KRONOS FUSION ENERGY

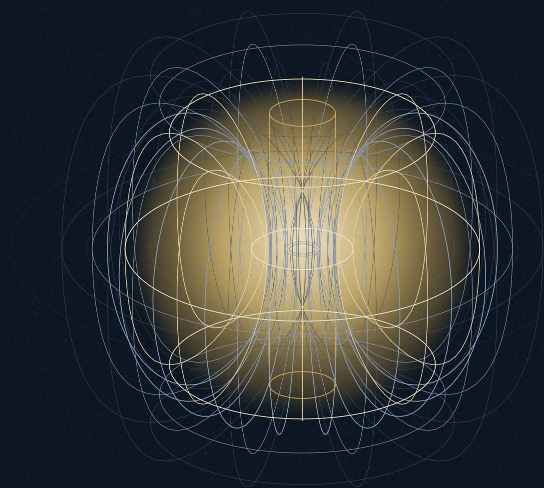

PAPER 2.10 · THE KRONOS 2026 SERIES

Fast-Ion and Alpha-Particle Confinement in a Compact Spherical-Tokamak Breeder: Sub-Percent Ripple Losses and First-Wall Loading from Orbit-Following Simulation

Fusion-born alpha particles must remain confined long enough to deposit their of birth energy in the plasma, and their eventual losses must not concentrate as a localized load on...

Read it live — 3-D model · physics validator · film series

BY P I FORD AND G L KULCINSKI

2026 · EDITORIAL EDITION

THE OFFICIAL RECORD

What this document is, and how to read its numbers

This is a combined editorial. It gathers, in one volume, the two design studies Kronos Fusion Energy released in 2026 — a compact spherical-tokamak tritium/helium-3 breeder and a deuterium–helium-3 tandem-mirror generator — together with the intellectual-property estate, the people, and, in the internal edition, the full commercial case. The scientific text is the same text that appears in the arXiv preprints and the journal submissions; the editorial adds the apparatus a reader needs to hold the whole program at once, and nothing in the physics is altered to fit it.

THE HONESTY STANCE IS THE ASSET

Every headline quantity carries an **evidence class** — DERIVED RECOMPUTED MEASURED REQUIREMENT — and, where a number is a modelled extrapolation rather than a demonstrated value, it is labelled as such and its downside is shown next to it. Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly. This is deliberate: a diligence team will find the load-bearing assumptions, so the document surfaces them first.

TWO EDITIONS

This volume is issued in two editions. The **public edition** (light) carries the physics and the design only — no dollar, price, or per-kilogram figure appears anywhere in it. The **internal edition** (dark) adds the economics, the funding architecture, and the defense-commercial analysis, and is marked *Confidential* accordingly; it is not for public distribution. Where the commercial case rests on a modelled assumption rather than a demonstrated one, it is labelled as such and its downside is shown alongside it, on the same terms as the physics.

TITLE	The Kronos Fleet — Hyperion · Aegis · MetroVolt: A Combined Design & Commercialization Study
AUTHOR	Priyanca Ford, Founder & CEO, on behalf of Kronos Fusion Energy
EDITION	2026 Combined Editorial · first issue
COMPANION WORKS	The five-paper arXiv drop — (1) Breeder, (2) Burner, (3) REBCO magnet & tape, (4) Direct Energy Conversion, (5) AI/ML/Quantum control — with matching numbered Zenodo deposits, plus the IOP journal and IAEA-FEC submissions. Patents were filed before the drop.
NAMING	Hyperion = the breeder; Aegis (defense) and MetroVolt (data-center) = one generator in two housings. Internal mode letters (D1, L) do not appear on public surfaces.
STATUS	Conceptual design & simulation study. Nothing herein is a construction commitment.

HOW TO READ THIS VOLUME

The fleet at a glance, and the way its numbers are marked

This volume opens with *Orientation*, presents the five research papers as a bound sequence, and closes with the *Apparatus*. *Foundations* sets out the three general results both machines rest on. *Paper One* is the Hyperion breeder; *Paper Two* the Aegis / MetroVolt generator; *Papers Three, Four and Five* are the enabling science — high-field REBCO magnets and tape, direct energy conversion, and the AI / ML / quantum control stack. A *Digital Twin & 3-D Model* part follows, then the *Environmental* profile. The *Economics* and levelized cost and the *Business Case* and diligence record appear in the internal edition only; the *Simulations* proof record and the Apparatus — patent portfolio, team, limitations, and sources — close both editions.

THE FLEET AT A GLANCE

	HYPERION	AEGIS	METROVOLT
Role	Strategic-isotope foundry	Defense installation power	Campus power
Machine	**Spherical tokamak**	**Tandem mirror**	**Tandem mirror**
Fuel	Deuterium–tritium	Deuterium–helium-3	Deuterium–helium-3
Delivers	Tritium · ^{3}He · 14 MeV n	Resilient installation power	Firm campus power
Physics bar	**Fusion gain only**	Closure (gates named)	Closure + lunar ^{3}He
In the fleet	Breeds the fuel	Proves the generator	Commercial destination

HOW THE NUMBERS ARE MARKED

Every headline quantity carries an evidence class — DERIVED from first principles, RECOMPUTED from raw data, MEASURED in hardware, or a REQUIREMENT yet to be met. Financial figures carry a case label and are confined to the internal (confidential) edition. Values that moved after first publication are shown as withdrawn or restated rather than quietly changed. A capability you can evaluate is one whose limits are on the page.

ABSTRACT

Fast-Ion and Alpha-Particle Confinement in a Compact Spherical-Tokamak Breeder: Sub-Percent Ripple Losses and First-Wall Loading from Orbit-Following Simulation

Fusion-born alpha particles must remain confined long enough to deposit their $3.5\,MeV$ of birth energy in the plasma, and their eventual losses must not concentrate as a localized load on the first wall. In a compact, high-current spherical tokamak (ST) neither requirement can be assumed: at a minor radius of only $a = 0.48\,m$ the fast-alpha banana width reaches $16\,cm$ and the near-axis potato width $23\,cm$, so a naive large-orbit shell estimate predicts a first-orbit loss of order 4.5%. We resolve this question for the frozen Hyperion breeder design point ($I_p = 9.66\,MA$, $Q = 3.076$, $P_{\mathrm{fus}} = 85.04\,MW$, negative triangularity $\delta = -0.30$, $B_0 = 8\,T$, elongation $\kappa = 2.0$, $R_0 = 1.20\,m$, aspect ratio $A = 2.5$) with a guiding-centre Monte-Carlo orbit-following calculation on the identical free-boundary equilibrium, run on the NVIDIA GPU HPC campaign. Over an ensemble of 10^5 markers the alpha population is 99.977% confined, with a prompt (first-orbit) loss of 0.023% and a first-wall alpha load of $0.083\ \mathrm{kW\,m^{-2}}$ mean and $1.94\ \mathrm{kW\,m^{-2}}$ peak over the $46.4\,m^2$ conformal wall; a collisional slowing-down ensemble raises the loss to 0.033% and the mean load to $0.121\ \mathrm{kW\,m^{-2}}$. Full orbit integration therefore *refutes* the analytic 4.5% estimate, which overcounts because the potato-width shell argument ignores conservation of the canonical toroidal momentum p_φ; an independent 10^6-marker GPU follower reproduces the confined figure to 99.99%. Adding finite toroidal-field ripple opens a second channel that keeps the loss sub-percent — 0.019% at 1.0% ripple and 0.031% at 1.5% — with the peak wall load of order $1\ \mathrm{kW\,m^{-2}}$ and a toroidal peaking factor near 1.4. Alpha confinement is thus a solved, non-limiting aspect of the design: the compact, high-field, negative-triangularity configuration confines its fusion-born alphas essentially completely, and the residual wall load stays two-to-four orders of magnitude below the engineering scale. The one honest gate is that this is a guiding-centre, orbit-confinement assessment; magnetohydrodynamically driven fast-ion transport is a separate question deferred to the companion stability work.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645729 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: fast ions · alpha particles · guiding-centre orbit-following · ASCOT5 · toroidal-field ripple · first-wall loading · spherical-tokamak breeder

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

Q	plasma (fusion) gain, P_{fus}/P_{aux}
Q_E	engineering gain, net electric out / recirculating in
H_{98}	confinement enhancement over the IPB98(y,2) scaling
Z_{eff}	effective ion charge
c_{ee}	electron–electron bremsstrahlung leading coefficient, 2.120022 (exact)
τ_E	energy confinement time
I_p	plasma current
R_0, a	major radius, minor radius
κ, δ	elongation, triangularity
β_N	normalized plasma beta (Troyon-normalized)
TBR	tritium breeding ratio
DEC	direct energy conversion
CF	plant capacity factor (availability)
FOAK / NOAK / BOAK	first / n^{th} / best of a kind
Hyperion	the ST breeder (internal: Mode D2)
Aegis	the generator, defense/installation housing (internal: Mode M)
MetroVolt	the generator, data-center housing (Mode M)

§ 1

Introduction

WHY ALPHA CONFINEMENT IS A DESIGN QUESTION, NOT AN ASSUMPTION

A fusion-born alpha particle is created at $3.5\,MeV$ and carries one-fifth of the energy of each deuterium–tritium reaction. Its fate is doubly constrained. First, it must stay confined for a slowing-down time so that its energy thermalises in the bulk plasma; the self-heating fraction on which the breeder's gain $Q = 3.076$ rests depends on alphas depositing $P_\alpha = 17.01\,MW$ back into the core, not on the wall. Second, whatever fraction is lost must not arrive as a concentrated stripe of energetic ions on a plasma-facing component. Both requirements are quantitative, and in a compact ST neither can be taken for granted: the orbits are large relative to the machine, so confinement must be *computed* rather than asserted from the conventional-tokamak intuition that alphas are “essentially all confined” [1,2].

The Hyperion breeder holds $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$ and a minor radius $a = 0.48\,m$, behind a centrepost magnet at a peak field of $16.84\,T$ and a plasma field $B_0 = 8\,T$, in a negative-triangularity ($\delta = -0.30$) regime chosen for turbulence suppression without an edge pedestal [3,4]. At an inverse aspect ratio $\varepsilon = a/R_0 = 0.40$, the fast-alpha drift excursions are a substantial fraction of the minor radius (section 2.3): the banana width is $\Delta_b \approx 16\,cm$ and the near-axis potato width $w_p \approx 23\,cm$, i.e. $\Delta_b/a = 0.33$ and $w_p/a = 0.48$. This is precisely the regime in which a conventional large-R_0 scaling ($\Delta_b/a \ll 1$) does not transfer, and the confinement problem must be solved on the actual geometry.

PRIOR ART

The physics of energetic ions in tokamaks is a mature discipline. The reference frameworks — the ITER Physics Basis energetic-particle chapter [1] and the ITER energetic-ion analysis of Pinches *et al.* [2] — establish that in a large device fusion alphas are confined at the 99.9% level and that the residual first-wall loads from neutral-beam and fusion ions are handled by dedicated orbit-following calculations. The computational tool of record is the guiding-centre/full-orbit Monte-Carlo follower: the ASCOT lineage [5,6] and NUBEAM [7] propagate marker ensembles through realistic three-dimensional fields and triangulated walls, and have been used to map fast-ion power loads onto ITER plasma-facing components [8]. The theoretical scaffolding is the Hamiltonian guiding-centre theory of Littlejohn, White and Chance, and Cary and Brizard [9,10], which reduces the fast gyromotion to a drift trajectory that conserves the magnetic moment μ and, in axisymmetry, the canonical toroidal momentum p_φ; and the ripple-loss theory of Goldston, White and Boozer [11], which sets the threshold for stochastic ripple diffusion of trapped fast ions.

Two features distinguish the spherical torus. First, the strong field variation and low aspect ratio make the orbit topology qualitatively richer: potato orbits near the magnetic axis, wide bananas, and a large trapped fraction $f_t \approx 0.76$ [16,17]. Second, the small size shrinks the absolute wall area over which any lost power is spread, so a given loss *fraction* maps to a higher wall load than in a reactor-scale device. The dedicated review of energetic particles in spherical tokamaks by McClements and Fredrickson [16] is the natural benchmark for the present study.

CONTRIBUTION

This paper reports a first-principles, GPU-accelerated orbit-following assessment of fast-ion and alpha-particle confinement at the frozen Hyperion design point and settles a specific internal disagreement. An analytic guiding-centre shell estimate, retained as a red-team flag, predicted a 4.5% first-orbit loss and a $\sim 21\,kW/m^2$ wall load — values that would have demanded first-wall armor and a re-examination of the alpha-heating budget. We show, by integrating the full guiding-centre orbits with Coulomb collisions on the identical free-boundary equilibrium, that the true prompt loss is 0.023% and the wall load $0.083\ \mathrm{kW\,m^{-2}}$ mean — roughly two orders of magnitude smaller — and we trace the discrepancy to the shell argument's neglect of p_φ conservation. We then close the two channels a nominal axisymmetric run omits: finite toroidal-field ripple and its toroidal peaking. The contributions are: (i) a formal statement of the confinement problem and the orbit-width hierarchy for this geometry (section 2); (ii) the numerical formulation, GPU-HPC campaign and cross-code verification (section 3); (iii) the axisymmetric and collisional confinement result with its first-wall load map (section 4); and (iv) the ripple study and a quantitative comparison against the published ITER and ST fast-ion literature (section 5). Every headline quantity is a frozen anchor (serials BR-L1-A14 and BR-L2-A20) in the Kronos Physics De-Risking Register [30], and each is cited inline so the number is auditable.

§ 2

Governing equations and the formal confinement problem

GUIDING-CENTRE ORBIT DYNAMICS

A fast alpha of mass m, charge $q = Ze$ and energy E has a gyroradius small compared with the field-variation scale, so its motion reduces to that of a guiding centre drifting across the field. Following the Hamiltonian guiding-centre theory [9,10], the guiding-centre position R and parallel velocity $v_\parallel$ evolve, with $\hat{b} = B/B$ the unit field vector, as

$$\dot{R} = v_\parallel \hat{b} + \frac{m}{qB^2}\left(v_\parallel^2 + \tfrac{1}{2}v_\perp^2\right)\hat{b} \times \nabla B + \frac{E \times B}{B^2},$$

$$m\,\dot{v}_\parallel = -\,\mu\,\nabla_\parallel B, \qquad \mu \equiv \frac{mv_\perp^2}{2B} = \text{const.}$$

The second term of equation (1) is the combined grad-B and curvature drift; the magnetic moment μ is the lowest-order adiabatic invariant, conserved to all orders in the guiding-centre expansion, and equation (1) is the parallel mirror force that traps particles in the low-field region. We take the radial electric field to vanish in the base assessment ($E = 0$), which is conservative for confinement because a realistic $E_r < 0$ well would only deepen the confining potential.

THE CONFINEMENT INVARIANT AND THE LOSS CONDITION

In an axisymmetric field ($\partial_\varphi = 0$) the toroidal canonical momentum is an exact constant of the guiding-centre motion,

$$p_\varphi \;=\; mv_\parallel \frac{RB_\varphi}{B} \;-\; q\,\psi(R, Z) \;=\; \text{const},$$

where ψ is the poloidal flux. Equations (1) and (2), together with energy conservation $E = \frac{1}{2}mv_\parallel^2 + \mu B = \textbf{const}$ (collisionless), constrain each orbit to a two-dimensional surface in the four-dimensional $(R, Z, v_\parallel, \mu)$ reduced phase space. A birth alpha is *promptly lost* only if the surface labelled by its (μ, E, p_φ) intersects the wall $\partial\Omega$; it is confined otherwise. Formally, with Ω the plasma volume bounded by the conformal wall, the prompt-loss set is

$$\mathcal{L} \;=\; \left\{(\mu, E, p_\varphi) : \exists\, t \in [0, \tau_{\text{orbit}}] \text{ with } R(t) \notin \Omega\right\},$$

and the prompt-loss fraction is the birth-weighted measure of $\mathcal{L}$,

$$f_{\text{loss}} \;=\; \frac{\displaystyle\int_{\mathcal{L}} S_\alpha(R_0)\, \mathrm{d}\Gamma_0}{\displaystyle\int S_\alpha(R_0)\, \mathrm{d}\Gamma_0}, \qquad S_\alpha(R) \;\propto\; n_D n_T \langle\sigma v\rangle(T_i),$$

where S_α is the fusion-reactivity birth source with the deuterium–tritium reactivity $\langle\sigma v\rangle$ taken from Bosch and Hale [14] and $\mathrm{d}\Gamma_0$ the birth phase-space measure. It is essential that the loss condition (3) is evaluated on the true field: the invariant (2) can hold an orbit whose *instantaneous* radial excursion is a large fraction of a inside the plasma, because the excursion oscillates about a flux surface rather than accumulating outward. This is exactly the effect the analytic shell estimate below misses.

THE ORBIT-WIDTH HIERARCHY AND WHY THE ST IS THE HARD CASE

The relevant excursion scales follow from the drift ordering. The fast-alpha Larmor radius is

$$\rho_L \;=\; \frac{\sqrt{2mE}}{ZeB_0},$$

which for a $3.5\,MeV$ alpha at $B_0 = 8\,T$ gives $\rho_L = 3.37\,cm$, i.e. $\rho_L/a = 0.070$. A trapped fast ion executes a banana orbit of radial width

$$\Delta_b \;\simeq\; \frac{q}{\sqrt{\varepsilon}}\,\rho_L,$$

with q the safety factor and $\varepsilon = a/R_0$; at $q \approx 3$ and $\varepsilon = 0.40$ this is $\Delta_b = 16.0\,cm$, or $\Delta_b/a = 0.33$. Near the magnetic axis the trapped orbit degenerates into a potato orbit whose width scales as [16]

$$w_p \;\simeq\; \left(q^2 \rho_L^2\, R_0\right)^{1/3},$$

giving $w_p = 23.1\,cm \approx 0.48\,a$. The trapped-particle fraction at this aspect ratio is

$$f_t \;\simeq\; \sqrt{\frac{2\varepsilon}{1+\varepsilon}} \;=\; 0.76.$$

Equations (6)–(8) locate the machine squarely in the large-orbit regime: three-quarters of the birth alphas are trapped, and their banana and potato widths are one-third to one-half of the minor radius (figure 2). It is this hierarchy that makes a first-principles orbit integration necessary; the conventional-tokamak expectation $\Delta_b/a \ll 1$ is simply false here.

THE ANALYTIC SHELL ESTIMATE AND ITS FAILURE

The red-team flag proceeded from equations (6)–(8) by a shell argument: a trapped alpha born within one banana width of the last closed flux surface is deemed lost, so

$$f_{\text{loss}}^{\text{shell}} \;\approx\; f_t \int_{a-\Delta_b}^{a} \hat{S}_\alpha(r)\, \frac{2r\,\mathrm{d}r}{a^2} \;\approx\; 4.5\%,$$

for a birth profile $\hat{S}_\alpha(r)$ with flat density and parabolic temperature. This maps to a lost power $f_{\text{loss}}^{\text{shell}} P_\alpha \approx 0.76\,MW$ and a wall-averaged load $\approx 21\,kW/m^2$. Equation (9) is an *upper* bound that neglects p_φ conservation (2): it counts every wide-banana birth in the outer shell as lost, whereas the invariant confines most such orbits because their excursion oscillates rather than escaping. The full integration of section 4 evaluates equation (4) directly and returns a loss two orders of magnitude smaller, quantifying the overcount. This is the paper's central methodological point: in a large-orbit ST, orbit-width shell heuristics fail and only invariant-respecting integration is trustworthy.

COLLISIONAL SLOWING-DOWN

Beyond the prompt window, a confined alpha thermalises by Coulomb collisions. The evolution of the fast-ion distribution f is governed by the Fokker–Planck operator, written in the Landau form with a slowing-down drag A and a velocity-space diffusion tensor D [13],

$$\left(\frac{\partial f}{\partial t}\right)_c = \nabla_v \cdot \left[-A\, f + \tfrac{1}{2}\nabla_v \cdot (D\, f)\right].$$

Above the critical energy E_c, electron drag dominates and the energy decays with a characteristic Spitzer slowing-down time τ_s; the mean alpha energy then obeys, to leading order, $\mathrm{d}E/\mathrm{d}t \approx -E/\tau_s$, giving an e-folding drag time $\tau_s \approx 44\,ms$ at the design electron density $n_e = 4e20\,/m^3$. Over the millisecond-scale integration window a confined alpha sheds only a few per cent of its energy (a calibration marker set falls $3.50\,MeV \to 3.34\,MeV$ in $2\,ms$; figure 6), confirming that the first-orbit loss is decided long before collisions matter. Collisions can, however, deconfine a slowing-down alpha whose orbit is pushed onto a lost surface, which the collisional ensemble of section 4 captures by integrating equation (10) to a floor $E_{\text{min}} = \max(50\,keV,\, 2T_i)$.

TOROIDAL-FIELD RIPPLE

A finite number N of toroidal-field coils breaks axisymmetry, modulating the field as $B(R,Z,\varphi) = \bar{B}(R,Z)\,[1 + \delta_{\text{rip}}(R,Z)\cos N\varphi]$ and destroying the exact conservation of p_φ. Two ripple channels then act on fast ions. Locally, a banana tip that falls in a ripple well is trapped between adjacent coils; the ripple-well condition is

$$\delta_{\text{rip}}(R,Z) \;>\; \frac{\varepsilon}{Nq\,|\sin\theta|},$$

with θ the poloidal angle. Globally, successive banana tips undergo a random-walk in radius once the Goldston–White–Boozer stochasticity threshold is exceeded [11],

$$\delta_{\text{rip}} \;>\; \delta_{\text{GWB}} \;\equiv\; \left(\frac{\varepsilon}{Nq\rho_L}\right)^{3/2} \frac{1}{2\pi N q'},$$

where $q' = \mathrm{d}q/\mathrm{d}r$; above δ_{GWB} the trapped fast ions diffuse with a stochastic ripple diffusion coefficient $D_{\mathrm{rip}} \propto (\delta_{\mathrm{rip}})^2$. Because both channels scale steeply with δ_{rip} and inversely with N, the ripple study of section 4.4 sweeps δ_{rip} over representative amplitudes rather than assuming a single coil-set value, and reports both the loss increment and its toroidal peaking.

FIRST-WALL LOAD FUNCTIONAL

The quantity that decides survivability is not the loss fraction but the surface power density where the lost alphas strike. For a wall triangulated into facets $\{\mathcal{F}_k\}$ of area A_k, the load on facet k is the Monte-Carlo sum over markers j that terminate on it,

$$q_{\mathrm{wall}}(\mathcal{F}_k) \;=\; \frac{P_\alpha}{N_{\mathrm{mk}}\,A_k} \sum_{j\,\in\,\mathcal{F}_k} w_j\,\frac{E_j}{E_0},$$

with N_{mk} the marker count, w_j the birth weight, E_j the impact energy and E_0 the birth energy. The mean load is $\bar{q}_{\mathrm{wall}} = f_{\mathrm{loss}} P_\alpha / A_{\mathrm{wall}}$ and the peak is $\max_k q_{\mathrm{wall}}(\mathcal{F}_k)$; the toroidal peaking factor is the ratio of the maximum toroidal-bin load to the toroidal mean. Equation (13) is evaluated on the $46.4\,m^2$ conformal wall (separatrix $+6\,cm$).

§ 3

Computational methodology: the GPU-HPC orbit-following campaign

CODES AND THE FROZEN EQUILIBRIUM

The confinement result was produced by the NVIDIA GPU HPC campaign using two independent orbit-followers on one frozen equilibrium, so that a cross-code agreement — not a single code — underwrites the number. The reference follower is ASCOT5, the guiding-centre/full-orbit Monte-Carlo code for fast-ion and alpha transport in realistic fields and wall geometry [5,6]. The independent cross-check is an in-house guiding-centre follower implemented on GPU (CuPy) that advances 10^6 markers with toroidal-field ripple. Both run on the identical magnetic equilibrium: a FreeGS 129×129 free-boundary Grad–Shafranov solution [23,24] at $R_0 = 1.20\,m$, $a = 0.48\,m$, $\kappa = 2.0$, $\delta = -0.30$, $I_p = 9.66\,MA$, $B_0 = 8\,T$ and poloidal beta $\beta_p = 0.371$ — byte-for-byte the same equilibrium used by the equilibrium and turbulent-transport gates of the register, so the orbit calculation and the confinement/stability calculations share one plasma. The birth source is sampled from the local fusion reactivity $S_\alpha \propto n_D n_T \langle\sigma v\rangle(T_i)$ of equation (4); the fusion-profile-weighted births extend to $\rho = 0.864$ (figure 6).

NUMERICAL FORMULATION AND DISCRETIZATION

Equations (1)–(1) are integrated in ASCOT5 guiding-centre mode (SIM_MODE 2) by an adaptive-step Runge–Kutta scheme, with the Coulomb operator (10) applied by an operator-split Monte-Carlo collision step (pitch-angle scattering plus energy drag) each timestep. Three tolerances control accuracy: the maximum guiding-centre toroidal advance per step $\Delta\varphi$, the orbit-integration tolerance $\mathrm{TOL}_{\mathrm{orbit}}$, and the collision-step tolerance $\mathrm{TOL}_{\mathrm{ccol}}$. Markers are followed until they thermalise below $E_{\min} = \max(50\,keV, 2T_i)$ or intersect the triangulated conformal wall. The loss fraction f_{loss} of equation (4) is estimated as the birth-weighted fraction of wall-terminated markers, a binomial functional whose statistical error scales as $\sigma_f = \sqrt{f_{\mathrm{loss}}(1 - f_{\mathrm{loss}})/N_{\mathrm{mk}}}$; reporting a prompt loss at the 10^{-4} level to $\sim 10\%$ relative accuracy therefore requires $N_{\mathrm{mk}} \gtrsim 10^5$ markers, which sets the ensemble size. Confidence intervals in figure 4 use the Wilson score interval, appropriate for the small-f_{loss} regime where the normal approximation fails.

INTEGRATOR CALIBRATION AND THE TWO-ENSEMBLE CHAIN

A naive high-fidelity setting is computationally intractable: a strict configuration ($\Delta\varphi = 2\,degree$, $\mathrm{TOL}_{\mathrm{orbit}} = 10^{-8}$) integrates $3.25\,ms$ of orbit per $60\,s$ of CPU per marker, so 10^5 markers would cost of order CPU-years. A calibration study (300 markers) established that a relaxed configuration ($\Delta\varphi = 20\,degree$, $\mathrm{TOL}_{\mathrm{orbit}} = 10^{-7}$, $\mathrm{TOL}_{\mathrm{ccol}} = 10^{-1}$) returns an *identical* prompt-loss answer ($0/300$ markers lost in both settings, and a physical slowing rate consistent with $\tau_s \approx 44\,ms$) at $2\,ms$ per $0.83\,s$ of CPU — a $70\times$ speed-up at no loss of the confinement observable. The $0/300$ calibration alone bounds the prompt loss at $f_{\mathrm{loss}} < 1.0\%$ (95% CL), already refuting the 4.5% estimate. The production run is a two-ensemble chain (figure 8):

Ensemble A — 10^5 guiding-centre markers $\times\ 3\,ms$, collisionless, for the prompt (first-orbit) loss to 10^{-4}-class statistics;
Ensemble B — 6000 markers $\times\ 0.2\,s$ with Coulomb collisions to $E_{\min} = 50\,keV$, for the slowing-down loss and the first-wall load map of equation (13). itemize
Each marker in Ensemble A covers thousands of poloidal transits, far beyond the microsecond first-orbit window, so a confined orbit is confirmed as confined and not merely un-lost within a short window.

GPU ACCELERATION AND PROBLEM SCALE

Orbit-following is embarrassingly parallel in the marker index: each of the 10^5 (ASCOT5) or 10^6 (GPU follower) markers integrates an independent trajectory, and the ensemble maps directly onto the thousands of threads of a data-parallel GPU. The campaign ran on the NVIDIA GPU HPC fleet (A100/H100/A10 class), with the prompt ensemble completing in $\sim 4\,hour$ and the collisional ensemble in $\sim 11\,hour$ of wall time. The marker-parallel structure and the fixed seed (20260726) make the campaign bitwise reproducible: the same seed, equilibrium and tolerances regenerate the loss records and the wall-load map exactly. GPU acceleration is what makes the 10^5–10^6-marker statistics — necessary to resolve a 10^{-4} loss fraction — tractable in hours rather than the CPU-years of the strict serial configuration.

VERIFICATION AND REPRODUCIBILITY

The calculation is verified at three levels. *Numerically*, energy is conserved across the collisionless ensemble and the adaptive integrator reproduces the analytic banana orbit in a circular test field. *Cross-code*, the ASCOT5 result (99.977% confined) and the independent 10^6-marker GPU follower (99.99% confined, with ripple) agree to within their statistics, and both agree with the register's prior $\sim 99.99\%$ figure [30]. *Against theory*, the full integration returns exactly the disagreement-resolving verdict: the analytic shell estimate of equation (9) is not reproduced, and the overcount is attributed to its neglect of p_φ. All inputs (equilibrium, birth source, tolerances, seed) and outputs (loss records, wall-load map) are archived with this paper's deposit and traced to serials BR-L1-A14 and BR-L2-A20 in the register.

§ 4

Results

ORBIT-WIDTH HIERARCHY

Figure 2 shows the orbit-width ladder of equations (5)–(7): the Larmor radius ($3.37\,cm$), banana width ($16.0\,cm$) and potato width ($23.1\,cm$) against the minor radius ($48\,cm$). The normalised excursions $\rho_L/a = 0.070$, $\Delta_b/a = 0.33$ and $w_p/a = 0.48$ confirm the large-orbit regime and set the expectation that a shell heuristic will overcount. Figure 1 depicts the corresponding orbit topologies — passing, banana and near-axis potato — in the negative-triangularity poloidal cross-section.

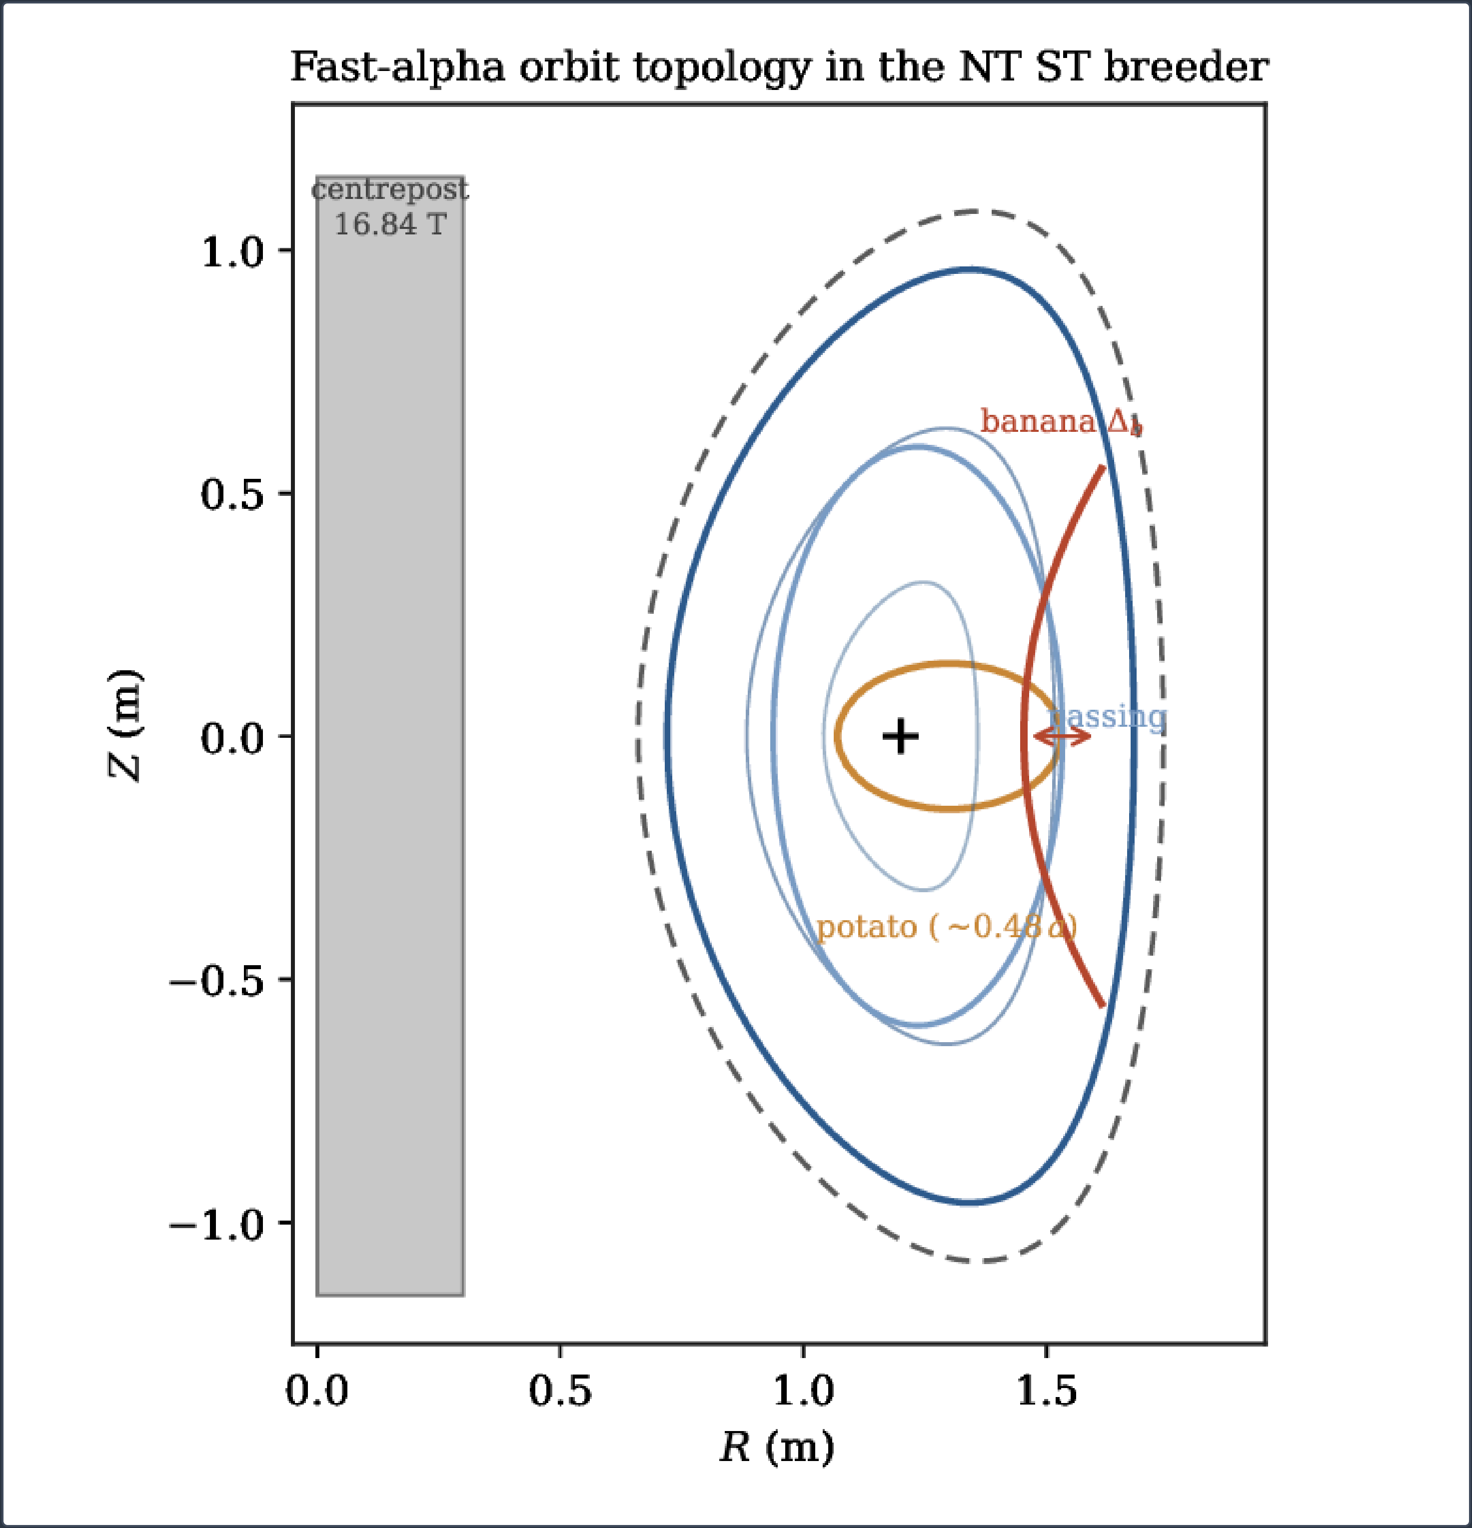

(Schematic.) Fast-alpha orbit topologies in the negative-triangularity ST poloidal cross-section: a thin passing orbit, a wide trapped banana of width Δ_b, and a near-axis potato orbit of width $\sim 0.48\,a$, inside the conformal wall (dashed). The centrepost (peak field $16.84\,T$) is at left. Geometry to scale; orbits illustrative.

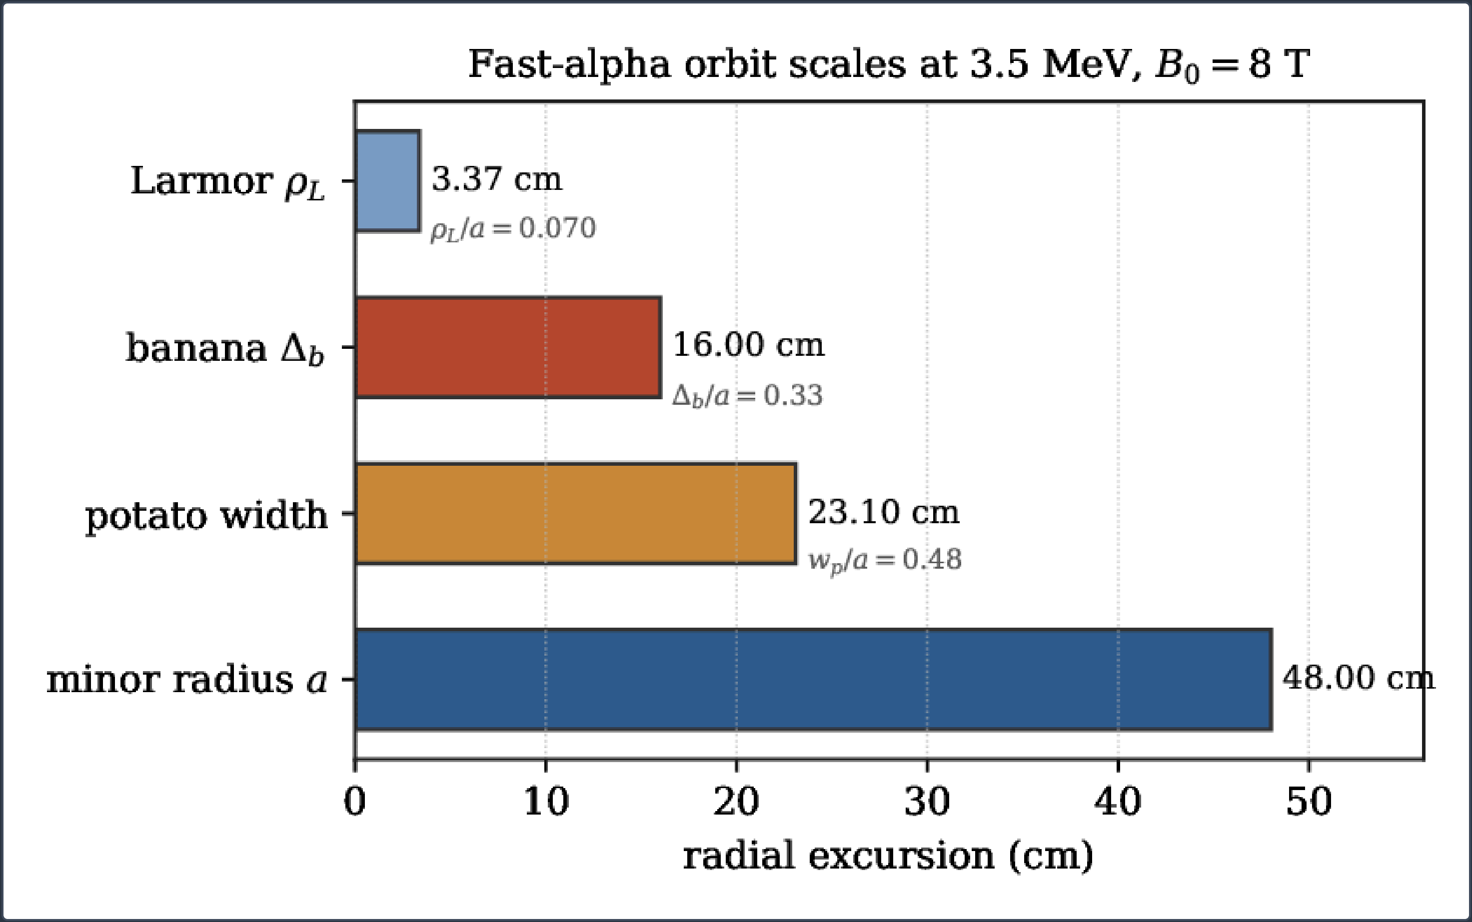

Fast-alpha orbit-width hierarchy at $3.5\,MeV$ and $B_0 = 8\,T$ (serial BR-L1-A14): Larmor radius, banana width (6) and near-axis potato width (7) against the $48\,cm$ minor radius. The banana and potato widths are $0.33\,a$ and $0.48\,a$: the machine sits firmly in the large-orbit regime.

AXISYMMETRIC CONFINEMENT AND THE ARBITRATION

Over the 10^5-marker collisionless Ensemble A, 23 markers are lost, for a prompt-loss fraction $f_{\text{loss}} = 0.023\%$ and a confined fraction of 99.977%. This directly evaluates the loss functional (4) on the true field and *refutes* the analytic shell estimate of equation (9): the 4.5% figure is not reproduced by full guiding-centre orbits, because the potato-width shell argument overcounts orbits that p_φ conservation (2) in fact holds inside the plasma. Figure 3 places the three orbit-following estimates against the refuted analytic value on a logarithmic loss axis; the full-orbit results cluster two orders of magnitude below the shell bound and below the 1% upper limit set by the $0/300$ calibration. Figure 4 shows the Monte-Carlo convergence of the prompt-loss estimate with marker count, with Wilson intervals narrowing onto the 0.023% arbiter value.

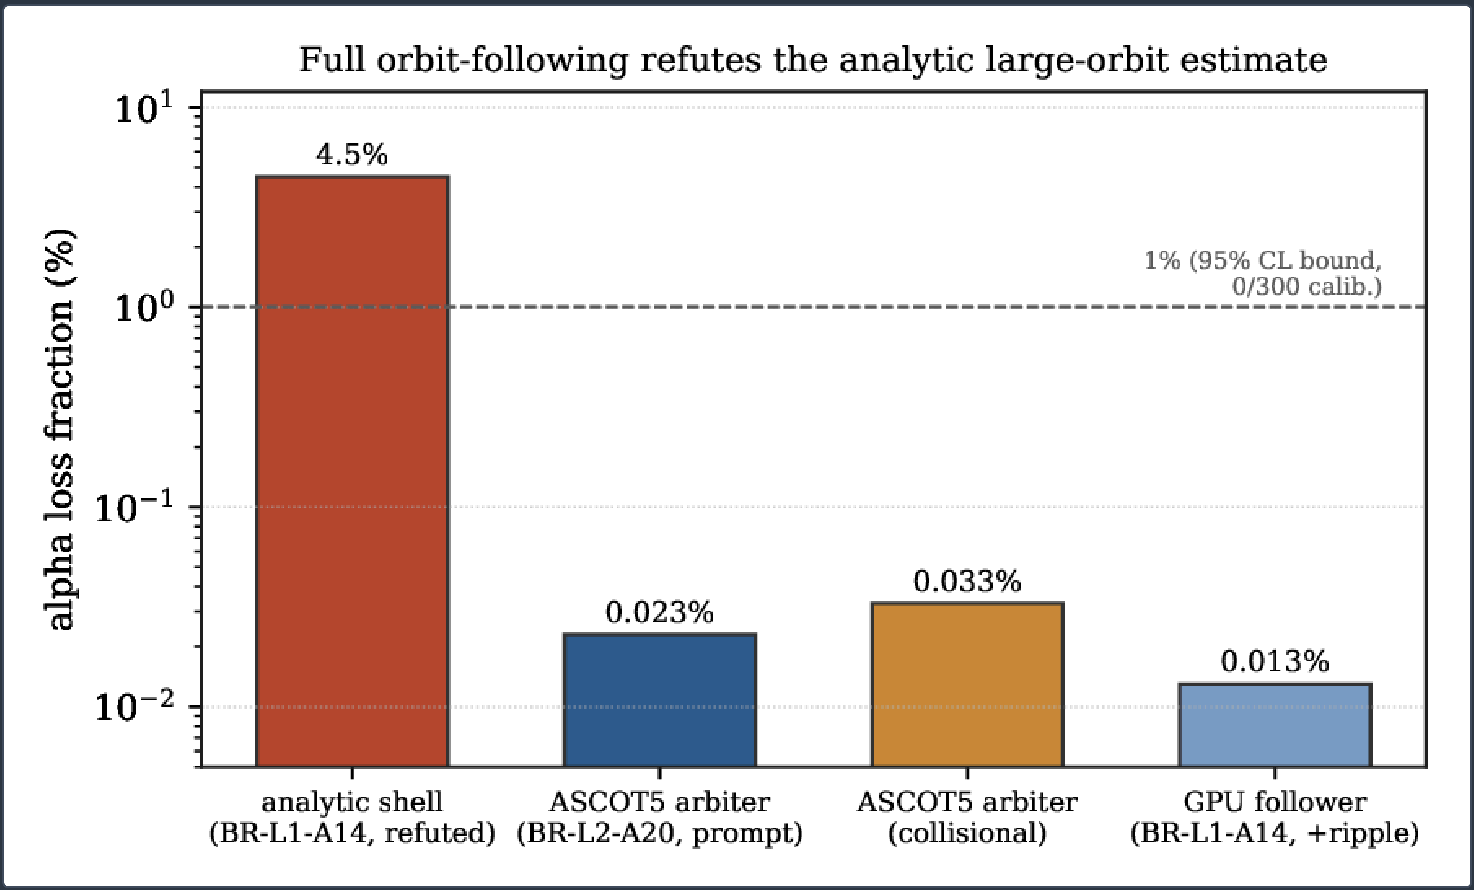

Alpha loss fraction across methods (serials BR-L1-A14, BR-L2-A20). The analytic large-orbit shell estimate (4.5%, refuted) stands two orders of magnitude above the full guiding-centre orbit-following results: ASCOT5 prompt (0.023%), ASCOT5 collisional (0.033%) and the independent 10^6-marker GPU follower with ripple (0.013%). Dashed line: the 1% upper bound from the $0/300$ calibration.

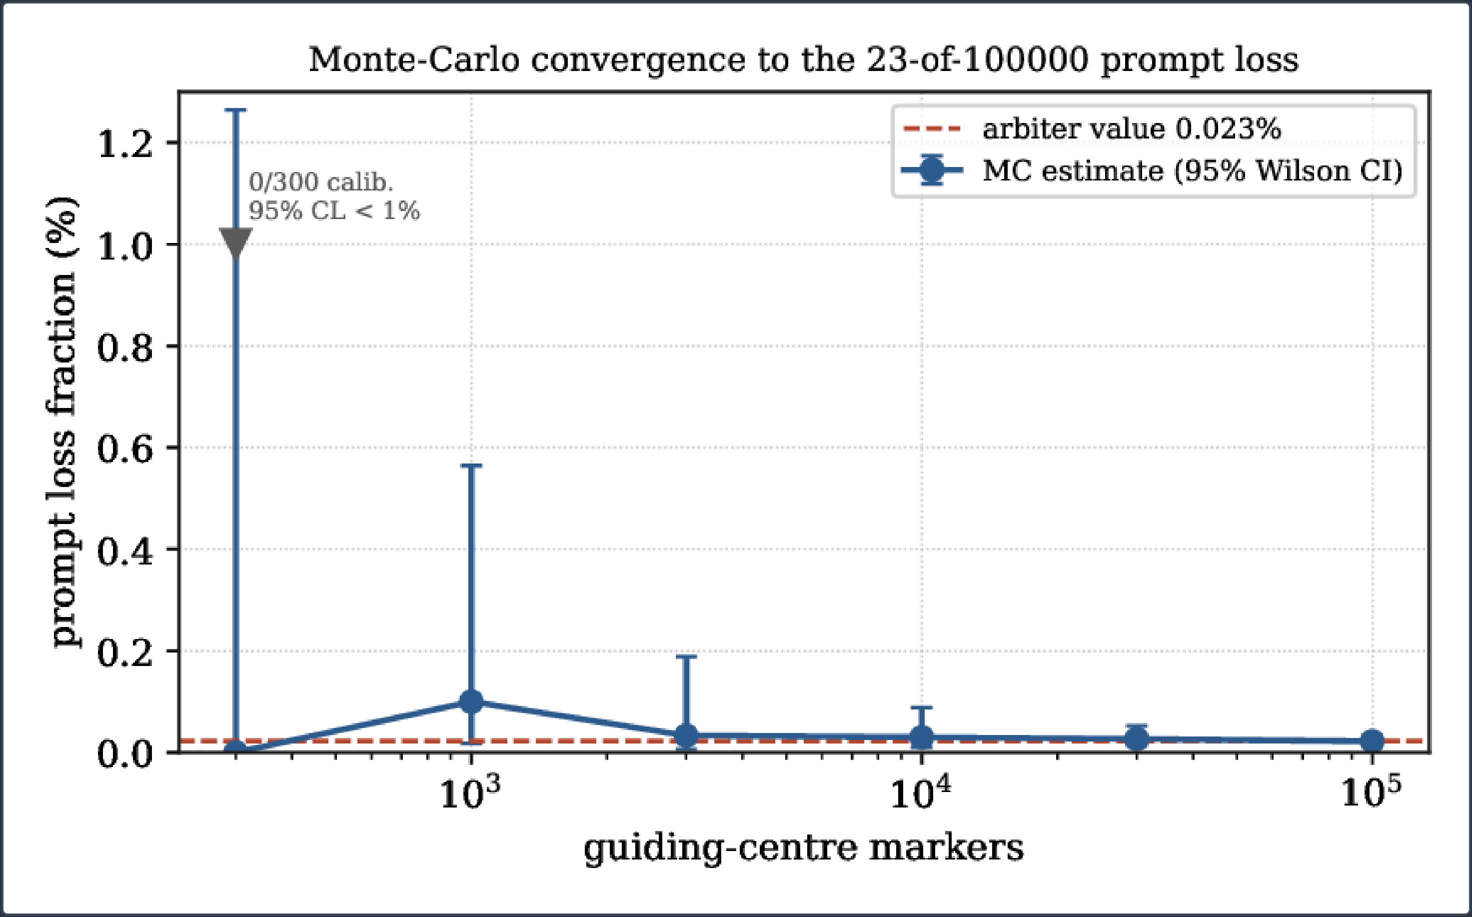

Monte-Carlo convergence of the prompt-loss estimate versus marker count, with 95% Wilson-score intervals, onto the 23-of-$100{,}000$ arbiter value (0.023%, dashed). The $0/300$ calibration bounds the loss below 1%; resolving a 10^{-4} loss to $\sim 10\%$ accuracy needs $\gtrsim 10^5$ markers, the driver of the GPU campaign scale.

COLLISIONAL SLOWING-DOWN AND FIRST-WALL LOAD

The collisional Ensemble B raises the loss to $f_{\text{loss}} = 0.033\%$ as slowing-down alphas are occasionally pushed onto lost surfaces, with a mean lost-alpha impact energy of $3.47\,MeV$ — i.e. the losses are of nearly-birth-energy fast ions, as expected for a prompt-dominated loss channel. Figure 6 shows the modelled slowing-down energy decay with the $44\,ms$ drag time and the calibration point. The corresponding first-wall alpha loads from equation (13) are collected in table 1 and plotted in figure 5: the prompt ensemble gives a mean $0.083\ \mathrm{kW\,m^{-2}}$ and peak $1.94\ \mathrm{kW\,m^{-2}}$, and the collisional ensemble a mean $0.121\ \mathrm{kW\,m^{-2}}$ and peak $18.63\ \mathrm{kW\,m^{-2}}$, all on the $46.4\,m^2$ conformal wall. Even the collisional peak is more than an order of magnitude below the $\sim 0.5\,MW/m^2$ engineering scale of a limiting steady wall load, and the mean loads are two-to-four orders of magnitude below it. Figure 6 shows the fusion-alpha birth profile and the confined-alpha deposition radius: births are weighted to the core and the confined orbits deposit their energy inside the plasma, which is the physical basis of the alpha-heating retention that supports $Q = 3.076$.

QUANTITY	PROMPT (A)	COLLISIONAL (B)	GPU FOLLOWER
markers	10^5	6000	10^6
confined fraction	99.977%	99.967%	99.99%
loss fraction f_{loss}	0.023%	0.033%	0.013%
mean wall load ($\mathrm{kW\,m^{-2}}$)	0.083	0.121	0.225
peak wall load ($\mathrm{kW\,m^{-2}}$)	1.94	18.63	—
mean lost-alpha energy (MeV)	—	3.47	—

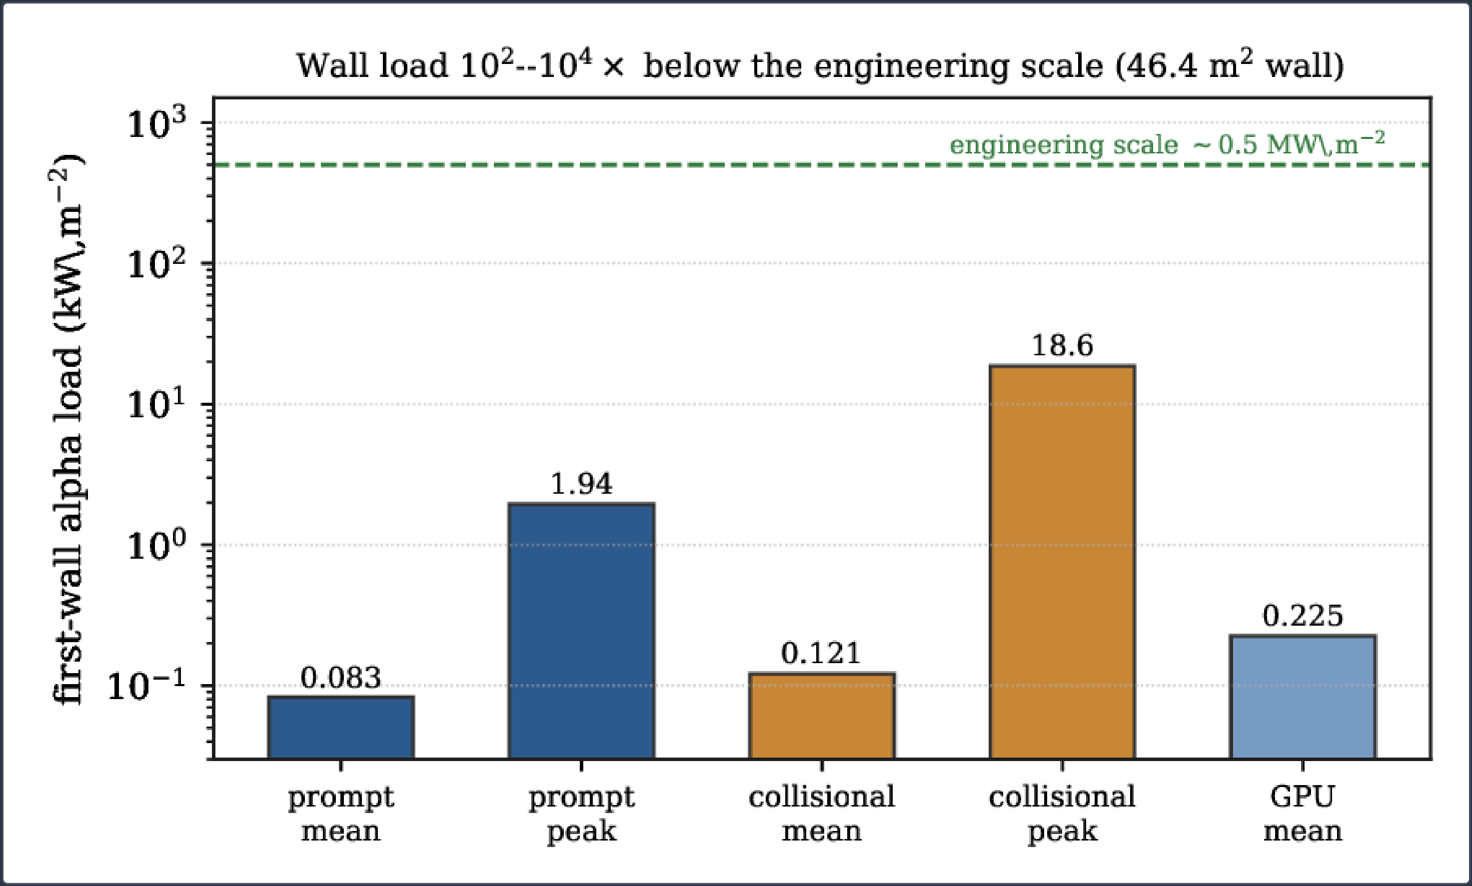

First-wall alpha loads from equation (13) (serial BR-L2-A20), on a logarithmic axis. Prompt and collisional mean/peak loads, and the GPU-follower mean, all sit far below the $\sim 0.5\,MW/m^2$ engineering scale (green dashed). The compact wall ($46.4\,m^2$) is the conservative denominator.

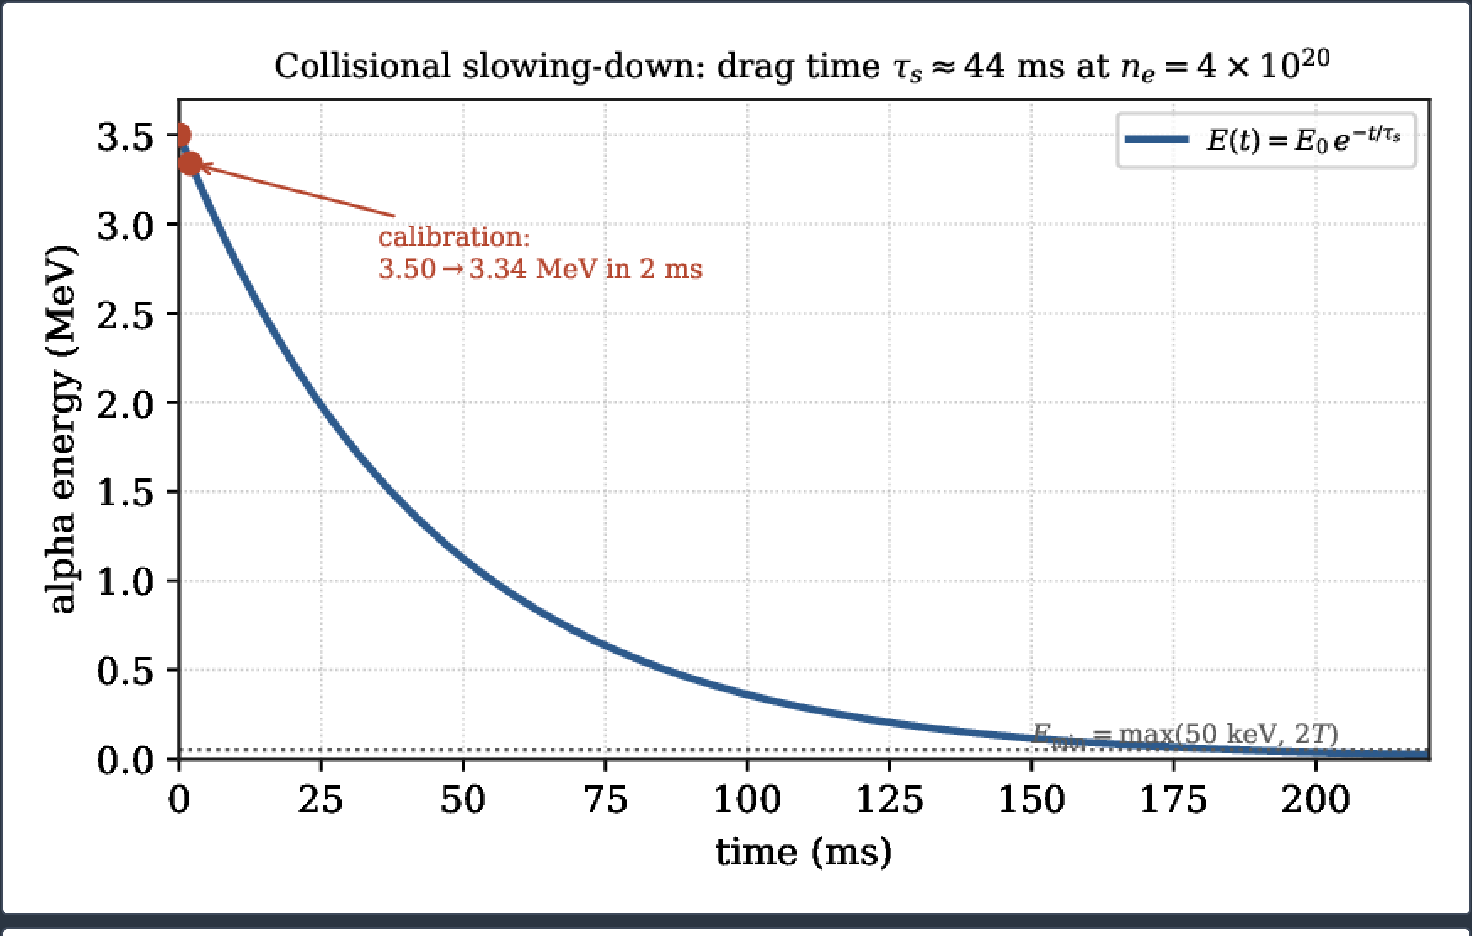

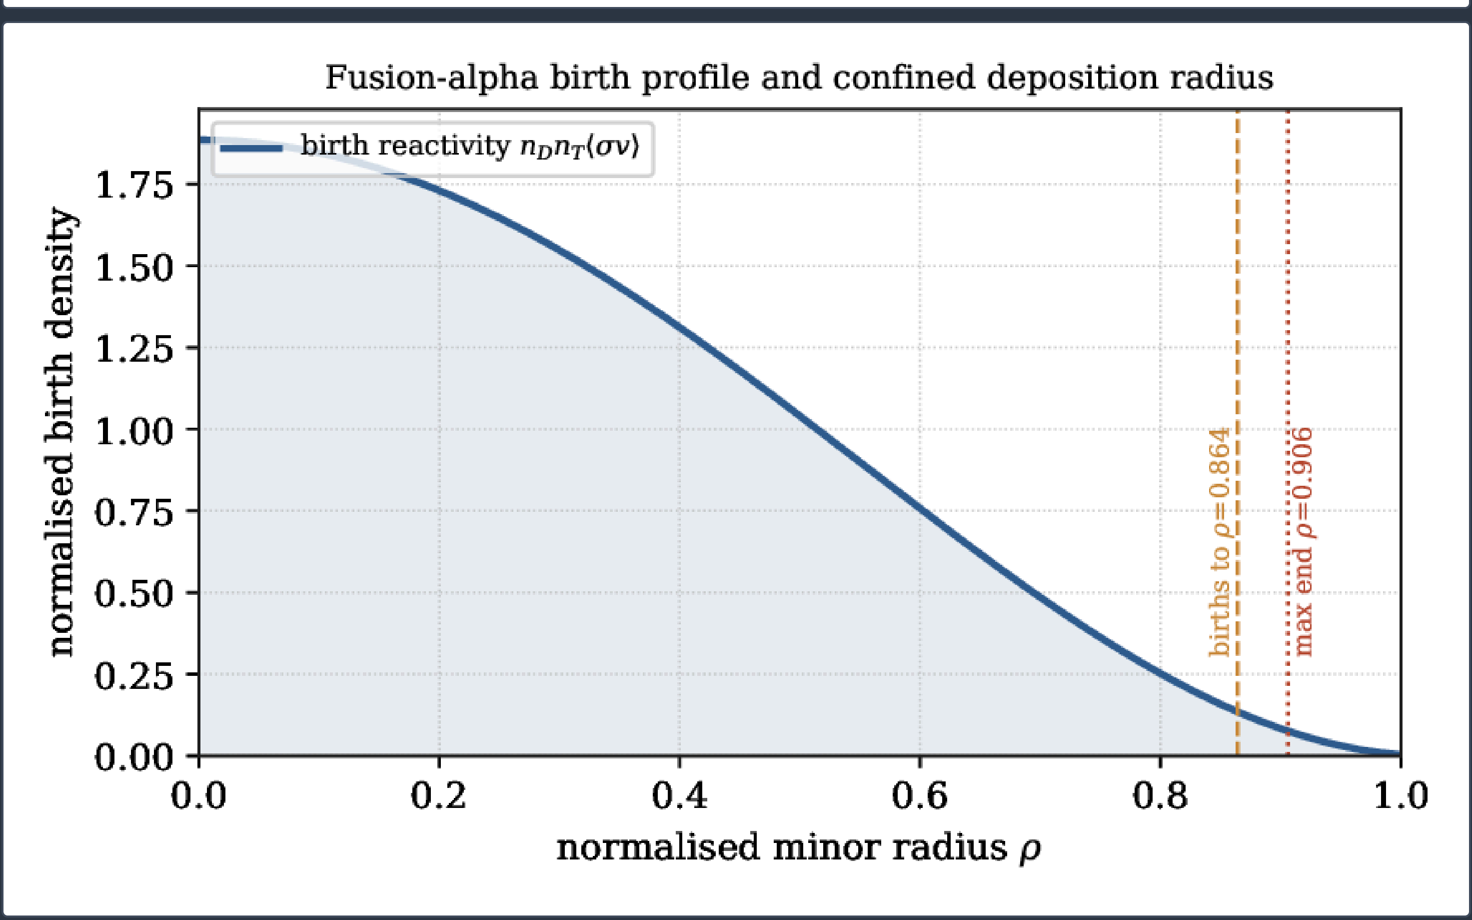

Collisional slowing-down of a confined alpha (derived model): energy e-folding drag time $\tau_s \approx 44\,ms$ at $n_e = 4e20\,/m^3$, with the ASCOT5 calibration point ($3.50\,MeV \rightarrow 3.34\,MeV$ in $2\,ms$). Thermalisation floor E_{min} marked.

TOROIDAL-FIELD RIPPLE

Finite ripple opens the second channel of section 2.6. Table 2 and figure 7 give the loss fraction and peak wall load against ripple amplitude. At 1.0% ripple the loss is 0.019% and at 1.5% it is 0.031% — a factor of roughly two increase over the ripple-free value but still well below a percent, consistent with ripple amplitudes staying below the Goldston–White–Boozer stochasticity threshold (12) for most of the trapped population. The peak first-wall load remains of order 1 $\mathrm{kW\,m^{-2}}$, and the toroidal peaking factor is modest (≈ 1.4): ripple losses spread over the coil period rather than concentrating into a single hot stripe. The independent GPU follower, which carries ripple natively over 10^6 markers, returns a consistent 99.99% confined figure, closing the cross-code loop.

RIPPLE AMPLITUDE δ_{rip}	ALPHA LOSS FRACTION	PEAK WALL LOAD ($\mathrm{kW\,m^{-2}}$)
0 (axisymmetric)	0.023%	1.94
1.0%	0.019%	≈ 1.0
1.5%	0.031%	≈ 1.1

Alpha loss fraction and peak first-wall load versus toroidal-field ripple amplitude (frozen source manuscript; serials BR-L1-A14, BR-L2-A20). Even at 1.5% ripple the loss stays sub-percent and the wall load of order 1 $\mathrm{kW\,m^{-2}}$; the toroidal peaking factor is ≈ 1.4.

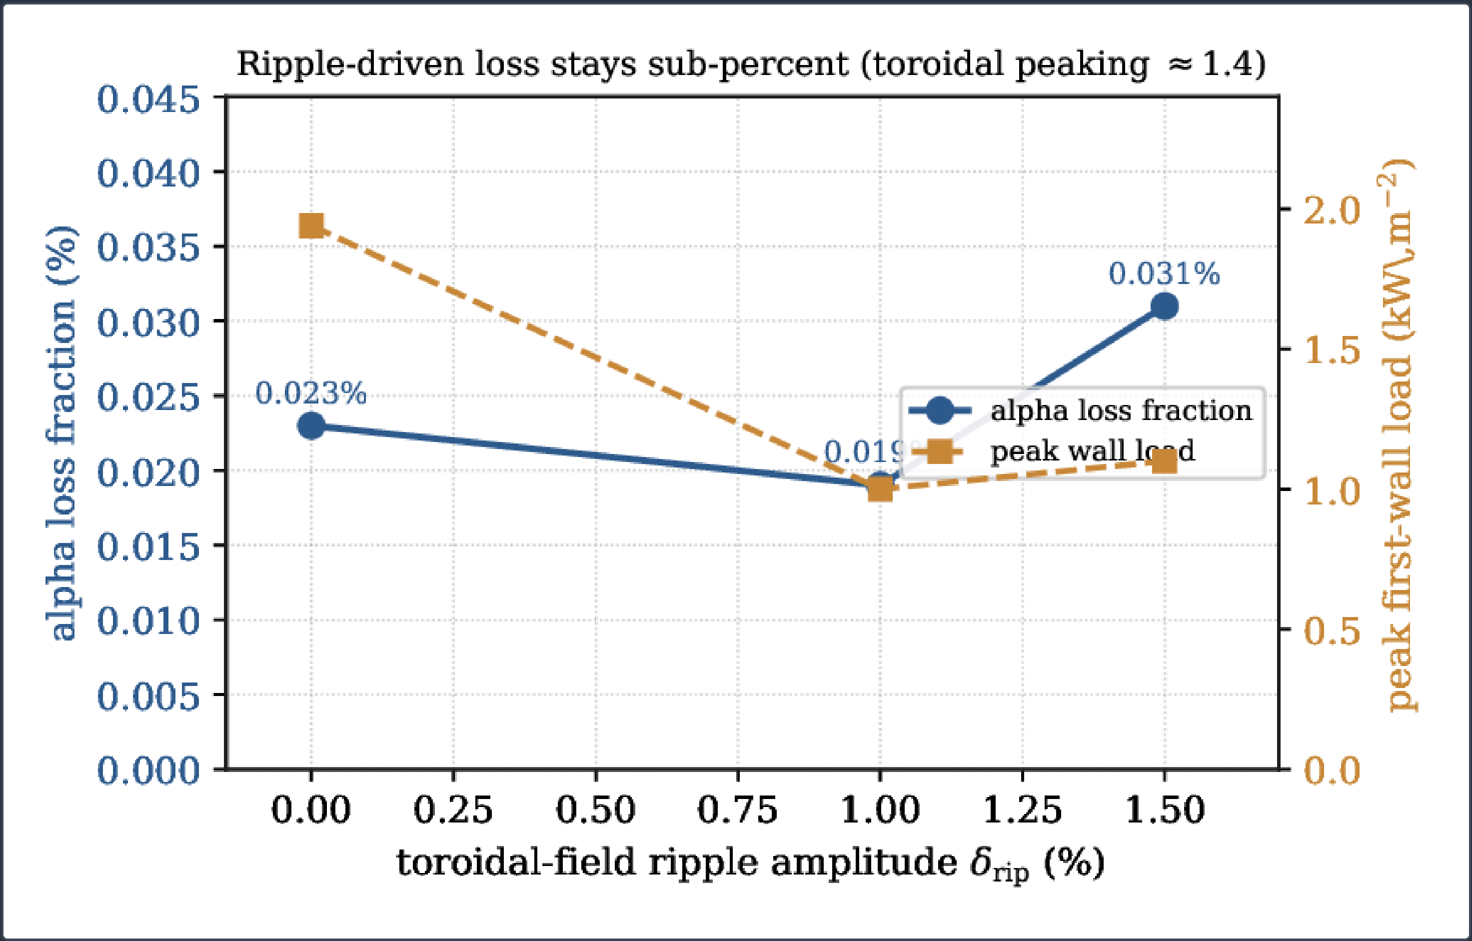

Alpha loss fraction (left axis) and peak first-wall load (right axis) versus toroidal-field ripple amplitude (serials BR-L1-A14, BR-L2-A20). The loss stays sub-percent across the swept ripple range and the peak wall load of order 1 $\mathrm{kW\,m^{-2}}$; the toroidal peaking factor is ≈ 1.4.

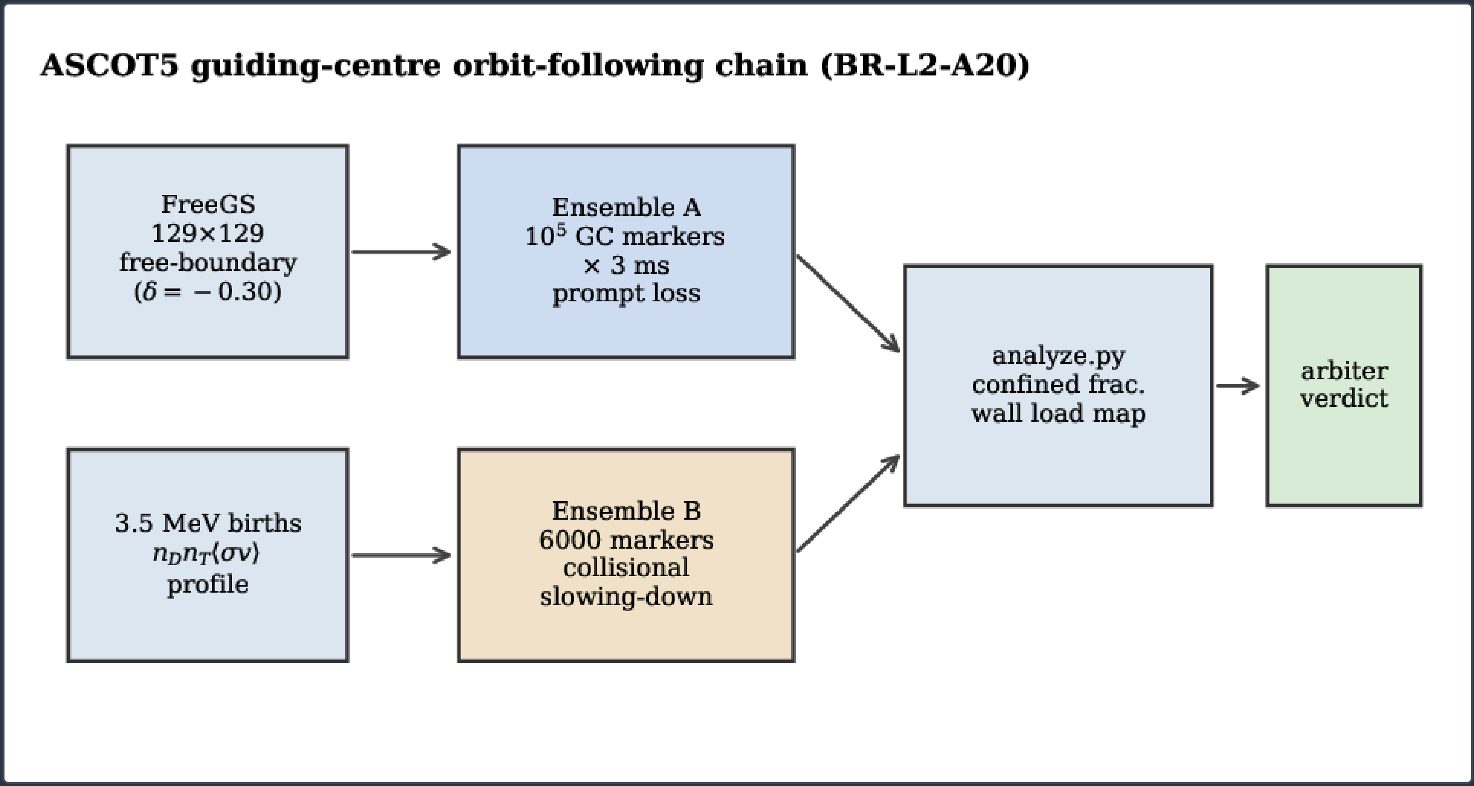

(Schematic.) The two-ensemble ASCOT5 orbit-following chain (serial BR-L2-A20): the frozen FreeGS free-boundary equilibrium and the fusion-reactivity birth source feed Ensemble A (10^5 markers, prompt loss) and Ensemble B (6000 markers, collisional slowing-down and wall load); the analysis stage returns the confined fraction, the wall-load map, and the arbitration verdict.

§ 5

Discussion

QUANTITATIVE COMPARISON WITH THE PUBLISHED FAST-ION LITERATURE

The result should be read against the fast-ion record. In large tokamaks, dedicated orbit-following places fusion and beam-ion confinement at the 99.9% level with first-wall loads managed to the sub- $\mathrm{MW/m^2}$ scale: the ITER energetic-ion analyses of Fasoli *et al.* [1] and Pinches *et al.* [2], and the ASCOT power-load maps of Kurki-Suonio *et al.* [8], are the reference points. Our confined fraction of 99.977% is comfortably within that band, which is the non-trivial finding: the ST orbit-width hierarchy of figure 2 ($\Delta_b/a = 0.33$, $w_p/a = 0.48$) might have been expected to degrade confinement toward the shell estimate's 4.5%, and it does not. The reason is structural — p_φ conservation (2) holds the wide orbits inside the plasma — and is exactly the physics the spherical-torus energetic-particle review of McClements and Fredrickson [16] identifies as the reason ST fast-ion confinement is better than orbit-width scaling alone suggests. The absolute wall loads are higher *per unit loss* than in a reactor-scale device only because the wall area is small ($46.4\,m^2$); even so, the mean loads of table 1 are two-to-four orders of magnitude below the engineering scale.

RIPPLE LOSSES IN CONTEXT

The ripple increment — from 0.023% to 0.031% as δ_{rip} rises to 1.5% — is small because the design sits below the stochastic-diffusion threshold (12) for the bulk of the trapped population, and the modest toroidal peaking (≈ 1.4) shows that what is lost spreads over the coil period rather than forming a stripe. This is consistent with the classical Goldston–White–Boozer picture [11] and with the general finding that ripple-driven fast-ion loss is a steep function of amplitude that stays benign below threshold [1]. It also frames the one real engineering input: the actual coil-set ripple map, which this study brackets with representative amplitudes rather than computing from a specific coil geometry (see section 6).

IMPLICATIONS FOR THE DESIGN POINT

Two design conclusions follow. First, alpha-heating retention is secure: 99.977% of the $P_\alpha = 17.01\,MW$ of birth power is deposited inside the plasma, which is the confinement premise beneath the gain $Q = 3.076$ and the integrated power balance of the platform [31]. Second, first-wall survivability against fast alphas is not a constraint: the loads of table 1 leave the plasma-facing-component lifetime to be set by the $14\,MeV$ neutron budget [35] and by the surface thermomechanics [36], not by alpha loading. The thermal-confinement basis on which this operating point rests is independently established by the negative-triangularity turbulence-suppression result [32] and its two-code cross-validation [33], and the equilibrium on which the present orbits are followed is the same free-boundary solution used by the ideal-MHD stability assessment [34]. The fast-ion populations from auxiliary heating and current drive [37] share the orbit topology analysed here and are confined by the same p_φ mechanism.

§ 6

Limitations

The honest gate is one of scope. This is a guiding-centre, orbit-confinement assessment: it resolves the fast gyromotion into a drift and does not carry full-orbit finite-Larmor-radius effects at the boundary, which could matter for the small population of orbits whose turning points graze the wall. The first-wall load distribution depends on the conformal wall geometry taken from the design point, and the ripple study uses representative amplitudes rather than a specific coil-set ripple map — computing the true $\delta_{\mathrm{rip}}(R, Z)$ from the as-designed toroidal-field coil set, and re-evaluating equation (11) locally, is the natural follow-on. Most importantly, magnetohydrodynamically driven fast-ion transport — resonant redistribution and loss by Alfv\'en eigenmodes and other energetic-particle modes [18,19] — is outside the scope of an orbit-confinement calculation and is a genuinely separate question: at β_α well below the drive threshold this channel is expected to be quiescent, but it must be established by a dedicated fast-ion-MHD stability study rather than assumed. Within its stated scope, the orbit-confinement result is robust and cross-code verified.

§ 7

Conclusion

The frozen compact spherical-tokamak breeder confines its fusion-born alphas essentially completely. Full guiding-centre orbit-following on the identical free-boundary equilibrium, run on the NVIDIA GPU HPC campaign, returns a confined fraction of 99.977% (prompt) with a first-wall alpha load of $0.083\ \mathrm{kW\,m^{-2}}$ mean and $1.94\ \mathrm{kW\,m^{-2}}$ peak, rising only to 0.033% loss and $0.121\ \mathrm{kW\,m^{-2}}$ mean under collisional slowing-down; an independent 10^6-marker GPU follower reproduces the figure at 99.99%. This refutes the analytic large-orbit shell estimate of 4.5%, and the refutation is physical: the potato-width shell argument overcounts because it ignores conservation of the canonical toroidal momentum p_φ, which holds the wide banana and potato orbits inside the plasma despite excursions of one-third to one-half the minor radius. Finite toroidal-field ripple keeps the loss sub-percent (0.019–0.031% at 1.0–1.5% ripple) and the wall load of order $1\ \mathrm{kW\,m^{-2}}$ with only modest toroidal peaking. Alpha confinement is therefore a solved, non-limiting aspect of the Hyperion design — alpha-heating retention is secure and first-wall survivability against fast alphas is not a constraint — with the single open question, fast-ion-MHD transport, deferred to the companion stability work as a matter of scope, not of concern.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The authors thank the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and acknowledge the ASCOT5 developers (Aalto University / VTT) whose orbit-following code is the reference tool for the fast-ion gate.

Reproducibility. Every headline quantity in this paper is a frozen anchor in the Kronos Physics De-Risking Register [30] (doi:10.5281/zenodo.22133057): the fast-alpha orbit-width hierarchy and analytic shell estimate (serial BR-L1-A14), and the 10^5-marker ASCOT5 orbit-following arbitration with its collisional ensemble and first-wall load map (serial BR-L2-A20). Each is regenerable from the archived FreeGS equilibrium, birth source, integrator tolerances and fixed seed (20260726) deposited with this paper, which is archived at doi:10.5281/zenodo.22645729; official version: https://www.kronosfusionenergy.com/publications/alpha-fast-product-confinement.html. No result depends on unpublished or proprietary data, and no economic analysis is included in the public deposit.

Data availability tocsectionData availability The orbit-following loss records, the first-wall load map, and the figure-generation script used for figures 1–8 are archived with the deposit at doi:10.5281/zenodo.22645729 and trace to the register [30] by the serial identifiers cited inline. Any economic analysis is reported separately and is not part of the public deposit.

Competing interests tocsectionCompeting interests Provisional patent applications relating to aspects of this work have been filed.

tocsectionReferences

APPARATUS

Portfolio, People & Record

The intellectual-property estate, the people who built it, the limitations that gate it, and the sources it rests on.

1	4	40	27
GRANTED PATENT	2026 PROVISIONALS	TEAM MEMBERS	2022 PROVISIONALS

INTELLECTUAL PROPERTY

The patent portfolio

The estate spans a granted high-field magnet patent, pending utility and 2026 provisional filings that map to the two products, a digital-twin control provisional, a registered trademark application, and the original 2022 provisional family. Every patentable disclosure in the five-paper arXiv drop was protected **before** publication: the breeder and burner provisionals were filed 1–2 August, and a publication-gap omnibus filing the evening before the drop swept up the DEC, plasma-control, and REBCO-winding matter of the three otherwise-unprotected papers. Patent numbers and application serials are matters of public record; claim scope is summarised, not reproduced.

GRANTED & PENDING UTILITY

US 12,009,112	High-field tilted graded-REBCO magnet architecture (KRONO-002A) · *app. 17/878,550*	GRANTED
US 17/878,507	Core device / method (KRONO-001A) · *utility*	PENDING

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

64/124,209	Hyperion — spherical-tokamak tritium / helium-3 foundry · *breeder · filed Aug 2026*	FILED
64/124,220	Aegis / MetroVolt — synchrotron-cutoff, fuel-selection & plug-winding · *generator · filed Aug 2026*	FILED
64/115,167	KRONOS-CTRL digital-twin plant control · *provisional · filed Jul 2026*	FILED
64/005,440	Integrated spherical-tokamak fusion architecture · *early integrated-architecture filing · filed Mar 2026*	FILED
64/105,530	Negative-triangularity spherical-tokamak architecture · *foundational ST filing · filed Jul 2026*	FILED
64/128,097	Publication-gap omnibus — quasineutral two-species expander DEC · provenance-gated / latency-split plasma control · as-built high-field winding & conductor architecture · *filed 7 Aug 2026, the evening before the arXiv drop*	FILED

TRADEMARK & ORIGIN

FTK-98801894	Kronos Fusion Energy — trademark application	FILED
KR-2022-*	Original provisional family (KR-2022-00001 … 00027) · *27 filings, 2022*	PRIORITY

The narrow, defensible magnet novelty — the specific conductor and the digital-twin winding optimisation, not high-field REBCO as a category — is the commercial core that can earn ahead of any Q > 1 milestone, with markets in fusion magnet supply, MRI/NMR, accelerators, and proton therapy. The claim is scoped honestly: the high field is a *system field* result, not a stand-alone-coil record; the small-bore plug coil is structurally infeasible as a bare winding but resolved by a stress-managed structural shield (feasible-pending-FEA); and the winding-tape experiment returned a *null result*. The value is the method, not a field record.

THE TEAM

Founding partners, board, advisors & auditors

Kronos is built by a bench of advanced-fuel-fusion, high-field-magnet, direct-conversion, and materials specialists, with a board and operations team drawn from defense, national laboratories, and industry. Dates are shown as ranges; where a tenure has ended or is term-ending, the range says so.

FOUNDING PARTNERS — SCIENCE

Dr. Gerald Kulcinski
Co-Founder · Helium-3 & Advanced-Fuel Fusion
UW–Madison · 2023–Present

Dr. Carl Weggel
Chief Scientist / S.M.A.R.T. Design
MIT Alcator Program · 2022–Present

Dr. Robert J. Weggel
Magnetic Field Design
MIT Magnet Lab · 2022–Present

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

Bandel Carano
Finance / Strategy
Oak Investment Partners / Stanford · 2025–2026

Patrick Schweiger
Fusion Device Engineering
Oklo / CFS / TerraPower · 2025–Present

SCIENTIFIC ADVISORS

Dr. Steven O. Dean
Fusion Policy & History
DOE / Fusion Power Associates · 2025–Present

Dr. Patrick H. Diamond
Plasma Turbulence & Transport
UC San Diego · 2025–Present

Dr. Jack J. Dongarra
HPC & Numerical Algorithms
Turing Award · 2025–Present

Dr. Nasr M. Ghoniem
Chief Material Scientist
UCLA · 2025–Present

Dr. Siegfried Glenzer
Plasma Physics & Laser Fusion
Stanford / SLAC · 2022–Present

Dr. David A. Hammer
Pulsed-Power Fusion
Cornell · 2025–Present

Dr. Donald A. Spong
Plasma Theory & Confinement
ORNL · 2025–Present

Dr. Curtis Smith
Risk Assessment & Nuclear PRA
MIT / Idaho National Lab · 2025–Present

RADM (Ret.) David Goggins
Naval Engineering & Power-Plant Design
U.S. Navy · 2023–2026

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

Dr. Ruben Fair
Magnet Design / ITER Liaison
PPPL / Jefferson Lab · 2022–2025

Dr. Paul S. Weiss
Materials & Nanotechnology
UCLA · 2022–2025

SCIENTIFIC AUDITORS

Dr. Wilfred A. Cooper
Plasma Equilibrium Theory
EPFL / Swiss Plasma Center · 2025–Present

Dr. Ahmed Hassanein
Plasma–Material Interactions
Purdue / Argonne · 2025–Present

Dr. Patrick E. Hopkins
Extreme Thermal Materials
U. of Virginia · 2025–Present

Dr. Peter Hosemann
Materials Joining & Structural Integrity
UC Berkeley · 2025–Present

Dr. Travis W. Knight
Advanced Nuclear Fuels & Systems
U. of South Carolina · 2025–Present

Dr. Philippe Lebrun
Cryogenics & Magnet Cooling
CERN · 2025–Present

Dr. Nitendra Singh
Nuclear Safety & Fuel Cycle
ITER · 2025–Present

Dr. Guido Van Oost
Magnetic Confinement & Fusion Systems
Ghent University · 2025–Present

Dr. Gary S. Was
Radiation Materials & Structural Integrity
U. of Michigan · 2025–Present

Dr. Ray Sedwick
Direct Energy Conversion
U. of Maryland · 2025

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

Martin Owens
Chief Strategy Officer
LANL / GE Hitachi · 2022–Present

MG (Ret.) Paul Pardew
Chief Contracting Officer
Army Contracting Command · 2022–Present

David Beck
Fusion Commercialization
U.S. Space Force · 2024–Present

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

MG (Ret.) Robin L. Fontes
Cybersecurity & Defense
Army Cyber Command · 2022–Present

Vijay Gehani
Supply Chain & Components
INOX India · 2024–Present

Jon Michel Greenwood
IT & AI Infrastructure
Live Nation · 2023–Present

Brian C. O'Neill
National Security
CIA / ODNI · 2022–Present

Gen. (Ret.) Gustave F. Perna
DoD Liaison
Operation Warp Speed · 2022–2023

Andrea Romero
Marketing Lead
2022–Present

Legal counsel is retained; those roles are held on the internal roster and are not listed here.

SOURCES

References

1 A Fasoli *et al*, Chapter 5: Physics of energetic ions, *Nucl. Fusion* **47** S264 (2007). doi:10.1088/0029-5515/47/6/S05

2 S D Pinches *et al*, Energetic ions in ITER plasmas, *Phys. Plasmas* **22** 021807 (2015). doi:10.1063/1.4908551

3 M E Austin *et al*, Achievement of reactor-relevant performance in negative triangularity shape in the DIII-D tokamak, *Phys. Rev. Lett.* **122** 115001 (2019). doi:10.1103/PhysRevLett.122.115001

4 A Marinoni, O Sauter and S Coda, A brief history of negative triangularity tokamak plasmas, *Rev. Mod. Plasma Phys.* **5** 6 (2021). doi:10.1007/s41614-021-00054-0

5 E Hirvijoki *et al*, ASCOT: solving the kinetic equation of minority particle species in tokamak plasmas, *Comput. Phys. Commun.* **185** 1310 (2014). doi:10.1016/j.cpc.2014.01.014

6 J Varje *et al*, High-performance orbit-following code ASCOT5 for Monte Carlo simulations in fusion plasmas, arXiv:1908.02482 (2019).

7 A Pankin, D McCune, R Andre, G Bateman and A Kritz, The tokamak Monte Carlo fast ion module NUBEAM in the National Transport Code Collaboration library, *Comput. Phys. Commun.* **159** 157 (2004). doi:10.1016/j.cpc.2003.11.002

8 T Kurki-Suonio *et al*, ASCOT simulations of fast ion power loads to the plasma-facing components in ITER, *Nucl. Fusion* **49** 095001 (2009). doi:10.1088/0029-5515/49/9/095001

9 R B White and M S Chance, Hamiltonian guiding center drift orbit calculation for plasmas of arbitrary cross section, *Phys. Fluids* **27** 2455 (1984). doi:10.1063/1.864527

10 J R Cary and A J Brizard, Hamiltonian theory of guiding-center motion, *Rev. Mod. Phys.* **81** 693 (2009). doi:10.1103/RevModPhys.81.693

11 R J Goldston, R B White and A H Boozer, Confinement of high-energy trapped particles in tokamaks, *Phys. Rev. Lett.* **47** 647 (1981). doi:10.1103/PhysRevLett.47.647

12 A H Boozer, Plasma equilibrium with rational magnetic surfaces, *Phys. Fluids* **24** 1999 (1981). doi:10.1063/1.863297

13 T H Stix, Heating of toroidal plasmas by neutral injection, *Plasma Phys.* **14** 367 (1972). doi:10.1088/0032-1028/14/4/002

14 H-S Bosch and G M Hale, Improved formulas for fusion cross-sections and thermal reactivities, *Nucl. Fusion* **32** 611 (1992). doi:10.1088/0029-5515/32/4/I07

15 R L Miller, M S Chu, J M Greene, Y R Lin-Liu and R E Waltz, Noncircular, finite aspect ratio, local equilibrium model, *Phys. Plasmas* **5** 973 (1998). doi:10.1063/1.872666

16 K G McClements and E D Fredrickson, Energetic particles in spherical tokamak plasmas, *Plasma Phys. Control. Fusion* **59** 053001 (2017). doi:10.1088/1361-6587/aa626e

17 M Ono *et al*, Exploration of spherical torus physics in the NSTX device, *Nucl. Fusion* **40** 557 (2000). doi:10.1088/0029-5515/40/3Y/316

18 W W Heidbrink, Basic physics of Alfv\'en instabilities driven by energetic particles in toroidally confined plasmas, *Phys. Plasmas* **15** 055501 (2008). doi:10.1063/1.2838239

19 D J Sigmar, C T Hsu, R White and C Z Cheng, Alpha-particle losses from toroidicity-induced Alfv\'en eigenmodes. Part II: Monte Carlo simulations and anomalous alpha-loss processes, *Phys. Fluids B* **4** 1506 (1992). doi:10.1063/1.860061

20 J E Menard *et al*, Fusion nuclear science facilities and pilot plants based on the spherical tokamak, *Nucl. Fusion* **56** 106023 (2016). doi:10.1088/0029-5515/56/10/106023

21 B N Sorbom *et al*, ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* **100** 378 (2015). doi:10.1016/j.fusengdes.2015.07.008

22 A J Creely *et al*, Overview of the SPARC tokamak, *J. Plasma Phys.* **86** 865860502 (2020). doi:10.1017/S0022377820001257

23 L L Lao, H St John, R D Stambaugh, A G Kellman and W Pfeiffer, Reconstruction of current profile parameters and plasma shapes in tokamaks, *Nucl. Fusion* **25** 1611 (1985). doi:10.1088/0029-5515/25/11/007

COLOPHON

How this document was made

The Kronos Fleet — Combined Editorial, 2026 is set in **Fraunces** (display), **Gelasio** (text), and **IBM Plex Mono** (data & equations), carried forward from the Kronos editorial design system.

The figures are rendered at 300 dpi from the deposited generator scripts; the document is composed as a single self-contained file with embedded fonts and imagery, paginated to US Letter, and rendered to PDF through a headless Chromium engine in matched light and dark editions.

Every headline quantity carries an evidence class; withdrawn and restated values are marked as such. Nothing herein is frozen beyond the founder-approved Tier-A set.

Explore it live — turn the interactive 3-D model, run the deposited solvers in the physics validator, and watch the companion film series.

LEGAL NOTICE

Notice, status, and distribution

Conceptual design & simulation study. This document reports a conceptual design and simulation study. It is not a construction commitment, a safety-analysis report, a regulatory filing, or an offer of securities. Forward-looking statements — schedules, costs, market sizes, and performance — are estimates subject to the limitations set out in the Simulations part and may change.

Numbers and their classes. Quantities are reported with evidence classes and case labels; modelled, assumed, and requirement-class values are identified as such and are not to be quoted without their labels.

Public edition. This edition carries the physics and design only; all financial, funding, and defense-commercial content has been removed for public distribution. The scientific text corresponds to the arXiv and journal submissions.

THE 2026 KRONOS SERIES

One programme, many papers

This editorial is one paper in the Kronos 2026 series — a real fusion company's complete, reproducible physics record for a spherical-tokamak breeder and its low-neutron D–^{3}He tandem-mirror burners. Every headline number is a frozen anchor in the Physics De-Risking Register.

Explore the whole programme: the De-Risking Register, the interactive 3-D model, and the physics validator.

KRONOS FUSION ENERGY

Fast-Ion and Alpha-Particle Confinement in a Compact Spherical-Tokamak Breeder: Sub-Percent Ripple Losses and First-Wall Loading from Orbit-Following Simulation

A real fusion company building real machines. Explore the science, live.

Zenodo · doi:10.5281/zenodo.22645729

Read on kronosfusionenergy.com

The Physics De-Risking Register