KRONOS FUSION ENERGY

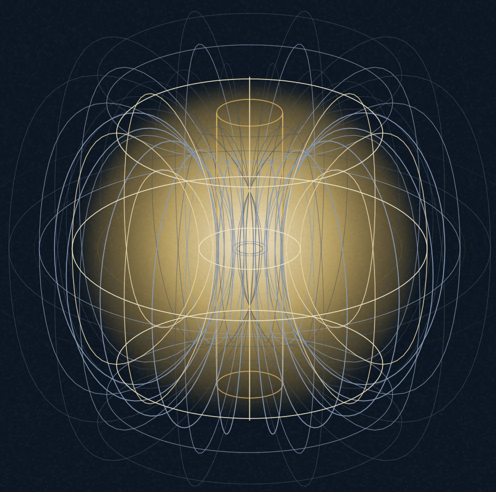

PAPER 1.1 · THE KRONOS 2026 SERIES

Hyperion as a Strategic-Materials Platform: Gain, Tritium Breeding, He-3 Co-Production, and 14 MeV Neutron Output from a Compact Spherical-Tokamak

Fusion breeders have historically foundered on two coupled obstacles: net electricity forces a large plasma and a large blanket, and a self-sufficient closed fuel cycle at...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Hyperion as a Strategic-Materials Platform: Gain, Tritium Breeding, He-3 Co-Production, and 14 MeV Neutron Output from a Compact Spherical-Tokamak

Fusion breeders have historically foundered on two coupled obstacles: net electricity forces a large plasma and a large blanket, and a self-sufficient closed fuel cycle at power-plant scale implies a dangerous standing tritium inventory. Both are consequences of demanding electricity from a device whose neutrons are physically its most valuable output. We present the integrated physics of Hyperion, a compact negative-triangularity spherical-tokamak breeder sized to a strategic-materials mission — tritium, helium-3, and 14 MeV neutrons — and show that once the net-electricity requirement is dropped, both obstacles fall below their binding thresholds. The frozen design point ($I_p = 9.66$ MA, on-axis $B_0 = 8.0\,\mathrm{T}$, centrepost peak field $16.84\,\mathrm{T}$, $\delta = -0.30$, $\kappa = 2.0$, $R_0 = 1.20\,\mathrm{m}$, $A = 2.5$) closes its zero-dimensional power balance at a plasma gain $Q = 3.076$ and fusion power $P_{\mathrm{fus}} = 85.04$ MW, carried openly as a driven, sub-ignition design-point target reached at demonstrated confinement. We state the governing equations of every subsystem as explicit mathematics — the Bosch–Hale power balance, the free-boundary Grad–Shafranov equilibrium and its ideal-MHD margins, the Sauter neoclassical bootstrap, the nonlinear-gyrokinetic transport ratio, the neutron-transport tritium-breeding functional, the tritium and helium-3 mass-balance ordinary differential equations, and the Rosenbluth–Putvinski runaway avalanche — carry each derivation from its first principle to its discretized numerical form, and evaluate them on a verification-first computational stack (Kronos toolkit, FreeGS/FreeGSNKE, NEO, nonlinear CGYRO, TGYRO, ASCOT5, OpenMC on ENDF/B-VIII.0, and a NIMROD-verified disruption model), with the nonlinear gyrokinetics and Monte-Carlo neutronics executed on the NVIDIA GPU HPC campaign. Negative triangularity is the confinement lever: converged nonlinear gyrokinetics show it suppressing ion-scale transport by 40–45% and electron-scale (ETG) transport by 79–80% relative to a positive-triangularity twin, with the operating point subcritical and the advantage robust (56–66%) across a factor-sixteen collisionality scan. The free-boundary equilibrium is Mercier-, ballooning- and external-kink-stable with a $1.64\times$ normalized-beta margin to the no-wall Troyon limit; the current is 85% externally driven with a 15.5% bootstrap fraction; fusion-born alphas are 99.98% confined; and tritium breeding rises along a confirmed ladder from a net surplus of 0.29 to $\sim 2.0 kg-T\,\mathrm{yr}^{-1}$ per unit at the advanced blanket (net breeding ratio 1.42), making a single advanced unit a tritium exporter and a national-scale $4 kg-T\,\mathrm{yr}^{-1}$ supply an explicit fleet property. The inboard centrepost is a quantified consumable cartridge, and the disruption response is answered directly and is demanding: a severe runaway-electron avalanche makes shattered-pellet mitigation a hard design requirement. We close with a dedicated sensitivity/uncertainty section and an experiment-and-theory benchmarking section. Every headline number is a frozen anchor with a gate serial, independently re-run and stamped in the accompanying physics de-risking register.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645691 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: spherical tokamak · tritium breeder · helium-3 · negative triangularity · materials platform · 14\ MeV neutron source · consumable centrepost · disruptions · bootstrap current · Grad--Shafranov equilibrium · uncertainty quantification

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

Fusion breeders are not a new idea, and they have failed for a consistent reason. To make net electricity a device must be large — a large plasma to hold the power, a large blanket to intercept it — and to close its own fuel cycle at that scale it must hold a standing tritium inventory measured in kilograms, with all the licensing, safety, and proliferation weight that inventory carries [51,37,38]. Both burdens are downstream of one requirement: net electricity from a device whose neutrons are, physically, the most valuable thing it makes. Strip that requirement away, size the machine instead to the *materials* its neutrons produce — tritium for the fusion economy, helium-3 for national-security and quantum applications, and a fast 14 MeV neutron flux for isotope production and materials qualification — and the two classical obstacles both drop below the thresholds that stopped the earlier concepts. The plasma can be compact, and the per-unit tritium inventory a fraction of a power plant's.

The scale of the tritium problem is worth stating quantitatively at the outset, because it motivates the entire design. The world civilian tritium inventory is of order 20–25 kg and declining, sourced almost entirely from a handful of ageing heavy-water fission reactors; a single deuterium–tritium power plant of ~ 3 GW fusion power would burn of order $150 kg - T\,\mathrm{yr}^{-1}$ and must therefore breed essentially all of its own fuel with a required breeding ratio comfortably above unity once one accounts for radioactive decay, sequestration in structures, processing losses, and the doubling inventory needed to fuel the next unit [37,38]. The consensus of the tritium fuel-cycle literature is that self-sufficiency for a power plant is achievable but tight, and that it forces both a large tritium-breeding blanket and a large in-vessel inventory. A device that must supply its own gigawatt-scale burn is therefore boxed in from two sides at once. Hyperion escapes the box not by solving the power-plant fuel cycle but by declining to pose it: a strategic-materials breeder burns $\Lambda_T = 4.77 kg - T\,\mathrm{fpy}^{-1}$, roughly thirty times less than a power plant, so its inventory is small, its blanket need only clear unity by a modest margin to become a net *exporter*, and its success is measured in kilograms of product rather than megawatts to a grid.

The earlier breeder concepts illustrate the trap and its shape. Fission-suppressed and tritium-assisted fusion breeders were proposed to leverage a modest fusion neutron source into a large fissile or tritium yield, but were sized to net-power or satellite-power missions and inherited either a large machine or a large fuel-cycle inventory [51,53]; inertial-electrostatic and other low-power-density D–^{3}He breeder concepts avoided the inventory problem but could not reach a useful volumetric reaction rate [52]. The spherical-tokamak pilot-plant studies [1,5,4] showed that the ST could deliver a compact, high-fluence neutron source, but framed the device as a step toward net electricity and therefore re-imported the scale burden. The distinguishing move here is neither a new confinement concept nor a new blanket chemistry but the mission reframe: sizing the ST to the smallest self-consistent materials point, and treating breeding as an export lever rather than a self-sufficiency constraint, is what lets the two obstacles fall together (§2).

Hyperion is the compact spherical tokamak (ST) that results. A spherical tokamak confines a large plasma current in a small volume behind a high-field inboard centrepost, which is exactly what makes it an efficient neutron source and breeder platform [1,2,3,4,5]. The ST line traces to Peng and Strickler's recognition [4] that retaining only the indispensable inboard conductor produces a very-low-aspect-ratio configuration with naturally high elongation, high toroidal beta, and large plasma current per unit field — properties that are liabilities for a reactor that must shield a superconducting central solenoid but are assets for a compact, high-fluence neutron source. We choose negative triangularity ($\delta = -0.30$) because it accesses good core confinement without the edge pressure pedestal and the edge-localised-mode cycle of conventional H-mode [6,7,8,9,11]; the quiet, ELM-free edge is an operational asset for a nuclear device whose first wall cannot tolerate repeated transient loads. The result is a machine that closes its power balance at demonstrated confinement, with the plasma gain carried honestly as a design-point target rather than a demonstrated one.

Contributions and structure. Relative to a conventional design note, this paper is a review-article-length treatment intended to make every headline number reconstructible from stated physics. It (i) states the governing physics of every subsystem as explicit mathematics, deriving each result from its first principle to its discretized form, and fixes a self-consistent compact operating point at $Q = 3.076$, $P_{\mathrm{fus}} = 85.04$ MW (§2–§5); (ii) describes the verification-first computational stack, its numerical schemes, and the role of the NVIDIA GPU HPC campaign (§4); (iii) demonstrates, with converged nonlinear gyrokinetics, that negative triangularity suppresses turbulence at both the ion and the electron scale into a quiet subcritical core, and that the advantage survives a collisionality scan (§6); (iv) shows the point is an ideal-MHD-stable free-boundary equilibrium sustained at an 85%-driven current (§7–§8); (v) confines the fusion-born alphas (§9); (vi) develops the tritium breeding ladder, the mass-balance surplus law, and the startup inventory that together make a single advanced unit a net exporter (§11); (vii) treats the centrepost as a quantified consumable and answers the disruption question directly (§12–§13); (viii) places helium-3 and the strategic products, and shows the deuterium–helium-3 loop does not close indigenously (§14); and (ix) quantifies sensitivity and uncertainty (§15) and benchmarks each subsystem against published experiment and theory (§16). Every number is a frozen anchor, independently re-run and stamped in the de-risking register[R0].

Scope. This is a strategic-materials paper. Its success metric is fusion gain, breeding ratio, helium-3, and 14 MeV neutrons — not megawatts to a grid. No net-electricity claim is made for the breeder, and no economic analysis appears here; economics is reported separately and is not part of the public record of this work.

§ 2

The strategic-materials mission and the design reframe

Before writing equations it is worth stating precisely what mission the equations are being solved against, because the mission — not any single physics choice — is what places both classical obstacles below threshold.

WHY NET ELECTRICITY IS THE WRONG REQUIREMENT FOR A BREEDER

A deuterium–tritium reaction releases $E_{DT} = 17.59$ MeV, of which $E_\alpha = 3.52$ MeV (20%) is carried by the charged alpha and $E_n = 14.06$ MeV (80%) by the neutron. In a power reactor the neutron energy is a means to an end: it is thermalized in a blanket, converted to heat, and run through a Rankine or Brayton cycle at $\sim$35–40% efficiency to make electricity, and the tritium bred in the same blanket exists only to refuel the plasma. In a strategic-materials breeder the ranking is inverted. The neutron itself is the product — it transmutes ^{6}Li into tritium, it can be captured to make medical and industrial isotopes, and it qualifies structural materials at fusion-relevant damage rates — and the tritium and helium-3 are bankable commodities. Demanding electricity from such a device would force a larger plasma (to raise absolute power) and a fully self-sufficient blanket with a large inventory, re-importing exactly the two burdens that the materials framing removes. The design logic of Hyperion is therefore to size the plasma to the smallest self-consistent point that (a) closes its power balance at demonstrated confinement and (b) produces a materials-relevant 14 MeV flux, and to let the blanket breeding ratio be a *lever* tuned for export rather than a fixed self-sufficiency constraint.

THE TWO OBSTACLES, QUANTIFIED

Table 1 states the two historical obstacles and the threshold each must clear, together with where Hyperion lands. The scale obstacle is expressed through the fusion power and major radius; the inventory obstacle through the per-unit burn rate and the in-vessel inventory. In both cases the materials mission moves the operating point from “above threshold” (binding) to “below threshold” (non-binding), and the rest of the paper is the demonstration that the compact point so obtained is physically self-consistent.

@L>X>X@ OBSTACLE	POWER-PLANT FRAMING (BINDING)	HYPERION MATERIALS FRAMING (BELOW THRESHOLD)
machine scale	GW-class P_{fus}, $R_0 \gtrsim 3$–6 m to hold and intercept the power	85.04 MW, $R_0 = 1.20$ m compact ST; power set by materials flux, not by grid demand
tritium inventory	kg-class standing inventory; self-sufficiency mandatory and tight [37,38]	burn $4.77 kg - T\,\text{fpy}^{-1}$ ($\sim 30\times$ smaller); inventory a fraction of a plant's; export not self-sufficiency is the goal
blanket	TBR > 1 mandatory with decay/loss/doubling margin [37]	TBR is a design lever; a modest surplus makes one unit an exporter

The two classical breeder obstacles and how the strategic-materials reframe places each below threshold.

WHY A SPHERICAL TOKAMAK, AND WHY NEGATIVE TRIANGULARITY

Two configuration choices follow from the mission. The ST is chosen because, at fixed engineering field on the inboard conductor, its low aspect ratio maximizes plasma current and toroidal beta and therefore the volumetric fusion (and neutron) rate per unit machine size [4,1,5]: a compact, high-fluence neutron source is precisely what a materials breeder wants. The price — no room for a shielded central solenoid, and an inboard conductor exposed to the source — is paid by making the centrepost a consumable cartridge (§12), which is an acceptable trade for a device whose mission is neutron production rather than decades of uninterrupted electricity. Negative triangularity is chosen because it delivers core confinement comparable to H-mode while remaining in an L-mode-like, ELM-free edge [6,7,8]: it suppresses turbulent transport (§6) without the pedestal and the periodic edge-localised-mode heat pulses that would otherwise threaten the first wall of a nuclear device. The two choices are complementary — the ST gives the neutron economy, negative triangularity gives the quiet, high-confinement core — and together they are what let a 1.20-m machine close its power balance.

§ 3

Governing equations and the formal problem

The design is the simultaneous solution of a coupled set of physics constraints. We state each governing equation here so that every headline number in §5–§14 traces to a defined quantity rather than an assertion, and we carry each derivation to the point where its numerical evaluation is unambiguous.

FUSION POWER BALANCE, THE LAWSON CRITERION, AND GAIN

On a deuterium–tritium plasma of total density $n = n_D + n_T$ with $n_D = n_T = n/2$, the volumetric fusion power density is $p_{\text{fus}} = n_D n_T \langle \sigma v \rangle_{DT} E_{DT} = \frac{1}{4} n^2 \langle \sigma v \rangle_{DT} E_{DT}$, so the total fusion power is

$$P_{\text{fus}} = \int_V n_D n_T \, \langle \sigma v \rangle_{DT}(T_i) \, E_{DT} \, \mathrm{d}V = \tfrac{1}{4} \, \overline{n^2 \, \langle \sigma v \rangle_{DT}} \, E_{DT} \, V, \qquad E_{DT} = 17.59 \quad \text{MeV},$$

where $\langle \sigma v \rangle_{DT}(T_i)$ is the Maxwellian-averaged reactivity and the overbar denotes a profile-weighted volume average. Writing profiles as $n(\rho) = n_0(1-\rho^2)^{\alpha_n}$ and $T(\rho) = T_0(1-\rho^2)^{\alpha_T}$ over the normalized minor radius $\rho = r/a$, the volume average that enters 1 is

$$\overline{n^2 \langle \sigma v \rangle} = \frac{1}{V} \int_0^1 n(\rho)^2 \langle \sigma v \rangle\big(T(\rho)\big) \, \frac{\mathrm{d}V}{\mathrm{d}\rho} \, \mathrm{d}\rho, \qquad \frac{\mathrm{d}V}{\mathrm{d}\rho} = 4\pi^2 R_0 a^2 \kappa \, \rho,$$

which defines the profile form factor $g \equiv \overline{n^2 \langle \sigma v \rangle} / (\bar{n}^2 \langle \sigma v \rangle (\bar{T}))$ used in the POPCON of §5. We evaluate $\langle \sigma v \rangle_{DT}$ with the Bosch–Hale parameterization [12], valid over 0.2–100 keV,

$$\langle \sigma v \rangle_{DT} = C_1 \, \theta \sqrt{\frac{\xi}{m_r c^2 \, T_i^3}} \, \exp(-3\xi), \qquad \theta = \frac{T_i}{1 - \dfrac{T_i(C_2 + T_i(C_4 + T_i C_6))}{1 + T_i(C_3 + T_i(C_5 + T_i C_7))}}, \qquad \xi = \left(\frac{B_G^2}{4\theta} \right)^{1/3},$$

with Gamow constant $B_G = 34.3827$ keV$^{1/2}$, reduced mass $m_r c^2 = 1.124656 \times 10^6$ keV, and the tabulated $C_1 \ldots C_7$ (Appendix ?); the curve peaks near 64 keV (Fig. 3), while the economically accessible operating window sits near 13–17 keV where $\langle \sigma v \rangle_{DT} \propto T_i^2$ approximately and the pressure-limited fusion power is maximized. The fraction of fusion energy carried by the alpha particle,

$$f_\alpha = \frac{E_\alpha}{E_{DT}} = \frac{3.52}{17.59} = 0.200,$$

heats the plasma; the complementary 0.800 leaves as 14.06 MeV neutrons — the device's product. Writing the plasma stored energy $W = \frac{3}{2} \int_V n(T_e + T_i) \, \mathrm{d}V$ and the energy-confinement time $\tau_E = W / P_{\text{loss}}$, the steady-state power balance is

$$\underbrace{f_\alpha P_{\text{fus}}}_{\text{alpha heating}} + P_{\text{ext}} = P_{\text{loss}} + P_{\text{rad}} = \frac{W}{\tau_E} + P_{\text{rad}},$$

and the plasma gain is $Q \equiv P_{\text{fus}} / P_{\text{ext}}$. Combining 5 with the definition of Q and neglecting radiation for the moment gives the compact relation

$$\frac{1}{Q} = \frac{P_{\text{ext}}}{P_{\text{fus}}} = \frac{W/\tau_E}{P_{\text{fus}}} - f_\alpha,$$

which shows explicitly that gain is set by the ratio of stored energy to fusion power divided by the confinement time, offset by the alpha self-heating fraction f_α. Writing $W = 3\bar{n}\bar{T}V$ (for $T_e = T_i = \bar{T}$) and $P_{\text{fus}} = \frac{1}{4} g \bar{n}^2 \langle \sigma v \rangle E_{DT} V$, 6 becomes the Lawson-type condition [13,14]

$$\bar{n} \tau_E = \frac{12 \, \bar{T}}{g \langle \sigma v \rangle E_{DT}} \left(\frac{1}{Q} + f_\alpha \right)^{-1},$$

so that the triple product $\bar{n}\bar{T}\tau_E$ required for a target Q is minimized near the temperature that maximizes $\langle \sigma v \rangle / \bar{T}^2$ — the physical origin of the ∼13–15 keV operating window. Equations 1–7 close the zero-dimensional operating point once the confinement quality (the H-factor multiplying a confinement scaling for τ_E) and the profiles are fixed; the resulting frozen values are given in §5 and their derivation to a self-consistent point is gate BR-L1-A1.

FREE-BOUNDARY EQUILIBRIUM AND STABILITY LIMITS

The axisymmetric equilibrium is the Grad–Shafranov equation for the poloidal flux $\psi(R,Z)$. It follows from force balance $J\times B=\nabla p$ together with $\nabla\cdot B=0$ and Amp\`ere's law; writing the axisymmetric field as $B=\nabla\phi\times\nabla\psi+F(\psi)\nabla\phi$ with $F=RB_\phi$ and using that p and F are flux functions, the toroidal component of force balance reduces to

$$\Delta^*\psi\equiv R\frac{\partial}{\partial R}\left(\frac{1}{R}\frac{\partial\psi}{\partial R}\right)+\frac{\partial^2\psi}{\partial Z^2}=-\mu_0R^2p'(\psi)-F(\psi)F'(\psi),\qquad F=RB_\phi,$$

solved as a free-boundary problem so that the coil currents, not a prescribed boundary, set the shape; the plasma boundary is the outermost closed flux surface consistent with the external poloidal-field-coil currents. The relevant global stability limit is the Troyon normalized-beta limit [27]: with $\beta=2\mu_0\langle p\rangle/B_0^2$ and

$$\beta_N=\frac{\beta\,[\%]\,a\,B_0}{I_p\,[\mathrm{MA}]},$$

ideal-MHD sets a no-wall ceiling $\beta_N^{\max}$; the operating margin is $\beta_N^{\max}/\beta_N$. Local interchange stability is tested by the Mercier criterion, which for a flux surface can be written as $D_M=D_I+D_R>0$ with the interchange term

$$D_I=\frac{\mu_0p'}{\langle B^2\rangle\,(q')^2}[\ldots]-\frac{1}{4},$$

where the bracket collects the field-line-averaged magnetic-well and current terms; the s–α ballooning condition on each flux surface tests high-n pressure-driven modes, and the low-n external kink is tested by an energy-principle metric $\delta W>0$. We evaluate all three (gate BR-L2-A19) and report the margins in §7.

NEOCLASSICAL BOOTSTRAP CURRENT

A solenoid-free ST must sustain its current non-inductively, and the pressure-gradient-driven bootstrap current sets the floor. The bootstrap current arises from the momentum transferred between trapped and passing electrons across a density/temperature gradient; it scales as $j_{\rm bs}\sim-\epsilon^{1/2}(1/B_\theta)\,\mathrm{d}p/\mathrm{d}r$ in the banana regime. We use the Sauter neoclassical formulation [22,23], in which the flux-surface-averaged parallel bootstrap current is

$$\langle j_\parallel B\rangle_{\rm bs}=-I(\psi)\,p_e\left[L_{31}\frac{p}{p_e}\frac{\partial\ln p}{\partial\psi}+L_{32}\frac{\partial\ln T_e}{\partial\psi}+L_{34}\,\alpha\,\frac{p_i}{p_e}\frac{\partial\ln T_i}{\partial\psi}\right],$$

with the coefficients $L_{31},L_{32},L_{34}(\nu_\star,\epsilon)$ and $\alpha(\nu_\star)$ functions of collisionality and inverse aspect ratio (Appendix ?); the bootstrap fraction is $f_{\rm bs}=I_{\rm bs}/I_p$, and the balance $I_p-I_{\rm bs}$ is supplied by external current drive. We evaluate 11 with the drift-kinetic code NEO [21] (gate BR-L2-A6), cross-checked against the analytic Sauter formulas (gate BR-L2-A16).

TURBULENT TRANSPORT AND THE TRIANGULARITY LEVER

Core confinement is set by drift-wave turbulence, governed by the nonlinear gyrokinetic equation for the perturbed distribution δf_s of species s,

$$\frac{\partial\delta f_s}{\partial t}+(v_\parallel\,b+v_d+v_E)\cdot\nabla\delta f_s+\dot v_\parallel\frac{\partial\delta f_s}{\partial v_\parallel}=C[\delta f_s]+S_s,\qquad v_E=\frac{c}{B}\,b\times\nabla\langle\phi\rangle,$$

closed by the gyrokinetic Poisson–Amp\`ere field equations, where v_d is the magnetic drift and C the collision operator. The turbulent heat flux of species s, $Q_s=\langle\int\mathrm{d}^3v\,\frac{1}{2}m_sv^2\,\delta f_s\,v_E\cdot\nabla\psi\rangle$, is reported in gyro-Bohm units $Q_{\rm gB}=n_eT_ec_s(\rho_s/a)^2$ (with $c_s=\sqrt{T_e/m_i}$, $\rho_s=c_s/\Omega_i$). Because drift-wave transport is stiff — the flux rises steeply once the driving gradient exceeds a critical value $a/L_{T,\rm crit}$ — the design lever is the effect of shaping on both the critical gradient and the stiffness. We quantify it through the ratio of the saturated flux between a negative- and a positive-triangularity plasma at matched gradients,

$$\mathcal{R}_s=\left.\frac{Q_s^{\delta<0}}{Q_s^{\delta>0}}\right|_{\text{matched }a/L_T},\qquad s\in\{i,e\},$$

evaluated separately at the ion (ITG+TEM) and electron (ETG) scales with nonlinear CGYRO [16,15] (gates BR-L2-A1c, REG-P13), and cross-checked against the quasilinear TGLF model [18] (gate BR-L2-A17).

NEUTRONICS AND THE TRITIUM BREEDING FUNCTIONAL

The neutron field obeys the steady linear Boltzmann transport equation

$$\Omega\cdot\nabla\phi(r,E,\Omega)+\Sigma_t(r,E)\,\phi=\int_{4\pi}\int_0^\infty\Sigma_s(r,E'\to E,\Omega'\to\Omega)\,\phi(r,E',\Omega')\,\mathrm{d}E'\,\mathrm{d}\Omega'+\frac{S_n(r)}{4\pi},$$

with a 14.06 MeV source S_n set by the neutron production of 1. The net tritium breeding ratio is the tritium produced per source neutron,

$$\mathrm{TBR}=\frac{1}{S_n}\int_V\int_0^\infty\left[\Sigma^{(n,\alpha)t}_{^6\mathrm{Li}}(r,E)+\Sigma^{(n,n'\alpha)t}_{^7\mathrm{Li}}(r,E)\right]\Phi(r,E)\,\mathrm{d}E\,\mathrm{d}V,$$

with $\Phi=\int_{4\pi}\phi\,\mathrm{d}\Omega$ the scalar flux. The exothermic $^6\mathrm{Li}(n,\alpha)t$ channel dominates at thermal energy while the threshold $^7\mathrm{Li}(n,n'\alpha)t$ channel adds tritium and a returned neutron above $\sim$2.5 MeV; a neutron multiplier (beryllium, via $(n,2n)$ above $\sim$1.85 MeV, or $\mathrm{Be_{12}Ti}$) raises the available neutron budget so that $\mathrm{TBR}>1$ is reachable despite parasitic absorption and the inboard coverage lost to the centrepost. We solve 14–15 with OpenMC [33] on ENDF/B-VIII.0 [34] (gates BR-L1-A5, BR-L2-A7), with a CAD-faithful STEP⟶DAGMC check (gate BR-L2-A22).

TRITIUM AND HELIUM-3 MASS BALANCE

The per-unit tritium burn rate follows from 1: one triton is consumed per E_{DT} of fusion, so

$$\Lambda_T=\frac{P_{\mathrm{fus}}}{E_{DT}}\,m_T\,(3.156\times10^7\,\mathrm{s})=4.77kg-T\,\mathrm{fpy}^{-1}\text{ per unit},$$

and the in-vessel tritium inventory evolves as

$$\frac{\mathrm{d}N_T}{\mathrm{d}t}=(\mathrm{TBR}-1)\,\dot N_{\mathrm{burn}}-\lambda_T N_T,\qquad \lambda_T=\frac{\ln 2}{t_{1/2}},\quad t_{1/2}=12.32\text{ yr},$$

whose steady net surplus rate is $S=(\mathrm{TBR}-1)\,\Lambda_T$ once the decay term is small against the source (Appendix ?). Helium-3 is bred not in the blanket but through the beta decay of banked surplus tritium,

$$^3\mathrm{H}\xrightarrow{\beta^-,\,t_{1/2}=12.32\text{ yr}}{}^3\mathrm{He},\qquad \frac{\mathrm{d}N_{^3\mathrm{He}}}{\mathrm{d}t}=\lambda_T N_T,$$

so the helium-3 production rate is capped by the decay of the banked inventory — a slow source, as §14 makes quantitative. The startup inventory needed to fill the plasma and processing loops before breeding overtakes burn scales with the tritium residence (particle-confinement plus processing) time τ_p; integrating 17 with a required in-vessel hold-up gives the startup values quantified in §11 (gate BR-L1-A6).

DISRUPTION DYNAMICS AND THE RUNAWAY AVALANCHE

A disruption terminates the current on the current-quench time τ_{CQ}; the induced toroidal electric field can seed and then avalanche a runaway-electron beam. The threshold is the Connor–Hastie critical field [42]

$$E_c=\frac{n_e e^3\ln\Lambda}{4\pi\varepsilon_0^2 m_e c^2},$$

and above it the runaway population grows by close-collision (secondary) avalanche [41,43],

$$\frac{\mathrm{d}I_{\mathrm{RE}}}{\mathrm{d}t}\sim\frac{I_{\mathrm{RE}}}{\tau_a},\qquad \tau_a^{-1}\propto\frac{E/E_c-1}{\ln\Lambda}\,\frac{c}{L},$$

so that the number of e-foldings scales with (E/E_c-1) times the current-decay time. Because the avalanche multiplication is exponential in the flux consumed, and the flux consumed scales with I_p, high-current devices are especially exposed. As §13 shows, at $E/E_c\approx1128$ this avalanche is severe, and the mitigation requirement is a first-class design fact [45,46,44].

§ 4

Configuration, codes, and computational methodology

CONFIGURATION

Hyperion is a compact ST: major radius $R_0 = 1.20$ m, minor radius $a = 0.48$ m, aspect ratio $A = 2.5$, elongation $\kappa = 2.0$, on-axis field $B_0 = 8.0$ T with a centrepost peak field of 16.84 T, plasma current $I_p = 9.66$ MA, and negative triangularity $\delta = -0.30$ (Fig. 2). The plasma volume implied by the shape is $V = 2\pi^2 R_0 a^2 \kappa \approx 10.9\,\mathrm{m}^3$. The configuration label is 22021; it is the frozen design point of the series. The high inboard field is delivered by a no-insulation REBCO centrepost magnet, treated in the companion conductor paper [71]; it is a distinct device from, and operates at a lower field than, the 26.49 T plug coil of the separate D–^{3}He tandem-mirror burner products [72]. The two magnets are never conflated: the breeder centrepost is a toroidal-field conductor at 16.84 T, the burner plug is a mirror coil at 26.49 T.

VERIFICATION-FIRST COMPUTATIONAL STACK

Every headline quantity is a frozen anchor carried in a declared-parameter registry, regenerated by a named script, and tagged by evidence class (Fig. 1). The physics codes and their gate serials are: the Kronos toolkit reduced-order evaluator for the zero-dimensional power and particle balance (gate BR-L1-A1); the free-boundary Grad–Shafranov solvers FreeGS and FreeGSNKE [24,25] for the equilibrium (BR-L1-A12), with EFIT-style reconstruction conventions [26]; Mercier, s–α ballooning and a reduced external-kink metric for ideal-MHD stability, with the full n=1 energy-principle solve carried as owed work (BR-L2-A19); the drift-kinetic code NEO [21] for the neoclassical bootstrap current (BR-L2-A6); nonlinear CGYRO [16,15] at ion and electron scale for turbulent transport (BR-L2-A1c, REG-P13), with TGYRO [17] for the flux-driven profile check (BR-L1-A11) and TGLF [18] for a quasilinear cross-check (BR-L2-A17); the ASCOT5 guiding-centre orbit follower [31,32] for alpha confinement (BR-L2-A20); OpenMC [33] on ENDF/B-VIII.0 [34] for neutronics, breeding, damage (NRT-dpa convention [35,36]), and activation (BR-L1-A13, BR-L1-A13b, BR-L2-A7u); a CAD-faithful DAGMC neutronics chain (BR-L2-A22); the systems-code cross-check with cfspopcon/PROCESS-class tools [30] (BR-L2-A30); and a reduced-order disruption model on a NIMROD-verified [39] equilibrium, cross-referenced to DREAM [40] for the runaway population (BR-L2-A10). Nuclear-data processing and uncertainty use NJOY/TALYS (BR-L2-A28).

THE NVIDIA GPU HPC CAMPAIGN

The computationally dominant workloads — the converged nonlinear gyrokinetic transport runs (ion-scale at the real electron-mass ratio and electron-scale ETG), the Monte-Carlo neutron transport for the breeding, damage, and centrepost-life scans, and the orbit-following alpha study — were executed on the NVIDIA GPU HPC campaign, an A100/H100/A10 accelerator fleet. Nonlinear CGYRO is a spectral Eulerian solver whose flux-tube and global runs map efficiently onto GPU accelerators; the neutronics scans are embarrassingly parallel over independent particle histories and shielding-thickness cases; the orbit-following study parallelizes over an ensemble of 10^5 markers. The reduced-order evaluators, the free-boundary equilibrium, and the neoclassical solve run on commodity CPU. A subset of the physics gates — and, for the neutronics, the breeding ratio — was independently re-run on separate hardware and on an independent nuclear-data evaluation (ENDF/B-VIII.0 versus FENDL), reproducing the frozen values within stated tolerance. The full per-gate provenance, including the byte-level record and the honest gates with their named next experiments, is the physics de-risking register[R0] [54].

NUMERICAL SCHEMES, VERIFICATION, AND CONVERGENCE

Each physics module is verified against a discretization or statistical convergence criterion before its result is frozen. The free-boundary equilibrium is solved by Picard iteration of 8 on a 129×129 (R, Z) grid with the toroidal-current density $J_\phi = Rp'(\psi) + FF'(\psi)/(\mu_0 R)$ updated from the profile functions until the normalized flux residual falls below 10^{-6}; the reported shape parameters are grid-converged (gate BR-L1-A12). The nonlinear gyrokinetic runs are advanced until the gyro-Bohm heat flux reaches a statistically stationary saturated state, with the reported flux the time average over the saturated window and the numerical resolution (radial and binormal wavenumber grids, parallel and velocity-space points) refined until the saturated flux changes by less than the run-to-run scatter (gate BR-L2-A1c). The Monte-Carlo neutronics are run to 1.2×10^8 source histories, at which the statistical standard error on the CAD-faithful breeding ratio is ± 0.004 (Fig. 20); the $1/\sqrt{N}$ convergence is explicit and the tally uncertainty is propagated into the TBR ladder. The orbit-following alpha study uses 10^5 markers, sufficient to resolve a sub-percent loss fraction with adequate counting statistics on the small lost population. Appendix ? collects the discretization parameters.

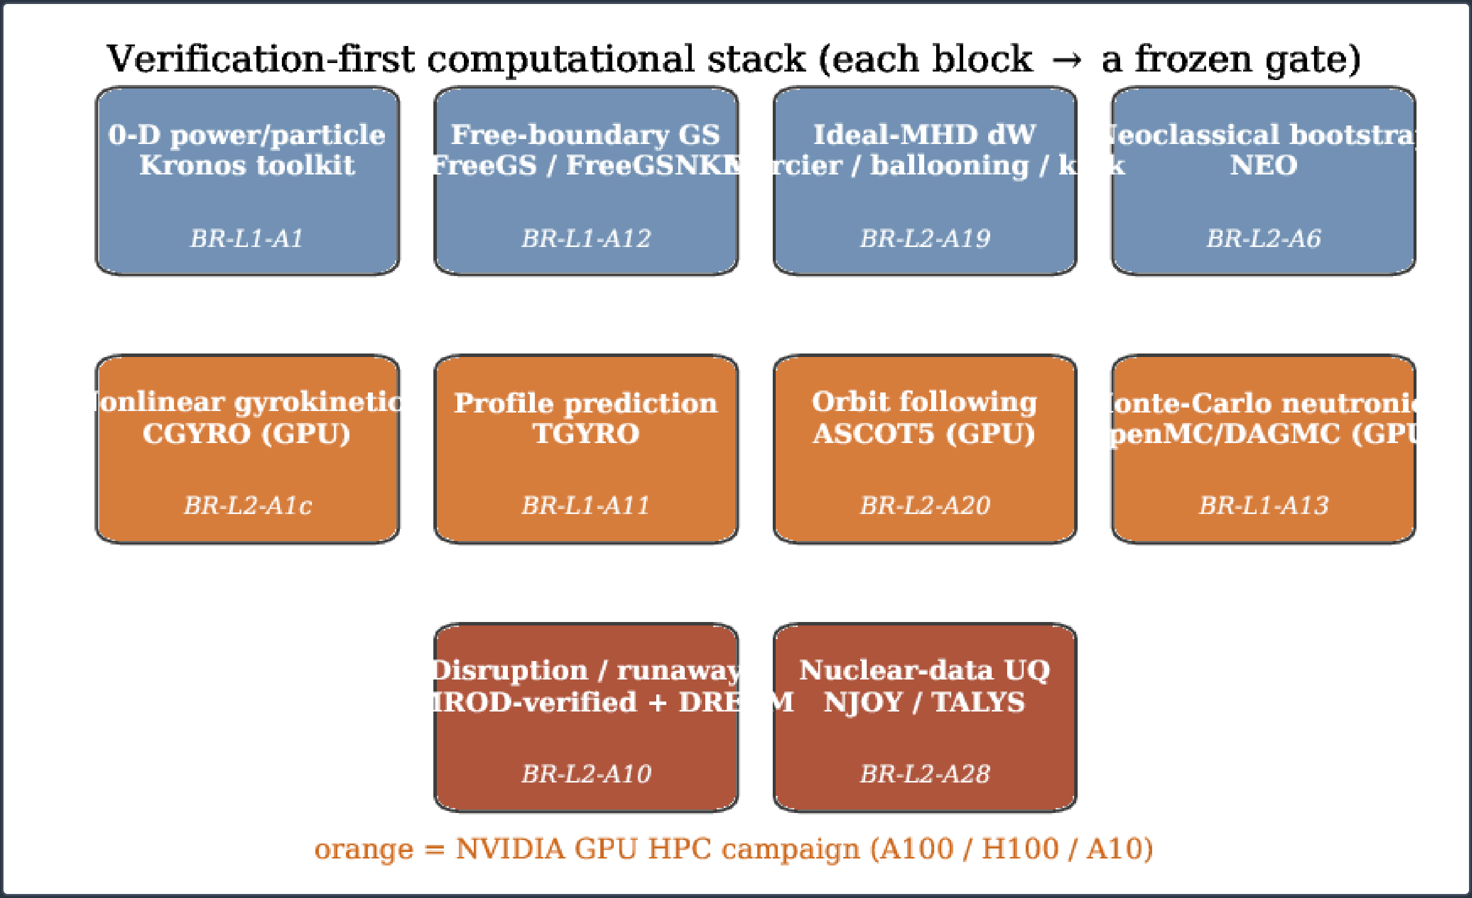

(Schematic.) The verification-first computational stack: each physics module maps to a frozen gate serial in the de-risking register[R0]. The orange blocks — nonlinear gyrokinetics (CGYRO), Monte-Carlo neutronics (OpenMC/DAGMC) and orbit-following (ASCOT5) — are the computationally dominant workloads executed on the NVIDIA GPU HPC campaign; the reduced-order, equilibrium, and neoclassical solves (blue) run on commodity CPU.

§ 5

The design point and power balance

Solving 1–7 at the frozen configuration returns $Q = 3.076$ at $P_{\text{fus}} = 85.04$ MW (Table 2, gate BR-L1-A1). Of the fusion power, $f_\alpha = 0.200$ (4) is retained as alpha heating and 0.800 leaves as 14.06 MeV neutrons (Fig. 5); the neutron power is 67.79 MW and the volumetric source strength $S_n = 3.01 \times 10^{19}\,\mathrm{n\,s^{-1}}$ (gate BR-OUT-3). The neutron wall loading is 1.885 $\mathrm{MW\,m^{-2}}$ and the first-wall damage rate $18.85\,\mathrm{dpa\,fpy^{-1}}$ (BR-L1-A13; Fig. 7, right). An independent systems-code evaluation (cfspopcon, gate BR-L2-A30) places the configuration inside the feasible operating region with a $3\times$ Greenwald-density margin ($f_{\text{GW}} = 0.30$), $P_{\text{SOL}}/P_{\text{LH}} = 1.81$ above the L–H threshold [29], and a comfortable normalized-beta margin, returning $P_{\text{fus}} = 90.8$ MW within 7% of the toolkit anchor — an independent-code consistency check on the operating point.

THE OPERATING WINDOW AND POPCON

Figure 6 (right) places the frozen point on a zero-dimensional POPCON of fusion-power contours in the density–temperature plane, computed from 1–2 with the profile form factor of §3.1. The 85.04 MW contour passes through the anchor at $\bar{n} \approx 4.0 \times 10^{20}\,\mathrm{m^{-3}}$ ($f_{\text{GW}} = 0.30$) and $T_i \approx 13$ keV, well inside the Greenwald density limit [28]. A re-optimization of the mission point (gate BR-L1-A2) shows the neutron source strength and the wall loading trade one-for-one (an invariant), with a peak accessible source strength of $5.39 \times 10^{19}\,\mathrm{n\,s^{-1}}$ at higher density; the frozen point is chosen for margin rather than peak flux. The confinement time entering the balance is anchored twice: a reduced-order evaluation gives $\tau_E = 0.358\,\mathrm{s}$ at $H_{98} = 1.0$ (gate BR-L1-A4) and a flux-driven TGYRO evaluation gives $\tau_E = 0.410\,\mathrm{s}$ at effective $H_{98} \approx 0.94$ (gate BR-L1-A11); the two agree to within the model spread (Fig. 7).

NEUTRON SOURCE STRENGTH, WALL LOADING, AND DISPLACEMENT DAMAGE

The strategic product is quantified by three derived quantities that follow directly from the power balance. The 14.06 MeV neutron production rate is one neutron per fusion reaction, so

$$S_n = \frac{P_{\text{fus}}}{E_{DT}} = \frac{85.04 \times 10^6\,\mathrm{W}}{17.59 \times 1.602 \times 10^{-13}\,\mathrm{J}} = 3.01 \times 10^{19}\,\mathrm{n\,s^{-1}},$$

with the neutron power $P_n = 0.800\,P_{\text{fus}} = 67.79$ MW (gate BR-OUT-3). Averaging the neutron power over the first-wall area A_{fw} gives the neutron wall loading

$$\Gamma_{\text{wall}} = \frac{P_n}{A_{\text{fw}}} = 1.885 \quad \mathrm{MW\,m^{-2}},$$

which is the pacing quantity for first-wall and centrepost life. The wall loading maps to a displacement-damage rate through the material's damage cross section: in the NRT convention [35] the displacements per atom per full-power year are $\dot{d} = \phi_{\text{fast}}\,\sigma_{\text{dpa}}\,t_{\text{fpy}}$, and for the frozen first-wall spectrum this evaluates to $18.85\,\mathrm{dpa\,fpy^{-1}}$ (gate BR-L1-A13), a factor ~ 2 above ITER first-wall rates and characteristic of a compact, high-fluence source. The more recent arc-dpa/rpa models [36] would revise the absolute displacement count downward by the athermal-recombination factor but leave the wall-loading and fast-fluence quantities of Eqs. 21–22 — and hence the centrepost coverage-versus-life trade of §12 — unchanged, because those depend on the fluence rather than on the displacement bookkeeping. The wall loading is deliberately not minimized: a materials breeder wants flux, and the re-optimization of gate BR-L1-A2 shows the source strength and the wall loading trading one-for-one, so the frozen point is a margin choice rather than a flux-maximizing one.

GAIN CARRIED HONESTLY

We carry the gain as a *design-point target*: no magnetic device has exceeded $Q \approx 0.7$ in a controlled D–T experiment (Fig. 6) [14], and the value is reached at demonstrated confinement rather than demonstrated on a machine. Critically, the target is a driven, sub-ignition gain — a materially easier requirement than net electricity, and the one a breeder actually needs, since from 6 alpha heating supplies only $f_\alpha Q = 0.62$ of the external power at this point and the plasma is nowhere near an ignition runaway. The distinction matters for a materials mission: a breeder does not need to reach ignition, it needs to sustain a 14 MeV flux economically, and a sub-ignition driven point is both easier to reach and easier to control.

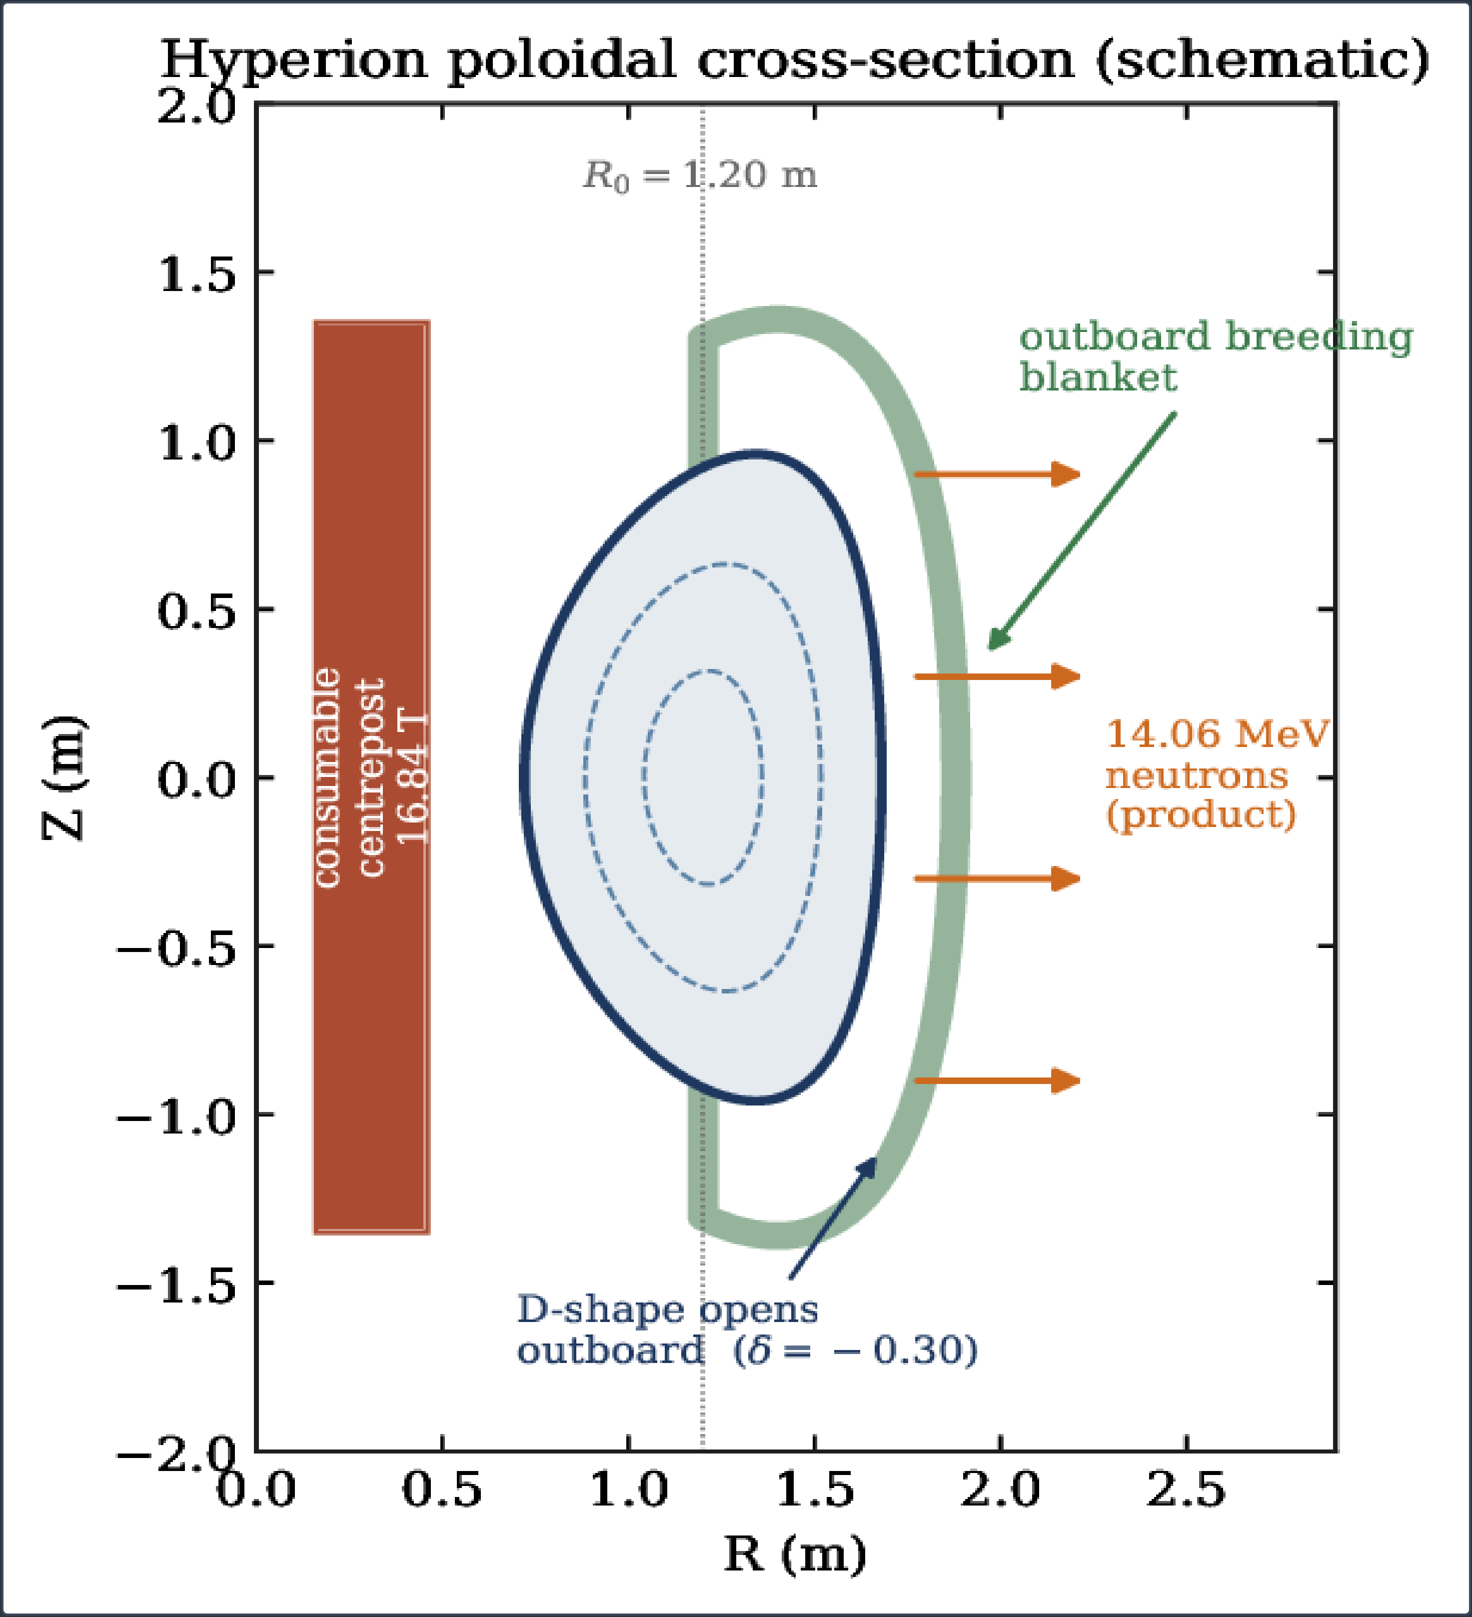

(Schematic.) Hyperion poloidal cross-section: negative-triangularity ($\delta = -0.30$) separatrix with the D-shape opening outboard, the high-field consumable centrepost (peak 16.84 T) on the inboard side, the outboard breeding blanket, and the 14.06 MeV neutron product. Geometry is to the frozen design point; no field data are plotted.

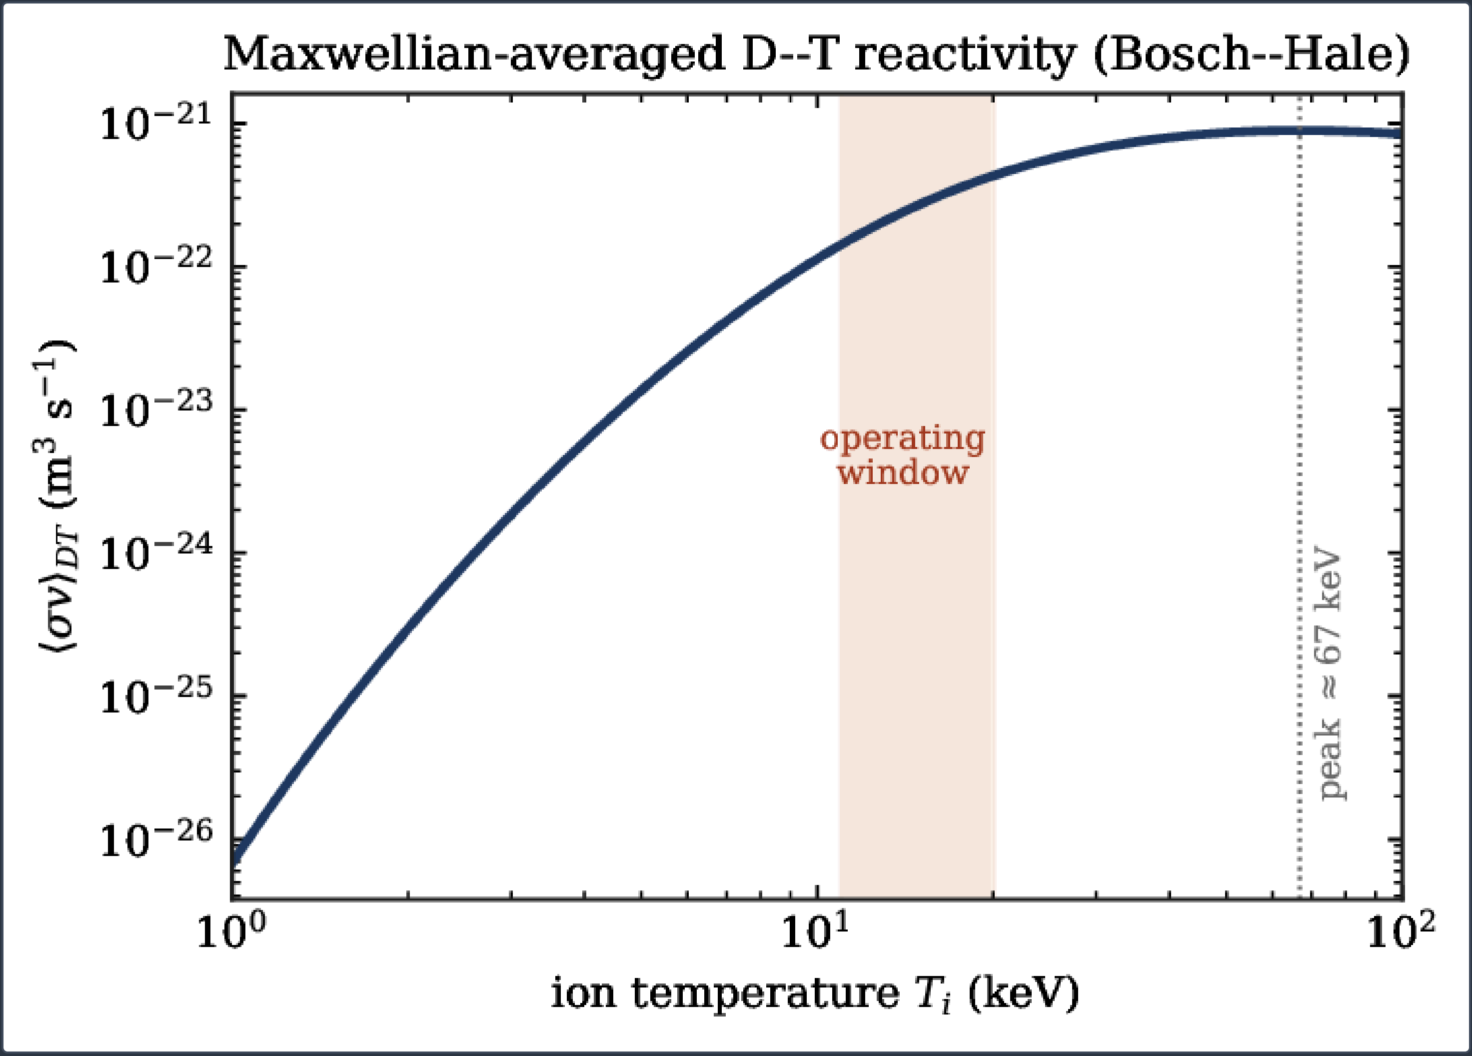

Maxwellian-averaged D–T reactivity $\langle\sigma v\rangle_{DT}(T_i)$ from the Bosch–Hale parameterization [12] of Eq. 3, peaking near **64** **keV**; the shaded band marks a representative operating window. The reactivity curve is textbook physics; the operating point of Table 2 is the frozen anchor.

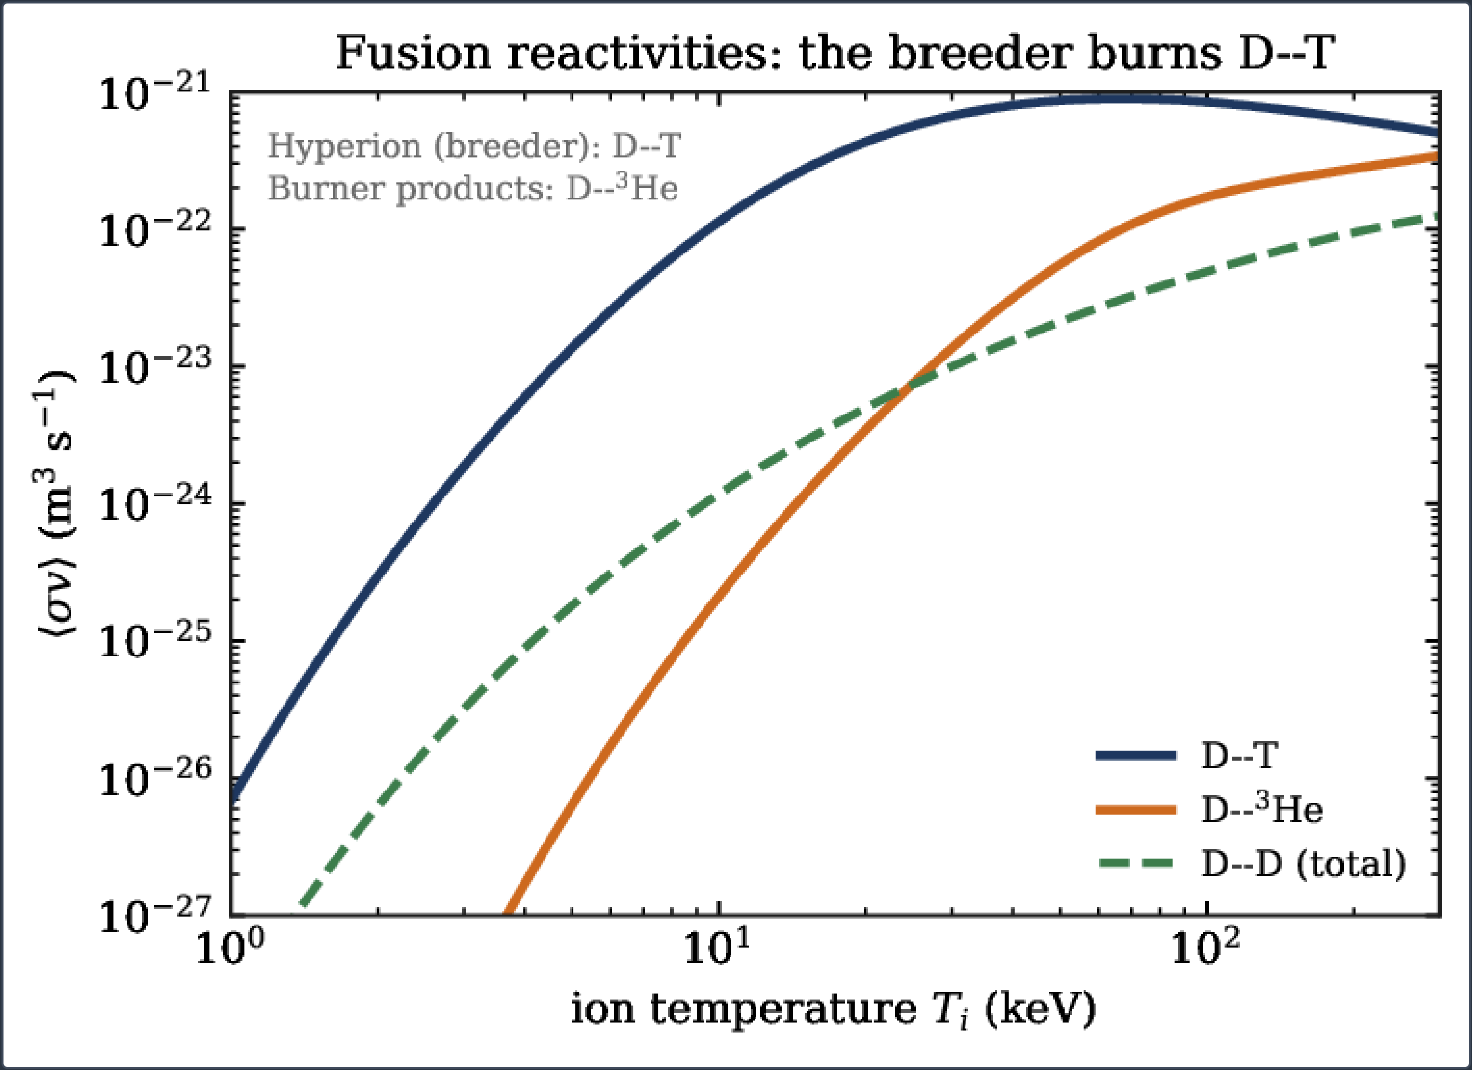

The D–T reactivity against D–^{3}He and D–D, showing why the breeder burns D–T while the companion burner products burn D–^{3}He [72,49]. Textbook physics; included to place the breeder fuel choice in context.

QUANTITY	VALUE	GATE
major / minor radius	1.20 / 0.48 m	22021
aspect ratio A / elongation κ	2.5 / 2.0 (2.47 free-boundary)	BR-L1-A12
triangularity δ	−0.30 (−0.334 free-boundary)	BR-L1-A7
plasma current I_p	9.66	BR-L1-A1
on-axis field B_0 / centrepost peak	8.0 / 16.84 T	BR-L1-A12
plasma volume V	10.9 m^3	Eq. 2
plasma gain Q	3.076	BR-L1-A1
fusion power P_{fus}	85.04	BR-L1-A1
neutron fraction / power	0.800 / 67.79	Eq. 1, BR-OUT-3
source strength S_n	3.01×10^{19} n s^{-1}	BR-OUT-3

table continued

wall loading / damage rate	1.885 MW m^{-2} / 18.85 dpa fpy^{-1}	BR-L1-A13
confinement time τ_E	0.358–0.410 s (H_{98} 0.94–1.0)	BR-L1-A4/A11
Greenwald fraction f_{GW}	0.30 (3× margin)	BR-L2-A30
normalized beta β_N	2.13	BR-L1-A12
bootstrap fraction f_{bs}	0.155	BR-L2-A6

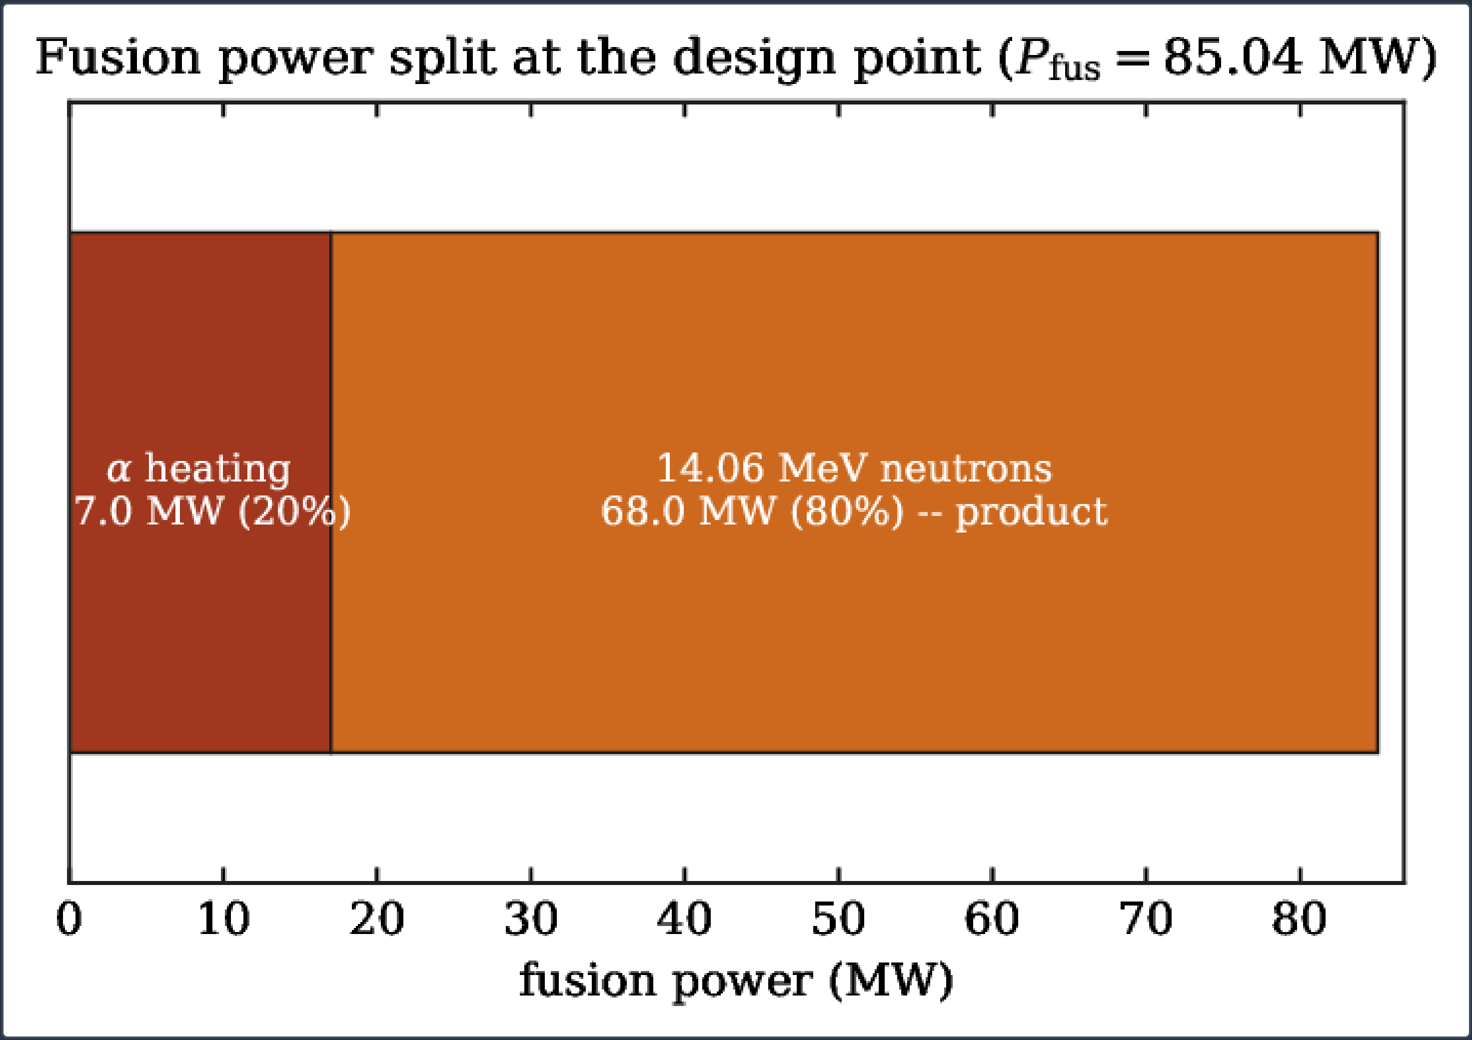

Fusion power split at the design point (Eq. 1, 4): 20% of the fusion energy is alpha heating retained in the plasma, 80% leaves as 14.06 MeV neutrons — the strategic product.

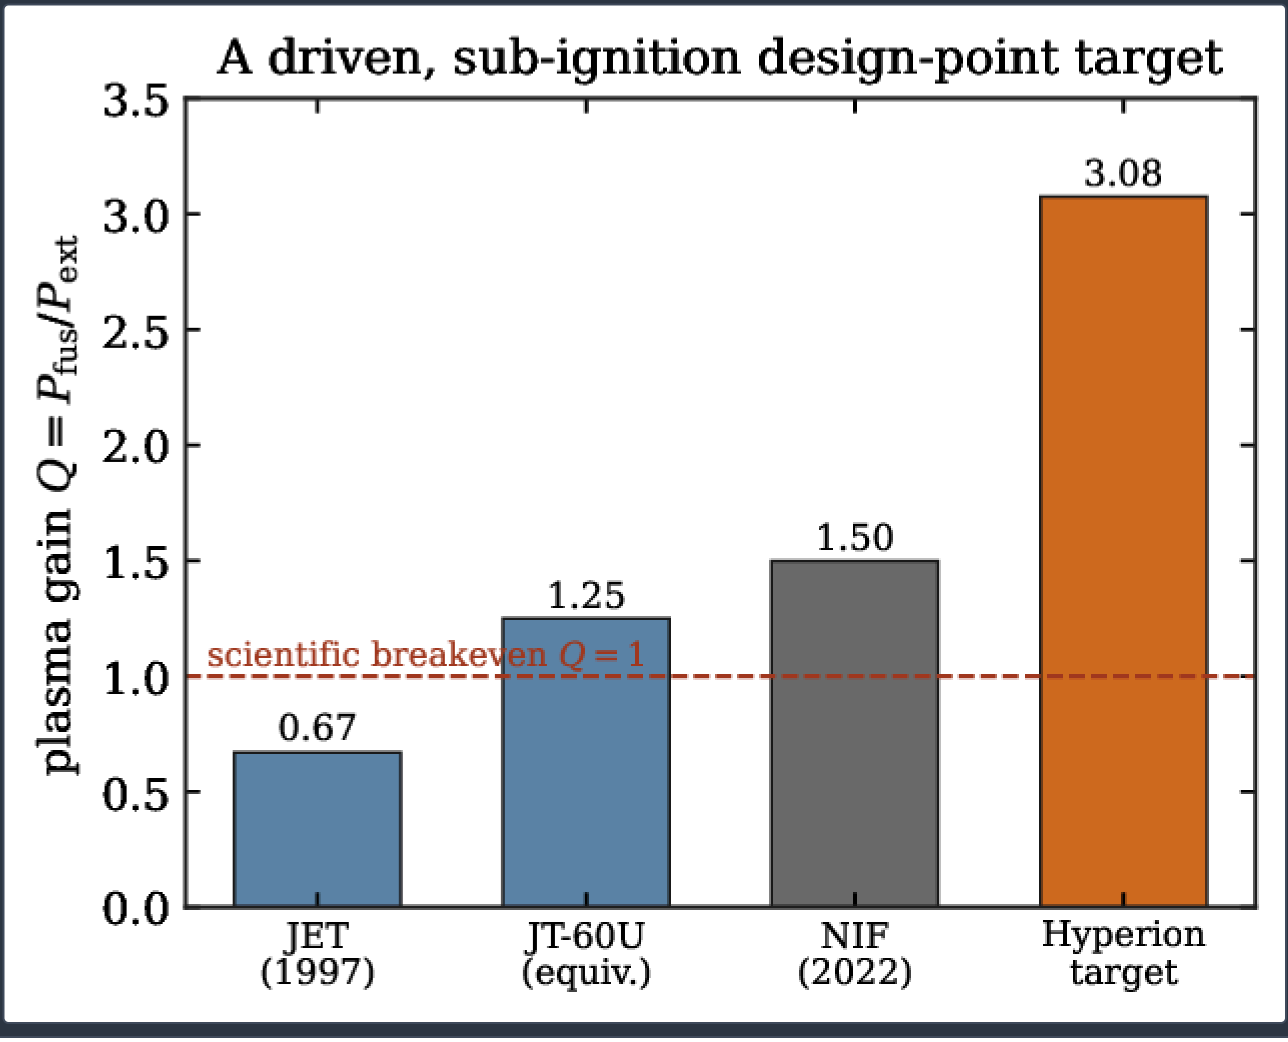
A driven, sub-ignition design-point target
plasma gain $Q = P_{fus}/P_{ext}$
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
scientific breakeven $Q = 1$
0.67
1.25
1.50
3.08
JET (1997)
JT-60U (equiv.)
NIF (2022)
Hyperion target

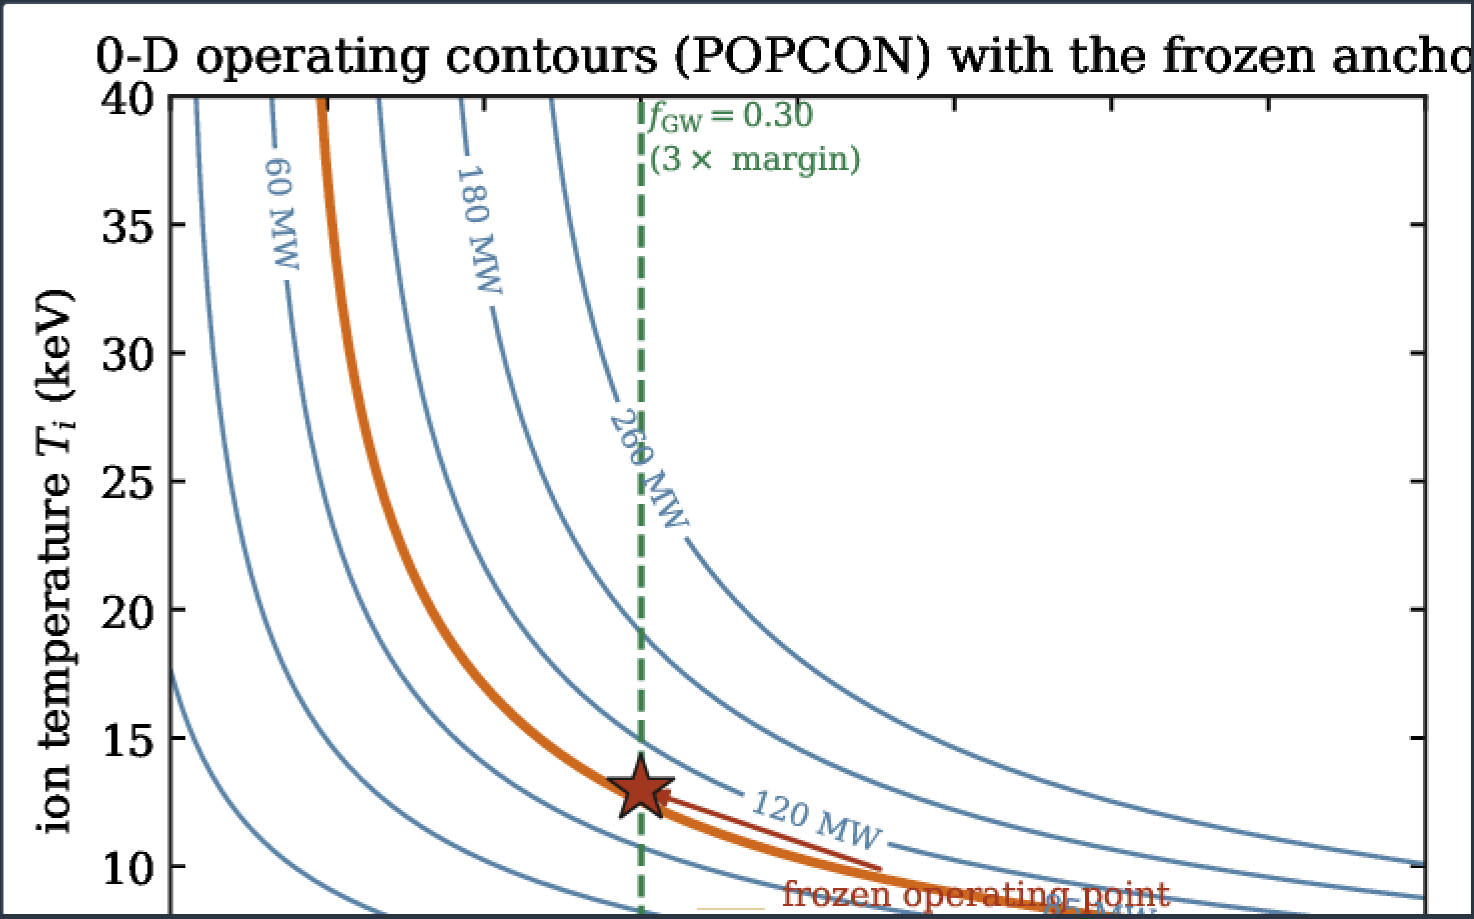
0-D operating contours (POPCON) with the frozen ancho
ion temperature T_i (keV)
40
35
30
25
20
15
10
$f_{GW} = 0.30$
(3 × margin)
60 MW
180 MW
260 MW
120 MW
frozen operating point

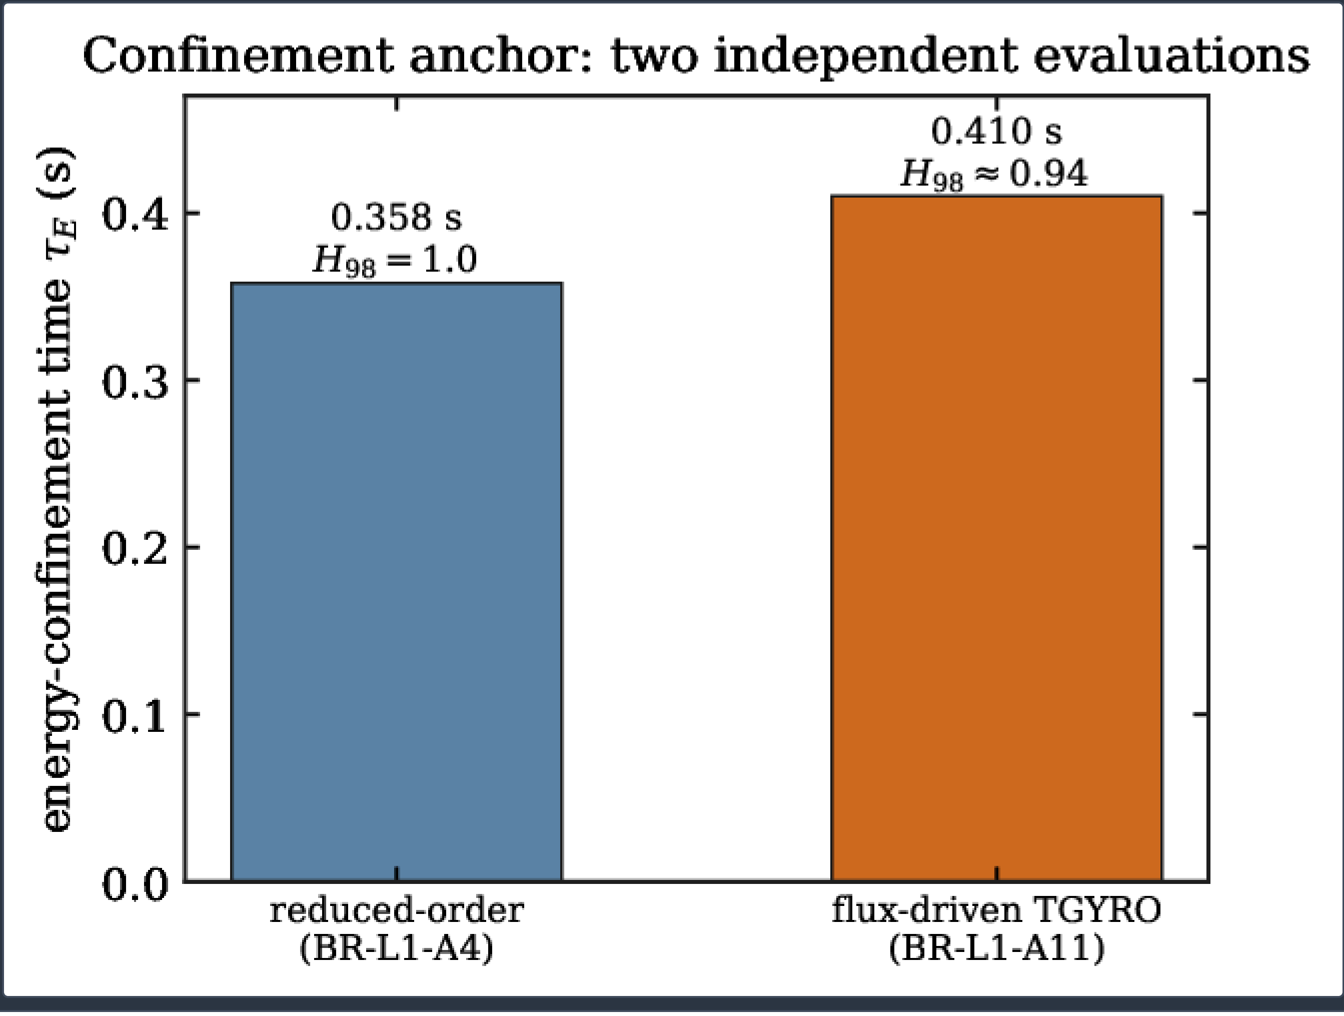
Confinement anchor: two independent evaluations
energy-confinement time τ_E (s)
0.0
0.1
0.2
0.3
0.4
0.358 s
$H_{98} = 1.0$
0.410 s
$H_{98} \approx 0.94$
reduced-order
(BR-L1-A4)
flux-driven TGYRO
(BR-L1-A11)

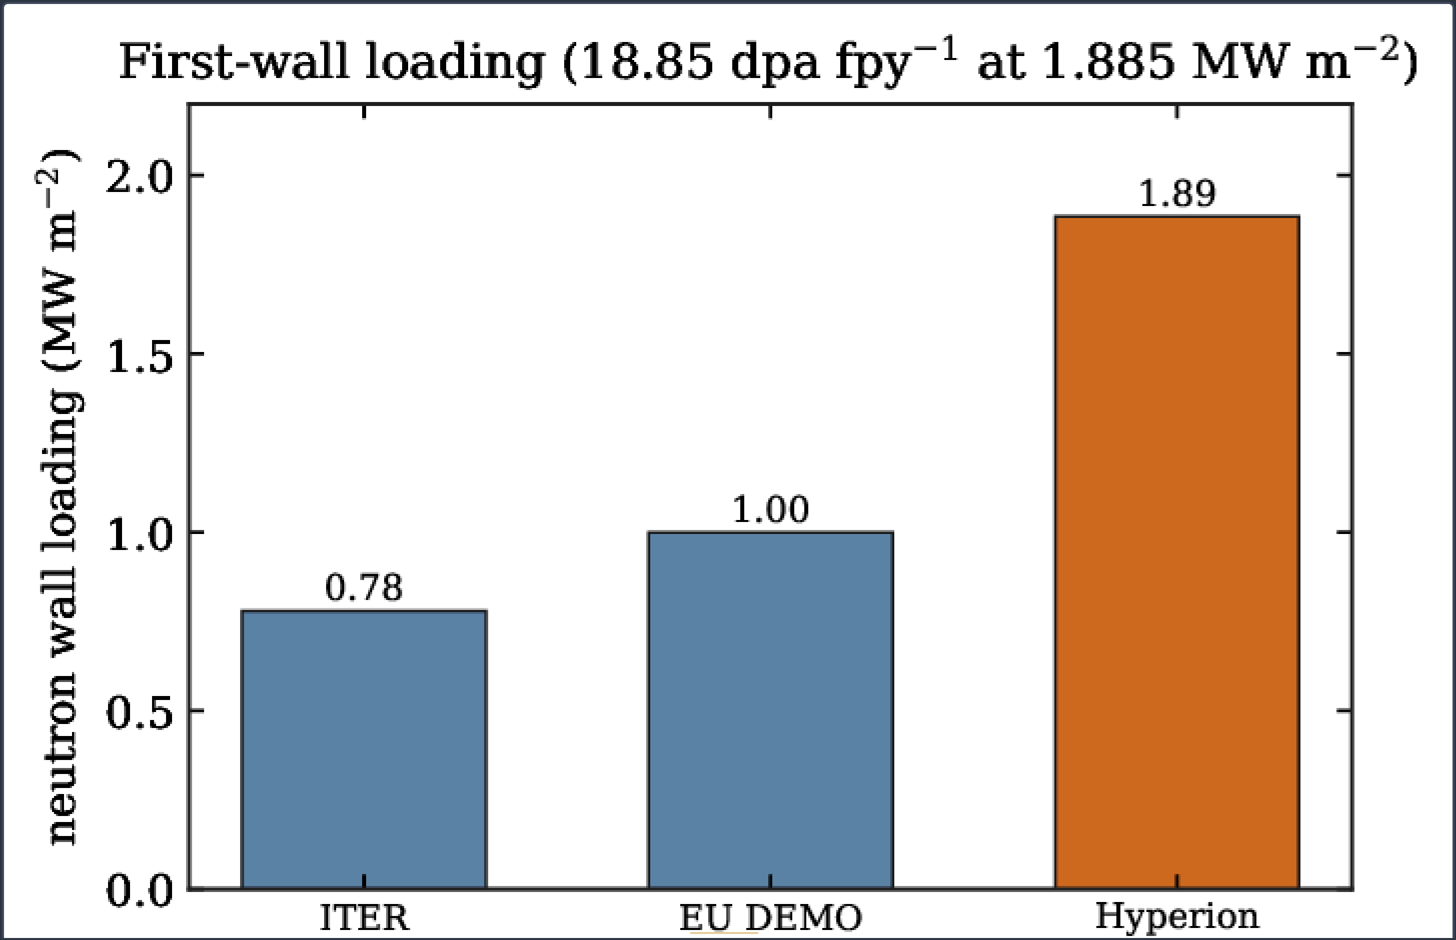
First-wall loading (18.85 dpa fpy^{-1} at 1.885 MW m^{-2})
neutron wall loading (MW m^{-2})
0.0
0.5
1.0
1.5
2.0
0.78
1.00
1.89
ITER
EU DEMO
Hyperion

§ 6

Confinement: the negative-triangularity lever

The power balance closes only if turbulent transport does not cap the pressure below the operating gradient. Negative triangularity is the design's confinement lever, and we establish it at both turbulence scales with nonlinear gyrokinetics, computing the flux ratio $\mathcal{R}_s$ of Eq. 13.

METHOD

We evolve the nonlinear gyrokinetic equations 12 with CGYRO [16] at the frozen equilibrium, carrying kinetic electrons, and compare a negative-triangularity case against a positive-triangularity twin at matched gradients. Ion-scale runs resolve the ITG and TEM branches at the real electron-mass ratio ($\mu = m_i/m_e = \mathbf{3672}$; gate BR-L2-A1c, converged and saturated); electron-scale runs resolve the ETG branch at the standard reduced mass ratio $\mu = \mathbf{400}$ (REG-P13). Transport is reported as the saturated gyro-Bohm-normalized heat flux, and the negative-to-positive ratio $\mathcal{R}_s$ is the headline. A linear stability spectrum (gate BR-L2-A18) locates the dominant branch and its growth rate before the nonlinear runs.

RESULTS: ION AND ELECTRON SCALE

At the operating gradient the converged ion-scale run returns $Q_i^{\delta>0} = \mathbf{64.2}$ and $Q_i^{\delta<0} = \mathbf{37.2}$ (gyro-Bohm units), a suppression of $\mathcal{R}_i = \mathbf{0.58}$, i.e. $\mathbf{42\%}$, squarely within the frozen $\mathbf{40–45\%}$ band (gate BR-L2-A1c; Fig. 8 left); the operating point sits below the nonlinear threshold, so the core is quiet and subcritical, the confinement-favourable regime. At the electron scale, negative triangularity suppresses ETG transport by $\mathbf{79–80\%}$ (gate REG-P13; Fig. 8 right): the electron heat flux falls by roughly a factor of five at both sampled gradients. The direction and approximate magnitude are corroborated against the TCV and DIII-D negative-triangularity campaigns [9,6,10,11] and are the subject of a dedicated gyrokinetic companion paper [55], with a two-code quasilinear TGLF cross-check reported separately.

STIFFNESS AND THE UP-SHIFTED THRESHOLD

The physical mechanism is a combination of a raised critical gradient and reduced stiffness. Figure 9 plots the schematic stiff-transport curves through the two converged points, with the critical gradient $a/L_{T,\mathrm{crit}} \approx \mathbf{2.1}$ located by the two-dimensional flux map of gate BR-L2-A1e; negative triangularity shifts the threshold outward, so at the fixed operating gradient the plasma sits lower on a shallower curve. A flux-driven TGYRO evaluation returns $\tau_E = \mathbf{0.41}\,\mathrm{s}$ ($H_{98} \approx \mathbf{0.94}$), consistent with the zero-dimensional balance (BR-L1-A11), and TGLF quasilinear transport at $\rho = \mathbf{0.6}$ reproduces the direction of the shaping effect (BR-L2-A17).

ROBUSTNESS TO COLLISIONALITY

A confinement lever is only useful if it is not an artifact of a single operating regime. Scanning the collisionality over a factor of sixteen ($\nu_\star/\nu_{\star,0} = 0.25$–$4$; gate BR-L2-A13) shows the negative-triangularity transport advantage persisting at 56–66% across the whole range (Fig. 9) — larger than the conservatively frozen 40–45% headline, confirming that the benefit is structural rather than a collisionality coincidence. We carry the smaller, saturated ion-scale value as the design number and report the scan as the robustness envelope.

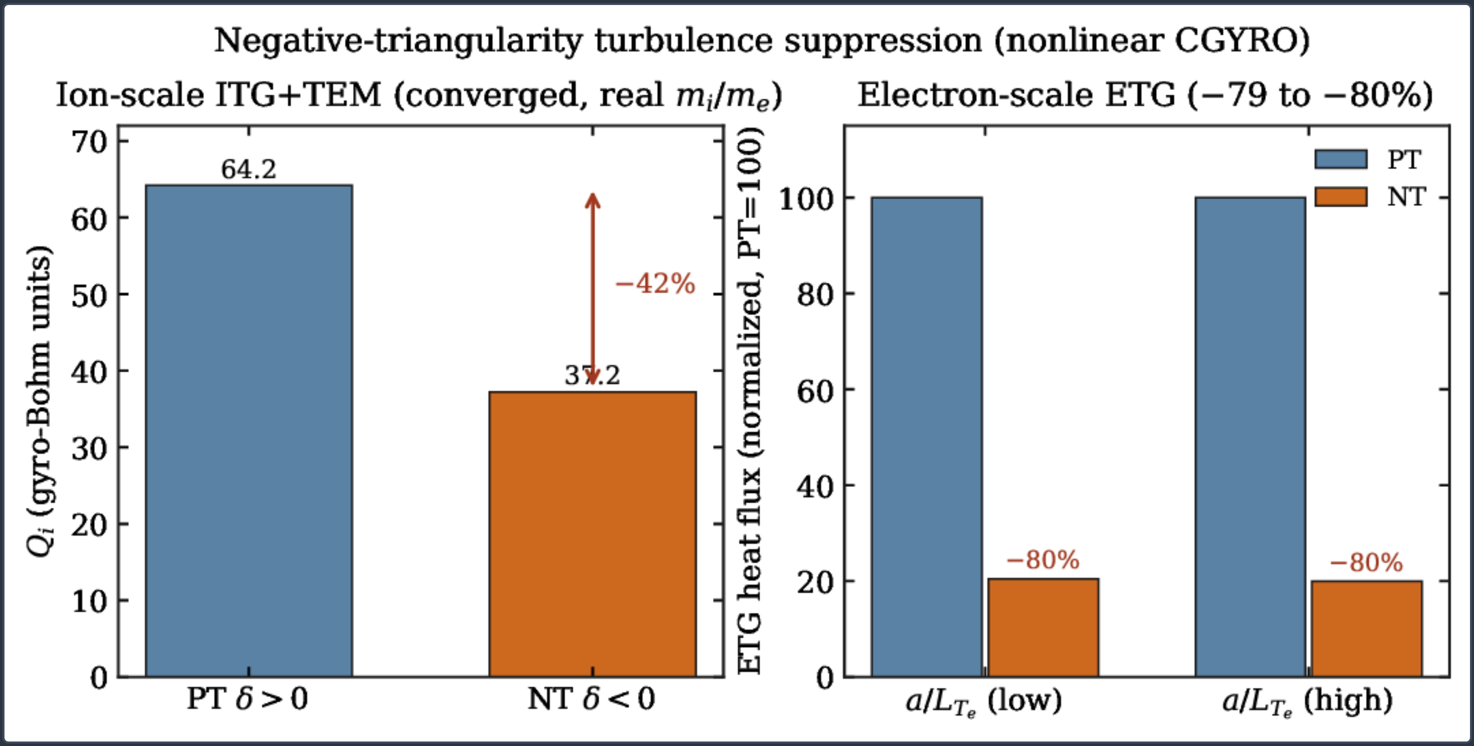

Negative-triangularity (NT) versus positive-triangularity (PT) turbulent transport. **Left:** converged ion-scale (ITG+TEM) heat flux from nonlinear CGYRO at the real mass ratio (gate BR-L2-A1c) — a 42% suppression. **Right:** electron-scale (ETG) saturated heat flux at two normalized gradients (gate REG-P13) — a 79–80% suppression. The operating point is subcritical at both scales.

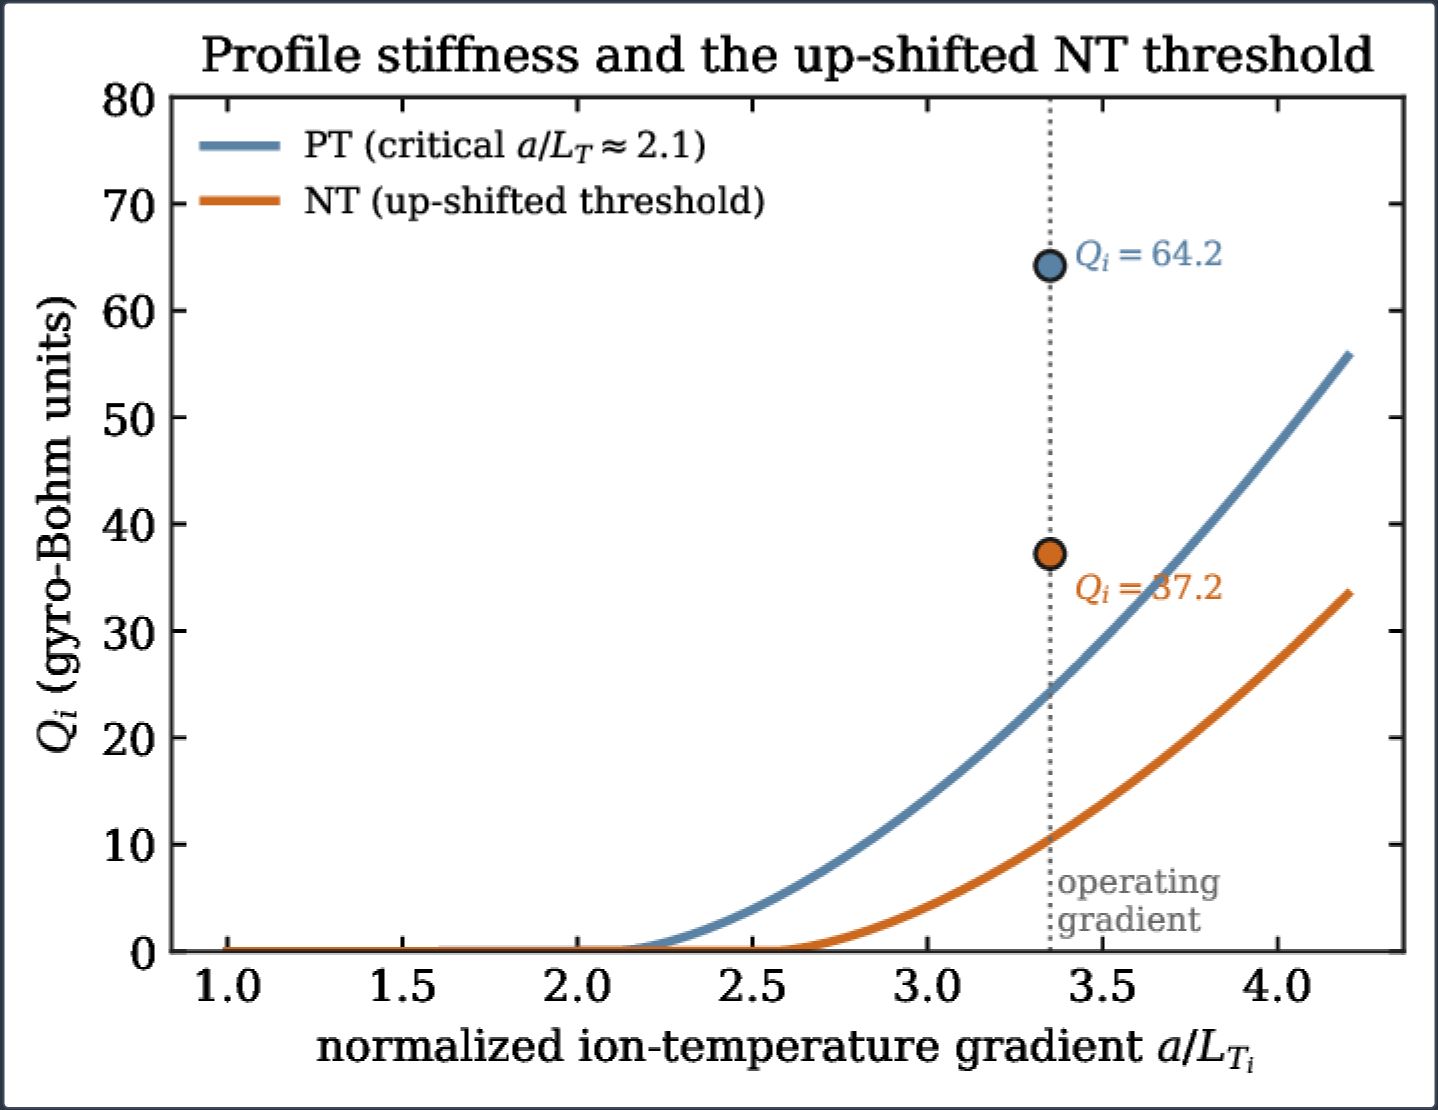
Profile stiffness and the up-shifted NT threshold
PT (critical $a/L_T \approx 2.1$)
NT (up-shifted threshold)
$Q_i = 64.2$
$Q_i = 37.2$
operating gradient
Q_i (gyro-Bohm units)
normalized ion-temperature gradient a/L_{T_i}
80
70
60
50
40
30
20
10
0
1.0
1.5
2.0
2.5
3.0
3.5
4.0

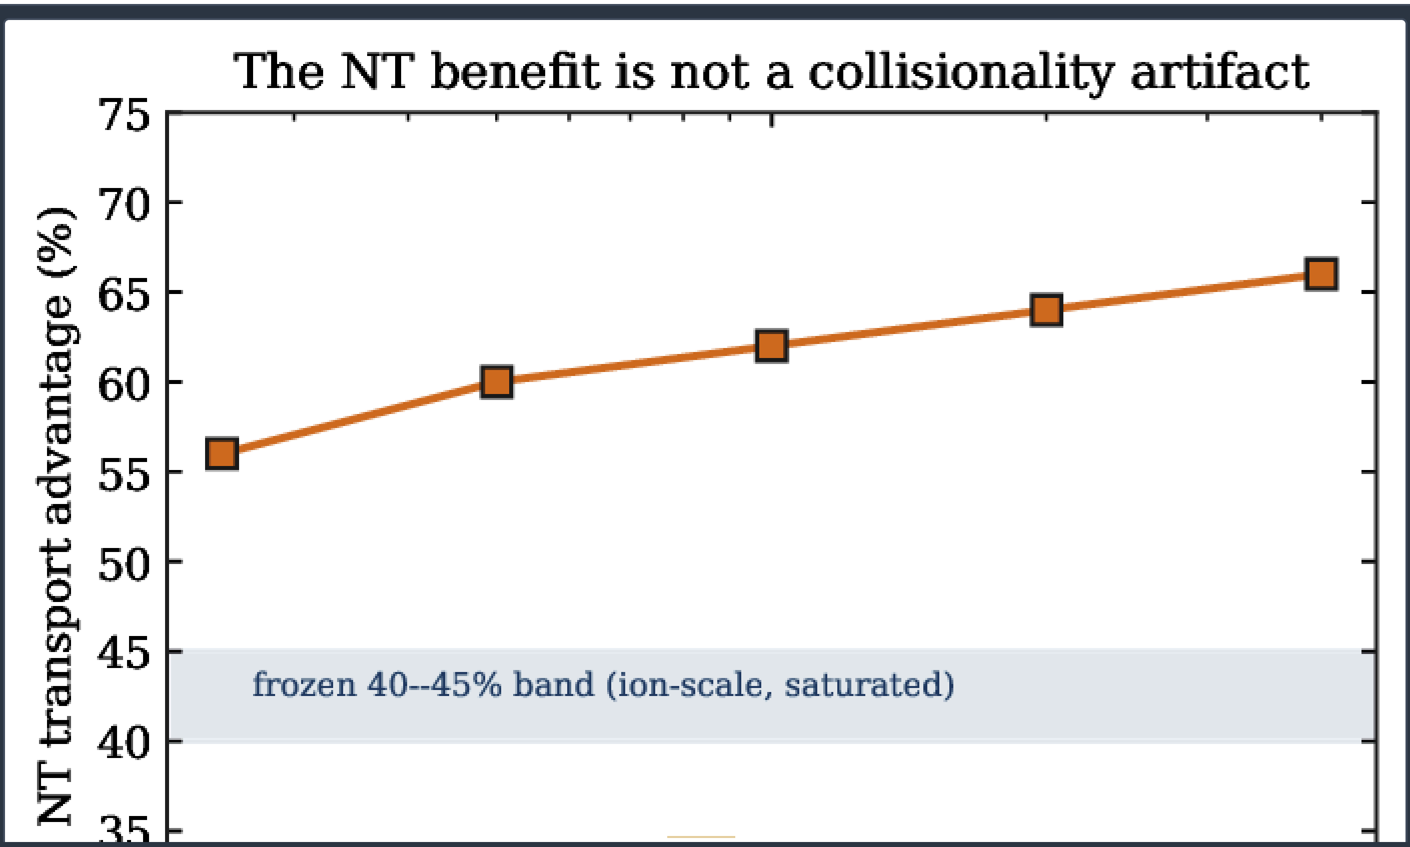
The NT benefit is not a collisionality artifact
NT transport advantage (%)
frozen 40--45% band (ion-scale, saturated)
75
70
65
60
55
50
45
40
35

§ 7

Free-boundary equilibrium and ideal-MHD stability

A free-boundary Grad–Shafranov solution of Eq. 8 [24,25] reproduces the design point on a 129×129 grid, with poloidal beta $\beta_p = 0.371$, edge safety factor $q_{95} = 2.90$, internal inductance $\ell_i = 1.40$, free-boundary triangularity $\delta = -0.334$, and normalized beta $\beta_N = 2.13$ (Eq. 9). Against the no-wall Troyon limit $\beta_N^{\max} \approx 3.5$ this is a factor 1.64 of margin (gate BR-L1-A12; Fig. 10). The equilibrium is Mercier- ($D_M > 0$, Eq. 10) and ballooning-stable on all confinement surfaces ($q > 1$; a core $q < 1$ sawtooth pocket is the only exception), and the low-n external kink carries a margin of 1.15–$1.31\times$ on the reduced metric evaluated over $\beta_N = 2.13$–2.44 against a no-wall limit of ≈ 2.8 (BR-L2-A19). Two items are flagged rather than smoothed: the free-boundary elongation reaches $\kappa = 2.47$, above the nominal 2.0, coupling to vertical stability; and the kink margin is a reduced metric pending a full resistive-wall n=1 energy-principle (δW) solve. The equilibrium and stability study is a companion paper [56], and the vertical-stability question is treated separately in the series.

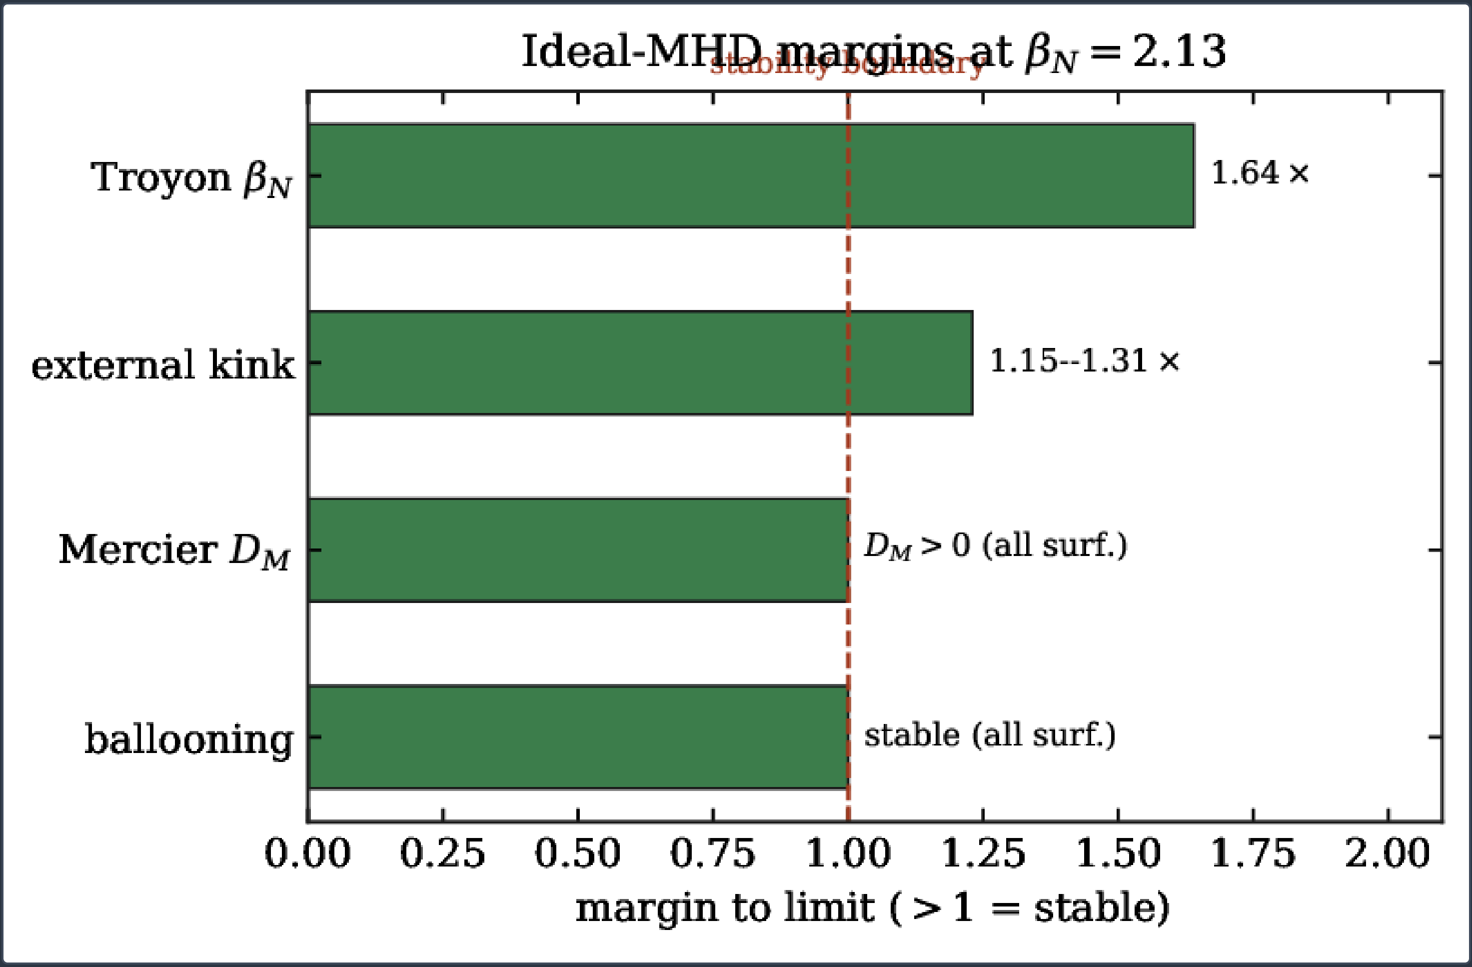

Ideal-MHD margins at $\beta_N = 2.13$ (gates BR-L1-A12, BR-L2-A19): the Troyon normalized-beta margin ($1.64\times$), the external-kink margin over the $\beta_N = 2.13$–2.44 band (1.31–$1.15\times$), and Mercier/ballooning stability on all confinement surfaces. All quantities clear the stability boundary.

§ 8

Current drive, bootstrap, and sustainment

A solenoid-free ST must sustain its current non-inductively. Evaluating the Sauter bootstrap integral 11 with NEO [21] at the design profiles returns $f_{bs} = 0.155$ (1.50 MA; gate BR-L2-A6), cross-checked against the analytic Sauter formulas (gate BR-L2-A16), so the plasma current $I_p = 9.66$ MA is 85% externally driven (8.16 MA) plus the 15.5% bootstrap (Fig. 11). The driven fraction is supplied by neutral-beam and radio-frequency current drive; the bootstrap fraction, while modest at this beta, aligns favourably with the negative-triangularity pressure profile. The steady-state current-drive study — source technology, coupling, and the path to a higher bootstrap fraction — is developed separately [67].

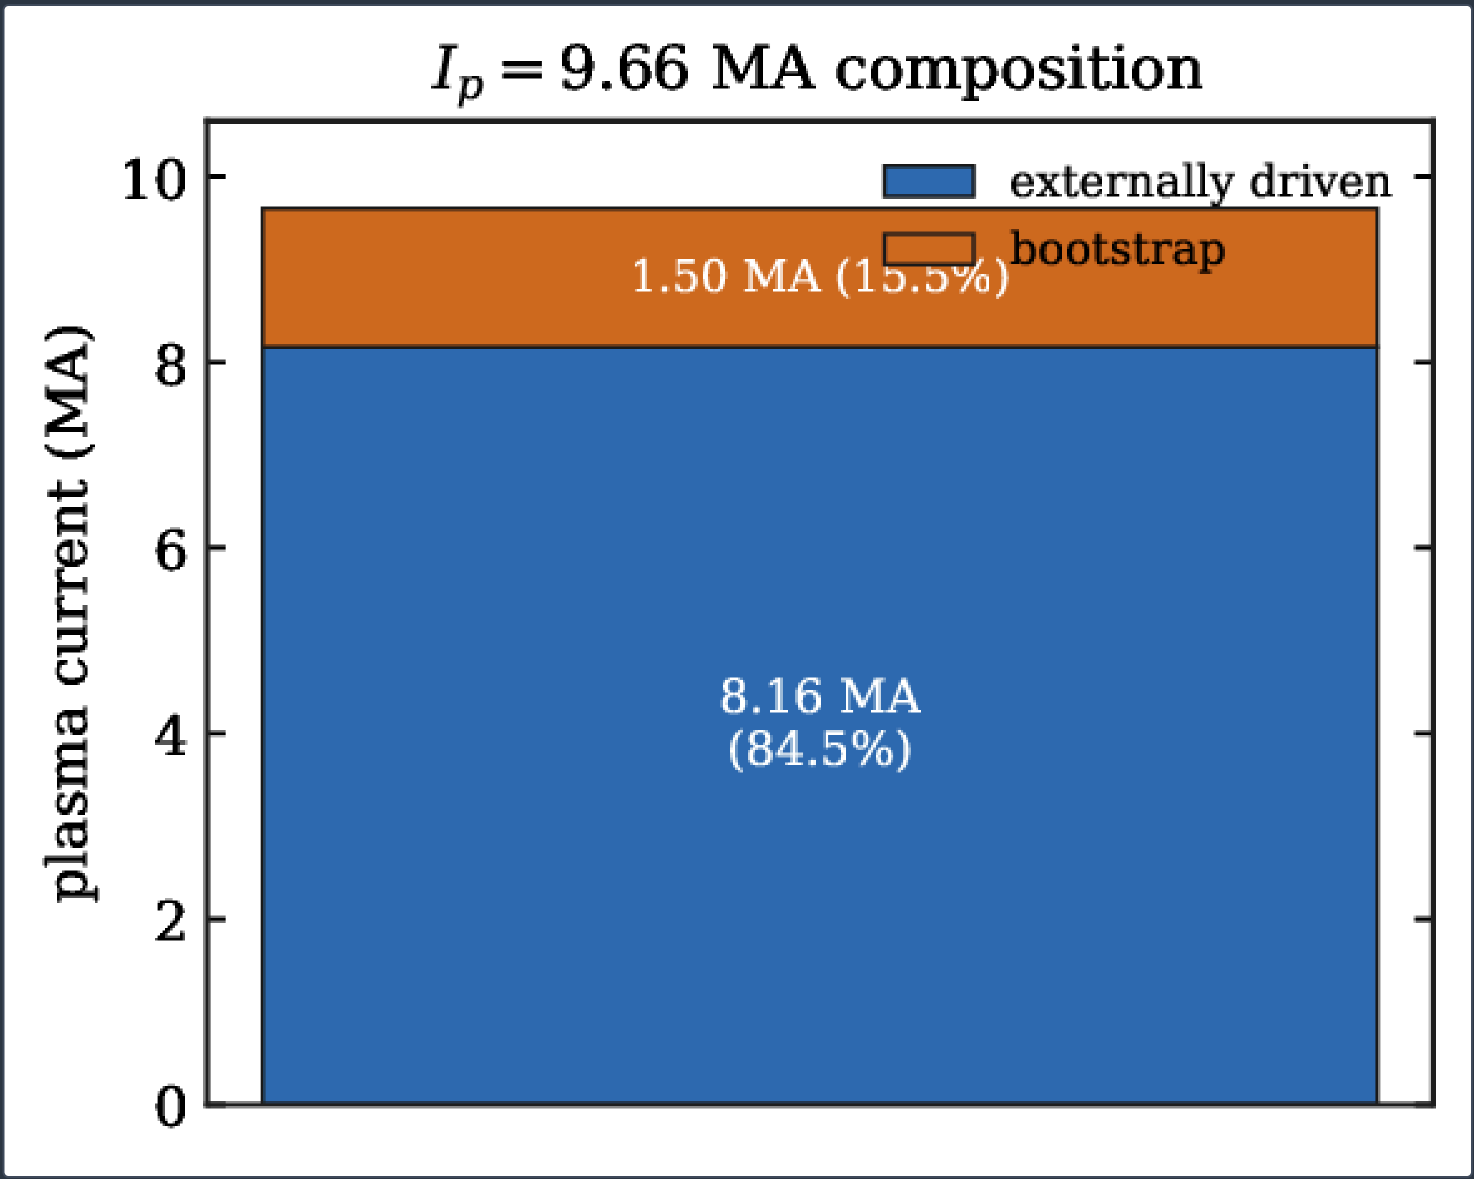

Plasma-current composition (gate BR-L2-A6): 8.16 MA externally driven plus a 1.50 MA (15.5%) bootstrap fraction from the Sauter neoclassical evaluation of Eq. 11.

§ 9

Fast-ion and alpha-particle confinement

Fusion-born alphas must deposit their energy in the plasma and must not concentrate on the first wall. A guiding-centre orbit-following calculation with ASCOT5 [31,32] over 10^5 markers returns **99.98%** confinement with a prompt loss of **0.023%** and a mean first-wall alpha load of **0.083 kW m^{-2}** (peak **1.94**), far below any engineering limit (gate BR-L2-A20). Finite toroidal-field ripple opens a second channel but keeps the loss sub-percent (Fig. 12): **0.019%** at **1.0%** ripple and **0.031%** at **1.5%**, with the peak wall load of order **1 kW m^{-2}**, and a collisional channel contributing a further **0.033%**. This closes a prior analytic estimate that had suggested a **∼4.5%** loss; the orbit-following result refutes it and agrees with an independent GPU evaluation (BR-L1-A14). Alpha confinement is not a limiting concern; the detailed study is a companion paper [57].

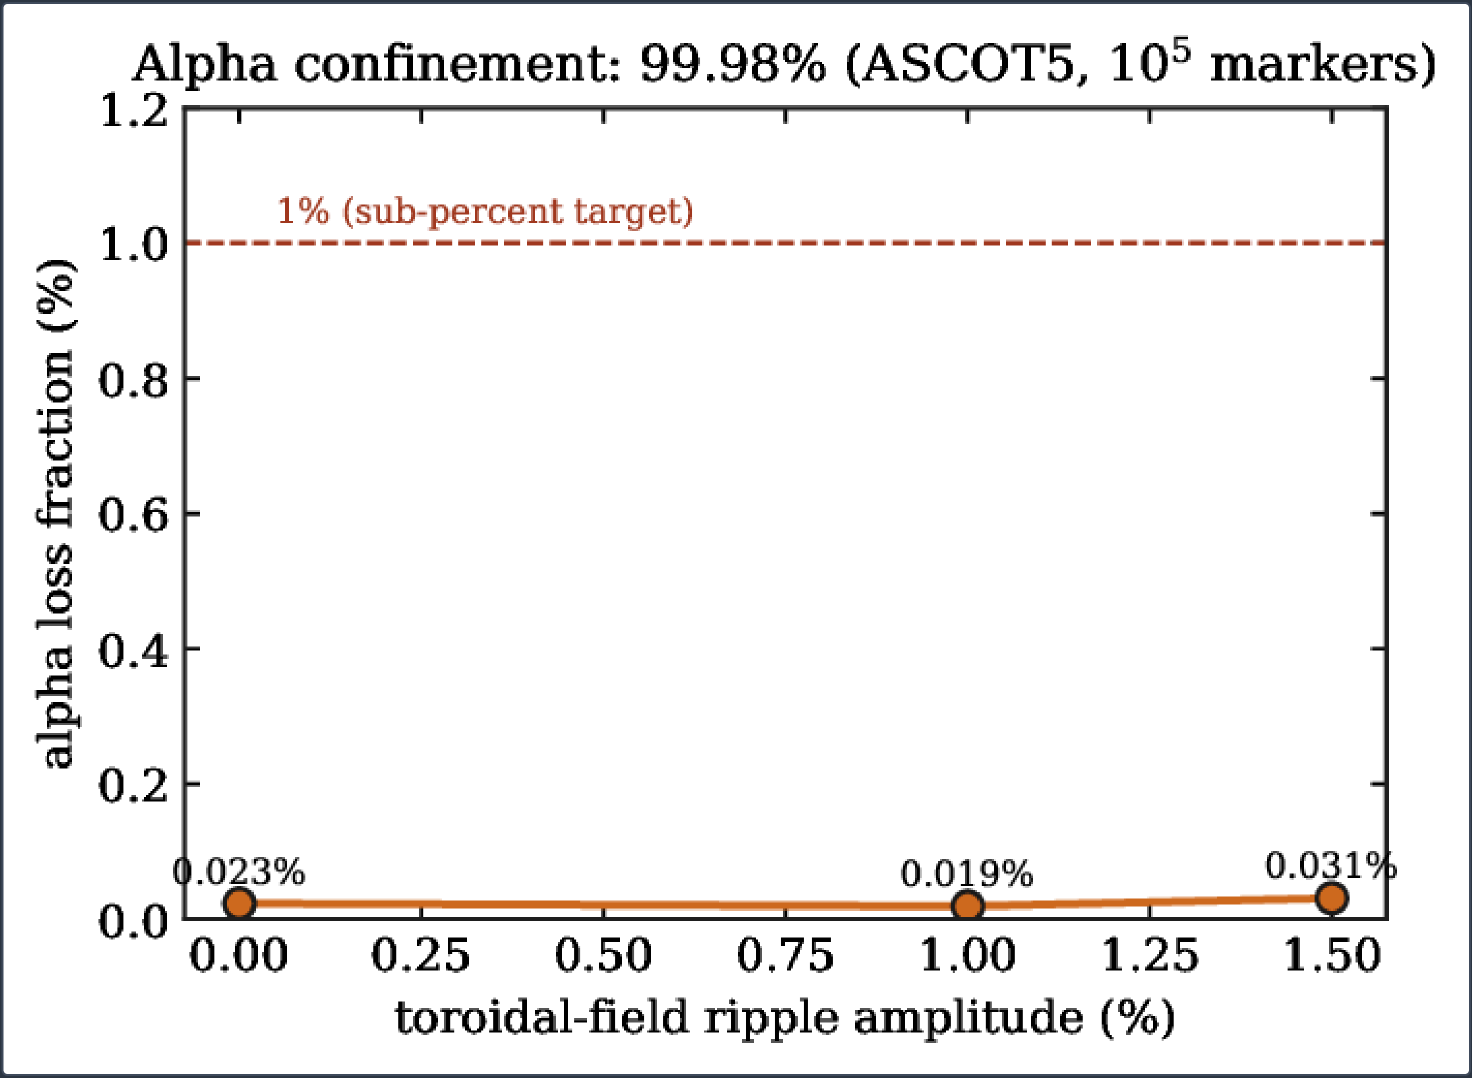

Alpha loss fraction versus toroidal-field ripple amplitude (gate BR-L2-A20, ASCOT5). Even at **1.5%** ripple the loss stays sub-percent against **99.98%** axisymmetric confinement.

§ 10

Power exhaust, the density limit, and the radiative boundary

The negative-triangularity edge removes the transient edge-localised-mode heat pulse, but the steady power crossing the separatrix must still be handled. The scrape-off-layer power is $P_{\mathrm{SOL}} = f_\alpha P_{\mathrm{fus}} + P_{\mathrm{ext}} - P_{\mathrm{rad}}$, and the exhaust challenge is conventionally posed through the figure of merit $P_{\mathrm{SOL}} B_0/(R_0 q_{95})$ against the achievable target heat flux. At the operating point $P_{\mathrm{SOL}}/P_{\mathrm{LH}} = 1.81$ (gate BR-L2-A30) places the plasma above the L–H threshold scaling [29] while remaining L-mode-edge, and the density sits at a Greenwald fraction $f_{\mathrm{GW}} = \bar{n}/n_{\mathrm{GW}} = 0.30$ with $n_{\mathrm{GW}} = I_p/(\pi a^2) = 13.3 \times 10^{20}\,\mathrm{m}^{-3}$, a factor ~ 3 below the empirical density limit [28]. This margin is deliberate: it keeps the core well away from the radiative density-limit boundary and leaves headroom for the edge-radiation solution.

Edge modelling (SOLPS-ITER and UEDGE, gates BR-L2-A11, BR-L2-A21) finds that detachment is seeding-mandatory: a one-dimensional SOL scan with impurity seeding returns target electron temperatures of $14.4/5.3/1.05\,\mathrm{eV}$ across the seeding sequence, with an upstream-to-target density ratio $n_t/n_u \approx 31$ — consistent with a partially detached target once an argon-class seed is added. The impurity-radiation budget (gate BR-L2-A29) supplies the cooling: the tungsten and argon cooling factors of Fig. 21 give a tungsten tolerance of $\sim 0.19\ \mathrm{MW\,ppm}^{-1}$ and a bremsstrahlung floor of $\sim 15\%$ of the radiated power, so a modest seed fraction radiates the required boundary power without quenching the core. The design consequence is that the divertor is a mandatory, stacked-lever detachment solution — super-X, double-null, and snowflake geometries each reduce the required seed fraction — rather than an optional refinement; the exhaust portfolio and its two-dimensional arbiter are companion work [68]. The exhaust question is therefore not a show-stopper but a firmly bounded engineering requirement, and it is the counterpart to the ELM-free-edge asset that negative triangularity buys.

§ 11

Blanket, tritium, and the breeding ladder

Tritium breeding is a design lever along a confirmed ladder, not a fixed number. The per-unit burn rate is $\Lambda_T = 4.77kg - T\,\mathrm{fpy}^{-1}$ from Eq. 16 (gate BR-OUT-2), and the net surplus per unit is $S = (\mathrm{TBR} - 1)\,\Lambda_T$ from the mass balance 17.

THE BREEDING LADDER ACROSS BLANKET CONFIGURATIONS

Evaluating the breeding functional 15 with OpenMC on ENDF/B-VIII.0 across blanket configurations (gates BR-L2-A7, BR-L2-A11h) traces the ladder of Figs. 13–15. The literal frozen recipe (beryllium packing $9.6–11.8\,\mathrm{vol\%}$) gives net $\mathrm{TBR} = 1.026$ (Be$_{0.65}$) to 1.062 (Be $_{0.80}$); adding a dedicated 5 cm solid-beryllium multiplier raises it to $\mathrm{TBR} = 1.1775–1.2039$; thickening the breeder to 65 and 80 cm reaches 1.2407 and 1.2545; and the advanced blanket — a void-reduced breeder with a solid-beryllium multiplier and a steel reflector — reaches net $\mathrm{TBR} = 1.42$. The 1 kg-T yr^{-1} export bar corresponds to $\mathrm{TBR} = 1.2097$ at 85 MW, so every rung above the 65 cm breeder is already an exporter. This $\mathrm{TBR} = 1.42$ is the confirmed per-unit ceiling across the surveyed blanket configurations: a lithium–lead alternative does not beat it, and the ST centrepost caps the inboard coverage that sets the ceiling. The CAD-faithful DAGMC check returns $(n, Xt) = 1.014 \pm 0.004$ (gate BR-L2-A22), and the cross-library check (ENDF/B-VIII.0 versus FENDL) reproduces the breeding ratio within tolerance and is the load-bearing verification of the export claim.

THE SURPLUS LAW AND STARTUP INVENTORY

The mass balance 17 gives the surplus as a linear function of TBR at fixed burn (Fig. 14): the literal blanket yields $S = 0.29$, the solid-beryllium rung $S = 0.9$, and the advanced rung $S \approx 2.0kg - T\,\mathrm{yr}^{-1}$, with the two register anchors $S = 1.6215$ (TBR 1.34) and $S = 3.8152$ (TBR 1.80) falling on the same line (gates BR-L1-A6, BR-OUT-1). The startup inventory needed before breeding overtakes burn scales with the tritium residence time: 2.38, 4.77, and 9.54 kg for $\tau_p = 0.5$, 1, and 2 yr (Fig. 14, right; gate BR-L1-A6) — a fraction of a power plant's standing inventory, which is the inventory-obstacle result of §2 made quantitative.

ONE UNIT EXPORTS; NATIONAL SCALE IS A FLEET PROPERTY

A single advanced unit is therefore a tritium exporter at the $\sim 2.0 kg - T\,\mathrm{yr}^{-1}$ scale (Fig. 15, left). A national-scale $4kg - T\,\mathrm{yr}^{-1}$ supply is explicitly a *fleet* number (3.6–$4.5kg - T\,\mathrm{yr}^{-1}$ from four to five units, or two advanced units), never a single-unit claim; the frozen fleet law returns $T = 3.68kg - T\,\mathrm{yr}^{-1}$ at the intermediate rung (gate BR-OUT-1). The blanket and fuel-cycle study, the beryllium-multiplier down-select, and the fleet path to self-sufficiency are companion papers [58,59,60].

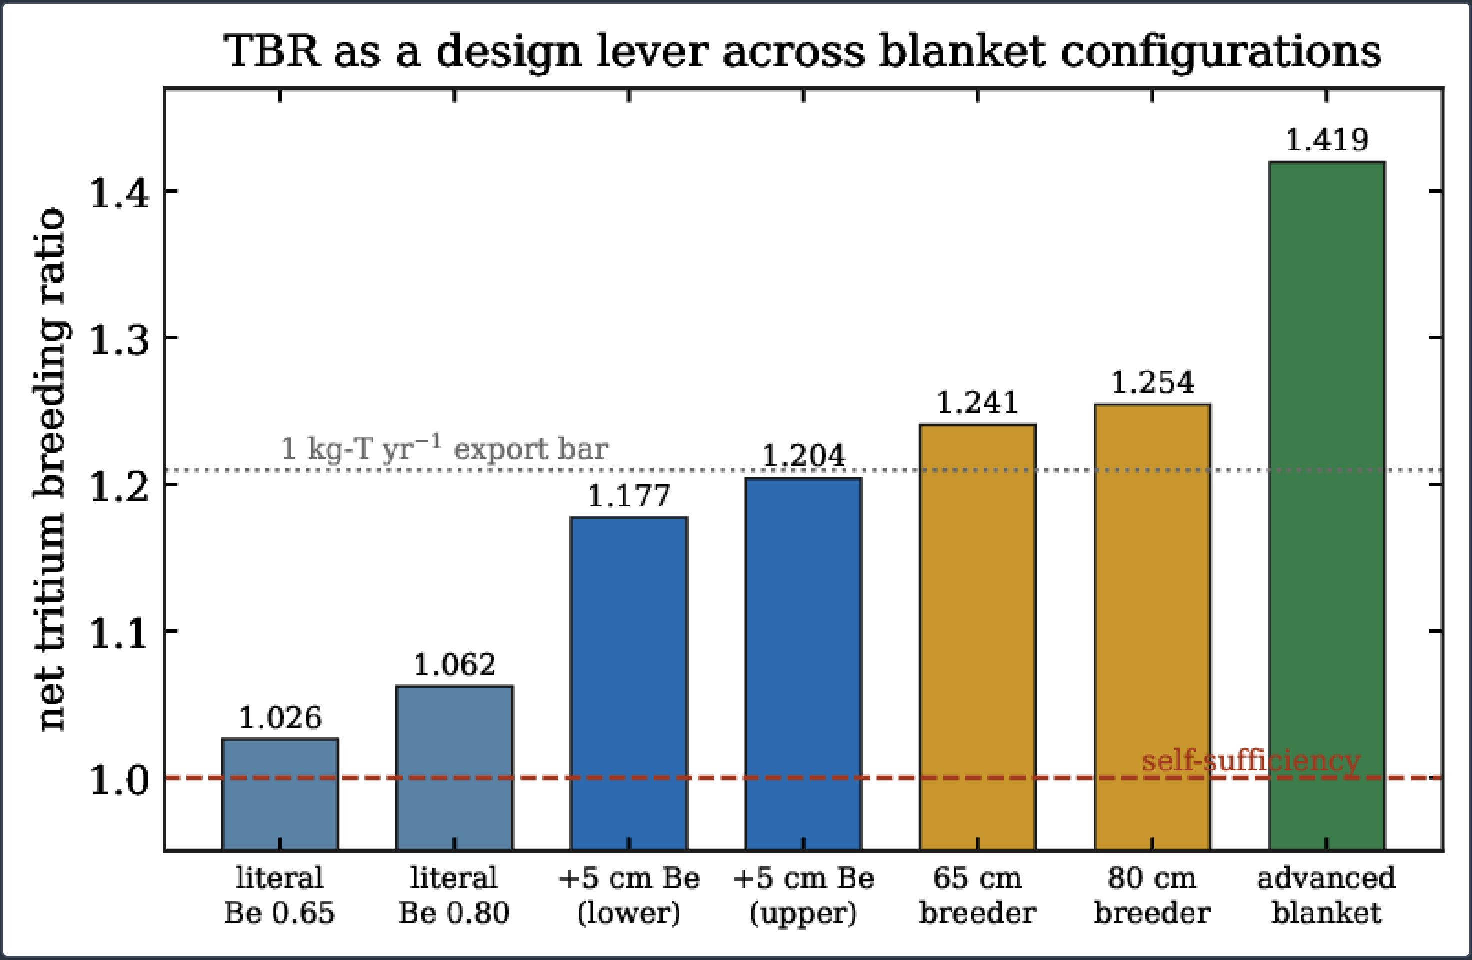

Net tritium breeding ratio across blanket configurations (gates BR-L2-A7, BR-L2-A11h): from the literal frozen recipe (1.026–1.062) through a 5 cm solid-beryllium multiplier (1.18–1.20) and thicker breeder zones (1.24–1.25) to the advanced blanket (1.42). The 1 kg-T yr^{-1} export bar and the self-sufficiency line are marked.

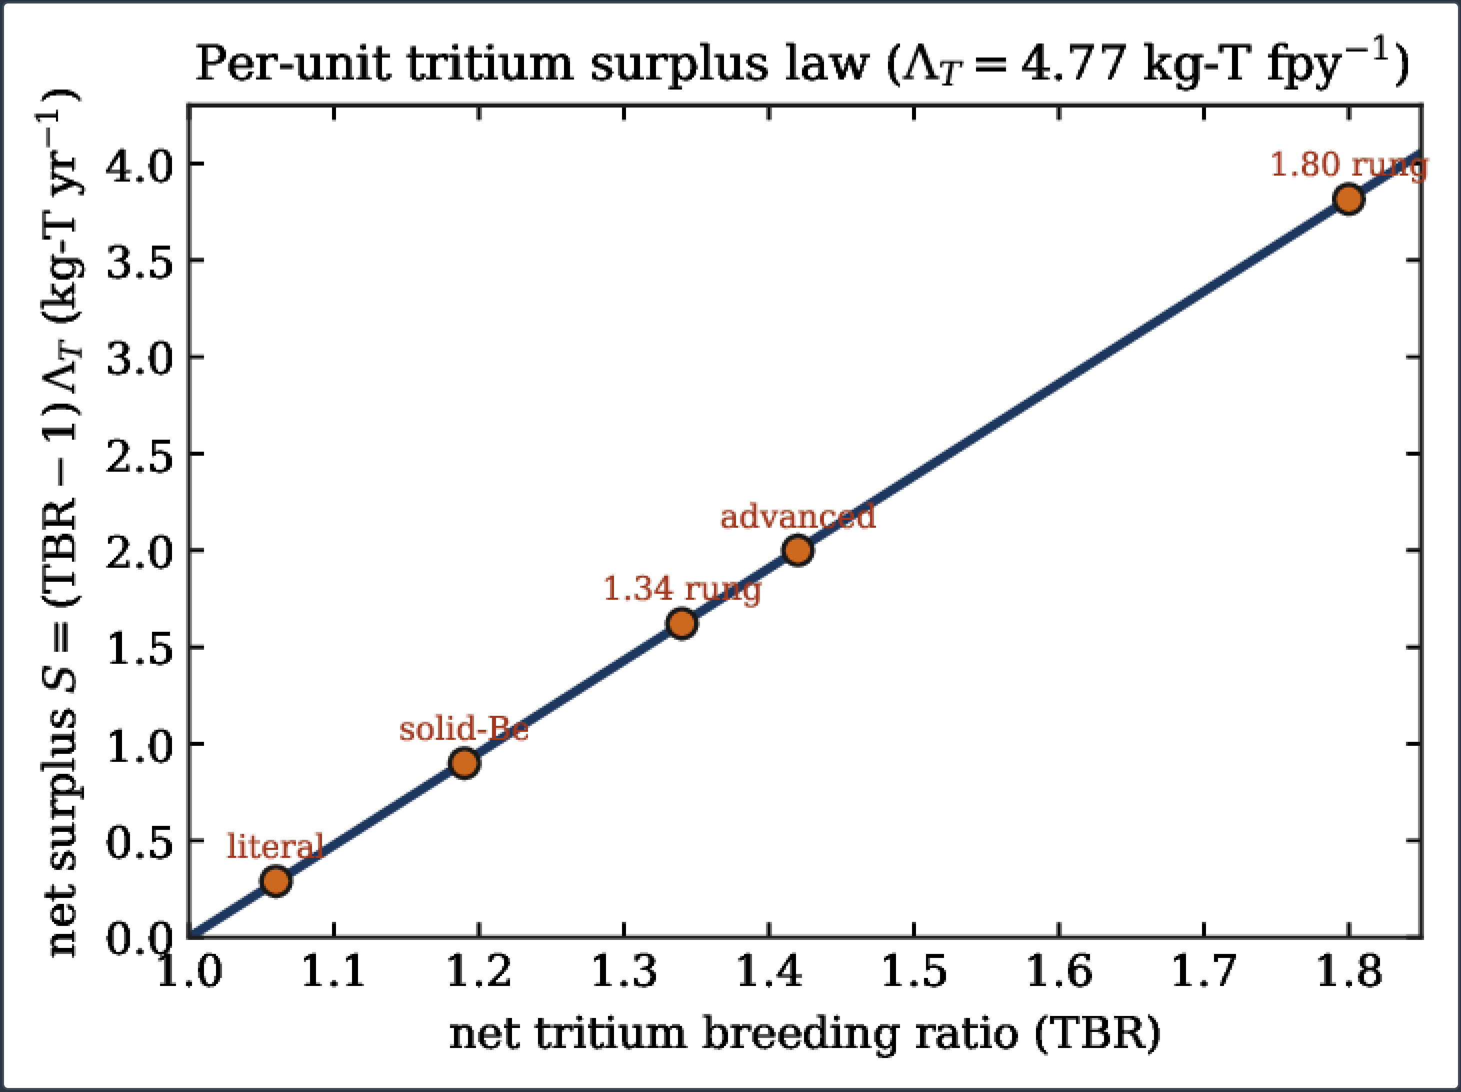
Per-unit tritium surplus law ($\Lambda_T = 4.77$ kg-T fpy^{-1})
net surplus $S = (\mathrm{TBR} - 1)\,\Lambda_T$ (kg-T yr^{-1})
net tritium breeding ratio (TBR)
literal
solid-Be
1.34 rung
advanced
1.80 rung
0.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
4.0
1.0
1.1
1.2
1.3
1.4
1.5
1.6
1.7
1.8

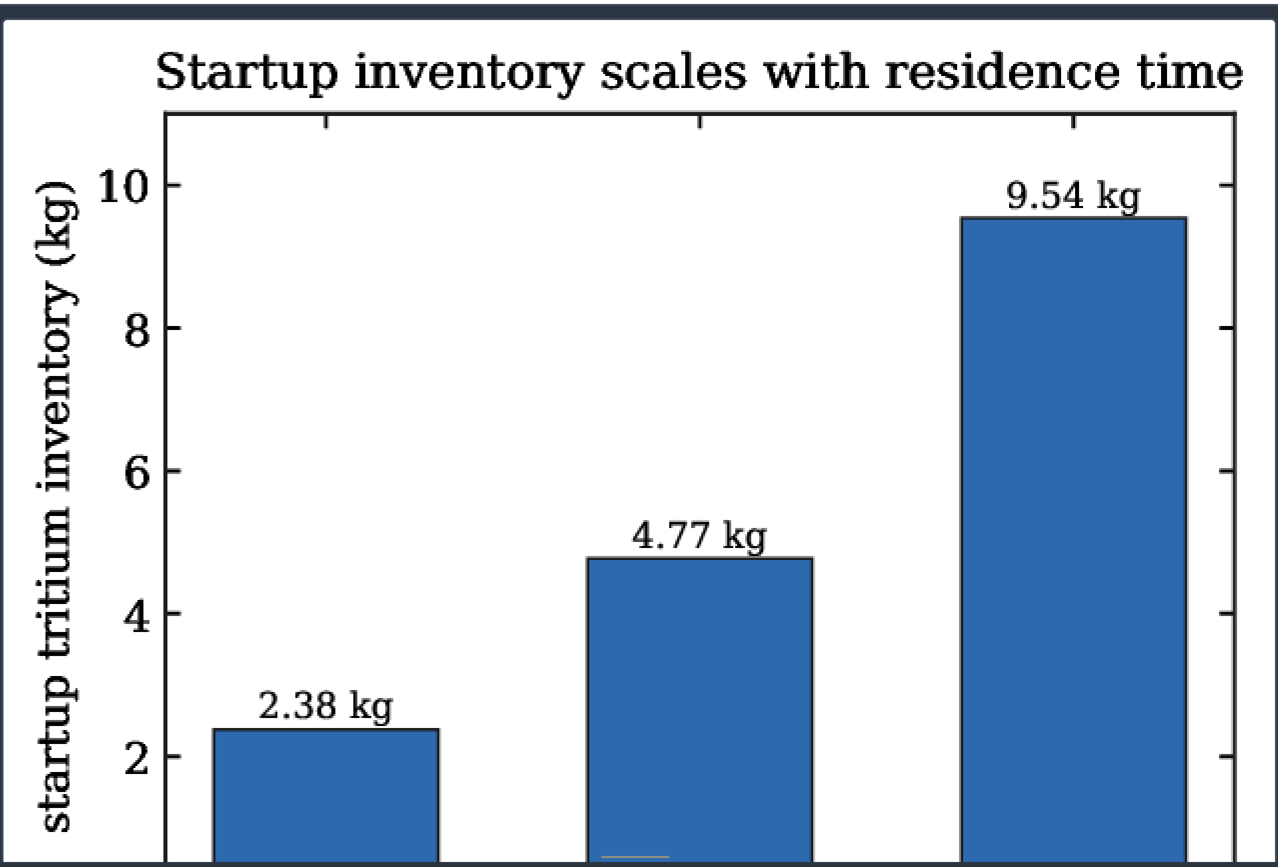
Startup inventory scales with residence time
startup tritium inventory (kg)
2.38 kg
4.77 kg
9.54 kg
2
4
6
8
10

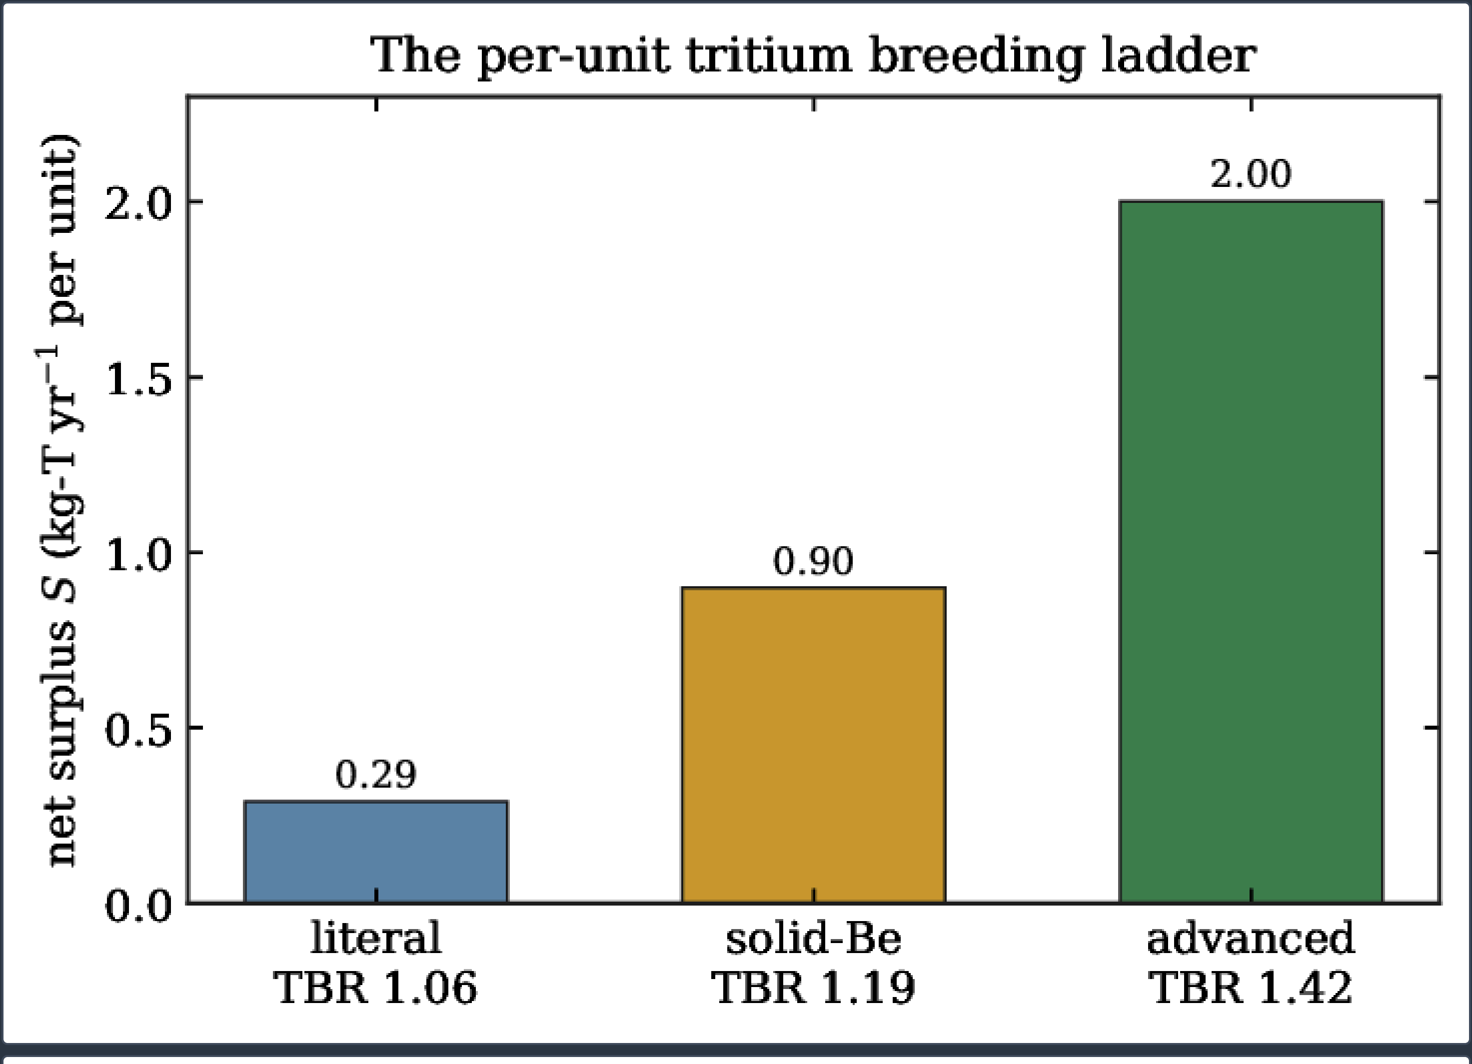

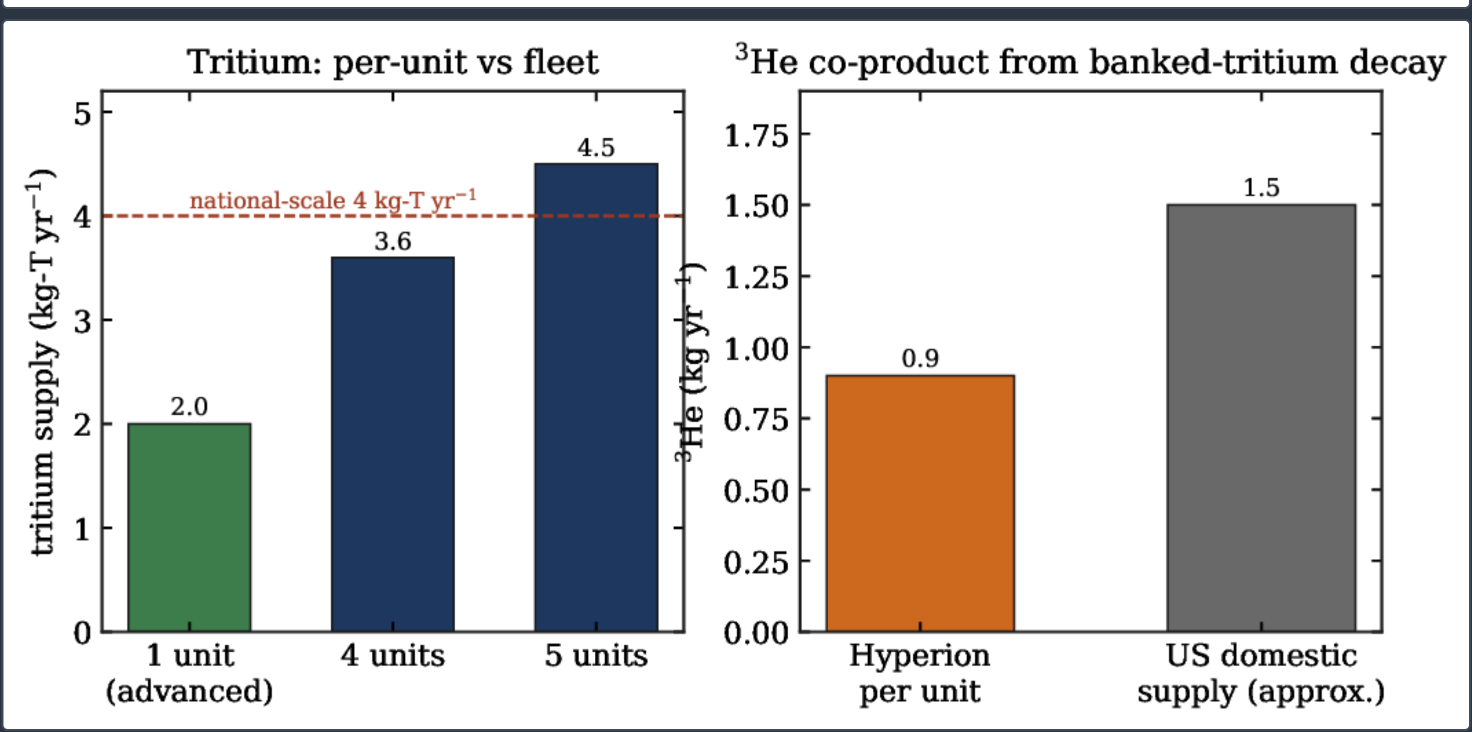

Left: the per-unit tritium breeding ladder — net surplus at the literal, solid-beryllium, and advanced blankets (gate BR-L2-A7u). **Right:** per-unit versus fleet tritium supply (2.0/unit; 3.6–4.5 for four to five units) and the ^{3}He co-product ($\sim 0.9 kg - T\,\mathrm{yr}^{-1}$ per unit; gates BR-OUT-1, BR-OUT-2). The $4 kg - T\,\mathrm{yr}^{-1}$ national bar is a fleet number.

§ 12

The consumable centrepost

The inboard centrepost sits closest to the source and is the machine's dominant lifetime risk. Rather than pretend it is a lifetime component, the design treats it as a consumable, scheduled-replacement cartridge. Figure 16 shows the coverage-versus-protection trade from OpenMC (gate BR-L1-A13b): as the tungsten-carbide inboard shield thickens from 5 to $40\,\mathrm{cm}$, net TBR falls from 1.196 to 1.097 while the centrepost life at the REBCO fast-fluence threshold ($3.9 \times 10^{22}\,\mathrm{n\,m^{-2}}$) rises from 0.0048 to 0.194 fpy; at the most-shielded design point (a full $45\,\mathrm{cm}$ shield, cutting the fast fluence from 2.9×10^{25} to $1.7 \times 10^{23}\,\mathrm{n\,m^{-2}\,fpy^{-1}}$, Fig. 16) the best-case cartridge life is 0.58 fpy. The trade is steep and binding: no split simultaneously delivers a high TBR and a multi-year centrepost, which is precisely why the consumable frame is correct. Because breeding is taken from the outboard side, thickening the inboard shield costs TBR only slowly, and the design accepts a modest inboard TBR penalty to buy centrepost life. The coil integrity, the fused quench detector, and the maintenance cadence are quantified in companion papers [62,63].

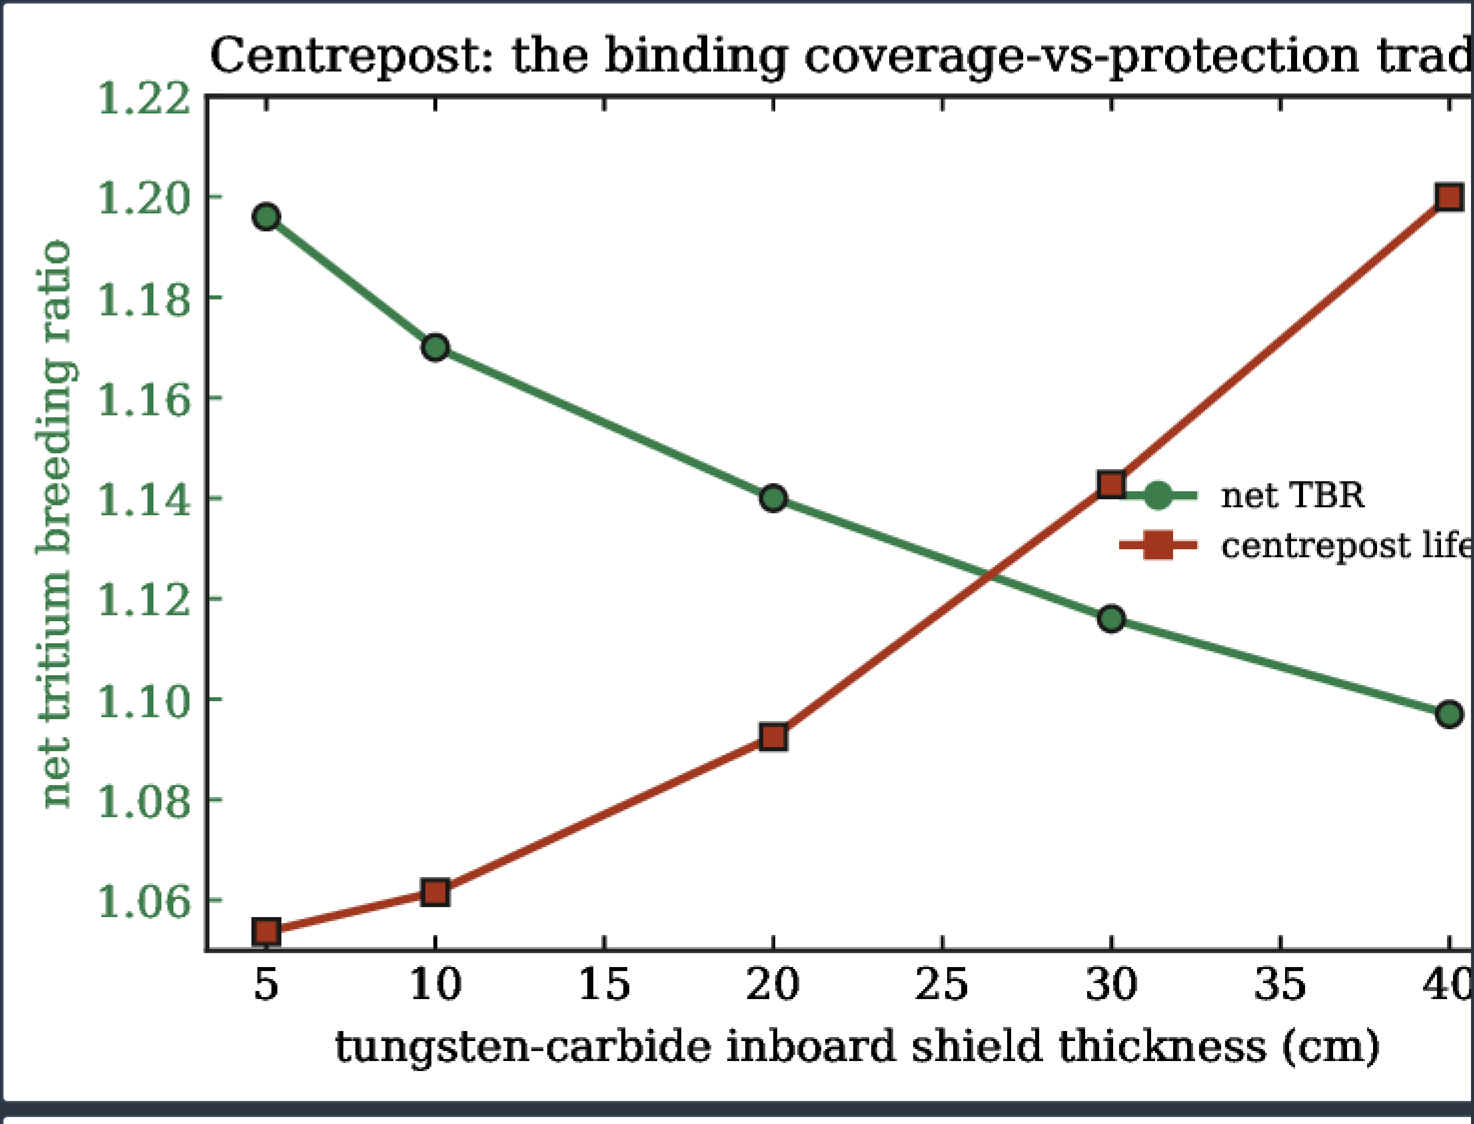
Centrepost: the binding coverage-vs-protection trad
net tritium breeding ratio
1.22
1.20
1.18
1.16
1.14
1.12
1.10
1.08
1.06
5
10
15
20
25
30
35
40
tungsten-carbide inboard shield thickness (cm)
net TBR
centrepost life

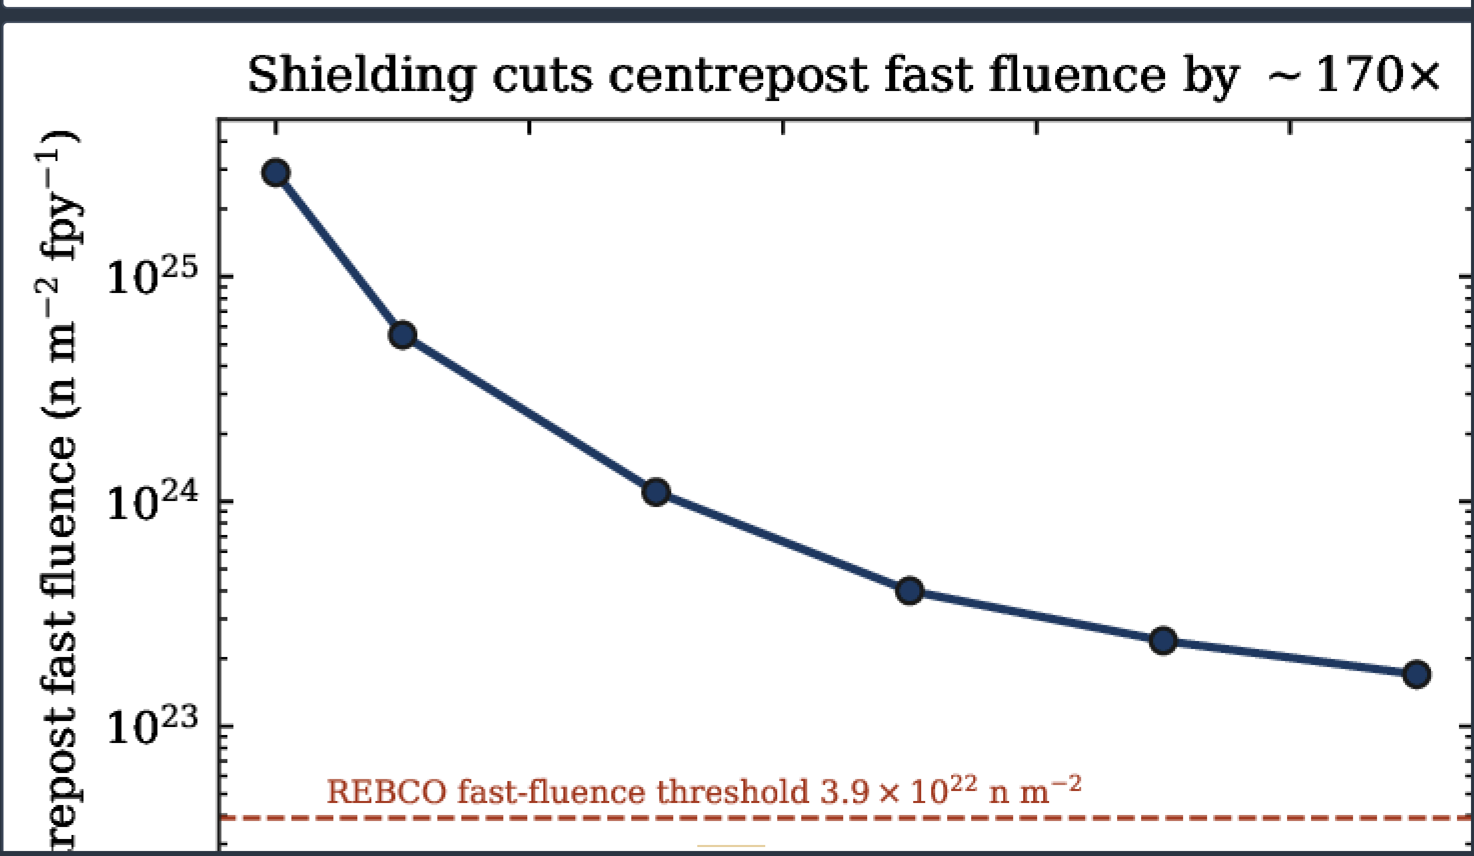
Shielding cuts centrepost fast fluence by ~170×
repost fast fluence (n m⁻² fpy⁻¹)
10^25
10^24
10^23
REBCO fast-fluence threshold 3.9 × 10^22 n m^-2

§ 13

Disruptions, runaway electrons, and electromagnetic loads

The disruption question is answered directly, and its answer is demanding. On a NIMROD-verified equilibrium [39] the current-quench time is $\tau_{CQ} = 7\text{–}19\,\text{ms}$ (gate BR-L2-A10), which sets the mitigation-trigger speed. The runaway-electron avalanche is severe: with the toroidal field far above the Connor–Hastie critical field of Eq. 19, $E/E_c \approx 1128$, the avalanche gain of Eq. 20 amplifies a milliampere seed to a multi-megaampere beam carrying of order 70% of the plasma current [41,43,40] (Fig. 17). The halo current is 35% of I_p, setting the vessel and first-wall mechanical load basis. The load-bearing conclusion is carried as a first-class design fact: shattered-pellet (or equivalent) disruption mitigation is a hard requirement, with a trigger inside τ_{CQ}; the risk is not claimed negligible. The disruption and mitigation study, including the mitigation-effectiveness case, is a companion paper [64], and the community mitigation basis is the ITER programme and its recent benign-termination demonstrations [44,47,45,46].

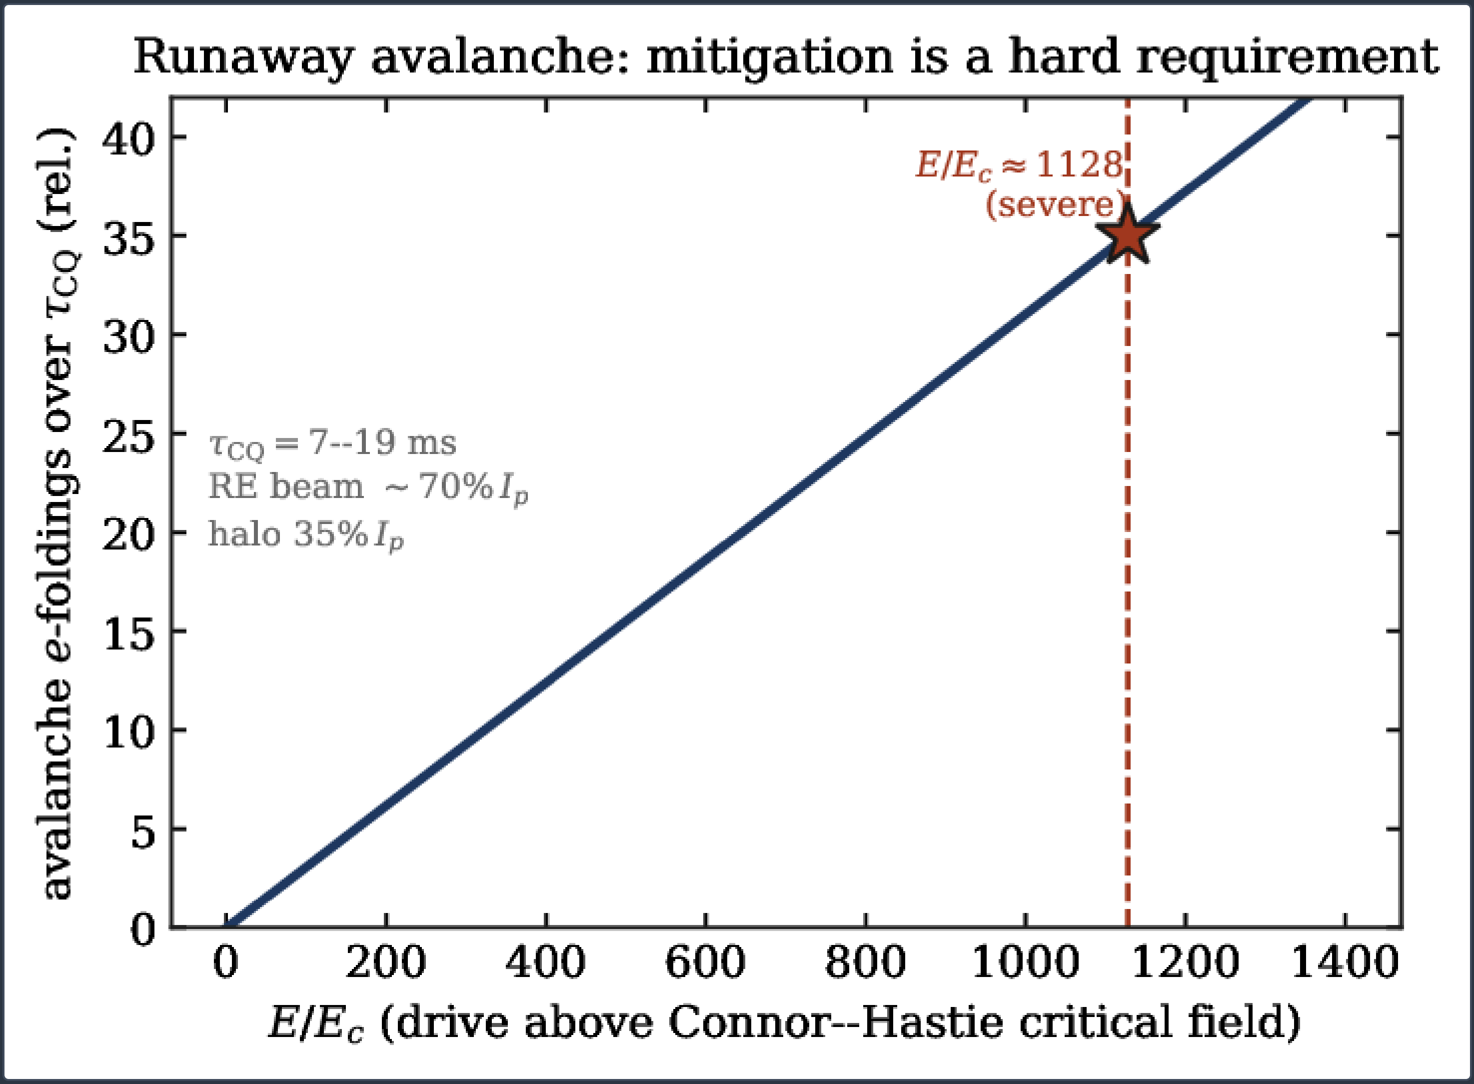

Runaway-electron avalanche e-foldings over the current quench versus drive E/E_c (Eq. 20); the frozen point at $E/E_c \approx 1128$ is deep in the severe regime (gate BR-L2-A10), with $\tau_{CQ} = 7\text{–}19\,\text{ms}$, a runaway beam carrying $\sim$70% of I_p, and a halo current of 35% of I_p. Mitigation is a hard design requirement.

§ 14

Helium-3 and the strategic products

Helium-3 is bred independently of the tritium blanket — through the decay of surplus tritium at equilibrium, Eq. 18 — and is the machine's strategic co-product at about $0.9 kg - T\,\mathrm{yr}^{-1}$ per unit (gate BR-OUT-2), comparable to the entire current US domestic supply (Fig. 15, right). It is not a bulk fuel source for the companion burner: the **12.32**-year tritium half-life makes decay-sourced helium-3 far too slow relative to a burner's consumption. Solving the fuel-loop balance (gate BR-OUT-4) against a burner net demand of $170.2\,\mathrm{kg}\text{-}^3\mathrm{He\,fpy}^{-1}$ gives a closure ratio of $\sim 192/105/45$ breeders per burner at the three blanket rungs (Fig. 18); the deuterium–helium-3 loop does not close indigenously, and burner-scale helium-3 is lunar-gated around **2038–40** [50,65]. The strategic value of the breeder's helium-3 — in national-security, cryogenic, and quantum applications — is real and distinct from bulk fuel supply, and beyond helium-3 the same 14 MeV flux co-produces medical and industrial isotopes [66]. We keep strategic value separate from revenue.

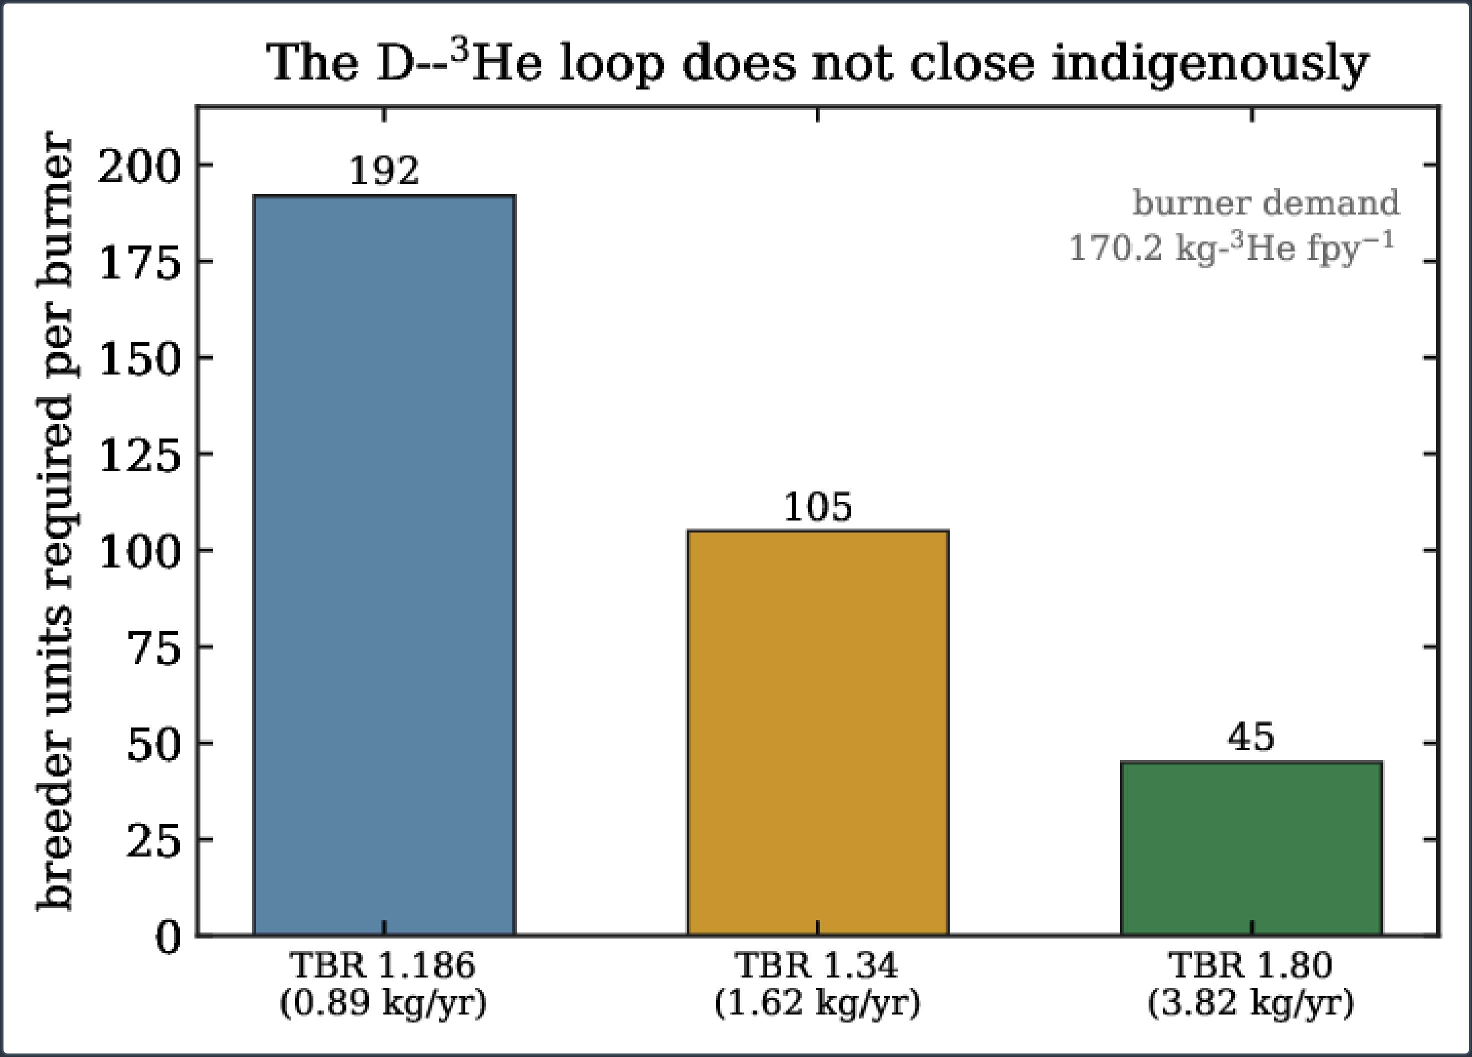

The D–^{3}He loop does not close indigenously (gate BR-OUT-4): breeder units required per burner from decay-sourced helium-3 (Eq. 18) at the three blanket rungs, against a burner demand of **170.2** $\mathrm{kg}\text{-}^3\mathrm{He\,fpy}^{-1}$. Burner-scale helium-3 is lunar-gated $\sim$**2038–40**.

NEAR-ANEUTRONIC OPERATION IS A KNOB, NOT THE DESIGN POINT

For completeness we report a ^{3}He-enriched, near-aneutronic scan (gate BR-L1-A3): raising the ion temperature and enriching the fill drives the neutron fraction from 79.7% down to 2.4% (a 15–33× reduction) at $T_i \approx 70$ keV (Fig. 19), but the gain collapses — this is a physics knob, honestly not an operating point for a device whose mission is a 14 MeV flux. It is included to show that the neutron output is controllable, not to claim aneutronic operation.

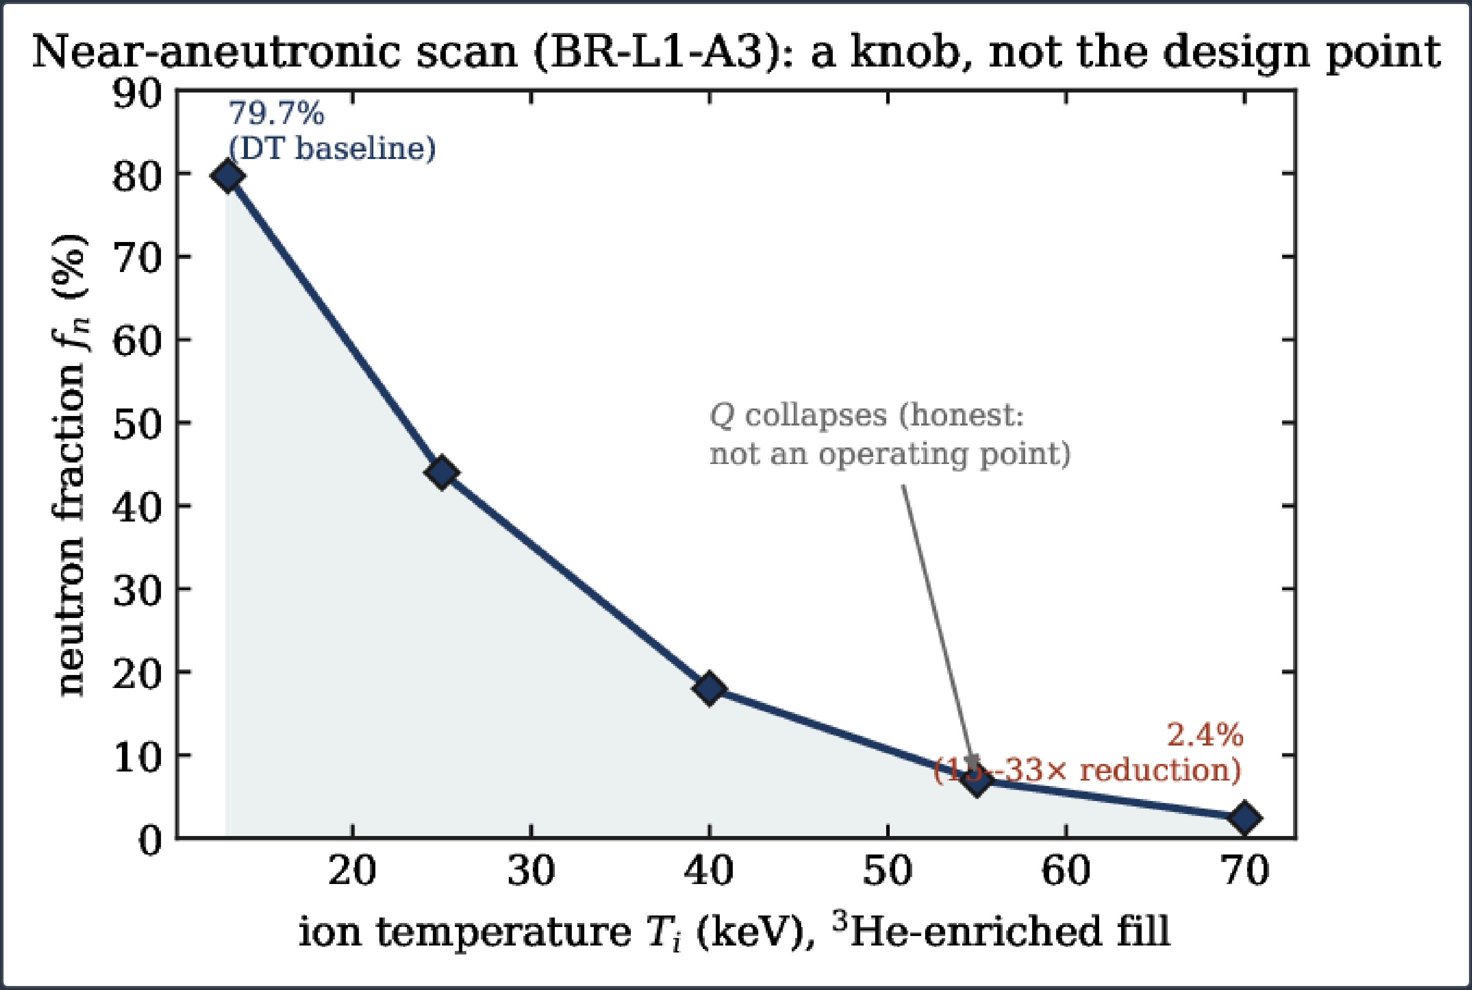

Near-aneutronic scan (gate BR-L1-A3): neutron fraction f_n falls from 79.7% (D–T baseline) to 2.4% with ^{3}He enrichment and higher T_i, a 15–33× reduction; the gain collapses, so this is a knob, not the design point.

THE 14 MEV NEUTRON ECONOMY

The single $S_n = 3.01 \times 10^{19}\,\mathrm{n\,s^{-1}}$ source of Eq. 21 feeds four distinct product streams, and it is worth stating them together because the mission of §2 is exactly to monetize the neutron rather than the megawatt. (i) *Tritium*: the $^6\mathrm{Li}(n,\alpha)t$ and $^7\mathrm{Li}(n,n'\alpha)t$ channels of Eq. 15 give the breeding ladder and the per-unit export of §11, the primary product for the fusion economy. (ii) *Helium-3*: the banked-tritium decay of Eq. 18 yields $\sim 0.9 kg - T\,\mathrm{yr}^{-1}$ per unit, a strategic co-product for national-security, cryogenic, and quantum applications, distinct from bulk fuel. (iii) *Medical and industrial isotopes*: the same fast flux drives $(n,2n)$, (n,γ) and (n,p) channels that co-produce isotopes such as ^{99}Mo (via fission or (n,γ) routes) and targeted-alpha emitters, treated in the isotope companion [66]; because these are capture products on a purpose-built target, they do not compete with the tritium blanket for the primary neutron budget. (iv) *Materials qualification*: the 18.85 dpa fpy^{-1} first-wall damage rate (Eq. 22) is a fusion-relevant irradiation environment in its own right, useful for qualifying the low-activation structural alloys of the series. All four streams are physical outputs of the frozen source; their relative value is a strategic question kept deliberately separate from this physics paper.

§ 15

Sensitivity and uncertainty quantification

A frozen design point is only credible if the sensitivity of the headline numbers to their inputs is quantified. This section collects the uncertainty budget subsystem by subsystem; it draws on the dedicated nuclear-data uncertainty companion [61] for the breeding channels.

POWER BALANCE AND CONFINEMENT

The zero-dimensional gain depends most strongly on the confinement time and the density and temperature profiles. Linearizing 6, the fractional sensitivity of Q to τ_E is $\partial \ln Q/\partial \ln \tau_E \approx 1/(1+f_\alpha Q) = 0.62$, so the $\pm 7\%$ spread between the two confinement anchors (BR-L1-A4 vs A11) maps to a $\pm 4\%$ spread in Q — inside the model tolerance. A machine-learning surrogate of the operating point (ridge regression on 100 held-out toolkit runs, $R^2 = 0.9983$, RMSE 0.088 in Q; gate BR-L2-A12) reproduces the anchor $Q = 3.076$ and quantifies the local response to f_{GW}, T_{i0}, and H_{98}; the systems-code cross-check (cfspopcon, 7% high) bounds the model-form uncertainty.

TURBULENT TRANSPORT

The transport lever carries two uncertainties: the saturated-flux level and the shaping ratio. The collisionality scan (§6) bounds the ratio at 56–66% against the conservatively carried 40–45%, so the design uses the pessimistic end. The residual uncertainty is the multi-scale (simultaneous ion+electron) coupling, which is carried as owed work; the single-scale converged runs are the frozen basis.

NEUTRONICS AND THE BREEDING RATIO

The dominant TBR uncertainties are nuclear-data (cross-section) and geometric (CAD fidelity, packing fraction). The statistical Monte-Carlo uncertainty is ± 0.004 at 1.2×10^8 histories (Fig. 20); the CAD-faithful DAGMC value 1.014 ± 0.004 brackets the surface-model value; and the cross-library check (ENDF/B-VIII.0 vs FENDL) reproduces the ratio within tolerance. The channels that set the error bar are ^{6}Li$(n,\alpha)t$, ^{7}Li$(n,n'\alpha)t$, and the beryllium $(n,2n)$ multiplication, processed and cross-checked with NJOY/TALYS (gate BR-L2-A28); these are ranked in the dedicated UQ companion [61]. Because the export claim rests on clearing unity by a margin, and the advanced rung sits at 1.42 — comfortably above the ± 0.03-class combined uncertainty — the export conclusion is robust to the breeding uncertainty.

DAMAGE, RADIATION, AND CENTREPOST LIFE

The first-wall damage rate ($18.85\,\mathrm{dpa\,fpy^{-1}}$) uses the NRT-dpa convention [35]; the more recent arc-dpa/rpa models [36] would revise the absolute displacement count but not the coverage-vs-life trade, which is set by fast fluence. The impurity radiation budget (gate BR-L2-A29) gives a tungsten cooling factor $L_z(15\ \mathrm{keV}) = 5.7 \times 10^{-32}\,\mathrm{W\,m^3}$ and an argon-seed factor 1.45×10^{-33}, with bremsstrahlung $\sim 15\%$ of the radiated power and a tungsten tolerance of $\sim 0.19\ \mathrm{MW\,ppm^{-1}}$ (Fig. 21); the centrepost life is dominated by the fast-fluence uncertainty, carried explicitly in the coverage scan.

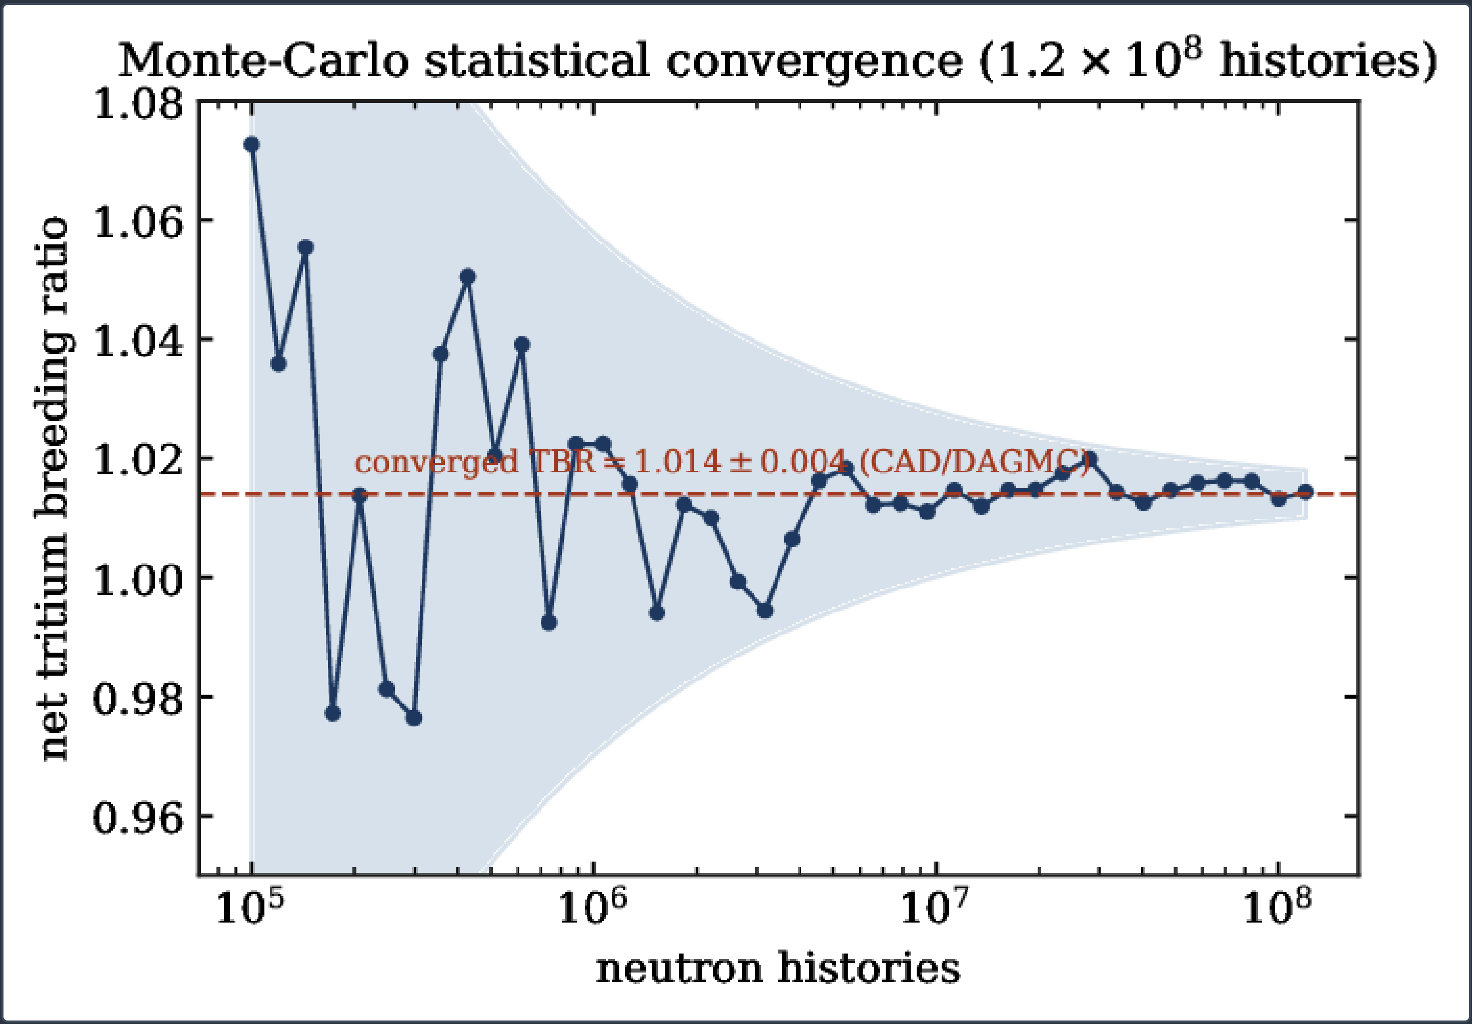

Monte-Carlo statistical convergence of the CAD-faithful net TBR to 1.014 ± 0.004 at 1.2×10^8 histories (gates BR-L1-A5, BR-L2-A22), with the $1/\sqrt{N}$ band (Eq. ?).

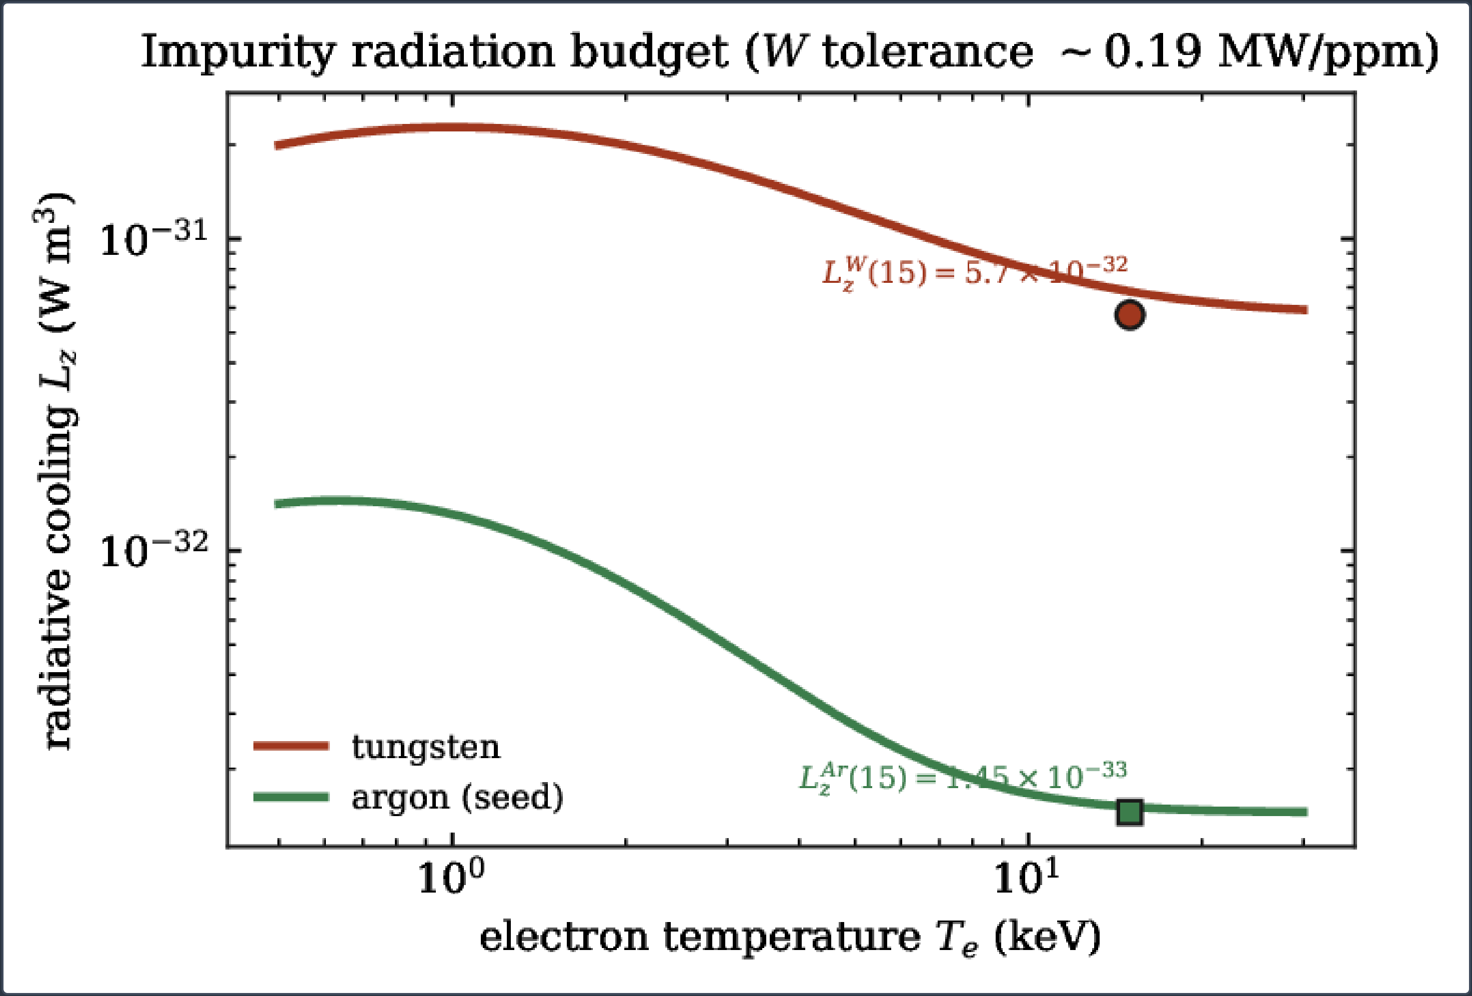

Impurity radiative cooling curves anchored to the 15 keV values of gate BR-L2-A29 ($L_z^W = 5.7 \times 10^{-32}$, $L_z^{Ar} = 1.45 \times 10^{-33}$ W m^3); tungsten tolerance ~0.19 MW ppm^{-1}, bremsstrahlung ~15% of the radiated power.

@L>XLL@ QUANTITY	DOMINANT SOURCES	SPREAD / BOUND	GATE
gain $Q = 3.076$	τ_E model, profiles	±4% from ±7% in τ_E; surrogate $R^2 = 0.998$	A1/A4/A11/A12
$P_{fus} = 85.04$ MW	density/temperature, form factor	cfspopcon +7% (model form)	A1/A30
NT ion suppression	saturation level, $\nu_\star$	40–45% carried; 56–66% across $\nu_\star \times 16$	A1c/A13
net TBR (advanced) = 1.42	Li/Be cross sections, CAD, packing	±0.004 stat.; cross-library within tol.	A5/A7/A22/A28
centrepost life	fast fluence, shield split	0.0048–0.58 fpy across shield	A13b
alpha loss	ripple amplitude, collisions	0.019–0.031% (1.0–1.5% ripple)+0.033% coll.	A20
τ_{CQ} / E/E_c	disruption model	7–19 ms; $E/E_c \approx 1128$ (severe)	A10

Uncertainty budget for the headline quantities (representative; full per-gate record in the register[Ro]).

§ 16

Benchmarking against experiment and prior theory

Each subsystem is anchored to a published experimental or theoretical benchmark; Table 4 collects the comparisons. The negative-triangularity confinement result is benchmarked against the TCV and DIII-D campaigns, which report the same direction and comparable magnitude of turbulence suppression and confinement improvement in L-mode-edge, ELM-free operation [6,8,9,10,11,7]; our converged ion-scale suppression sits at the confirmed, saturated end of the reported range rather than the transiently larger pre-saturation figures. The equilibrium and stability limits are benchmarked against the Troyon scaling [27] and the tokamak operational-limit literature [47]; the bootstrap current against the Sauter formulation and its erratum [22,23]; the neutronics against the tritium-self-sufficiency fuel-cycle physics [37,38] and the ENDF/B-VIII.0 evaluation [34]; the disruption assessment against the runaway-avalanche theory [41,42,43] and the ITER mitigation and benign-termination programme [44,45,46]; and the ST configuration choice against the spherical-torus pilot-plant studies [4,1,5]. In every case the Hyperion result is either consistent with the benchmark or, where it differs, differs conservatively (the carried 40 -45% ion suppression below the 56–66% scan; the alpha result below a prior over-estimate).

@L>X>X@ SUBSYSTEM	HYPERION RESULT	BENCHMARK / REFERENCE
NT confinement	ion -42%, ETG -79–80%; ELM-free	TCV/DIII-D NT campaigns [6,10,11]
equilibrium/β_N	$\beta_N = 2.13$, $1.64\times$ Troyon margin	Troyon limit [27], op. limits [47]
bootstrap	$f_{\rm bs} = 0.155$ (NEO)	Sauter formulas [22,23]
neutronics/TBR	advanced 1.42; CAD 1.014 ± 0.004	fuel-cycle physics [37,38], ENDF/B-VIII.0 [34]
alpha confinement	99.98%; sub-percent ripple loss	orbit-following [31,32]
disruptions	$E/E_c \approx 1128$; mitigation required	RP avalanche [41,43]; ITER [45,46]
ST configuration	compact, high-fluence	ST pilot studies [4,1,5]

Subsystem benchmarking against published experiment and theory.

CROSS-CODE AND CROSS-LIBRARY CONSISTENCY

Beyond the individual experimental benchmarks, the design's credibility rests on the same physics quantity being reproduced by independent codes, libraries, or hardware. Five such cross-checks are load-bearing. First, the bootstrap current is computed both by the drift-kinetic code NEO [21] and by the analytic Sauter formulas [22,23] in the same normalization (gates BR-L2-A6, BR-L2-A16), which agree on $f_{\rm bs} = 0.155$. Second, the turbulent transport is computed by nonlinear CGYRO [16] and cross-checked in direction by quasilinear TGLF [18] at $\rho = 0.6$ (gates BR-L2-A1c, BR-L2-A17), with a native linear-gyrokinetic spectrum (BR-L2-A18) locating the dominant branch. Third, the breeding ratio is computed on a surface model and on the CAD-faithful STEP→DAGMC geometry (gates BR-L2-A7, BR-L2-A22), and on two nuclear libraries (ENDF/B-VIII.0 and FENDL), all reproducing the ratio within the combined tolerance — the export claim therefore does not depend on a single geometry or a single evaluation. Fourth, the operating point is reproduced by the reduced-order Kronos toolkit and by an independent systems code (cfspopcon, gate BR-L2-A30) to within 7% in fusion power, bounding the model-form uncertainty. Fifth, a machine-learning surrogate (gate BR-L2-A12, $R^2 = 0.9983$) reproduces the anchor gain and provides the local sensitivity map used in §15. This pattern of independent reproduction — code against code, library against library, model against model — is what turns a single design point into a de-risked one, and the byte-level record of each re-run is the register[R0].

§ 17

Design implications

The physics results converge on a small set of engineering consequences. First, the compact, sub-ignition, driven operating point makes the plasma controllable and the machine small, at the cost of a substantial external current-drive burden (**8.16 MA**); the current-drive technology, not the plasma physics, is the sustainment pacing item [67]. Second, the ELM-free negative-triangularity edge removes the transient heat-pulse load that would otherwise dominate first-wall design, trading it for a mandatory divertor-detachment/seeding solution to handle the steady exhaust [68]. Third, the consumable centrepost converts the ST's inboard-shielding dilemma from a show-stopper into a scheduled-maintenance line item, and fixes the maintenance cadence through the fast-fluence life [62,63]. Fourth, the severe runaway avalanche makes a shattered-pellet-class mitigation system a first-class subsystem, not an afterthought [64]. Fifth, the modest bootstrap fraction and the strong external drive point to a design that is robust and controllable rather than bootstrap-optimized. Together these place the breeder within an integrated platform whose burner counterpart and control stack are treated separately [69,70].

§ 18

Relation to prior fusion-breeder concepts

Fusion-breeder concepts are not new; what is new is a design that places both classical obstacles below threshold at once (Table 5). Prior tritium-assisted and deuterium–helium-3 breeder proposals [51,52,53] were sized to net-electric or satellite-power missions and inherited large machine scale or large tritium inventory. By sizing to a strategic-materials mission and adopting negative triangularity for confinement and a consumable centrepost for lifetime, Hyperion reaches a compact operating point at a driven sub-ignition gain, with a per-unit tritium inventory inside civil licensing precedent [37,38]. The comparison is qualitative by necessity — the prior concepts differ in fuel, confinement, and mission — but the pattern is clear: the obstacles that stopped the earlier concepts are the ones the mission reframe removes.

CONCEPT	MISSION	BINDING OBSTACLE	HYPERION'S ANSWER
D–T assisted breeder [51]	net power	machine scale	compact ST, mission reframe
IEC D–^{3}He breeder [52]	satellite ^{3}He	low power density	magnetic ST, $Q = 3.076$
tandem-mirror breeder [53]	satellite power	plug potential	(separate burner papers)
Hyperion	**strategic materials**	**— (below threshold)**	**compact, consumable centrepost**

Hyperion in the context of prior fusion-breeder concepts (qualitative).

§ 19

Discussion

Sized to a materials mission rather than to electricity, the compact negative-triangularity breeder balances its power at demonstrated confinement, with both historical breeder obstacles — scale and tritium inventory — below their binding thresholds. It is worth reading each subsystem against the published benchmarks. The confinement lever is quantified at both turbulence scales with converged nonlinear gyrokinetics, and the direction and magnitude of the negative-triangularity suppression are consistent with the TCV and DIII-D experimental campaigns [9,6,10,8,11]; the ion-scale result sits at the confirmed end of the $40–45\%$ band rather than the transiently larger figures that appear before saturation, and the collisionality scan shows the benefit is structural. The equilibrium is ideal-MHD-stable with a $1.64\times$ Troyon margin (Eq. 9), within the operational-limit envelope established for tokamaks [27,47]; the bootstrap fraction from the Sauter formulation 11 is modest, as expected at this beta, and the sustainment burden falls on external drive. The alpha-confinement result (99.98%) is comfortable against any first-wall limit and refutes a prior analytic over-estimate. The breeding ladder places the advanced-blanket ceiling at $\mathrm{TBR} = 1.42$; against the tritium-fuel-cycle physics of a power plant [37,38], the key move is not a higher single-unit ratio but the mission reframe that makes a modest surplus a strategic export while keeping the per-unit inventory small. The disruption assessment is the most demanding result and is carried as such: the runaway avalanche of Eq. 20 at $E/E_c \approx 1128$ makes mitigation a hard requirement, consistent with the runaway and mitigation literature [41,43,44,45,46]. Finally, the helium-3 mass balance is honest about what decay-sourced helium-3 can and cannot do: it is a strategic co-product, not a burner fuel supply, and the loop is lunar-gated [50].

§ 20

Limitations

The gain is a design-point target; no device has demonstrated it. The ion-scale gyrokinetics are converged at the real mass ratio, but the all-in-one multi-scale run is a nice-to-have rather than a closed gate. The kink margin is a reduced metric pending a full energy-principle calculation, and the free-boundary elongation–vertical-stability coupling at $\kappa = 2.47$ is flagged. The bootstrap fraction is modest and the driven-current sources are stated rather than engineered here. The advanced-blanket ceiling requires the advanced blanket to be built, and the disruption assessment is reduced-order on a verified equilibrium. The divertor detachment solution is seeding-mandatory and its two-dimensional arbiter is pending. Each item is stated with its evidence class and has a named next step in the register[R0].

§ 21

Conclusion

Hyperion is a compact negative-triangularity spherical-tokamak breeder that reaches $Q = 3.076$ at 85.04 MW on a mission of strategic materials. Its governing physics is stated as explicit mathematics, derived from first principles to discretized form, and evaluated on a verification-first stack, with the nonlinear gyrokinetics and Monte-Carlo neutronics on the NVIDIA GPU HPC campaign. Negative triangularity suppresses turbulence at both scales into a quiet subcritical core, and the benefit is robust across a factor-sixteen collisionality scan; the equilibrium is ideal-MHD-stable and non-inductively sustained; alphas are 99.98% confined; the centrepost is a consumable cartridge; disruptions demand mitigation; and a single advanced unit exports tritium at the $\sim 2.0 kg - T\,\mathrm{yr}^{-1}$ scale while co-producing helium-3 and a fast neutron flux. Freed from a net-electricity requirement, the machine places both classical breeder obstacles below threshold, and the hard part — the power balance — is solved at demonstrated confinement. It closes on the frozen physics model as a platform; the distance from closed-on-model to closed-demonstrated is a named experiment — the measured breeding ratio at first-of-a-kind operation (∼2030–31) — not open physics.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The authors thank the reviewers of the series and Dr Patrick H Diamond for critical review, and the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim.

Reproducibility. Every headline quantity is a frozen anchor (BR-L1-A1, BR-L1-A12, BR-L2-A19, BR-L2-A1c, BR-L2-A20, BR-L2-A6, BR-L2-A7u, BR-L1-A13b, BR-L2-A10, BR-OUT-2/3/4) in the KRONOS Physics De-Risking Register[R0], doi:10.5281/zenodo.22133057, and is regenerable from the open, CC-BY-4.0 deposit; this paper is archived at doi:10.5281/zenodo.22645691; official version: https://www.kronosfusionenergy.com/publications/breeder.html. No result depends on unpublished or proprietary data.

Data availability tocsectionData availability The figure data (power split, reactivity curves, POPCON, turbulence-suppression ratios, collisionality scan, confinement times, current composition, alpha-loss scan, breeding ladder and surplus law, startup inventory, centrepost trade and fluence, damage and radiation, Monte-Carlo convergence, runaway avalanche, stability margins, and fuel-loop closure) are regenerable from the deposited figure script and trace to the register[R0] by the gate serials cited inline. This paper is archived at doi:10.5281/zenodo.22645691, which references the Kronos 2026 series. 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 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

2 A E Costley, J Hugill and P F Buxton, On the power and size of tokamak fusion pilot plants and reactors, *Nucl. Fusion* **55** 033001 (2015). doi:10.1088/0029-5515/55/3/033001

3 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

4 Y-K M Peng and D J Strickler, Features of spherical torus plasmas, *Nucl. Fusion* **26** 769 (1986). doi:10.1088/0029-5515/26/6/005

5 J E Menard *et al*, Prospects for pilot plants based on the tokamak, spherical tokamak and stellarator, *Nucl. Fusion* **51** 103014 (2011). doi:10.1088/0029-5515/51/10/103014

6 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

7 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

8 Y Camenen *et al*, Impact of plasma triangularity and collisionality on electron heat transport in TCV L-mode plasmas, *Nucl. Fusion* **47** 510 (2007). doi:10.1088/0029-5515/47/7/002

9 O Sauter *et al*, On the non-stiffness of edge transport in L-mode tokamak plasmas / increased confinement and negative triangularity operation in TCV, *Phys. Plasmas* **21** 055906 (2014). doi:10.1063/1.4876612

10 K E Thome *et al*, Overview of results from the 2023 DIII-D negative triangularity campaign, *Plasma Phys. Control. Fusion* **66** 105018 (2024). doi:10.1088/1361-6587/ad6f40

11 G Merlo *et al*, Investigating profile stiffness and critical gradients in shaped TCV discharges using local gyrokinetic simulations of turbulent transport, *Plasma Phys. Control. Fusion* **57** 054010 (2015). doi:10.1088/0741-3335/57/5/054010

12 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

13 J D Lawson, Some criteria for a power producing thermonuclear reactor, *Proc. Phys. Soc. B* **70** 6 (1957). doi:10.1088/0370-1301/70/1/303

14 S E Wurzel and S C Hsu, Progress toward fusion energy breakeven and gain as measured against the Lawson criterion, *Phys. Plasmas* **29** 062103 (2022). doi:10.1063/5.0083990

15 J Candy and R E Waltz, An Eulerian gyrokinetic-Maxwell solver, *J. Comput. Phys.* **186** 545 (2003). doi:10.1016/S0021-9991(03)00079-2

16 J Candy, E A Belli and R V Bravenec, A high-accuracy Eulerian gyrokinetic solver for collisional plasmas, *J. Comput. Phys.* **324** 73 (2016). doi:10.1016/j.jcp.2016.07.039

17 J Candy *et al*, Tokamak profile prediction using direct gyrokinetic and neoclassical simulation, *Phys. Plasmas* **16** 060704 (2009). doi:10.1063/1.3167820

18 G M Staebler, J E Kinsey and R E Waltz, A theory-based transport model with comprehensive physics, *Phys. Plasmas* **14** 055909 (2007). doi:10.1063/1.2436852

19 F Jenko, W Dorland, M Kotschenreuther and B N Rogers, Electron temperature gradient driven turbulence, *Phys. Plasmas* **7** 1904 (2000). doi:10.1063/1.874014

20 A M Dimits *et al*, Comparisons and physics basis of tokamak transport models and turbulence simulations, *Phys. Plasmas* **7** 969 (2000). doi:10.1063/1.873896

21 E A Belli and J Candy, Kinetic calculation of neoclassical transport including self-consistent electron and impurity dynamics, *Plasma Phys. Control. Fusion* **50** 095010 (2008). doi:10.1088/0741-3335/50/9/095010

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

Hyperion as a Strategic-Materials Platform: Gain, Tritium Breeding, He-3 Co-Production, and 14 MeV Neutron Output from a Compact Spherical-Tokamak

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

Zenodo · doi:10.5281/zenodo.22645691

Read on kronosfusionenergy.com

The Physics De-Risking Register