hexagon:2611.00928v1

CERTIFICATES OF ABSENCE: VALIDATED CONVERSE KAM FOR FLUX SURFACES IN COIL-DEFINED FIELDS

AAPO LIPPONEN

Abstract. Flux surfaces are invariant tori of a divergence-free magnetic field-line flow in three dimensions. We give validated exclusions of C^1 tori that are radial graphs about a specified reference curve. For the distributed 70-coil Wendelstein 7-X (W7-X) filament model in the outward-shifted configuration, a symmetry-free conefield certificate excludes such tori about the validated magnetic axis from a nominal $10\,\mathrm{nm} \times 10\,\mathrm{nm}$ rectangle in a mid-radius resonant island chain. On an exactly symmetrised outward model, a separate certificate excludes graphs defined relative to that model's validated axis through a 20 nm symmetry-line interval. On the symmetrised standard model, a point is certified, and the exclusion at that point holds for every current vector in a restricted five-current box of relative half-width 10^{-8}, relative to a fixed fitted reference curve. An interval-jet Biot–Savart kernel and a Lohner integrator enclose the flow and its tangent; an independent exact-rational checker verifies the witness inequalities. Lean 4 verifies abstract converse-KAM implications, radial-graph transversality and finite conefield algebra, but not the enclosure implementation, the conefield geometric theorem or its identification with the coil field. These local exclusions concern the stated vacuum filament models, not macroscopic island boundaries or engineering current tolerances.

1. Introduction

1.1. **The problem and the principal result.** Field-line plots suggest where magnetic flux surfaces break into islands and chaotic layers, but they cannot show that a surface is absent. This paper shows that a finite validated computation from stellarator coil data can do so for a specified class of surfaces: it produces inequalities whose satisfaction rules out every surface of the class through a chosen set.

Our principal result concerns a nominal $10\,\mathrm{nm} \times 10\,\mathrm{nm}$ rectangle inside the mid-radius 5/5 resonant island chain of the outward-shifted Wendelstein 7-X (W7-X) filament model: no C^1 invariant torus that is a radial graph about the validated magnetic axis passes through any point of the rectangle or of its orbit. The certificate uses the distributed 70-coil file and requires no reversing symmetry. Its area, $10^{-16}\,\mathrm{m}^2$, is set by computational cost;

2020 *Mathematics Subject Classification.* Primary 37J40, 37M21; Secondary 65G40, 37C10, 78A30, 68V15.

Key words and phrases. converse KAM, invariant tori, computer-assisted proof, validated numerics, interval arithmetic, magnetic field lines, stellarator, formal verification.

the island itself is millimetre-scale, and its boundary is not located here. The achievement is a uniform computer-assisted exclusion from coil-defined ODEs, not a practically sized map of absent surfaces.

We use cylindrical coordinates (R, φ, Z): distance from the device's vertical axis, toroidal angle and height. A poloidal section is a plane $\varphi =$ constant, with coordinates (R, Z). A closed reference curve $a(\varphi) = (R_a(\varphi), Z_a(\varphi))$ supplies a moving centre in these sections. Writing $R - R_a = r\cos\vartheta$, $Z - Z_a = r\sin\vartheta$ defines local polar coordinates. The excluded *radial-graph tori* have $r = f(\vartheta, \varphi) > 0$ with f of class C^1: each section meets every ray from the reference exactly once. A *validated axis* here means a locally unique periodic field line, located in an interval-Newton box centred on the numerically located magnetic axis; uniqueness is not global. Converse KAM tests how the linearised flow transports possible tangent directions to such a surface. If validated bounds leave no admissible direction, no surface of the class can pass through the initial point. Failure to certify is inconclusive. Establishing that surfaces persist nearby requires a separate existence argument.

1.2. **Results and contributions.** Table 1 separates the three W7-X results by model and criterion. The *distributed* model uses the coil file as supplied. The *symmetrised* model constructs exact mirror partners for the coils, allowing a criterion that requires exact reversibility. Both outward results use their own model's validated magnetic axis; the standard result uses a single fixed fitted reference curve, including throughout its current box.

TABLE 1. Principal W7-X exclusions. Sizes are nominal; §4 gives decimal bounds and witness data, and the cited release records contain the exact endpoints. Every conclusion concerns C^1 radial graphs about the reference curve specified in the text.

Model	Criterion	Certified set	Reference curve	Record
Outward, distributed	Conefield	10 nm × 10 nm rectangle	Validated axis of this model	[D7]
Outward, symmetrised	Reversible	20 nm interval	Validated axis of this model	[D10]
Standard, symmetrised	Reversible	Point; five-current box, relative half-width 10^{-8}	Fixed fitted curve	[D2]–[D3]

The mathematical contribution is validated converse KAM for a continuous-time flow evaluated directly from coil geometry. Section coordinates about a recorded reference curve remove the need for a global flux function (Theorem 1).

Numerically, a Lohner integrator and an interval-jet Biot–Savart kernel enclose the flow and its tangent (§3). The same machinery carries initial sets and fixed parameter intervals, yielding one conclusion for every member (§2.5).

The applications comprise the W7-X exclusions in Table 1 and a reduced spectral-field example with a rigorously bounded truncation tail (Supplement S6.2). Section 1.4 states how the numerical enclosures, independent checks and formal proofs support these conclusions.

1.3. **Prior work.** *Maps.* For twist maps, converse KAM already has rigorous computer implementations and parameter-family certification. MacKay and Percival [1] developed interval-arithmetic criteria for rotational invariant circles of twist maps, including a conefield formulation. Muldoon [2] treated interval parameters in extended phase space. Stark [3] proved robustness and exhaustiveness under restrictions on the map and its net flux, without a finite computation bound; no corresponding exhaustiveness result is known for three-dimensional flows. For the standard map, rigorous non-existence [4] and persistence [5] closely bracket the transition.

Flows. MacKay [6] set out the programme of supplying rigorous trajectory and tangent bounds, demonstrating it by hand for a closed-form Hamiltonian. Rigorous non-existence for flows also exists without interval arithmetic, for geodesic flows on tori with large bumps, an idea of Bangert [7, §9] rediscovered for plasmas by Kaiser and Salat [8] (see [9]). MacKay's later criterion [10] excludes invariant surfaces transverse to a chosen direction field, and Martínez-del-Río and MacKay [11] supply the symmetry-free conefield theorem used here. Applications of these criteria so far use floating-point integration: they study the choice of direction field [12], magnetic field lines [13] and the restricted three-body problem [14]. On the existence side, Valvo and Locatelli [15] apply validated a posteriori KAM methods to a field-line model.

The gap. Validated exclusion for vector fields defined by realistic coil geometry requires flow and tangent enclosures computed directly from realistic coil data; that is what this paper supplies. Existing stellarator tools answer related but different questions. Coil-tolerance studies use shape gradients [16], stochastic optimisation [17], Hessian rankings [18, 19], current sampling [20] and adjoint island sensitivities [21]; coil currents can also enlarge the edge islands substantially [22]. Fast diagnostics such as weighted Birkhoff averages [23, 24], alignment indices [25], turnstile fluxes [26] and persistent homology [27] locate candidate regions at much lower cost. Validated exclusion complements them with a finite sufficient test whose numerical error is controlled.

1.4. **Scope of verification.** An audited, outward-rounded interval implementation encloses the flow and its tangent; an independent exact-rational checker checks the finite inequalities; and a mathematical criterion supplies non-existence (Figure 1). The local justification uses only the hypothetical

torus's tangent planes along the certified finite segment (§2.2). Lean 4 proves the abstract general-point and reversible implications, radial-graph transversality and finite conefield algebra under explicit enclosure and field/flow/torus hypotheses. The conefield implication proved below, coil-field identification and checker-to-predicate correspondence remain outside Lean. Supplement S7.11 records the formal inventory and the formalisation's stronger, complete-flow hypothesis. No W7-X enclosure has been reproduced by a second validated integrator: the coil enclosures rest on one implementation, cross-checked by GPU/host identity and Arb spot checks; the exact-rational checker verifies witness algebra, not the enclosures.

Bracketed identifiers such as [D7] name data records in the archived package; Supplement S4 and the package manifest identify their files and SHA-256 digests. Throughout, $[\cdot,\cdot]$ denotes a validated enclosure. Floating-point ranges are written "from a to b" and labelled as such.

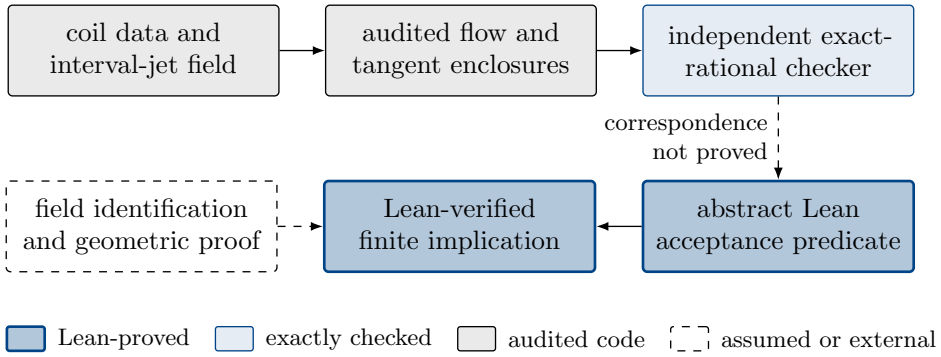


FIGURE 1. What is proved and what is trusted in a certificate. Boxes are filled by trust level: audited interval code produces the enclosures and finite inequalities; the independent exact-rational checker re-verifies the witness algebra exactly; Lean proves the abstract implications under explicit hypotheses; the field identification and the geometric proof are external to Lean. Dashed arrows mark steps that are assumed, not proved: the executable checker has no proved equivalence to the Lean predicate. Enclosure soundness is assumed by the checker and by Lean throughout.

2. The criterion, and what a certificate means

2.1. **Geometry and reference curves.**

Constructing the section chart. At each φ, translate the section origin to the recorded curve and write

$$R = R_a(\varphi) + r\cos\vartheta, \qquad Z = Z_a(\varphi) + r\sin\vartheta, \qquad \psi = r^2/2.$$

This chart is regular for $r > 0$, with angular coordinates understood locally modulo 2π. Here ψ measures distance from the reference curve; it is unrelated

to any flux function of the field. A positive radial graph is $r = f(\vartheta, \varphi) > 0$; its section tangent is $f_\vartheta e_r + f e_\vartheta$, whereas $\xi = r e_r$. Their section determinant is $rf = f^2 > 0$ on the graph, proving transversality without assuming that the reference curve is a field line. Using a validated axis fixes a physically distinguished reference; using a fitted curve defines a different, explicitly recorded graph class. Noncircular or elongated sections can satisfy the radial-graph condition. Surfaces with folds that give multiple intersections, island surfaces surrounding a separate elliptic periodic point (an O-point), and other topologies are outside the class, and a certificate says nothing about whether they exist.

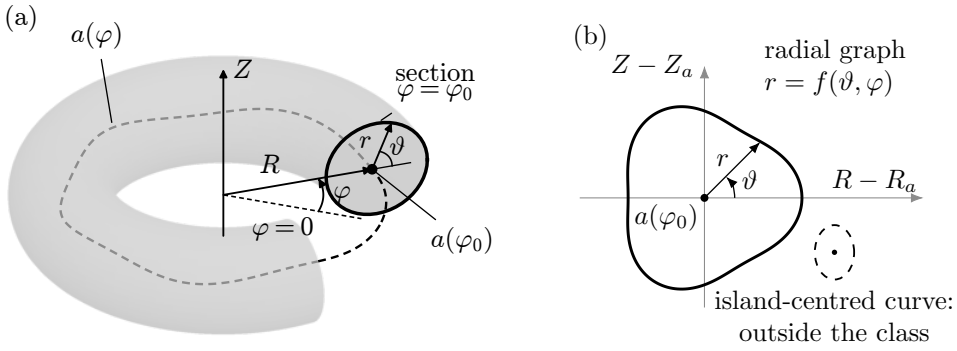


FIGURE 2. The coordinates and the class of tori that a certificate excludes (schematic, not to scale). (a) A toroidal domain cut open at the section plane $\varphi = \varphi_0$: the closed reference curve $a(\varphi)$ (dashed) pierces the cut face at $a(\varphi_0)$; R is the distance of that point from the Z axis, φ the toroidal angle from the dashed ray $\varphi = 0$, and (r, ϑ) are polar coordinates about $a(\varphi_0)$ in the section. (b) One section. A positive C^1 radial graph meets each ray once; the radial direction $\xi = r e_r$ is transverse to it. A small island-centred curve (dashed) does not enclose the reference and is outside this class. Neither drawn curve is asserted to be invariant.

Changing the reference changes the graph class; Theorem 1 requires an C^1 closed curve, not a magnetic axis; the fitted and validated curves used here are analytic. The certificates use stellarator-symmetric Fourier fits with 48 harmonics for the standard configuration and 24 for the outward configuration; the floating-point fit residuals are below 9×10^{-8} m (Supplement S1.2). The standard results retain this fixed *fitted curve*.

The outward results use each model's own *validated axis*. An interval-Newton calculation proves a full-2π periodic orbit, unique within a box of half-width 2×10^{-7} m (§4.1). Uniqueness is local to that box; the name "axis" records the choice of Newton box centred on the floating-point axis, not a continuation calculation. For the symmetrised model, exact reversibility and local uniqueness also prove that this orbit is symmetric.

TABLE 2. Notation used in the criterion and its coil implementation. The spectral coordinate ψ is a flux coordinate; its coil counterpart is defined geometrically.

Symbol	Meaning
B, B_φ	Magnetic field and its physical toroidal component; $B^\varphi = B_\varphi/R$.
V, φ	Field normalised to toroidal-angle time, $V = RB/B_\varphi$; $V^\varphi = 1$.
(R_a, Z_a)	Recorded reference curve; a validated magnetic axis only where stated.
ψ, ϑ	Coil section coordinates: half the squared reference distance and its polar angle.
ξ	Direction transverse to the excluded torus class; radial about the reference in the coil chart.
η	Transported displacement satisfying the variational equation; its initial direction depends on the criterion.
β, λ	Flux two-form and auxiliary coorientation one-form; their scalar tests appear below.
ε, $[\cdot, \cdot]$	Fixed parameter vector and validated enclosure, respectively.

The field-line flow uses toroidal angle φ as time, with $V^\varphi = 1$. One *transit* is an advance of 2π. The spectral single-displacement system has state $(\psi, \vartheta, \eta^\psi, \eta^\vartheta)$; the reversible coil runs instead integrate (R, Z, η_R, η_Z). Conefield runs carry both section tangent columns: mathematically two state coordinates and four tangent components, implemented as two four-dimensional coil systems advanced in lockstep.

2.2. **The criterion for a field given as coils.** The construction needs no known global flux function. Its right-hand side is a sum over thousands of filament segments, evaluated with validated interval arithmetic in §3.2. The suspension adjoins the periodic time coordinate φ to the section state: in the order (R, φ, Z) its vector field is $(V^R, 1, V^Z)$.

Let B be a smooth divergence-free field with $B_\varphi \neq 0$ on an open set $U \subset \{R > 0\}$ of the suspended space, so that field lines in U obey $\mathrm{d}R/\mathrm{d}\varphi = RB_R/B_\varphi$, $\mathrm{d}Z/\mathrm{d}\varphi = RB_Z/B_\varphi$ with φ as time. Let $(R_a(\varphi), Z_a(\varphi))$ be *any* C^1 closed reference curve, and set $\psi = \frac{1}{2}[(R - R_a)^2 + (Z - Z_a)^2]$ and choose the section-radial direction field $\xi = \nabla_{R,Z}\psi = (R - R_a)\,\partial_R + (Z - Z_a)\,\partial_Z$, with zero toroidal component. Write B_φ for the physical toroidal component of B, so that the contravariant component in the coordinates (R, φ, Z) is $B^\varphi = B_\varphi/R$. With φ as time the invariant section density is $\mu = RB^\varphi = B_\varphi$, and set β to the flux form with this fixed section orientation. The section coorientation extends to the suspension as $\mathrm{d}_{R,Z}\psi - \mathrm{d}_{R,Z}\psi(V)\,\mathrm{d}\varphi$; it annihilates $(V^R, 1, V^Z)$ and is positive on ξ away from the reference.

Theorem 1 (general-point and reversible criteria in section coordinates; [13, Theorems 3.1 and 3.2])**.** *For the field and C^1 reference curve just defined, work away from $\xi = 0$ and suppose the exact orbit and tangent passage maps*

exist on the finite segment used. For a section displacement η the flux and coorientation tests are

$$\beta(\eta,\xi) = B_\varphi[\eta_R(Z-Z_a) - \eta_Z(R-R_a)], \tag{1}$$

$$\lambda(\eta) = \eta_R(R-R_a) + \eta_Z(Z-Z_a). \tag{2}$$

For a displacement with a flow component, these formulas apply to its section representative; both forms annihilate the flow direction. Either of the following exact-flow conditions excludes every C^1 invariant radial-graph torus about the reference through s_0 or any point of its orbit.

General-point criterion. *Start the transported displacement at $\eta(0) = c\xi(s_0) + dB(s_0)$, with $c > 0$. The scalar $\beta(\eta,\xi)$ has a strict sign change on a segment on whose whole bracketing window $\lambda(\eta) < 0$.*

Reversible (symmetry-line) criterion. *The field obeys $B(\sigma x) = -D\sigma B(x)$ for $\sigma(R,\varphi,Z) = (R,-\varphi,-Z)$, and the reference satisfies $R_a(-\varphi) = R_a(\varphi)$, $Z_a(-\varphi) = -Z_a(\varphi)$, so $\xi(\sigma x) = D\sigma\xi(x)$. The initial point lies on the fixed line, $\eta(0)$ is antisymmetric under σ and independent of $B(s_0)$, and $\beta(\eta,\xi)$ has a strict sign change at positive time. No λ condition is needed. The guard $B_\varphi \neq 0$ holds on the tested segment and its reflection; the identity $B_\varphi \circ \sigma = B_\varphi$ carries the guard to the latter.*

Proof. The cylindrical volume form is $\Omega = R\,\mathrm{d}R \wedge \mathrm{d}\varphi \wedge \mathrm{d}Z$. The restriction of $\iota_B\Omega$ to a section has coefficient $-B_\varphi$; we use its negative for β in (1). Divergence-freeness gives

$$\partial_\varphi B_\varphi + \partial_R(B_\varphi V^R) + \partial_Z(B_\varphi V^Z) = 0,$$

so the section passage maps transport this form, with no extra factor of R. The stated extension of λ annihilates B and has $\lambda(\xi) = |\xi|^2 > 0$.

Suppose a C^1 invariant radial graph passes through s_0. Its tangent planes have a continuous coorientation ν with $\nu(\xi) > 0$; invariance transports these planes along the finite segment. In the general-point case, $\nu(\eta(0)) > 0$, and invertibility of the tangent maps prevents a transported nontangent vector from becoming tangent. Thus $\nu(\eta) > 0$ throughout. At the zero supplied by the sign change, the kernel of β is the field direction, so $\eta = c_t\xi + d_tB$. Since $\lambda(\eta) < 0$, one has $c_t < 0$, contradicting $\nu(\eta) = c_t\nu(\xi) > 0$.

In the reversible case, reflection relates the tangent vectors at the positive zero time t_* and its negative: if $\eta(t_*) = c_t\xi + d_tB$, then $\eta(-t_*) = -c_t\xi + d_tB$. Independence of $\eta(0)$ and $B(s_0)$, preserved by the tangent maps, gives $c_t \neq 0$. A tangent initial vector would remain tangent, already contradicting this expression and transversality. Otherwise $\nu(\eta)$ cannot vanish, yet its values at the two ends have opposite signs. This is the required contradiction. Only C^1 regularity of the reference, not analyticity, enters the direction field and coorientation argument. □

Validated enclosures supply these exact-flow inequalities; they are not additional hypotheses of the geometric criterion.

Two remarks on the φ-dependence. The subscript in $\nabla_{R,Z}$ is essential: ξ is the radial direction within each poloidal section, with coefficients depending on φ; it is not the full three-dimensional gradient of ψ. This is exactly the direction transverse to radial graphs about the moving curve. And since ξ vanishes on the reference curve itself, β carries information only away from it: the working annulus excludes the curve, and every certified orbit here stays bounded away from it (for the W7-X probes, about 0.3 m).
Where $B_\varphi \neq 0$ is needed. The contradiction is derived on a *finite orbit segment* of the probe: an invariant torus through the probe contains the segment, and only its tangent plane along the segment is used, whatever the torus does elsewhere. So the criterion requires $B_\varphi \neq 0$ only on a neighbourhood of the certified segment, which is the domain guard the runs enforce: one strict sign of B_φ on every step's a-priori enclosure, covering all intermediate times. The usual hypothesis on a whole annulus is stronger than needed. If a step box cannot exclude zero from a filament-distance denominator, the interval kernel raises a failure flag; the coil drivers reject that step and halve its length rather than accept an enclosure across a filament. Supplement S7.9 gives this finite-segment argument; it is not formalised (§1.4).

2.3. **The symmetry-free conefield criterion.** The criterion transports both tangent columns, the images of ∂_R and ∂_Z at s_0, in both time directions. For the displacement η_v started at $\eta_v(0) = v$, a combination of the two columns, put $\widehat{\alpha}_t(v) = -\beta(\eta_v(t), \xi)/|B_\varphi|$, that is $-\operatorname{sgn}(B_\varphi)$ times the bracket of (1) (the omitted factor $|B_\varphi|$ is positive), and let $a_t = \widehat{\alpha}_t(\xi_0)$ and $b_t = \widehat{\alpha}_t(\xi_0^\perp)$, where $\xi_0 = \xi(s_0)$ and $\xi_0^\perp$ is ξ_0 turned through a right angle in the sense fixed by the sign of B_φ. A candidate section tangent line transverse to ξ_0 can be written as $\operatorname{span}(\xi_0^\perp - \varsigma\xi_0)$; thus ς parametrises the possible tangent line, not a spatial coordinate. A slope ς is *admissible* if $\varsigma\, a_t \leq b_t$ at every time t of either direction. The times at which a_t has a certified sign form two families, T_+ ($a_t > 0$, where the inequality bounds ς above) and T_- ($a_t < 0$, where it bounds ς below). Two times $t_\pm \in T_\pm$ with $b_+a_- - b_-a_+ > 0$, where $a_\pm = a_{t_\pm}$ and $b_\pm = b_{t_\pm}$, prove that no admissible slope exists: the cone is empty. The two times with their enclosures form the certificate's *witness*. The test is division-free, and its acceptance predicate is formalised in Lean. That a torus of the class through s_0 yields an admissible slope is the geometric content of the following criterion; we give its finite-segment proof in the setting used here.

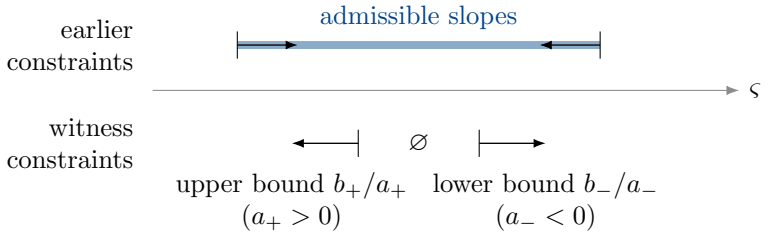


FIGURE 3. How a witness empties the cone of admissible slopes (schematic). Each signed step bounds the same candidate tangent slope ς. Intersecting more constraints shrinks the interval of admissible slopes (blue); the witness upper bound lies below the witness lower bound, leaving no admissible slope. Positions are illustrative, not numerical data.

Theorem 2 (conefield criterion; [11, Theorem 3.1])**.** *Let B be a C^1 field with $B_\varphi \neq 0$ on a domain, and let Φ_t be the local flow of the suspended field $V = RB/B_\varphi$, with $V^\varphi = 1$, along an orbit $x(t)$, defined for $t \in [-\tau, \tau]$. Let ξ be a C^1 direction field and h an auxiliary section field, with $\Omega(h, \xi, B) > 0$ for a volume form Ω. Put $\xi_0 = \xi(x_0)$, $h_0 = h(x_0)$ and*

$$\alpha_t(v) = \Omega_{x(t)}\big(\xi(x(t)), B(x(t)), D\Phi_t v\big), \qquad q_t = \frac{\alpha_t(h_0)}{\alpha_t(\xi_0)}$$

where the denominator is nonzero. If there are times $t_\pm$ with $\pm\alpha_{t_\pm}(\xi_0) > 0$ and $q_{t_+} \leq q_{t_-}$, then no C^1 invariant torus transverse to ξ passes through x_0.

Proof. Suppose such a torus T passes through x_0. Its tangent plane contains $B(x_0)$. Since $\Omega(h_0, \xi_0, B(x_0)) > 0$, this plane meets $\operatorname{span}\{h_0, \xi_0\}$ in a line, not in $\operatorname{span}\{\xi_0\}$ by transversality. Hence the line is spanned by $u = h_0 - s_*\xi_0$ for a unique s_*. The same nonzero volume determinant shows that u is independent of $B(x_0)$.

Invariance puts the existing finite orbit segment in T, and $D\Phi_t u$ in its tangent plane. The tangent maps are invertible and carry the field direction to the field direction: indeed $V = RB/B_\varphi$ and $D\Phi_t V(x_0) = V(x(t))$. Thus $B(x(t))$ and $D\Phi_t u$ remain independent and span that plane. The value $\alpha_t(u) = \Omega(\xi, B, D\Phi_t u)$ can vanish only if $\xi(x(t))$ lies in the same plane, contrary to transversality. It is continuous and nowhere zero on $[-\tau, \tau]$. At zero, $\alpha_0(\xi_0) = 0$ and the cyclic permutation gives $\alpha_0(u) = \Omega(\xi_0, B(x_0), h_0) = \Omega(h_0, \xi_0, B(x_0)) > 0$. Consequently

$$\alpha_t(h_0) - s_*\alpha_t(\xi_0) > 0 \quad \text{for all } t \in [-\tau, \tau].$$

The positive and negative denominator cases give respectively $s_* < q_t$ and $s_* > q_t$. At the two witness times these require $q_{t_-} < s_* < q_{t_+}$, contradicting $q_{t_+} \leq q_{t_-}$. □

The volume form is used only to define α and its orientation; this local argument does not require volume preservation. The coil application remains divergence-free, as required for the flux-form identities above. No completeness outside the stated finite segment is assumed. Its hypotheses are met as follows.

(1) *Field and time.* B is the filament field, analytic and divergence-free on the coil-free working domain. The runs use φ as time; the resulting change of passage time adds to the transported vector only a multiple of B, which the covector $\iota_B \iota_\xi \Omega$ annihilates, so the pulled-back covector is unchanged. The domain guard $B_\varphi \neq 0$ on every step justifies the reparametrisation, and backward time integrates $-V$, not reversed currents.

(2) *Direction and horizontal fields.* ξ is the section-radial field of Theorem 1, analytic away from the reference curve, from which every run stays bounded away. h is ξ turned through a right angle in the section with the sense set by $\operatorname{sgn} B_\varphi$, with $h(x_0) = \xi_0^\perp$, which gives $\Omega(h, \xi, B) = |B_\varphi|\,|\xi|^2 > 0$.

(3) *Torus class.* A C^1 radial graph about the reference curve is transverse to ξ (the local computation is proved in Lean, §2.1), so the theorem excludes the C^1 members of the radial-graph class; the Lipschitz extension of [11, §5] is not used.

(4) *Inequalities and orientation.* With the section orientation of (1), $\widehat{\alpha}_t = \alpha_t / |B_\varphi(x(t))|$. Taking $h = \operatorname{sgn}(B_\varphi)(-\xi_Z, \xi_R)$ in section coordinates gives $a_0 = 0$ and $b_0 = |\xi_0|^2 > 0$. This initial positivity fixes the positive, rather than negative, rescaling. With $a_t = \widehat{\alpha}_t(\xi_0)$ and $b_t = \widehat{\alpha}_t(h_0)$, the witness's $b_+ a_- - b_- a_+ > 0$ is $q_{t_+} < q_{t_-}$, strictly stronger than the theorem requires.

Lean uses `ConeCriterionInput` as a hypothesis ("a torus through s_0 yields an admissible slope") and proves the finite implication from it. Both conefield results below rest on Theorem 2.

2.4. **The general-point and reversible alternatives.** The reversible criterion requires exact reversibility under $\sigma(R, \varphi, Z) = (R, -\varphi, -Z)$. Neither distributed coil file has this property. We therefore use two routes: the symmetry-free conefield criterion on the distributed outward model, and the reversible criterion on models whose mirror coils are generated exactly. Symmetrisation changes the model; its construction and segment-level audit are given in Supplement S7.5. The reversible runs cannot instead be read as general-point certificates, because their initial tangent ∂_Z is not the required radial tangent; the same supplement gives an explicit counterexample.

The corresponding symmetry-line runs on the distributed outward file ([D11], Table S6) satisfy every recorded inequality, including against the validated axis, but establish no exclusion: the criterion they use requires exact reversibility, which that file lacks.

Here $\beta(\eta, \xi)$ is the invariant 2-form of the flow evaluated on the transported displacement and the direction field. The auxiliary 1-form λ satisfies $\lambda(B) = 0$ and $\lambda(\xi) > 0$; for $\xi = \nabla\psi$ one may take $\lambda = \mathrm{d}\psi - kB^\flat$, with $k = B \cdot \nabla\psi/|B|^2$ ([13], §3.2). The reversor $\sigma(R, \varphi, Z) = (R, -\varphi, -Z)$ reverses the field-line flow; its fixed line on the initial section $\varphi = 0$ is $Z = 0$.

Every reversible coil certificate here uses $\eta(0) = \partial_Z$, which is antisymmetric under $(R, \varphi, Z) \mapsto (R, -\varphi, -Z)$ and is *not* the radial $\xi(s_0)$. Each run establishes the sign change by bracketing β between two enclosures of strictly opposite sign. The proof of the reversible criterion uses only a zero of β at a positive time; Lean proves both forms (Supplement S7.2).

For the spectral example, flux coordinates reduce the test to two scalars: a sign change of $\beta(\eta, \xi)$ is a sign change of η^ϑ, and at that moment $\operatorname{sign}\lambda(\eta) = \operatorname{sign}\eta^\psi$, provided $B^\varphi \neq 0$, which every run certifies as a domain guard on each step (Supplement S1.4). The reduction needs only $V^\varphi \equiv 1$ and a diagonal metric, so it survives the pull-back to the coordinates a spectral field is given in (Supplement S6); it is not formalised. No coil certificate uses a λ condition: each is a reversible or a conefield certificate (§2.3), and Lean proves the sign and zero equivalences between $\beta(\eta, \xi)$ and the scalar bracket of Theorem 1 that they check.

2.5. **Certification over parameter families.** The parameter is made a state variable with zero dynamics. For the reversible coil system, augment $Y = (R, Z, \eta_R, \eta_Z, \varepsilon)$ with $\dot{\varepsilon} = 0$; slices of the augmented flow at fixed ε are exactly the single-parameter systems, and a validated enclosure of the augmented flow contains every member trajectory.

Strict inequalities on an enclosure hold for every member it contains. The W7-X reversing-symmetry results use the following two cases of this argument.

Remark 3 (reversible criterion over a parameter box or an initial set)**.** Let $\eta(0)$ be antisymmetric under the reversing symmetry and independent of B, and let every field considered be exactly reversible under σ, $B(\sigma x) = -D\sigma\, B(x)$, with the fixed reference curve σ-symmetric, so that $\xi(\sigma x) = D\sigma\, \xi(x)$. Assume the domain guards of Theorem 1 on all enclosed segments; exact reversibility carries the B_φ guard to their reflections. *(a) Parameter box.* If s_0 lies on the symmetry line and the integrator, started from $\{(s_0, \eta(0))\} \times P$, produces enclosures of $\beta(\eta, \xi)$ of strictly opposite sign on two accepted steps, then for every $\varepsilon \in P$ no C^1 radial-graph torus of the field with parameter ε, relative to the single fixed reference curve, passes through s_0 or any point of its orbit. *(b) Initial set.* If X is a set of points of the symmetry line and the integrator, started from $X \times \{\eta(0)\}$ (or from $X \times H$ with H enclosing the admissible tangent at every point of X), produces enclosures of $\beta(\eta, \xi)$ of strictly opposite sign on two accepted steps, then for every $s_0 \in X$ no C^1 radial-graph torus passes through s_0 or any point of its orbit. No λ, and hence no η^ψ, hypothesis is required in either case.

Justification. The enclosures contain the flow image of every initial point and every parameter, in particular the member trajectory and transported tangent launched from each $s_0 \in X$ with each $\varepsilon \in P$, and a strict inequality on a set holds at each of its elements. For each member the two strict signs give a sign change of β on the intervening window, and the member is exactly reversible with s_0 on its symmetry line and $\eta(0)$ admissible; that is the hypothesis of the reversible alternative of Theorem 1, whose conclusion holds for that member.

Case (a) fixes the point and varies the field; case (b) fixes the field and varies the initial point; neither implies the other. The current box of §4.3 uses (a), and the symmetry-line interval of §4.1 uses (b). In (a) exact reversibility must hold for every member, not only the nominal field.

Remark 4 (conefield certificates over an initial box)**.** Let X be a box of initial points in the section $\varphi = 0$, bounded away from the reference curve. Assume the domain hypotheses of Theorem 2 on every enclosed forward and backward segment, with $R > 0$, one strict sign of B_φ on every step's a-priori enclosure, and $b_0 = |\xi_0|^2 > 0$ for every initial point. Let the conefield runs started from X, with both tangent columns started from ∂_R and ∂_Z, produce enclosures whose two witness times give $a_+ > 0$, $a_- < 0$ and $b_+a_- - b_-a_+ > 0$ when $a_\pm$, $b_\pm$ are formed in interval arithmetic from the enclosed brackets and an interval containing the radial offset ξ_0 of every point of X. Then, given Theorem 2, no C^1 torus of the radial-graph class passes through any point of X or its orbit.

Justification. The tangent columns start from the constant vectors ∂_R and ∂_Z, so the enclosure of the run from X contains, for every $s \in X$, the trajectory from s, both columns of its fundamental matrix and hence the brackets $\beta(M(t)\partial_R, \xi)$ and $\beta(M(t)\partial_Z, \xi)$ at every step. Only the initial offset $\xi_0 = (R - R_a(0), Z - Z_a(0))$, which fixes a_t and b_t, varies over X; the witness forms a_t and b_t in interval arithmetic from an offset interval containing every point's offset, so $a_{t_\pm}(s)$ and $b_{t_\pm}(s)$ lie in the emitted intervals for each s. The strict signs and the positive lower bound on $b_+a_- - b_-a_+$ are properties of those intervals, so the cone is empty at every $s \in X$, and Theorem 2 applies point by point.

General parameter-box and covering statements are given in Supplement S7.6.

For a parameter box P, the set certified uniformly is $\bigcap_{\varepsilon \in P} E_\varepsilon$, where E_ε is the exclusion set for the field with parameter ε. This intersection is over fields, not over the initial points in a spatial box.

The computational difficulty lies in producing sufficiently narrow enclosures over a non-trivial parameter box. The first-order parameter dependence is carried as an exactly transported sensitivity against a fixed interval that is never re-boxed (Supplement S2.3).

3. The validated implementation and its audit

3.1. **Validated integration, ablation, and the CAPD audit.** The integrator takes Lohner steps, with Taylor coefficients computed by jet arithmetic, a Lagrange remainder on a Picard a-priori enclosure valid for the whole step, and a QR-reorganised frame to control the wrapping effect. CAPD, an independent validated-integration library, gave overlapping enclosures at every recorded step of the four-dimensional analytic Case A test-field system, on integrable and near-separatrix orbits, not on W7-X. Ablation attributed the reduction in enclosure growth to the frame rather than Taylor order (Supplement S2.6). Trajectory and coordinate-tangent enclosures can be reused when changing the reference curve: only the reference-dependent witness algebra is recomputed. Supplement S1.2 and S5.1 give the axis rechecking procedure and its independent checks.

3.2. **A validated right-hand side for a coil field.** Validated evaluation of the Biot–Savart field is the main new component and the principal computational cost. A straight segment contributes

$$B = \frac{\mu_0 I}{4\pi} \frac{2(r_i + r_f)(\mathbf{r}_i \times \mathbf{r}_f)}{r_i r_f[(r_i + r_f)^2 - L^2]},$$

where $\mathbf{r}_i, \mathbf{r}_f$ point from its endpoints to the field point, r_i, r_f are their lengths and L is the segment length. Interval power-series arithmetic encloses each algebraic operation, including the square roots defining the distances; outward rounding also encloses the geometry constants. The sum over 4 800 to 6 720 total segments must be evaluated as an interval-valued truncated power series in φ to Taylor integration order 12–14, twice differentiated in (R, Z) for the tangent flow, with every operation outward-rounded. On one core through the arbitrary-precision path this is ∼60 s per validated step, so a single certificate would take months. We therefore wrote the segment sum as a CUDA kernel in interval-jet arithmetic with directed rounding, with the geometry reorganised so that the expensive part is linear in the segment data and the trigonometric series in φ is shared across all segments. For the single-displacement reversible coil runs, elapsed time per accepted step is 0.09–0.12 s on one RTX 5090 (Supplement S3.1), some 500–700 times faster than the arbitrary-precision path, and point runs using that first-order formulation of 3.7×10^4 to 5.1×10^4 steps take about one to two hours. The rectangle certificate of §4.1 requires both tangent columns advanced in lockstep in each time direction. Its 104 829 accepted steps count paired-column advances over both directions, not single-displacement steps: its recorded direction elapsed times sum to 8.00 hours, excluding axis rechecks (Supplement Table S13). These are recorded run elapsed counters, including rejected attempts and overhead, not kernel-only timings; the single-displacement rate is not a rate for this workload. At second order in the parameter a step costs 2.2–3.0 times as much on the host path ([M1]).

The GPU kernel is audited by an identity test: the same kernel source compiles against an emulation backend on the host, and the two must agree *bit for bit* on every series coefficient. The kernel is also checked against the arbitrary-precision ball path and against the analytic on-axis field of a circular loop. An error shared by both implementations passes the identity test, so conversion boundaries are checked separately against higher-precision ball arithmetic. Supplement S3.2 records the defects found and the build-specific coverage, including the earlier identity audit of the third-derivative path used by the second-order parameter run. The enclosure implementation remains unverified formally.

3.3. **An independent certificate verifier.** Each run emits the accepted step list, witness endpoint and a-priori enclosures, transported frame and domain guards. An independent checker uses only the Python standard library and exact rational interval arithmetic to check the finite criterion, sharing no code with the generator. For conefield witnesses it recomputes the signs and determinant; for the reversible criterion it verifies the strict sign change and its step bookkeeping. A separate validated-axis checker recomputes the reference offsets and witness algebra (Supplement S5.1).

Every W7-X finite witness listed in Supplement S7.11 passes the appropriate checker. As described in §1.4, the checker trusts the flow enclosures and does not evaluate the magnetic field. The package's `docs/certificate-format.md` records tests in which the released witnesses pass and eleven tampered variants are rejected.

4. Results

The principal result excludes radial-graph tori about the validated axis of the distributed 70-coil outward W7-X model through a nominal 10 nm × 10 nm rectangle (conefield box [D7]). We then compare the symmetrised outward and standard models.

4.1. **Outward-shifted configuration: certificates at mid-radius.** Geometry of the target. The rotational transform ι is the mean number of poloidal turns per toroidal turn. We use the outward current choice, row H of [28], with the current convention checked against [29, 20]. This choice moves the $\iota = 1$ resonance into the plasma, producing a mid-radius 5/5 resonant island chain. Its O-point (the elliptic periodic point at an island centre) is at $R = 6.276843\,\mathrm{m}$ on $Z = 0$, $\varphi = 0$. Greene's residue $\mathrm{Res} = (2 - \operatorname{tr} M)/4$, for the area-preserving periodic return matrix M, is 5.369×10^{-2}. The probe has relative depth $s = (\Delta R/w)^2 = 0.05$, where ΔR is its offset from the O-point and $w = 6.74\,\mathrm{mm}$ is the pendulum half-width ($s = 1$ at the separatrix). Floating-point surveys place it at about 0.69 ± 0.03 of the confined outboard radius and find radial-graph sections in all 40 sampled planes. These diagnostics select the target but enter no enclosure.

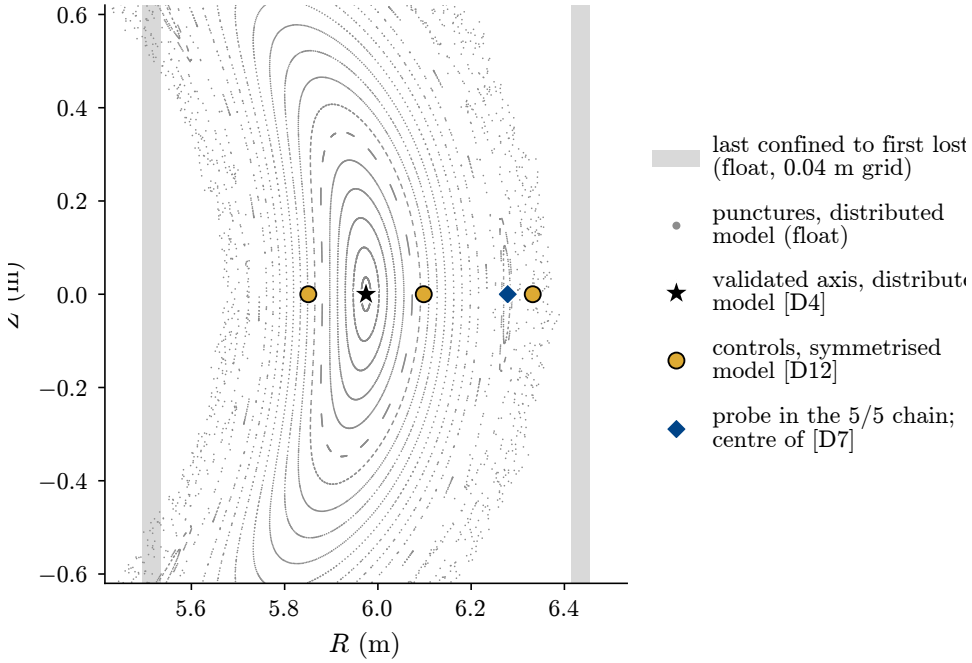


FIGURE 4. Where the certificates of the outward configuration sit in the section $\varphi = 0$. Grey dots are floating-point punctures traced on the distributed model, and the grey bands mark the floating-point confinement boundary; both are context. The star marks the validated axis of the distributed model [D4]. The blue diamond marks the probe inside the 5/5 chain, which is the point certificates' initial point and the centre of the principal 10 nm × 10 nm rectangle [D7] on the distributed model; the rectangle and the symmetrised-model symmetry-line interval are far smaller than the symbol (Figure 5 zooms in). The yellow circles are the controls on the symmetrised model [D12], which do not detect within the tested horizon.

Supplement S7.8 gives the current vector, survey and comparison with the standard configuration's exterior island chain; Figure 4 gives orientation.

For the distributed 70-coil outward model, the symmetry-free conefield criterion excludes every C^1 invariant torus that is a radial graph about the validated magnetic axis through a nominal 10 nm × 10 nm rectangle centred at the probe $R = 6.278350162606\,\mathrm{m}$, $Z = 0$, $\varphi = 0$ (conefield box [D7]). Remark 4 applies to the entire rectangle, with no exact-reversibility assumption. Figure 5 locates the rectangle and shows its size.

One validated integration per time direction carries the complete initial-condition rectangle (radial endpoints displayed rounded inwards to twelve

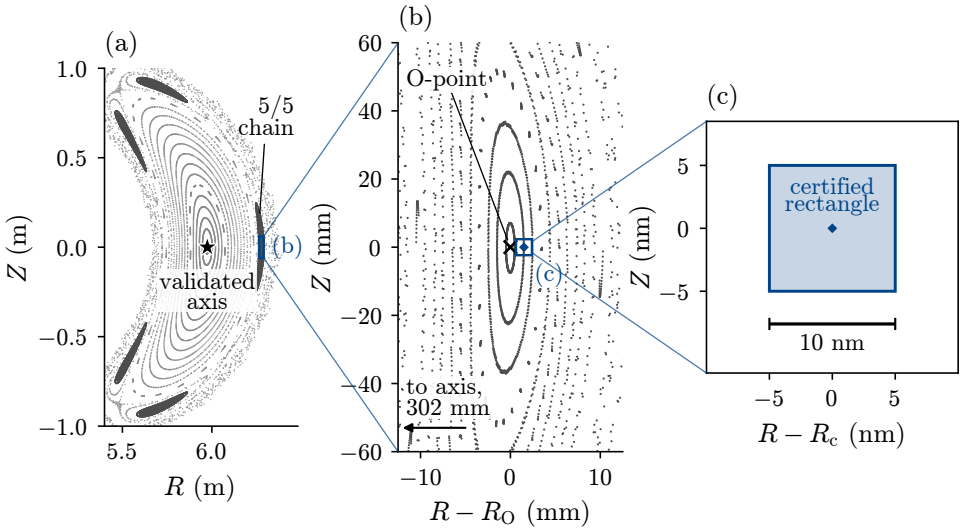


FIGURE 5. Where the principal certified rectangle [D7] lies and how small it is: three successive zooms in the section $\varphi = 0$ of the distributed 70-coil outward model. (a) Floating-point punctures of field lines of that model, recorded once per field period: survey lines seeded from the axis outward (grey) and lines seeded across the 5/5 island chain (dark); the star marks the validated axis and the box the region of panel (b). (b) The central part of the outboard island about its floating-point O-point (cross), with different scales on the two axes; the square marks the rectangle, which is far smaller than the symbol, and the arrow points towards the validated axis. (c) The certified rectangle, drawn to scale at its stated nominal size of 10 nm × 10 nm about its centre $R_{\mathrm{c}} = 6.278350162606\,\mathrm{m}$, $Z = 0$. The punctures are context and enter no enclosure; no symmetrised-model data are drawn.

decimal places)

$$R \in [6.278350157607, 6.278350167606]\ \mathrm{m},$$
$$Z \in [-5\times10^{-9}, 5\times10^{-9}]\ \mathrm{m}, \qquad \varphi = 0,$$

of nominal half-widths 5×10^{-9} and $5\times10^{-9}\,\mathrm{m}$ and area $10^{-16}\,\mathrm{m}^2$, to a horizon of approximately 3.1401 transits. The final witness, formed against this model's validated axis, is

$$\begin{array}{c|c|c} t_+ \approx 3.140 \text{ transits} & t_- \approx -3.140 \text{ transits} & \Delta = b_+a_- - b_-a_+ \\ a_+ \in [1.075, 1.257] & a_- \in [-1.258, -1.074] & \Delta > 0.00827 \end{array}$$

The exact-rational axis checker accepts these bounds over the whole rectangle. An admissible slope would obey $\varsigma \le b_+/a_+$ and $\varsigma \ge b_-/a_-$, but the difference of these bounds is $\Delta/(a_+a_-) < 0$. The cone is therefore empty at every

initial point, and Theorem 2 excludes radial-graph tori about the validated axis. The independent checker also accepts the fitted-reference witness, with exact-rational determinant lower bound at least 0.009388; the curves of Figure 6 illustrate that intermediate reference, and only its two black bars come from the axis recheck.

A half-width of 5 nm succeeds; at 10 nm, enclosure growth stops the run at 3.10 transits, before the cone empties. These trials measure the enclosure method, and need not be monotone in the width; Supplement S7.11 reports the widths tried and their costs.

The validated axis. An interval-Newton calculation encloses the full-2π return map and its derivative and proves a unique fixed point in the recorded box (outward axis [D4]; Supplement S1.2 gives the box, Newton image and determinant). Accepted-step enclosures cover the complete orbit in the coil-free domain, where the filament field is analytic. This analytic reference curve satisfies the C^1 hypothesis of Theorem 1. A standard-library checker verifies

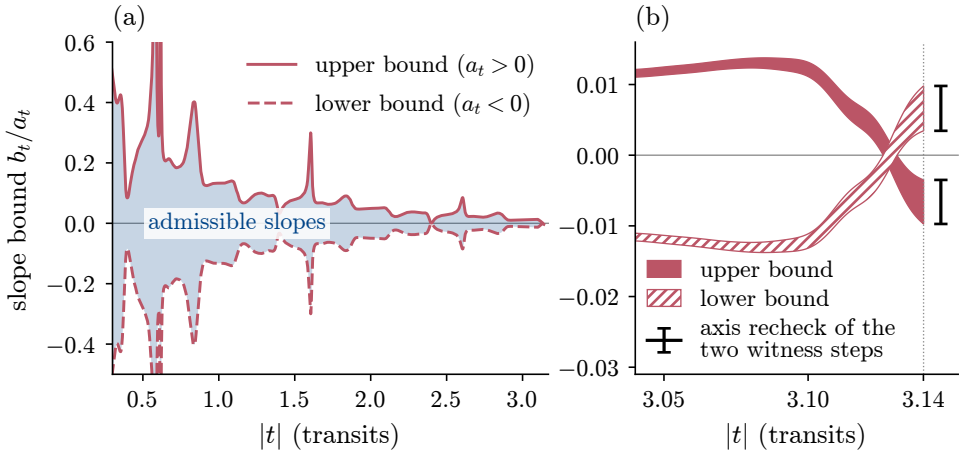


FIGURE 6. The cone of admissible slopes closing on the certified rectangle [D7], illustrated with the per-step enclosures against the fitted reference curve; against the validated axis only the two witness steps are re-enclosed (black bars), and the axis recheck verifies those two steps. Steps with $a_t > 0$ bound the slope from above by b_t/a_t and steps with $a_t < 0$ bound it from below; the two directions of integration are drawn against $|t|$. (a) The two bounds, with the slopes still admissible shaded blue. (b) The enclosures of b_t/a_t over the last 0.1 transit, where the upper bound falls strictly below the lower one. The black bars, drawn beside the horizon, are the quotient ranges formed from the axis recheck's recorded enclosures of $a_\pm$ and $b_\pm$: b_+/a_+ (lower bar) and b_-/a_- (upper bar). The finite certificate is the exact-rational check of the two witness steps.

the interval inverse and strict Newton inclusion; it trusts the return-map enclosures.

The point certificate, two ways. Both routes of §2.4 certify the probe $R = 6.278350162606\,\mathrm{m}$, $Z = 0$, $\varphi = 0$. Under the conefield criterion on the *distributed* model the cone empties, with determinant enclosure

$$[0.0172400, 0.0173565].$$

The independent checker accepts its exact-rational lower bound, and the witness is also accepted against the validated axis (outward conefield point [D5]). This is a certificate about the distributed model itself, by Theorem 2. Under the reversible criterion on the symmetrised model, whose field differs by 7×10^{-9} relative at the probe, detection occurs at 3.1283028 transits and the checker accepts (outward symmetrised point [D8]). Both outward models have 6720 directed segments in total: all 6720 have exact symmetry partners in the symmetrised model, whereas 4716 non-planar segments are exactly partnered in the distributed file.

Enclosure representation. The outward certificates use the split frame, which applies the QR reorganisation separately to the state and tangent blocks, as the C^r-Lohner algorithm [30] prescribes for the Jordan-like variational block (§S2.2). A Lohner step's soundness does not depend on the frame: every accepted step satisfies the same validation conditions, and the split frame only reduces enclosure growth.

Controls. Three control probes on $Z = 0$, in the core, inboard and outboard at $R = 6.0983716$, 5.8507716 and $6.3320716\,\mathrm{m}$, run on the symmetrised model with the settings, reference curve and representation of the symmetrised point above. None meets the criterion within the 4.382-transit horizon (outward controls [D12]). The controls check that the criterion does not fire everywhere against this reference curve. Figure 4 locates the controls.

Certified symmetry-line interval. On the symmetrised model, one initial-condition-set integration strengthens the point statement to every initial point in the set

$$R \in [6.278350152607, 6.278350172606]\ \mathrm{m}, \qquad Z = 0, \qquad \varphi = 0.$$

Its half-width is $10^{-8}\,\mathrm{m}$ and its total length is $2 \times 10^{-8}\,\mathrm{m}$. Every point of it lies on the reversor's fixed line, so the reversible criterion's hypothesis holds exactly at each of them, and the same antisymmetric initial tangent applies throughout. Detection occurred at 3.1284458 transits, the independent finite checker accepted, and the witness signs remained strict against the model's own validated axis (symmetrised interval [D10]). The independent exact-rational `axiswitnesscheck.py` also accepts these two brackets, taking the trajectory/tangent and propagated-axis enclosures as inputs; it does not reproduce axis propagation.

The interval's 20 nm width is more than five orders of magnitude below the island half-width. Strict finite witness margins persist under sufficiently small changes by continuity, within the hypotheses of the criterion; the reversible

criterion also requires preservation of exact reversibility and its symmetry-line initial conditions. Neither this nor the open-map robustness results of [3] gives a quantitative macroscopic width for this three-dimensional flow.

4.2. **Standard configuration: a point in the island divertor.** On the symmetrised standard model, the reversible criterion excludes every C^1 invariant torus of the radial-graph class through $R = 6.2398244\,\mathrm{m}$, $Z = 0$, $\varphi = 0$, and its orbit (standard point [D2]). The reference is a fixed fitted Fourier curve. The probe lies in the 5/5 island-divertor chain.

The configuration is the public, as-designed 50-coil non-planar set of the W7-X standard configuration (MAKEGRID filament file distributed with FOCUSADD [31, 32]; provenance and SHA-256 in the package's `docs/reproducibility.md`), each coil at 1.62×10^6 ampere-turns (the file's 15 kA per turn over 108 turns).

Only current ratios affect field lines, so the file's higher-field current (on-axis 2.911 T against the nominal 2.5 T) changes nothing. From the filament file alone our evaluation reproduces the axis position and core rotational transform (§S7.8) and locates the O-point of the 5/5 island-divertor chain on the symmetry line at $R = 6.237141\,\mathrm{m}$, with Greene residue $\mathrm{Res}_{5/5} = 8.167 \times 10^{-2}$: a strongly driven chain, for which the heuristic scaling $t^* \propto \mathrm{Res}^{-1/2}$, with t^* the detection time, suggests detection within a few transits (analytic-model evidence in Supplement S7.7); it is not a rigorous time bound.

The probe lies inside the 5/5 island on the $\varphi = 0$ symmetry line, at $R = 6.2398244\,\mathrm{m}$, $Z = 0$, which is 2.68 mm outboard of the O-point. A floating-point estimate of the island half-width, 27 to 33 mm, puts its relative depth at $s \approx 0.007$ to 0.010 (the package's `docs/certificate-stacks.md`). The reversible criterion of Theorem 1 applies relative to the fitted Fourier curve. The distributed standard file is reversible only to 10^{-14} relative in field, so the certificate is computed on its symmetrised model (§2.4), whose field differs from the distributed model by about 10^{-17} relative. This model-to-model difference is distinct from the distributed field's reversibility defect. There the reversible criterion inequalities hold at 2.1547107 transits and the independent checker accepts (Figure 7). The bracket enclosures at steps 36 743 and 36 744 are $[-3.0988, -2.7488] \times 10^{-5}$ and $[3.1936, 6.6948] \times 10^{-6}$, respectively.

4.3. **The coil-current box.** The standard point exclusion extends to a five-current box of nominal relative half-width 10^{-8} on the symmetrised model (standard current box [D3]), by Remark 3(a). Every member uses the same probe and fixed fitted reference curve as the point certificate.

Each non-planar coil type receives an independent current multiplier. The certified endpoints are the exact recorded binary64 values; an inward-rounded decimal interval is $[0.999999990000001, 1.000000009999999]$ for each of the five types. The second-order parameter-expansion run reaches the witness at

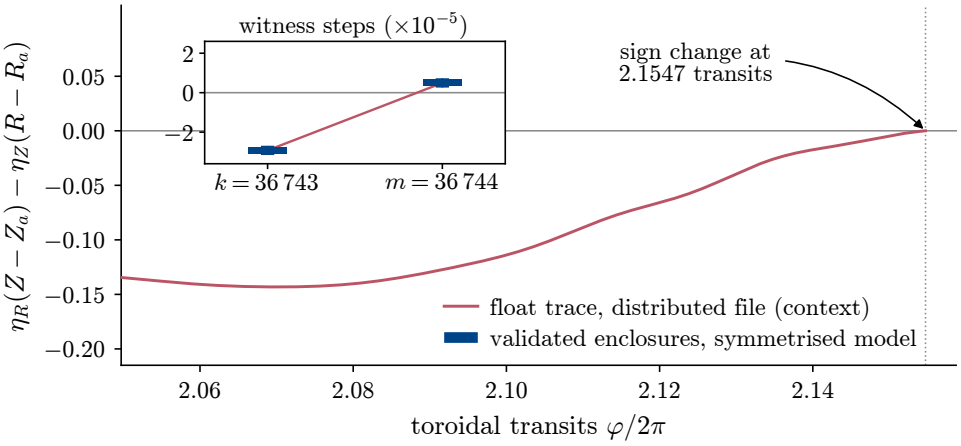


FIGURE 7. The standard point certificate: the section bracket of (1) changes sign at detection. The blue bars in the inset are the validated enclosures at the two witness steps on the symmetrised model; the red RK4 trace on the distributed file is floating-point context.

2.154708 transits with 45 752 accepted steps and an accepted exact-rational finite replay. Every σ-pair shares a coil type, so exact reversibility holds throughout the box. Supplement S2.4 describes the second-order protocol, including its current-Hessian terms, and the widths tried.

Current multipliers by coil type preserve five-fold periodicity and stellarator symmetry. They therefore cannot represent the $m = 1$, $n = 1$ error field (one poloidal and one toroidal oscillation) produced by coil-position errors and corrected by W7-X's trim coils. The common-scaling direction, $B \mapsto cB$, leaves field lines unchanged; the four directions transverse to common scaling in the five-dimensional current space change current ratios and control the useful box width. The certificate does not show that the structure at the probe is the same 5/5 island chain throughout the box.

Engineering limitation. Coil tolerances are a major engineering concern [33], assessed through measurement [34] and sampling [20]. The certified family is restricted and uncalibrated: the roughly 10^{-5} measurement in [34] concerns one resonant error-field component, and comparing it with these five current multipliers would require a calibrated current-to-error-field map. Wider runs that failed reflect enclosure growth (Supplement S2.4 and S7.4).

Secondary spectral-field example. The model problem of [13] is reproduced in Supplement S6. The reduced Tore Supra ergodic-divertor model certifies a current interval of ±4% (Tore Supra current box [D1], eight parameter subintervals). The independent checker accepts all eight parameter subintervals and their exact-rational coverage of the requested interval.

5. Discussion and limitations

Model. Every W7-X certificate concerns a specified vacuum filament model and does not incorporate plasma response or coil-placement uncertainty. On the standard configuration only the symmetrised model is certified; no symmetry-free certificate on the distributed standard file is established here, and repeating the standard-configuration point on the 70-coil file would be a check of dependence on the coil-data representation. An equilibrium for row H's currents from a code that can represent islands, such as SPEC [35], would place the certified region next to an independent account of the same chain; a VMEC solution, which assumes nested surfaces, can serve only as a nested-surface control.

Size and cost. The W7-X sets are far smaller than the millimetre-scale islands. At the present cost of about eight GPU-hours per $10^{-16}\,\mathrm{m}^2$, tiling cannot reach a physically informative scale: a square micrometre would take some 10^5 GPU-hours. Widening the certified region therefore needs tighter enclosures.

Torus class. A rigorous full-field twist or definite-shear certificate, together with the invariant-domain and section hypotheses required to apply Birkhoff's graph theorem, would make the conditional extension in Supplement S7.1 available for the coil field and upgrade the excluded class from radial graphs to all tori homotopic to a constant-radius coordinate torus. These additional hypotheses are not established here, and no certificate uses this extension. The existence side, a posteriori KAM for invariant tori of the same coil fields, is the subject of separate work.

Formal verification. End-to-end verification would need verified interval arithmetic, a verified interval-jet Biot–Savart evaluation and a verified Lohner step with proved Taylor remainders, at the scale of Immler's Isabelle/HOL verification [36] of the integrator for Tucker's Lorenz proof [37].

6. Conclusions

A finite converse-KAM certificate can be computed directly from realistic stellarator coil data. On the distributed outward-shifted W7-X model, the conefield certificate excludes C^1 radial-graph tori about the validated axis from a nominal $10\,\mathrm{nm} \times 10\,\mathrm{nm}$ rectangle inside the mid-radius 5/5 island chain (conefield box [D7]), without a reversing-symmetry hypothesis. The symmetrised outward model supplies a complementary 20 nm interval excluding graphs defined relative to its own axis; on the symmetrised standard model, a point and a restricted five-current box are certified about a fixed fitted curve; and the reduced spectral model supplies a separate $\pm 4\%$ current query. Each conclusion rests on three separately checkable layers: validated enclosures from audited code, an independent exact-rational check of the finite inequalities, and a mathematical criterion: the conefield implication proved above combined with Lean-verified finite algebra, or the Lean-verified abstract general-point or reversible implication.

Acknowledgments

The author thanks Robert MacKay, David Martínez-del-Río and Nikos Kallinikos for discussion and feedback.

Use of AI tools

The author used AI language-model systems throughout the work, including code, exposition and mathematical derivations, and retains sole responsibility for all claims. The verification scope of every result is stated in the text; no AI system is an author.

Conflicts of interest

The author declares no conflicts of interest.

ORCID iDs

Aapo Lipponen https://orcid.org/0009-0009-4613-2969

Data and code availability

Every number in this paper is tied to fixed input data and scripted checks. Tables and figures cite stable data records ([D1]–[D12], [M1], [V1]; contents in Table S7), and the release manifest maps each record to its files, roles, sizes and SHA-256 digests. The package's `docs/reproducibility.md` describes its contents: source, coil files with provenance, toolchain versions, verification gates, the independent verifier, the Lean development and a minimal reproduction. The complete package, including every certificate file, the manifests, the checker, the Lean sources and the coil files, is archived on Zenodo at https://doi.org/10.5281/zenodo.24369283, so that readers and referees can re-check the certificates. The two W7-X coil files are third-party data, used unchanged at pinned upstream commits: the 70-coil file from Proxima Fusion's VMEC++ test data [38] (MIT licence) and the 50-coil standard-configuration file from FOCUSADD [31, 32] (BSD-3-Clause licence); both carry the same as-designed filament geometry, and the package reproduces their licence notices.

References

[1] R. S. MacKay and I. C. Percival, *Converse KAM: theory and practice*, Commun. Math. Phys. **98**, 469–512 (1985), doi:10.1007/bf01209326.

[2] M. R. Muldoon, *Ghosts of order on the frontier of chaos*, PhD thesis, California Institute of Technology (1989); arXiv:1610.09170.

[3] J. Stark, *An exhaustive criterion for the non-existence of invariant circles for area-preserving twist maps*, Commun. Math. Phys. **117**, 177–189 (1988), doi:10.1007/bf01223588.

[4] I. Jungreis, *A method for proving that monotone twist maps have no invariant circles*, Ergodic Theory Dynam. Systems **11**, 79–84 (1991), doi:10.1017/s0143385700006027.

[5] J.-L. Figueras, A. Haro and A. Luque, *Rigorous computer-assisted application of KAM theory: a modern approach*, Found. Comput. Math. **17**(5), 1123–1193 (2017), doi:10.1007/s10208-016-9339-3.

[6] R. S. MacKay, *A criterion for non-existence of invariant tori for Hamiltonian systems*, Physica D **36**, 64–82 (1989), doi:10.1016/0167-2789(89)90248-0.

[7] V. Bangert, *Mather sets for twist maps and geodesics on tori*, in: U. Kirchgraber and H. O. Walther (eds.), Dynamics Reported, Vol. 1 (Wiley, 1988), pp. 1–56, doi:10.1007/978-3-322-96656-8_1.

[8] R. Kaiser and A. Salat, *Surface current equilibria from a geometric point of view*, Phys. Plasmas **1**(2), 281–295 (1994), doi:10.1063/1.870830.

[9] R. S. MacKay, *Invariant tori for the pressure-jump Hamiltonian*, Phys. Plasmas **32**, 052507 (2025); arXiv:2501.10353, doi:10.1063/5.0258987.

[10] R. S. MacKay, *Finding the complement of the invariant manifolds transverse to a given foliation for a 3D flow*, Regul. Chaotic Dyn. **23**, 797–802 (2018), doi:10.1134/S1560354718060126.

[11] D. Martínez-del-Río and R. S. MacKay, *Conefield approach to identifying regions without flux surfaces for magnetic fields*, Plasma Phys. Control. Fusion **67**, 055005 (2025); arXiv:2501.06796, doi:10.1088/1361-6587/adc59b.

[12] N. Duignan and J. D. Meiss, *Nonexistence of invariant tori transverse to foliations: an application of converse KAM theory*, Chaos **31**, 013124 (2021), doi:10.1063/5.0035175.

[13] N. Kallinikos, R. S. MacKay and D. Martínez-del-Río, *Regions without flux surfaces of given class for magnetic fields in toroidal geometry*, Plasma Phys. Control. Fusion **65**, 095021 (2023); arXiv:2304.09613, doi:10.1088/1361-6587/acea3f.

[14] N. Kallinikos, R. S. MacKay and T. Syndercombe, *Regions without invariant tori of given class for the planar circular restricted three-body problem*, Physica D **434**, 133216 (2022); arXiv:2202.12691, doi:10.1016/j.physd.2022.133216.

[15] L. Valvo and U. Locatelli, *Hamiltonian control of magnetic field lines: computer assisted results proving the existence of KAM barriers*, J. Comput. Dyn. **9**(4), 505–527 (2022); arXiv:2101.07785, doi:10.3934/jcd.2022002.

[16] M. Landreman and E. Paul, *Computing local sensitivity and tolerances for stellarator physics properties using shape gradients*, Nucl. Fusion **58**, 076023 (2018), doi:10.1088/1741-4326/aac197.

[17] J.-F. Lobsien, M. Drevlak, T. Sunn Pedersen and the W7-X Team, *Stellarator coil optimization towards higher engineering tolerances*, Nucl. Fusion **58**, 106013 (2018), doi:10.1088/1741-4326/aad431.

[18] C. Zhu, S. R. Hudson, S. A. Lazerson *et al.*, *Hessian matrix approach for determining error field sensitivity to coil deviations*, Plasma Phys. Control. Fusion **60**, 054016 (2018).

[19] C. Zhu, D. A. Gates, S. R. Hudson, H. Liu, Y. Xu, A. Shimizu and S. Okamura, *Identification of important error fields in stellarators using the Hessian matrix method*, Nucl. Fusion **59**, 126007 (2019); arXiv:1904.04147, doi:10.1088/1741-4326/ab3a7c.

[20] R. Davies, C. B. Smiet, C. Batzdorf, J. Geiger, J. Loizu and S. A. Henneberg, *Characterisation of X- and O-points in Wendelstein 7-X with respect to coil currents*, J. Plasma Phys. **92**, E27 (2026); arXiv:2511.15641, doi:10.1017/s0022377826101287.

[21] A. Geraldini, M. Landreman and E. Paul, *An adjoint method for determining the sensitivity of island size to magnetic field variations*, J. Plasma Phys. **87**, 905870302 (2021); arXiv:2102.04497, doi:10.1017/s0022377821000428.

[22] R. Davies, D. Boeyaert, A. Wolfmeister, B. Geiger, G. Harrer and J. Geiger, *Computational studies of giant edge islands and unpaired X-points in HSX and W7-X by manipulating coil currents*, arXiv:2607.14722 (2026).

[23] N. Duignan and J. D. Meiss, *Distinguishing between regular and chaotic orbits of flows by the weighted Birkhoff average*, Physica D **449**, 133749 (2023); arXiv:2212.10783, doi:10.1016/j.physd.2023.133749.

[24] M. Ruth and D. Bindel, *Finding Birkhoff averages via adaptive filtering*, Chaos **34**, 123109 (2024); arXiv:2403.19003, doi:10.1063/5.0215396.

[25] H. T. Moges, Y. Antonenas, G. Anastassiou, Ch. Skokos and Y. Kominis, *Kinetic versus magnetic chaos in toroidal plasmas*, Phys. Plasmas **31**, 012302 (2024); arXiv:2311.06213, doi:10.1063/5.0173642.

[26] C. B. Smiet, L. Rais, J. Loizu and R. Davies, *Turnstile flux as a measure for chaotic transport in magnetic confinement fusion devices*, Chaos **35**, 073129 (2025); arXiv:2503.15792, doi:10.1063/5.0275878.

[27] N. Bohlsen, V. Robins and M. Hole, *Automatic classification of magnetic field line topology by persistent homology*, Physica D **475**, 134595 (2025); arXiv:2408.09298, doi:10.1016/j.physd.2025.134595.

[28] T. Andreeva, J. Kisslinger and H. Wobig, *Characteristics of main configurations of Wendelstein 7-X*, Problems of Atomic Science and Technology **2002**(4), Series: Plasma Physics (7), 45–47 (2002); companion of IPP-Report III/270.

[29] J. Geiger, C. D. Beidler, Y. Feng, H. Maaßberg, N. B. Marushchenko and Y. Turkin, *Physics in the magnetic configuration space of W7-X*, Plasma Phys. Control. Fusion **57**, 014004 (2015), doi:10.1088/0741-3335/57/1/014004.

[30] D. Wilczak and P. Zgliczyński, *C^r-Lohner algorithm*, Schedae Informaticae **20**, 9–46 (2011).

[31] N. McGreivy, S. R. Hudson and C. Zhu, *Optimized finite-build stellarator coils using automatic differentiation*, Nucl. Fusion **61**, 026020 (2021), doi:10.1088/1741-4326/abcd76.

[32] N. McGreivy, *FOCUSADD*, file `focusadd/initFiles/w7x/coils.std_mc` at commit `a05524dd`, https://github.com/nickmcgreivy/FOCUSADD (BSD-3-Clause).

[33] R. Strykowsky, T. Brown, J. Chrzanowski, M. Cole, P. Heitzenroeder, G. H. Neilson, D. Rej and M. Viol, *Engineering cost & schedule lessons learned on NCSX*, 23rd IEEE/NPSS Symposium on Fusion Engineering (SOFE), San Diego (2009), doi:10.1109/FUSION.2009.5226449.

[34] T. Sunn Pedersen, M. Otte, S. Lazerson, P. Helander, S. Bozhenkov, C. Biedermann, T. Klinger, R. C. Wolf, H.-S. Bosch and the Wendelstein 7-X Team, *Confirmation of the topology of the Wendelstein 7-X magnetic field to better than 1:100,000*, Nat. Commun. **7**, 13493 (2016), doi:10.1038/ncomms13493 (open access).

[35] Z. S. Qu, S. R. Hudson, R. L. Dewar, J. Loizu and M. J. Hole, *On the non-existence of stepped-pressure equilibria far from symmetry*, Plasma Phys. Control. Fusion **63**, 125007 (2021); arXiv:2108.02056, doi:10.1088/1361-6587/ac2afc.

[36] F. Immler, *A verified ODE solver and the Lorenz attractor*, J. Automated Reasoning **61**, 73–111 (2018), doi:10.1007/s10817-017-9448-y.

[37] W. Tucker, *A rigorous ODE solver and Smale's 14th problem*, Found. Comput. Math. **2**, 53–117 (2002), doi:10.1007/s002080010018.

[38] Proxima Fusion GmbH, *Large C++ tests for VMEC++*, file `test_data/coils.w7x` at commit `394a09a6`, https://github.com/proximafusion/vmecpp_large_cpp_tests (MIT).

Supplementary Material for "Certificates of Absence: Validated Converse KAM for Flux Surfaces in Coil-Defined Fields"

Aapo Lipponen
Independent researcher, Finland
E-mail: aapo.lipponen@gmail.com

Section, equation, table and figure numbers written as §n, (n), Table n or Figure n refer to the main paper; items numbered Sn are in this document. This material supplies the detail the main paper compresses: the criterion in the coordinates the code works in (§S1), the validated integrator as an algorithm (§S2), the coil-field kernel, its audit and the soundness defects it found (§S3), the settings and outcomes of every certificate the paper quotes and the provenance of each record (§S4), the independent verifier (§S5), validation on analytic and spectral fields (§S6), extended treatments of points the main text states briefly (§S7). Cost measurements are reported as calibration evidence of [M1]; no companion manuscript is required to interpret this release.

The *general-point criterion* and *reversible (symmetry-line) criterion* refer respectively to Theorems 3.1 and 3.2 of [3]; tables abbreviate these names to "general-point" and "reversible".

S1. Notation and section coordinates

S1.1. Section coordinates, and the B_φ convention

The coil construction requires no known global flux function: the excluded class is whatever the choice of ξ makes it. Theorem 1 is the criterion in that setting; here is the computation behind it, and the one convention that can lose a factor of R silently.

Take the section (R, Z) at toroidal angle φ, with φ as time:

$$\frac{\mathrm{d}R}{\mathrm{d}\varphi} = \frac{R\,B_R}{B_\varphi}, \qquad \frac{\mathrm{d}Z}{\mathrm{d}\varphi} = \frac{R\,B_Z}{B_\varphi}, \qquad \mu \;=\; B_\varphi \;=\; RB^\varphi, \tag{S1}$$

valid wherever $B_\varphi \neq 0$, which the floating-point survey checks and the validated run re-checks per step as a domain guard. Here B_φ is the *physical* toroidal component, the one a Biot–Savart evaluation returns in the orthonormal cylindrical frame, while the *contravariant* component in (R, φ, Z) is $B^\varphi = B_\varphi/R$: the two differ by exactly the factor R in the right-hand side, and on W7-X by $R \approx 6\,\mathrm{m}$. The section flow preserves $\mu\,\mathrm{d}R \wedge \mathrm{d}Z$, which is $\nabla \cdot B = 0$ in cylindrical coordinates. For the positive-density chart

used in Lean we normalise the fixed section orientation: when $B_\varphi < 0$, replace the flux form by its negative, giving density $-B_\varphi > 0$. This global sign choice leaves zeros and strict sign changes unchanged.

With a C^1 closed reference curve $(R_a(\varphi), Z_a(\varphi))$ and

$$\psi = \tfrac{1}{2}\big[(R-R_a)^2 + (Z-Z_a)^2\big], \qquad \xi = \nabla_{R,Z}\psi = (R-R_a,\ Z-Z_a) \tag{S2}$$

in the Euclidean metric of the section (and with zero toroidal component), the flux form on poloidal pairs is $\beta(\eta,\xi) = \mu\,[\eta_R(Z-Z_a) - \eta_Z(R-R_a)]$, which is eq. (1). Because B_φ keeps one sign on the working annulus, which the domain guard checks per step, a sign change of β is a sign change of the bracket alone, that is, of $\eta \times (x - x_a)$: the displacement rotating past the reference-radial direction, geometrically identical to the η^ϑ sign change in S1.4. The second condition of the nonsymmetric theorem, on λ formed from this direction field and B, is not evaluated here. The reversible coil-field certificates are of the symmetric form (criterion `th32-coil`): the line $Z = 0$ at $\varphi = 0$, on which every such run starts, is taken as available and the λ condition dropped on it, as Theorem 3.2 of [3] permits for an exactly reversible field. The distributed W7-X coil files are only approximately stellarator-symmetric, by amounts that differ per file and must not be quoted across them: for the outward file the field identity holds to 2.2×10^{-8} and its planar-coil geometry only to 1.6×10^{-5} m, while the standard file carries no planar coils and its identity holds to 10^{-14}, limited by last-digit rounding of the vertices. Neither is exact, so that hypothesis is not discharged for them as distributed. §2.4 of the main text gives the two outward-configuration replacements that close it: a conefield certificate, which needs no reversor, and a reversible certificate against a model that is exactly σ-invariant by construction. The standard configuration has its own symmetrised model (standard point [D2]), at a field change of about 10^{-17} because its file carries no planar coils. What is new is the right-hand side (S1) and its interval jets (§S3).

S1.2. The reference curve, and the FourierAxis fit

The reference curve fixes *which* class is excluded and is data of the certificate. Two properties follow from (S2): classes for different curves are incomparable, neither implying the other, so a certificate is always relative to the curve recorded with it; and the quality of the fit affects only provenance, since Theorem 1 needs only that the curve be C^1 and closed. A poor choice shrinks the physical interest of the class, never the validity of the certificate.

The curves used here are stellarator-symmetric Fourier polynomials,

$$R_a(\varphi) = \sum_{k=0}^{K} c_k^R \cos(kN_{\mathrm{fp}}\varphi), \qquad Z_a(\varphi) = \sum_{k=1}^{K} s_k^Z \sin(kN_{\mathrm{fp}}\varphi), \tag{S3}$$

implemented as a common `FourierAxis` type and evaluated through the same generic numeric dispatch as the field, so identical code composes with floats, Arb balls and Arb power series. A finite sum of this form is analytic, hence satisfies the theorem's C^1

hypothesis, and the coefficients are part of the emitted certificate (§S5), so the verifier re-evaluates the curve rather than trusting a number from the generator. They come from a float fit: Newton on the one-period return map from $(R_0, 0)$ at $\varphi = 0$, then RK4 sampling over a full field period and a least-squares fit of (S3) to the samples. No part of that fit is claimed to be rigorous. Recorded with each curve are the Newton closure, the maximum residual and the excursion of the axis from a constant circle, the last being the measurement that retired the constant-circle class (Table S1).

Table S1. The fitted Fourier curves used by the coil-field certificates. "Excursion" is the maximum distance of the fitted axis from the constant circle through its bean-plane point within one field period, which is why a constant circle defines a physically near-empty class for these devices. All entries are float diagnostics; nothing but the coefficients of (S3) enters a certificate, and every certificate carries its curve's coefficients, which the verifier re-evaluates. "Closure" is the Newton closure of the axis on which the curve is fitted, and the residual column is $\max(\mathrm{resid}_R, \mathrm{resid}_Z)$ over the traced points. Each curve is fitted on the distributed coil file and used unchanged on its symmetrised model. Used by: [D2] and [D3] (standard); [D5]–[D12] (`outward`).

field	N_{fp}	K	closure (m)	max residual (m)	excursion (m)
W7-X standard, 50-coil file	5	48	1.3×10^{-14}	8.1×10^{-8}	0.744
W7-X `outward`, 70-coil file	5	24	5.6×10^{-15}	8.7×10^{-8}	0.727

Distributed-model interval-Newton axis enclosure. We validate the magnetic axis of this filament field as a periodic orbit. Time was normalised by $\varphi = 2\pi\tau$, with 2π represented by an `arb` interval, and the same Lohner integrator enclosed the full-period return map and both columns of its derivative on the box

$$R \in [5.9745714182, 5.9745718181]\ \mathrm{m}, \qquad Z \in [-2.25698, 1.74301] \times 10^{-7}\ \mathrm{m}.$$

The interval-Newton image lies strictly inside that box:

$$R \in [5.974571637244, 5.974571637387]\ \mathrm{m}, \qquad Z \in [-2.33254, -2.30350] \times 10^{-8}\ \mathrm{m}.$$

The determinant of the enclosed fixed-point Jacobian is $[0.57018, 0.57102]$, bounded away from zero. Thus the full-2π return map has a unique fixed point in the box, and the accepted-step a-priori enclosures cover its complete orbit. The orbit stays in the coil-free working domain, where the finite-filament field is analytic, so this periodic solution is an admissible analytic reference curve for Theorem 1. A standard-library checker recomputes the 2×2 interval inverse and strict inclusion from outward-padded binary64 endpoints (outward axis [D4]).

Validated-axis symmetry and witness rechecking. The reversible criterion also needs a σ-symmetric curve (Theorem 1). The fitted Fourier curves are symmetric by construction. So is the symmetrised model's own validated axis: its interval-Newton image at $\varphi = 0$ has $|Z| < 2.5 \times 10^{-10}\,\mathrm{m}$, so the mirror image $(R, -Z)$ of its fixed point lies in the uniqueness box; exact reversibility makes that mirror image a fixed point of the inverse return map,

hence of the return map, and uniqueness forces $Z = 0$ and an orbit that σ maps to itself (outward symmetrised point [D8]). Because trajectory and tangent enclosures do not depend on the reference curve, each outward witness can be re-evaluated against enclosures of the validated axis.

Changing the curve changes ξ and hence the excluded class, yet re-evaluating the witness suffices, because neither criterion the coil runs use lets the curve into the integration. The curve enters only through ξ in $\beta(\eta, \xi)$ and, for the general-point criterion, through $\eta(0) = \xi(s_0)$; the reversible criterion starts from $\eta(0) = \partial_Z$, admissible by antisymmetry and independence from B alone, and the conefield runs from ∂_R and ∂_Z. So the base trajectory, the transported tangents and every step's enclosure are unchanged by a change of curve; only scalars assembled from them change. *The reversible criterion.* Re-evaluating the two witness brackets against the new ξ re-establishes the whole hypothesis of the reversible criterion. *Conefield.* The curve also fixes the initial offset ξ_0, and with it a_t and b_t. The recheck encloses the validated axis at $\varphi = 0$ and forms ξ_0 from it (over the whole initial box for a set-valued run, refusing a set not bounded away from the orbit), propagates the orbit's interval-Newton box to the two witness angles, recomputes both columns' brackets there from the stored enclosures, which the independent checker has matched to the domain-checked step streams, and re-establishes $a_+ > 0$, $a_- < 0$ and $b_+a_- - b_-a_+ > 0$; an admissible slope would have to satisfy the inequality at those two times, so the cone is empty for the validated-axis class. Both rechecks run in Arb ball arithmetic in the generator's code. For reversible witnesses, including [D10], the independent exact-rational `axiswitnesscheck.py` recomputes the two brackets and checks their opposite strict signs. For conefields, `coneaxischeck.py` additionally recomputes the initial offset, covectors and determinant (§S5.1). Both trust the axis propagation and trajectory/tangent enclosures. Under the general-point criterion the move would *not* be legitimate, since $\eta(0)$ was started along the old ξ. Only the outward results therefore support the magnetic-axis-centred statement; the standard-configuration result remains relative to its fitted Fourier curve.

S1.3. Validated background-profile shear

Table S2 encloses the derivative of the unperturbed transform profile ι. It does not verify hypothesis (i) of Corollary S1 for the perturbed field. For Hamiltonian modes with amplitudes a_{mn}, radial profiles $p_{mn}(\psi)$ and constant phases ζ_{mn}, primes denoting ψ-derivatives, the full shear includes $\sum a_{mn}p''_{mn}\cos(m\vartheta - n\varphi + \zeta_{mn})$ and the discarded spectrum. A uniform sign bound on these terms and the remaining global Birkhoff hypotheses have not been established for any configuration here.

Table S2. Validated background-profile derivative (calibration [M1]); each configuration decides in a single interval evaluation. On the two analytic models the unperturbed transform is $\iota(\psi) = w_1 + 2w_2\psi$, so $\iota' = 2w_2$. "Low shear" in the stellarator model refers to the narrow *range* of ι (0.500–0.556), not to ι' approaching zero. ι' carries units of inverse flux, so its magnitude depends on the flux-unit convention while its *sign* does not. These background bounds alone do not supply full-field twist. The check is not available for coil fields, where ι is not represented as a profile. All certified results retain the radial-graph restriction.

configuration	ι' on the annulus	background shear
model problem (Case A)	$2w_2 = 2$ exactly	definite (analytic)
island-divertor model	$2w_2 = 0.2$ exactly	definite (analytic)
Tore Supra ED, $I_d = 4.5\,\mathrm{kA}$	$[-1.856,\ -0.8836]$	**definite**
stellarator model (case (c))	$[0.1924,\ 0.2587]$	**definite**

S1.4. The criterion in section coordinates

The field-line flow is taken with the toroidal angle φ as time. With V the normalised field, $V^\varphi \equiv 1$, the integrated system is

$$\dot{s} = V(s,\varphi), \qquad \dot{\eta} = DV(s,\varphi)\,\eta, \qquad \eta(0) = \xi(s_0), \tag{S4}$$

in the four scalars $(\psi, \vartheta, \eta^\psi, \eta^\vartheta)$: the φ-component of $DV\eta$ vanishes, so a displacement that starts poloidal stays poloidal, and that is the only structural fact the reduction uses.

Theorem 1 is stated with the flux 2-form $\beta = i_B\Omega$, whose kernel is $\mathrm{span}(V)$, and the 1-form $\lambda = \mathrm{d}\psi - kB^\flat$, $k = B\cdot\nabla\psi/|B|^2$. On poloidal pairs both collapse: with $\xi = 2B_0\psi\,\partial_\psi$ and $\beta = \mathrm{d}\psi\wedge\mathrm{d}\vartheta$,

$$\beta(\eta,\xi) = \eta^\psi\xi^\vartheta - \eta^\vartheta\xi^\psi = -2B_0\,\psi\,\eta^\vartheta, \qquad \lambda(\eta)\Big|_{\eta^\vartheta=0} = \eta^\psi\Big(1 - \frac{(B^\psi)^2}{2B_0\psi\,|B|^2}\Big), \tag{S5}$$

with B_0 the reference field strength that fixes the flux unit of ψ, the second bracket positive by Cauchy–Schwarz whenever $B^\varphi \neq 0$. (Cauchy–Schwarz alone gives $g_{\psi\psi}(B^\psi)^2 \le |B|^2$, with equality exactly when $B^\vartheta = B^\varphi = 0$; the domain guard $B^\varphi \neq 0$ excludes that set, and without it the reduction gives no sign at all.) So on $\psi > 0$ a sign change of $\beta(\eta,\xi)$ is a sign change of η^ϑ, and at such a moment $\operatorname{sign}\lambda(\eta) = \operatorname{sign}\eta^\psi$. The whole test is the one the code runs: integrate four scalars; each time η^ϑ changes sign, check the sign of η^ψ on the a-priori enclosure of the bracketing step. The reduction needs $V^\varphi \equiv 1$ and a diagonal metric, not canonical coordinates, which is why it survives the pull-back to the coordinates a spectral field is given in. In the coil chart of Supplement S1.1 the two are equal, and that equality is proved in the Lean development (Supplement S5.1); the spectral reduction is not formalised.

S2. The validated integrator

What follows is the contract each step meets, and the equations it meets it with.

S2.1. The Lohner step and its a-priori enclosure

Ball arithmetic wrapped around a Runge–Kutta scheme encloses the iterates of a discrete map, not the flow. A Lohner step encloses the flow. For $\dot{y} = F(\varphi, y)$, one step from t_0 over h is

$$\begin{aligned} y(t_0 + h) \;&\in\; \sum_{k \le p} c_k(t_0, y(t_0))\, h^k \;+\; c_{p+1}(T, Z)\, h^{p+1}, \\ \text{subject to} \qquad & Y + [0, h] \cdot F(T, Z) \;\subseteq\; Z, \end{aligned} \tag{S6}$$

with $T = [t_0, t_0 + h]$, the c_k the ODE Taylor coefficients at t_0, computed by jet arithmetic (truncated power series pushed through the *unchanged* field code, coefficients read off by $c_{k+1} = (F\text{-coefficient } k)/(k+1)$), and the last term the exact Lagrange remainder. The reversible and spectral certificates use $p = 14$, the conefield runs $p = 12$ (Table S5), and all arithmetic is Arb midpoint–radius ball arithmetic [19] at a stated precision, every operation outward-rounded.

The inclusion on the right is the Picard–Banach condition making Z an a-priori enclosure: every solution starting in the current set stays in Z for the whole step. Z is found by iterating that map with inflation, halving h if the inclusion will not close, and is deliberately crude, since it enters only through the remainder, which carries h^{p+1}.

Z does a second job that is a hypothesis of the theorem rather than a convenience. The detection event falls at an unknown time inside a step, and the λ condition must hold on the whole bracketing window, not at its endpoints: an endpoint-only reading admits no intermediate-value proof, a point the Lean development made explicit and named. Because Z contains the state at *every* time in the step, certifying $\eta^{\psi} < 0$ on Z certifies $\lambda < 0$ at the crossing time itself.

We do not use CAPD to produce the certificates, although it is tighter, because none of the three things this paper turns on is available in it: interval-valued field parameters carried as exactly-transported sensitivity generators (§2.5) and their second-order extension; an interval-jet Biot–Savart right-hand side on a GPU behind a bit-identity gate (§3.2); and an acceptance contract stated in a proof assistant (§1.4). The geometry reorganisation that makes the kernel's expensive part linear in the segment data is set out in §S3.1.

S2.2. Frames: QR and the split frame

After each step the reachable set is (new centre) $+ S\cdot$(old set), with S an interval enclosure of the step Jacobian on the hull of the current set. Storing the set as a box and multiplying by S forces a re-boxing that overestimates by up to the box diagonal, and compounding that is the wrapping effect. The set is therefore a parallelepiped $Y = \hat{y} + Ar$, and after each step the new frame is $A' = Q$ from the QR factorisation of $\mathrm{mid}(SA)$, columns pre-sorted by width contribution. The new radii are an interval enclosure $r' \supseteq (Q^{-1}SA)\, r + Q^{-1}d$, with d the centre/remainder defect. Choosing Q from QR of $\mathrm{mid}(SA)$ makes $Q^{-1}SA$ nearly upper triangular, so re-boxing r' overestimates far less than re-boxing in the original coordinates: QR reduces the wrapping effect, it does

not eliminate it. Q^{-1} is a rigorous interval matrix, so no rigour is lost. Soundness does not depend on which invertible frame is chosen, since the new radii are formed from the whole row of SA either way. The split frame below therefore changes enclosure growth without affecting soundness.

Why a split is needed is the *structure* of the variational matrix. The field-line flow does not see the tangent vector η, so the 4×4 variational matrix is

$$S = \begin{pmatrix} M & 0 \\ N & M \end{pmatrix}, \qquad N = D^2\Phi \cdot \eta, \tag{S7}$$

block lower triangular *with equal diagonal blocks*. Its eigenvalues are those of M, each doubled, and N generically supplies generalised eigenvectors: the extended system is then Jordan-like, and its solutions pick up the secular $t\lambda^t$ growth that no single orthogonal frame follows. The remedy is not ours: the C^r-Lohner algorithm [9] controls the wrapping effect "separately for x-variables and variables $D_a\varphi$", and warns that this separation "is not present in the naïve application of C^0 algorithm to the system of variational equations and it turns out that this often practically switches off the control of the wrapping effect on x-variables, as various choices used in this control become dominated by the $D_a\varphi$-variables". A single pivoted frame over the combined system is precisely that case.

The *split frame* takes the QR reorganisation on the (R, Z) and (η_R, η_Z) blocks separately instead of on the 4×4; the cross term Nr_{state} of (S7) then enters the tangent block as an ordinary interval and wraps additively, rather than being followed by a single orthogonal frame that cannot follow it. Two tests check that the split frame contains the true trajectory of a point initial condition (against a Runge–Kutta reference with steps 1250 times smaller than the Lohner step) and is no wider than the 4×4 frame in the state rows. Because the criterion is one-sided and every invertible frame gives a valid enclosure, no choice of frame can manufacture a certificate: a frame option can extend an enclosure's reach, never produce a false detection.

S2.3. The parametric extension

The parameter is a state variable with zero dynamics: for a box P the augmented system is $Y = (\psi, \vartheta, \eta^\psi, \eta^\vartheta, \varepsilon)$ with $\dot{Y} = (V(\cdot;\varepsilon),\ DV(\cdot;\varepsilon)\eta,\ 0)$. Slices at fixed ε are invariant and are exactly the single-parameter systems, so a validated enclosure of the augmented flow contains every member trajectory, which is the whole content of Theorem S1.

The difficulty is the representation. Passing the fat interval P into every field evaluation re-boxes the parameter spread at every step, forgetting the state–parameter correlations exactly as a box forgets shear. Carrying ε as an extra QR-frame coordinate preserves them and is sound, but was measured unstable here: the true set is a one-dimensional curve whose components have wildly heterogeneous scales, and the QR frame built on such a degenerate collection was measured to amplify each step's remainder (super-exponential blowup by transit 2 at $w = 10^{-2}$). The adopted representation is an

explicit rank-one parameter term,

$$Y \;=\; \hat{y} \;+\; \textstyle\sum_j s_j\,[-\delta_j,\delta_j] \;+\; A\,r, \qquad s_j \leftarrow S\,s_j, \tag{S8}$$

with s_j the sensitivity of the whole state to parameter j, initialised to the unit vector on the ε_j slot and updated exactly by the step Jacobian, and $[-\delta_j,\delta_j]$ fixed for the whole run and never re-boxed. The first-order parameter dependence is then carried with zero wrapping by construction, and Ar carries only remainders and curvature, homogeneous small scales that are what the QR machinery is good at. Per step this is the mean-value form $f(\varepsilon_0)+(\partial f/\partial\varepsilon)(P-\varepsilon_0)+$remainder with the sensitivity accumulated exactly rather than re-derived.

S2.4. Second order in the parameter

What (S8) loses is visible in the radius recursion, $\mathrm{rad}(s_j) \leftarrow |S|\,\mathrm{rad}(s_j) + \mathrm{rad}(S)\,|\mathrm{mid}\,s_j|$ with $\mathrm{rad}(S) \approx D^2\Phi\cdot$(hull width): the genuine second-order term is carried as a symmetric radius on the first-order column and propagated by $|S|$, where the true recursion is signed and cancels. No tighter evaluation repairs that. The second-order integrator therefore extends the state by $u_j = \partial y/\partial\varepsilon_j$ (the letter w is reserved for the box half-width), which satisfies $\dot{u}_j = DF(y;\varepsilon)u_j + \partial F/\partial\varepsilon_j$ with $u_j(0)=0$, and runs the first-order integrator of §S2.3 on (y,u); its Jacobian needs one field derivative more (D^3V, D^2V_ε), and, for non-affine parameter dependence, the pure current Hessians $V_{\varepsilon_j\varepsilon_k}$ and $DV_{\varepsilon_j\varepsilon_k}$. For linear coil fields, put $q_j = B_{\varphi,j}/B_\varphi$; quotient differentiation gives

$$V_{\varepsilon_j\varepsilon_k} = -V_{\varepsilon_j}q_k - V_{\varepsilon_k}q_j.$$

Spatial differentiation of this identity supplies $DV_{\varepsilon_j\varepsilon_k}$. These terms enter the u_j-row Jacobian columns through $F_{\varepsilon_j\varepsilon_k}$, including their action on the transported tangent. The historical D3 generator omitted them; that witness remains withdrawn. The repaired run archived in standard current box [D3] includes them and restores the $\pm10^{-8}$ current-box exclusion at D2's exact probe, with an accepted finite replay. The repaired interface requires either these derivatives or an explicit affine-parameter contract, with scalar and power-series difference tests on a non-affine two-current coil field. After each step the sensitivity columns s_j (the y-rows of the parameter term) are reset to the centre $\hat{u}_j$ of the u-rows and the second-order displacement is reconstructed fresh from the transported second-order sensitivity columns $s_{2,jk}$, the u_j-rows of the sensitivity column for parameter k (enclosing $\partial u_j/\partial\varepsilon_k$), by

$$E \;=\; \sum_j s_{2,jj}\,[0,\,\delta_j^2/2] \;+\; \sum_{j\neq k} s_{2,jk}\,[-\delta_j\delta_k,\,\delta_j\delta_k] \;+\; \sum_j (Ar)_{u_j}\,[-\delta_j,\delta_j], \tag{S9}$$

kept *outside* the propagated remainder. The ordered off-diagonal sum overbounds the half-weighted quadratic term. The reset is legitimate because $s \leftarrow Ss$ is a mean-value statement valid for any set of the form $\hat{y} + s\tau + Ar$ that encloses the members, and (S9) gives such a set; keeping E outside matters because pushing it into the propagated

remainder as an increment double-counts the whole second-order displacement every step. The frame is block-diagonal here for the reason of §S2.2 one level up: a single orthogonal frame over (y, u) cannot follow the non-normal map $\tilde{S}_{u,y} \sim D^2\Phi \cdot u$ and was measured to wrap at $\times 1.3$ per step.

The repaired W7-X five-current experiment (standard current box [D3]) reaches relative half-width 10^{-8} with the complete Jacobian. Its historical first-order comparison remains exploratory and does not measure a speedup of the repaired run. For a point certificate a box of width zero has no second order to carry. This was checked directly: at zero width the field enclosure with seven current groups differentiated is identical to the last bit, at three precision tiers, to the enclosure with all coils in one group.

Historical width trials and their interpretation. An uncertified run on the distributed file, made before the last kernel repair (calibration [M1]), lost a box of half-width 10^{-6} at second order at 1.66 transits. A box that fails bounds the reach of the protocol, not the widest box for which the conclusion is true, and the gap can be large. On a controlled Hamiltonian family (§S7.4) the certifiable half-width falls by four decades, from 9.85×10^{-4} to 4.19×10^{-8}, as the amplitude ratio grows from 10 to 100, while the largest half-width w_{phys} at which floating-point tracing shows the probe still librating about the island O-point at both box edges stays ≥ 0.3. This is a numerical indicator, not a validated exclusion or existence result. The null direction shows the same on W7-X: the diagonal has the same field lines for every member and is certifiable for any half-width below 1, yet in uncertified pre-repair runs (calibration [M1]) a first-order box of half-width 10^{-6} along it, against a constant reference circle, was lost at 1.39 transits, while at second order the same box detects at the point's 2.1547 transits.

S2.5. What may and may not be compared across runs

Soundness and comparability are independent properties of an adaptive validated integrator, and only the first is what a certificate asserts. The arithmetic kernel is inclusion-isotone, but the method as a whole is a hybrid map whose control guards are not: the Picard first-success index, step acceptance and step selection, a width-dependent QR column ordering, and any thresholded split or fallback each partition the input space, and two nested inputs may land in different cells. Freezing the column *permutation* is not enough; the actual output frame must be shared, since QR depends on the represented generators. The effect is not hypothetical. Correcting one convolution so that it is outward-rounded per operation (the fourth row of Table S4) makes the field enclosure strictly wider for every input, yet on uncertified `outward` island runs the corrected stack produced final widths $8\times$ smaller at 1.0 transits and $500\times$ smaller at 2.1, with an identical count of rejected steps; equal rejection counts do not imply equal accepted step sizes, Picard indices, controller states or frames. A single ordering flip can account for a factor of this size, because it moves remainder width between directions that the downstream flow contracts at different rates: over the $\sim 2 \times 10^4$ steps of these runs a persistent anisotropy of 0.031% per step suffices. Enclosure width is therefore frame-path

sensitive at the 10–100× level.

That an interval algorithm can be sound without being inclusion-monotone is not a new observation: [12, §2.4] separates inclusion monotonicity from convergence and exhibits a convergent inclusion function that is not inclusion-monotonic, and [13] gives explicit examples in which narrower interval inputs produce wider outputs. What is specific here is that the guards causing it are the adaptive controls of a validated ODE solver rather than the arithmetic.

Consequently every enclosure reported here contains the true flow, conditional on the soundness of the audited generator as stated in §1.4, and every certificate drawn from one stands on that condition; but comparative statements of the form "this representation reaches farther, therefore the underlying problem is easier for it" are *not* supported unless the two runs share a common execution trace, and reach figures should be read as implementation-performance observations rather than as measurements of the mathematics. Two comparisons survive without that apparatus. A strict-sign certificate transfers *downward*: if a sound enclosure of a parameter box proves the criterion at time t, every field in the box satisfies it at t, whatever an independently adaptive run on a smaller input reports. And for a family of inputs ordered by inclusion, replacing each saved certification time by the minimum over its supersets, $\tilde{\tau}_i = \min\{\hat{\tau}_j : X_i \subseteq X_j\}$, yields a valid and inclusion-monotone family of certificate times from saved artefacts alone. For any other comparison the remedy is inexpensive: run the adaptive controller once on an outer envelope, record the accepted step sizes, a-priori sets, centres, frames, split trees and fallback choices, and replay every compared input with that trace frozen. Every operation in the replay is then inclusion-isotone, so nested inputs give nested outputs, at the cost of about one extra pass.

S2.6. Ablation and the CAPD audit

Ball arithmetic wrapped around a Runge–Kutta scheme is not enough: it encloses the iterates of a discrete map, not the flow. The integrator here takes Lohner steps: Taylor coefficients of the flow by jet arithmetic through the unchanged field code, a Lagrange remainder evaluated on a Picard a-priori enclosure $\mathcal{Z}$ valid for the whole step, and a QR-reorganised frame to control the wrapping effect. Because $\mathcal{Z}$ covers all intra-step times, it also closes the gap between checking λ at step endpoints and at the crossing itself.

The design claim is measured by ablation. On integrable orbits, enclosure growth falls from 1.79 to 0.20 bits per transit (×9) when the QR frame is enabled; with the frame disabled but the Taylor method and rigorous remainder kept, growth stays at 1.17–1.59 bits per transit, the ball-RK4 level. The gain comes from the frame rather than from the order. On numerically classified chaotic orbits the QR integrator gives ×1.51 per transit against a floating-point Lyapunov growth rate of ×1.41: it removes the representation excess and, as it must, cannot cross the dynamical floor. The ablation tests one frame policy, not all (calibration [M1]).

The implementation is cross-audited against CAPD on the identical four-dimensional analytic Case A test-field system, not on a W7-X coil field. The audit uses an independent arithmetic stack (Arb balls against filib-style double intervals), algorithm (Lohner–QR against Hermite–Obreshkov), and set representation (CAPD 2.5.1). All four components overlap at every recorded step over ten transits on both an integrable and a near-separatrix orbit, with agreement to ∼12 digits and CAPD's enclosures ∼10× tighter and contained in ours (verification bundle [V1], which carries both audit artefacts and the two CAPD drivers). CAPD does not produce the certificates, although it is tighter, because it lacks three things this paper turns on (§S2.1); as an independent implementation whose disagreement would have falsified ours, it is the largest component of the audit that substitutes for formal verification of the generator. The parametric extension of §2.5 is audited the same way. Because the adaptive step controller is not inclusion-monotone, enclosure widths of two runs are compared only on a frozen execution trace (§S2.5). The rigorous computation agrees with a floating-point replica of its criterion to four digits, at about a thousand times the cost. The rigour buys three things a replica cannot supply. The soundness defects of §S3 all produced plausible, wrong intervals, exposed only by comparing independent implementations or auditing conversion boundaries. A float detection records what one discretised orbit saw; the certificate establishes the sign change for the enclosed flow. And the parameter-box statement is a universal quantifier: no number of float samples establishes a property of every field in a box, and the widest box a given protocol certifies is a number no sampling method defines.

Table S3 prices that difference at one probe, against the fastest of the float diagnostics of §1.3. At the standard probe the weighted Birkhoff calculation took 345 s on one CPU core and classified one numerical orbit; the certificate took 3 238 s on one RTX 5090 and excluded radial-graph tori about the fitted curve through the probe. The two results are consistent. The weighted Birkhoff average reads the certified standard-configuration probe as a regular orbit whose rotation number is locked at $\iota = 1$ to ten digits. This is an island orbit, not a rotational torus about the axis. The diagnostic also reads the certified outward probe the same way and all three outward controls of §4.1 as apparently regular, with estimated rotation numbers $\iota = 0.8862,\ 0.8882,\ 1.0565$. These estimates do not prove irrationality or torus existence. The certificate's controls survive (outward controls [D12]).

S3. The coil-field kernel and its audit

S3.1. Interval-jet Biot–Savart

The right-hand side (S1) is evaluated directly from filament geometry, each straight segment contributing in the Hanson–Hirshman form

$$B \;=\; \frac{\mu_0 I}{4\pi}\,\frac{2\,(r_i + r_f)\,(\mathbf{r}_i \times \mathbf{r}_f)}{r_i r_f\big[(r_i + r_f)^2 - L^2\big]}, \tag{S10}$$

Table S3. The cost of rigour at the standard-configuration probe ($R = 6.2398244\,\mathrm{m}$, $Z = 0$, $\varphi = 0$), both rows measured. Float row: weighted Birkhoff average [5, 6] over 10^4 field-period returns (2 000 transits) of a fixed-step RK4 trace (400 steps per period; halving the step to 800 reproduces the rotation number to ten digits), classified by $\mathrm{dig}(N) = -\log_{10}|\mathrm{WB}_N - \mathrm{WB}_{N/2}|$ against the floor $\log_{10} N$: regular above floor+2.5, chaotic below floor+1.5. At a stochastic-edge point of the outward-shifted configuration ($R = 6.41\,\mathrm{m}$) the same diagnostic returns chaotic (dig = 3.9, floor 3.3). Data: calibration [M1], standard point [D2].

method	verdict at the island probe	wall	what the verdict claims
weighted Birkhoff average, float RK4	regular; rotation number locked at $\iota = 1$ to 1.7×10^{-10}; dig = 9.6 against a floor of 4.0	345 s, one core of an i3-1315U	classification of one numerical orbit, at the stated step and horizon, as quasi-periodic and resonant
converse-KAM certificate	DETECTED at 2.1547 transits, 36 745 validated steps	3 238 s, one RTX 5090	that no invariant torus of the radial-graph class passes through the point, nor through any point of its orbit, in the field of the symmetrised filament model

with $\mathbf{r}_i, \mathbf{r}_f$ the vectors from the segment's endpoints to the field point, $r_i = \|\mathbf{r}_i\|$, $r_f = \|\mathbf{r}_f\|$ and L the segment length: algebraic, hence interval-evaluable, and stable near the filament. For a validated step the sum over 4 800 segments must be evaluated as an interval-valued truncated power series in φ to order 14, twice differentiated in (R, Z) for the tangent flow and three times for the second-order parametric path, every operation outward-rounded. On one core through the arbitrary-precision path that is ∼60 s per validated step, which prices a certificate in months.

The segment sum is therefore a CUDA kernel in interval-jet arithmetic with directed rounding, the geometry reorganised so the expensive part is linear in the segment data: the cross product obeys $(x-a)\times(x-b) = x\times(a-b) + a\times b$, and the segment distances expand as $|x-a|^2 = (R^2+Z^2) - 2R(c\,a_0 + s\,a_1) - 2Z a_2 + |a|^2$ with the trigonometric series in φ shared across all segments. The certificates cost 0.088 s (standard point [D2]) and 0.116 s (outward symmetrised point [D8], narrow symmetrised interval [D9], symmetrised interval [D10]) per validated first-order step on one RTX 5090, over 3 238–5 891 s of wall clock. The withdrawn historical run of the current box ([D3]) took 0.28 s per step over 12 904 s, rejected steps included; these are not timings of the repaired run.

S3.2. The bit-identity gate, and the soundness defects

The audit of the kernel is an identity test, not a tolerance test. The same kernel source compiles against an emulation backend on the host, and the two must agree *bit for bit* on every series coefficient, at every partition of the segment set the certificates use and on the third-derivative path (verification bundle [V1]: bit identical, zero differing slots, worst difference 0 ulp, at 4 800 and at 960 segments; that artefact predates the fifth and sixth repairs of Table S4). Before the re-emitted certificates were computed, the

repaired-kernel gate passed second-order GPU/host identity at 4 800 and 960 segments, host interval conversions under all three tested rounding modes, and the GPU current-parameter battery in its fast mode (verification bundle [V1]); that pass did not repeat the third-derivative identity audit. The kernel is additionally checked against the arbitrary-precision ball path and against the analytic on-axis field of a circular loop, and every validated run re-verifies its own field and reference curve against an independent float evaluator before certifying anything.

The gate found four soundness defects. A fifth and a sixth were outside its reach and were found by review. Each of the six produced well-formed intervals that need not contain the true value, and none was visible in the answer. Table S4 describes each.

The fifth defect shows what an identity gate is worth. Host and device shared the conversion, so they agreed bit for bit on the same wrong number; an identity reference generated with the defect present pins the defect rather than catching it. Twelve of its 78 recorded components are not supersets of the same quantity re-evaluated at 240 bits (the worst by 1.4×10^{-10}), whereas all 78 regenerated components are; failure to nest under increased precision is the audit alarm we use for inward rounding (two sound enclosures need not nest, so an alarm is resolved against a sound reference enclosure). An identity gate detects only *divergence* between two implementations, never a shared error, so the reference is now accompanied by a check that consults no reference at all: every recorded component must contain the same quantity evaluated at 240 bits. This 78-component check is finite audit evidence for this path, not a proof of the conversion implementation.

The sixth defect. A certificate computed with rounded derived constants would be a statement about the field of "binary64 vertices plus rounded constants", not the filament model; the discrepancy ($\lesssim 10^{-13}$ relative) is far below every enclosure width but was not enclosed. Every coil record in this paper is computed with enclosed constants (§S4.1).

S4. Certificate data

S4.1. Which stack each record was computed on

Per-record provenance is documented in the package's `docs/certificate-stacks.md`.

Table S4. The six soundness defects found in the enclosure generator. Every coil record in this paper is computed with all six corrected (§S4.1). The derivation notes, identity artefact and corrected implementations are grouped in the verification bundle [V1].

defect	how it was found	fix
Geometry coefficients converted to doubles under round-to-nearest and thereafter treated as the exact segment endpoints: the enclosure was of the wrong field, by a relative 10^{-16} per coefficient	emulation/CUDA bit-identity test	the conversion carried outward-rounded on both sides
A rounding mode never installed: the library's directed-rounding self-test ran only from a `main()` that a library build does not have, so `FE_DOWNWARD` was never set in the process that used it. Every interval came out plausible and slightly too narrow	the same gate, once run against the library rather than the executable	the self-test moved into the library's own initialisation path
A barrier-divergent block: the kernel spreads each chunk's segments over thread rows that synchronise inside the per-segment routine, so a chunk shorter than the row count is undefined behaviour on the GPU. The per-coil-group sums the current-box certificates use (960 segments in 600 chunks) hit it; the full 4 800-segment set (8 per chunk) never did	the identity test, once extended past the full segment set to every partition the certificates use	the kernel refuses an unsafe partition and the host chooses a safe one; the $w = 10^{-8}$ point certificate rerun on the fixed kernel reproduces to every printed digit
The truncated interval convolution forming R^2 and Z^2 accumulates $4(k+1)$ rounded products and k rounded additions per coefficient; outward rounding is required per operation rather than once at the end, and the gate that establishes it is agreement with 400-bit ball arithmetic on all 1800 coefficients at series cap 15, under both rounding modes	transcription of the host interval loops into C, then a coefficient-by-coefficient check against `arb` at 400 bits	every rounded product and partial sum padded outward, in the Python, the C twin and the self-test copy in lockstep; 0 of 4500 coefficients fail
The 256-bit `arb` cosine and sine coefficients exported to the kernel through `float(b.lower())` and `float(b.upper())`: both conversions can round to the same binary64, so the exported interval need not contain the `arb` enclosure. At angle 1 both endpoints collapsed onto 0.5403023058681398. A neighbouring cancellation in the adjacent-float helper had the same character	by review of the conversion boundary (the identity gate cannot detect it, since host and device agree bit for bit on the same wrong number), and then confirmed against 240-bit ball arithmetic	outward conversion per coefficient, endpoints formed in `arb` before the binary64 conversion, and the unvalidated fallback removed
The derived segment constants — $\lvert a-b\rvert^2$, $2\times10^{-7}I$, $\lvert a\rvert^2$, $\lvert b\rvert^2$, $a\times b$ and $a-b$ — computed in binary64 on the host and passed to the kernels and their host twins as zero-width intervals. The vertices themselves were enclosed; the constants derived from them were not, by at most 2.6×10^{-16} relative (4.9×10^{-12} on small components of $a\times b$, $3.4\times10^{-15}\,\mathrm{m}^2$ absolute)	*not* by the identity gate, for the reason of the fifth: every backend packed the same doubles. Found by review of the packing boundary, then measured against `arb` from the same binary64 vertices	all derived constants built as `arb` enclosures and passed as lower/upper pairs in every kernel and host twin; kernel identity gates re-passed; every coil record recomputed (§S3.2)

Table S5. Settings and outcomes of the certificate runs quoted in the main paper. The reversible runs (criterion `th32-coil`, on the $Z = 0$, $\varphi = 0$ symmetry line) use Taylor order 14, tolerance 10^{-15} and $h_0 = 10^{-3}$; the conefield runs use order 12, tolerance 10^{-13} and $h_0 = 10^{-3}$, carry both physical tangent columns in both time directions to 3.14 transits and always use the split frame; all use precision 53 bits. The spectral Tore Supra run uses $h_0 = 0.25$ and the radial-graph class in flux coordinates of §S1.4. Outward rows are centred at $R_0 = 6.278350162606\,\mathrm{m}$. "ref" is the reference curve of §S1.2 with K harmonics; "initial set" is the half-width of the initial interval or box (R only on the symmetry line, R and Z for the conefield boxes) or of the current box; frames are those of §S2.2, "base" the 4×4 pivoted QR and "split" the block-diagonal frame. Horizon and transit are in toroidal transits; for a conefield run the transit is where the cone empties, for [D1] it is the total validated integration over its 15 runs, and a dash marks a value the run summary does not report.

record, run	probe R (m)	initial set	criterion	frame	hor.	transit	outcome
[D1] Tore Supra ED $(10, 6)$, $I_d = 4.5\,\mathrm{kA}$	ψ=0.28332	current ±4%	general-point	base	25	147.8	CERTIFIED
W7-X standard, symmetrised 50-coil model, ref axis $K = 48$							
[D2] point	6.2398244	point	reversible	base	3.344	2.15471	DETECTED
[D3] repaired current box, 2nd order	6.2398244	current $\pm 10^{-8}$	reversible	base	3.344	2.154708	DETECTED
W7-X `outward`, symmetrised 70-coil model, ref axis $K = 24$							
[D8] point	R_0	point	reversible	split	4.382	3.12830	DETECTED
[D9] interval	R_0	$\pm 10^{-10}\,\mathrm{m}$	reversible	split	4.382	3.12831	DETECTED
[D10] interval	R_0	$\pm 10^{-8}\,\mathrm{m}$	reversible	split	4.382	3.12845	DETECTED
W7-X `outward`, distributed 70-coil file, ref axis $K = 24$							
[D11] interval	R_0	$\pm 10^{-10}\,\mathrm{m}$	reversible	split	4.382	3.12831	DETECTED
[D11] interval	R_0	$\pm 10^{-8}\,\mathrm{m}$	reversible	split	4.382	3.12845	DETECTED
[D5] point	R_0	point	conefield	split	3.14	—	cone empty
[D6] box	R_0	$\pm 10^{-9}\,\mathrm{m}$	conefield	split	3.14	—	cone empty
[D7] boxes	R_0	$\pm 3 \times 10^{-9}$ and $\pm 5 \times 10^{-9}\,\mathrm{m}$	conefield	split	3.14	—	cone empty

Table S6. Finite interval witnesses on the distributed `outward` file (distributed-file intervals [D11]). Acceptance requires detection, an independent finite check and strict witness signs against the validated axis. Exact reversibility is not established, so these checks imply no exclusion. The transit records detection or where a failed run stopped.

half-width (m)	transit	finite check	validated axis	status
10^{-10}	3.1283	pass	pass	finite witness accepted
10^{-8}	3.1284	pass	pass	finite witness accepted

Table S7: Stable data records used by the paper. Exact files, roles, sizes and SHA-256 digests are in the release manifest.

ID	Record	Contents and status
D1	Reduced Tore Supra certified query and island-divertor results	Reduced three-family, rounded-input Tore Supra model: repaired current-tail scaling, eight accepted leaves and an independently checked exact-rational cover of the +/-4% current interval. Historical predictions, timings and defective verdict are preserved separately. The no-tail island-divertor results remain generator-level evidence. Status: Tore Supra certified with leaf and coverage replay; island-divertor evidence generator-level.

ID	Record	Contents and status
D2	Sigma-symmetrised standard W7-X model and its branch 3.2 certificate	The standard configuration's own coil file (coils.w7x_std_mc_focusadd, 50 modular coils, no planar coils) made exactly reversible by the same sign-bit mirror construction as D8, and the branch 3.2 certificate emitted against it at the 5/5 island probe R = 6.2398244 m with the 48-harmonic fitted Fourier reference curve, with the independent finite check. Symmetrising this file moves the field by about 1e-17 relative, so the discharged hypothesis costs essentially no model fidelity. Status: certified; reversor hypothesis exact by construction.
D3	Five coil-type current box on the sigma-symmetrised standard model	Repaired second-order five coil-type current box, relative half-width 1e-8, at D2's exact probe R=6.2398244 m with its fixed fitted reference. Complete pure current Hessians, exact memberwise reversibility, fresh enclosures and independent finite replay. The defective historical witness is retained separately and supports no exclusion. Status: certified after complete current-Hessian repair; fixed fitted reference; exact reversor memberwise.
D4	Validated outward magnetic axis	Interval-Newton periodic-orbit certificate of the outward magnetic axis and its independent check; the conefield witnesses of D5 and D6 and the interval witnesses of D11 are rechecked against this axis. Status: certified.
D5	Symmetry-free conefield certificate, outward W7-X	The distributed filament model under the conefield criterion, which carries no reversing-symmetry hypothesis: both physical tangent columns in both time directions, the combined analysis, and the exact-rational independent check. Status: certified; no reversor hypothesis; validated-axis recheck accepted; uses the published geometric theorem (external to Lean).
D6	Set-valued symmetry-free conefield box, outward W7-X	Every initial point of the section rectangle R0 +/- 1e-9 m, Z = 0 +/- 1e-9 m at phi = 0 of the distributed 70-coil filament model, under the conefield criterion (no reversing-symmetry hypothesis): both tangent columns in both time directions from the initial box, the combined analysis and the exact-rational independent check, and the recheck of the witness against the validated axis. The +/-1e-8 m box was not certified: its enclosure broke down at 3.10 transits, before the cone emptied. Status: certified positive-area section box (4e-18 m^2); validated-axis recheck accepted; uses the published geometric theorem (external to Lean).

ID	Record	Contents and status
D7	Set-valued conefield width ladder, outward W7-X	The symmetry-free conefield ladder on the distributed 70-coil outward filament model: both R0 +/- 3e-9 m, Z +/- 3e-9 m and R0 +/- 5e-9 m, Z +/- 5e-9 m at phi = 0 pass the exact-rational finite checker and their validated-axis rechecks. The widest rectangle has area 1e-16 m^2, 25 times D6. Both directions reach 3.140 transits. Status: certified positive-area section box (1e-16 m^2); validated-axis recheck accepted; uses the published geometric theorem (external to Lean).
D8	Sigma-symmetrised outward W7-X model and its branch 3.2 certificate	The exactly reversible filament model built by mirroring one representative of each sigma-pair with a sign-bit flip, its invariance and residual audit, and the branch 3.2 certificate emitted against it with the independent finite check. The reversor hypothesis is discharged by the model construction, not by the certificate. Status: certified; reversor hypothesis exact by construction; rechecked against the symmetrised model's own validated axis.
D9	Sigma-symmetrised outward symmetry-line interval certificate	R0 +/- 1e-10 m on Z = 0 on the exactly reversible outward model of D8, branch 3.2, with the independent finite check. Every point of the interval lies on the reversor's fixed line, so branch 3.2's hypothesis holds exactly at each of them; the checker's initial-set clauses accept a nonzero-width segment. Status: certified one-dimensional section interval; reversor hypothesis exact by construction; rechecked against the symmetrised model's own validated axis (D8).
D10	Sigma-symmetrised outward symmetry-line interval certificate, 1e-8	R0 +/- 1e-8 m on Z = 0 on the exactly reversible outward model of D8, branch 3.2, with the independent finite check: the exact-reversor companion of D11's 1e-8 distributed-model interval, as D9 is of the 1e-10 one. Status: certified one-dimensional section interval; reversor hypothesis exact by construction; rechecked against the symmetrised model's own validated axis (D8).
D11	Finite interval witnesses on the distributed outward model	R0 +/- 1e-10 and 1e-8 m on Z = 0 of the distributed 70-coil model, branch 3.2: both widths detected, accepted by the independent checker and rechecked against the validated axis (D4). Finite-check acceptance establishes the stored inequalities only. Exact reversibility is not established for the distributed file, so no branch 3.2 non-existence conclusion is inferred; D10 supplies the exclusion certificate on the exactly reversible companion model. The original run summary predates the separately archived validated-axis checks; the status shown includes them. Status: finite interval witnesses accepted; exact reversibility not established; no exclusion inferred (see D10).

ID	Record	Contents and status
D12	Selectivity controls on the sigma-symmetrised outward model	Three control probes on Z = 0 of D8's exactly reversible outward model, in the core, inboard and outboard, run with D8's criterion, settings and reference curve: each survives the 4.382-transit horizon without detecting, which is what makes them controls. Survival to the horizon claims nothing about tori. Status: control, not a certificate; all three survive the horizon.
M1	Method calibration and cost measurements	Background-profile shear bounds (not full-field twist), weighted-Birkhoff comparison, the independent coil-field and derivative cross-check, the segment-level reversing-symmetry audit of the coil file, and representative higher-order timing data. Also the standard-file reversor audit, the frame ablation, the inverse-problem and amplitude-sweep runs, the W7-X validation and target-selection float data, and historical current-box and null-direction runs; none of these is a certificate or a measurement of the repaired second-order protocol. Status: validated checks and empirical measurements.
V1	Independent verification bundle	GPU/host identity artefacts, including the repaired-kernel gate reports and job transcript, standard-library certificate checkers, the CAPD cross-audit on an independent arithmetic stack, and the Lean certificate semantics. Status: verification support.

S5. The independent verifier

The certificate format and a full transcript are documented in the package's `docs/certificate-format.md`.

S5.1. What `certcheck` *checks, and what it trusts*

The independent checker uses only the Python standard library: no `flint`, no `numpy`, and no imports from the enclosure generator. Every emitted bound is a double and therefore an exact rational, so the criterion's algebra is redone in exact rational interval arithmetic and there is no rounding in the checker to argue about; the one transcendental step, the reference curve's Fourier series, is evaluated from a 100-digit bound on π with an alternating-series remainder rather than from the platform `libm`. Exit status 0 is accepted, 1 rejected, 2 the file records no detection and is therefore not a certificate.

The checks are written conjunct by conjunct against `Checker.lean`. Structurally: every step length positive; $t_0[j+1]$ enclosing $t_0[j] + h[j]$, so the step times are the running sums the type demands; each witness's time being that step's end time; each witness's endpoint enclosure lying inside the a-priori box of the same step; $k < m$

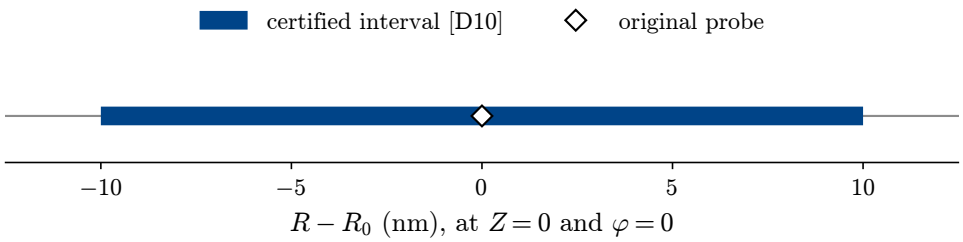


Figure S1. The certified interval on the $Z = 0$, $\varphi = 0$ symmetry line, in local radial coordinates about the original probe (symmetrised interval [D10] on the symmetrised model). The distributed-file intervals [D11] contain finite comparison witnesses on the distributed file with no exclusion inferred, because exact reversibility is not established. Every initial point on the blue segment is covered by one validated integration; the white diamond is the original probe, whose point certificates are reported separately (outward conefield point [D5], outward symmetrised point [D8]).

and $t_m > 0$. For `th32-coil`: $Z_0 = 0$ exactly, the reversor's fixed line at $\varphi = 0$; $\eta_0 = (0, \eta_Z)$, so $\mathcal{R}\eta_0 = -\eta_0$; B_φ strictly signed and of one sign over every step up to the detection; and a *recomputation* of $\beta(\eta, \xi)/B_\varphi = \eta_R(Z - Z_a(\varphi)) - \eta_Z(R - R_a(\varphi))$ from the emitted endpoint enclosures and the checker's own reference curve, with 0 excluded, the recomputed bracket required to overlap the emitted one and the two strict signs required to be opposite. The spectral criteria get the same structural checks with the criterion read on η^ϑ, plus, for Theorem 3.1, that the emitted window is exactly the steps $k \dots m$ with no gaps, that $\eta^\psi < 0$ on *every* a-priori box in it rather than at the endpoints, and that $\psi > 0$ there. Each check prints its margin. For a conefield certificate over an initial box (Remark 4 of the main text) the conefield checker verifies, in exact rationals, the two signs and the determinant bound, the sign of B_φ and the positivity of R on every step of both directions, that the offset interval has a component of strict sign (so ξ_0 vanishes nowhere on the box), that the recorded box lies in the section $\varphi = 0$, and that the offset interval is no narrower than the box. It does not recompute the offset from the fitted reference curve. For the validated-axis claim, `coneaxischeck.py` additionally binds the source witness and axis artefact by their SHA-256 hashes, replays the source cone check and interval-Newton inequality, recomputes the initial offset from the recorded Newton box, checks the forward/backward phase bookkeeping against a rational enclosure of 2π, and recomputes both witness brackets, covectors and the strict determinant. Every recomputed interval must be contained in its reported enclosure. Axis propagation, return-map and trajectory/tangent enclosures remain inputs, as does their identification with the same field model. This is a separate implementation of the finite bridge arithmetic, not an independent reproduction of the flow.

What the verifier trusts must be quoted with any use of it. The emitted enclosures are taken on faith: that the endpoint enclosure of step k contains the true state and transported tangent is `EncloseSound`, an assumption in `Checker.lean` too, discharged by the Lohner construction of §S2 and the bit-identity audit of §S3.2 and re-checked nowhere here; a generator that emitted correct arithmetic on unsound enclosures would pass. So is the hypothesis that makes criterion `th32-coil` applicable at all: the checker

verifies the acceptance predicate of the reversible criterion, and takes on faith that the field is exactly reversible, which holds by construction for the symmetrised models and only approximately for the distributed coil files (§S1; §2.4 of the main text). The magnetic field is never evaluated: the B_φ enclosure of the domain guard is read from the file, and what is checked is the sign logic that consumes it. The statement is about the filament model in the named coil file, or the spectral field with its tail status, which is upstream of every check. The global step from section homotopy to a Lipschitz radial graph runs through Birkhoff and a definite-shear hypothesis and is external to the Lean development; the local C^1 radial-graph transversality determinant is proved there. The relation $\operatorname{sign}\lambda(\eta) = \operatorname{sign}\eta^\psi$, which licenses the checker's testing η^ψ where the criterion names λ, is now proved: in this chart the two are not merely of one sign but *equal*, because $\psi = \frac{1}{2}|x-a|^2$ makes the section gradient $\xi = x - a$ and the coorientation's value on a section tangent is $\langle x-a, \eta\rangle$ (`lamXiCoil_eq_etaPsi`). That proof is for the coil section chart; the spectral Theorem 3.1 path, the only one that uses a λ condition, uses the different coorientation of §S1.4, whose reduction is not formalised. No coil certificate uses a λ condition. For the latter, Lean proves that positivity, negativity and vanishing of $\beta(\eta, \xi)$ are equivalent to the corresponding property of the scalar bracket checked in the certificate, using the chart's positive section density (`w_pos`). For $B_\varphi < 0$ the coil chart normalises orientation to use the positive density $-B_\varphi$ (§S1.1); a global sign flip cannot affect a strict sign change in any case. What the verifier removes from the trusted set is the detection logic, the reference-curve evaluation, the criterion's algebra, the domain-guard bookkeeping, the structural coherence of the step ledger and the rounding of all of those: the places a bug in the *criterion*, as opposed to one in the integrator, would live.

S6. Validation on analytic and spectral fields

S6.1. The model problem of Kallinikos, MacKay and Martínez-del-Río

On the helical perturbation of an axisymmetric toroidal field of [3], the problem the criterion was developed on, rigorous detections follow the published floating-point behaviour: the detection time scales as $t^* \propto \mathrm{Res}^{-1/2}$ in Greene's residue Res, with measured slope -0.502 at every probe depth (eq. (15) of [3]), and the prefactor agrees to 3% deep in the island (§S7.7). The inverse problem on this model is reported in §S7.3.

A field given as a spectrum is certifiable when the represented object is the field-line flux function, both field components are exact derivatives of one scalar, and the truncation tail enters as a per-step Grönwall defect; the tail is then provable, hypothesised (and stated as a hypothesis of every downstream theorem), or zero by construction (§S7.7).

S6.2. Tore Supra ergodic divertor: model and certified query

We define a reduced model inspired by [18]: a cylindrical power-law spectrum with only the three toroidal families $n = 6, 18, 30$, applied also to low poloidal harmonics.

Families $n \geq 42$ are excluded from the model, not enclosed in its tail. For $1 \leq |m| \leq 130$ the amplitudes and equilibrium polynomial inputs are the recorded binary64 values interpreted exactly; the poloidal completion for $|m| > 130$ uses the analytic sinc formula with the recorded constants interpreted exactly. This convention defines a model, not an enclosure of the ideal published spectrum or of the device. Its poloidal tail has $|a_{mn}| \leq (K/|m|)\zeta^{|m|}$, $\zeta < 1$, including the trimmed candidate modes.

The repaired run certifies exclusion at current $\pm 4\%$, with eight sub-boxes, 15 runs, 147.8 transits and 2464 s (Figure S2; Tore Supra current box [D1]). It multiplies the discarded-spectrum defect bounds by $\sup_{\varepsilon \in P} |\varepsilon|$, correcting the historical tail-scaling omission. Every successful leaf is emitted with the exact-input model ledger's hash. The independent checker accepts all eight leaves and their exact-rational coverage of the requested current interval. This restores the exclusion for the defined reduced model; it supplies no full-field twist certificate or engineering tolerance. The historical 2768 s report remains archived separately. Settings are in Table S5.

S7. Extended treatments

The main text states each of the following points briefly; the arguments, data and qualifications are given here in full.

S7.1. The coil-field upgrade route: the twist reduction and the cylinder theorem

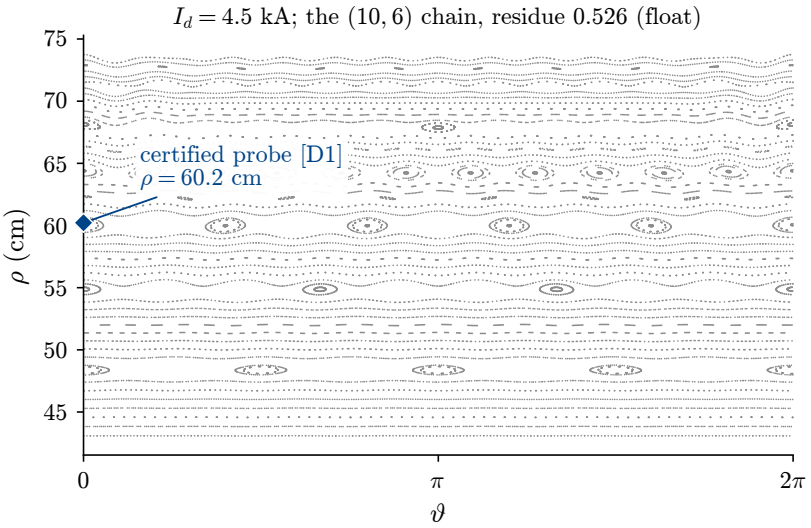


Figure S2. The reduced three-family Tore Supra ED model section at the nominal query current (grey punctures, floating-point context), with the probe (blue diamond) inside the $(10, 6)$ chain; its resonance reduces to 5/3, so it appears as five islands. The eight repaired leaves cover the current interval, not an initial-condition region in this section, so no subdivision is drawn. Data: Tore Supra current box [D1].

Corollary S1 (section-homotopic extension; [2] §5)**.** *Let $\mathcal{T}$ be the torus of angles (ϑ, φ) and c a fixed radial coordinate in an interval. Assume, on the domain $\mathcal{T} \times$ interval: (i) definite shear for the full perturbed field, with the twist, invariant-domain and section hypotheses required by the Birkhoff argument; and (ii) the hypotheses under which [2, §5] invokes Birkhoff's theorem in the form of [21, §44], namely that the method "applies equally well to all Lipschitz tori topologically transverse to the given direction field" and that "every invariant torus that is homotopic to $\mathcal{T} \times c$ is transverse to ξ and Lipschitz". Then the conclusion of Theorem 1 upgrades from "no radial-graph torus through this point" to "no invariant torus homotopic to the section through this point".*

Hypothesis (ii) is quoted, not re-derived. The formalisation proves only the local C^1 geometry (for a positive radial graph the determinant of the radial direction with its angular tangent is $r^2 > 0$); the global bridge needs twist-map theory outside the formal scope. No certified statement depends on the corollary. Table S2 bounds only the background profile ι', not the shear of the full field. Write a spectral perturbation of the field-line Hamiltonian as modes with amplitudes a_{mn}, radial profiles $p_{mn}(\psi)$ and constant phases ζ_{mn}, with primes denoting ψ-derivatives (§S1.3). Then $\partial_\psi V^\vartheta = \iota' + \sum_{m,n} a_{mn} p''_{mn} \cos(m\vartheta - n\varphi + \zeta_{mn})$ also contains the retained modes and the tail. Its sign, and the global hypotheses needed for the application of Birkhoff's graph theorem, have not been certified. This restriction applies to the spectral models as well as W7-X: all exclusion statements concern the radial-graph class alone. For coil fields, an enclosed return-map twist together with the section-density, invariant-cylinder and all-phase hypotheses would be a possible route through [17, §2]; endpoint shear alone does not discharge these obligations.

The obstruction to Corollary S1 on a coil field is the verification of hypothesis (i), not

the theorem, and it does not in fact need an ι profile. In the canonical section coordinates determined by the field itself, $\omega = \rho\,\mathrm{d}\theta \wedge \mathrm{d}r$ with $\rho = JB^{\varphi} > 0$ and $I(\theta, r) = \int \rho\,\mathrm{d}u$, the chart derivative is lower triangular, $D\kappa = \left(\begin{smallmatrix} 1 & 0 \\ q & \rho \end{smallmatrix}\right)$. Write $M = (m_{ij})$ for the endpoint variational matrix expressed in the (θ, r) section chart, obtained from the integrated (R, Z) matrix by conjugating with the polar chart derivatives at the endpoints. Direct multiplication gives

$$\partial_I \Theta \;=\; \frac{m_{12}}{\rho_0}. \tag{S11}$$

The *sign* of the twist is therefore decided by the endpoint shear m_{12} alone — a quantity the runs of §4.2 already integrate — together with a positive lower bound on ρ. Errors in q, in the endpoint chart, and in any rectangular normalisation enlarge the constants in the Lipschitz bound, but cannot change whether a positive-twist certificate is attainable. Two cautions attach to such a calculation, both elementary. First, an instantaneous shear condition does not suffice: on a cylinder, $H = \frac{1}{2}I^2 + \frac{1}{4}(1 - \cos 2\pi\theta)$ has $H_{II} \equiv 1 > 0$, yet the linearisation at the equilibrium is $\delta\theta'' + \pi^2\delta\theta = 0$, so $\delta\theta(1) = (\sin\pi/\pi)\,\delta I(0) = 0$ and the time-one map has $\partial_I\Theta = 0$ there; the certifiable object is the endpoint quantity (S11), not a sign condition on the field. Second, twist certified on one section constrains that section only: an exact symplectic suspension can have a positive-twist return map while its suspended torus acquires vertical tangencies at an intermediate phase, so propagation to all phases needs either a transport-cone bound along the one-turn tube or a bound on the spatial derivatives of the field-line vector field. We have not carried out that computation, and the W7-X statement in the main text is accordingly made for the radial-graph class alone; we record the reduction because it replaces an apparently expensive hypothesis with two enclosures the existing machinery already produces.

A second point concerns which form of Birkhoff's theorem the upgrade should invoke. The classical statements are for an annulus: [20, §3] assumes an analytic, area-preserving map *of a closed annulus into itself*, and [21, §44] likewise. Neither hypothesis is available for a coil field, where no invariant annulus is known and assuming one would be close to circular in an argument whose conclusion is non-existence. The form that does apply is Mather's [17, §2], stated on the cylinder: for a C^1 area-preserving, orientation- and end-preserving diffeomorphism of $\mathbb{T}^1 \times \mathbb{R}$ that is tilted, and an open g-invariant $U \cong \mathbb{T}^1 \times \mathbb{R}$ with $\mathbb{T}^1 \times (-\infty, a] \subseteq U \subseteq \mathbb{T}^1 \times (-\infty, b)$, the frontier of U is a Lipschitz graph; and a homotopically non-trivial invariant circle supplies such a U as a component of its complement, by Schoenflies. The invariance the argument consumes is therefore the invariance of the candidate circle itself, not of an ambient annulus, and the regularity delivered is Lipschitz, which is the class the criterion needs. Mather notes that this version is not a consequence of the classical ones though it is proved similarly; a detailed proof of a close variant is in [14, Chap. I]. What remains for a coil field is then an extension rather than an invariant annulus: a C^1, area-preserving, orientation- and end-preserving tilted diffeomorphism of the cylinder agreeing with the validated return map on the certified region, together with the coverage condition that the candidate circle lie in that region. We flag this as the route, not as a result of the main paper.

S7.2. The two formal statements of Theorem 3.2

Theorem 3.2 of [3] is proved in the Lean development in both forms, the published sign-change form and the vanishing form its proof uses. The two formal statements are:

```
theorem converse_kam_th32_zero (F : Foliation S) (x0 : R3)
    (rev : Reversor S F) (hmem : x0 in S.U)
    (hline : x0 in rev.symmetryLine)
    (eta0 : R3) (hanti : rev.R eta0 = -eta0)
    (hindep : forall c : Real, eta0 != c * S.B x0)
    (T : Real) (_hT : 0 < T) (hev : Event32 F x0 eta0 T) :
    not (exists Tor : TransverseTorus S F, x0 in Tor.carrier)

theorem converse_kam_th32 (F : Foliation S) (x0 : R3)
    (cert : Certificate32 F x0) :
    not (exists T : TransverseTorus S F, x0 in T.carrier)
```

(Unicode transliterated: `in` for $\in$, `forall`/`exists`/`not` for $\forall/\exists/\neg$, `*` for scalar multiplication.) Here `Event32` is the vanishing of $\beta(\eta(t),\xi)$ at some $t \in (0,T]$ along the transported displacement started at η_0, `Certificate32` packages the symmetry hypotheses with a strict sign change, and the second theorem is proved from the first by the lemma that a sign change contains a zero.

S7.3. The inverse problem: certified tolerances

The engineering question inverts the certificate: at a target point y, the widest tolerance $w^*(y)$ for which y is certified to lie on no radial-graph torus for every member of the box. The complementary statement, "the no-surface region does not extend beyond y", is not certifiable by this method and is not made.

A scalar half-width is a convenient special case of the output. In particular, d parameters do not cost exponentially many boxes. The integrator does not tile the parameter space: it carries one exactly-propagated sensitivity vector per parameter, a zonotope linear in d.

Proposition S1 (certified tolerance body)**.** *Orient the finite witness inequalities so their central values are positive, and suppose that on a centrally symmetric convex $P \subset \mathbb{R}^d$ they have validated first-order forms $c_\ell(p) = m_\ell + g_\ell^{\mathsf{T}} p + e_\ell(p)$ with $m_\ell > 0$ and $|e_\ell(p)| \le r_\ell(P)$, for $\ell = 1,\dots,L$. With support function $h_P(g) = \sup_{p\in P} g^{\mathsf{T}} p$, the strict conditions $h_P(g_\ell) + r_\ell(P) < m_\ell$ certify every parameter in P. In the exact linear model $r_\ell = 0$ the largest centred convex admissible set is the intersection of slabs $\bigcap_\ell \{p : |g_\ell^{\mathsf{T}} p| < m_\ell\}$.*

Proof. Central symmetry gives $\inf_{p\in P} g_\ell^{\mathsf{T}} p = -h_P(g_\ell)$, so $c_\ell(p) \ge m_\ell - h_P(g_\ell) - r_\ell(P) > 0$ throughout P and every strict witness sign is preserved. With the remainder vanishing, every admissible centred convex set lies in the slab intersection, and every point p of the intersection is certified (take P the segment from $-p$ to p), which is maximality. □

For an axis-aligned box $P = \prod_j [-w_j, w_j]$ this reduces to $h_P(g_\ell) = \sum_j |g_{\ell j}| w_j$, which is where the scalar w^* below comes from. For engineering use the natural output is the body, not the scalar: tolerances on different coil types are not interchangeable, and the slab intersection says which combinations are admissible rather than forcing a cube.

Bisection with the certificate as oracle computes a certified w^*, a lower bound on the widest certifiable width since the adaptive oracle is not monotone in the width (§S2.5); on the model problem an exact answer exists for comparison, from the analytic island edge of the least-perturbed member, because the certified-for-all set is an intersection over the box. The single-box oracle is 70–95× conservative and merely restates the measured enclosure wall. The subdivision oracle achieves 14–22% relative conservatism, roughly constant over six targets and two nominal amplitudes, with the sources identified: finite horizon, the decoherence-scale sub-box floor near threshold, and bisection resolution. Numerically, the generator reports ±24.6% validated against an exact ±29.8%, from sub-box certificates in 603 s on one core (calibration [M1]; the four-probe figure is Figure S3).

S7.4. The certifiable and physical half-widths across the amplitude-ratio sweep

Across a sweep of the amplitude ratio ϱ on a controlled Hamiltonian family (calibration evidence of [M1]), the certifiable half-width falls by four decades while the physical threshold w_{phys}, defined in §S2.4 by float tracing alone, does not measurably move (Table S8).

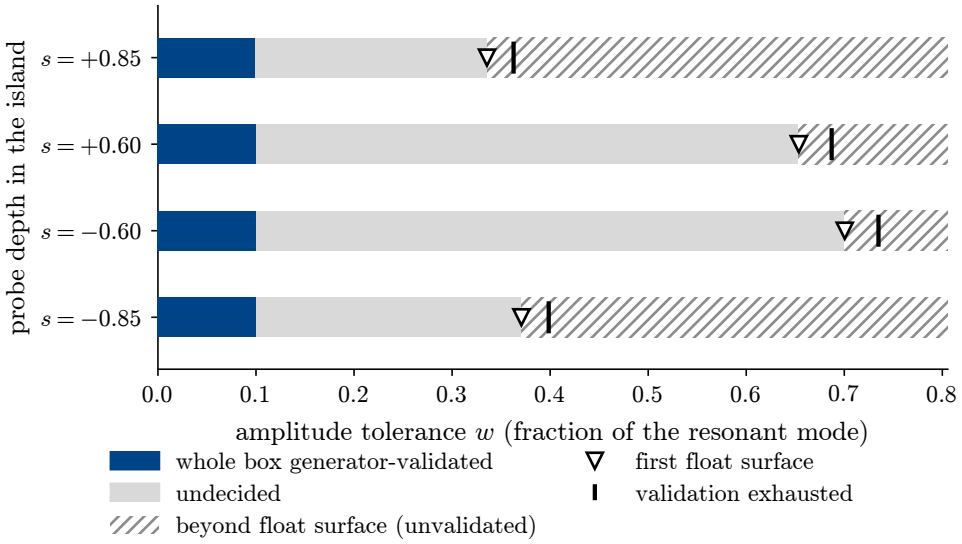


Figure S3. Generator-validated two-sided tolerance at four probe depths in the island-divertor model: what the certificate reaches (blue), where a float section first shows a surface (triangle), and the interval between them that neither side decides (grey); the hatched band beyond the first float surface is not validated. The generator-validated tolerance is a lower bound on the true limit and is reported as one; the existence-side statement, that the no-surface region does not extend beyond the probe, is not certifiable by this method and is not made. Data: Tore Supra current box [D1].

Table S8. The physical threshold w_{phys}, censored at the top of the tested range, against the certifiable half-width w_{sub} across the amplitude ratio ϱ of the controlled family; the ratio is the gap between them.

ϱ	w_{phys}	w_{sub} (certifiable)	ratio
10	≥ 0.3	9.85×10^{-4}	$\geq 3.0 \times 10^{2}$
30	≥ 0.3	3.44×10^{-5}	$\geq 8.7 \times 10^{3}$
70	≥ 0.3	4.13×10^{-7}	$\geq 7.3 \times 10^{5}$
100	≥ 0.3	4.19×10^{-8}	$\geq 7.2 \times 10^{6}$

S7.5. The reversor audit, the counterexample and the two routes

The conefield criterion is defined in §2.3 of the main text.

The reversor is not exact for these coil files, and what that costs. The reversible criterion is available only because the field has an *exact* reversing symmetry, and we have checked the distributed filament models segment by segment against $\sigma(R, \varphi, Z) = (R, -\varphi, -Z)$, counting a directed segment (a, b, I) as matched when $(\sigma a, \sigma b, -I)$, or equivalently $(\sigma b, \sigma a, I)$, is also in the model. The symmetrised models replace one coil of each σ-pair by the σ-image of the other, traversed in reverse order at the same current. *The two configurations use different coil files and the answer differs between them; the figures are not transferable.* The outward configuration runs `coils.w7x_vmecpp_tests` [11]: 70 coils in seven groups, of which 20 are planar at current factors -0.14 and $+0.14$ of nominal. Its 50 modular coils are symmetric to file rounding, 10^{-14} to 10^{-12} m, but all 1920 planar segments lack an exact partner, by up to 1.6×10^{-5} m, and the field identity $B(\sigma x) = -D\sigma\, B(x)$ consequently holds only to 4×10^{-9}–2.2×10^{-8} relative at interior points (calibration [M1]). That mismatch is not an artefact of how the file discretises a symmetric curve: projecting each σ-image vertex onto the partner coil's polyline leaves 10^{-5} m, and so does removing the best rigid motion. It is a difference of shape (outward symmetrised point [D8]). The standard configuration runs a different file, `coils.w7x_std_mc_focusadd` [10]: 50 modular coils in five groups at 1.62 MA and no planar coils at all. There 4716 of its 4800 directed segments already have an exact partner (the same count as the non-planar part of the outward file) and the misses are last-digit file rounding, 10^{-14} to 10^{-12} m, so its field identity holds to 6×10^{-17}–1×10^{-14} relative. That is six orders tighter than the outward file, for an unrelated reason, but still not the exact reversibility that the reversible criterion requires. Both audits are one script over the distributed files and take seconds to repeat (calibration [M1]).

An approximate reversor does not weaken the reversible criterion's conclusion; it removes the criterion's hypothesis. Nor can a reversible run be reinterpreted under the general-point criterion: that criterion requires the transported displacement to be started along $\xi(s_0)$, up to a positive multiple and an added component along B, and these runs start at ∂_Z, which is not of that form; §S7.5.1 gives an explicit counterexample. We close the gap in two independent ways, which answer different questions and neither of which supersedes the other.

(i) Change the criterion, keep the model. The conefield criterion (§2.4) carries no symmetry hypothesis. On the distributed model, carrying both tangent columns in both time directions to 3.14 transits, the cone is empty on a determinant enclosure $[0.0172400, 0.0173565]$, and the independent checker recomputes its lower bound as an exact rational and accepts (outward conefield point [D5]). The certificate uses the geometric criterion of [2], proved in the main text but external to Lean. This geometric argument is distinct from the missing exact-reversibility hypothesis of the distributed-file intervals [D11].

(ii) Change the model, keep the criterion. σ acts on the 70 coils as a fixed-point-free involution, so one representative of each of the 35 pairs generates the other. Negating the y and z coordinates of a representative is a sign-bit flip, exact in IEEE-754, so the generated segment list is exactly σ-invariant as a mathematical object: all 6720 directed segments have an exact partner, against 4716 of 6720 for the distributed file, so the reversible criterion's hypothesis holds exactly. This is a different model: the mirrored half of each pair moves by up to $1.4 \times 10^{-5}\,\mathrm{m}$. Its field differs from the distributed one by 7×10^{-9} relative at the certified point, and the reversible criterion certifies on it at 3.12830 transits, with the independent finite check accepting (outward symmetrised point [D8]). The reversor hypothesis here is established by the model construction and its exact combinatorial audit, *not* by the certificate, which cannot see it; the checker's transcript continues to say so.

The exact-rational determinant lower bound for [D5] is

$$895159250693784536996563186479 31/2^{112}.$$

Why a transported vector, and where the conefield is used. Most certificates in this paper transport a single vector from a single initial condition, and the subdivision of §2.5 is in the parameter. We chose the vector because it is the object the validated integrator transports exactly, the tangent being one more row of the Lohner system, and because the criterion then reduces to signs of two scalars. Sets of initial conditions are carried by the same enclosure machinery: the symmetry-line intervals [D9], [D10] and [D11] transport one vector from every point of an interval on $Z = 0$. The conefield formulation of [2], which carries a set of directions and needs no symmetry, is used for the point certificate [D5] and for the boxes [D6] and [D7], all on the distributed file. Its reach is set by the wrapping of the variational block: a $\pm 10^{-8}\,\mathrm{m}$ box in R and Z was lost at 3.10 transits, just before the point's cone empties in a double-precision run at 3.1285, while boxes of half-widths 10^{-9}, 3×10^{-9} and $5 \times 10^{-9}\,\mathrm{m}$ certify, each with an accepted validated-axis recheck. The widest has area $10^{-16}\,\mathrm{m}^2$ (conefield box [D7]). Whether a cone transport can pay for a physically sized box on a coil field remains open.

S7.5.1. An arbitrary initial tangent does not suffice for the general-point criterion Since the coil files as distributed are not exactly reversible, one might keep the runs' $\eta(0) = \partial_Z$, note that the λ condition holds, and appeal to the general-point criterion with a general

B-independent initial tangent. That is false; the following counterexample is explicit and is checked symbolically and numerically (verification bundle [V1]).

On $\mathbb{T}^3$ with volume form $\Omega = dx \wedge dy \wedge dz$, put $s(x,y) = \sin 2\pi(x-y)$ and

$$B \;=\; s\,(\partial_x + \partial_y) \;+\; \partial_z, \qquad \xi \;=\; \partial_x.$$

B is nowhere zero and divergence free. The torus $T = \{x - y = 0\}$ is invariant, since $d(x-y)(B) = s - s = 0$, it is transverse to ξ, since $d(x-y)(\xi) = 1$, and it passes through the origin.

Take the initial point $s_0 = 0$ and the initial tangent $\eta(0) = \partial_y$, the analogue of the ∂_Z that the coil runs use. The orbit of the origin is $(0, 0, t)$, on which $s = 0$, so the variational equation has the constant coefficient matrix $DB = 2\pi \left(\begin{smallmatrix} 1 & -1 & 0 \\ 1 & -1 & 0 \\ 0 & 0 & 0 \end{smallmatrix}\right)$, whose square vanishes. Hence the transported tangent is exactly

$$\eta(t) \;=\; (-2\pi t,\; 1 - 2\pi t,\; 0),$$

and on the orbit $\beta(\eta, \xi) = -\eta_y = 2\pi t - 1$ and $\lambda(\eta) = \eta_x = -2\pi t$. So β changes sign strictly at $t_* = 1/2\pi$, and $\lambda < 0$ for every $t > 0$: *both* inequalities of the general-point criterion hold at $\eta(0) = \partial_y$, and an invariant torus of the class passes through the initial point regardless.

The mechanism is that $\eta(0)$ starts on the wrong side of the hypothetical torus and stays there; at t_* the transported tangent has reached $-\xi$, not ξ. Choosing $\eta(0) = \xi(s_0)$ is what fixes the side without knowing the torus, and the only admissible freedom is $\eta(0) = a\,\xi_0 + b\,B_0$ with $a > 0$. Here ∂_y is not in the span of $\xi_0 = \partial_x$ and $B_0 = \partial_z$ at all, so no such (a, b) exists. This is why the approximate-reversor problem of §2.4 could not be dissolved by re-reading the existing runs as the general-point criterion, and why it needed either a symmetry-free criterion or an exactly symmetric model.

S7.6. The special cases of the parameter-box theorem, and its representation

The main text retains the two W7-X cases, Remark 3 for reversing symmetry and Remark 4 for conefields. The general-point statement and the covering argument are recorded here.

Theorem S1 (parameter-box soundness)**.** *Let $P = [\varepsilon_{\mathrm{lo}}, \varepsilon_{\mathrm{hi}}]$ (or a product of such intervals) and let the validated integrator, started from $\{(s_0, \xi(s_0))\} \times P$, produce the strict sign certificates of §2 on its enclosures: $\eta^{\vartheta} > 0$ and $\eta^{\vartheta} < 0$ on consecutive accepted steps, with $\eta^{\psi} < 0$ on the a-priori enclosure of the bracketing step. Then for every $\varepsilon \in P$, no radial-graph torus of the field with parameter ε, where the class is defined relative to the single fixed reference curve supplied to the computation, passes through s_0, nor through any point of its orbit.*

Proof. The enclosures rigorously contain the flow image of every initial point, in particular the member trajectory $(s_\varepsilon(t), \eta_\varepsilon(t), \varepsilon)$ for each $\varepsilon \in P$; strict inequalities on a set hold for each element. By the slice property those inequalities are, member by member, exactly the single-field certificate of Theorem 3.1 of [3] in the reduction of §2. □

Corollary S2 (initial-condition sets, for either criterion)**.** *Let X be a set of initial points and P a parameter box, and let the integrator, started from $\{(s, \eta(s)) : s \in X\} \times P$ with $\eta(s)$ admissible for the criterion invoked at each s, produce enclosures satisfying that criterion's strict inequalities. Then for every $s \in X$ and every $\varepsilon \in P$, no radial-graph torus of the field with parameter ε passes through s or any point of its orbit. Finite unions of such sets are certified, and their total measure is a lower bound on the measure of the no-torus region.*

Proof. As for Remark 3: the enclosures contain each member $(s, \eta(s), \varepsilon)$, a strict inequality on a set holds at each element, and a union of certified sets is certified pointwise. □

The augmented five-variable system was re-audited against CAPD: all 80 checkpoints overlap in all five components, including the two certificate quantities and the tangent invariant $W = \mathrm{d}\Psi(\eta)$, where Ψ is the helical-flux first integral of the single-mode audit field.

Every query returns one of three verdicts, mutually exclusive by construction: CERTIFIED (Theorem S1 for the box), VIOLATED (an exhibited member for which the property fails), or EXHAUSTED (an explicit list of uncertified sub-boxes). These are generator verdicts; independently checked status additionally requires archived finite evidence and coverage. A VIOLATED verdict carries a grade: "exact" when the violation is established in the integrable normal form, "numeric" when it rests on float integration. Grade-numeric violations are float *evidence*, labelled as such at every use, never certificates (§1.4).

Pointwise justification of the initial-box statement. The pointwise enclosure proof of Remark 4 is given in the main text.

S7.7. The model problem per depth, and when a spectrum is certifiable

Table S9 gives the model-problem numbers per depth behind the slope, prefactor and conservatism quoted in §S6.1.

When a spectrum is certifiable. Three conditions precede any certificate about a configuration given as data rather than a formula. The represented object must be the field-line flux function, not $|B|$ (a Boozer $|B|$ spectrum does not determine field lines). Both field components must be exact derivatives of one scalar, so the Hamiltonian structure is exact for every member of a parameter box. In this implementation the tail enters through per-step defect bounds (a Grönwall argument); differentiating a value-only tail bound would not be justified, because a bound on f does not bound f', and the certificate consumes two derivatives of the discarded profiles. The logical status of each certificate is then fixed by a trichotomy: (a) *provable tail*; (b) *hypothesised tail*, stated as a hypothesis of every downstream theorem, in those words; (c) *zero tail by construction*, meaning that the certified object is the model defined by the spectrum file.

Table S9. The model problem of [3], rigorous code against the published float behaviour, at three relative depths s inside the island (three to four residues per depth, Res $= 10^{-4}$ to 10^{-1}). Slope: fitted exponent of the rigorous detection time in the residue (float ancestor $-1/2$, eq. (15) of [3]); ratio: rigorous detection time over the float prediction across the residues at that depth; α: fitted exponent of the certifiable half-width in the residue with its OLS error (at $s = 0.9$ the fit has three points and one residual degree of freedom, so its formal error is not a meaningful uncertainty; the pooled value is the one used). Data: Tore Supra current box [D1].

depth s	slope of t^*(Res)	$t^*_{\mathrm{rig}}/t^*_{\mathrm{float}}$	α
0.05	-0.502	1.01–1.03	0.474 ± 0.013
0.5	-0.502	1.27–1.29	0.482 ± 0.011
0.9	-0.502	1.41–1.42	$0.452 \pm 2 \times 10^{-5}$
pooled			0.4735 ± 0.0069

Validated spectral evaluation itself is an established technique ([15, 16, 19]; [1] §3); the contribution is the field-line-Hamiltonian spectrum with a rigorously bounded truncation tail consumed by a converse-KAM certificate.

S7.8. The W7-X validation and target selection

Validation of the standard configuration. From the filament file alone, our own evaluation places the magnetic axis at $R = 5.94859\,\mathrm{m}$ in the bean plane with a one-period return residual of 4.8×10^{-15}, and returns a flat core rotational transform rising through $\iota = 0.857$ to $\iota = 1$ near $r = 0.30\,\mathrm{m}$, both independent reproductions of known properties of the standard configuration; the second is the $\iota = 1$ resonance on which the island divertor is built. Refining the one-turn section map on the symmetry line locates that chain's O-point at $R = 6.237141\,\mathrm{m}$ with closure 5.5×10^{-14} and Greene residue $\mathrm{Res}_{5/5} = 8.167 \times 10^{-2}$, a strongly driven chain, two to three decades above the quiet-core band. The on-axis field is $2.911\,\mathrm{T}$ rather than the nominal $2.5\,\mathrm{T}$, because the file's $15\,\mathrm{kA/turn}$ is a higher-field variant of the same geometry. This changes nothing: the field-line equation $\mathrm{d}R/\mathrm{d}\varphi = RB_R/B_\varphi$ depends on B only through $\hat{B}$, so trajectories, the axis, ι, the island chain and its residue are unchanged by an overall current scale, and only the ratios between coil groups carry dynamics. Both public W7-X vacuum equilibria stop below $\iota = 1$ (standard 0.856342 to 0.967619, high mirror 0.870032 to 0.961754), so the standard 5/5 chain, with its O-point outside the equilibrium boundary at $R = 6.209381\,\mathrm{m}$, is the island divertor and a certificate there contradicts no nested-surface claim.

Selection of the ***`outward`*** *target.* A negative planar current lifts the transform profile and drags the $\iota = 1$ surface inward from the divertor into the plasma, so the 5/5 chain, a fixed point of the one-turn map, can be moved into the confined region by a published choice of currents. We scanned the nineteen configurations of Andreeva, Kisslinger and Wobig [7, Table 1], cross-checking the current reading three ways (the high-mirror row reproduces the ratios of Geiger *et al.* [8]; the public VMEC input of that configuration carries the same ratios to all printed digits; the resulting filament currents lie inside the operating ranges of Davies *et al.* [4]) and calibrating the planar

sign convention against Geiger's statement of which sign lowers ι. Because the knob is the planar current, the `outward` records use the 70-coil filament file. The target is row H, per-coil-type factors $[1.04, 1.04, 1.01, 0.96, 0.96, -0.14, +0.14]$ at a nominal filament current of 1.46 MA; its 5/5 O-point is at $R = 6.276843\,\mathrm{m}$, $Z = 0$, $\varphi = 0$, with Greene residue 5.369×10^{-2} and pendulum half-width 6.74 mm, and the probe at relative depth $s = 0.05$ is $R = 6.278350162606\,\mathrm{m}$.

Where the probe sits. On the outboard $Z = 0$ line of the floating-point survey the chain's predicted resonant radius is 0.625 of the contiguously confined radius (the distance from the axis to the outermost point inside which every traced line stayed confined), and the probe itself, at 0.304 m from the axis, is at 0.69; in area-equivalent radius $\sqrt{A/A_{\mathrm{out}}}$ the chain sits at 0.71. The denominator is a coarse float quantity: 60-turn traces on an offset grid of spacing 0.04 m resolve it only to $\pm 0.04\,\mathrm{m}$, the fraction to about ± 0.03, and a longer trace can only move it inward. Nothing in any certificate depends on this number. The class the certificates exclude is radial graphs about the axis in every section, so we checked that the surfaces here have that shape: traced over one field period and sampled in 40 sections, both the surface just inside the chain and the outermost traced surface are single-valued in polar angle about the fitted Fourier curve in every sampled section, with neighbouring radii differing by at most 10.1% and no angular gap above 0.17 rad (floating-point context, provenance rather than certificate data).

S7.9. Finite-segment application of the criteria

The following finite-segment argument explains the physical application without assuming a complete flow on the whole numerical working domain; it is a mathematical argument outside the current Lean development.

Fix a hypothetical C^1 invariant radial-graph torus through s_0 and a finite tested segment lying in an open coil-free domain where $B_\varphi \neq 0$ and the reference distance is positive. On a neighbourhood of that segment the filament field and its normalised field are smooth. Local ODE existence and uniqueness give differentiable passage maps along the segment and invertible tangent maps. Invariance transports the hypothetical torus's tangent plane along these maps. Choose its continuous coorientation α_t with $\alpha_t(\xi) > 0$; transversality supplies this choice along the compact segment.

For the general-point criterion, start with $\eta_0 = a\xi(s_0) + bB(s_0)$, $a > 0$. Then $\alpha_0(\eta_0) > 0$. Tangent-plane invariance and invertibility imply $\alpha_t(\eta_t) \neq 0$ along the segment, so continuity preserves its positive sign. At the certified zero of $\beta(\eta, \xi)$, the one-dimensional kernel of β gives $\eta = c\xi + dB$. Since $\lambda(B) = 0$, $\lambda(\xi) > 0$ and $\lambda(\eta) < 0$, one has $c < 0$, contradicting $\alpha_t(\eta_t) = c\alpha_t(\xi) > 0$.

For the reversible criterion, use the segment together with its reflection under the exact reversor; equivariance of the recorded reference curve is also required. At a zero at time $t_* > 0$, write $\eta_{t_*} = c\xi + dB$. If $c = 0$, invertibility and transport of the field direction force η_0 to be parallel to $B(s_0)$, contrary to the initial-vector hypothesis. If $c \neq 0$, the reversal laws give $\eta_{-t_*} = -c\xi + dB$. The coorientation evaluations at the

two ends therefore have opposite signs. They cannot vanish unless η_0 is tangent to the torus; in that case they vanish throughout, contradicting $c \neq 0$. Both cases exclude the hypothetical torus. No existence outside the reflected finite segment is used.

Changing from physical field-line time to toroidal-angle time changes a tangent passage vector by a multiple of B; the flux covector annihilates that multiple. The nonzero B_φ guard makes this passage-time change regular. The conefield application likewise uses only its finite forward and backward segments and the geometric criterion proved in the main text. These arguments explain which local hypotheses must be checked; they do not verify the numerical enclosures, construct a Lean local flow, or prove checker equivalence. The Birkhoff/Lipschitz and 3D-to-section bridge that §5 attributes to [2, §5] is intentionally not declared as a Lean theorem. The formal sources and their machine manifest are included in verification bundle [V1].

S7.10. Hypotheses the formalisation made explicit

These are the places where the informal statements of [3, 2], or the code, left something implicit that the Lean type checker forced into the open. The Lean development's documentation (its README, "Findings") gives each with the declarations involved.

(i) *The class is defined by transversality, not by being a graph.* Theorem 3.1 excludes tori transverse to ξ (`TransverseTorus`); the radial graphs are a separate predicate (`PrincipalTorus`). Smooth positive-radius radial graphs are transverse to ξ by the local determinant computation, with no shear hypothesis; that computation is proved in Lean. The construction of the invariant tangent plane and coorientation and the 3D-to-section correspondence remain external. Birkhoff's theorem and a definite-shear hypothesis enter only to enlarge the class to section-homotopic (Lipschitz-graph) tori (§5).

(ii) *ξ must be independent of B wherever the criterion is evaluated*, as a pointwise condition on the whole window; the runs enforce it as a domain guard.

(iii) *The λ window is structural*: the intermediate-value step supplies a zero somewhere in the bracketing step, at an unknown time, so $\lambda(\eta) < 0$ is required at every time of that step. The runs certify $\eta^\psi < 0$ on the step's a-priori enclosure; at the crossing time this gives $\lambda < 0$, because there $\operatorname{sign}\lambda(\eta) = \operatorname{sign}\eta^\psi$ (S1.4).

(iv) *Theorem 3.2's published hypothesis is stronger than its proof needs*: vanishing of $\beta(\eta, \xi)$ suffices; the code tests the sign change, which is sound either way.

(v) *An orientation convention is needed and is stated nowhere*: the covector cutting out the torus's tangent plane and λ must both be normalised positive on ξ; the coorientation is a field of the torus structure.

(vi) *Regularity*: the proof is a continuity argument on the tangent plane, so it needs a C^1 torus; the Lipschitz case of [2] is external.

(vii) *Divergence-freeness enters only through flow-invariance of β*, and for a coil field through transport of the section density B_φ.

(viii) *Completeness of the flow* is assumed; the code instead carries a domain guard and stops.

(ix) *The sign of the section density*: the bracket of (1) has the sign of $\operatorname{sgn}(B_\varphi)\,\beta$, so its zeros and sign changes are those of β, only because B_φ has one strict sign on the working annulus.

(x) *The reference curve is a curve, and it fixes the class*: ξ inherits its φ-dependence, classes for different curves are incomparable, and the criterion needs the curve only C^1, analyticity being a property of the Fourier representation.

(xi) *The parameter-box theorem has no content beyond the base theorem*: it is a one-line quantifier statement.

(xii) *The coil β and λ reductions are proved; the spectral λ reduction is not*, and it does not enter the coil path.

(xiii) *Lipschitz in which variable*: resolved as Lipschitz in ϑ uniformly in φ and continuous in φ.

(xiv) *Continuity of the tangent flow in time* must be assumed; without it there is no intermediate-value argument at all.

(xv) *Continuity of β in the base point* must likewise be assumed.

(xvi) *What the proofs consume*: flow-invariance of β is never used in the proof of Theorem 3.1, and checker acceptance does not imply the independence of ξ and B at the initial point, which is an explicit hypothesis of `checker_sound31`.

(xvii) *"Reversible" means more than the published hypotheses say*: Theorem 3.2 needs three flow-level consequences of reversibility, stated as fields of `Reversor`; in exchange it needs no independence of ξ and B; it still requires $\eta_0 \notin \operatorname{span} B(x_0)$ (`hindep`, Supplement S7.2). Its proof has two branches where the published sketch describes one.

(xviii) *The application of Birkhoff's graph theorem is external*, for want both of twist-map theory in Mathlib and of a proved correspondence between a three-dimensional invariant torus and an invariant circle of the return map.

S7.11. Verification obligations and detailed evidence

Formal-development inventory. The development has 2 506 lines on Lean 4.33.1 and Mathlib, contains no `sorry`, and an axiom audit (`Lean.collectAxioms`) of all 181 theorem constants in its compiled modules finds only the standard axioms `propext`, `Classical.choice` and `Quot.sound`.

The abstract Lean flow is assumed complete for every real time. The finite-segment argument of §S7.9 justifies the local mathematical application but is not formalised. Equivalence of the executable checker and the Lean predicates is not proved.

Distributed-file axis comparison. At the two signing steps of the reversible interval witness on the distributed file (distributed-file intervals [D11], half-width 10^{-8} m), targeted

propagation of the Newton box to the corresponding phases gives

$$\beta_k \in [-9.630, -1.421] \times 10^{-4}, \qquad \beta_m \in [1.165, 9.467] \times 10^{-4}.$$

The signs therefore change strictly when the reference curve is the validated axis itself.

These accepted signs are finite evidence only: [D11] lacks the exact reversibility required for exclusion by the reversible criterion.

The scope of verification in §1.4 of the main paper is expanded in Tables S10 and S11. Table S12 records the model, reference and criterion for each W7-X result.

The last row tests for shared-formula defects. A second, independently written coil-field evaluator parses the coil file and applies the current convention itself; its B agrees with ours to 4.1×10^{-15} of $|B|$ at five sampled points, including the outward probe. Its derivatives, Richardson-extrapolated central differences of that B, share no differentiation formula, jet algebra or indexing with our analytic jets, and agree to 4.8×10^{-10} relative for ∂B and 1.9×10^{-10} for the derivatives of the transported flow field $V = R\mathbf{B}/B_\varphi$ (calibration [M1]). These are not a precision claim: the difference error (step 10^{-4} m, halved once) is not bounded rigorously, so only the conservative comparison scale of 10^{-8} relative is used. These sampled checks can expose derivative defects but supply no universal detection threshold: cancellation, unexercised parameter directions and shared conventions can hide a defect. Pure current Hessians are now tested separately against differences of the extended right-hand side (§S2.4).

Table S13 reports the recorded width ladder, horizon, sample counts, enclosure growth and costs, including the unsuccessful wider box. In the two D7 runs on the same build, increasing the initial half-width from 3 to 5 nm increases the selected radial endpoint enclosure width from about 10 to 173 μm and the tangent-column width from

Table S10. Validation methods and remaining obligations. The per-result model and initial set are in Table S12.

Evidence	Validation	Remaining obligations
Abstract implications	Lean general-point and reversible proofs and finite conefield algebra	Field/flow/torus structure laws; sound enclosures; conefield geometric input external to Lean
W7-X exclusions: [D2]–[D3], [D5]–[D10]	Audited enclosures; independent finite checks; outward axis rechecks	Model identification external to Lean; exact symmetrised-model reversor or published conefield theorem; no section-homotopic upgrade
Reduced spectral query: [D1]	Fresh tail-scaled enclosures; eight finite leaf checks and exact input coverage	Recorded rounded-input, three-family model; enclosure and tail soundness external to Lean
Repaired current box: [D3]	Complete current Hessians; fresh enclosures and finite replay	Fixed fitted reference; symmetry-preserving five-current subspace; historical witness withdrawn
Distributed interval: [D11]	Finite and axis checks accepted	Exact reversibility not established; no exclusion inferred

Table S11. What each cross-check is independent of, and the defects it can therefore detect. No gate is independent of the criterion's formulation; a defect in the filament model or in the reference-curve convention would pass all five.

gate	independent of	shares with the generator	detects
CAPD audit	implementation, arithmetic stack, step and frame algorithm	the ODE right-hand side we supply it	coding errors in our integrator: stepping, tangent transport, enclosure bookkeeping
GPU/host bit identity	compiler, backend, hardware	algorithm, arithmetic specification, field formula	rounding-mode and backend defects; nothing above the arithmetic
Arb at higher precision	working precision	algorithm, field formula, indexing	conditioning and precision loss; not a shared formula error
Independent checker	implementation, arithmetic (exact rationals), language runtime; shares no code with the generator	the emitted enclosures, which it trusts	errors in the final inequalities and in the witness bookkeeping; it found one
Independent field evaluator	implementation, coil-file parse, current convention, and — through finite differences — the differentiation itself	the filament model and the coil file	a wrong field formula, a transposed index, a mis-signed term, in B and in its derivatives

0.025 to 0.481, at similar sample counts and recorded elapsed times. Adaptive frames and step selection prevent an inclusion-monotone scaling law. The limiting observation is enclosure growth before a strict witness, rather than an increase in sample count alone. No macroscopic reach is inferred.

Table S12. What each record certifies: model, reference curve, criterion and initial set. Records whose sets are contained in a listed one, on the same model and criterion and under the same conditions (small conefield box [D6], narrow symmetrised interval [D9], the smaller boxes of [D7]), are in Table S5.

record	model	reference curve	criterion; initial set	status
Standard configuration, 5/5 *edge chain*				
[D2]	symmetrised 50-coil model	fitted Fourier curve	reversible, reversor exact; point	certified
[D3]	symmetrised 50-coil model	fixed fitted Fourier curve	Remark 3(a), reversible, exact reversor memberwise; five currents $\pm 10^{-8}$	certified after current-Hessian repair
Outward configuration, 5/5 *chain at mid-radius*				
[D5]	distributed 70-coil file	fitted Fourier curve; validated axis ([D4]) by recheck	conefield, Thm. 2 (proved in the main text); point	certified; no reversor hypothesis
[D8]	symmetrised 70-coil model	fitted Fourier curve; its own validated axis by recheck	reversible, reversor exact; point	certified
[D10]	symmetrised 70-coil model	fitted Fourier curve; its own validated axis by recheck	Remark 3(b), reversible, reversor exact; interval $\pm 10^{-8}$ m on $Z = 0$	certified
[D11]	distributed 70-coil file	fitted Fourier curve; validated axis by recheck	reversible finite witness; intervals $\pm 10^{-10}$ and $\pm 10^{-8}$ m on $Z = 0$	finite check accepted; exact reversibility not established; no exclusion inferred
[D7]	distributed 70-coil file	fitted Fourier curve; validated axis by recheck	conefield, Remark 4; box $\pm 5 \times 10^{-9}$ m in R and Z	certified; no reversor hypothesis

Table S13. Archived symmetry-free box reach (small conefield box [D6], conefield box [D7]). w is the common initial half-width; N counts both directions. ΔR and $\Delta\eta_R$ are the largest endpoint enclosure widths at the two selected witnesses, for the transported ∂_R column. Time is the sum of recorded direction elapsed counters, excluding axis rechecks; separate builds and restarts preclude a scaling fit. The failed run lacks complete summaries; its logs place loss near 3.10 transits.

w (nm)	horizon	N	ΔR (μm)	$\Delta\eta_R$	time (h)	cone
1	3.140	108400	3.28	0.0069	8.30	empty
3	3.140	105351	9.99	0.0249	8.08	empty
5	3.140	104829	173.48	0.4808	8.00	empty
10	~ 3.10	115700	—	—	—	no decision

References

[1] J.-L. Figueras, A. Haro and A. Luque, *Rigorous computer-assisted application of KAM theory: a modern approach*, Found. Comput. Math. **17**(5), 1123–1193 (2017), doi:10.1007/s10208-016-9339-3.

[2] D. Martínez-del-Río and R. S. MacKay, *Conefield approach to identifying regions without flux surfaces for magnetic fields*, Plasma Phys. Control. Fusion **67**, 055005 (2025); arXiv:2501.06796, doi:10.1088/1361-6587/adc59b.

[3] N. Kallinikos, R. S. MacKay and D. Martínez-del-Río, *Regions without flux surfaces of given*

class for magnetic fields in toroidal geometry, Plasma Phys. Control. Fusion **65**, 095021 (2023); arXiv:2304.09613, doi:10.1088/1361-6587/acea3f.

[4] R. Davies, C. B. Smiet, C. Batzdorf, J. Geiger, J. Loizu and S. A. Henneberg, *Characterisation of X- and O-points in Wendelstein 7-X with respect to coil currents*, J. Plasma Phys. **92**, E27 (2026); arXiv:2511.15641, doi:10.1017/s0022377826101287.

[5] N. Duignan and J. D. Meiss, *Distinguishing between regular and chaotic orbits of flows by the weighted Birkhoff average*, Physica D **449**, 133749 (2023); arXiv:2212.10783, doi:10.1016/j.physd.2023.133749.

[6] M. Ruth and D. Bindel, *Finding Birkhoff averages via adaptive filtering*, Chaos **34**, 123109 (2024); arXiv:2403.19003, doi:10.1063/5.0215396.

[7] T. Andreeva, J. Kisslinger and H. Wobig, *Characteristics of main configurations of Wendelstein 7-X*, Problems of Atomic Science and Technology **2002**(4), Series: Plasma Physics (7), 45–47 (2002); companion of IPP-Report III/270.

[8] J. Geiger, C. D. Beidler, Y. Feng, H. Maaßberg, N. B. Marushchenko and Y. Turkin, *Physics in the magnetic configuration space of W7-X*, Plasma Phys. Control. Fusion **57**, 014004 (2015), doi:10.1088/0741-3335/57/1/014004.

[9] D. Wilczak and P. Zgliczyński, *C^r-Lohner algorithm*, Schedae Informaticae **20**, 9–46 (2011).

[10] N. McGreivy, *FOCUSADD*, file `focusadd/initFiles/w7x/coils.std_mc` at commit `a05524dd`, https://github.com/nickmcgreivy/FOCUSADD (BSD-3-Clause).

[11] Proxima Fusion GmbH, *Large C++ tests for VMEC++*, file `test_data/coils.w7x` at commit `394a09a6`, https://github.com/proximafusion/vmecpp_large_cpp_tests (MIT).

[12] L. Jaulin, M. Kieffer, O. Didrit and E. Walter, *Applied Interval Analysis*, Springer (2001), §2.4.

[13] M. T. Mizukoshi, W. A. Lodwick, M. Ceberio and V. Kreinovich, *Standard interval computation algorithm is not inclusion-monotonic: examples*, in Uncertainty, Constraints, and Decision Making, Springer (2023), 423–440.

[14] M. R. Herman, *Introduction à l'étude des courbes invariantes par les difféomorphismes de l'anneau*, Astérisque **103–104** (1983), Chap. I.

[15] M. Breden, J. Gonzalez and J. D. Mireles James, *Validated enclosure of renormalization fixed points via Chebyshev series and the DFT*, arXiv:2409.20457 (2024).

[16] J. L. Aurentz and B. Hashemi, *The Laurent–Horner method for validated evaluation of Chebyshev expansions*, arXiv:2409.14952 (2024).

[17] J. N. Mather, *Non-existence of invariant circles*, Ergodic Theory Dynam. Systems **4**(2), 301–309 (1984), §2. (The version of Birkhoff's theorem stated on the cylinder rather than the annulus; Mather notes it is not a consequence of [21] or [20].)

[18] S. S. Abdullaev, T. Van Rompuy, K. H. Finken, M. Lehnen and M. Jakubowski, *Modelling of magnetic field structure of the Ergodic Divertor of Tore Supra and comparison to the Dynamic Ergodic Divertor of TEXTOR*, arXiv:0805.1477 (2008).

[19] F. Johansson, *Arb: efficient arbitrary-precision midpoint-radius interval arithmetic*, IEEE Trans. Computers **66**, 1281–1292 (2017).

[20] G. D. Birkhoff, *Sur quelques courbes fermées remarquables*, Bull. Soc. Math. France **60**, 1–26 (1932), §3.

[21] G. D. Birkhoff, *Surface transformations and their dynamical applications*, Acta Math. **43**(1), 1–119 (1920), §44. (This, not the 1927 *Dynamical systems*, is the reference [2] §5 cites as [12] for "a theorem of Birkhoff (sec. 44 of [12])".)

Independent researcher, Finland
Email address: aapo.lipponen@gmail.com