filament is compact; when the resonance breaks, the confinement goes and the defect is released, its worldline now free to wander the web. Second, the invariant and the locus are different things: the charge is the slice linking data of Section 4.3, local to each intersection with a hypersurface and indifferent to what the filament does globally, so it is unchanged by the release, while the geometry of the locus is not. Third, the loci at the moment of breaking are nested, not coincident: the defect sits at the amplitude collapse $I_{\mathbf{w}} = 0$ at the centre of the resonance, and the destroyed structure occupies finite resonant amplitude around it. What remains posited is the release itself at $n = 3$: that the chaotic layer of a broken resonance connects to the transport web and carries the defect with it, which rests on the Arnold-diffusion hypotheses stated in Section 2. The phase identification $\phi_{\mathbf{w}} = -\theta_{\mathbf{w}}$ is derived, the persistence is a theorem, and the release is the hypothesis. Appendix B illustrates the dynamical half in a computable model with the same number of degrees of freedom. Dense toroidal packing geometry may model the local adjacency graph of these cells: the packing [35] supplies the contact complex, while the section $\Psi_{\mathbf{w}}$ supplies the phase map.

4.3 Linking structure of the defect sector

We now record the gauge-invariant content of the defect sector in terms of linking numbers on a comoving slice. It sharpens the reading of the map just given: the invariant content is linking, and linking is carried by homology classes, not by position.

Mixed helicity is charge-weighted linking The singular part of the dyality curvature is delta-supported on the defects, $\Delta F^{\phi} = \mathrm{d}a_{\phi} = 2\pi\sum_{\alpha}\nu_{\alpha}\,\delta^{(2)}(I^{(\alpha)})$ with integer winding ν_{α}, as in (40). Restricting to Σ_t, $F^{\phi}|_{\Sigma} = 2\pi\sum_{\alpha}\nu_{\alpha}\,\eta_{\gamma_{\alpha}}$, where $\eta_{\gamma_{\alpha}}$ is the two-form Poincaré dual of the defect string $\gamma_{\alpha} = I^{(\alpha)} \cap \Sigma_t$. Define the mixed helicity of the slice,

$$\mathcal{H}_{\Theta\phi}(\Sigma_t) = \frac{1}{2\pi}\int_{\Sigma_t} A^{\Theta}\wedge F^{\phi} = \sum_{\alpha,\beta}\nu_{\alpha}m_{\beta}\;\mathrm{lk}\big(\gamma^{\phi}_{\alpha},\gamma^{\Theta}_{\beta}\big), \tag{43}$$

the mutual linking of ϕ-defect strings with Θ-flux lines, weighted by their windings ν_{α}, m_{β}. This is a cross linking of the two bundle sectors, the slice avatar of the bulk mixed Chern-2 invariant (17); it is not a self-linking of either field. The elementary crossing carries weight $q_{\Theta}q_{\phi} = \hbar c/2$ by Dirac quantisation, so the linking unit is locked to the same condition that fixes $|C_2| = \frac{1}{6}$ in (22) and the phase velocity in (79).

The phase-linking dictionary The identification of phase winding with geometric linking is the following statement. Let γ^{ϕ} be a closed dyality-defect loop on Σ_t around which the dyality phase winds by 2π, and let γ^{Θ} be a closed Θ-flux line. Then

$$\mathrm{lk}\big(\gamma^{\phi},\gamma^{\Theta}\big) = \frac{1}{2\pi}\oint_{\gamma^{\Theta}}\mathrm{d}\phi, \tag{44}$$

where ϕ is the phase sourced by γ^{ϕ}. The content is standard: the phase sourced by a unit-winding loop is the solid-angle function of that loop, its gradient is the Gauss linking kernel, and the circulation of $\mathrm{d}\phi$ around γ^{Θ} counts the signed crossings of γ^{Θ} through a Seifert surface of γ^{ϕ}, which is the linking number. The Hopf link gives ±1 with the sign set by orientation, matching the sign convention of (22); unlinked loops give zero. Both statements were verified numerically against the Gauss double integral in the development of this subsection.

Conservation, reconnection, and the bulk Chern number Differentiating (43) along the foliation,

$$\frac{\mathrm{d}}{\mathrm{d}t}\,\mathcal{H}_{\Theta\phi}(\Sigma_t) = \big(\text{signed reconnection density}\big), \qquad \int \mathrm{d}t\,\frac{\mathrm{d}\mathcal{H}}{\mathrm{d}t} = \frac{1}{8\pi^2}\int_{M^4} \mathcal{F}_\Theta \wedge \mathcal{F}_\phi, \tag{45}$$

up to the charge normalisation of (16): the bulk mixed Chern number is the net reconnection count of the defect strings over the evolution. Between reconnection events the slice linking is constant, so reconnection-free, foliation-preserving evolution, the projectable class selected by the medium's comoving frame in Section 8.2, conserves $\mathcal{H}_{\Theta\phi}$ exactly. The protection is topological and dynamical at once. The only route to a change of linking is a worldsheet intersection in M^4, and on the dynamical side that intersection requires two released defects to drift into contact through the web, which for $n \geq 3$ has no dividing barrier and is exponentially slow by Nekhoroshev's theorem. The suppression is a statement about the slowness of the drift, not about a barrier: nothing in the web forbids the meeting, it is only reached on exponentially long times. Topological conservation and dynamical slowness therefore protect the same object, the released zero of Section 4.2.

One caveat is required for gauge invariance. The helicity (43) is gauge invariant for closed defect loops on Σ_t. Open strings ending on the slice boundary require a choice of framing, and for them the linking number is the gauge-invariant content. All charge statements in the main text concern closed elementary classes and are unaffected.

A remark on the spaces involved is due before the kinematics are closed. The Chern–Ishida construction lives on the configuration space of codimension-two submanifolds of an ambient manifold carrying a volume form: physical defect strings in a slice Σ_t, or worldsheets in M^4, with the corresponding volume form. The cell torus T^3 supplies oscillator labels and is not itself that ambient manifold, and the mixed four-form $F_\Theta \wedge F_\phi$ is neither a three-dimensional volume form nor guaranteed non-degenerate. We use the construction for the holonomy description of defect phases and the linking data above; the charge unit of (22) rests on the integrality condition (23) together with the stated signed-basis truncation and the isotropy of the inertial match, not on the holonomy theorem itself.

This linking structure closes the kinematics. The defect distribution now carries static, topologically protected quadrupole data: the charges, positions, and orientation tensors of the elementary classes. When the web is driven by the fast carrier of Section 2, that quadrupole becomes the coherent $\ell = 2$ source. Which channels of the medium respond to it, and which propagate, is the subject of the next section.

5 The collective spectrum of the ordered medium

The channels of the medium are now developed in turn, so that the physical content of every sector is derived rather than assigned. Spin-2 selection proceeds in two steps: the local bilinear (20) removes the trace and antisymmetric components, and the transverse-traceless projection removes the longitudinal ones. What must then be established is which channels carry propagating modes, and why. The answer occupies this section: the source feeds all channels equally, the tensor channel alone carries a propagating stiffness, the defect distribution supplies a coherent quadrupole source in that channel, the vector channel is empty by a hypersurface-orthogonality theorem, the phase carries no mode, and the surviving scalar is the condensate amplitude mode.

5.1 Channel decomposition

Fix $\hat{\boldsymbol{k}}$ and the orthonormal helicity basis $e^{(\lambda)}_{ij}$, $\lambda = \pm 2, \pm 1, 0$, together with the trace direction $\delta_{ij}/\sqrt{3}$. Two exact statements follow. (i) The spin-quadrupole vertex of a triplet source has equal strength in every helicity channel: with $Q_{ij} = \frac{1}{2}(S_iS_j + S_jS_i) - \frac{1}{3}S(S{+}1)\delta_{ij}$ for $S = 1$, the state-averaged weight $W_\lambda = \frac{1}{3}\,\mathrm{Tr}\big(Q^{(\lambda)}Q^{(\lambda)\dagger}\big)$ with $Q^{(\lambda)} = e^{(\lambda)}_{ij}Q_{ij}$ equals $\frac{1}{3}$ for all five λ, as required by the Wigner–Eckart theorem. The source does not select . (ii) The identities underlying Section 8.1 hold per channel and only where claimed: the longitudinal contraction $|\hat{k}^i e^{(\lambda)}_{ij}|^2$ vanishes for $\lambda = \pm 2$ and equals $\frac{1}{2}$ for $\lambda = \pm 1$ and $\frac{2}{3}$ for $\lambda = 0$, while the trace vanishes on all $e^{(\lambda)}$ and is maximal on the trace direction. It will become evident in Section 8.2 that the anisotropic and trace blocks of the polysupermetric are therefore activated on every channel except $\lambda = \pm 2$; those channels are constraint-coupled, and their physical content is the scalaron ($\ell = 0$) and gauge ($\ell = 1$) sectors derived below. Only the $\lambda = \pm 2$ block carries the isotropic propagating stiffness. In this decomposition all channels share the microscopic spectral support, and the $J = 2$ channel alone develops the pole because it alone possesses a propagating kinetic term.

5.2 Nonzero TT variance for balanced configurations

The bilinear $\mathcal{O}_{ij}^{(2)}$ splits in (21) into a smooth co-rotation term and a defect term, and the two are not alternatives: they are the same object evaluated on the two sectors of Section 2. The co-rotation term lives on the surviving tori, at finite doublet amplitude $|\Psi_{\mathbf{w}}| \neq 0$ where $\Delta F^{\phi} = 0$. The defect term lives at the punctures $|\Psi_{\mathbf{w}}| = 0$ of (37), the broken-torus sector. The variance below is built on the punctures, and the reason the co-rotation term does not enter is that it is not static.

On the Beltrami mode the doublet $(T^e_{\mathbf{w}}, T^m_{\mathbf{w}})$ rotates in quadrature (34) at the carrier $\omega_{\mathbf{w}} = c\kappa_{\mathbf{w}}$, so the cell $\mathbf{E}$ and $\mathbf{B}$ oscillate as $\cos\omega_{\mathbf{w}}t$ and $\sin\omega_{\mathbf{w}}t$. Every component of the co-rotation tensor is then a pure second harmonic of the carrier with vanishing mean,

$$E_iB_j + B_iE_j - \tfrac{2}{3}\delta_{ij}(\mathbf{E}\cdot\mathbf{B}) \;\propto\; \sin(2\omega_{\mathbf{w}}t), \qquad \big\langle\, E_iB_j + B_iE_j - \tfrac{2}{3}\delta_{ij}(\mathbf{E}\cdot\mathbf{B}) \,\big\rangle_{\text{cycle}} = 0 \tag{46}$$

It carries no static quadrupole and so drops out of the $\mathbf{k} \to 0$ source (48)–(50); its spectral weight sits at $2\omega_{\mathbf{w}}$, away from the dyality-plasmon pole at $\omega^2 = c_h^2(\mathbf{k}^2 + m_h^2)$. The restriction of (47)–(51) to the defect sector is this spectral separation, not a truncation.

The two sectors are coupled, but only one way at leading order. The carrier is fast; the web action drifts slowly. Averaging the coupling εH_1 over the carrier phase gives no secular drift of the web action off resonance, and a nonzero drift $\propto \sin\delta$ only on a phase-locked resonance $\mathbf{w}\cdot\omega(I) = 0$—precisely the broken cycles whose defects have been released. The carrier therefore pumps the defect web through its resonances, while the web feeds back on the carrier only through the slow frequency shift $\partial\omega_{\mathbf{w}}/\partial I$, an adiabatic $\mathcal{O}(\varepsilon)$ correction with no secular part. The smooth backbone sources the defect sector; the defect drift does not source the carrier. This is what lets the scattered punctures act as one coherent $\ell = 2$ source over the EFT window: the convex web of Section 2 keeps them dynamically joined, driven by the backbone and held together by the Nekhoroshev-slow drift [26].

Consider N codimension-2 defects at positions $\{x_\alpha\}$ in the spatial slice, with mixed Chern charges $\{C_\alpha = \pm\frac{1}{6}\}$ and local orientation tensors $M^{(\alpha)}_{kl}$, where $M^{(\alpha)}_{kl}$ encodes the mixed normal structure of the α-th defect worldsheet. The Fourier-space TT operator is

$$\widetilde{\mathcal{O}}^{(2)}_{ij}(\boldsymbol{k}) = \Pi^{\mathrm{TT}}_{ijkl}(\hat{k}) \sum_\alpha C_\alpha\, e^{\,\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}_\alpha}\, M^{(\alpha)}_{kl} \tag{47}$$

Expanding the exponential in powers of $\boldsymbol{k}$:

$$\widetilde{\mathcal{O}}^{(2)}_{ij}(\boldsymbol{k}) = \Pi^{\mathrm{TT}}_{ijkl}(\hat{k}) \left[\sum_\alpha C_\alpha\, M^{(\alpha)}_{kl} + \mathrm{i}k_m \sum_\alpha C_\alpha x^m_\alpha M^{(\alpha)}_{kl} - \tfrac{1}{2}k_mk_n \sum_\alpha C_\alpha x^m_\alpha x^n_\alpha M^{(\alpha)}_{kl} + O(k^3) \right] \tag{48}$$

The constant and linear terms involve the tensor sums $\mathcal{M}_{kl} = \sum_\alpha C_\alpha M^{(\alpha)}_{kl}$ and $\mathcal{D}^m_{kl} = \sum_\alpha C_\alpha x^m_\alpha M^{(\alpha)}_{kl}$. Scalar charge balance $\sum_\alpha C_\alpha = 0$ does not by itself make $\mathcal{M}$ vanish when the orientation tensors differ. The balanced ensemble is defined here by the vanishing of the TT projections of both sums, a stated property of the ensemble rather than a consequence of charge balance, realised by the centrosymmetric, orientation-averaged arrangement assumed. The leading surviving contribution is then the quadrupole:

$$\widetilde{\mathcal{O}}^{(2)}_{ij}(\boldsymbol{k}) \;\xrightarrow{k\to 0}\; -\tfrac{1}{2}\, k_mk_n\, \Pi^{\mathrm{TT}}_{ijkl}(\hat{k})\, Q^{mn}_{kl} \;+\; O(k^3) \tag{49}$$

where

$$Q^{mn}_{kl} = \sum_\alpha C_\alpha\, x^m_\alpha x^n_\alpha\, M^{(\alpha)}_{kl} \tag{50}$$

is the quadrupole tensor of the defect distribution. The spin-2 spectral strength therefore satisfies

$$S_2(k) \;\sim\; k^4 \left|Q^{\mathrm{TT}}\right|^2 \qquad (k \to 0) \tag{51}$$

where $|Q^{\mathrm{TT}}|^2 = Q^{mn}_{kl} Q^{mn}_{k'l'} \Pi^{\mathrm{TT}}_{klk'l'}$. This is nonzero for generic positions and orientations within the balanced ensemble. It is a conditional spatial multipole scaling; a nonzero numerator does not by itself establish a pole in frequency, which is the content of Section 6.2. The analogy with classical multipole theory is precise: a charge-neutral ($\sum q_i = 0$) but spatially extended distribution has zero monopole field but generally nonzero quadrupole radiation. The Chern balance condition eliminates the spin-0 monopole channel; the quadrupole moment of the defect distribution constitutes the coherent spin-2 source. The spin-2 polarization is unique because the operator is pure spin-2 by construction, and the quadrupole moment of a balanced configuration is generically nonzero.

5.3 Vector sector: empty by hypersurface orthogonality

On the projectable foliation class of Section 8.2 the medium frame u^μ is foliation-defining: admissible evolution carries a global time function T and $u_\mu = N\,\partial_\mu T$. A frame of this form is hypersurface-orthogonal, so its twist vanishes identically. In Einstein–aether theory the spin-1 modes are carried by the twist components of the aether; in the hypersurface-orthogonal (khronometric) limit those components are absent as an identity, not as a gauge choice [36, 37]. Combined with the transversality constraint $\hat{k}^i e^{(\pm1)}_{ij} \neq 0$ established in Section 5.1, the $\ell = 1$ channel therefore contains no local propagating degree of freedom. The khronon's own remnant on the projectable class is the spatially constant lapse $N(t)$: a single global mode with no local dynamics. This is a stronger statement than the stiffness argument and is the form in which the vector sector should be understood.

5.4 Phase scalar: eliminated by constraint

The phase sector of the order parameter has the medium-frame form $\mathcal{L}_\phi = \frac{1}{2}\rho_s\,(\partial\phi - q_\phi A^\phi)^2$ with no bulk kinetic term for A^ϕ (Section 3). Varying A^ϕ yields an algebraic constraint that locks the phase gradient to the electromagnetic sector up to the defect terms of (14): the phase is pure gauge in the bulk and carries no pole. This is the Anderson–Higgs mechanism with a constrained, non-dynamical gauge field: the would-be Goldstone is removed without a massive vector appearing, consistent with the Stückelberg absorption of the dyality phase at m_Z in Section 7. There is accordingly no second-sound branch in this medium. The scalaron is the amplitude mode, not a phase branch.

5.5 Amplitude scalar: the scalaron

The surviving $\ell = 0$ excitation is the amplitude (Higgs) mode of the condensate, $\sigma \propto \delta|\Psi|$. Its gap is fixed by the same condensate structure as the tensor branch, and two readings bracket it, in the conditional register of Section 11: the pair-breaking reading places the gap at the depairing energy scale, $M_s \simeq c_h m_h/c$ in energy units, a static range of order 10^2 μm; the scalar-channel closure reading gives $M_s \simeq m_h$ with unit channel velocity and a static range of order 10^0 cm. The two differ by exactly c_h/c, and which applies depends on whether the amplitude gap is kinematic (pair-breaking) or collective (closure); resolving this is deferred. In either reading the emergent $f(R)$ coefficient $\frac{1}{2}f''(0) = 1/(6M_s^2)$, that is $f''(0) = 1/(3M_s^2)$, is fixed once the reading is, which is the sense in which the spectrum side of the $f(R)$ correspondence is derived.

5.6 Parity protection of the scalar coupling

The matter coupling of this scalaron is not the universal $f(R)$ coupling, and the reason is structural. The heavy-fermion loop generates only the mixed weight $Z_{\Theta\phi} \propto q_\Theta q_\phi$ of (72), so the amplitude mode couples through the mixed bilinear of (20), and its scalar channel is the trace of that bilinear, the metric contraction $F^{\Theta\,\mu\nu}F^{\phi}_{\mu\nu}$. On the locked background $F^\phi = \star F^\Theta$ this contraction is

$$F^{\Theta\,\mu\nu}(\star F^\Theta)_{\mu\nu} = F^{\Theta\,\mu\nu}\widetilde{F}^{\Theta}_{\mu\nu} \;\propto\; \mathbf{E}\cdot\mathbf{B} \tag{52}$$

a pseudoscalar, in agreement with the trace of the spatial form (21). Unpolarised bulk matter has $\langle \mathbf{E}\cdot\mathbf{B}\rangle = 0$ identically, so the linear coupling of the scalaron to unpolarised matter vanishes by parity.

The parity assignment here inverts the single-field intuition, and the point is worth stating explicitly because the wedge and the contraction exchange roles in the mixed sector. For one field strength, $F \wedge F$ is the parity-odd topological density and $F^{\mu\nu}F_{\mu\nu}$ the parity-even Maxwell scalar. In the mixed sector the Hodge star appears once, through the constraint, so the assignments swap: the wedge $F^\Theta \wedge F^\phi$ reduces on the locked background to $F^\Theta \wedge \star F^\Theta \propto F^{\Theta}_{\mu\nu}F^{\Theta\,\mu\nu}$, the parity-even Maxwell term already identified in (15), while the contraction acquires the star and is parity-odd as in (52). The wedge is the topological density of the mixed Chern-2 invariant and is not the operator the scalar channel couples through; substituting it here would give the parity-even combination $B^2 - E^2$ and reverse the conclusion. Three consequences follow. First, inverse-square-law tests of gravity are not directly sensitive to this scalar at leading order, so the short-range fifth-force constraints that would exclude a universally coupled scalar of this mass do not apply as stated; the residual coupling through parity-violating admixtures in matter is strongly suppressed, and its quantitative confrontation is deferred. Second, the surviving matter corrections to M_s are second order, negative in sign, and negligible in magnitude, so the scalaron mass is environment-independent for all laboratory purposes: no screening mechanism operates, and none is needed. Third, the laboratory face of the scalar branch is spin-dependent: polarised media carry

the pseudoscalar densities the mode couples to, so the scalaron mediates short-range forces between polarised spins at the range scales above, the constraint class of polarised torsion-pendulum experiments. A quantitative comparison with those limits is deferred with the other experimental material.

5.7 Spectrum summary

channel	local dof	gap	speed
$\ell = 2$ (plasmon)	2	m_h	c_h
$\ell = 1$	0	–	–
$\ell = 0$ phase	0 (constraint-eliminated)	–	–
$\ell = 0$ amplitude (scalaron)	1	M_s (two readings)	$\sim c$
khronon	0 local (global $N(t)$)	–	–

6 Effective field theory

We introduce a symmetric tensor field h_{ij}^{TT}, the *dyality plasmon*, describing the collective spin-2 oscillation of the dyality order parameter F^{ϕ} against the electromagnetic background F^{Θ}. In the language of Volovik [29, 38], this mode is analogous to zero sound in the tensor ($\ell = 2$) channel of the dyality-ordered medium. This is a collisionless collective excitation of the χ sector whose propagation is governed by the stiffness of the condensate rather than by a new microscopic light cone. The full branch content of the medium is derived in Section 5: the vector channel is empty, the phase carries no propagating mode, and the surviving scalar is the condensate amplitude mode. The dyality plasmon is the massive $\ell = 2$ branch of that spectrum.

The relation of h_{ij}^{TT} to the microscopic variables should be stated before its action is written. It is not a new entity: it is the coarse-grained collective coordinate of the defect-web quadrupole density, standing to the microscopic defect variables as an acoustic phonon field stands to atomic displacements. The action is written first; Section 6.2 then derives its resonance from the microscopic defect response rather than positing it.

The field h_{ij}^{TT} is not a graviton and defines no metric; the massive-gravity consistency question and the $f(R)$ context are treated once, in Section 10.

The projector presupposes a spatial slicing, hence a timelike unit vector u^{μ} that fixes the condensate rest frame, and we define the TT projector $\Pi_{ijkl}^{\mathrm{TT}}(u, k)$ once, for all sectors, with respect to it.

The status of u^{μ} should be fixed here, because it does three distinct jobs in what follows. First, it is a *preferred background*: the unit timelike vector, $u_{\mu}u^{\mu} = 1$, that defines the projector and the frame in which the operator content is organised. It is a background datum of the effective theory. It is not varied, and no action is supplied for it in this paper. Second, the canonical analysis requires more than a vector: the admissible foliation class of Section 8 needs u^{μ} hypersurface-orthogonal, which is an additional restriction and not implied by the first role. Third, the *velocity of matter relative to the medium* is a separate question, settled by the comoving condition $u^{\mu}_{\mathrm{matter}} = u^{\mu}_{\mathrm{medium}}$ of Section 10, which is an explicit assumption of the coupling, not a property of u^{μ} itself. The Einstein–aether construction [36] is the natural covariant completion in which u^{μ} becomes dynamical; every u-dependent structure below is written in covariant form so that such a completion is available, but it is not carried out here.

Writing the frame dependence covariantly has one economy worth stating in advance. The frame enters the quadratic theory in exactly one algebraic place, the rank-one term $(Z_t - Z_s)u^{\rho}u^{\sigma}$ of the kinetic form of (88), the covariant kinetic form derived in Section 8. Every frame-dependent consequence in this paper is that single term: the velocity ratio $c_h^2 = Z_s/Z_t$, the departure of the constitutive cone from the light cone, and the obstruction to full refoliation invariance are all proportional to $c_h^2/c^2 - 1 = (Z_s - Z_t)/Z_t$, and all vanish together at the acoustic degeneracy. The medium frame is then not a violation of a fundamental Lorentz symmetry. It is the substrate frame of an emergent theory, in which Lorentz invariance is recovered in the acoustic limit exactly as in analogue gravity [39]. Microscopic causality is carried by the parent electromagnetic and defect degrees of freedom. The local symmetric-traceless curvature bilinear is not, by itself, a conserved Fierz–Pauli source; the source coupled to h_{ij}^{TT} is its transverse-traceless projection with respect to this same u^{μ}.

6.1 Medium EFT assumptions and power counting

The theory is a medium effective theory, not a relativistic particle EFT, and we state its rules before writing operators. It is defined in the condensate rest frame fixed by the unit timelike field u^{μ}, $u_{\mu}u^{\mu} = 1$. The microscopic Lorentz symmetry is emergent and acoustic, recovered only in the long-wavelength limit, not imposed. The unbroken symmetries organising the operator content are

therefore the spatial rotations $SO(3)$ in the medium frame and the electromagnetic gauge invariance $U(1)_\Theta$; boosts relate the medium frame to others but are not a symmetry of the EFT. The field h_{ij}^{TT} is a collective response variable of the medium, transverse and traceless with respect to u^μ, carrying two polarisations. It is not a dynamical metric, so no diffeomorphism gauge symmetry is assumed and no metric self-interactions are generated. For definiteness, the effective-theory data are: ultraviolet cutoff $\Lambda = m_\chi \sim \xi^{-1}$, the defect-core scale; expansion parameters k/m_χ, ka, and m_h/k within the window (53); operator hierarchy graded by the counts stated below, with the leading operators of dimensions 4 and 5 listed explicitly; power counting anisotropic (Lifshitz), with ∂_t and ∂_ℓ graded independently and their ratio defining c_h; and symmetry principles $SO(3)$ rotations in the medium frame, $U(1)_\Theta$ gauge invariance, and the foliation-preserving class of Section 8, with boosts emergent and acoustic.

The building blocks and their mass dimensions, in units $\hbar = c = 1$, are:

object	mass dimension	role
$A_\mu^\Theta,\ A_\mu^\phi$	1	EM connection; constrained dyality connection
$F_{\mu\nu},\ \mathcal{F}_{\mu\nu}$	2	curvature; charge-normalised curvature, integer periods
$\mathcal{O}_{ij}^{(2)}$	4	TT mixed-curvature bilinear source (portal $h\,\mathcal{O}^{(2)}$: dimension 5)
h_{ij}^{TT}	1	collective tensor response field (canonical)
u^μ	0	medium frame, $u^2 = 1$; background datum, not varied
m_h	1	mode gap
c_h	0	constitutive velocity
g_{hO}	-1	portal coupling, fixed by matching (Section 7)
$g_{he},\ g_{h\nu}$	-1	matter couplings

Operators are graded by three counts: the number of fields h_{ij}^{TT}, the number of insertions of the composite source $\mathcal{O}_{ij}^{(2)}$, each carrying dimension 4 and suppressed by the heavy scale m_χ, and the numbers of temporal and spatial derivatives. Because Lorentz boosts are not a symmetry, ∂_t and ∂_ℓ are counted separately; their relative coefficient is not fixed to unity but defines the velocity c_h. This independent grading of temporal and spatial derivatives is anisotropic (Lifshitz) scaling; the gauge structure it selects is foliation-preserving rather than fully diffeomorphism-invariant, a statement made precise canonically in Section 8. The derivative expansion is controlled within the window

$$|\mathbf{k}| \ \ll\ \xi^{-1} \sim m_\chi, \qquad \omega \ \lesssim\ c_h\,\xi^{-1}, \qquad ka \ll 1 \tag{53}$$

The gap term is a relevant operator of the expansion, not a boundary of it, so the domain extends through the crossover $k \sim m_h$ to the $k = 0$ gap used in Section 11. where ξ is the defect-core coherence length and a is the cell size. The second condition, $ka \ll 1$, is the sub-wavelength requirement that justifies the toroid-dipole truncation of Section 4: the toroid quadrupole and higher members of the toroidal tower enter at $O(ka)$, so the $n = 3$ description and the operators below are leading order in ka. This requirement has a physical consequence. It confines the entire construction to the sub-wavelength near field, which is consistent with, and indeed forces, the experimental setting of Section 10: the dyality plasmon cannot be resolved by a far-field radiative probe, where $ka \gtrsim 1$ both invalidates the truncation and returns the system to the Siegert-reduced sector in which the toroidal content is absent.

At leading order the allowed operators are then few. The two-derivative quadratic terms $(\partial_t h^{\mathrm{TT}})^2$, $(\partial_\ell h^{\mathrm{TT}})^2$, and the gap $m_h^2 (h^{\mathrm{TT}})^2$ are dimension 4; they constitute the medium-frame action below. The leading coupling is the portal $h_{\mathrm{TT}}^{ij}\mathcal{O}_{ij}^{(2)}$, dimension 5, with coefficient g_{hO} of mass dimension -1 fixed by the matching of Section 7. The contact term $\mathcal{O}_{ij}^{(2)}\mathcal{O}_{ij}^{(2)}$, dimension 8, is the infrared remnant of the heavy sector with coefficient $\sim 1/m_\chi^4$. Higher spatial-derivative terms such as $(\partial_\ell^2 h^{\mathrm{TT}})^2$ are suppressed by ξ^2 and lie outside the window (53). Cubic and higher self-interactions of h^{TT} are absent at this order because h^{TT} is a response variable rather than a metric. The quadratic action that follows is therefore the leading term of this expansion, and the matching of Section 7 fixes its coefficients.

Let

$$Y_{ij}(x) = F_{i\ell}^\Theta(x) F_{j\ell}^\phi(x) + F_{i\ell}^\phi(x) F_{j\ell}^\Theta(x) \tag{54}$$

The physical spin-2 source is defined in momentum space by

$$\widetilde{\mathcal{O}}_{ij}^{(2)}(\mathbf{k}) = \Pi_{ijkl}^{\mathrm{TT}}(\mathbf{k})\,\widetilde{Y}_{kl}(\mathbf{k}) \tag{55}$$

where Π_{ijkl}^{TT} is the spatial transverse-traceless projector defined in Appendix A. It obeys

$$k^i \widetilde{\mathcal{O}}_{ij}^{(2)}(\mathbf{k}) = 0, \qquad \delta^{ij}\widetilde{\mathcal{O}}_{ij}^{(2)}(\mathbf{k}) = 0 \tag{56}$$

Thus the scalar and vector components are not removed by algebraic tracelessness alone. They are absent for the reason established in Section 5: the vector and phase channels carry no propagating mode, and the surviving amplitude scalar belongs to a separate branch. Applying Π^{TT} to the source is a kinematic decomposition: it selects the $\ell = 2$ component, while which channels carry poles follows from the response kernel. Section 5 settled that question, and the effective theory below is written for the surviving branch.

The medium-frame quadratic action for the collective mode is

$$\mathcal{L}_h^{\mathrm{TT}} = \frac{1}{2}\left[\frac{1}{c_h^2}\left(\partial_t h_{ij}^{\mathrm{TT}}\right)\left(\partial_t h_{ij}^{\mathrm{TT}}\right) - \left(\partial_\ell h_{ij}^{\mathrm{TT}}\right)\left(\partial_\ell h_{ij}^{\mathrm{TT}}\right) - m_h^2 h_{ij}^{\mathrm{TT}} h_{ij}^{\mathrm{TT}}\right] \tag{57}$$

where c_h is a response coefficient of the dyality-ordered condensate. It is not the microscopic causal speed of a fundamental Lorentz-invariant field. The condensate defines a preferred medium frame for the collective excitation, while microscopic causality is governed by the parent electromagnetic and defect degrees of freedom.

The interaction Lagrangian is

$$\mathcal{L}_{\mathrm{int}} = g_{hO}\, h^{ij}_{\mathrm{TT}} \mathcal{O}^{(2)}_{ij} + g_{he}\, h^{ij}_{\mathrm{TT}} T^{(e),\mathrm{TT}}_{ij} + g_{h\nu}\, h^{ij}_{\mathrm{TT}} T^{(\nu),\mathrm{TT}}_{ij} \tag{58}$$

where the matter bilinears are understood to be projected onto the same spatial TT sector, that is, with respect to the same u^μ that defines the medium frame. This is the comoving condition $u^\mu_{\mathrm{matter}} = u^\mu_{\mathrm{medium}}$. In a fully emergent theory, in which the matter fields descend from the same substrate, it holds automatically, as all low-energy modes share one u^μ [29]. Here the Standard Model matter is coupled in by hand through g_{he} and $g_{h\nu}$, so the condition is imposed explicitly and defines the validity window of the EFT. Configurations boosted away from u^μ lie outside that window. The corresponding propagator is

$$D_{ijmn}(\omega, \mathbf{k}) = \frac{\Pi^{\mathrm{TT}}_{ijmn}(\mathbf{k})}{\omega^2/c_h^2 - \mathbf{k}^2 - m_h^2 + \mathrm{i}\epsilon} \tag{59}$$

The portal coupling g_{hO} plays a role analogous to an inverse stiffness scale, but it couples h^{TT}_{ij} to the composite mixed-curvature source (55), not to the full stress-energy tensor. The driving operator is therefore not a fundamental field strength; it is the TT-projected response of the near-field mixed-curvature sector, and h^{TT}_{ij} is its effective collective mediator.

6.2 From the defect response to the pole

The propagating pole of h^{TT}_{ij} is not inserted by hand. This subsection obtains it within a collective-response model of the defect web, specified by a single-cell response and an assumed inter-cell kernel, once the capacitive inertia of the smooth backbone and the Josephson stiffness of the dyality coupling are supplied. The argument is the standard collective-mode resummation made explicit: a pole-free single-cell response, dressed by the inter-cell coupling, develops a collective pole whose location is the zero of the dressed denominator rather than a term added to the action.

The defect distribution of Section 4 carries, on a comoving slice, the static quadrupole tensor $M^{(\alpha)}_{kl}$ of the elementary classes, whose coherent sum is the projected quadrupole density

$$\mathcal{O}^{(2)}_{ij}(\omega, \mathbf{k}) = \Pi^{\mathrm{TT}}_{ijkl}(\mathbf{k}) \sum_\alpha C_\alpha\, M^{(\alpha)}_{kl}\, e^{\mathrm{i}\mathbf{k}\cdot\mathbf{x}_\alpha} \tag{60}$$

the Fourier transform of (47). When the web is driven by the fast carrier of Section 2 this density is the source for the $\ell = 2$ channel, and the object of interest is the retarded TT autocorrelation

$$\chi^{\mathrm{TT}}_{ij,mn}(\omega, \mathbf{k}) = \mathrm{i}\int d^4x\, e^{\mathrm{i}(\omega t - \mathbf{k}\cdot\mathbf{x})}\, \theta(t)\, \langle [\, \mathcal{O}^{(2)}_{ij}(x), \mathcal{O}^{(2)}_{mn}(0)\,] \rangle_{\mathrm{TT}} \tag{61}$$

Consider first a single cell with the inter-cell coupling switched off. The TT collective coordinate of one cell, the orientation quadrupole of its defect content, then responds as a local oscillator,

$$\chi_0(\omega) = \frac{\rho_{\mathrm{def}}}{Z_t\left[\omega^2_{\mathrm{cell}} - (\omega + \mathrm{i}0)^2\right]} \tag{62}$$

with ρ_{def} the spatial density of defect cores, Z_t the capacitive inertia, the rearrangement cost of the H_0 field energy of (1) under a time variation of the quadrupole, $Z_t \propto q_\Theta^2$ by the charge assignment

of Section 7, and $\omega_{\rm cell}$ the local pinning frequency of the isolated cell, which for a stiff cell lies above the band retained. The bare response has a local resonance but no low-energy collective pole.

The collective pole is generated by the coupling between cells. The dyality term εH_1 of (2) supplies an inter-cell alignment coupling in this channel through the phase coherence of the defect web. Assume it separable, with the small-wavevector form

$$J(\mathbf{k}) = \frac{J_0 - J_2\,\mathbf{k}^2}{\rho_{\rm def}} + O(k^4), \qquad J_2 > 0 \tag{63}$$

The signs have a simple lattice realisation. If the quadratic energy contains $\frac{1}{2}Z_t\omega_{\rm cell}^2\sum_\alpha b_\alpha^2 - \frac{1}{2}\sum_{\alpha\beta} J_{\alpha\beta}b_\alpha b_\beta$ for the cell coordinates b_α, a positive short-range alignment coupling has $J(\mathbf{k})$ maximal at $\mathbf{k}=0$; on a cubic nearest-neighbour example of spacing d, $J(\mathbf{k}) = 2J_{\rm nn}\sum_i \cos(k_i d) = 6J_{\rm nn} - J_{\rm nn}d^2k^2 + O(k^4)$. The example fixes the signs; the actual coupling of toroidal cells remains to be calculated. Its gradient stiffness is the quantity the charged-rotor assignment of Section 7 assigns, $J_2 \equiv Z_s \propto q_\phi^2$, with $q_\phi \gg q_\Theta$ so that $Z_s \gg Z_t$; the assignment is the constitutive input recorded in Table 2, and enters here and nowhere else in the dispersion.

Resumming the ladder of inter-cell couplings gives the dressed susceptibility as the geometric series

$$\chi^{\rm TT}(\omega,\mathbf{k}) = \frac{\chi_0(\omega,\mathbf{k})}{1 - J(\mathbf{k})\,\chi_0(\omega,\mathbf{k})} \tag{64}$$

which is the random-phase (Dyson) resummation of (61) with the separable kernel (63). The collective pole is the zero of the dressed denominator,

$$1 = J(\mathbf{k})\,\chi_0(\omega,\mathbf{k}) \tag{65}$$

With (62) and (63) the dressed response is $\chi^{\rm TT} = \rho_{\rm def}/\big[r + J_2\mathbf{k}^2 - Z_t(\omega+{\rm i}0)^2 + O(k^4)\big]$ with $r = Z_t\omega_{\rm cell}^2 - J_0$, which follows from the geometric series or by inverting the quadratic kernel; in an interacting theory it is an RPA-type approximation whose control would have to be justified. A stable gapped mode requires $Z_t > 0$, $r > 0$, and $J_2 > 0$, and the pole is then

$$Z_t\,\omega^2 = r + Z_s\,\mathbf{k}^2 \equiv Z_s\big(\mathbf{k}^2 + m_h^2\big), \qquad m_h^2 = \frac{r}{Z_s} \tag{66}$$

The gap is produced by softening: $0 < r \ll Z_t\omega_{\rm cell}^2$, a near-cancellation of the local pinning against the uniform alignment coupling, so that the collective gap lies far below the local resonance. This is a dynamical condition on J_0, not something removed by normalising the field; if $r<0$ the assumed background is unstable. The closure of Section 11 supplies r: the condensate amplitude fixes the residual pinning $r = Z_s m_h^2$, and the gap relation (87) is its expression. Equation (66) is therefore obtained from the resummation under the stated kernel conditions, rather than assumed. Expanding about the pole, the composite correlator carries residue $\rho_{\rm def}/Z_t$; absorbing the density into the definition of the canonically normalised collective coordinate fixes

$$\chi^{\rm TT}_{ij,mn}(\omega,\mathbf{k}) = Z_h\,\frac{\Pi^{\rm TT}_{ijmn}(\mathbf{k})}{\omega^2/c_h^2 - \mathbf{k}^2 - m_h^2 + {\rm i}\epsilon}, \qquad c_h^2 = \frac{Z_s}{Z_t}, \qquad Z_h = \frac{1}{Z_t} \tag{67}$$

where $c_h^2 = Z_s/Z_t$ reproduces (80) and $Z_h = 1/Z_t$ is the residue in that canonical normalisation, the value $Z_h \simeq 1$ used in Section 11. The overlap of any other microscopic operator with the mode introduces its own factor; Z_h as quoted refers to the canonical collective coordinate, not to an arbitrary composite. The portal coupling g_{hO} does not enter this residue; it appears only when the mode is read out, multiplying the response (67) by g_{hO}^2 in any observable that couples an external probe to $\mathcal{O}_{ij}^{(2)}$ through (58).

The pole at $\omega^2 = c_h^2(\mathbf{k}^2 + m_h^2)$ is therefore a consequence of the stiffness-to-inertia ratio of the ordered medium, obtained within the collective-response model specified above. The inter-cell coupling of toroidal cells remains to be calculated, and the control of the RPA-type resummation in an interacting theory would require separate justification. The velocity c_h is fixed by the charge assignment and Dirac quantisation of Section 7.3; the gap m_h is the $k\to 0$ pinning of the same kernel, closed self-consistently in Section 11; the residue is fixed by the inertia Z_t. The promotion of the composite source $\mathcal{O}_{ij}^{(2)}$ to the propagating field $h_{ij}^{\rm TT}$ is thus justified after the fact: $h_{ij}^{\rm TT}$ is the collective coordinate whose linear response to the defect quadrupole is exactly (67), and the quadratic action (57) is the Lagrangian whose linear response reproduces it. The microscopic content is the defect-web quadrupole; the macroscopic manifestation is the gapped spin-2 pole.

7 Constitutive inputs and a candidate ultraviolet completion

The sector below is a candidate completion, an existence argument that a dual-charged heavy state generates the required mixed-curvature structure. It is not an established ultraviolet completion, and we state at the outset what it does not supply: the gauge group and symmetry-breaking pattern, the topological condition supporting the monopole or dyon, the representation of the heavy fermion, anomaly cancellation, the ultraviolet status of A^{ϕ}_{μ}, the full one-loop matching, and the normalisation and mixing of the induced operators. What the sector does establish is narrower and is all the construction uses: the charge weights of the mixed-curvature bilinear, fixed by the double charge of χ, and the natural defect-core scale, fixed by the BPS bound. The mass estimate below likewise presupposes an explicit gauge model with the appropriate breaking topology and should be read as the natural core scale of such a model, not as a derived spectrum.

As a schematic realisation of the portal, and not as a constructed gauge model, we posit a heavy fermion χ carrying dyality charge,

$$\mathcal{L}_{\chi} = \bar{\chi}(\mathrm{i}\gamma^{\mu} D_{\mu} - m_{\chi})\,\chi \tag{68}$$

with covariant derivative

$$D_{\mu} = \partial_{\mu} + \mathrm{i}q_{\Theta} A^{\Theta}_{\mu} + \mathrm{i}q_{\phi} A^{\phi}_{\mu} \tag{69}$$

The fermion χ carries both electromagnetic charge q_{Θ} and dyality charge q_{ϕ}: it is a Witten dyon [40] bound to the core of the codimension-2 defects $I^{(\alpha)}$ of Section 4. Throughout, χ enters only through these two charges and its mass; the group-theoretic data listed at the head of this section, the gauge group and breaking pattern, the fermion representation, anomaly cancellation, and the one-loop matching, remain open, and the manipulations below should be read as the schematic consequences of a dual-charged heavy state rather than as results of a specified model. The dyality angle ϕ plays the role of the theta-parameter, and the Witten effect induces electric charge $q_{\Theta} = q_{\phi}\phi/2\pi$ on the magnetically charged defect core, relating the two charges without additional assumptions.

The connection A^{ϕ}_{μ} appearing in (69) should not be interpreted as a propagating dyality photon. It is the constrained relative-phase connection of the near-field sector, the Zwanziger dual potential of Section 3.1, and the Lagrangian (68) with both charges in one covariant derivative is well defined in that two-potential formalism [14]. Because it carries no independent Maxwell kinetic term in the bulk, integrating it out fixes the mixed-curvature order parameter rather than producing a new spin-1 excitation. The propagating field introduced in Section 6 is instead the coarse-grained TT collective mode of this order parameter.

7.1 Constrained connection and matching prescription

The matching must respect the constrained status of A^{ϕ}_{μ}. It is not integrated over as a propagating gauge field. Consistent with the bulk constraint (14), it is treated as a compact background connection that sources the defect-sector dyality current J^{ϕ} on the worldsheets $I^{(\alpha)}$. With both A^{Θ} and A^{ϕ} held as backgrounds, the heavy fermion is integrated out to give the one-loop response functional

$$W[A^{\Theta}, A^{\phi}] = -\,i \ln \det\!\big(i\gamma^{\mu} D_{\mu} - m_{\chi}\big) \tag{70}$$

with D_{μ} as in (69). The mixed second variation $\delta^2 W/\delta A^{\Theta}\,\delta A^{\phi}$ is the induced mixed-curvature response, and its spatial TT projection is the susceptibility χ_{TT} of (96). Only after this step is the TT-projected mixed response promoted to the collective field h^{TT}_{ij}. No propagating dyality photon enters at any stage: A^{ϕ} appears solely as a background source, and the single propagating collective degree of freedom is h^{TT}_{ij}. Expanded to quadratic order in the backgrounds, W is the induced vacuum polarisation. Because each external leg carries its vertex charge, it factorises as

$$W^{(2)} = \frac{1}{2}\int \frac{d^4q}{(2\pi)^4}\, q_a q_b \left(q^2 g^{\mu\nu} - q^{\mu}q^{\nu}\right) \Pi(q^2)\, \widetilde{A}^{a}_{\mu}(-q)\, \widetilde{A}^{b}_{\nu}(q), \qquad a,b \in \{\Theta,\phi\} \tag{71}$$

with a single charge-independent scalar $\Pi(q^2)$, the one-loop Dirac vacuum polarisation of (83), and the weight matrix

$$Z_{\Theta} : Z_{\Theta\phi} : Z_{\phi} = q_{\Theta}^2 : q_{\Theta}q_{\phi} : q_{\phi}^2 \tag{72}$$

Two features follow from the double charge alone, independent of the unspecified model data. The mixed weight $Z_{\Theta\phi} \propto q_{\Theta}q_{\phi}$ is non-zero only because χ carries both charges; this is what lets the loop generate the mixed bilinear $F^{\Theta}F^{\phi}$, rather than $F^{\Theta}F^{\Theta}$ or $F^{\phi}F^{\phi}$ separately, and it is the microscopic origin of $\mathcal{O}^{(2)}_{ij}$. And the tensor structure is transverse, $q_{\mu}(q^2 g^{\mu\nu} - q^{\mu}q^{\nu}) = 0$, so $W^{(2)}$ is

Lorentz invariant and gives $c_h = c$ on its own. The loop therefore fixes the charge weights and the common normalisation $\Pi(q^2)$; it does not split the spatial from the temporal stiffness. That split is supplied by the condensate frame in Section 7.3.

Because χ is a topological bound state rather than a free bulk excitation, it does not appear in the asymptotic spectrum and can be integrated out below m_χ. Its mass sets the defect core size $\xi_{\rm core} \sim \hbar c/m_\chi$; the hierarchy $m_\chi \gg m_h$ is the condition that the core is small relative to the TT coherence length, the standard criterion for a well-defined vortex in the superfluid analogy of Ref. [29, 38]. As the elementary dual-charged carrier, χ is the UV source of the toroidal topology whose infrared remnant is $\mathcal{O}^{(2)}_{ij}$.

At one loop, integrating out χ generates the curvature portal. The topological charge relation fixes the dimensionless vertex normalisation, while the Wilson coefficient retains the ordinary loop and threshold factors. Schematically,

$$C_{OO}(\mu) \sim \frac{1}{16\pi^2} \frac{\left(g^{\rm top}_{h\chi}\right)^2}{m_\chi^4} \mathcal{T}\left(\frac{\mu^2}{m_\chi^2}\right) \tag{73}$$

where $\mathcal{T}$ is a dimensionless threshold function. At lower energies this matching contributes to the local operator

$$\mathcal{L}_{\rm eff} \supset \frac{g^2_{hO}}{2m_h^2}\, \mathcal{O}^{(2)}_{ij} \mathcal{O}^{(2)}_{ij} \tag{74}$$

with the understanding that g_{hO} is the renormalised low-energy portal coupling, not merely the bare topological charge product.

7.2 Scale input and velocity prediction

The dyality condensate forms at the electroweak symmetry breaking scale m_W. The $\mathrm{U}(1)_\phi$ fibre has no independent meaning above m_W, where it is embedded in the unbroken $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$; below m_W it condenses with coherence length $\xi_{\rm core} \sim \hbar c/m_W$. The Dirac quantisation condition is evaluated at the Z-pole scale m_Z. This is the physically motivated choice: the Z boson couples to the axial-vector current $\bar\psi\gamma^\mu\gamma^5\psi$, which carries precisely the same interaction structure as the toroid dipole moment [41]. The Z pole is therefore the scale at which the axial-vector (toroidal) interaction is resonant, and $\alpha^{-1}(m_Z^2)$ is directly measured rather than inferred.

The Witten dyon mass is not an independent input but a scale set by the assumed breaking, and is therefore quoted as a benchmark of such a model rather than a derived spectrum. For a magnetic soliton in a gauge theory broken at scale m_W, the 't Hooft–Polyakov result [42, 43] gives the mass as the integrated soliton energy, yielding the BPS lower bound

$$m_\chi \;\geq\; \frac{4\pi\, m_W}{g_W^2} \;=\; \frac{m_W}{\alpha_W^{(2)}} \;\approx\; 29\, m_W \;\approx\; 2.3\ \mathrm{TeV} \tag{75}$$

where $\alpha_W^{(2)} = g_W^2/(4\pi) \approx 1/29$ is the $\mathrm{SU}(2)_L$ coupling and the non-BPS (Higgs self-coupling) range extends to ~ 10 TeV. The hierarchy is $m_h \ll m_W \ll m_\chi$, the standard EFT consistency condition.

The running fine-structure constant at the Z pole, measured at LEP to be [44]

$$\alpha^{-1}(m_Z^2) = 128.945 \pm 0.053 \tag{76}$$

The mixed Chern-2 quantisation condition (18) reduces to the Dirac quantisation condition for the χ dyon charges:

$$q_\Theta\, q_\phi = \frac{\hbar c}{2} \quad\Longrightarrow\quad q_\phi = \frac{e}{2\alpha(m_Z^2)} \approx 64.5\, e \tag{77}$$

The quantisation condition used deserves a precision. The condition (77) is the Dirac–Schwinger–Zwanziger pairing specialised to a purely electric probe: for two dyons (q_1, g_1) and (q_2, g_2) the quantised combination is the antisymmetric product $q_1 g_2 - q_2 g_1$, an integer multiple of the unit of (77), and with the probe carrying $(e, 0)$ and the χ sector its dyality charge, the antisymmetric form reduces to the Dirac half-unit form quoted. The distinction is load-bearing twice in this construction. First, the Witten assignment $q_\Theta = q_\phi\, \phi/2\pi$ used above gives the dyon a ϕ-dependent electric charge; a naive symmetric pairwise condition would forbid this, while in the antisymmetric combination the ϕ-dependent parts cancel between any two dyons of the same vacuum, so promoting ϕ to a dynamical phase is consistent precisely because the quantisation is of Dirac–Schwinger–Zwanziger form. Second, the antisymmetric condition fixes the product only up

to an integer n; equation (77) takes the minimal $n = 1$, and a completion whose minimal consistent unit is larger rescales q_ϕ, and hence c_h, by n, so the integer itself is fixed by the same dispersion measurement that fixes c_h (Section 11).

The corresponding topological vertex normalisation is

$$g_{h\chi}^{\rm top} = \frac{q_\Theta q_\phi}{\hbar c} = \frac{1}{2} \tag{78}$$

The phase velocity of the dyality plasmon is the stiffness-to-inertia ratio of the tensor channel, $c_h/c = q_\phi/q_\Theta$. Section 7.3 derives this and shows that, at the Z pole, Dirac quantisation fixes

$$c_h = \frac{c}{2\alpha(m_Z^2)} \approx 64.5\, c \tag{79}$$

The vertex normalisation (78) should not be read as the full low-energy Wilson coefficient of the spin-2 portal. The latter is obtained only after matching the heavy χ sector onto the infrared EFT and therefore carries the usual loop, threshold, and operator-normalisation factors. Thus $g_{h\chi}^{\rm top} = 1/2$ is a topological normalisation fixed by the Dirac condition at the Z-pole scale m_Z, while the renormalised portal coupling g_{hO} is an EFT parameter determined by low-energy matching. The phase velocity c_h is neither. It is a constitutive ratio of the medium, fixed in value by the same Dirac condition only once the stiffness normalisation of Section 7.3 is supplied. The dyality angle ϕ parameterises a neutral SO(2) rotation on (E, B) and (ρ_e, ρ_m) simultaneously; within the unbroken $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ it belongs to the neutral (A_μ^3, B_μ) sector. Electroweak symmetry breaking assigns mass $m_Z = m_W/\cos\theta_W$ to the neutral broken generator: ϕ is its longitudinal degree of freedom, absorbed at scale m_Z by the Stückelberg mechanism. The Dirac condition is therefore evaluated at m_Z exactly. This identification is corroborated by Dubovik and Kuznetsov [41], who show that the toroid dipole moment of the Majorana neutrino is proportional to $G_F = \sqrt{2}\, g^2/(8m_W^2)$; by the Sirlin relation [45], $G_F \propto \alpha(m_Z^2)/m_Z^2$, so G_F is naturally an m_Z-scale quantity. Dubovik and Kuznetsov further show that the Majorana TDM equals twice the Dirac TDM, because both particle and antiparticle currents contribute to the electromagnetic vertex; fundamentally a statement about the unreduced near-field.

The completion is schematic, but the gap it leaves is bounded rather than open. Three of the ingredients a full model must supply are already constrained by consistency with the rest of the construction, independent of which gauge model realises them. The dyality charge is fixed to $q_\phi = e/2\alpha(m_Z^2)$ by Dirac quantisation (77), so any admissible completion must reproduce this value. The heavy state must be dual-charged, since a singly charged fermion gives a vanishing mixed weight $Z_{\Theta\phi}$ and no portal at all; this excludes the Standard Model content and requires a genuine monopole or dyon sector. And the core scale is bounded below by the BPS relation (75), so the hierarchy $m_h \ll m_W \ll m_\chi$ is a consequence, not an assumption. What remains genuinely open is the model that meets these three constraints simultaneously, together with its anomaly cancellation and one-loop matching. We therefore present the sector as a schematic completion that the construction constrains but does not determine, and treat the charge weights (72) as motivated by the double charge rather than derived from a specified spectrum.

7.3 Derivation of the phase velocity

Write the medium-frame action (57) with the two stiffnesses left general, in the condensate frame u^μ,

$$\mathcal{L}_h^{\rm TT} = \tfrac{1}{2}\Big[\, Z_t\,(\partial_t h_{ij}^{\rm TT})^2 - Z_s\,(\partial_\ell h_{ij}^{\rm TT})^2 - M^2\,(h_{ij}^{\rm TT})^2\,\Big], \qquad \frac{c_h^2}{c^2} = \frac{Z_s}{Z_t} \tag{80}$$

so the velocity is fixed entirely by the ratio of the spatial rigidity to the temporal susceptibility. The pair (Z_t, Z_s) will reappear in Section 8 as the temporal and spatial eigenvalues of a covariant kinetic quadratic form, the medium polysupermetric; we use the name from here on so that the canonical development introduces no new object. Equation (57) is the case $Z_t = 1/c_h^2$, $Z_s = 1$.

The velocity is not radiative in origin. As the transverse structure of $W^{(2)}$ in (71) shows, a loop that respects the Ward identity produces only $q^2 g_{\mu\nu} - q_\mu q_\nu$, that is $\omega^2 - \mathbf{k}^2$, and the dispersion is then $\omega^2 = \mathbf{k}^2 + M^2/Z$ with $c_h = c$. The loop by itself cannot generate $c_h \neq c$. The deviation comes only from the preferred-frame term supplied by the condensate, which adds a contribution $\propto (q\cdot u)^2$ to (80) and splits Z_t from Z_s. In the language of Section 8, this rank-one $(q\cdot u)^2$ piece is the sole anisotropic part of the medium polysupermetric, and Z_s, Z_t are its spatial and temporal eigenvalues. Because the spin-2 channel admits the single tensor structure $\Pi^{\rm TT}$, any geometric loop factor is common to the time and space terms and cancels in the ratio Z_s/Z_t.

The matching of Section 7.1 already fixed the charge weights (72) and the common normalisation $\Pi(q^2)$. What it did not fix is which weight multiplies the spatial term and which the temporal one, because the loop is frame-symmetric. The condensate supplies this; the assignment is selected by which charge the order-parameter phase carries. The spatial term in (80) is the energy cost of a spatial variation of the order parameter. The propagating order parameter is the dyality curvature F^{ϕ} oscillating against the electromagnetic background F^{Θ}, and at the cell level its phase is the dyality angle $\phi_{\mathbf{w}} = \arg \Psi_{\mathbf{w}}$. By the dyality transformation law (10), (30) this phase is charged under q_ϕ and under q_ϕ alone: a dyality rotation shifts $\arg \Psi_{\mathbf{w}}$ by ϕ, whereas a passive electromagnetic duality with the sources held fixed leaves it untouched (the passive relabelling of Section 3). The stiffness Z_s of a spatial gradient of this phase is therefore necessarily $\propto q_\phi^2$; it cannot be assigned q_Θ^2 without making an electromagnetic rotation shift the dyality angle, in direct contradiction with (10). The temporal term is the inertia of the doublet amplitude. On the Beltrami mode the amplitude oscillation is the free-Maxwell exchange (34) at $\omega_{\mathbf{w}} = c\kappa_{\mathbf{w}}$, governed by the field energy $(1/8\pi)\int E^2$ of H_0, and is therefore $\propto q_\Theta^2$. The orientation of the split is fixed by which charge generates the phase rotation. Of the three components of the assignment defined below, the first two are properties of the bundle and the third is the constitutive input:

$$Z_s \propto q_\phi^2, \qquad Z_t \propto q_\Theta^2 \tag{81}$$

Within this identification the reversed assignment is not available: it would charge the phase under q_Θ, which (10) forbids. The constraint fixes which charge labels which sector; it does not by itself fix the magnitudes of the kinetic coefficients, which is the constitutive input recorded below. The structure is that of a charged quantum-rotor (Josephson) array, in which the capacitive charging energy that resists a change of amplitude is the temporal inertia, $C \leftrightarrow Z_t \leftrightarrow q_\Theta^2$, the phase-coupling energy between cells is the spatial stiffness, $E_J \leftrightarrow Z_s \leftrightarrow q_\phi^2$, and the collective phase mode has speed $c_h^2 = E_J/C$ up to the cell spacing. The strong dyality charge makes the medium stiff in space and the weak electromagnetic coupling makes it light in inertia, so the mode is fast; the reversed array would put the small charge on the stiffness and give $c_h = 2\alpha(m_Z^2)\, c \approx c/64.5$, a slow mode, excluded by the same fact that fixes the rotor. This is the stiffness-side reading of the Hamiltonian split of Section 2: the temporal inertia is carried by H_0—the free-Maxwell part, the q_Θ field energy, the locked $\mathbf{E}\cdot\mathbf{B}$ term of (21)—and the spatial stiffness by εH_1—the dyality coupling (2), the q_ϕ sector, the defect web. The split is not an independent input but the same $H_0/\varepsilon H_1$ decomposition viewed through the polysupermetric. The common $\Pi(q^2)$ cancels in the ratio, leaving

$$\frac{c_h^2}{c^2} = \frac{Z_s}{Z_t} = \left(\frac{q_\phi}{q_\Theta}\right)^2 = \frac{1}{4\alpha^2} \tag{82}$$

using Dirac quantisation $q_\Theta q_\phi = \hbar c/2$ with $q_\Theta = e$. At the Z pole this is (79). The value $64.5\,c$ is therefore not a consequence of topology alone, nor of the loop alone: the loop fixes the weights, the charge of the dyality phase fixes their assignment (81), and what remains genuinely conditional is the running prescription below, that q_ϕ is fixed at m_Z while $q_\Theta = e(q^2)$ runs to the probe scale. The picture is physical: the medium is stiff in space, through the strong dyality charge, and light in inertia, through the weak electromagnetic coupling, hence fast.

By the *charged-rotor constitutive assignment* we mean, precisely, three statements of three different statuses.

(i) *Derived.* Each near-field cell is identified with a charged quantum rotor whose internal compact $U(1)$ is the dyality phase; the cell map established $\phi_{\mathbf{w}} = -\theta_{\mathbf{w}}$ in the coherent Beltrami sector (36).

(ii) *Topological.* The rotor's conjugate variable is integer-quantised because the phase is compact, and that integrality is the mixed-Chern period condition (18); labelling the internal $U(1)$ as dyality bestows the topological constraint, and the Dirac condition $q_\Theta q_\phi = \hbar c/2$ emerges rather than being inserted.

(iii) *Constitutive input.* The rotor's capacitive (inertial) energy is assigned to the q_Θ sector, the field energy of H_0, and its Josephson (stiffness) energy to the q_ϕ sector, the dyality coupling εH_1, with no loop, condensate, or mixing contamination of the split.

Components (i) and (ii) are properties of the bundle; component (iii) alone is the constitutive input, and it is what Table 2 records as the model hypothesis. The one-loop matching fixes the charge weights (72); it does not fix the kinetic coefficients of the medium action, which may in

general receive contributions from loop integrals, field normalisation, condensate density, cell geometry, threshold corrections, and operator mixing. The assignment (81) is the statement that these corrections respect the bundle split of stiffness and inertia; it is motivated by the transformation law (10), and it is not derived from an explicit response calculation. The value $64.5\,c$ is therefore the benchmark that follows from this assignment together with Dirac quantisation, not a parameter-free prediction.

The residual normalisation is the one-loop vacuum polarisation, which we take exactly rather than as a scalar estimate. For a Dirac fermion of charge q_f and mass m the transverse Feynman weight is $\int_0^1 x(1-x)\,dx = \frac{1}{6}$ and the finite-part constant is $\int_0^1 x(1-x)\ln[x(1-x)]\,dx = -\frac{5}{18}$, giving

$$\hat{\Pi}(q^2) \xrightarrow{|q^2|\gg m^2} \frac{\alpha}{3\pi}\ln\frac{|q^2|}{m^2} - \frac{5\alpha}{9\pi} \tag{83}$$

The inertia Z_t couples with $q_\Theta = e(q^2)$, which runs to the probe scale, while q_ϕ is fixed once at m_Z by Dirac quantisation. The running is that of the physical electromagnetic coupling, which sums (83) over the full charged spectrum and is therefore larger than a single-fermion estimate. The velocity then runs with the probe scale as

$$\frac{c_h(q^2)}{c} = \frac{1}{2\sqrt{\alpha(m_Z^2)\,\alpha(q^2)}} \tag{84}$$

At $q^2 = m_Z^2$ this returns $c_h = c/2\alpha(m_Z^2) \approx 64.5\,c$. In the infrared, where $\alpha \to \alpha(0)$ with $\alpha^{-1}(0) = 137.036$, it rises to $c_h/c = 1/[2\sqrt{\alpha(m_Z^2)\,\alpha(0)}] \approx 66.5$. The mesoscopic near-field window lies below m_e, where α is frozen at the Thomson value, so within the window c_h is flat at $\approx 66.5\,c$ to better than 0.1%; the full excursion from the near field to the Z pole is about 3%.

Equation (84) rests on the assignment (81), in which the inertia runs to the probe scale while q_ϕ is fixed at m_Z. This is the standard expectation that a coupling is evaluated at the scale of the process, and it is the modelling input of this section. If instead the whole ratio is treated as a single m_Z matching quantity, $c_h = 64.5\,c$ at every scale and there is no dispersion. The two cases differ observably: the presence or absence of the dispersion (84) is itself a measurement.

It is useful to define the effective spin-2 density and local TT amplitude:

$$\rho_{\rm TT}(\mathbf{x}) \equiv \mathcal{O}^{(2)}_{ij}(\mathbf{x})\,\mathcal{O}^{(2)}_{ij}(\mathbf{x}) \tag{85}$$

$$g_{\rm TT}(\mathbf{x}) \equiv \sqrt{\rho_{\rm TT}(\mathbf{x})} = \left|\mathcal{O}^{(2)}_{ij}(\mathbf{x})\right| \tag{86}$$

$g_{\rm TT}$ is the local amplitude of the spin-2 condensate; its vacuum expectation value $g_{\rm TT,0} \equiv \langle g_{\rm TT}\rangle$ sets the condensate stiffness. Equation (74) is the infrared remnant of the ultraviolet box diagrams shown in Fig. 1(a–b).

The mediator mass is set by the UV scale and the condensate amplitude,

$$m_h^2 \;\sim\; \frac{\Lambda^4}{g_{\rm TT,0}}, \qquad \Lambda = m_\chi \;\geq\; 2.3\ \text{TeV} \tag{87}$$

where $g_{\rm TT,0} = \langle\sqrt{\rho_{\rm TT}}\rangle$ is the vacuum TT amplitude, of mass dimension two, so both sides carry dimension two. This is a gap-equation relation: a larger condensate amplitude softens the mode. The hierarchy $m_h \ll m_W \ll m_\chi$ is ensured for $g_{\rm TT,0} \gg \Lambda^2$.

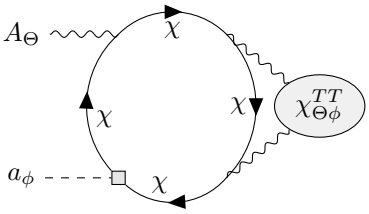
8 Canonical structure and admissible foliations

The effective theory of Section 6 was written as a medium-frame action with two independent stiffnesses and an independently graded time derivative. We now show that this is the transverse-traceless (TT) projection of a covariant field-strength Hamiltonian structure of the De Donder–Weyl type, and use the associated hypersurface-deformation algebra to fix the foliation class on which the phase velocity $c_h \neq c$ is consistent with covariant evolution. This places the construction within the covariant Hamiltonian quantisation recently developed for the teleparallel equivalents to general relativity [46, 47, 48], in which gauge field strengths play the role of generalised velocity fields and the Legendre transform of a quadratic action is regular.

8.1 The medium polysupermetric

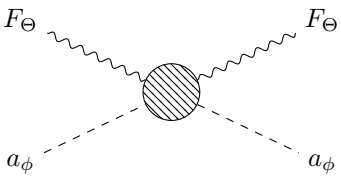
In a field-strength Hamiltonian formulation the kinetic term of a field with velocity field strength F is governed by a polysupermetric $\mathcal{M}$, with $\mathcal{L} = \frac{1}{2}\,\mathcal{M}\,FF - V$, and the phase velocity of a mode is

(a) UV mixed response



Heavy dyality-charged fermions generate the mixed-curvature susceptibility. The dashed a_ϕ insertion is the constrained background data from the local cell, not a propagating dyality photon.

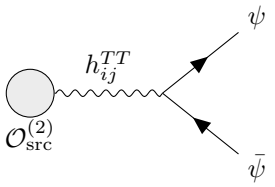
(b) Induced Wilson response



$$\mathcal{L}_{\text{eff}} \supset \frac{g_{hO}^2}{2m_h^2}\,\mathcal{O}_{ij}^{(2)}\mathcal{O}_{ij}^{(2)}$$

Below the heavy scale, the mixed response reduces to the local Wilson operator $\mathcal{O}_{ij}^{(2)}\mathcal{O}_{ij}^{(2)}$. The a_ϕ legs denote constrained dyality insertions.

(c) Collective exchange



Only the collective tensor h_{ij}^{TT} is drawn as a propagating mediator. It transmits the mixed-curvature source sector to matter bilinears.

Figure 1. Hierarchy of the mixed-curvature spin-2 sector. Panel (a): representative UV response diagram. The heavy dyality-charged fermion χ generates a mixed susceptibility $\chi^{TT}_{\Theta\phi}$ in the presence of an electromagnetic field A_Θ and a constrained dyality insertion a_ϕ. Panel (b): induced infrared Wilson operator after integrating out the heavy response. Panel (c): matter-sector exchange mediated by the propagating collective field h_{ij}^{TT}. Dashed a_ϕ lines denote constrained background or defect-supported dyality insertions.

the corresponding eigenvalue ratio of $\mathcal{M}$ [46]. For the collective field the velocity is $\partial_\mu h_{ij}^{\rm TT}$, and the medium-frame action (57), with the general stiffnesses of (80), is reproduced by

$$\mathcal{L}_h^{\rm TT} = \tfrac{1}{2}\,\mathfrak{g}^{\rho\sigma}\big(\partial_\rho h_{ij}^{\rm TT}\big)\big(\partial_\sigma h_{ij}^{\rm TT}\big) - \tfrac{1}{2}m_h^2\, h_{ij}^{\rm TT} h_{ij}^{\rm TT}, \qquad \mathfrak{g}^{\rho\sigma} = Z_s\, g^{\rho\sigma} + (Z_t - Z_s)\, u^\rho u^\sigma. \tag{88}$$

In the rest frame $\mathfrak{g}^{\rho\sigma} = \mathrm{diag}(Z_t, -Z_s, -Z_s, -Z_s)$: a single timelike eigenvalue Z_t along u^μ and a triply degenerate spacelike eigenvalue $-Z_s$. The contraction returns $Z_t(\partial_t h)^2 - Z_s(\partial_\ell h)^2$, and the dispersion gives

$$c_h^2 = \frac{Z_s}{Z_t} = \frac{\text{spatial eigenvalue of } \mathfrak{g}}{\text{temporal eigenvalue of } \mathfrak{g}}. \tag{89}$$

The constitutive split (81) is the statement that these two eigenvalues are carried by the dyality and electromagnetic charges, $Z_s \propto q_\phi^2$ and $Z_t \propto q_\Theta^2$, so that c_h is a property of the polysupermetric rather than of any microscopic light cone.

The rank-one term $(Z_t - Z_s)u^\rho u^\sigma$ is the only anisotropic structure in $\mathfrak{g}$, and it is absent from any Lorentz-invariant, general-relativity-equivalent theory. This is a property of the TT projector, not an assumption. For the symmetric teleparallel equivalent of general relativity, whose action is quadratic in the nonmetricity $Q_{\rho\mu\nu} = \nabla_\rho g_{\mu\nu}$, evaluation on a TT mode removes every anisotropic term: the permutation contraction $\frac{1}{2}Q_{\rho\mu\nu}Q^{\mu\nu\rho} \propto |k^\mu h_{\mu\nu}^{\rm TT}|^2$ vanishes by transversality, and both nonmetricity traces vanish on a traceless mode, leaving only the isotropic $-\frac{1}{4}Q_{\rho\mu\nu}Q^{\rho\mu\nu}$ block. The TT projection of that polysupermetric is therefore $Z_s\, g^{\rho\sigma}$ with $Z_s = Z_t$ and $c_h = c$ identically. The dyality-ordered medium differs from this general-relativity-equivalent case by the single rank-one condensate term in (88); the $f(R)$-type long-wavelength context of Section 10 is its isotropic limit.

8.2 Hypersurface-deformation algebra and the admissible foliations

Three facts frame the scope of this subsection: the field $h_{ij}^{\rm TT}$ is not a metric, the theory does not possess full diffeomorphism invariance, and the TT projector is spatially nonlocal. The constraint structure of the medium and matter sectors is not provided here; u^μ contributes none, being a background datum rather than a varied field (Section 6.1). A fully covariant treatment in which u^μ is dynamical would add its own constraints, and is not attempted. What follows is therefore a consistency check of Tomonaga–Schwinger evolution on a fixed foliation class, not the construction of a fully covariant constrained system. The computation identifies the unique cone-sensitive obstruction and the foliation class on which it vanishes; full refoliation invariance is claimed only in the acoustic limit, where the obstruction vanishes identically.

A covariant Hamiltonian quantisation evolves a hypersurface wavefunctional through a Tomonaga–Schwinger equation generated by hypersurface deformations [46, 49], without a preferred time coordinate. Its consistency is controlled by the closure of the hypersurface-deformation (Dirac–Teitelboim) algebra, whose only cone-sensitive bracket is the

normal–normal one. For a sector propagating on the cone of speed c_h, with normal generator $H_\perp$ smeared by lapses N, M, a direct computation on a spatial slice Σ of induced metric h_{ij} gives

$$\{H_\perp[N], H_\perp[M]\} = c_h^2 \int_\Sigma d^3x \, \left(N\,\partial_i M - M\,\partial_i N\right) h^{ij}\, H_j^{\rm coll}, \qquad H_j^{\rm coll} = \hat{\Pi}^{\alpha\beta}\partial_j h_{\alpha\beta}, \tag{90}$$

with the mass and potential terms cancelling identically and the slice-curvature contributions ($\partial_i h_{jk}$) cancelling in the antisymmetric combination, so that the geometry of the leaf generates no projectability-surviving term. The structure function is $c_h^2 h^{ij}$: the speed enters through the spatial gradient term.

The matter and gravitational sectors propagate on the ordinary cone, entering (90) with $c_h \to c$. Coupling both to a single foliation, the total bracket closes on the tangential generator only if the structure function is common; the deviation from single-diffeomorphism closure is the single term

$$\mathcal{A}[N, M] = \left(\tfrac{c_h^2}{c^2} - 1\right) \int_\Sigma d^3x \, \left(N\,\partial_i M - M\,\partial_i N\right) h^{ij}\, H_j^{\rm coll}. \tag{91}$$

The coefficient is the bare cone ratio minus one, with no loop or solid-angle factor. On a u^μ-adapted foliation the antisymmetrised lapse gradient is the foliation tilt, $N\partial_i M - M\partial_i N = NM\,(a_i^{(M)} - a_i^{(N)})$ with $a_i = D_i \ln N$ a spatial component of ∇u, so $\mathcal{A} \propto (c_h^2/c^2 - 1)\,\langle H_i^{\rm coll}, (\nabla u)^i \rangle$.

Two limits annihilate $\mathcal{A}$ exactly. In the acoustic limit $c_h \to c$ the coefficient vanishes and the full Dirac–Teitelboim algebra closes: emergent Lorentz invariance coincides with emergent full refoliation symmetry. Alternatively, for leaf-homogeneous lapse, $\partial_i N = 0$, the smearing factor vanishes for any c_h; this is the projectability condition, and the admissible deformations form the foliation-preserving subalgebra FDiff = Diff(Σ) $\ltimes$ Reparam(t) of Hořava–Lifshitz type [37]. Unlike Hořava gravity, however, the present field carries no scalar mode and defines no metric, so the strong-coupling and projectability difficulties of the Hořava scalar sector do not arise here: $h_{ij}^{\rm TT}$ supplies only the two tensor polarisations, and FDiff enters solely as the gauge algebra of the admissible foliations. The anisotropic grading of Section 6.1, in which ∂_t and ∂_ℓ scale independently with ratio c_h, is exactly the regime for which this is the natural gauge algebra. The comoving condition $u^\mu_{\rm matter} = u^\mu_{\rm medium}$ of Section 6.1, imposed from the source side by the non-covariance of the toroidal moment (Section 10), is the selection of one zero-tilt foliation shared by the collective and matter sectors, on which $\mathcal{A}$ is never activated. The same mechanism disposes of the aether's own propagating sectors: a mode of speed s_a contributes a term $\propto (s_a^2 - c^2)(N\partial_i M - M\partial_i N)$, again purely tilt-sourced, vanishing on the projectable class.

We state the canonical result in the conditional spirit of (C4) and (P1)–(P2). The structural obstruction to covariant evolution at $c_h \neq c$ is the single term (91); it is large in coefficient but exactly removable, either acoustically or on the projectable foliation class selected by the medium frame. This is distinct from the ultraviolet (Schwinger) anomaly of the regularised generator, an $\mathcal{O}(\hbar^2)$ coincidence-limit question independent of c_h and shared with all hypersurface-deformation quantisations [46], which we do not address. Within the projectable class the dyality plasmon admits a covariant, non-frozen, many-fingered Tomonaga–Schwinger evolution, with full refoliation invariance emerging in the acoustic limit.

8.3 Regge matching and the defect charge unit

The comoving frame that selects the zero-tilt foliation is also the frame in which the charge normalisation of Section 4 was derived, and this subsection discharges that hypothesis. In the acoustic limit of this section the comoving projection matches the rest-frame cell to the flat inertial metric $\eta_{\mu\nu}$. Flatness of the matched cell is the vanishing of the total Regge deficit over the hinges of a closed cell, which is the balance condition $\sum_\alpha C_\alpha = 0$; isotropy of the same match leaves the three cycle directions of the surviving T^3 equivalent, fixing a common $|C_2|$; and integrality of (23) then distributes one unit over six co-oriented hinges, $|C_2| = \frac{1}{6}$, as derived in Section 4. The reading of the mixed charge as a deficit angle is an interpretation adopted in this limit and conditional on it; the balance and isotropy conditions it supplies, together with integrality, are what fix $|C_2|$. The anisotropic departure from the match is the same rank-one condensate term $(Z_t - Z_s)u^\rho u^\sigma$ of the polysupermetric that lifts c_h from c: the charge symmetry and the velocity anomaly are controlled by one structure, read in two limits.

9 Causality, front velocity, and the infrared limit

The foliation analysis of Section 8 fixed where the effective evolution is defined; this section establishes what it can and cannot transmit: the retarded propagator, the velocity hierarchy, the

two-cone separation of the constitutive cone from the signal cone, a microscopic causal analogue, and the infrared limit.

9.1 Retarded propagator and velocity hierarchy

The retarded propagator of the medium action (57) is

$$D^R_{ijmn}(\omega,\boldsymbol{k}) = \frac{\Pi^{\mathrm{TT}}_{ijmn}(\boldsymbol{k})}{(\omega+\mathrm{i}0^+)^2/c_h^2 - \boldsymbol{k}^2 - m_h^2}, \tag{92}$$

analytic in the upper half ω plane, so the susceptibility (96) obeys the Kramers–Kronig relations by construction. On the dispersion $\omega^2 = c_h^2(\boldsymbol{k}^2 + m_h^2)$ the phase and group velocities are

$$v_p(k) = c_h\sqrt{1+m_h^2/k^2} \;\geq\; c_h, \qquad v_g(k) = \frac{c_h\,k}{\sqrt{k^2+m_h^2}} \;<\; c_h, \tag{93}$$

with $v_g \to c_h$ from below at large k. The characteristic speed of the truncated wave operator, and hence its front velocity, is c_h. We state this plainly rather than obscure it: within the effective description, disturbances of h^{TT}_{ij} propagate on the c_h cone.

9.2 Two cones: the constitutive cone and the signal cone

The c_h cone is not a causal cone of the theory, and two separate arguments offer support, one from the dispersion relation and one from the operators.

Take the dispersion relation first. The group velocity approaches c_h already at $k \sim$ few $\times\, m_h$, well inside the window (53), so the ultraviolet ceiling cannot be invoked here. The front is a different matter. It is fixed by the high-frequency response, which a truncated effective equation does not determine, just as the characteristic cone of an acoustic metric is not the light cone of the fluid's electrodynamics. The operator half is where care is needed. The composite h^{TT}_{ij} is not a local operator: the TT projector carries inverse Laplacians, and Chu and Liu [**?**] have shown that even in linearised general relativity the TT variable receives support from outside the observer's past null cone and cannot be a standalone observable; there the gauge-invariant scalars and vector cancel the acausal part inside every observable, which is the linearised Riemann tensor. We therefore claim microcausality only where it is a theorem: for the local parent operators F^{Θ}, a_ϕ, and the defect variables, and hence for every local observable built from them, including the receiver-port fields of Section 11. No microcausality claim is made for the projected composite itself.

This yields a two-cone structure. The signal cone of every local observable is luminal, governed by the parent sector. The c_h cone is the characteristic cone of the frame-anchored, spatially nonlocal constitutive kernel of the TT response, the analogue of a Pippard kernel in a superconductor: its superluminal support is response-level structure, and no front rides it. The departure $c_h^2/c^2 - 1 = (Z_s - Z_t)/Z_t$ measures exactly the frame-anchored content of the kinetic form, and at the acoustic degeneracy $Z_s = Z_t$ the constitutive cone collapses onto the signal cone and the covariant local case is recovered. In general relativity the nonlocal TT content is gauge bookkeeping, cancelled in observables; here the frame is physical, so the medium selects a preferred TT projection and resolves the TT/tt ambiguity of [**?**], at the price that the nonlocality becomes physical. The price is affordable because it is constitutive: it resides in the response kernel, where a medium carries nonlocality as a matter of course, and never in a signal front. The same frame-anchored nonlocality is the $(n\cdot\partial)^{-1}$ structure of the two-potential formalism of Section 3.1; the medium supplies the frame both require. Because no foliation enters the parent-sector statement, it holds on every slicing, tilted or not. This is the one element of this section that does not need the projectable class.

9.3 A microscopic causal analogue

The analogue in this subsection concerns the projectable, zero-tilt foliation class of Section 8.2, on which the effective theory and its retarded propagator are defined; there the medium is fixed, homogeneous, and passive, and the analogy is exact in kind. The tilted case is different in character and is treated in the next subsection. The coexistence of superluminal constitutive velocities with a luminal operator front is the classic Sommerfeld–Brillouin result [**?**] for the Lorentz oscillator medium: near an anomalous-dispersion band both v_p and v_g exceed c over a finite frequency interval, while the forerunner front propagates at exactly c, as required by the upper-half-plane analyticity of the refractive index. The demonstration is elementary enough to record. For the single-resonance Lorentz medium, $n^2(\omega) = 1 + \omega_p^2/(\omega_0^2 - \omega^2 - \mathrm{i}\gamma\omega)$, the phase and group velocities exceed c throughout the anomalous band above ω_0, yet the retarded response is

analytic in the upper half plane and $n(\omega) \to 1$ as $|\omega| \to \infty$, so the Green function vanishes identically outside the light cone: the front travels at exactly c whatever the band velocities do. The constitutive velocities and the front are controlled by different parts of the response, the anomalous band and the high-frequency limit respectively. The present construction is of the same type. The candidate completion of Section 7 is a local, causal Lagrangian; the response functional (70) built from it is retarded and analytic; and the constitutive coefficient c_h extracted in the window (53) plays the role of the anomalous-band group velocity. The window terminates before the front-velocity regime is probed, so no contradiction can arise within the domain of validity, and outside it the parent front is luminal by the operator argument above.

9.4 The tilted case

Away from the projectable class the situation must be stated separately, because the passive-medium analogy above no longer applies. On tilted foliations the hypersurface-deformation algebra does not close (Section 8.2), so the effective theory is not defined there as a closed canonical system. What the obstruction term (91) describes, when read as an interaction rather than as an algebra deficit, is a coupling of the collective momentum density to the foliation tilt with coefficient $(c_h^2/c^2 - 1)$. This is the canonical counterpart of the vacuum-Cherenkov channel discussed in Section 10: a tilted configuration exchanges energy between the collective and matter sectors, so the medium is driven rather than passive. The consequences are matters of energetics and stability, emission, absorption, and amplification of the mode, not of signal speed. The front is unaffected: the front-velocity argument requires only retarded response and free high-frequency behaviour, and neither is a passivity condition. Amplifying media are the standard demonstration, with gain confined strictly behind a front that remains luminal.

The magnitudes place the tilted case cleanly outside the vacuum setting of this paper. For a gravitational tilt $a_i = \partial_i \ln N \sim g/c^2 \simeq 1.1 \times 10^{-16}\,\mathrm{m}^{-1}$, the ratio of the drive to the mode gap is $(c_h^2/c^2 - 1)\, a/m_h \sim 4 \times 10^{-15}$, so the zero-tilt idealisation used throughout is exact for any gravitational environment. The drive becomes comparable to the gap only for lapse or frame gradients on scales below roughly forty metres, which no gravitational field supplies but which structured media supply easily; in such settings the tilt term acts as a source for the collective mode rather than as a modification of the causal structure, and lies outside the scope of this paper.

Finally, even if propagation at c_h within the medium is granted, no causal paradox can be assembled from it. The standard construction of a closed signal curve from superluminal propagation requires two emitters in relative motion, each using its own simultaneity, which is precisely a refoliation. Admissible evolution here is defined on a preferred foliation carrying a global time function (Section 8.2), and this is the usual causal-protection statement for theories of the Hořava and Einstein–aether type [37, 36]: a global time ordering exists on the admissible class, so no admissible evolution closes a signal curve, at any c_h.

9.5 Infrared limit

Three decoupling statements complete the analysis. The first is electrodynamic. For probes below the gap, $\omega \ll c_h m_h$, the pole in (92) is unreachable and the susceptibility reduces to the contact term $\chi_{\mathrm{TT}} \to -Z_h\, \Pi^{\mathrm{TT}}/(c_h^2 m_h^2)$ up to $\mathcal{O}(k^2/m_h^2)$ corrections. The collective response is then a rigid, non-propagating renormalisation confined to distances below $L_{\mathrm{coh}} = 1/m_h$, and conventional near-field electrodynamics is recovered. Far-field observables were never modified in the first place, since the Siegert reduction is exact for the S-matrix (Section 1). The second concerns the tensor sector. The massive $\ell = 2$ mode decouples in the same limit, and h_{ij}^{TT} is not the graviton, so its decoupling leaves the massless tensor sector of the emergent $f(R)$ reading untouched. The third concerns the scalar, which decouples on its own terms rather than with the tensor. It is the $\ell = 0$ amplitude mode of Section 5, gapped at M_s and therefore governed by that scale, not by m_h. What keeps it invisible at long wavelength is not its gap but its parity: the scalar portal is the pseudoscalar contraction (52), whose linear coupling to unpolarised matter vanishes identically, leaving a residual coupling that is short-range and spin-dependent. The anticipated long-wavelength behaviour is therefore recovered by decoupling, not by tuning.

10 Relation to massive gravity and emergent gravity

The dyality plasmon is easily mistaken for a Lorentz-violating spin-2 field of the type catalogued in the Standard Model extension. It is not a graviton of nonlinear massive gravity. It is an emergent, medium-type tensor response. This section makes the distinction explicit and draws out its consequences.

10.1 Massive gravity

Nonlinear Fierz–Pauli and de Rham–Gabadadze–Tolley theories propagate two spin-2 modes and require the dRGT tuning to keep the sixth, Boulware–Deser mode non-dynamical. That analysis tests a dynamical massive metric. It is not the relevant test here, because h_{ij}^{TT} is not a metric: it defines no line element, introduces no second metric, is not coupled universally to $T_{\mu\nu}$, and carries only the two TT polarisations of a spatial $\ell = 2$ collective mode. The scalar constraint whose failure produces the Boulware–Deser ghost is therefore absent by construction.

The relation to $f(R)$ gravity must be stated at three separate levels. First, the vacuum-elasticity analogy: the medium carries a gapped tensor branch and a distinct scalar branch, the qualitative spectrum of $f(R)$-type vacuum elasticity in the sense of Volovik [9], ghost-free for $f''(R) > 0$ [50, 51]. This analogy is claimed. Second, the emergent-gravity interpretation: the scalar of that spectrum, the scalaron, is the $\ell = 0$ collective branch of Section 6, and the dyality plasmon is the distinct, gapped $\ell = 2$ branch. This reading is motivated and is the target of the programme, but it rests on the conditional elements of Table 2. Third, a derived $f(R)$ effective action: none is derived in this paper, and h_{ij}^{TT} is explicitly not a metric perturbation. Claims of the third kind are not made, and the two branches are not conflated. The sector analysis of Section 5 derives the spectrum side of this statement, including the identification of the scalaron with the condensate amplitude mode, and shows that the universal $f(R)$ matter coupling is not realised: the emergent theory carries the $f(R)$ spectrum without the $f(R)$ coupling.

10.2 Emergent Lorentz invariance

The medium frame u^μ is not a fundamental preferred frame. In an emergent theory the substrate carries a rest frame, and Lorentz invariance is recovered as a low-energy, acoustic-limit property, exactly as quasiparticles in a superfluid see an emergent metric with an emergent light cone [39, 29]. The ratio $c_h/c = q_\phi/q_\Theta$ is the stiffness-to-inertia ratio of the tensor channel, derived in Section 7.3. It is a medium response coefficient, not a signal speed, and it runs with the probe scale through (84). Microscopic causality stays with the parent electromagnetic and defect degrees of freedom.

10.3 Comoving matter and vacuum Cherenkov

A collective mode at $c_h \neq c$ coupled to matter on the ordinary light cone would, if the two sectors carried independent and relatively boosted frames, radiate by a vacuum-Cherenkov process. The emergent picture removes this by referring both sectors to a single u^μ. We implement the condition covariantly, referring every TT projection to a single unit timelike u^μ written in Einstein–aether form [52, 36]; as stated in Section 6.1, u^μ is a background datum here and the dynamical completion in which it is varied is not supplied. The same dual-potential formalism also fixes the frame from the source side: under a boost the toroidal moment does not transform as a clean tensor but mixes with the quadrupole and mean-square-radius moments [17, Eq. (4.9)], so a pure toroid dipole, and with it the three-component decomposition $\mathbf{q}$ of Section 2, is defined only in the medium rest frame. The comoving condition is therefore imposed by the non-covariance of the toroidal moment, not merely assumed. In a fully emergent theory, in which the matter fields descend from the same substrate, the comoving condition $u^\mu_{\mathrm{matter}} = u^\mu_{\mathrm{medium}}$ is automatic. In the present theory the Standard Model matter is coupled by hand, so the condition is an explicit assumption and the constraint applies to the emergent dispersion of the coupled system rather than to two sectors with mismatched cones. Demonstrating that the matter sector co-emerges from the dyality vacuum, which would make the condition automatic, is left to future work. Canonically, this single shared frame is the selection of one common foliation on which the collective and matter sectors evolve together; Section 8 identifies it as the foliation-preserving (projectable) class on which the $c_h \neq c$ evolution is consistent, the obstruction otherwise being proportional to $(c_h^2/c^2 - 1)$ and to the foliation tilt.

10.4 The 4-form order parameter

The mixed Chern-2 invariant $\int F_\Theta \wedge F_\phi$ is a 4-form. It plays the role of the conserved 4-form field strength q of Klinkhamer–Volovik q-theory [11, 12], in which the vacuum is a self-tuning medium with an $f(R)$-like, variable-Λ response. Its stiffness sets the emergent tensor response. The vacuum is a medium whose macroscopic elasticity is the emergent spin-2 sector. The electroweak scale enters as the medium-formation scale, in line with the emergent electroweak gravity of McElrath [13], and the dyality plasmon is a pre-quantum precursor of the emergent metric rather than a competing graviton.

11 Conditional pole location and experimental prospects

This section evaluates a conditional benchmark: it derives a conditional location for the tensor pole from the relic condensate and then states the experimental programme it defines.

11.1 Pole location from the relic condensate

The emergent-gravity target of Section 10 identifies the vacuum condensate with the relic neutrino sector in the sense of McElrath [13]. Adopting this identification as a working hypothesis fixes the condensate density by cosmology, $n_{\rm pair} = \frac{1}{2} n_\nu$ with $n_\nu = 336\,{\rm cm}^{-3}$ and relic temperature $T_\nu = 1.945$ K, leaving no fitted density. Three results then follow from the response analysis, the first two unconditional given the hypothesis. (i) The particle-hole continuum of the relic sea terminates three to four decades below the dispersion $\omega^2 = c_h^2(k^2 + m_h^2)$, so ${\rm Im}\,\chi = 0$ on the pole: Landau damping is forbidden and the intrinsic linewidth vanishes, leaving inhomogeneity as the only broadening. The residue is the internal value $Z_h = 1/Z_t \simeq 1$ of (67) in the canonical normalisation, with pair-recoil corrections of order $(m_h/2m_\nu)^2 \sim 10^{-9}$. (ii) On the pole the normal-sea response is inertial, suppressed by $(\Delta_{\rm ph}/\omega)^2 \sim 10^{-8}$ relative to its static value: the normal component renormalises the stiffness negligibly and generates no mass. (iii) The gap is therefore generated in the condensed-pair zero mode, whose response carries no particle-hole denominator; taking the zero-mode self-energy denominator to be the gap itself closes the system in the form

$$m_h^3 = 2\, g_Q^2\, \gamma\, S_Q\, n_{\rm pair}\, \phi_c, \qquad S_Q = \tfrac{5}{3}, \tag{94}$$

with g_Q the portal matching coefficient, taken dimensionless in this normalisation so that the two sides of (94) carry equal dimension, γ an $\mathcal{O}(1)$ geometric factor, and $\phi_c \in (0, 1]$ the pair-correlated fraction. The relation of g_Q to the dimensionful portal $g_{h\mathcal{O}}$ requires the matching integral and is not supplied here. Equation (94) is independent of the neutrino mass and cube-root robust in every uncertain input.

The status of ϕ_c is fixed by two results. First, equilibrium pairing of today's relic sea fails: the exchange kernel is confined to momentum transfers $q \lesssim m_h \ll p_F$, and the resulting forward-scattering coupling $\lambda \simeq 0.1$ gives an equilibrium gap orders of magnitude below the relic kinetic scale. The condensate cannot form now. Second, a condensate formed at the electroweak epoch, where Section 7 places the condensation of the $U(1)_\phi$ fibre, cannot decay now: thermal depairing costs the zero-mode gap against a kinetic scale five hundred times smaller, the relic collision rate is some twenty-nine orders below the Hubble rate, Landau damping is forbidden by result (i), and the paired fraction is comoving-conserved. The absence of this particular damping channel bounds one contribution to the linewidth; it does not establish that the total intrinsic width vanishes, and inhomogeneous broadening is treated below as the dominant term. The fraction ϕ_c is therefore a primordial initial condition frozen at formation. The two results just stated are statements of local stability, and together they fix the shape of the free energy of the paired fraction at the present epoch: the first says the unpaired state is locally stable, $F''(0) > 0$; the second says the paired state is locally stable against depairing, $F''(\phi_c) > 0$. Two local minima of one functional require a barrier between them, so the landscape is a double well, whose minimal Landau form is

$$F(\phi) = a\,\phi^2 - b\,\phi^4 + c\,\phi^6 \tag{95}$$

with $a > 0$ the net cost of pairing from the trivial state under the phase-space restriction $q \lesssim m_h \ll p_F$, $-b\,\phi^4$ the mode-mediated gain, quartic because both the pairing and its mediator require pairs, and $c > 0$ the Pauli saturation. Two minima exist for $b^2 > 3ac$, at $\phi_c^2 = [b + (b^2 - 3ac)^{1/2}]/3c$, the condensed well is global for $b^2 > 4ac$, and the saturated well $\phi_c = 1$ corresponds to $2b = 3c$. The signs of a, b, c are fixed by the two results; their magnitudes, and hence the barrier height, the global minimum, and the selection between wells at the electroweak formation epoch where the coefficients differ, are not computed here. Bistability is therefore established in the structural sense, two locally stable states separated by a barrier, and ϕ_c is retained as a conditional initial-state parameter taking one of two values, zero or the condensed value, with the saturated value adopted for the numbers below. Adopting $\phi_c = 1$, the residual uncertainty is g_Q and γ alone, and the pole lies at $m_h = (1.6\text{–}2.9) \times 10^{-5}$ eV, a k=0 gap of 1.0–1.9 meV at the infrared velocity of (84), in the 250–470 GHz band, with coherence length $L_{\rm coh} = 1/m_h \simeq 0.7\text{–}1.3$ cm. These numbers are conditional on the condensate identification, the portal coefficient, and the zero-mode closure, in that order of importance, and are classified accordingly in Table 2.

A cross-check closes the source question. The alternative coupling route, through the neutrino toroid dipole moment of Dubovik and Kuznetsov [41] with $\tau \sim G_F$, carries the dimensionless

source weight τq^2: at pole and near-field kinematics this is smaller than the pair-quadrupole weight by some thirty orders, while at $q \sim m_Z$ it is dominant. The two routes divide cleanly by kinematics, with no tension: the toroid dipole fixes the evaluation scale of the Dirac condition in Section 7, and the pair condensate fixes the pole. The formation-epoch calculation of ϕ_c and the quantitative probe design built on these ranges are deferred.

11.2 A resonant near-field search

The natural setting is not far-field scattering. It is mesoscopic and condensed-matter near-field systems in which source-current topology, toroidal multipoles, and reactive electromagnetic structure are retained, the same regime in which anapole and toroidal metamaterial responses are measured [6]. The phase velocity is conditional and falsifiable: $c_h = c/(2\alpha(m_Z^2)) \approx 64.5\,c$ at the Z pole, running to $\approx 66.5\,c$ in the near-field window through (84). The dispersion itself is a signature: its presence or absence distinguishes the running inertia of Section 7.3 from a static matching value. The object a near-field probe measures is not h_{ij}^{TT} itself but the induced TT response of the medium, the dressed susceptibility derived in Section 6.2,

$$\chi_{\mathrm{TT}}(\omega, \mathbf{k}) = \langle \mathcal{O}_{ij}^{(2)}\, \mathcal{O}_{kl}^{(2)} \rangle_{\mathrm{TT}} \sim \Pi_{ijkl}^{\mathrm{TT}}(\hat{\mathbf{k}})\, \frac{Z_h}{\omega^2/c_h^2 - \mathbf{k}^2 - m_h^2 + \mathrm{i}\epsilon} \tag{96}$$

a tensorial near-field susceptibility with a pole on the dyality-plasmon dispersion. The signature that distinguishes it from ordinary toroidal or anapole electrodynamics is this pole: a resonant $\ell = 2$ response tracking $\omega^2 = c_h^2(\mathbf{k}^2 + m_h^2)$ with the velocity (84), absent from a purely radiative or static near-field description. Because the ordered medium possesses a rest frame, that of the relic condensate, the search is most naturally cast as an aether-frame null experiment of Michelson–Morley type: the observable is the frame-dependent modulation of the resonance, not an absolute rate.

The architecture follows from that framing. If the condensate is present it already carries the mode, so the apparatus does not have to generate one. A transmission experiment, with a source metasurface driving a receiver metasurface, would pass through the portal twice and its rate would scale as $g_{h\mathcal{O}}^2$, the one quantity Section 7 leaves unfixed. A single resonant receiver passes through it once and responds linearly in $g_{h\mathcal{O}}$. The candidate apparatus is therefore one dielectric anapole disc per resonance frequency, a high-resistivity silicon resonator at its anapole condition supplying the unreduced toroidal near-field content the coupling requires, with no source disc.

The static field is supplied by a Halbach stack of axial period Λ. It furnishes the magnetic factor of the $E{\cdot}B$ coupling without AC power, and its periodicity fixes the wavevector to which a disc is sensitive: the stack acts as a grating that momentum-matches an ambient mode to the disc's resonance. A periodic rather than single-harmonic field matches along the Fourier comb $k_n = 2\pi n/\Lambda$, so an array of discs probes the dispersion at a set of wavevectors,

$$f_n(\Lambda) = \frac{c_h}{2\pi}\sqrt{\left(\frac{2\pi n}{\Lambda}\right)^2 + m_h^2} \tag{97}$$

rather than at one, the disc tuned to f_n reading tooth n. The comb samples the $\ell = 2$ dispersion from $k \simeq 0$, the k=0 gap band, out to $k \sim$ few $\times\, m_h$ at centimetre periods, where the crossover from gapped to linear dispersion is maximal. The spacing between adjacent teeth, $\mathrm{d}f/\mathrm{d}n = (c_h/\Lambda)\, k_n/\sqrt{k_n^2 + m_h^2}$, approaches the constant c_h/Λ for $k_n \gg m_h$ and reads c_h directly there, while the departure of f_2/f_1 from two reads the gap-to-momentum ratio; the gap band is reached as the large-Λ limit of the fundamental rather than as an $n = 0$ tooth. Table 1 collects the conditional parameters and the quantities an experiment must fix; every entry inherits the hypothesis chain of Section 11 and the classifications of Table 2. Absolute sensitivity, and therefore the choice of receiver and the noise budget, depends on the achievable coupling to the disc and is left to a dedicated instrument study; the falsifiable content here is the dispersion (97) and the frame-dependent modulation below, both independent of the coupling normalisation.

Four amplitude-free discriminators are proposed to separate a genuine signal from conventional electromagnetic responses. Structured media possess resonances, cutoffs, and nontrivial dispersion of their own, so each discriminator is a test against a calculated background response of the apparatus rather than a proof of exclusion on its own. The primary one is the aether-frame signature. The resonance on the comb teeth must carry the frame-dependent sidereal modulation (98) derived in the next subsection, with the readout and template conditions stated there; the first-order shift vanishes at the $k = 0$ gap, itself a check, since the gap line must not modulate while the teeth do. The depth on the teeth is kinematic, fixed by known velocities and

Table 1. Conditional parameters of the near-field search. All entries assume the condensate identification, the portal normalisation, and the zero-mode closure.

quantity	conditional value
pole mass m_h	$(1.6\text{–}2.9)\times10^{-5}$ eV, eq. (94) at $\phi_c = 1$
k=0 gap	1.0–1.9 meV; 250–470 GHz at the infrared velocity of (84)
residue Z_h	$\simeq 1$; pair-recoil corrections $\mathcal{O}(10^{-9})$
intrinsic width	Landau damping kinematically forbidden; linewidth inhomogeneity-limited
coherence length $L_{\rm coh} = 1/m_h$	0.7–1.3 cm; cell size $a \ll 1/k$ within the window (53)
frequency and wavevector range	k=0 gap 257–466 GHz; comb teeth (97) sample $k \simeq 0$ to $k \sim 3\,m_h$ for Λ = 2–10 cm, reaching into the THz
source geometry	Halbach stack, axial period Λ = 2–10 cm, B_0 = 1–5 T. Detector: Si metasurface with an anapole resonance at the target $k = 0$ gap frequency.
portal coupling	not predicted; the weights are fixed but the overall normalisation is nonperturbative (Section 7). For reference, an electroweak-suppressed portal would give $1/v_{\rm EW} \simeq 4\times10^{-12}$ eV^{-1}
measurable	sidereal frequency modulation of the comb teeth, depth $(v/c_h)k_n/\sqrt{k_n^2+m_h^2} \sim 2\times10^{-5}$, travelling-wave readout; the scattering channel is the cross-polarised TT response (96)

independent of the coupling, so a null at that depth constrains the condensate rest frame directly, in the manner of the classical aether-drift experiments. The remaining three are corroborating. The resonance must track the comb (97) as Λ is changed, a locus no mode propagating at c can occupy; it must null under scalar and vector drive patterns; and it must survive exchange of every apparatus material, the sample being the vacuum.

11.3 Laboratory motion and the frame test

If the medium frame is that of the relic sea, laboratory motion supplies the decisive test. For a travelling Fourier component the first-order shift is $\delta\omega = -\mathbf{v}\cdot\mathbf{k}$, so

$$\frac{\delta\omega}{\omega} = -\frac{v}{c_h}\,\frac{k}{\sqrt{k^2+m_h^2}}\cos\vartheta \tag{98}$$

The relic sea is cosmological rather than galactic, so the relevant velocity is the laboratory's motion through the cosmic microwave background frame, $v \simeq 370\,\mathrm{km\,s^{-1}}$, giving $v/c_h \simeq 1.9\times10^{-5}$ at $c_h = 66.5c$. This is a directional scale, not a universal modulation depth. The first-order term vanishes at the gap, which is itself a check, since the gap line must not modulate while the teeth do, and it cancels in a reciprocal standing-wave measurement, so a travelling-wave readout is required. The disc's orientation and the time dependence of ϑ must be folded in before a sidereal template is specified, and the measured phase would distinguish motion through the cosmological frame from motion through the galactic halo.

This modulation is the primary discriminator; those listed above corroborate it. A null across the band therefore limits the coupling-weighted response at the achieved sensitivity, and, the theory being bistable (Section 11), a clean null at design sensitivity excludes the formation branch conditional on the stated chain. It does not exclude an arbitrarily weakly coupled condensate.

12 Discussion and summary

Table 2 summarises the epistemic status of the principal statements of the paper, gathered here as a closing inventory. Terms such as fixed and derived are used only for entries in the derived class, and only relative to the stated assumptions of that entry; the conditional entries are the paper's open hypotheses, and the deferred quantitative confrontations (the parity-violating residual coupling of the scalar branch, the polarised-media force bounds, and the formation-epoch calculation of the condensed fraction) are listed with them.

We have constructed a medium effective theory for a transverse-traceless $\ell = 2$ collective response of the pre-reduced electromagnetic near-field sector. The result is a framework with a core derived from stated assumptions and a sharply bounded conditional part: the core is that the unreduced toroidal sector supports the collective $\ell = 2$ response, that its defect charge unit follows from the comoving inertial match, and that the phase-velocity ratio $c_h^2 = (q_\phi/q_\Theta)^2$ is selected by the charged-rotor constitutive assignment; the conditional part is narrowed to the dyon ultraviolet completion, the emergent-gravity stiffness normalisation, and the scale-running prescription that together set the numerical value of c_h, each resting on a stated hypothesis. The central distinction is between the transverse radiative Fock vacuum and the local near-field vacuum before the Siegert reduction is imposed. The former is the correct vacuum of the reduced asymptotic photon Hilbert

space. The latter retains the reactive multipole content and the boundary, edge, or topological source data responsible for local near-field structure. In this sense the electromagnetic vacuum remains source-free in the ordinary Maxwell sense,

$$j^{\mu}_{\text{ext}} = 0$$

but it need not be structure-free once the source-bound sector has not been quotiented out.

The dyality degree of freedom is meaningful only in this pre-reduced sector. Ordinary electromagnetic duality rotates the field strengths while treating the source sector as fixed or absent. Dyality, by contrast, is the simultaneous rotation of fields and sources appropriate to dyonic source-field configurations. When the source information is erased, dyality reduces to a passive relabelling. When that information is retained as near-field, boundary, or defect data, it can be represented by a constrained relative connection A^{ϕ}_{μ}. Its curvature enters the mixed topological form

$$F_{\Theta} \wedge F_{\phi}$$

which supplies the geometric structure from which the spin-2 channel is extracted.

The local mixed-curvature bilinear contains more structure than a radiative photon state. Its symmetric-traceless part isolates the algebraic spin-2 numerator, while the physical source coupled in the EFT is its transverse-traceless projection,

$$\widetilde{\mathcal{O}}^{(2)}_{ij}(\mathbf{k}) = \Pi^{\mathrm{TT}}_{ijkl}(\mathbf{k})\widetilde{Y}_{kl}(\mathbf{k})$$

The effective field h^{TT}_{ij} is therefore not a graviton and not a nonlinear massive-gravity field. It is the coarse-grained collective mediator of the projected spin-2 response of the near-field medium, an emergent, $f(R)$-type tensor mode in the sense of Volovik. Because the emergent sector carries a single extra scalar rather than a second spin-2 mode, the Boulware–Deser problem of massive gravity does not arise. The medium frame u^{μ} is the substrate frame of an emergent theory, not a fundamental preferred frame, and the matter coupling is referred to that same frame.

The schematic ultraviolet completion, dyality-charged Witten dyons bound to the codimension-2 defect sector, motivates the charge structure. The Dirac condition fixes the topological charge normalisation,

$$g^{\text{top}}_{h\chi} = \frac{q_{\Theta}q_{\phi}}{\hbar c} = \frac{1}{2}$$

and, given the emergent-gravity stiffness normalisation, the same condition fixes the collective phase velocity at the Z pole,

$$c_h = \frac{c}{2\alpha(m_Z^2)} \approx 64.5c$$

which runs to $\approx 66.5\,c$ in the near field through (84). The charge normalisation $g^{\text{top}}_{h\chi}$ is a topological quantity, not a substitute for low-energy matching: the renormalised portal coupling and Wilson coefficients retain the usual loop, threshold, and operator-normalisation factors. The Witten dyon mass satisfies the 't Hooft–Polyakov bound $m_{\chi} \gtrsim 2.3\,\text{TeV}$ as the natural defect-core scale of the model.

Beyond the classical EFT, the medium-frame action carries a covariant Hamiltonian structure. Its quadratic form is the transverse-traceless projection of a De Donder–Weyl field-strength polysupermetric of the kind used in the covariant quantisation of the teleparallel equivalents to general relativity. The phase velocity $c_h^2 = Z_s/Z_t$ is its eigenvalue ratio, degenerate ($c_h = c$) in the Lorentz-invariant case and lifted by a single rank-one condensate term. The associated Tomonaga–Schwinger evolution is consistent on a definite foliation class: the obstruction to full refoliation invariance is one term proportional to $(c_h^2/c^2 - 1)$ and to the foliation tilt ∇u, removed by either of two independent conditions: the projectable foliations selected by the medium's comoving frame, which suffice for any c_h, or the acoustic limit $c_h \to c$. The preferred medium frame is therefore not in tension with covariant evolution; it fixes the admissible foliations, placing the construction in the Hořava–Lifshitz class for which a foliation-preserving canonical evolution is available. Whether the residual ultraviolet anomaly of the regularised hypersurface generator can be removed is, as for general relativity, left open.

These results are not independent. The same rank-one condensate term $(Z_t - Z_s)\,u^{\rho}u^{\sigma}$ that lifts the polysupermetric eigenvalue ratio away from $Z_s = Z_t$ is the anisotropy of the comoving inertial match, and that match is what fixes the defect charge in the first place. In the flat-isotropic (acoustic) limit the term vanishes: $c_h \to c$, and the inertial cell is flat and isotropic. Under the Regge reading of Section 8.3 this flatness and isotropy give the balance $\sum_{\alpha} C_{\alpha} = 0$ and a common

magnitude $|C_2| = \frac{1}{6}$, a geometric normalisation. Turning on the condensate is the single deformation that both lifts c_h above c and breaks boost invariance, preserving the rest-frame spatial rotations. The charge and the velocity are therefore controlled by one structure, read in two limits: the symmetric limit, where the charge unit is derived, and the anisotropic regime, where $c_h \neq c$ becomes observable. What carries the result between them is quantisation. The charge $|C_2| = \frac{1}{6}$ is a topological invariant fixed by integrality of the total Chern number, so it cannot drift as the anisotropy is switched on; the value derived in the isotropic limit persists, unchanged, into the driven sub-wavelength regime where the velocity signature lives. The acoustic limit is thus not a separate theory from the anisotropic one but its long-wavelength face: the quantised charge is the bridge, and what this paper determines is the structure of the anisotropic limit. The departure from $c_h = c$ is one-dimensional, tied to the dyality charge, and visible only where a probe resolves the medium's stiffness, the same sub-wavelength near field $ka \ll 1$ that defines the construction's validity window.

The resulting physical picture is that the standard radiation limit discards a local tensor response that can survive in the unreduced near field. Retaining the local phase-space structure of electromagnetic oscillator cells before radiative projection gives a medium-like vacuum response with a transverse-traceless collective channel. The dyality plasmon is the massive tensor ($\ell = 2$) excitation of that channel, and its experimental setting is the mesoscopic near-field regime discussed in Section 10 rather than far-field scattering.

Acknowledgments
The author is grateful for discussions with Bob McElrath, Andrey Miroshnichenko, Vladlen Shvedov and Giovanni Modanese for the critical reading of early drafts. The author used Claude to refine the language and formatting of this manuscript and assumes full responsibility for the content.

Funding
This research was partially supported by the Commonwealth of Australia as represented by the Defence Science and Technology Group of the Department of Defence.

Data availability
No new experimental data were created or analysed in this study. Numerical scripts reproducing the conditional estimates of Sections 5 and 11 are available from the author on reasonable request.

Supplementary data

A Spatial transverse-traceless projector

Let $\boldsymbol{k} \neq 0$ and define $\hat{k}_i = k_i/|\boldsymbol{k}|$. The transverse projector on spatial indices is

$$\Pi_{ij}(\boldsymbol{k}) = \delta_{ij} - \hat{k}_i \hat{k}_j \tag{99}$$

The rank-four transverse-traceless projector is

$$\Pi^{\mathrm{TT}}_{ijkl}(\boldsymbol{k}) = \frac{1}{2}[\Pi_{ik}(\boldsymbol{k})\Pi_{jl}(\boldsymbol{k}) + \Pi_{il}(\boldsymbol{k})\Pi_{jk}(\boldsymbol{k})] - \frac{1}{2}\Pi_{ij}(\boldsymbol{k})\Pi_{kl}(\boldsymbol{k}) \tag{100}$$

where the coefficient of the last term is $1/(d-1) = 1/2$, the transverse subspace being $(d-1) = 2$ dimensional for $d = 3$ spatial indices. With this normalisation $\Pi^{\mathrm{TT}}_{ijij} = 2$, the number of TT polarisations, and the projector is idempotent. It satisfies

$$k_i\, \Pi^{\mathrm{TT}}_{ijkl}(\boldsymbol{k}) = 0, \qquad \delta_{ij}\, \Pi^{\mathrm{TT}}_{ijkl}(\boldsymbol{k}) = 0 \tag{101}$$

Given a symmetric tensor $X_{kl}(\boldsymbol{k})$, its transverse-traceless part is

$$X^{\mathrm{TT}}_{ij}(\boldsymbol{k}) = \Pi^{\mathrm{TT}}_{ijkl}(\boldsymbol{k})\, X_{kl}(\boldsymbol{k}) \tag{102}$$

Applying this to the mixed-curvature bilinear

$$Y_{kl}(x) = F^{\Theta}_{km}(x)F^{\phi}_{lm}(x) + F^{\phi}_{km}(x)F^{\Theta}_{lm}(x) \tag{103}$$

one obtains the momentum-space definition

$$\widetilde{\mathcal{O}}^{(2)}_{ij}(\boldsymbol{k}) = \Pi^{\mathrm{TT}}_{ijkl}(\boldsymbol{k})\, \widetilde{Y}_{kl}(\boldsymbol{k}) \tag{104}$$

which is the precise form of the operator used in the spin-2 spectral density (51).

B Numerical illustration at three degrees of freedom

The identification of Section 4.2 has two halves. The topological half is that the winding of a dyality defect survives the breaking of its resonance; the dynamical half, the posited one, is that the breaking releases the defect into a connected web. This appendix exhibits both in a system where they can be computed, at the same number of degrees of freedom as the cell. What the system does not supply is the value of the charge, and the scope paragraph says why.

Model. We use the Froëschlé map of two coupled standard maps,

$$\begin{aligned} y_1' &= y_1 - \tfrac{K_1}{2\pi}\sin 2\pi x_1 - \tfrac{b}{2\pi}\sin 2\pi(x_1+x_2), \\ y_2' &= y_2 - \tfrac{K_2}{2\pi}\sin 2\pi x_2 - \tfrac{b}{2\pi}\sin 2\pi(x_1+x_2), \end{aligned} \qquad x_i' = x_i + y_i' \tag{105}$$

the standard model of Arnold diffusion [53, 54]. It is a four-dimensional symplectic map, the section of a flow with three degrees of freedom, so its invariant two-tori are the sections of the invariant three-tori of the cell and carry the same codimension two. Each degree of freedom is a compact phase with its conjugate momentum, so each carries a doublet in the sense of Section 4.2, and the coupling $b\sin 2\pi(x_1+x_2)$ is a mixed term between the two phases. We take (x_1, y_1) as the ϕ sector and (x_2, y_2) as the Θ sector. The coupling b plays the role of the perturbation ε of Section 2: raising it destroys resonant tori in order of their arithmetic quality.

Method. Orbits are classified by the fast Lyapunov indicator, the growth of a tangent vector under the linearised map, which separates regular from chaotic orbits without reference to any visual criterion. The phase space is displayed as a three-dimensional slice [55]: a point is recorded each time an orbit returns to the surface $x_2 = 0$ and plotted at (x_1, y_2, y_1). In this representation a regular orbit on an invariant two-torus draws a closed loop, a nested family of tori sweeps a tube, and a chaotic orbit fills a volume; the tubes of Berry's picture of $n = 3$ phase space [19] appear directly. The slice has a second meaning here. Holding $x_2 = 0$ holds the Θ phase fixed, so the winding of a ϕ-loop counted on the slice is a winding at fixed Θ, which is the slice linking datum by which Section 4.3 defines the charge; the innermost loop of each nested family sits at the elliptic centre of its resonance, where the ϕ doublet vanishes, and is the defect locus at fixed Θ. Identical seeds are integrated at two couplings, and the orbit reported below is one selected automatically as regular at the lower coupling and chaotic at the higher.

Result. Figure 2 shows the outcome. At the lower coupling the surviving tori appear as loops and tubes with the gaps between them connected, as the codimension forces, and each family closes on a central loop, the defect locus. At the higher coupling many tori are gone. The central loops of the families that survive are unchanged, which is the persistence: destruction of the surrounding tori has not moved the defect or altered its winding. The selected orbit, whose torus does not survive, no longer returns to a loop: its returns spread through the connected region between the tori that remain, and would fill it densely in the limit of long times. This is the release. Both halves of the identification are therefore visible in one computation.

Scope. The model carries both sectors and the mixed structure between them, which is why the slice linking geometry of the charge is reproduced and not merely suggested. What it does not reproduce is the normalisation. The integrality condition and the signed-basis truncation that fix $|C_2| = \frac{1}{6}$ are properties of the cell's Chern structure, and the map's mixed coupling is a dynamical resonance term rather than a topological density. The computation therefore illustrates the geometry of the charge and the dynamics of persistence and release at three degrees of freedom; the value of the charge comes from the cell.

References

[1] Siegert A J F 1937 *Physical Review* **52** 787–789 ISSN 0031-899X URL https://link.aps.org/doi/10.1103/PhysRev.52.787

[2] Naus H W L 1997 *Physical Review C* **55** 1580–1582 ISSN 0556-2813, 1089-490X URL https://link.aps.org/doi/10.1103/PhysRevC.55.1580

[3] Hadjimichael E 1979 *PHYSICS LETTERS* **85**

[4] Dubovik V M and Tugushev V V 1990 *Physics Reports* **187** 145–202 ISSN 0370-1573 URL https://www.sciencedirect.com/science/article/pii/037015739090042Z

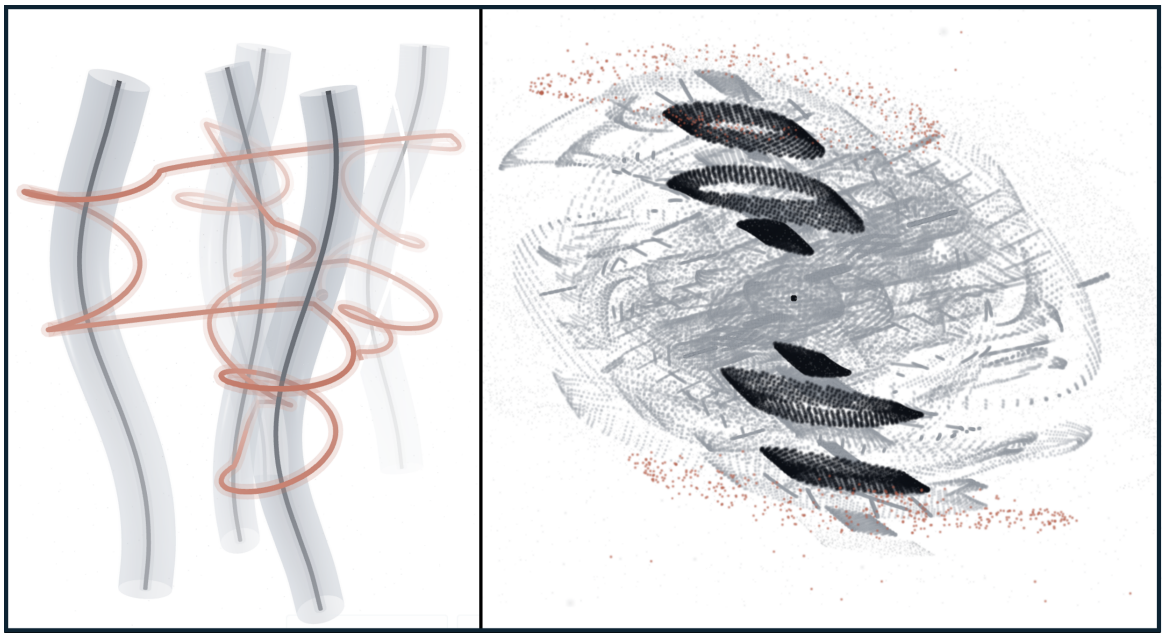


Figure 2. The defect–resonance identification of Section 4.2, drawn schematically (left) and computed from the model flow (105) (right). *Left:* the invariant tori of the cell are three-dimensional and the energy surface is five-dimensional, so in projection a torus appears as a tube and the gaps between tori form a single connected region [19]. The filament along each tube is the locus of the dyality defect (37) of that cycle, confined while its torus holds. The highlighted curve is one such locus after its resonance has broken: released, it threads and links the surviving tubes, and on longer times it returns arbitrarily close to every part of the connected region. Its winding is unchanged by the transition; only the geometry of its locus is. *Right:* the three-dimensional phase-space slice of the Froëschlé map described in the text, at $K_1 = 0.7$, $K_2 = 0.85$, coupling $b = 0.4$ and 1.2×10^6 returns per orbit. Each point is a return of an orbit to $x_2 = 0$, which holds the Θ phase fixed, plotted at (x_1, y_2, y_1); a regular orbit on an invariant two-torus draws a closed loop, a nested family sweeps a tube, and a chaotic orbit fills a volume. The loops and tubes are surviving tori; the innermost loop of each family is the ϕ-defect locus at fixed Θ, the counterpart of the filaments on the left. The highlighted points are a single orbit whose torus does not survive this coupling: its returns no longer lie on a loop but spread through the connected region between the tori that remain. Nothing in the right panel is special to the dyality construction. Persistence and release are properties of near-integrable dynamics at three degrees of freedom with compact phases, and they appear here in a system chosen for being standard rather than for being ours. What the cell contributes is not the mechanism but the charge the defect carries.

[5] Zel'Dovich I B

[6] Miroshnichenko A E, Evlyukhin A B, Yu Y F, Bakker R M, Chipouline A, Kuznetsov A I, Luk'yanchuk B, Chichkov B N and Kivshar Y S 2015 *Nature Communications* **6** 8069 ISSN 2041-1723 URL https://www.nature.com/articles/ncomms9069

[7] Han M Y and Biedenharn L C 1971 *Il Nuovo Cimento A* **2** 544–556 ISSN 0369-3546, 1826-9869 URL http://link.springer.com/10.1007/BF02899873

[8] Chern A and Ishida S 2025 Implicit representations of codimension-2 submanifolds and their prequantum structure arXiv:2507.11727 [math] URL http://arxiv.org/abs/2507.11727

[9] Volovik G E 2024 *JETP Letters* **119** 564–571 ISSN 0021-3640, 1090-6487 URL https://link.springer.com/10.1134/S0021364024600526

[10] Volovik G E 2012 The Superfluid Universe arXiv:1004.0597 [cond-mat] URL http://arxiv.org/abs/1004.0597

[11] Klinkhamer F R and Volovik G E 2017 *Modern Physics Letters A* **32** 1750103 ISSN 0217-7323, 1793-6632 URL https://www.worldscientific.com/doi/abs/10.1142/S0217732317501036

[12] Klinkhamer F R and Volovik G E 2008 *Physical Review D* **77** 085015 ISSN 1550-7998, 1550-2368 URL https://link.aps.org/doi/10.1103/PhysRevD.77.085015

[13] McElrath B 2008 Emergent Electroweak Gravity arXiv:0812.2696 [gr-qc] URL http://arxiv.org/abs/0812.2696

[14] Zwanziger D 1971 *Physical Review D* **3** 880–891

[15] Blagojević M and Senjanović P 1988 *Physics Reports* **157** 233–346

[16] Hinterbichler K 2012 *Reviews of Modern Physics* **84** 671–710 ISSN 0034-6861, 1539-0756 URL https://link.aps.org/doi/10.1103/RevModPhys.84.671

[17] Dubovik V M, Martsenyuk M A and Saha B 2000 *Physical Review E* **61** 7087–7097 ISSN 1063-651X, 1095-3787 URL https://link.aps.org/doi/10.1103/PhysRevE.61.7087

[18] Hubbard J and Ilyashenko Y 2003 *Discrete and Continuous Dynamical Systems* **10** 367–385 ISSN 1078-0947 URL http://www.aimsciences.org/journals/displayArticles.jsp?paperID=146

[19] Berry M V 1978 Regular and irregular motion *AIP Conference Proceedings* vol 46 (AIP) pp 16–120 ISSN 0094243X URL https://pubs.aip.org/aip/acp/article/46/1/16-120/684628

[20] Berry M V and Keating J P 1990 *Journal of Physics A: Mathematical and General* **23** 4839–4849 ISSN 0305-4470, 1361-6447 URL https://iopscience.iop.org/article/10.1088/0305-4470/23/21/024

[21] Aubry S and Le Daeron P 1983 *Physica D: Nonlinear Phenomena* **8** 381–422 ISSN 01672789 URL https://linkinghub.elsevier.com/retrieve/pii/0167278983902336

[22] Mather J N 1982 *Topology* **21** 457–467 ISSN 00409383 URL https://linkinghub.elsevier.com/retrieve/pii/0040938382900234

[23] Greene J M 1979 *Journal of Mathematical Physics* **20** 1183–1201 ISSN 0022-2488, 1089-7658 URL https://pubs.aip.org/jmp/article/20/6/1183/449401/A-method-for-determining-a-stochastic-transition

[24] Mather J N 1991 *Mathematische Zeitschrift* **207** 169–207 ISSN 0025-5874, 1432-1823 URL http://link.springer.com/10.1007/BF02571383

[25] Xue J 2023 *Science China Mathematics* **66** 1649–1712 ISSN 1869-1862 URL https://doi.org/10.1007/s11425-022-2118-1

[26] Nekhoroshev N N 1977 *Russian Mathematical Surveys* **32** 1–65 ISSN 0036-0279, 1468-4829 URL https://www.mathnet.ru/eng/rm3304

[27] Chirikov B V 1979 *Physics Reports* **52** 263–379 ISSN 03701573 URL https://linkinghub.elsevier.com/retrieve/pii/0370157379900231

[28] Mackay R, Meiss J and Percival I 1984 *Physica D: Nonlinear Phenomena* **13** 55–81 ISSN 01672789 URL https://linkinghub.elsevier.com/retrieve/pii/0167278984902707

[29] Volovik G E 2009 *The universe in a helium droplet* repr ed (*International series of monographs on physics* no 117) (Oxford: Oxford University Press) ISBN 978-0-19-850782-6 978-0-19-956484-2

[30] Pasti P, Sorokin D P and Tonin M 1997 *Physical Review D* **55** 6292–6298 _eprint: hep-th/9611100

[31] Sorokin D P 2025 _eprint: 2501.17918

[32] Maznytsia A, Preitschopf C R and Sorokin D P 1999 *Nuclear Physics B* **539** 438–452 _eprint: hep-th/9805110

[33] Marsden J and Weinstein A 1974 *Reports on Mathematical Physics* **5** 121–130 ISSN 00344877 URL https://linkinghub.elsevier.com/retrieve/pii/0034487774900214

[34] Dubovik V M, Saha B and Rubin J L 2000 *Ferroelectrics Letters Section* **27** 1–6 ISSN 0731-5171, 1563-5228 URL http://www.tandfonline.com/doi/abs/10.1080/07315170008204647

[35] Gabbrielli R, Jiao Y and Torquato S 2014 *Physical Review E* **89** 022133 ISSN 1539-3755, 1550-2376 arXiv:1402.6745 [cond-mat] URL http://arxiv.org/abs/1402.6745

[36] Jacobson T 2008 Einstein-aether gravity: a status report arXiv:0801.1547 [gr-qc] URL http://arxiv.org/abs/0801.1547

[37] Hořava P 2009 *Physical Review D* **79** 084008 ISSN 1550-7998, 1550-2368 URL https://link.aps.org/doi/10.1103/PhysRevD.79.084008

[38] Volovik G E 2022 *Journal of Experimental and Theoretical Physics* **135** 663–670 ISSN 1090-6509 URL https://doi.org/10.1134/S106377612211019X

[39] Barceló C, Liberati S and Visser M 2011 *Living Reviews in Relativity* **14** 3 ISSN 2367-3613, 1433-8351 URL http://link.springer.com/10.12942/lrr-2011-3

[40] Witten E 1979 *Physics Letters B* **86** 283–287 ISSN 03702693 URL https://linkinghub.elsevier.com/retrieve/pii/0370269379908384

[41] Dubovik V M and Kuznetsov V E 1998 *International Journal of Modern Physics A* **13** 5257–5277 ISSN 0217-751X, 1793-656X arXiv:hep-ph/9606258 URL http://arxiv.org/abs/hep-ph/9606258

[42] Hooft G 1974 *Nuclear Physics B* **79** 276–284 ISSN 05503213 URL https://linkinghub.elsevier.com/retrieve/pii/0550321374904866

[43] Polyakov A M 1974 *Soviet Journal of Experimental and Theoretical Physics Letters* **20** 194 ISSN 0021-3640 aDS Bibcode: 1974JETPL..20..194P URL https://ui.adsabs.harvard.edu/abs/1974JETPL..20..194P

[44] Collaboration T A, Collaboration t D, Collaboration t L, Collaboration t O, Collaboration t S, Group t L E W, electroweak t S and groups h f 2006 *Physics Reports* **427** 257–454 ISSN 03701573 arXiv:hep-ex/0509008 URL http://arxiv.org/abs/hep-ex/0509008

[45] Sirlin A 1980 *Physical Review D* **22** 971–981 ISSN 0556-2821 URL https://link.aps.org/doi/10.1103/PhysRevD.22.971

[46] Chester D and Pandey V 2026 Covariant Hamiltonian quantization of teleparallel equivalents to general relativity version Number: 1 URL https://arxiv.org/abs/2603.26961

[47] Kanatchikov I V 2013 *Journal of Physics: Conference Series* **442** 012041 ISSN 1742-6596 URL https://iopscience.iop.org/article/10.1088/1742-6596/442/1/012041

[48] Günther C 1987 *Journal of Differential Geometry* **25** ISSN 0022-040X URL https://projecteuclid.org/journals/journal-of-differential-geometry/volume-25/issue-1/The-polysymplectic-Hamiltonian-formalism-in-field-theory-and-calculus-of/10.4310/jdg/1214440723.full

[49] Teitelboim C 1973 *Annals of Physics* **79** 542–557 ISSN 00034916 URL https://linkinghub.elsevier.com/retrieve/pii/0003491673900961

[50] Sotiriou T P and Faraoni V 2010 *Reviews of Modern Physics* **82** 451–497 ISSN 0034-6861, 1539-0756 URL https://link.aps.org/doi/10.1103/RevModPhys.82.451

[51] De Felice A and Tsujikawa S 2010 *Living Reviews in Relativity* **13** 3 ISSN 2367-3613, 1433-8351 URL http://link.springer.com/10.12942/lrr-2010-3

[52] Jacobson T and Mattingly D 2001 *Physical Review D* **64** 024028 ISSN 0556-2821, 1089-4918 URL https://link.aps.org/doi/10.1103/PhysRevD.64.024028

[53] Froeschle C 1972 *Astronomy and Astrophysics* **16** 172 ISSN 0004-6361 aDS Bibcode: 1972A&A....16..172F URL https://ui.adsabs.harvard.edu/abs/1972A&A....16..172F

[54] Lega E, Guzzo M and Froeschlé C 2003 *Physica D: Nonlinear Phenomena* **182** 179–187 ISSN 01672789 URL https://linkinghub.elsevier.com/retrieve/pii/S0167278903001210

[55] Lange S, Richter M, Onken F, Bäcker A and Ketzmerick R 2014 *Chaos: An Interdisciplinary Journal of Nonlinear Science* **24** 024409 ISSN 1054-1500, 1089-7682 URL https://pubs.aip.org/cha/article/24/2/024409/280732/Global-structure-of-regular-tori-in-a-generic-4D

Table 2. Status of the principal statements. Derived means the result follows from the stated assumptions of the cited section. Conditional means the result additionally depends on a named hypothesis listed in the final column.

Statement	Status	Basis and named hypotheses
Dyality as simultaneous source-field rotation, (10)	definition	Section 3; reduces to passive duality when source data are absent.
A^{ϕ}_{μ} as constrained relative-phase connection, (13), (14)	definition and derived constraint	the Zwanziger dual potential of two-potential electrodynamics, Section 3.1; the bulk constraint follows from the absence of an independent kinetic term together with the constitutive relation and decay of the residual curvature at spatial infinity.
Phase identification $\phi_{\mathbf{w}} = -\theta_{\mathbf{w}}$, (36)	derived, leading order	coherent Beltrami sector, Section 4.2; topological content is independent of the identification.
Defect-resonance identification, (42)	persistence derived; release posited	the dyality defect carries integer winding, which survives the breaking of its resonance (theorem); that the broken resonance's chaotic layer releases the defect into the transport web at $n = 3$ degrees of freedom is the hypothesis, resting on the Arnold-diffusion assumptions of Section 2; dynamical half illustrated in Appendix B.
Charge unit $\lvert C_2\rvert = \frac{1}{6}$ and balance $\sum_{\alpha} C_{\alpha} = 0$, (22), (23)	derived, conditional	flatness and isotropy of the comoving inertial match, integrality of the total Chern number, and the truncation of the elementary classes to the six signed basis classes; Sections 4 and 8.3. Hypotheses: the acoustic-limit inertial match, that truncation, and the Regge reading of the charge, which is a geometrical interpretation supplying a normalisation. Not obtained from the Chern–Ishida holonomy theorem.
Propagating $\ell = 2$ pole	derived under stated kernel conditions	channel selection in Section 5.1; pole obtained by resummation of the defect response, Section 6.2, conditional on the separability, stability and softening conditions stated there; conditional pole location in Section 11.
Scalar and vector sector content	derived; scalaron mass conditional	vector channel empty by hypersurface orthogonality, phase eliminated by constraint, amplitude mode as scalaron with parity-protected coupling (52); Section 5.
Stiffness assignment $Z_s \propto q_{\phi}^2$, $Z_t \propto q_{\Theta}^2$, (81)	named constitutive assumption	charged-rotor assignment, Section 7.3; motivated by (10), not derived from a response calculation. Enters the dispersion at one point only, the gradient stiffness of the inter-cell kernel.
Phase velocity $c_h = c/2\alpha(m_Z^2) \simeq 64.5\, c$, (79)	conditional prediction	Dirac quantisation plus the charged-rotor assignment, the stiffness normalisation, and the running prescription of Section 7.3.
Witten-dyon sector, $m_{\chi} \gtrsim 2.3$ TeV	schematic completion	motivates the charge weights; q_{ϕ} fixed by quantisation, a dual-charged state mandatory, and the hierarchy $m_h \ll m_W \ll m_{\chi}$ a consequence of the BPS bound; gauge group, breaking pattern, anomaly cancellation and matching not supplied; Section 7.
Medium frame u^{μ}	background datum	unit timelike vector defining the TT projector and the frame of the operator content; not varied and no action supplied; hypersurface orthogonality is an additional restriction, and the comoving condition $u^{\mu}_{\text{matter}} = u^{\mu}_{\text{medium}}$ a separate assumption of the coupling; Section 6.1.
Hypersurface-deformation result, (90), (91)	consistency check	evolution on the projectable foliation class; Section 8.2. No claim of full refoliation invariance away from the acoustic limit.
Causal status of c_h, (96)	derived, scoped	microcausality holds for the parent operators and for every local observable built from them; no such claim is made for the nonlocal TT composite; the c_h cone is the characteristic cone of the constitutive kernel and the signal cone is luminal; Section 9.
Relation to $f(R)$ gravity	interpretive context	vacuum-elasticity analogy and motivated emergent-gravity reading; Section 10. No $f(R)$ action is derived.