

Rigidity of Fourier Multipliers for Universal Euler Convective Cancellations

Rael Nogueira Pires

Abstract

We prove an operator-level rigidity theorem for universal Euler convective cancellations on the three-torus. Let $R(k) : k_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}^3$, $k \in \mathbb{Z}^3 \setminus \{0\}$, be an arbitrary reality-compatible Fourier symbol, with no boundedness, locality, isotropy, self-adjointness, or range assumption. If $\langle Rw, (w \cdot \nabla)w \rangle_{L^2} = 0$ for every real, mean-zero, divergence-free trigonometric polynomial w , then, and only then,

$$R = A + B \operatorname{curl} + \operatorname{curl} B, \quad A, B \in \operatorname{Sym}(3, \mathbb{R}).$$

Thus universal cancellation forces an arbitrary symbol to be the restriction of a constant-coefficient self-adjoint operator of order at most one. If the output is also divergence-free, the family reduces to $R = cI + d \operatorname{curl}$. The zero-mode block is invisible and remains arbitrary when constant fields are admitted. The proof combines phase-polarized triads, analytic propagation through $\mathbb{Z}^3$, and exact finite rank certificates. We also obtain sharp finite test counts under explicit ordering hypotheses, identify the isosceles degeneracy caused by Leray projection, and pull the classification back to universal strain–vorticity cancellations. The certificates are reproduced by exact-arithmetic programs; the Lean 4 formalization and its compiler-trusted finite-rank dependency are documented separately. No application to PDE regularity or existence is claimed.

Contents

1	Introduction	3
1.1	The classification problem	3
1.2	Main results	3
1.3	Method, and what is checked by machine	5
1.4	Relation to prior work	5
1.5	Outline	6
2	Setting, notation and conventions	7
2.1	Fields and symbols	7
2.2	The canonical transverse frame	7
2.3	The pairing	8
2.4	Coordinate invariance of the certificates	8
2.5	Triad conventions	8

3	The theorem package	9
3.1	Unrestricted-output classification	9
3.2	Solenoidal classification and automatic self-adjointness	10
3.3	Pullback to strain–vorticity cancellations	11
3.4	Sharp finite certificates	12
4	Sufficiency: why the displayed generators cancel	13
5	Triad reduction	13
5.1	Isolating a complex triad coefficient	13
5.2	The raw recipient map	14
5.3	Leray projection and the isosceles defect	15
6	Exact finite seeds	16
6.1	The solenoidal seed	16
6.2	The unrestricted-output seed	16
6.3	Certificate summary	17
7	Propagation and proofs of the classification theorems	17
7.1	Solenoidal one-triad propagation	18
7.2	Raw two-triad propagation, and the proof of Theorem A	18
7.3	Coverage check for the propagation routes	19
7.4	Structural comparison	20
7.5	Proof of Theorem D	20
8	Singular values and quantitative rigidity	21
8.1	The raw interaction spectrum	21
8.2	The projected interaction spectrum	21
8.3	Stable and unstable propagation routes	22
8.4	A finite stability estimate	24
9	Strain–vorticity applications	24
9.1	The curl pullback and the proof of Corollary C	24
9.2	Anisotropic longitudinal completion	26
9.3	What does not follow	27
10	Computer-assisted components, reproduction and formalization	28
10.1	What is analytic and what is finite	28
10.2	Reproduction	28
10.3	Formalization and the trust boundary	29
11	Discussion and open problems	31
A	Certificate data	32

1 Introduction

1.1 The classification problem

Work on the three-torus $\mathbb{T}^3 = (\mathbb{R}/2\pi\mathbb{Z})^3$ with normalized Haar measure, and call a field *admissible* if it is a real, mean-zero, divergence-free trigonometric polynomial. The incompressible Euler equation transports such a field by the quadratic convective nonlinearity $(w \cdot \nabla)w$. Conservation of kinetic energy and of helicity are the statements that the two particular operators I and curl are orthogonal to that nonlinearity,

$$\int_{\mathbb{T}^3} w \cdot (w \cdot \nabla)w \, dx = 0, \quad \int_{\mathbb{T}^3} \text{curl } w \cdot (w \cdot \nabla)w \, dx = 0, \quad (1)$$

for every admissible w . We ask for the converse. Let

$$R(k) : k_{\mathbb{C}}^{\perp} \longrightarrow \mathbb{C}^3, \quad k \in \mathbb{Z}^3 \setminus \{0\}, \quad (2)$$

be an arbitrary complex-linear symbol, assumed only to be compatible with reality, and define the cubic functional

$$\mathcal{Q}_R(w) := \int_{\mathbb{T}^3} R w \cdot (w \cdot \nabla)w \, dx. \quad (3)$$

Which symbols satisfy $\mathcal{Q}_R(w) = 0$ for every admissible w ?

The hypotheses omit the assumptions that usually make such a question tractable. The symbol may depend on k arbitrarily; it need not be bounded, homogeneous, isotropic, local, Hermitian, or transverse in its range. Each fiber is an independent complex 3×2 matrix, so the symbol space is infinite-dimensional and has no prescribed structure. Nevertheless, the universal cancellation condition forces the symbol into a twelve-dimensional real family obtained by restricting constant-coefficient, formally self-adjoint operators of order at most one, and a two-dimensional family once incompressibility of the output is imposed. Allowing longitudinal output enlarges the cancellation kernel; imposing the solenoidal range condition collapses it to the energy–helicity family. Within the reality-compatible, translation-invariant Fourier-multiplier class with solenoidal output considered here, the two operators in (1) exhaust the possibilities.

1.2 Main results

Write $C_k v = ik \times v$ for the curl symbol, so that C_k is Hermitian on $k_{\mathbb{C}}^{\perp}$ and $C_k^2 = |k|^2 I$ there. The main classification is the following; it is restated and proved as Theorem A in Sections 3 and 7.

Theorem A. A complex-linear, reality-compatible symbol $R(k) : k_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}^3$ satisfies $\mathcal{Q}_R(w) = 0$ for every admissible w if and only if there are unique constant real symmetric matrices A, B with

$$R(k)v = Av + BC_k v + C_k Bv, \quad v \in k_{\mathbb{C}}^{\perp}, \quad k \neq 0.$$

The result has four useful consequences.

Hermiticity is a conclusion. The pairing in (3) is the real L^2 pairing; in Fourier variables it is the complex-bilinear extension of the Euclidean dot product, pairing the coefficients at k and $-k$. No adjoint is inserted anywhere, and the hypothesis is not that a quadratic form is conserved.

A non-Hermitian symbol has a genuine skew sector which *does* contribute to $\mathcal{Q}_R$, and the phase polarization of Lemma 5.1 is exactly what detects it. Every classified symbol is nevertheless the restriction to $k_{\mathbb{C}}^{\perp}$ of a formally self-adjoint constant-coefficient operator on the full vector-field space. In the solenoidal class this is sharp: Corollary B below concludes $R = cI + d \text{curl}$ with no Hermitian assumption made anywhere.

Order and locality are conclusions. The hypothesis allows nonlocal symbols of arbitrary growth. The conclusion is a differential operator of order at most one with constant coefficients.

The longitudinal directions are cancellations, not invariants. The functional (3) is cubic in w and pairs R against the quadratic convection. When R is self-adjoint on the solenoidal phase space, $\mathcal{Q}_R$ is the Euler derivative of the quadratic functional $\frac{1}{2}\langle w, Rw \rangle$. For a general $R: k_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}^3$ it is only a convective cancellation: longitudinal output can couple to the pressure, so the ten extra directions in the twelve-dimensional family must not be advertised as new Euler invariants. We return to this point in Section 9.3.

The zero mode is invisible, not rigid. If the mean-zero hypothesis is dropped, Theorem A⁰ shows that the classification is unchanged on every nonzero mode while $R(0)$ remains an arbitrary real 3×3 matrix. No field evaluation can see it, because periodic convection has zero spatial mean. The total kernel dimension is then $12 + 9 = 21$.

Imposing incompressibility of the output collapses the family:

Corollary B. If in addition $R(k)k_{\mathbb{C}}^{\perp} \subset k_{\mathbb{C}}^{\perp}$ for every $k \neq 0$, then $\mathcal{Q}_R \equiv 0$ holds if and only if $R(k) = cI + dC_k$ for two global real constants c, d . In particular every nonzero-mode fiber is Hermitian.

Because curl is invertible on mean-zero solenoidal periodic fields, the whole classification transports to the strain-vorticity pairing. Writing $\omega = \text{curl} u$ and $S = \nabla_{\text{sym}} u$, and setting $\mathcal{D}_T(u) = \langle \nabla_{\text{sym}}(Tu), \omega \otimes \omega \rangle$, we obtain:

Corollary C. $\mathcal{D}_T(u) = 0$ for every admissible u if and only if $T = A \text{curl} + B(-\Delta) + \text{curl} B \text{curl}$ with A, B constant real symmetric; with solenoidal output this reduces to $T = c \text{curl} + d(-\Delta)$.

The scalar case $T = -\Delta$ of the second formula gives the Laplacian identity $\langle -\Delta S, \omega \otimes \omega \rangle = 0$. Corollary C says that within the translation-invariant multiplier class this identity and the first-order one are the only universal strain-vorticity cancellations with solenoidal output, and it identifies precisely which longitudinal completions restore cancellation for an arbitrary anisotropic second-order scalar operator (Corollary E, Section 9.2). The scalar identity is due to Miller [15]; the classification here includes it as one member of the larger multiplier family.

Finally, the proof is quantitative. Under two explicit ordering hypotheses on a finite symmetric mode set with M unoriented classes, we prove the existence of test families with exactly $8M - 2$, $4M - 2$ and $12M - 12$ real field evaluations that certify the classification in the complex solenoidal, Hermitian solenoidal and unrestricted-output classes respectively, and we show that no smaller family can work (Theorem D). The machine-readable supplement exports the selected seed rows, pivot columns, mode order and parameter ordering used by the nonsingular-minor certificates. For a coordinate box of radius N these have the optimal $\Theta(N^3)$ size rather than the $\Theta(N^6)$ cost of scanning a dense collection of triads. No result on PDE existence, smoothness, or regularity is claimed anywhere in this paper.

1.3 Method, and what is checked by machine

The proof separates a finite exact calculation from analytic propagation over the lattice.

Sufficiency (Section 4) is a short analytic computation. For necessity, Lemma 5.1 shows that testing $\mathcal{Q}_R$ on real six-mode fields supported on a non-collinear triad $p + q + r = 0$, with total phase products 1 and i , recovers the full complex triad coefficient. Two local lemmas then describe what one triad sees at a fixed recipient mode. In the unrestricted class, one triad annihilates eight of the twelve real parameters of the recipient block and leaves a single complex image line (Lemma 5.2); a second triad with a non-parallel difference direction finishes the job. In the solenoidal class, a single triad with parents of unequal length already determines all eight real parameters of the recipient block, while equal parent lengths leave a four-dimensional residue (Lemma 5.3). This isosceles defect is created exactly by the Leray projection and is invisible before it; Proposition 8.2 computes the responsible singular value in closed form.

The induction that turns those local statements into a global theorem needs a base case that no triad argument supplies, and this is where exact computation enters. Two finite constraint matrices are built and their ranks are certified over $\mathbb{Q}$:

Class	Seed modes	Parameters	Rows	Rank	Nullity
Hermitian solenoidal	7	28	592	26	2
Complex solenoidal	7	56	720	54	2
Complex unrestricted output	13	156	680	144	12

In each case the matching kernel is exhibited explicitly by the classified generators, so rank and nullity together give an exact kernel *equality*, not merely a bound (equations (43) and (47)). Rank is certified from below by a nonsingular minor evaluated modulo two primes, and from above by the explicit kernel. No floating-point rank decision and no numerical tolerance enters at any point: the programs use integers, exact rational arithmetic, or reductions modulo primes. Two of the certificates are reproduced by algorithmically independent implementations that reconstruct the convolution by recipient wavevector rather than reusing the row builders.

The formal development in Lean 4 with Mathlib covers Theorem A and Corollary B together with four supporting components: the arbitrary real zero block in Theorem A⁰; the torus-integral realization of Corollary C and its extension to smooth solenoidal fields; the sharp scalar evaluation counts underlying Theorem D at the finite-dimensional matrix level; and the local singular spectra, conditioning bounds, bad family and stability modulo the exact kernel from Section 8. The source contains no admitted proofs. The two finite seed rank certificates use `native_decide`, which the machine-readable axiom report exposes as seven compiled-code axioms, and the twenty-six-row physical-to-matrix reconstruction for the Hermitian seed is an external exact check rather than a kernel theorem. Section 10 gives the result-by-result trust boundary and records what remains analytic.

1.4 Relation to prior work

The result sits next to several classification literatures. We organize the comparison by the object being classified rather than chronologically, since the hypotheses differ far more than the conclusions do.

Rugged quadratic invariants and Galerkin truncations. That energy and helicity are well-established quadratic invariants of three-dimensional Euler in Fourier-truncated settings. T. D. Lee’s statistical mechanics of Galerkin-truncated ideal flow [14], Kraichnan’s absolute-equilibrium theory and the Kraichnan–Montgomery review [10], Jon Lee’s triad-interaction

representation and isolating constants of motion [12, 11], his later study of mixing on constant energy–helicity surfaces [13], and Shebalin’s monograph [19] form the historical base. Shebalin states that three-dimensional Euler admits no further bilinear integral invariants. None of these formulates an arbitrary mode-dependent matrix symbol, retains longitudinal output, or treats the non-Hermitian sector classified here.

Local conservation laws. Serre classifies the first-order local invariants of three-dimensional Euler [18]; Vukobratovic derives the conservation laws and symmetric forms of the incompressible inviscid-fluid equations [20]; and Cheviakov and Oberlack construct hierarchies of conserved quantities from generalized Ertel theorems [4]. These are differential-density classifications under a local finite-order ansatz. The present hypothesis admits arbitrary nonlocal modewise symbols from the start.

Hamiltonian and Casimir uniqueness. Khesin and Chekanov [9], Enciso, Peralta-Salas and Torres de Lizaola [6], Khesin, Peralta-Salas and Yang [8], and Dullin, Meiss and Worthington [5] classify invariants under much stronger geometric hypotheses — invariance under volume-preserving diffeomorphisms, or the Lie–Poisson structure in Fourier space. Helicity uniqueness there is a statement about regular integral invariants of a group action. Our hypothesis is convective orthogonality alone, and energy is included in the conclusion rather than excluded by the symmetry requirement.

Helical and triad geometry. Waleffe’s helical decomposition makes the energy–helicity constraints on a single triad transparent [21]; Moffatt exhibits further invariants of an isolated triad [16]; Rathmann and Ditlevsen study pseudo-invariants of restricted helical interactions [17]; Alexakis and Brachet analyze fluxes in quasi-equilibrium [1]. Invariants of an isolated or decimated triad system are strictly more numerous than universal ones, which is exactly the distinction that propagation over the full lattice enforces here.

Triad-network nullspaces. Harper, Bustamante and Nazarenko [7] reduce the quadratic invariants of discrete clusters of weakly interacting waves to sparse matrix nullspaces. The broad methodology of computing invariants as nullspaces of triad-interaction matrices is therefore prior art, and we claim no novelty for it. What is different here is the symbol class, the passage from a finite cluster to the entire lattice, and the exact rank structure of that passage.

Kinematic strain–vorticity constraints. Betchov’s homogeneity relation [2] and the proof by Carbone and Wilczek that only two such constraints exist for isotropic turbulence [3] are the closest kinematic antecedents of the pullback in Section 9. The scalar Laplacian strain–vorticity identity is a direct antecedent of Corollary C; our claim concerns the classification and the anisotropic longitudinal completion.

The new results are the all-symbol implication (Theorem A), the automatic Hermiticity in the solenoidal class (Corollary B), and the twelve-dimensional unrestricted longitudinal kernel; the standard facts recalled above — conservation of energy and helicity, their triad manifestations, the helical decomposition, the local conservation-law and Casimir classifications, the triad-network nullspace methodology, and the scalar Laplacian strain–vorticity identity — are not claimed as new.

1.5 Outline

Sections 2–3 fix conventions and state the results. Sections 4–7 prove sufficiency, the triad lemmas, the finite seeds and the lattice propagation. Section 8 gives the singular spectra, Section 9 the strain–vorticity applications, and Section 10 the computer-assisted components and formalization. Section 11 lists open problems.

2 Setting, notation and conventions

2.1 Fields and symbols

Work on $\mathbb{T}^3 = (\mathbb{R}/2\pi\mathbb{Z})^3$ with normalized Haar measure. A *test field* is a real, mean-zero, divergence-free trigonometric polynomial,

$$w(x) = \sum_{k \in \mathbb{Z}^3 \setminus \{0\}} \widehat{w}(k) e^{ik \cdot x}, \quad k \cdot \widehat{w}(k) = 0, \quad \widehat{w}(-k) = \overline{\widehat{w}(k)}, \quad (4)$$

with finite Fourier support. We call such a field *admissible*. For each $k \neq 0$ set

$$C_k v = ik \times v. \quad (5)$$

On $k_{\mathbb{C}}^{\perp}$ the operator C_k is Hermitian and satisfies

$$C_k^2 = |k|^2 I. \quad (6)$$

A symbol R is *reality compatible* when

$$R(-k)\bar{v} = \overline{R(k)v}, \quad v \in k_{\mathbb{C}}^{\perp}. \quad (7)$$

Two symbol classes occur throughout,

$$\mathcal{R}_{\text{raw}} = \{R(k) : k_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}^3\}, \quad \mathcal{R}_{\text{sol}} = \{R(k) : k_{\mathbb{C}}^{\perp} \rightarrow k_{\mathbb{C}}^{\perp}\}; \quad (8)$$

the first allows longitudinal output, the second preserves incompressibility. When the zero mode is admitted we use the convention $0_{\mathbb{C}}^{\perp} = \mathbb{C}^3$, so that reality compatibility says exactly that $R(0)$ is a real 3×3 matrix.

There is no boundedness or growth assumption on $R(k)$. Since every test field is a trigonometric polynomial, all sums below are finite.

2.2 The canonical transverse frame

To make (7) an unambiguous matrix identity, fix a frame convention. For $k \neq 0$ let $\kappa(k)$ be the member of $\{k, -k\}$ whose first nonzero coordinate is positive, let $j(\kappa)$ be the least index with $\kappa \times e_j \neq 0$, and set

$$f_1(\kappa) = \frac{\kappa \times e_{j(\kappa)}}{|\kappa \times e_{j(\kappa)}|}, \quad f_2(\kappa) = \frac{\kappa}{|\kappa|} \times f_1(\kappa). \quad (9)$$

For *both* signed modes k and $-k$ we use the same real orthonormal frame

$$\mathcal{F}(k) = \mathcal{F}(-k) = (f_1(\kappa(k)), f_2(\kappa(k))). \quad (10)$$

Thus $(f_1, f_2, \kappa/|\kappa|)$ is positively oriented at the canonical representative, and compatibility with reality — rather than positive orientation — is what is retained at its negative. In these frames (7) reads $R(-k) = \overline{R(k)}$.

2.3 The pairing

The pairing in (3) is the real L^2 pairing of two real fields. In Fourier coordinates it uses the complex-*bilinear* extension of the Euclidean dot product to pair the coefficient at k with the coefficient at $-k$. No adjoint and no conjugation is silently inserted.

This convention must be fixed before any row calculation, because it is what makes the automatic-self-adjointness statement meaningful. Had the sesquilinear pairing been used, the skew-Hermitian part of R would be annihilated by construction and Corollary B would be a statement about a quotient rather than about the symbol.

2.4 Coordinate invariance of the certificates

The exposition uses the normalized frames (9); the certificate programs use an integer, non-normalized transverse frame, which keeps all arithmetic exact over $\mathbb{Z}$ or $\mathbb{Q}$. The following lemma reconciles them.

Lemma 2.1 (coordinate and generated-row invariance). *Fix a finite geometric family of real triad-supported test fields. Replacing only the transverse frames used to express the symbol blocks preserves the rank, the nullity and the geometric kernel of the resulting constraint matrix.*

Separately, fix finitely many triads and, on each of them, generate the full row space using all tensor-product polarizations from a transverse basis at each mode together with the two total phase classes 1 and i . Replacing those polarization bases preserves this complete generated row space, hence its rank and nullity. No such statement is made for an arbitrary proper subset of the coordinate-generated rows.

Proof. If F_k and G_k are two real bases of $k^\perp$, then $G_k = F_k S_k$ with $S_k \in \text{GL}_2(\mathbb{R})$. For a fixed geometric test family, changing the domain coordinates of a raw block — or both domain and range coordinates of a solenoidal block — right-multiplies the complete constraint matrix by an invertible block-diagonal matrix. This preserves rank and nullity and merely re-expresses the same geometric kernel.

If the test polarizations are also replaced, trilinearity of $(w, w, w) \mapsto \mathcal{Q}_R(w)$ expresses the full tensor-product family in the new bases as an invertible multilinear combination of the full family in the old bases; Lemma 5.1 gives the same statement for the two real phase rows. The two complete generated row spaces therefore agree after the change of coordinates. An arbitrary selected subset of coordinate rows need not be carried to an equally indexed subset, which is why no such claim is included. $\square$

The dense seed matrices of Section 6 use the complete phase-polarization families, so the second part of Lemma 2.1 applies to them. Once particular minor rows have been selected, they are fixed geometric test fields, and changing only the symbol coordinates falls under the first part. The integer script convention therefore certifies the same geometric kernels as the normalized convention used here. Since the changes of coordinates need not be orthogonal, numerical singular values are *not* preserved; all singular values in Section 8 refer to the normalized convention.

2.5 Triad conventions

A *triad* is a relation $p + q + r = 0$ among lattice modes, and it is *non-collinear* when $p \times q \neq 0$. In propagation arguments the new mode is called the *target* and written t ; its known *parents* are

q, r with

$$t = q + r, \quad p = -t, \quad p + q + r = 0. \quad (11)$$

The local lemmas act at the *recipient* $p = -t$ and therefore prove $D(-t) = 0$ first; reality compatibility then gives $D(t) = 0$. Consequently $D(k) = 0$ if and only if $D(-k) = 0$, so every set of known modes below is automatically symmetric under $k \mapsto -k$, and we use this without further comment.

The raw convection received at p is $N = (U \cdot r)V + (V \cdot q)U$; the Leray projection $P_{p^\perp}$ is applied only in the solenoidal output class.

3 The theorem package

3.1 Unrestricted-output classification

Theorem A (unrestricted-output rigidity). *Let $R(k) : k_\mathbb{C}^\perp \rightarrow \mathbb{C}^3$, $k \neq 0$, be a complex-linear, reality-compatible symbol. Then*

$$\mathcal{Q}_R(w) = 0 \quad \text{for every admissible } w \quad (12)$$

holds if and only if there are unique constant real symmetric matrices A, B with

$$R(k)v = Av + BC_kv + C_kBv, \quad v \in k_\mathbb{C}^\perp, \quad k \neq 0. \quad (13)$$

Sufficiency is proved in Section 4, necessity in Section 7.2. Uniqueness is the following self-contained statement, which uses only the signed coordinate-axis modes and is independent of every finite-rank certificate.

Proposition 3.1 (uniqueness of the two matrix coefficients). *If $A, B \in \text{Sym}(3, \mathbb{R})$ satisfy $Av + BC_kv + C_kBv = 0$ for every $k \neq 0$ and every $v \in k_\mathbb{C}^\perp$, then $A = B = 0$. Consequently the pair (A, B) in Theorem A is unique.*

Proof. Evaluate the identity at k and at $-k$, recalling $C_{-k} = -C_k$, and add and subtract. For every $v \perp k$ this gives

$$Av = 0, \quad BC_kv + C_kBv = 0.$$

Taking $k = e_1$ and $k = e_2$ in the first identity shows that A annihilates all three coordinate vectors, so $A = 0$.

For the second identity write $C_k = iK_k$ with $K_kv = k \times v$ and divide by i . With $B = (b_{mn})$ symmetric, the four coordinate evaluations are

(k, v)	$BK_kv + K_kBv$
(e_1, e_2)	$(b_{13}, 0, b_{22} + b_{33})^T$
(e_1, e_3)	$(-b_{12}, -b_{22} - b_{33}, 0)^T$
(e_2, e_1)	$(0, -b_{23}, -b_{11} - b_{33})^T$
(e_3, e_1)	$(0, b_{11} + b_{22}, b_{23})^T$.

Setting these four vectors to zero gives $b_{13} = b_{12} = b_{23} = 0$ and $b_{22} + b_{33} = b_{11} + b_{33} = b_{11} + b_{22} = 0$, whence all three diagonal entries vanish as well. Thus $B = 0$. $\square$

Theorem A⁰ (all-mode completion). *Let $R(k) : k_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}^3$, $k \in \mathbb{Z}^3$, be complex-linear and reality compatible, with $0_{\mathbb{C}}^{\perp} = \mathbb{C}^3$. Then*

$$\mathcal{Q}_R(w) = 0 \quad \text{for every real divergence-free trigonometric polynomial } w \quad (14)$$

holds if and only if there are unique $A, B \in \text{Sym}(3, \mathbb{R})$ such that (13) holds for every $k \neq 0$, while $R(0)$ is an arbitrary real 3×3 matrix.

Proof. Necessity on the nonzero modes follows by restricting (14) to mean-zero fields and applying Theorem A; reality compatibility at $k = 0$ is exactly reality of the matrix entries.

For sufficiency write $w = a + u$ with $a \in \mathbb{R}^3$ constant and u mean zero. Periodicity and $\text{div } u = 0$ give

$$\int_{\mathbb{T}^3} (a \cdot \nabla) u \, dx = 0, \quad \int_{\mathbb{T}^3} (u \cdot \nabla) u \, dx = 0, \quad (15)$$

so the constant vector $R(0)a$ is orthogonal to the complete convective term. On the nonzero modes let $S = A + B \text{curl} + \text{curl } B$. This is a constant-coefficient self-adjoint operator and it commutes with the directional derivative $D_a = a \cdot \nabla$, whence

$$\langle Su, D_a u \rangle = \langle u, D_a Su \rangle = -\langle D_a u, Su \rangle = 0. \quad (16)$$

The remaining term is $\mathcal{Q}_S(u) = 0$ by Theorem A, or directly by Section 4. Thus (14) holds for every w , with no restriction on $R(0)$. $\square$

Three consequences deserve emphasis. The universal kernel has real dimension 12 on the mean-zero phase space and dimension 21 in the all-mode formulation — twelve rigid nonzero-mode parameters and nine invisible zero-mode ones. A completely arbitrary, potentially nonlocal symbol is forced to have constant coefficients and differential order at most one. And admitting constant velocities as tests neither creates new nonzero-mode symbols nor constrains the zero-mode block.

3.2 Solenoidal classification and automatic self-adjointness

Lemma 3.2 (scalarly from lattice directions). *Let $L \in \text{Sym}(3, \mathbb{R})$. If $Lk \parallel k$ for every $k \in \mathbb{Z}^3 \setminus \{0\}$, then $L = cI$ for some $c \in \mathbb{R}$.*

Proof. Applying the hypothesis to e_1, e_2, e_3 makes the coordinate vectors eigenvectors, so $L = \text{diag}(\lambda_1, \lambda_2, \lambda_3)$. Applying it to $e_i + e_j$ gives $\lambda_i = \lambda_j$ for each $i \neq j$. $\square$

Corollary B (solenoidal rigidity). *Under the hypotheses of Theorem A, assume in addition the range condition*

$$R(k)k_{\mathbb{C}}^{\perp} \subset k_{\mathbb{C}}^{\perp} \quad (k \neq 0). \quad (17)$$

Then (12) holds if and only if

$$R(k) = cI + dC_k \quad (18)$$

for two global real constants c, d . In particular the cubic identity forces every nonzero-mode block $R(k)$ to be Hermitian.

Proof. Sufficiency is contained in Theorem A, since $cI + dC_k$ is the case $A = cI$, $B = (d/2)I$ of (13) and preserves $k_{\mathbb{C}}^{\perp}$ by (6). For necessity, apply Theorem A to obtain (13) and impose (17). Let $v \perp k$. Since $C_k Bv = ik \times Bv$ is orthogonal to k , the range condition at k says

$$k \cdot Av + k \cdot BC_k v = 0,$$

and at $-k$, using $C_{-k} = -C_k$ and the same frame (10), it says $k \cdot Av - k \cdot BC_k v = 0$. Adding and subtracting gives

$$k \cdot Av = 0, \quad k \cdot BC_k v = 0 \quad \text{for all } v \in k_{\mathbb{C}}^{\perp}.$$

By (6) the map $C_k : k_{\mathbb{C}}^{\perp} \rightarrow k_{\mathbb{C}}^{\perp}$ is onto, so the second identity says $k \cdot Bw = 0$ for every $w \in k_{\mathbb{C}}^{\perp}$. Symmetry of A and B now converts both statements into

$$Ak \parallel k, \quad Bk \parallel k$$

for every lattice direction k . Lemma 3.2 forces $A = cI$ and $B = bI$, and (13) becomes $R = cI + 2bC_k$. $\square$

In the all-mode formulation of Theorem A⁰ the same argument gives $R(k) = cI + dC_k$ for every $k \neq 0$, while $R(0)$ remains an arbitrary real 3×3 matrix: every constant field is divergence free, so the range condition is vacuous there. Universal cancellation forces Hermiticity on every nonzero fiber, and cannot force it at zero.

Corollary B also admits a direct and sharper proof inside the solenoidal class, using a seven-mode seed and one unequal-length triad per new mode. That route is retained in Sections 6 and 7 because it yields the optimal $8M - 2$ certificate and the projected conditioning law, neither of which follows from the algebraic reduction above.

3.3 Pullback to strain–vorticity cancellations

Let u be admissible and write $\omega = \text{curl } u$, $S = \nabla_{\text{sym}} u$. For a translation-invariant vector multiplier $T(k) : k_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}^3$ define

$$\mathcal{D}_T(u) = \langle \nabla_{\text{sym}}(Tu), \omega \otimes \omega \rangle. \quad (19)$$

Corollary C (complete strain–vorticity classification). $\mathcal{D}_T(u) = 0$ for every admissible u if and only if

$$T = A \text{curl} + B(-\Delta) + \text{curl } B \text{curl}, \quad A = A^T, \quad B = B^T, \quad (20)$$

for a unique pair of constant real symmetric matrices. If T is required to preserve divergence-free output, this reduces to

$$T = c \text{curl} + d(-\Delta), \quad c, d \in \mathbb{R}. \quad (21)$$

The proof is in Section 9.1. As in Theorem A⁰, the mean-zero hypothesis on u may be dropped and $T(0)$ included as an arbitrary real 3×3 block: $\omega = \text{curl } u$ is always mean zero, and $\nabla_{\text{sym}}(T(0)\hat{u}(0)) = 0$. Thus (20)–(21) classify every nonzero block and the zero block is exactly invisible.

3.4 Sharp finite certificates

Let $\Lambda \subset \mathbb{Z}^3 \setminus \{0\}$ be a finite symmetric set of nonzero modes, write $[k] = \{k, -k\}$, and let $M = |\Lambda/\{\pm 1\}|$ be its number of unoriented modes. The signed Boolean seed $\mathcal{B}_7 \cup (-\mathcal{B}_7)$ and the coordinate unit cube Λ_1 are defined in (40) and (45).

Definition 3.3 (unequal-triad ordering). A finite symmetric set Λ has an *unequal-triad ordering* if it contains the signed Boolean seed and its remaining unoriented classes can be ordered as $[t_1], \dots, [t_{M-7}]$ so that at step j there are signed modes q_j, r_j in the seed or in earlier classes with

$$t_j = q_j + r_j, \quad q_j \times r_j \neq 0, \quad |q_j| \neq |r_j|.$$

The full six-mode support $\{\pm t_j, \pm q_j, \pm r_j\}$ of each witnessing triad is required to lie in Λ .

Definition 3.4 (two-triad ordering). A finite symmetric set Λ has a *two-triad ordering* if it contains Λ_1 and its remaining unoriented classes can be ordered as $[t_1], \dots, [t_{M-13}]$ so that at step j there are earlier signed modes q_j, r_j, q'_j, r'_j with

$$t_j = q_j + r_j = q'_j + r'_j, \quad q_j \times r_j \neq 0, \quad q'_j \times r'_j \neq 0, \quad r_j - q_j \nparallel r'_j - q'_j.$$

Again the full six-mode support of both witnessing triads must lie in Λ .

Theorem D (sharp finite certificates). *Let Λ be a finite symmetric set with M unoriented modes.*

1. *If Λ has an unequal-triad ordering, then a general complex solenoidal symbol on Λ has $8M$ real parameters, and there exists a family of exactly $8M - 2$ real triad-supported field evaluations certifying the energy-helicity form (18).*
2. *Under the same ordering hypothesis, there exists a certifying family of exactly $4M - 2$ evaluations in the Hermitian solenoidal subclass, and this cardinality is minimal.*
3. *If Λ has a two-triad ordering, then a general unrestricted-output symbol on Λ has $12M$ real parameters, and there exists a family of exactly $12M - 12$ evaluations certifying (13).*

No scalar test family with fewer evaluations can work in any of these settings, because one real field evaluation supplies at most one real linear equation while the unavoidable kernels have dimensions 2, 2 and 12 respectively.

The theorem makes no claim for a finite symmetric set lacking one of the two orderings. Its proof is in Section 7.5. For every coordinate box $\Lambda_N = \{k \neq 0 : |k_j| \leq N\}$ with $N \geq 1$ both ordering hypotheses hold, and

$$M_N = 4N^3 + 6N^2 + 3N, \tag{22}$$

so that the sharp counts are

symbol class	number of tests
Hermitian solenoidal	$16N^3 + 24N^2 + 12N - 2$
complex solenoidal	$32N^3 + 48N^2 + 24N - 2$
unrestricted output	$48N^3 + 72N^2 + 36N - 12$.

(23)

These families have the optimal $\Theta(N^3)$ size rather than the $\Theta(N^6)$ cost of scanning a dense collection of triads. They classify the nonzero blocks; if a zero-mode block is appended, no field evaluation can constrain its nine real parameters, so sharpness is understood modulo that exactly invisible block.

4 Sufficiency: why the displayed generators cancel

The sufficiency half of Theorem A is analytic and independent of every finite certificate. Let w be admissible and $A = A^T$ constant. Then

$$\mathcal{Q}_A(w) = \int_{\mathbb{T}^3} (Aw) \cdot (w \cdot \nabla)w \, dx = \frac{1}{2} \int_{\mathbb{T}^3} w \cdot \nabla(w \cdot Aw) \, dx = 0, \quad (24)$$

the last step because $\operatorname{div} w = 0$.

For the first-order family put $\omega = \operatorname{curl} w$ and $v = (w \cdot \nabla)w$. Since $\operatorname{div} w = \operatorname{div} \omega = 0$,

$$\operatorname{curl} v = (w \cdot \nabla)\omega - (\omega \cdot \nabla)w. \quad (25)$$

Using self-adjointness of curl and (25), for constant $B = B^T$,

$$\begin{aligned} \mathcal{Q}_{B \operatorname{curl} + \operatorname{curl} B}(w) &= \langle B\omega, v \rangle + \langle Bw, \operatorname{curl} v \rangle \\ &= \langle B\omega, (w \cdot \nabla)w \rangle + \langle Bw, (w \cdot \nabla)\omega \rangle - \langle Bw, (\omega \cdot \nabla)w \rangle. \end{aligned} \quad (26)$$

By symmetry of B the first two terms are together $\int (w \cdot \nabla)(w \cdot B\omega) \, dx$, which vanishes because $\operatorname{div} w = 0$. The third is

$$-\frac{1}{2} \int_{\mathbb{T}^3} \omega \cdot \nabla(w \cdot Bw) \, dx = 0, \quad (27)$$

because $\operatorname{div} \omega = 0$. Hence $\mathcal{Q}_{A+B \operatorname{curl} + \operatorname{curl} B} \equiv 0$ for all constant symmetric A, B , which is the sufficiency half of Theorem A, and by (15)–(16) also of Theorem A⁰.

When $A = cI$, (24) is the kinetic-energy cancellation; when $B = bI$, (26) is the helicity cancellation, with the normalization $B \operatorname{curl} + \operatorname{curl} B = 2b \operatorname{curl}$.

5 Triad reduction

5.1 Isolating a complex triad coefficient

Real test fields carry a reality constraint, so it is not immediate that they see the full complex triad coefficient of a possibly non-Hermitian symbol. They do, and the mechanism is a two-element phase family.

Lemma 5.1 (polarization by the phases 1 and i). *Let $p + q + r = 0$ be a non-collinear lattice triad and choose real transverse polarizations $A_0 \perp p$, $U \perp q$, $V \perp r$. Define the complex triad coefficient*

$$\begin{aligned} \mathfrak{T}_R(p, q, r; A_0, U, V) &= i[R(p)A_0 \cdot ((U \cdot r)V + (V \cdot q)U) \\ &\quad + R(q)U \cdot ((V \cdot p)A_0 + (A_0 \cdot r)V) \\ &\quad + R(r)V \cdot ((A_0 \cdot q)U + (U \cdot p)A_0)]. \end{aligned} \quad (28)$$

For phases $\alpha, \beta, \gamma \in \mathbb{C}$ let w be the real six-mode field determined by $\widehat{w}(p) = \alpha A_0$, $\widehat{w}(q) = \beta U$, $\widehat{w}(r) = \gamma V$ and $\widehat{w}(-k) = \overline{\widehat{w}(k)}$. Then

$$\mathcal{Q}_R(w) = 2 \operatorname{Re}(\alpha\beta\gamma \mathfrak{T}_R(p, q, r; A_0, U, V)). \quad (29)$$

Consequently, if $\mathcal{Q}_R$ vanishes for the two phase products $\alpha\beta\gamma = 1$ and $\alpha\beta\gamma = i$, then

$$\mathfrak{T}_R(p, q, r; A_0, U, V) = 0.$$

Proof. Expanding the Fourier coefficient of $(w \cdot \nabla)w$ and grouping the terms by recipient mode gives the three summands in (28). It remains to check that no other signed zero-sum triple contributes.

Non-collinearity means $p \times q \neq 0$, so p, q, r are nonzero and pairwise non-collinear: if for instance $p \parallel r$, then $q = -p - r \parallel p$, a contradiction. Consider any three modes from $\{\pm p, \pm q, \pm r\}$ summing to zero. If an unoriented mode is repeated, then either an opposite pair cancels and leaves a nonzero third mode, or the relation has the form $\pm 2x \pm y = 0$ or $\pm 3x = 0$; all alternatives are impossible by nonvanishing and pairwise non-collinearity. So the three unoriented modes are p, q, r . Writing their signs as $\varepsilon, \eta, \theta \in \{\pm 1\}$ and using $r = -p - q$ gives

$$(\varepsilon - \theta)p + (\eta - \theta)q = 0,$$

and linear independence of p, q forces $\varepsilon = \eta = \theta$. Hence the only signed zero-sum triples are permutations of (p, q, r) and $(-p, -q, -r)$. The positive triad contributes $\alpha\beta\gamma\mathfrak{T}_R$ and, by reality compatibility, the negative one contributes its complex conjugate; their sum is (29).

Writing $\mathfrak{T}_R = x + iy$, the phase product 1 gives $2x = 0$ and the phase product i gives $-2y = 0$. $\square$

For a fixed polarization triple the eight choices $(\alpha, \beta, \gamma) \in \{1, i\}^3$ used by the certificate programs have total products in $\{1, i, -1, -i\}$; they are redundant up to row sign and span exactly the same two real constraints as the products 1 and i .

Lemma 5.1 is the reason a non-Hermitian symbol is detectable at all: the skew-adjoint sector contributes to (3), even though it would be invisible in a real quadratic form.

5.2 The raw recipient map

The convective output received at the mode p , up to a fixed nonzero scalar, is

$$N_{q,r}(U, V) = (U \cdot r)V + (V \cdot q)U. \quad (30)$$

Lemma 5.2 (raw image line and two-triad determination). *Let $p + q + r = 0$ be non-collinear and suppose D satisfies the universal cancellation (12) and vanishes at q and at r . Then*

$$D(p)(p^\perp) \subset \text{span}_{\mathbb{C}}\{r - q\}. \quad (31)$$

One raw triad therefore has real rank eight on the twelve-parameter target block. If a second known-parent triad at the same recipient has difference direction not parallel to $r - q$, the two triads force $D(p) = 0$ and their stacked target block has real rank twelve.

Proof. Put $z = q \times r$, $Z = |z|$ and $Jx = z \times x$, and write $U = u_0z + \alpha Jq$, $V = v_0z + \beta Jr$. A direct computation gives

$$N_{q,r}(U, V) = Z^2[(\alpha v_0 - \beta u_0)z + \alpha\beta J(r - q)], \quad (32)$$

so that

$$\text{range } N_{q,r} = \text{span}\{z, J(r - q)\}. \quad (33)$$

Since $z \times J(r - q) = -Z^2(r - q)$, the orthogonal complement of this range is exactly $\text{span}\{r - q\}$.

Lemma 5.1, with the two parent terms in (28) equal to zero, gives

$$D(p)A_0 \cdot N_{q,r}(U, V) = 0 \quad (34)$$

for all polarizations, which is (31).

A general complex 3×2 target block has twelve real parameters. One raw triad imposes rank eight and leaves the four-real-dimensional family

$$D(p) = (r - q) \otimes \ell, \quad \ell : p_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}. \quad (35)$$

If a second known-parent decomposition of the same target produces an image line not parallel to $r - q$, the two conditions (31) intersect only at zero, so $D(p) = 0$ and the stacked block has rank twelve. $\square$

5.3 Leray projection and the isosceles defect

In the solenoidal class the recipient is projected:

$$N_p(U, V) = P_{p^{\perp}} [(U \cdot r)V + (V \cdot q)U]. \quad (36)$$

Lemma 5.3 (projected one-triad determination). *Suppose D is solenoidal, satisfies (12), and vanishes at q and at r . If $|q| \neq |r|$, then $D(p) = 0$. If $|q| = |r|$, the target ranks are exactly four for an arbitrary complex solenoidal block and three for a Hermitian block.*

Proof. With the notation of Lemma 5.2, the choices $(U, V) = (z, Jr)$ and $(U, V) = (Jq, Jr)$ give

$$N_p = -|z|^2 z, \quad N_p = |z|^2 P_{p^{\perp}} J(r - q). \quad (37)$$

The second vector has a nonzero in-plane component precisely when

$$J(r - q) \cdot Jp = |z|^2 (|q|^2 - |r|^2) \neq 0. \quad (38)$$

So if $|q| \neq |r|$ the two outputs span $p^{\perp}$ and determine all eight real parameters of an arbitrary complex 2×2 target symbol. If $|q| = |r|$, the projected range is only $\text{span}\{z\}$: one complex output row is annihilated and the other remains free.

Concretely, in an orthonormal output basis whose first vector is parallel to z , the unequal-length case tests both complex output rows of $D(p)$, while the equal-length case tests only the first. The latter gives two complex, hence four real, constraints on a general block. For a Hermitian block the first row consists of one real diagonal entry and one arbitrary complex off-diagonal entry, hence three real constraints. $\square$

The exact local ranks are therefore

	$ q \neq r $	$ q = r $
arbitrary complex solenoidal block	8	4
Hermitian solenoidal block	4	3.

(39)

The raw map (30) has no isosceles degeneracy at all. Equal parent lengths become exceptional only after the longitudinal direction has been discarded by the Leray projection. Section 8 computes the singular value that vanishes in the limit and shows that it is controlled by the logarithmic separation of the two parent shells.

6 Exact finite seeds

The necessity proofs begin with exact finite-dimensional certificates. Each script symbol block is expressed in the integer, non-normalized frame covered by Lemma 2.1. Test fields use two integer transverse basis polarizations at each mode, the phase choices described after Lemma 5.1, and conjugate negative-mode amplitudes. Expanding (3) produces a real rational matrix in the symbol parameters, and Lemma 2.1 transfers its rank and nullity to the normalized convention used here.

6.1 The solenoidal seed

Use the seven unoriented Boolean modes

$$\mathcal{B}_7 = \{e_1, e_2, e_3, e_1 + e_2, e_1 + e_3, e_2 + e_3, e_1 + e_2 + e_3\}. \quad (40)$$

A general complex 2×2 block has eight real parameters, so the seed symbol space has dimension 56. Write $\mathbf{M}_{\mathcal{B}_7}$ for the resulting exact matrix. It has

$$720 \text{ nonzero rows,} \quad \text{rank} = 54, \quad \text{nullity} = 2. \quad (41)$$

A fixed 54×54 minor has determinant residues

$$190,194 \quad \text{and} \quad 891,695,525 \quad (42)$$

modulo 1,000,003 and 1,000,000,007 respectively. All row denominators lie in $\{1, 2, 3, 6\}$, so neither prime annihilates a denominator, and either nonzero residue proves $\text{rank} \geq 54$ over $\mathbb{Q}$. The two linearly independent kernel vectors given by the restrictions of I and C_k prove the matching upper bound. Rank-nullity then yields the exact kernel *equality*

$$\ker \mathbf{M}_{\mathcal{B}_7} = \left\{ \left((cI + dC_k)|_{k_{\mathbb{C}}^\perp} \right)_{k \in \mathcal{B}_7 \cup (-\mathcal{B}_7)} : c, d \in \mathbb{R} \right\}. \quad (43)$$

Consequently every reality-compatible solenoidal symbol annihilating the seed tests agrees on the signed seed with one and only one symbol $cI + dC_k$.

In the Hermitian subclass the same seed gives

$$592 \text{ nonzero rows,} \quad 28 \text{ parameters,} \quad \text{rank} = 26. \quad (44)$$

Since I and C_k are Hermitian, rank 26 and nullity 2 give the same kernel (43) restricted to the Hermitian parameter space, so the Hermitian seed determines the same unique pair c, d .

An exhaustive exact search through all 4,095 subsets of one through six modes of the coordinate unit cube shows that none attains the required Hermitian rank $4M - 2$. Exactly four seven-mode subsets attain rank 26: the nonzero subset sums of $e_1, \pm e_2, \pm e_3$, modulo total sign. This is minimality within a finite domain only, and is not a statement about arbitrary six-mode configurations elsewhere in the lattice.

6.2 The unrestricted-output seed

Use the full coordinate unit cube

$$\Lambda_1 = \{k \in \mathbb{Z}^3 \setminus \{0\} : |k_j| \leq 1\}, \quad (45)$$

which contains thirteen modes modulo total sign. A general complex 3×2 block has twelve real parameters, so the parameter space has dimension 156. Enumerating the 44 non-collinear triads up to ordering and total sign gives an integer matrix

$$\mathbf{M}_1 \in \mathbb{Z}^{680 \times 156}, \quad \text{rank } \mathbf{M}_1 = 144, \quad \text{nullity } \mathbf{M}_1 = 12. \quad (46)$$

The twelve explicit generators obtained from bases of $\text{Sym}(3, \mathbb{R})$ for A and B are linearly independent and annihilate every row, so rank-nullity again gives an exact kernel equality,

$$\ker \mathbf{M}_1 = \left\{ \left((A + BC_k + C_k B)|_{k_{\mathbb{C}}^{\perp}} \right)_{k \in \Lambda_1} : A, B \in \text{Sym}(3, \mathbb{R}) \right\}. \quad (47)$$

Proposition 3.1 uses only the signed coordinate-axis modes, which belong to Λ_1 , so the matrices A, B in (47) are unique. Thus every unrestricted-output symbol annihilating the unit-cube tests agrees there with one and only one generator of the form (13).

A fixed 144×144 minor has residues

$$272,144 \quad \text{and} \quad 290,707,620 \quad (48)$$

modulo the same two primes. An independent convolution-first implementation reconstructs the 680 rows and fixes the canonical digest

$$\text{a4fc3b45a6fc94defd38d992dbbe27f4413254b3af307e41d1c7e25839e9c30d}$$

(SHA-256 of the canonically ordered row data).

The seven-mode Boolean seed does *not* suffice in this larger class: it has 84 parameters, rank 66 and nullity 18. The six additional finite-network freedoms disappear on the full unit cube. This is a concrete illustration of the difference between invariants of a restricted triad network and universal ones.

6.3 Certificate summary

Class	Seed	Parameters	Nonzero rows	Rank	Kernel
Hermitian, solenoidal output	7 modes	28	592	26	2
Complex, solenoidal output	7 modes	56	720	54	2
Complex, unrestricted output	13 modes	156	680	144	12

The construction above records the mathematical generation of these matrices and the modular-minor argument used in the proof. The programs are a reproducible certificate; the finite objects, and the implication from their ranks to the theorems, are stated here explicitly and do not depend on reading the code.

7 Propagation and proofs of the classification theorems

Throughout this section we use the target/parent/recipient convention (11): the new mode is $t = q + r$, the recipient at which the local lemmas act is $p = -t$, and reality transports $D(-t) = 0$ to $D(t) = 0$.

7.1 Solenoidal one-triad propagation

By the kernel identity (43) there is a unique $cI + dC_k$ agreeing with R on $\mathcal{B}_7 \cup (-\mathcal{B}_7)$. Subtract it and call the difference D , so that D is a reality-compatible solenoidal symbol satisfying (12) and vanishing on the fourteen signed seed modes. We show $D \equiv 0$.

Six unequal-length steps recover the missing unit-cube modes. First

$$\begin{aligned} e_1 + e_2 - e_3 &= (e_1 + e_2) - e_3, \\ e_1 - e_2 + e_3 &= (e_1 + e_3) - e_2, \\ e_1 - e_2 - e_3 &= e_1 - (e_2 + e_3), \end{aligned} \tag{49}$$

whose parents have squared lengths 2 and 1; then

$$\begin{aligned} e_1 - e_2 &= (e_1 - e_2 + e_3) - e_3, \\ e_1 - e_3 &= (e_1 - e_2 - e_3) + e_2, \\ e_2 - e_3 &= (e_1 + e_2 - e_3) - e_1, \end{aligned} \tag{50}$$

whose parents have squared lengths 3 and 1. In each equality the left-hand side is the target t and the two right-hand vectors are q, r ; the parents are non-collinear with distinct lengths, so Lemma 5.3 gives $D(-t) = 0$, and reality gives $D(t) = 0$. Permuting coordinates covers the remaining unit-cube classes.

All axis modes follow from

$$2e_i + e_j = e_i + (e_i + e_j), \quad 2e_i = (2e_i + e_j) - e_j, \tag{51}$$

and inductively, for $n \geq 2$,

$$(n+1)e_i + e_j = ne_i + (e_i + e_j), \quad (n+1)e_i = ((n+1)e_i + e_j) - e_j. \tag{52}$$

Again each left-hand side is a target and each right-hand side its known-parent decomposition, with non-collinear parents of distinct lengths.

For a remaining non-axis target t outside the recovered coordinate unit cube, so that $|t|_1 \geq 3$, induct on $|t|_1$. Choose a nonzero coordinate, set $s = \text{sgn}(t_j)e_j$, and write

$$t = (t - s) + s. \tag{53}$$

For example, $t = (2, 1, 0)$ is obtained from $(1, 1, 0) + e_1$; the same construction applies to every non-axis target. Indeed, $t - s$ has ℓ^1 -norm $|t|_1 - 1 \geq 2$, so its squared Euclidean norm is at least 2, while $|s|^2 = 1$. The two parents are non-collinear: t has a nonzero coordinate other than the one used to define s , and that coordinate remains in $t - s$. Thus the parents have distinct lengths and Lemma 5.3 at $p = -t$ determines $D(-t)$; reality determines $D(t)$. Since the parents lie in earlier ℓ^1 -shells or in the recovered seed, this is a genuine induction. It reaches every non-axis mode, and together with the axis steps above proves $D \equiv 0$ and $R = cI + dC_k$. This proves Corollary B directly, without passing through Theorem A.

7.2 Raw two-triad propagation, and the proof of Theorem A

For Theorem A the kernel identity (47) supplies the unique generator (13) agreeing with R on the unit cube. Subtract it and again call the difference D , a reality-compatible symbol with unrestricted output that satisfies (12) and vanishes on $\Lambda_1 \cup (-\Lambda_1)$.

First recover the shell-2 axis modes $2e_i$ by the two decompositions (56) at $n = 2$, whose four parents $e_i \pm e_j$ and $e_i \pm e_\ell$ all lie in Λ_1 ; every non-axis mode in shells 1 and 2 already lies in the coordinate unit cube. Then induct on $|t|_1 \geq 3$. Within each shell, process all non-axis targets before the axis targets.

Let t be a non-axis target with $|t|_1 = n \geq 3$, choose two nonzero coordinates i, j , and set $s_i = \text{sgn}(t_i)e_i$, $s_j = \text{sgn}(t_j)e_j$. Use the two parent decompositions

$$t = (t - s_i) + s_i = (t - s_j) + s_j. \quad (54)$$

Here $t - s_i$ and $t - s_j$ lie in shell $n - 1$ while s_i, s_j lie in the unit cube, so all four parents are already known; each pair is non-collinear because t has another nonzero coordinate. At the common recipient $p = -t$, the two surviving image lines from (31) have directions

$$d_i = 2s_i - t, \quad d_j = 2s_j - t. \quad (55)$$

They are non-parallel. Indeed, if t has a third nonzero coordinate, comparison in that coordinate forces the proportionality factor to be one, whereas the i coordinate would then require $2\text{sgn}(t_i) - t_i = -t_i$, which is impossible. If t has exactly two nonzero coordinates, of magnitudes a, b , then

$$d_i \times d_j = 2(2 - a - b) s_i \times s_j \neq 0$$

because $a + b = |t|_1 \geq 3$. By Lemma 5.2 the two raw triads force $D(-t) = 0$, and reality gives $D(t) = 0$.

Once the non-axis part of shell n is known, take the positive axis target $t = ne_i$ with $n \geq 2$ and use

$$\begin{aligned} ne_i &= ((n-1)e_i + e_j) + (e_i - e_j), \\ ne_i &= ((n-1)e_i + e_\ell) + (e_i - e_\ell), \end{aligned} \quad (56)$$

with $j \neq \ell$ both different from i . Here $(n-1)e_i + e_j$ and $(n-1)e_i + e_\ell$ are non-axis modes of the current shell, already handled, while $e_i - e_j$ and $e_i - e_\ell$ belong to the unit cube. Hence every parent in (56) is known at this stage, including when $n = 2$. The difference directions

$$(2-n)e_i - 2e_j, \quad (2-n)e_i - 2e_\ell$$

are non-parallel, so Lemma 5.2 at $p = -ne_i$ gives $D(-ne_i) = 0$ and reality gives $D(ne_i) = 0$; the negative axis mode follows simultaneously. This completes the induction, so $D \equiv 0$ and (13) holds. Together with Section 4 and Proposition 3.1, Theorem A is proved, and with (15)–(16) so is Theorem A⁰.

7.3 Coverage check for the propagation routes

The case distinctions above can be checked independently of the rank certificates. For the solenoidal route, the bootstrap identities recover the three missing body diagonals and then the three missing face diagonals of the unit cube. The axis identities recover each positive axis in increasing ℓ^1 -order, and the general step recovers every target with at least two nonzero coordinates. Its parents have ℓ^1 -norm one and one less than that of the target, so the induction is well founded. Signs are transported by reality. For the raw route, every non-axis target has at least two nonzero coordinates and therefore admits the two decompositions in (54); after these targets are processed, the decompositions in (56) recover each axis, starting from the shell-2 axes $2e_i$, whose parents all lie in Λ_1 . Thus no target is omitted by either ordering.

As an exact finite audit of these assertions, the verification programs replay both orderings in the coordinate box $|k_j| \leq 10$. The solenoidal replay reaches all 9,260 signed nonzero modes. The raw replay checks all 4,617 target blocks, verifies that both parent pairs are known when used, and checks that the two difference directions have nonzero cross product and that the stacked target block has exact rank twelve in every case. These computations are finite checks of the stated formulas; the preceding inductions extend the argument from the checked box to all of $\mathbb{Z}^3 \setminus \{0\}$.

7.4 Structural comparison

The two propagation arguments are summarized by one geometric dichotomy:

output	one triad leaves	what kills the target	(57)
$\mathbb{C}^3$	one complex image line	a second non-parallel line	
$p^\perp$	no freedom if $ q \neq r $	one unequal-length triad.	

Projection makes generic propagation stronger — one triad suffices instead of two — but it creates the exact isosceles defect, which is the subject of Section 8.

7.5 Proof of Theorem D

Definitions 3.3 and 3.4 isolate exactly what the preceding inductions use.

For an unequal-triad ordering, exact elimination selects 54 independent rows on the complex solenoidal seed. At each later target, Lemma 5.3 supplies eight rows whose restriction to the new target block has rank eight. Ordering rows and columns by propagation makes the matrix block lower triangular, of rank

$$54 + 8(M - 7) = 8M - 2. \quad (58)$$

Restricting the same construction to Hermitian blocks selects 26 seed rows and four rows at each later target, giving rank

$$26 + 4(M - 7) = 4M - 2, \quad (59)$$

which proves clause 2 as the Hermitian specialization of the same triangular construction.

For a two-triad ordering, select the 144 independent unit-cube rows. The two non-parallel image lines at each subsequent target supply a rank-twelve diagonal block, so the total rank is

$$144 + 12(M - 13) = 12M - 12. \quad (60)$$

All selected fields are supported in Λ by the corresponding ordering definition. The energy-helicity plane lies in the kernel of the first two matrices and the twelve symbols (13) lie in the kernel of the third, so the displayed ranks give exactly those kernels. Conversely, one real field evaluation supplies at most one real linear equation, so the parameter dimensions $8M$, $4M$, $12M$ and the kernel dimensions 2, 2, 12 give the matching lower bounds. This proves Theorem D.

For a coordinate box Λ_N , the bootstrap steps (49)–(50), the axis steps (51)–(52) and the induction (53) all remain inside Λ_N and give an unequal-triad ordering; the two-parent construction (54)–(56) likewise remains inside Λ_N and gives a two-triad ordering. Counting unoriented modes gives (22) and hence (23). Finite symmetric sets without one of these orderings lie outside the theorem.

8 Singular values and quantitative rigidity

Set

$$\lambda = |q|^2, \quad \mu = |r|^2, \quad \kappa = |p|^2, \quad z = q \times r, \quad Z = |z|, \quad Jv = z \times v.$$

All singular values in Sections 8.1–8.3 are *local*: they refer to the Euclidean structures induced by orthonormal frames and unit tensor-product polarizations, for the maps

$$\mathbf{L}_{p;q,r}^{\text{raw}}(U \otimes V) = (U \cdot r)V + (V \cdot q)U, \quad \mathbf{L}_{p;q,r}^{\text{sol}} = P_{p^\perp} \mathbf{L}_{p;q,r}^{\text{raw}}. \quad (61)$$

Their singular spectra are invariant under orthonormal changes of basis. Lemma 2.1 explains why the non-normalized certificate coordinates preserve the corresponding ranks but not these numerical values.

8.1 The raw interaction spectrum

Proposition 8.1 (raw singular spectrum). *For the unprojected map (30), the two nonzero squared singular values are*

$$\tilde{\sigma}_n^2 = \frac{Z^2(\lambda + \mu)}{\lambda\mu}, \quad \tilde{\sigma}_d^2 = \frac{Z^2|r - q|^2}{\lambda\mu}. \quad (62)$$

Both are positive for every non-collinear triad.

Proof. The vectors

$$n = \frac{z}{Z}, \quad t_q = \frac{Jq}{Z\sqrt{\lambda}}, \quad t_r = \frac{Jr}{Z\sqrt{\mu}}, \quad t_d = \frac{J(r - q)}{Z|r - q|}$$

are orthonormal frames in the two parent planes and in the raw two-dimensional image. Formula (32), applied first to the unnormalized vectors z, Jq, Jr , gives

$$N_{q,r}(z, z) = 0, \quad N_{q,r}(z, Jr) = -Z^2 z, \quad N_{q,r}(Jq, z) = Z^2 z, \quad N_{q,r}(Jq, Jr) = Z^2 J(r - q).$$

Hence in the ordered tensor-product input basis $(n \otimes n, n \otimes t_r, t_q \otimes n, t_q \otimes t_r)$ and the output basis (n, t_d) , the raw interaction has matrix

$$\begin{pmatrix} 0 & -Z/\sqrt{\mu} & Z/\sqrt{\lambda} & 0 \\ 0 & 0 & 0 & Z|r - q|/(\sqrt{\lambda}\sqrt{\mu}) \end{pmatrix}. \quad (63)$$

Its two rows are orthogonal, and their squared Euclidean norms are exactly the quantities in (62). $\square$

8.2 The projected interaction spectrum

Proposition 8.2 (projected spectrum and target observation). *For the Leray-projected map (36), the squared singular values are*

$$\sigma_n^2 = \frac{Z^2(\lambda + \mu)}{\lambda\mu}, \quad \sigma_t^2 = \frac{Z^2(\lambda - \mu)^2}{\kappa\lambda\mu}. \quad (64)$$

The projection leaves the normal channel unchanged and multiplies the second row singular value by

$$\frac{|\lambda - \mu|}{\sqrt{\kappa} |r - q|}. \quad (65)$$

Let $\mathbf{O}_{p;q,r}(D(p))$ denote the complete local target observation formed from the coefficients $D(p)A_0 \cdot N_p(U, V)$ over orthonormal input polarizations. If, in the orthonormal frame (n, t_p) ,

$$D(p) = \begin{pmatrix} a & c \\ \bar{c} & b \end{pmatrix}, \quad a, b \in \mathbb{R}, \quad c \in \mathbb{C},$$

is Hermitian, then

$$\|\mathbf{O}_{p;q,r}(D(p))\|_F^2 = \sigma_n^2(a^2 + |c|^2) + \sigma_t^2(b^2 + |c|^2), \quad (66)$$

and in particular

$$\|\mathbf{O}_{p;q,r}(D(p))\|_F \geq \min(\sigma_n, \sigma_t) \|D(p)\|_{\text{HS}}. \quad (67)$$

For an arbitrary complex solenoidal block each singular value in (64) occurs with real multiplicity four in $\mathbf{O}_{p;q,r}$.

Proof. Replace t_d in Proposition 8.1 by the normalized tangent vector $t_p = Jp/(Z\sqrt{\kappa})$. The identities (37)–(38), together with

$$N_p(Jq, Jr) = Z^2 \frac{\lambda - \mu}{\kappa} Jp,$$

show that in the same input basis and the output basis (n, t_p) the projected interaction has matrix

$$\begin{pmatrix} 0 & -Z/\sqrt{\mu} & Z/\sqrt{\lambda} & 0 \\ 0 & 0 & 0 & Z(\lambda - \mu)/(\sqrt{\kappa}\sqrt{\lambda}\sqrt{\mu}) \end{pmatrix}. \quad (68)$$

Its rows are again orthogonal, and their squared norms give (64). Dividing the absolute value of the last entry of (68) by the corresponding entry of (63) gives (65), which exhibits the isosceles defect explicitly. Contracting the two orthogonal output rows against the two rows of the Hermitian matrix $D(p)$ gives (66) and hence (67). Without the Hermitian relation each of the two complex rows of $D(p)$ contains two complex coordinates, and separating real and imaginary parts gives real multiplicity four for each singular value. $\square$

These are conditional estimates on one target block with its parent blocks held fixed; Section 8.4 explains why they do not by themselves control the global certificate.

8.3 Stable and unstable propagation routes

Write cond for the ratio of the larger to the smaller nonzero singular value.

Corollary 8.3 (uniform local conditioning of the explicit route). *Every post-seed local projected interaction block in the sparse solenoidal ordering of Section 7.1 satisfies the sharp bounds*

$$\min(\sigma_n^2, \sigma_t^2) \geq \frac{1}{10}, \quad \text{cond} \leq \sqrt{15}. \quad (69)$$

Proof. The three bootstrap triads in (49) have, up to interchanging q and r , $(\lambda, \mu, \kappa, Z^2) = (1, 2, 3, 2)$, hence $(\sigma_n^2, \sigma_t^2) = (3, 1/3)$. The three triads in (50) have $(\lambda, \mu, \kappa, Z^2) = (3, 1, 2, 2)$, hence $(\sigma_n^2, \sigma_t^2) = (8/3, 4/3)$.

The shifted axis modes use $q = ne_i$, $r = e_i + e_j$ with $n \geq 1$. Direct substitution in (64) gives

$$\sigma_n^2 = \frac{n^2 + 2}{2}, \quad \sigma_t^2 = \frac{(n^2 - 2)^2}{2(n^2 + 2n + 2)}. \quad (70)$$

The second expression is at least $1/10$ for integers $n \geq 1$, with equality at $n = 1$ and value $1/5$ at $n = 2$; for $n \geq 3$, comparison with 1 reduces to $n(n^3 - 6n - 4) > 0$, which holds at $n = 3$ and is strictly increasing thereafter. Moreover

$$15(n^2 - 2)^2 - (n^2 + 2)(n^2 + 2n + 2) = (n - 1)(n - 2)(14n^2 + 40n + 28) \geq 0$$

for $n \geq 1$, so $\sigma_n^2/\sigma_t^2 \leq 15$, the cases $n = 1, 2$ being equalities of the displayed inequality.

The completion of the positive axes uses $q = ne_i + e_j$, $r = -e_j$ with $n \geq 2$, giving

$$\sigma_n^2 = \frac{n^2(n^2 + 2)}{n^2 + 1}, \quad \sigma_t^2 = \frac{n^4}{n^2 + 1}. \quad (71)$$

Here the smaller squared singular value is at least $16/5$ and the squared condition number is $(n^2 + 2)/n^2 \leq 3/2$.

Finally consider a general non-axis step (53). It chooses a unit vector r and a parent q with

$$m = |q|^2 \geq 2, \quad a = q \cdot r, \quad d = m - a^2 = |q \times r|^2 \geq 1, \quad 0 \leq a \leq \sqrt{m}.$$

Formula (64) becomes

$$\sigma_n^2 = \frac{d(m + 1)}{m}, \quad \sigma_t^2 = \frac{d(m - 1)^2}{m(m + 1 + 2a)}. \quad (72)$$

The normal value exceeds 1 . If $m = 2$, integrality forces $a \in \{0, 1\}$ and the tangent value is $1/3$ or $1/10$, with squared condition number 9 or 15 . If $m \geq 3$, then

$$\sigma_t^2 \geq \frac{(m - 1)^2}{m(\sqrt{m} + 1)^2} = \left(1 - \frac{1}{\sqrt{m}}\right)^2 > \frac{1}{10},$$

while

$$\frac{\sigma_n^2}{\sigma_t^2} = \frac{(m + 1)(m + 1 + 2a)}{(m - 1)^2} \leq \frac{m + 1}{(\sqrt{m} - 1)^2} \leq 4 + 2\sqrt{3} < 15.$$

The penultimate expression is decreasing in $\sqrt{m} > 1$, so its maximum for $m \geq 3$ occurs at $m = 3$ and equals $4 + 2\sqrt{3}$. This proves (69); both constants are attained in (70) at $n = 1$, equivalently in (72) at $m = 2$, $a = 1$. $\square$

Unequal parent lengths alone do not imply uniform conditioning. For

$$q_n = (n, 0, 0), \quad r_n = (0, n, 1), \quad p_n = -(q_n + r_n),$$

one has $Z^2 = n^2(n^2 + 1)$, $\lambda = n^2$, $\mu = n^2 + 1$ and $\kappa = 2n^2 + 1$, so that (64) yields

$$\sigma_n^2 = 2n^2 + 1, \quad \sigma_t^2 = (2n^2 + 1)^{-1}, \quad \text{cond} = 2n^2 + 1. \quad (73)$$

The parents are never of equal length, yet the conditioning degenerates quadratically. Choosing the propagation route well is therefore a genuine constraint, not a formality.

Finally, since $|\lambda - \mu|/\sqrt{\lambda\mu} = 2|\sinh(\frac{1}{2}\log(\mu/\lambda))|$, taking the square root of the second formula in (64) gives

$$\sigma_t = \frac{2Z}{\sqrt{\kappa}} \left| \sinh\left(\frac{1}{2}\log\frac{\mu}{\lambda}\right) \right|. \quad (74)$$

Observability of the tangential channel is thus governed by the logarithmic separation of the two parent shells.

8.4 A finite stability estimate

Fix the normalized frames of Section 2 and a normalization of the selected real test fields. Let A_Λ be the resulting *global* certificate matrix, let x_R be the real coordinate vector of the complete symbol, and let $K = \ker A_\Lambda$ be the appropriate classified kernel.

Proposition 8.4 (global stability, and its distinction from local conditioning). *If γ_Λ is the smallest positive singular value of A_Λ , then*

$$\text{dist}(x_R, K) \leq \gamma_\Lambda^{-1} \|A_\Lambda x_R\|_2. \quad (75)$$

In propagation order the rows at a new target have the schematic form $y_j = L_j x_j + \sum_{\ell < j} E_{j\ell} x_\ell$, where L_j is the local diagonal block and the $E_{j\ell}$ carry the already encountered parent columns. Full rank of the seed block and of every L_j proves the global rank by block triangularity. It does not bound γ_Λ in terms of $\min_j \sigma_{\min}(L_j)$ alone, because the off-diagonal blocks can accumulate earlier errors during forward substitution.

Proof. Restrict A_Λ to $K^\perp$; its least singular value there is γ_Λ , so $\|A_\Lambda x\|_2 \geq \gamma_\Lambda \|x\|_2$ for $x \in K^\perp$. Applying this to the orthogonal projection of x_R onto $K^\perp$ gives (75). The block formula follows from the propagation ordering: each selected row involves its current target and only earlier parent blocks. Triangularity therefore proves rank, whereas the global least singular value depends on the complete matrix, including the $E_{j\ell}$. $\square$

Thus, for each fixed finite certificate, approximate cancellation on its selected fields quantitatively forces proximity to the classified family. This is a finite-certificate stability estimate and nothing more: the local bounds (69) imply neither an N -independent nor a polynomial bound for γ_Λ^{-1} , and no such global bound is claimed here. Making one is Problem 3 below.

9 Strain–vorticity applications

This section transports the classification through the vorticity and records its exact structural consequence for anisotropic second-order operators. It is deliberately bounded: what follows are identities, not estimates.

9.1 The curl pullback and the proof of Corollary C

Let u be admissible, $\omega = \text{curl } u$ and $S = \nabla_{\text{sym}} u$, and let $T(k) : k_\mathbb{C}^\perp \rightarrow \mathbb{C}^3$ be complex linear and reality compatible. No solenoidal-output assumption is imposed. With $\mathcal{D}_T$ as in (19), symmetry of $\omega_i \omega_j$ gives

$$\mathcal{D}_T(u) = \int_{\mathbb{T}^3} \partial_i (Tu)_j \omega_i \omega_j dx,$$

and since $\partial_i \omega_i = 0$, periodic integration by parts gives

$$\mathcal{D}_T(u) = - \int_{\mathbb{T}^3} (Tu)_j \omega_i \partial_i \omega_j dx = - \langle Tu, (\omega \cdot \nabla) \omega \rangle. \quad (76)$$

Note that $\mathcal{D}_T(u)$ is *not* $\langle Tu, S(u)\omega \rangle$; that is a different cubic form.

On mean-zero solenoidal fields the curl is a bijection: for $k \neq 0$ and $v \perp k$ one has $C_k^2 v = |k|^2 v - k(k \cdot v) = |k|^2 v$ and hence $C_k^{-1} v = C_k v / |k|^2$, and this inverse preserves finite support, reality and transversality. Setting

$$w = \omega, \quad R = T \operatorname{curl}^{-1}, \quad (77)$$

equation (76) becomes the exact equivalence

$$\mathcal{D}_T(u) = -\mathcal{Q}_R(w). \quad (78)$$

The minus sign and the single factor $|k|^{-2}$ are both essential. No growth condition on either symbol is needed, because the tests are finite Fourier sums.

Proof of Corollary C. As u ranges over admissible fields, $w = \operatorname{curl} u$ ranges over all admissible fields, so by (78) the cancellation $\mathcal{D}_T \equiv 0$ is equivalent to $\mathcal{Q}_R \equiv 0$ for $R = T \operatorname{curl}^{-1}$. Theorem A gives $R = A + B \operatorname{curl} + \operatorname{curl} B$ with unique constant real symmetric A, B . Composing on the right with curl and using $\operatorname{curl}^2 = -\Delta$ on the mean-zero solenoidal space gives

$$T = A \operatorname{curl} + B(-\Delta) + \operatorname{curl} B \operatorname{curl},$$

which is (20); uniqueness is inherited from Theorem A and the invertibility of the curl. This is a bijective pullback of the classification, not an extra condition forcing A or B to be scalar.

If T has solenoidal output then so does $R = T \operatorname{curl}^{-1}$, and Corollary B gives $R = cI + d \operatorname{curl}$, hence $T = c \operatorname{curl} + d(-\Delta)$, which is (21). Conversely these operators preserve divergence-free output and satisfy the cancellation. Formula (21) requires the range assumption; (20) does not. $\square$

Taking $T = -\Delta$ in (21) recovers Miller's scalar Laplacian identity [15]

$$\langle -\Delta S, \omega \otimes \omega \rangle = 0. \quad (79)$$

It does *not* give $\langle S, \omega \otimes \omega \rangle = 0$: the field

$$u = 2(0, 1, 1) \cos x + 2(1, 0, 0) \cos y - 2(1, -1, 1) \sin(x + y)$$

is real, mean zero and divergence free, and satisfies $\int \omega^T S \omega = -2$ in normalized measure. Corollary C therefore says that, within the reality-compatible translation-invariant multiplier class on $\mathbb{T}^3$ and for the universal test family, no scalar, anisotropic, matrix-valued, nonlocal or non-self-adjoint solenoidal Fourier multiplier outside $\operatorname{span}\{\operatorname{curl}, -\Delta\}$ creates a new universal strain-vorticity cancellation.

For the quadratic strain energy

$$\mathcal{E}_T(u) = \frac{1}{2} \langle S, \nabla_{\operatorname{sym}}(Tu) \rangle = \frac{1}{4} \langle -\Delta u, Tu \rangle, \quad (80)$$

call T coercive when there is $c_0 > 0$ such that $\mathcal{E}_T(u) \geq c_0 \|u\|_{\dot{H}^2}^2$ for every admissible u . Within the classified solenoidal family, the complete coercive cone on the mean-zero 2π -torus is

$$T = a(-\Delta) + b \operatorname{curl}, \quad a > |b|. \quad (81)$$

Indeed, on a helical Fourier mode the curl eigenvalue is $\pm|k|$, so the contribution to (80) is

$$\frac{1}{4}|k|^2(a|k|^2 \pm b|k|)|\widehat{u}(k)|^2.$$

Because every nonzero lattice frequency satisfies $|k| \geq 1$, uniform positivity relative to $|k|^4|\widehat{u}(k)|^2$ is equivalent to $a > |b|$; necessity is already witnessed by the two helicities at $|k| = 1$. This is algebraic coercivity together with cancellation of one nonlinear channel. It is not a closed PDE estimate.

9.2 Anisotropic longitudinal completion

The unrestricted family (20) answers a question the solenoidal one cannot: given an arbitrary constant-coefficient anisotropic scalar second-order operator, what must be added to it to restore the universal strain–vorticity cancellation?

For a real symmetric matrix H set

$$B_H = \frac{\operatorname{tr} H}{2}I - H, \quad L_H = -\sum_{i,j} H_{ij}\partial_i\partial_j, \quad \Phi_H(u) = \operatorname{div}(B_H u). \quad (82)$$

The correspondence is invertible, with inverse $H = (\operatorname{tr} B)I - B$; the two formulas are different, and $H \mapsto B_H$ is not an involution on $\operatorname{Sym}(3, \mathbb{R})$.

Corollary E (anisotropic longitudinal completion). *Let $H \in \operatorname{Sym}(3, \mathbb{R})$ and put $T_B = B(-\Delta) + \operatorname{curl} B \operatorname{curl}$. Then on divergence-free fields*

$$T_{B_H} u = L_H u - \nabla \Phi_H(u), \quad PT_{B_H} u = L_H u, \quad (83)$$

and B_H is the unique symmetric matrix with this property. Consequently every constant-coefficient anisotropic scalar second-order operator has a unique homogeneous second-order completion of the form T_B inside the classified family, and the completed operator satisfies

$$\langle \nabla_{\operatorname{sym}}(L_H u - \nabla \Phi_H(u)), \omega \otimes \omega \rangle = 0, \quad (84)$$

and the defect of the bare anisotropic operator is exactly longitudinal:

$$\langle L_H S, \omega \otimes \omega \rangle = \langle \nabla^2 \Phi_H(u), \omega \otimes \omega \rangle. \quad (85)$$

Proof. A coordinate expansion for symmetric B gives, for every vector v ,

$$(BC_k + C_k B)C_k v = (k^T H k)v - (Hk \cdot v)k, \quad H = (\operatorname{tr} B)I - B. \quad (86)$$

On $v \perp k$ the left-hand side is $|k|^2 Bv + C_k B C_k v$, which is the symbol of T_B . One must not replace $C_k B C_k v$ by $|k|^2 Bv - k(k \cdot Bv)$: that purported matrix triple-product identity is false. For instance with $H = \operatorname{diag}(1, 2, 3)$, so $B_H = \operatorname{diag}(2, 1, 0)$, and $k = e_1, v = e_2$, one has $C_k B_H C_k v = 0$. Hence the actual full left-hand side of (86) is $B_H C_k^2 v + C_k B_H C_k v = e_2$. Replacing only the

second summand by the incorrect expression would instead make the full left-hand side equal to $e_2 + e_2 = 2e_2$, while the right-hand side of (86) is e_2 . Thus the proposed replacement fails when like quantities are compared.

For transverse v , symmetry and $B_H = (\text{tr } H/2)I - H$ give $k \cdot B_H v = -Hk \cdot v$, so (86) reads $T_{B_H} u = L_H u - \nabla \Phi_H(u)$ in physical space, and applying the Leray projection P kills the gradient. This is (83).

For uniqueness, suppose two completions have the same projection and let G be the difference of the corresponding H matrices. Formula (86) gives $(k^T G k)v = 0$ for every nonzero lattice k and every $v \perp k$. Taking $k = e_i$ forces the diagonal of G to vanish and taking $k = e_i + e_j$ forces its off-diagonal entries to vanish, so $G = 0$; the inverse correspondence then gives equality of the two B matrices.

Identity (84) is Corollary C applied to $T = T_{B_H}$, which is the case $A = 0, B = B_H$ of (20); it is not a direct application of Euler cancellation to T_B . Applying the symmetric gradient to (83) and substituting into (84) gives (85). $\square$

The longitudinal term cannot be discarded from the pairing in (85), because $(\omega \cdot \nabla)\omega$ need not be divergence free; an argument pairing a longitudinal vector against a transverse one is therefore insufficient to eliminate this defect.

Since $\text{div } u = 0$ and H is symmetric, $\Phi_H(u) = -H_{ab}S_{ab}$. Expanding (85) over all symmetric H gives the integrated tensor identity

$$\int_{\mathbb{T}^3} \omega_i \omega_j (\partial_a \partial_b S_{ij} - \partial_i \partial_j S_{ab}) dx = 0, \quad 1 \leq a, b \leq 3. \quad (87)$$

The tensor is symmetric in a, b . Its trace is Miller's scalar identity (79), up to the immaterial overall sign — it is not three times that identity — and its trace-free part supplies five anisotropic relations. For $H = cI$ one has $\Phi_H = 0$ and (85) reduces to (79).

9.3 What does not follow

Two boundaries must be stated explicitly.

The twelve directions are cancellations, not invariants. When R is self-adjoint on the solenoidal phase space, $\mathcal{Q}_R = 0$ is the statement that $\frac{1}{2}\langle w, R w \rangle$ is conserved by Euler. For general $R : k_{\mathbb{C}}^{\perp} \rightarrow \mathbb{C}^3$ the longitudinal output pairs against the pressure gradient in the full Euler equation, so the ten directions beyond $\text{span}\{I, \text{curl}\}$ are not, in general, new Euler invariants. The correct reading of Theorem A is a rigidity statement about the convective term, and that is how it is used in Corollary C.

The anisotropic completion is not a regularity criterion. Identity (85) is exact, and for fixed H the elementary bound

$$|\langle L_H S, \omega \otimes \omega \rangle| \leq \|\omega\|_2^2 \|\nabla^2 \Phi_H(u)\|_{L^\infty(\text{op})}$$

holds; it bounds the H -contracted expression, not every individual component for a single choice of H . The defect is a constant-coefficient third derivative of u , and standard energy control supplies no bound for the displayed norm. It is moreover a pressure-*like* longitudinal potential rather than a nonlinear pressure: Φ_H is linear in u , whereas a pressure term is determined quadratically, and we assert no formula identifying their Hessians. For nonscalar H the completion does not remove the analytic difficulty; it relocates it into the Hessian channel. No new a priori, scale-critical, time-integrable bound for that defect is proved here, and the geometric-depletion program is context and a possible direction, not a theorem of this paper.

10 Computer-assisted components, reproduction and formalization

10.1 What is analytic and what is finite

The proof is not a numerical experiment. Its logical division is as follows.

- *Analytic*: sufficiency of the generators (Section 4); phase separation (Lemma 5.1); the raw and projected triad lemmas (Lemmas 5.2 and 5.3); propagation through $\mathbb{Z}^3$ (Section 7); reduction by the range condition (Corollary B); uniqueness of A, B (Proposition 3.1); the curl pullback and the anisotropic completion (Section 9); and the singular-value formulas (Section 8).
- *Computer-assisted but exact*: the ranks of the two finite seeds, the selection of explicit nonsingular minors, the independent reconstruction of the constraint matrices, the unit-cube minimality search, and a large finite replay of the propagation order.
- *Finite-dimensional linear algebra*: the optimality and stability of the sparse certificates (Theorem D and Proposition 8.4).

No floating-point rank decision and no numerical tolerance enters the certificates. The computations use integers, exact rational arithmetic, or exact reductions modulo primes.

10.2 Reproduction

The exact programs are run from a single entry point. They execute in isolated subprocesses and report the interpreter version, the source revision, the working-tree state and the elapsed time; the run fails if any certificate is missing or any assertion fails. The checks cover

- the 720 solenoidal seed rows, exact rank 54/56, and the two minor residues (42);
- an independent replay of the same rows and of the energy–helicity kernel;
- all 264 signed non-collinear unit-box target triads, with the exact complex ranks 8/4 of (39);
- 5,199 solenoidal propagation steps, cumulatively over the boxes of radii 2, 3, 4, 10 (the radius-ten replay alone has 4,623 steps);
- the 680 unrestricted-output rows, exact modular rank 144/156, the twelve generators, and the two minor residues (48);
- an independent reproduction of the unrestricted rows and of their SHA-256 digest;
- 4,617 raw two-triad target blocks at exact rank twelve, reaching all 9,260 signed modes through radius ten;
- all 4,095 unit-cube subsets of at most six modes in the seven-mode minimality audit;
- sharp Hermitian certificates through radii 1, 2, 3, 4, 10, including the assembled rational rank $246 = 4 \cdot 62 - 2$ at radius two;

- an exact semantic reconstruction of all 26 selected Hermitian Boolean-seed rows from their signed triads, transverse polarizations and phase choices, checked row-by-row against the typed Lean matrix up to nonzero rational scaling;
- 56,064 signed non-collinear triads, cumulatively over the boxes of radii 1, 2, 3, against the exact singular-value formulas (62) and (64);
- the sharp local minimum $1/10$, the sharp maximum condition number $\sqrt{15}$ of (69), and the near-isosceles family (73) through $n = 100$;
- the curl-pullback identity (78), the explicit negative control after (79), and the corrected cancellations, in rational Gaussian arithmetic.

The selected seed rows, pivot columns, mode order and parameter order are exported in the accompanying machine-readable file `certificate-index.json` in the supplement directory. Row digests and modular-minor residues are asserted inside the individual programs and summarized in the repository certificate manifest. Two of the programs are deliberately independent at the algorithmic level: they import neither the corresponding Fourier row builders nor the row reducers, and instead reconstruct the convolution by recipient wavevector. Their agreement therefore checks the row construction through a distinct implementation path. Both are nevertheless written in the same language, so a re-implementation in a second system would strengthen the evidence further.

10.3 Formalization and the trust boundary

Theorem A, its all-mode completion, Corollaries B and C, the matrix-level finite-dimensional content supporting Theorem D, and the local spectral and conditioning formulas are formalized in Lean 4 with Mathlib. The four supporting results added after the original classification have the following precise scope.

- *Theorem A⁰*. Lean classifies every nonzero transverse fiber, identifies reality compatibility at frequency zero with a real 3×3 block, and proves that replacing this block by any real matrix does not change universal all-mode cancellation.
- *Theorem D*. Lean proves the rank-nullity lower bound for real scalar evaluation matrices, constructs sharp abstract coordinate certificates with counts $8M - 2$, $4M - 2$ and $12M - 12$, verifies the 26×28 Hermitian Boolean-seed matrix has nullity two, and checks the ordered-count identities. Lean also proves a general propagation theorem for arbitrary-length unequal-triad orderings of any mode-set size, instantiated at one additional mode beyond the Boolean seed with explicit named test fields whose Euler cubic pairing vanishes against every admissible symbol. The realization of the complete sharp geometric families for every admissible ordering remains the analytic and exact-program argument of Sections 3 and 7.
- *Corollary C*. For trigonometric polynomials, Lean evaluates the velocity and its vorticity as functions on the torus, differentiates them with the Mathlib derivative rather than by symbol manipulation, forms the pointwise contraction $\nabla_{\text{sym}}(Tu) : (\omega \otimes \omega)$, and proves that its normalized integral over $\mathbb{T}^3$ —an iterated interval integral over one full period in each coordinate, reduced by the character orthogonality relation $\int_0^{2\pi} e^{int} dt = 0$ for $n \neq 0$ —equals the finite-Fourier functional, together with the physical-space identity

$\int \nabla_{\text{sym}}(Tu) : (\omega \otimes \omega) = - \int Tu \cdot (\omega \cdot \nabla)\omega$. The classification transfers to that integral. A second layer extends the functional by an absolutely convergent resonant series to the smooth solenoidal fields, those whose Fourier coefficients decay faster than every polynomial, exhibits a member of that class with infinite spectrum, and proves the same classification there by truncation; the extended functional is tied to physical space as the limit of the torus integrals of its truncations. No Sobolev completion at finite regularity is asserted.

- *Section 8.* Lean derives the raw and projected Gram matrices, the multiplicity-four observation spectrum, the sharp $1/10$ and $\sqrt{15}$ bounds for every route family—the shifted-axis steps, the general steps with $m \geq 3$ and the two admissible $m = 2$ steps, and the two bootstrap triads, with the axis formulas obtained by substituting the axis-step parameters into the projected Gram spectrum rather than postulated separately—the quadratically ill-conditioned family, and the stability estimate (75) with no spectral hypothesis: positivity of γ_Λ on $K^\perp$ is derived from finite-dimensionality rather than assumed, and the estimate is instantiated at the Hermitian Boolean-seed certificate itself. The constant is per certificate; no bound uniform in the mode count is formalized, in line with Problem 3.

At the checked revision, the source contains no `sorry` or `admit`, and every declared module elaborates under the pinned toolchain. This does not make every certificate kernel-checked; the trust status splits as follows.

Kernel-checked, depending only on `propext`, `Classical.choice` and `Quot.sound`: the vector identities used throughout; Lemma 3.2 in a slightly stronger form that does not assume symmetry; Proposition 3.1 and the uniqueness of the two generator matrices on all nonzero transverse fibers; the finite-support Fourier definitions of admissible fields, symbols, convolution and the cubic form; reality compatibility of the classified family; the generator cancellations on every triad and the finite-sum argument connecting triad cancellation to the convolutional cubic form, that is, sufficiency for every admissible finite Fourier field including an arbitrary invisible zero-mode block; phase polarization, i.e. Lemma 5.1; the raw-recipient parametrization and the local two-triad determination rule of Lemma 5.2; reality transport of fiberwise vanishing; the propagation theorem in the ℓ^1 -shell order of Section 7.2, with concrete parent-pair certificates for every nonzero lattice mode outside the cube; the two-frame decomposition of a transverse fiber at each of the thirteen seed modes; a non-vacuity witness showing that $R(k) = \|k\|^2 I$ is reality compatible but not a generator, so that neither side of the classification is trivially satisfied; and the *semantic seed bridge*, which proves that a reality-compatible symbol with triad cancellation satisfies every selected constraint of the unit-cube matrix. The bridge is what makes the rank computation a statement about symbols rather than about an arbitrary integer matrix.

Compiler-trusted: the rank-144 claim over $\mathbb{Z}/101\mathbb{Z}$ and the explicit inverse certificate for the selected 144×144 modular minor are discharged by `native_decide`, which evaluates compiled code instead of reducing in the kernel, and therefore trusts the Lean compiler, its runtime and GMP. The formal statements of Theorem A and Corollary B inherit seven such axioms. The precise status is: *the analytic necessity and sufficiency arguments for Theorem A and Corollary B are kernel-checked, and the formal Theorem A holds modulo trusting the Lean compiler for one finite linear-algebra computation.* The certificate is roughly three million modular operations, which is not viable for kernel reduction at present; shrinking it is the route to removing this dependency. The authoritative check is `#print axioms` applied to the formal statements, not any prose list.

The later formalization has three deliberately separated trust layers. First, the all-mode theorem records the arbitrary real zero block; Corollary C is realized as a normalized integral over the three-torus of the pointwise contraction, and extended from trigonometric polynomials to the smooth solenoidal fields; the geometric frame identities, local Gram spectra, multiplicity-four observation, route bounds, bad family and stability modulo the exact kernel are kernel-checked. Second, Lean proves that the typed 26×28 Hermitian Boolean-seed matrix has rank 26 and nullity 2 and proves the sharp $8M-2$, $4M-2$ and $12M-12$ scalar lower bounds at the finite-dimensional matrix level; the seed rank calculation uses the disclosed `native_decide` boundary. Third, the mandatory exact program `verify_solenoidal_seed_semantics.py` reconstructs every selected Hermitian row from an explicit triad-supported Euler test field and compares it with the Lean matrix. This last physical-to-matrix reconstruction is an exact executable bridge, not yet an internal kernel-checked Lean theorem.

The formal development stops at the smooth solenoidal fields. It does not assert a Sobolev completion at finite regularity, a multiplier-domain theorem or PDE solution theory, and on the smooth class the functional is tied to physical space as the limit of the torus integrals of its truncations rather than by an integral of the smooth field itself. Likewise, the Lean matrix-count theorem does not by itself construct the full geometric observation family for every ordering in Theorem D; that construction is the analytic and exact-certificate argument in Sections 3 and 7. These boundaries prevent an abstract coordinate matrix from being mislabeled as a fully internalized physical certificate.

11 Discussion and open problems

Universal orthogonality to three-dimensional Euler convection is highly restrictive. Translation invariance and reality alone force an arbitrary Fourier symbol into a finite-dimensional family of restrictions of constant-coefficient, formally self-adjoint operators: twelve dimensions before the pressure direction is removed and only two on the solenoidal phase space. Allowing longitudinal output creates the larger kernel, while imposing the solenoidal range condition leaves only the energy-helicity family. The isosceles defect is created by the Leray projection and is measured by (74); both statements admit sharp finite certificates verified in exact arithmetic. Through the vorticity, the same theorem classifies the universal Fourier-multiplier extensions of the scalar strain-vorticity cancellation and identifies the precise pressure-like completion that anisotropy requires.

The claims have clear limits. The paper does not address existence, regularity or blow-up. It does not prove a multiplier theorem on $\mathbb{R}^3$, and the zero Fourier block is invisible rather than rigid. No classification of nonlinear, time-dependent, or non-translation-invariant functionals is attempted; the multiplier class considered throughout is translation-invariant by construction.

Open problems include the following.

1. Can the finite seed ranks be replaced by a short symbolic argument that preserves the unrestricted hypotheses, removing the computer-assisted step entirely?
2. Does a six-mode solenoidal seed exist outside the coordinate unit cube? The minimality result of Section 6.1 is confined to that cube.
3. Can the smallest positive singular value γ_Λ of the complete sparse certificate be bounded uniformly, or at least polynomially, in the box radius? Proposition 8.4 shows that the local bounds (69) do not suffice.

4. What is the optimal propagation graph for global conditioning, as opposed to the optimal number of tests?
5. Is there an analogous classification on $\mathbb{R}^3$ under natural regularity assumptions on a continuous symbol?
6. How does the classification change on other compact three-manifolds, or on irrational tori?
7. Can the logarithmic shell gap in (74) be coupled to a dynamical inter-shell decorrelation estimate?
8. Can the longitudinal completion (85) be exploited dynamically without losing control to the Hessian channel?

A Certificate data

The following constants are asserted by the verification programs and summarized in the machine-readable manifest that accompanies the source release.

Quantity	Value
Complex solenoidal seed	7 unoriented modes, 56 parameters
nonzero rows / rank / nullity	720 / 54 / 2
minor residues mod (1000003, 1000000007)	(190194, 891695525)
Hermitian solenoidal seed	7 unoriented modes, 28 parameters
nonzero rows / rank / nullity	592 / 26 / 2
Unrestricted-output seed	13 unoriented modes, 156 parameters
nonzero rows / rank / nullity	680 / 144 / 12
minor residues mod (1000003, 1000000007)	(272144, 290707620)
canonical row digest (SHA-256)	a4fc3b45...39e9c30d
Local projected ranks in the unit box	264 signed triads: 192 of rank 8, 72 of rank 4
Solenoidal propagation replay	5199 steps over radii 2, 3, 4, 10
Raw two-triad propagation replay	4617 target blocks, 9260 signed modes
Unit-cube minimality search	4095 subsets, exactly four minimizers
Local conditioning replay	56064 signed triads over radii 1, 2, 3
Sharp local bounds	$\sigma_{\min}^2 = 1/10$, $\text{cond} = \sqrt{15}$

The full digest is

a4fc3b45a6fc94defd38d992dbbe27f4413254b3af307e41d1c7e25839e9c30d

The solenoidal propagation count is cumulative over the listed boxes; the radius-ten ordering alone contributes 4623 steps. The local-conditioning count is likewise cumulative. The unit-cube minimality statement exhausts all subsets of at most six modes of Λ_1 and finds exactly four seven-mode minimizers; it is not a statement about arbitrary six-mode configurations elsewhere in the lattice.

Data and code availability

The exact-arithmetic certificate programs, the machine-readable certificate manifest and the Lean formalization accompany this paper in the GitHub repository. Every release is archived on

Zenodo under the concept DOI 10.5281/zenodo.22819121, which resolves to the most recent version and lists the version DOI of each individual snapshot; cite the version DOI to pin an exact snapshot. The archived materials are distributed under the Creative Commons Attribution 4.0 International license. The verification entry point and software versions are recorded in the accompanying reproduction documentation.

References

- [1] Alexandros Alexakis and Marc-Etienne Brachet. Energy fluxes in quasi-equilibrium flows. *Journal of Fluid Mechanics*, 884:A33, 2020.
- [2] Robert Betchov. An inequality concerning the production of vorticity in isotropic turbulence. *Journal of Fluid Mechanics*, 1(5):497–504, 1956.
- [3] Maurizio Carbone and Michael Wilczek. Only two Betchov homogeneity constraints exist for isotropic turbulence. *Journal of Fluid Mechanics*, 948:R2, 2022.
- [4] Alexei F. Cheviakov and Martin Oberlack. Generalized Ertel’s theorem and infinite hierarchies of conserved quantities for three-dimensional time-dependent Euler and Navier–Stokes equations. *Journal of Fluid Mechanics*, 760:368–386, 2014.
- [5] Holger R. Dullin, James D. Meiss, and Joachim Worthington. Poisson structure of the three-dimensional Euler equations in Fourier space. *Journal of Physics A: Mathematical and Theoretical*, 52(36):365501, 2019.
- [6] Alberto Enciso, Daniel Peralta-Salas, and Francisco Torres de Lizaur. Helicity is the only integral invariant of volume-preserving transformations. *Proceedings of the National Academy of Sciences*, 113(8):2035–2040, 2016.
- [7] Katie L. Harper, Miguel D. Bustamante, and Sergey V. Nazarenko. Quadratic invariants for discrete clusters of weakly interacting waves. *Journal of Physics A: Mathematical and Theoretical*, 46(24):245501, 2013.
- [8] Boris Khesin, Daniel Peralta-Salas, and Cheng Yang. The helicity uniqueness conjecture in 3D hydrodynamics. *Transactions of the American Mathematical Society*, 375(2):909–924, 2022.
- [9] Boris A. Khesin and Yuri V. Chekanov. Invariants of the Euler equations for ideal or barotropic hydrodynamics and superconductivity in D dimensions. *Physica D: Nonlinear Phenomena*, 40(1):119–131, 1989.
- [10] Robert H. Kraichnan and David Montgomery. Two-dimensional turbulence. *Reports on Progress in Physics*, 43(5):547–619, 1980.
- [11] Jon Lee. Isolating constants of motion for the homogeneous turbulence of two and three dimensions. *Journal of Mathematical Physics*, 16(7):1367–1373, 1975.
- [12] Jon Lee. The triad-interaction representation of homogeneous turbulence. *Journal of Mathematical Physics*, 16(7):1359–1366, 1975.

- [13] Jon Lee. Development of mixing and isotropy in inviscid homogeneous turbulence. *Journal of Fluid Mechanics*, 120:155–183, 1982.
- [14] T. D. Lee. On some statistical properties of hydrodynamical and magneto-hydrodynamical fields. *Quarterly of Applied Mathematics*, 10(1):69–74, 1952.
- [15] Evan Miller. On the interaction of strain and vorticity for solutions of the Navier–Stokes equation. *Pure and Applied Analysis*, 8(1):247–270, 2026.
- [16] H. K. Moffatt. Note on the triad interactions of homogeneous turbulence. *Journal of Fluid Mechanics*, 741:R3, 2014.
- [17] Nicholas M. Rathmann and Peter D. Ditlevsen. Pseudo-invariants causing inverse energy cascades in three-dimensional turbulence. *Physical Review Fluids*, 2(5):054607, 2017.
- [18] Denis Serre. Les invariants du premier ordre de l’équation d’euler en dimension trois. *Physica D: Nonlinear Phenomena*, 13(1–2):105–136, 1984.
- [19] John V. Shebalin. The statistical mechanics of ideal homogeneous turbulence. Technical Report NASA/TP-2002-210783, NASA, 2002. NASA Technical Publication.
- [20] Lubin G. Vulkov. Conservation laws and symmetrization of the equations of incompressible inviscid fluids. *Quarterly of Applied Mathematics*, 57(3):549–560, 1999.
- [21] Fabian Waleffe. The nature of triad interactions in homogeneous turbulence. *Physics of Fluids A: Fluid Dynamics*, 4(2):350–363, 1992.