KRONOS FUSION ENERGY

PAPER 2.6 · THE KRONOS 2026 SERIES

The Negative-Triangularity Edge: Heat-Flux Width, the ELM-Free Pedestal, and the Scrape-Off-Layer Boundary Condition for Power Handling in a Compact Spherical-Tokamak Breeder

Negative triangularity turns the plasma edge from the power-handling liability it is in a conventional tokamak into an engineering asset. At the Hyperion breeder's frozen...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

The Negative-Triangularity Edge: Heat-Flux Width, the ELM-Free Pedestal, and the Scrape-Off-Layer Boundary Condition for Power Handling in a Compact Spherical-Tokamak Breeder

Negative triangularity turns the plasma edge from the power-handling liability it is in a conventional tokamak into an engineering asset. At the Hyperion breeder's frozen operating point ($\delta = -0.30$, $I_p = 9.66\,MA$, $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $B_0 = 8\,T$, $\kappa = 2.0$, $R_0 = 1.20\,m$, aspect ratio $A = 2.5$), the edge is ELM-free by construction: the peeling–ballooning drive of an H-mode pedestal is never accessed, so the transient edge-localised-mode heat pulse that sizes conventional divertors is absent, and the confinement is carried by the core ($\tau_E = 0.410\,s$, effective $H_{98} \approx 0.94$). The power-handling problem therefore reduces cleanly to the steady scrape-off layer (SOL), which we treat as a formal boundary-value problem. We derive the coupled system — the Eich heat-flux width, the two-point parallel-transport model with power and momentum losses, the Lengyel impurity-radiation integral, and the sheath target boundary condition — and close it on the config-22021 equilibrium with a GPU-accelerated code campaign. A first-principles heat-flux width of $\lambda_q = 0.207\,mm$ (Eich regression) sets a parallel heat flux $q_{\|} = 27.4\,GW/squarem$ and an unmitigated attached target load of **1370** (conventional) and **171** $MW/squarem$ (Super-X, $f_x = 8$), both far above the $\sim 10\,MW/squarem$ material limit. A converged UEDGE quasi-1-D solve resolves the density-driven target collapse from **14.4** through **5.3** to **1.05** eV with a divertor/upstream compression $n_t/n_u = \mathbf{31}$ at onset, within $\sim 15\%$ of the independent 1-D arbiter (≈ 27); a converged SD1D (BOUT++) fluid-plus-neutral solve reproduces the same physics and detaches the target with $< 0.5\%$ nitrogen. Density alone cannot detach the leg at design exhaust, so seeding is mandatory; divertor geometry then walks the required seed fraction from **6.4%** in a single null to $\sim 0.5\%$ with a stacked advanced divertor. At the cold detached target, a binary-collision sputtering calculation gives a deuterium-on-tungsten peak yield of 1.35×10^{-2} atoms ion^{-1} at sub-monolayer-per-pulse erosion, and the sheath energy falls below the physical-sputtering threshold. The negative-triangularity edge is quantitative, bounded, and controllable, with a 2-D SOLPS-ITER divertor solve as the one open confirmation.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645830 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: negative triangularity · heat-flux width · scrape-off layer · divertor detachment · ELM-free edge · impurity seeding · plasma-material interaction · spherical-tokamak breeder

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

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

§ 1

Introduction

THE POWER-HANDLING PROBLEM OF A COMPACT BREEDER

A spherical tokamak (ST) confines a large plasma current in a small volume, which is exactly what makes it an efficient 14 MeV neutron source and tritium-breeding platform [1,2] and, at the same time, what concentrates its exhaust. The Hyperion breeder holds $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$ and a minor radius $a = 0.48\,m$; roughly 29 MW of heating power must ultimately cross the last closed flux surface and be exhausted through a scrape-off layer (SOL) whose width, as we show below, is of order a fifth of a millimetre. The resulting parallel heat flux is tens of gigawatts per square metre — an order of magnitude above the ITER value [3] — and the design must dissipate almost all of it before it reaches a material surface. Power handling, not confinement, is the binding edge question for a compact breeder.

WHY NEGATIVE TRIANGULARITY CHANGES THE QUESTION

In a conventional positive-triangularity (PT) tokamak the edge is governed by the H-mode pedestal and its periodic collapse. The pedestal is held on the boundary of the coupled peeling–ballooning (P–B) instability [4]; when the edge pressure gradient and bootstrap current push the plasma across that boundary, an edge-localised mode (ELM) expels a burst of energy onto the divertor in $\lesssim 1\,ms$ [5]. These transient loads, not the steady exhaust, size the plasma-facing components of a reactor-class PT device. Negative triangularity (NT) removes them. The inverted shape raises the edge magnetic shear and suppresses access to the ballooning-unstable second stability region, so a steep pedestal never forms and the edge stays in an ELM-free, L-mode-like state while the core still delivers reactor-relevant confinement [6,7,8]. The Hyperion point is chosen inside this regime at $\delta = -0.30$; the accompanying quiescent edge is documented in the companion peeling–ballooning study [31] and rests on the negative-triangularity turbulence suppression established by nonlinear gyrokinetics [32]. The practical consequence for this paper is decisive: with the transient ELM load removed *by construction*, the power-handling problem collapses to a single, tractable, steady-state boundary-value problem for the SOL and divertor.

PRIOR ART AND THE OPEN QUESTION

The physics of that steady problem is mature. The upstream heat-flux width obeys the Eich multi-machine regression [9] and the Goldston heuristic-drift (HD) model [10]; parallel transport along the SOL is captured by the two-point model [11,12]; divertor detachment — the partial extinction of the target plasma by volumetric power, momentum and particle losses — is described by the Lengyel impurity-radiation integral and its modern refinements [14,13,15], and is now routinely computed by 2-D edge codes such as SOLPS-ITER [20] and UEDGE [21] and by 1-D fluid-plus-neutral models such as SD1D [22,23]. Advanced divertor geometries — the Super-X [16], the snowflake [17], and their experimental realisations [18] — and highly radiative NT pilot-plant concepts [19] extend the detachment window. What has not been done is to close this chain *quantitatively on a compact high-field ST breeder*: to compute the actual heat-flux width and parallel flux at the frozen operating point, resolve the detachment threshold with a converged edge-code solve on the design equilibrium, quantify the seeding required and the geometry that relaxes it, and confirm that the resulting target is a benign sputtering environment.

CONTRIBUTION

This paper does exactly that. We (i) pose the NT-edge power-handling problem formally, deriving the heat-flux-width, two-point, detachment and sheath equations and stating why the P–B drive is not accessed at $\delta = -0.30$ (§2); (ii) close the system on the config-22021 equilibrium with a GPU-accelerated code campaign — FreeGS/TGYRO for the upstream boundary condition, a two-point/Eich reduced-order engine, UEDGE 8.1.2, SD1D on BOUT++, an independent 1-D solver, and RustBCA for plasma–material interaction (§3); (iii) report the results — $\lambda_q = 0.207\,mm$, $q_\parallel = 27.4\,GW/squarem$, the density- and seeding-driven target collapse, the four-method cross-check of the detachment threshold, the seeding-portfolio ladder, and the sputtering budget (§4); (iv) benchmark every number against the published SOL, detachment and NT literature (§5); and (v) name the one open confirmation (§6). Every quantitative claim traces to a frozen anchor in the Kronos Physics De-Risking Register [30]; serial identifiers (e.g.\ HX-14, BR-L2-A21) are given inline so each number is auditable.

§ 2

Governing equations and the formal edge problem

UPSTREAM BOUNDARY CONDITION: POWER ENTERING THE SOL

The SOL is driven by the power crossing the separatrix. From the frozen power balance,

$$P_{\mathrm{SOL}} = P_\alpha + P_{\mathrm{aux}} - P_{\mathrm{rad,core}} = \frac{P_{\mathrm{fus}}}{5} + \frac{P_{\mathrm{fus}}}{Q} - \left(P_{\mathrm{brem}} + P_{\mathrm{sync}}\right),$$

with $P_\alpha = P_{\mathrm{fus}}/5 = 17.0\,MW$ the charged-alpha heating (the neutron power $P_n = 67.8\,MW$ leaves the plasma directly), $P_{\mathrm{aux}} = P_{\mathrm{fus}}/Q = 27.7\,MW$, and core radiation $P_{\mathrm{brem}} + P_{\mathrm{sync}} = 3.12 + 1.55\,MW$ (register KX-L1-A3). Equation (1) gives $P_{\mathrm{SOL}} \approx 29\,MW$, comfortably above the L–H-like power threshold and, more to the point here, the quantity that the divertor must ultimately dissipate.

THE HEAT-FLUX WIDTH: EICH REGRESSION AND THE GOLDSTON HEURISTIC-DRIFT MODEL

The power entering the SOL is deposited in a narrow layer at the outboard midplane (OMP). Its *e*-folding width λ_q is set, for attached conditions, by the competition between cross-field transport and parallel losses. Two independent predictors bracket it. The Eich multi-machine regression [9] (regression 14, attached H-mode/L-mode database including STs) gives

$$\lambda_q\ [\mathrm{mm}] \;=\; 0.63\ B_{\theta,\mathrm{OMP}}^{-1.19},$$

with the outboard poloidal field $B_{\theta,\mathrm{OMP}} = 2.55\,T$. The Goldston HD model [10] derives the same width from first principles by balancing the ∇B and curvature drifts of ions into the SOL against near-sonic parallel outflow to the plates,

$$\lambda_q^{\mathrm{HD}} \;\sim\; 2\,\frac{a}{R_0}\,\rho_{\theta i}, \qquad \rho_{\theta i} = \frac{\sqrt{2m_i T_i}}{e\,B_{\theta,\mathrm{OMP}}},$$

with $\rho_{\theta i}$ the poloidal ion gyroradius. For the breeder edge ($T_i \sim 300\,eV$, $B_{\theta,\mathrm{OMP}} = 2.55\,T$) equation (3) yields a sub-millimetre width consistent with equation (2) to within the ~ 15–20% that separates the two predictors across the tokamak database [9,10]. Both scale inversely with B_θ; the high poloidal field of a high-current compact ST therefore forces λ_q well below a millimetre.

PARALLEL TRANSPORT: THE TWO-POINT MODEL WITH LOSSES

Mapping the OMP flux to the target requires the parallel heat flux and the parallel-transport balance. The parallel flux is the perpendicular flux divided by the field-line pitch,

$$q_\parallel = \frac{P_{\mathrm{SOL}}}{2\pi R_0\,\lambda_q} \cdot \frac{B}{B_\theta},$$

which at the operating point is tens of gigawatts per square metre. Along the field line of connection length $L_\parallel$, electron heat conduction is Spitzer–H\"arm, $q_\parallel = -\kappa_{0e} T_e^{5/2}\,\partial_\parallel T_e$ with [redacted]. Integrating the conduction equation between an upstream point "*u*" and the target "*t*" gives the first two-point relation,

$$T_{e,u}^{7/2} = T_{e,t}^{7/2} + \frac{7}{2}\,\frac{q_\parallel\,L_\parallel}{\kappa_{0e}}\,(1 - f_{\mathrm{cond}}),$$

where f_{cond} is the fraction of parallel power radiated upstream of the target. Total-pressure balance along the flux tube, degraded by charge-exchange and recombination momentum losses f_{mom}, gives the second,

$$2\,n_t T_{e,t} = (1 - f_{\mathrm{mom}})\,n_u T_{e,u}.$$

Closing with the Bohm sheath heat transmission $q_t = \gamma\,n_t T_{e,t}\,c_{s,t}$ (sheath coefficient $\gamma \approx 6$, sound speed $c_{s,t} = \sqrt{2T_{e,t}/m_i}$) yields the target electron temperature and load. In the conduction-limited, loss-free limit equations (5)–(6) reproduce the classic $T_{e,t} \propto n_u^{-2}$ collapse; the added momentum loss steepens it toward detachment.

THE DETACHMENT CRITERION AND THE LENGYEL INTEGRAL

Detachment is reached when the target electron temperature falls below $T_{e,t} \simeq 5\,eV$, at which point volume recombination and charge-exchange remove particles and momentum faster than they arrive and the target ion flux rolls over. The impurity concentration c_z required to radiate the necessary power fraction follows from the Lengyel integral of the parallel conduction equation with a line-radiation sink,

$$q_{\|}^{2} = 2\,\kappa_{0e}\,c_z \int_{T_t}^{T_u} T_e^{1/2}\,L_z(T_e)\,\mathrm{d}T_e,$$

where $L_z(T_e)$ is the impurity radiative loss function (here nitrogen, evaluated with ADAS atomic data). Equation (7) makes the seed requirement explicit: at fixed L_z, the concentration needed to detach scales as $c_z \propto q_{\|}^2$, so the extreme $q_{\|}$ of a compact ST would demand an intolerable seed unless $q_{\|}$ at the target is first reduced by divertor geometry (flux expansion, connection-length lengthening). This is the analytic origin of the seeding-plus-geometry portfolio quantified in §4.

WHY THE EDGE IS ELM-FREE AT $\delta = -0.30$

delta=-0.30 The above SOL problem is a *steady* one only because the edge does not build a collapsing pedestal. The peeling–ballooning drive is set by the normalised edge pressure gradient and edge current,

$$\alpha = -\,\frac{2\mu_0 R_0\,q^2}{B^2}\,\frac{\mathrm{d}p}{\mathrm{d}r}, \qquad \hat{J} = \frac{\langle J_{\|}\rangle_{\mathrm{edge}}}{\langle J\rangle},$$

and instability sets in when the operating point $(\alpha, \hat{J})$ crosses the P–B boundary [4]. Negative triangularity raises the local magnetic shear on the outboard side and denies access to the ballooning second-stability region, so the edge gradient saturates at low α: the operating point stays in the stable interior of the $(\alpha, \hat{J})$ plane (figure 10) rather than riding the boundary, and no ELM cycle exists to be triggered [6,8,31]. The confinement is then carried by the core, not by an edge barrier; the integrated transport solve gives $\tau_E = 0.410\,s$ and effective $H_{98} \approx 0.94$ (register BR-L1-A11). This is the load-bearing simplification of the whole analysis: the divertor is sized by the steady exhaust of equation (4), never by a transient.

TARGET SHEATH ENERGY AND PHYSICAL SPUTTERING

The last equation of the chain is at the surface. The energy of an ion crossing the sheath is $E_i \approx 2T_i + 3Z\,T_{e,t} \simeq 5\,T_{e,t}$, and the gross physical-sputtering rate of the tungsten target is $\Gamma_{\mathrm{ero}} = Y(E_i)\,\Gamma_{\mathrm{ion}}/n_W$, with Y the yield and n_W the atom density. Because $Y(E)$ has a threshold $E_{\mathrm{th}} \approx 209\,eV$ for deuterium on tungsten, a target held cold by detachment ($T_{e,t} \lesssim 2\,eV \Rightarrow E_i \lesssim 10\,eV$) sits below threshold and sputters negligibly. The detachment that protects the target thermally also protects it chemically — a point we quantify in §4.8.

§ 3

Computational methodology: the GPU-HPC campaign

SOLVER STACK AND PROBLEM SCALE

The edge problem was closed on the frozen config-22021 equilibrium by a coupled campaign run on the Kronos NVIDIA GPU HPC campaign (A100/H100/A10 partitions) together with the in-house digital-twin node; the codes and their roles are listed in table 1. The upstream boundary condition of equation (1) is inherited from the flux-driven integrated transport solve (TGYRO/KRONOS toolkit, register BR-L1-A11) on the GPU-accelerated free-boundary equilibrium (FreeGS, 129×129 grid, register BR-L1-A12), which fixes $P_{\rm SOL}$, R_0, a, κ, B_θ and the connection length. The heat-flux width and two-point/Lengyel reduced-order engine (register HX-14/HX-15) evaluates equations (2)–(7) at the design geometry.

CONVERGED EDGE-CODE SOLVES

The detachment threshold was resolved by two converged edge codes on the same equilibrium. *UEDGE 8.1.2* (register BR-L2-A21) solves the 2-D fluid plasma–neutral equations on a Cartesian-slab SOL (6×16 mesh, $L_\parallel \approx 6.5\,m$, recycling $R = 0.98$, analytic hydrogenic rates); every reported point is a true steady state reached by time-step-ramp continuation of the Newton–Krylov solver (`nksol`, termination flag `iterm=1`). *SD1D* on the *BOUT++* framework [23,22] (register HX-14b) solves the 1-D parallel fluid equations with recycling neutrals, nitrogen ADAS radiation (`scd/plt/prb96_n`), sheath coefficient $\gamma = 6$ and CVODE time integration, on a $L_\parallel = 28.3\,m$, $n_y = 120$ field line at a divertor-throat parallel flux of 5 GW/squarem. An independent, purpose-built 1-D fluid solver (Spitzer–H\"arm conduction, Voronov ionisation, recombination, charge-exchange friction, coronal nitrogen L_z, Bohm sheath) was run at the full OMP peak $q_\parallel = 27.4\,GW/squarem$ as a cross-check. Plasma–material interaction at the target used *RustBCA* [24] (register BR-L2-A27), a binary-collision-approximation code, across an 18-case deuterium-on-tungsten energy/angle scan (100/300/1000 eV, $0°/60°$).

GPU ACCELERATION, VERIFICATION AND REPRODUCIBILITY

The GPU HPC campaign carried the parts of the problem that are throughput-bound: the free-boundary and transport solves that fix the upstream boundary condition, the density- and seeding-scan sweeps that map the detachment window (each scan point an independent steady-state solve, embarrassingly parallel across the fleet), and the BCA particle ensembles. The reduced-order two-point/Lengyel engine and the 1-D solvers ran on the in-house node so that the boundary-condition sweeps never contended with the fleet. Verification is by *cross-method redundancy*: the detachment threshold is computed four independent ways — analytic two-point, analytic Lengyel, converged SD1D, and the independent 1-D solver — and by *cross-code agreement* with the UEDGE compression ratio against a 1-D arbiter (§4). Every figure is regenerated deterministically from the archived scan outputs by a single script deposited with this paper; the frozen anchors are catalogued in the de-risking register [30]. No proprietary or licence-gated data is required to reproduce the reduced-order and 1-D results; the 2-D SOLPS-ITER confirmation (§6) is the one step that is access-gated.

CODE / METHOD	ROLE IN THE EDGE PROBLEM	ANCHOR
FreeGS (free-boundary GS)	config-22021 equilibrium, B_θ, $L_\parallel$	BR-L1-A12
TGYRO / KRONOS toolkit	flux-driven $P_{\rm SOL}$, τ_E, H_{98}	BR-L1-A11
two-point + Eich engine	λ_q, $q_\parallel$, target load, dissipation	HX-14
Lengyel + ADAS (N)	seed fraction for radiated power	HX-15
UEDGE 8.1.2 (2-D slab)	converged density-driven detachment	BR-L2-A21
SD1D / BOUT++ (1-D)	converged fluid+neutral detachment	HX-14b
self-written 1-D solver	full-$q_\parallel$ cross-check	HX-14b
RustBCA (BCA)	D→W sputtering, target erosion	BR-L2-A27
portfolio synthesis	seeding/geometry ladder	BR-L2-A11-SUM / BR-SX-09

The GPU-HPC edge campaign: codes, roles, and the register anchor each supplies.

§ 4

Results

GEOMETRY AND THE UPSTREAM BOUNDARY CONDITION

Table 2 collects the config-22021 geometry and the SOL boundary parameters that follow from equations (1)–(4); figure 1 shows the power path from core to target. The frozen point delivers $P_{\mathrm{SOL}} = 29\, MW$ into a sub-millimetre layer.

QUANTITY	SYMBOL	VALUE
major / minor radius	R_0 / a	1.20 m / 0.48 m
aspect ratio, elongation, triangularity	A, κ, δ	2.5, 2.0, -0.30
toroidal field, plasma current	B_0, I_p	8 T, 9.66 MA
outboard poloidal field	$B_{\theta,\mathrm{OMP}}$	2.55 T
fusion power, gain	P_{fus}, Q	85.04 MW, 3.076
power to SOL	P_{SOL}	29 MW
heat-flux width (Eich)	λ_q	0.207 mm
parallel heat flux	$q_{\parallel}$	27.4 GW/squarem
upstream separatrix temperature	$T_{e,u}$	407 eV
energy confinement, *H*-factor	τ_E, H_{98}	0.410 s, ≈ 0.94
Greenwald density	n_{GW}	$1.34e21\ /cubicm$

Config-22021 geometry and the derived SOL boundary condition (frozen anchors; register HX-14, BR-L1-A11/A12).

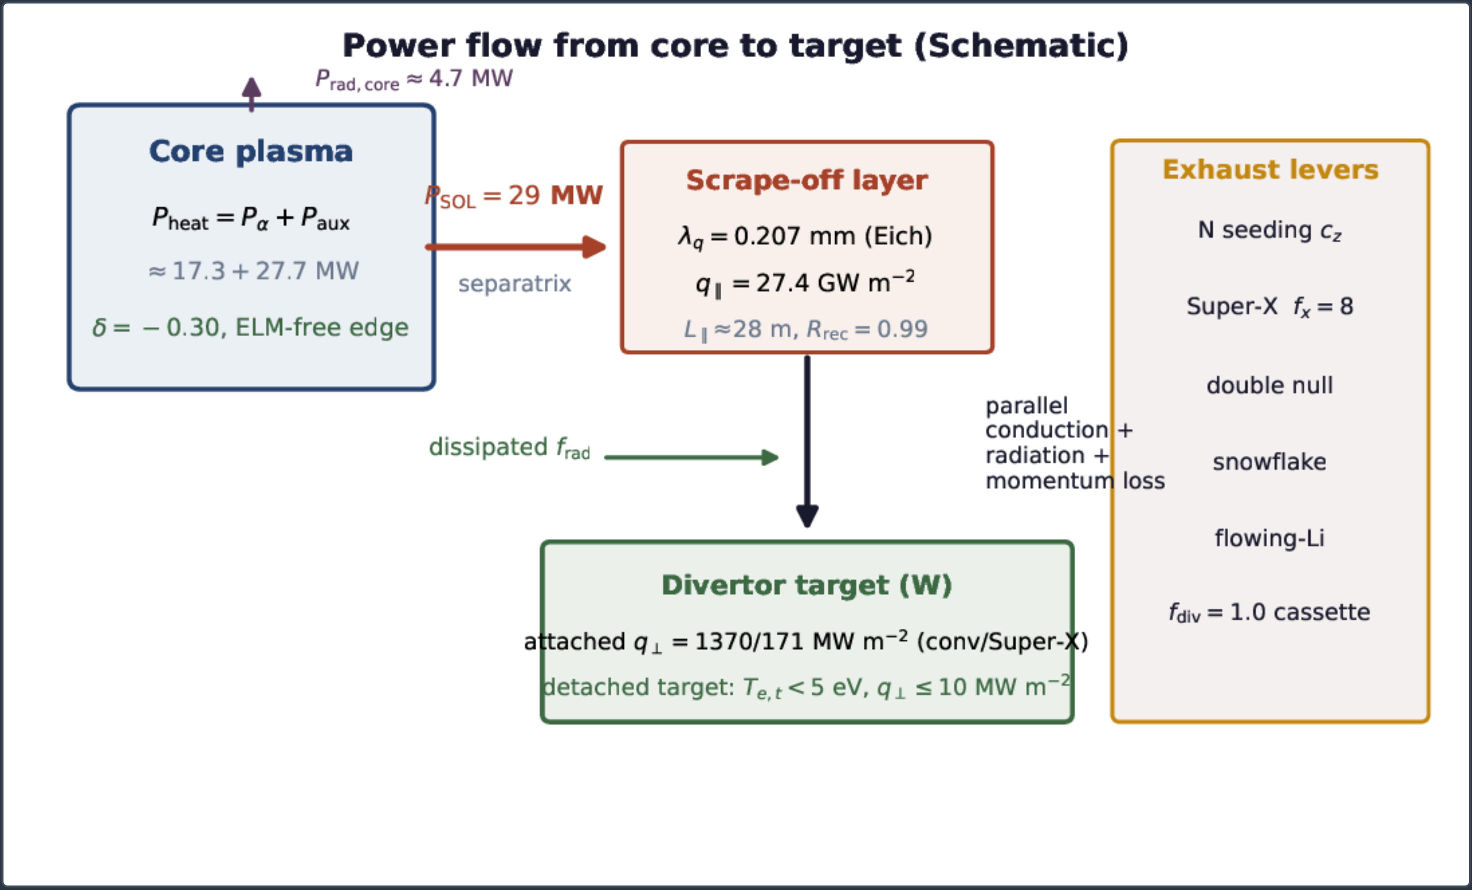

(Schematic.) Power flow from the ELM-free core across the separatrix into the sub-millimetre scrape-off layer and onto the divertor target, with the exhaust levers that dissipate it. Frozen values are annotated; the figure encodes structure, not a solved profile.

HEAT-FLUX WIDTH AND THE PARALLEL LOAD

Figure 2(a) places the breeder on the Eich regression of equation (2): the high outboard poloidal field $B_{\theta,\mathrm{OMP}} = 2.55\,T$ drives $\lambda_q = 0.207\,mm$, at the narrow end of the multi-machine database and consistent with the Goldston HD band of equation (3) to within ~ 15–20%. Mapping this width through equation (4) gives the parallel heat flux of figure 2(b): $q_\parallel = 27.4\,GW/squarem$ at the operating point, roughly $30\times$ the ITER value and the quantity that the divertor must dissipate. The narrow width is a direct consequence of the compact, high-current ST geometry, not a modelling choice.

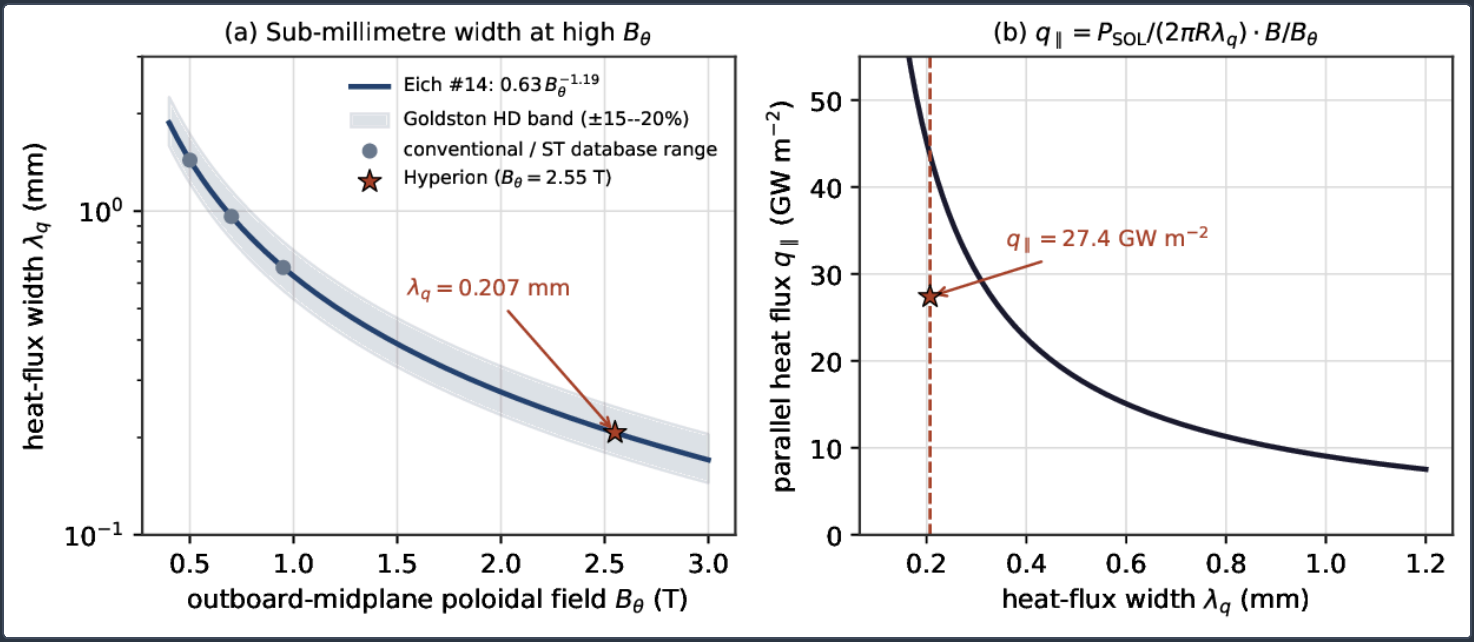

(a) Heat-flux width λ_q versus outboard poloidal field: the Eich regression (line) with the Goldston HD ± 15–20% band, the conventional/ST database range, and the Hyperion point at $B_\theta = 2.55\,T$, $\lambda_q = 0.207\,mm$ (star). (b) Parallel heat flux from equation (4) at fixed $P_{\mathrm{SOL}} = 29\,MW$; the operating point sits at $q_\parallel = 27.4\,GW/squarem$.

THE ATTACHED TARGET LOAD AND THE DISSIPATION REQUIREMENT

Equations (5)–(6) with the target grazing angle ($\sim 3°$) and poloidal flux expansion give the unmitigated attached target load: 1370 MW/squarem on a conventional single-null leg and 171 MW/squarem on a Super-X leg with flux expansion $f_x = 8$ (figure 3). Both are far above the $\sim 10\, MW/squarem$ steady material limit, so a large radiated fraction is mandatory: reaching the limit requires dissipating 99.3% of the parallel power on a conventional leg or 94.2% on a Super-X leg, and the target electron temperature crosses the 5 eV detachment threshold at 91.4% dissipation. Geometry alone (the Super-X flux expansion) buys almost an order of magnitude in target load but does not by itself close the problem; dissipation is required.

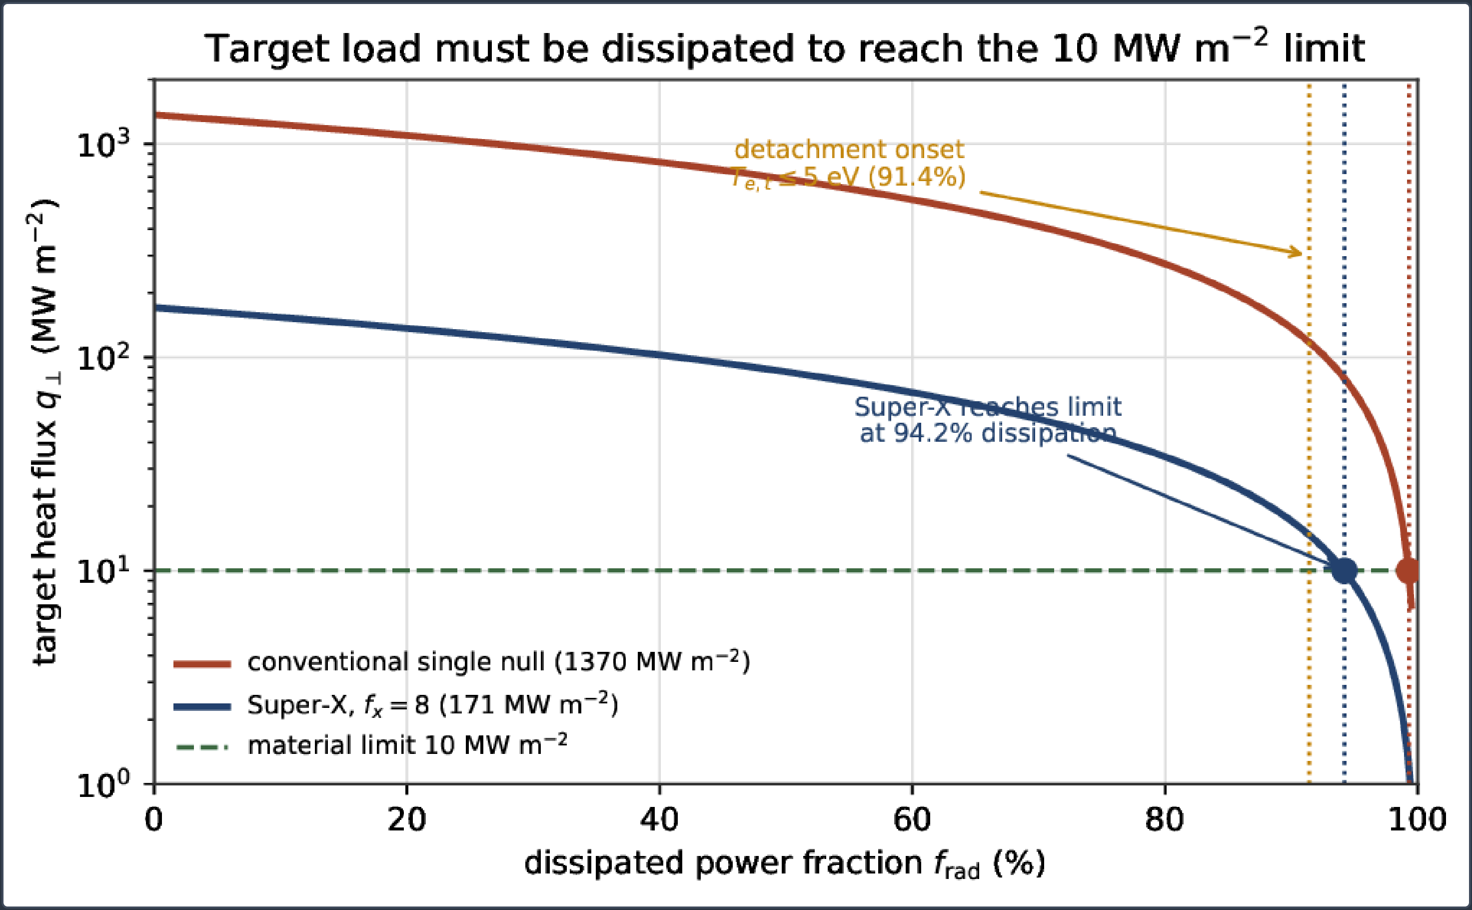

Two-point-model target heat flux versus dissipated power fraction for the conventional single-null (1370 MW/squarem attached) and Super-X ($f_x = 8$, 171 MW/squarem) legs. The 10 MW/squarem material limit is reached at $99.3\%/94.2\%$ dissipation; detachment onset ($T_{e,t} \leq 5\, eV$) is at 91.4% (register HX-14).

CONVERGED DETACHMENT: UEDGE DENSITY SCAN

The reduced-order dissipation requirement is confirmed by a converged 2-D edge-code solve. Table 3 and figure 4 report the UEDGE 8.1.2 three-point upstream-density scan: the target electron temperature collapses monotonically from 14.4 eV (attached) through 5.28 eV (marginal, exactly at the 5 eV threshold) to 1.05 eV at onset, where the 5 eV front reaches the plate. The collapse steepens from $T_{e,t} \propto n_u^{-1.7}$ on the attached branch to $n_u^{-2.8}$ into onset — the two-point $\sim n_u^{-2}$ scaling of equations (5)–(6) plus the added neutral momentum loss that marks detachment. The divertor/upstream compression reaches $n_t/n_u = 31$ at onset, agreeing with the independent 1-D detachment arbiter (≈ 27) to within $\sim 15\%$. Scaled to the design parallel flux, the upstream density required for density-only detachment approaches one-to-two times the Greenwald density and is therefore inaccessible: *density alone cannot detach the breeder leg at design exhaust*, a firm conclusion from a converged plasma solve rather than a scaling assertion.

n_{core} (M^{-3})	$n_{e,u}$ (M^{-3})	$T_{e,u}$ (EV)	$T_{e,t}$ (EV)	n_t/n_u	STATE
1×10^{20}	6.05×10^{19}	106.0	14.4	15.2	attached
2×10^{20}	1.11×10^{20}	84.9	5.28	19.9	marginal (5 eV)
4×10^{20}	1.96×10^{20}	66.7	1.05	31.0	detachment onset

Converged UEDGE 8.1.2 upstream-density scan on config-22021 (register BR-L2-A21). Each row is a true steady state; the target collapses from attached to detachment onset.

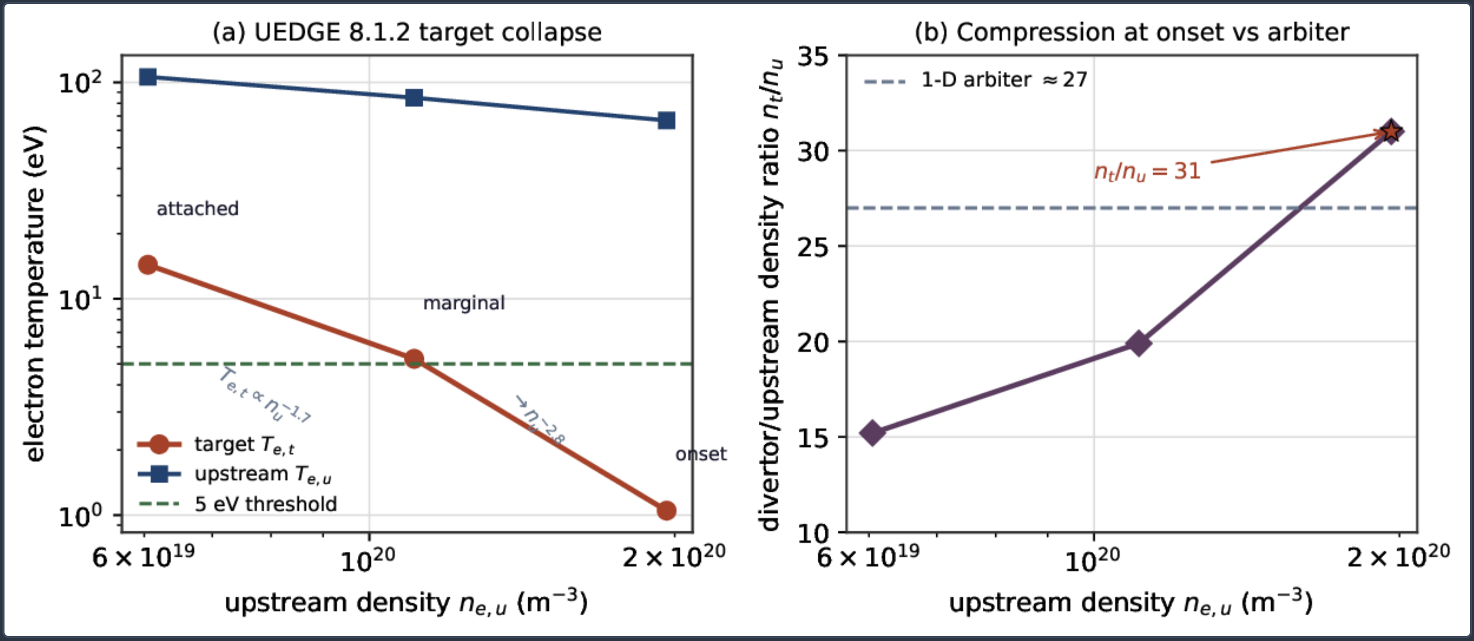

Converged UEDGE density scan (register BR-L2-A21). (a) Upstream and target electron temperature versus upstream density; the target crosses the 5 eV threshold with a $n_u^{-1.7} \to n_u^{-2.8}$ steepening. (b) Divertor/upstream density compression reaching $n_t/n_u = 31$ at onset, within $\sim 15\%$ of the 1-D arbiter.

CONVERGED DETACHMENT: SD1D FLUID-PLUS-NEUTRAL SOLVE

A second, independent converged code — SD1D on BOUT++ — reproduces the same physics with a full recycling- neutral model and ADAS nitrogen radiation. Figure 5 shows the density scan at zero seeding: the attached branch holds the target near 19 eV and 172 MW/squarem up to $\sim 0.4\, n_{\mathrm{GW}}$, then detaches at $\approx 0.6\, n_{\mathrm{GW}}$, where the target temperature collapses to $\sim 1.9\, eV$, the load drops below 20 MW/squarem, and a cold front lifts $0.47\, m$ off the plate. Figure 6 shows the nitrogen-seeding scan at fixed upstream density $0.40\, n_{\mathrm{GW}}$: a sharp detachment transition in which *less than* 0.5% *nitrogen* fully detaches the target ($T_{e,t} \approx 1.9\, eV$, $q_\perp \lesssim 10\, MW/squarem$), with 2–3% holding deep detachment. Both scans confirm the reduced-order picture: the attached load is far above the limit, density alone at accessible values does not protect the target, and a modest nitrogen seed detaches it within the impurity-tolerance budget.

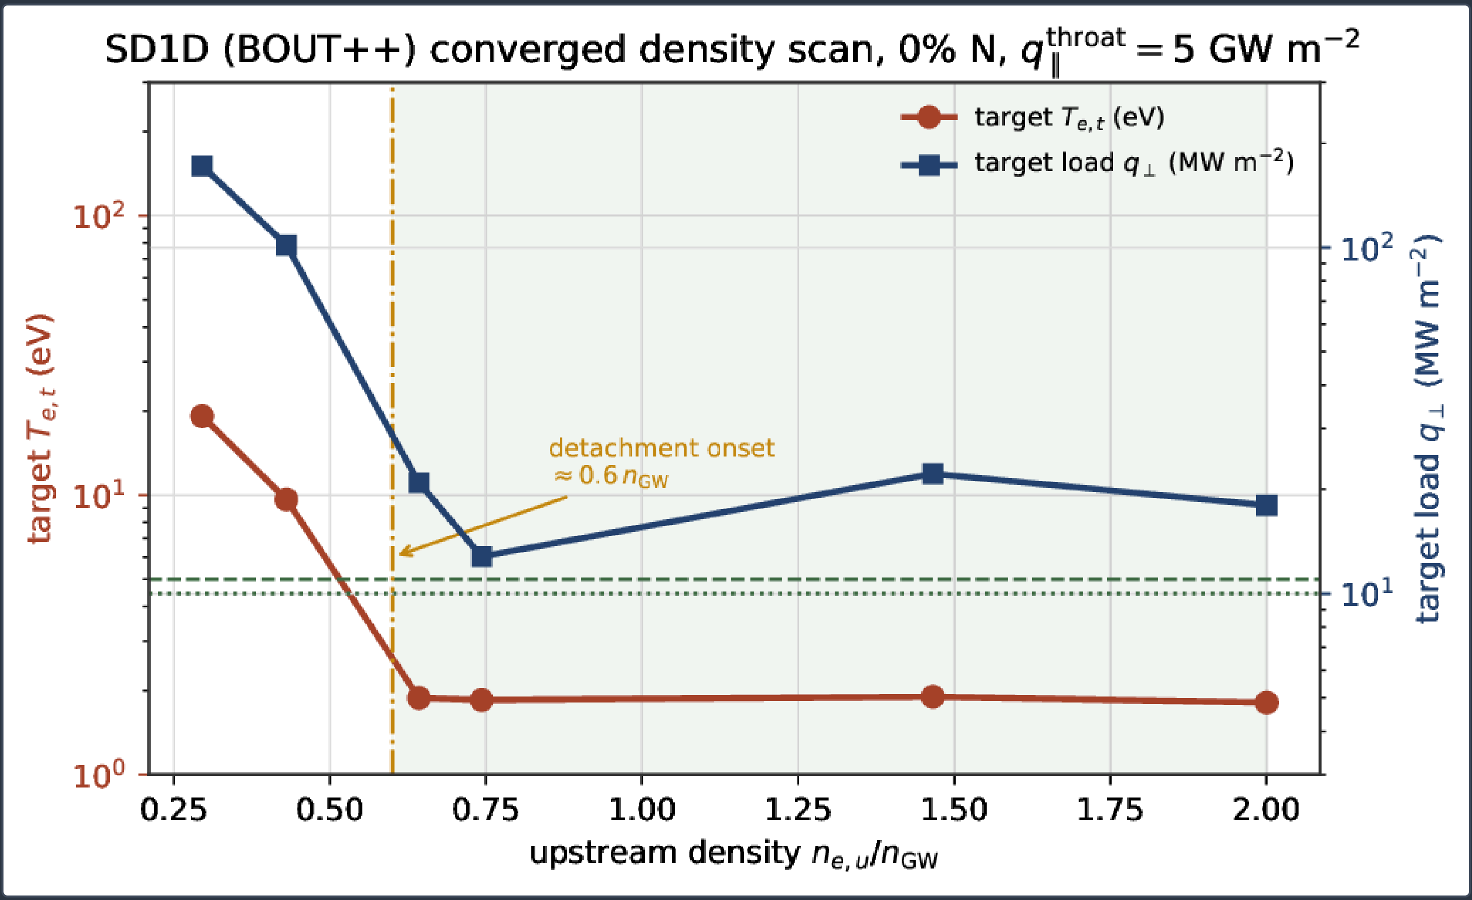

Converged SD1D (BOUT++) density scan at zero seeding (register HX-14b): target electron temperature (left axis) and target load (right axis) versus upstream density in Greenwald units. Detachment onset at $\approx 0.6\, n_{\mathrm{GW}}$ (shaded); the throat parallel flux is 5 GW/squarem (see §6).

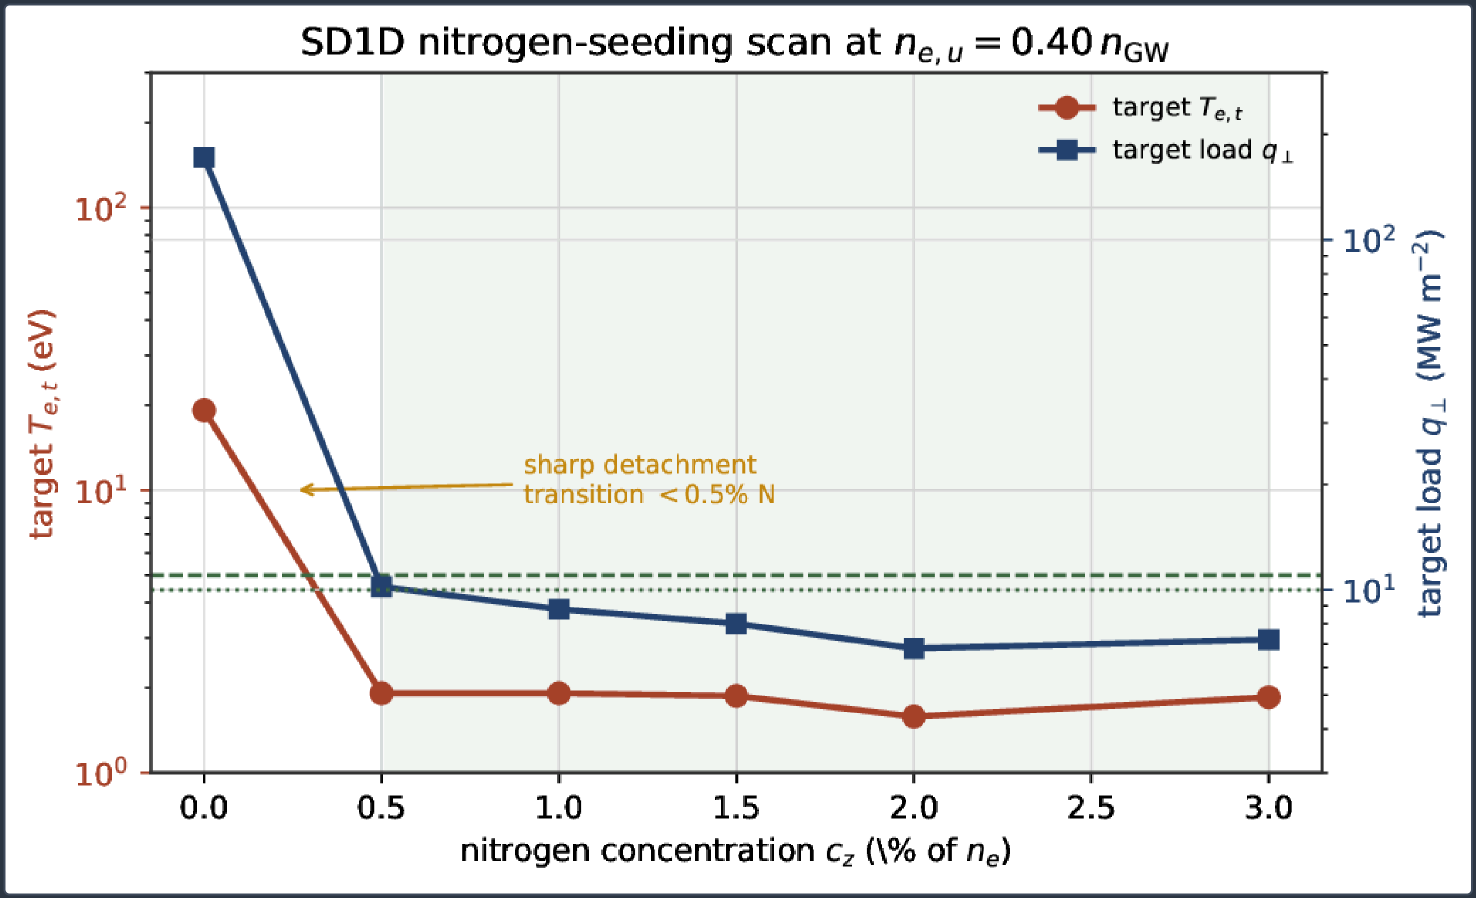

Converged SD1D nitrogen-seeding scan at $n_{e,u} = 0.40\, n_{\mathrm{GW}}$ (register HX-14b): a sharp transition detaches the target at $< 0.5\%$ nitrogen, with the load falling below the 10 MW/squarem limit.

CROSS-METHOD DETACHMENT THRESHOLD

The four methods span the parallel-flux range and agree on the physics (table 4, figure 7). All four give an attached target load far above the material limit — 1370–1401 MW/squarem at the full $q_\parallel$, 180 MW/squarem at the 5 GW/squarem SD1D throat flux — and all four detach the target with a sub-1%-to-\ 1% nitrogen seed: $\sim 0.8\%$ (Lengyel), $\sim 1.0\%$ (self-written 1-D at the full $q_\parallel = 27.4\, GW/squarem$), $< 0.5\%$ (SD1D at the throat flux), with the two-point model requiring 94% dissipation. The absolute seed fraction scales with $q_\parallel$ as equation (7) predicts (the SD1D value is lower because it is evaluated at the lower divertor-throat flux), but the conclusion is method-independent: the seed needed is comfortably inside the impurity budget.

METHOD	CODE	$q_\parallel$ (GW M^{-2})	N FOR FULL DETACHMENT
two-point model	analytic (HX-14)	27.4	94% dissipation
Lengyel integral	analytic (HX-15)	27.4	$\sim 0.8\%$
self-written 1-D	`sol1d.py`	27.4	$\sim 1.0\%$
SD1D	BOUT++	5.0	$< 0.5\%$

Cross-method detachment threshold (register HX-14/HX-14b/HX-15). All four methods agree that the attached target $§I10

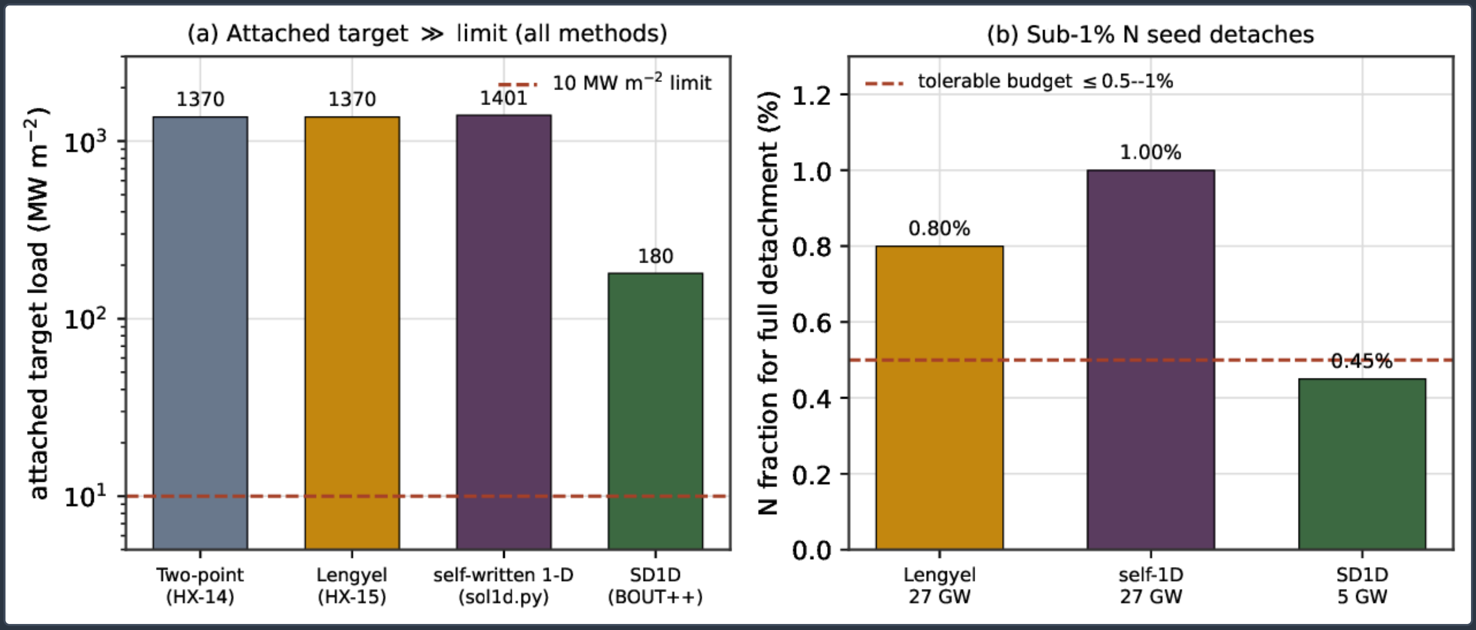

Cross-method verification (register HX-14/14b/15). (a) Attached target load, all methods far above the 10 MW/squarem limit. (b) Nitrogen fraction for full detachment; all methods land inside the tolerable seed budget, with the absolute value scaling with the evaluated $q_\parallel$.

THE SEEDING-AND-GEOMETRY PORTFOLIO

Because seeding is mandatory, the design question is *how much*, and divertor geometry is the lever that lowers the answer (figure 8). A seeded conventional single-null leg detaches at a tolerable fraction of order 6.4%; the Super-X flux expansion relaxes it to the $2–3\%$ range (a single lever does not by itself close the leg), a double-null power split brings it to $\sim 1\%$, and adding a snowflake reaches $\sim 0.5\%$ — within the fraction the core can tolerate without diluting the fusion power. A Monte-Carlo closure study of the stacked geometry (register BR-SX-09, $N = 4000$) finds the stacked Super-X$\times$double- null$\times$snowflake closes at $\leq 0.5\%$ argon with 95% probability. A flowing-lithium vapor-box leg is a second, independent closing path that removes seeding entirely (register BR-L2-A11c; caps the target load $\sim 31 \rightarrow 0.16\, MW/squarem$ with in-loop tritium extraction), documented in the companion flowing-lithium study [35]. The exhaust is therefore not a single-point solution but a portfolio of stacking levers, detailed in the companion detachment-portfolio [33] and divertor-down-select studies [34], and regulated in real time by the companion detachment controller [36]; crucially, the fully-bred divertor cassette ($f_{\text{div}} = 1.0$) that spans the full azimuth is confirmed, so the divertor choice does not reopen the tritium-breeding gate.

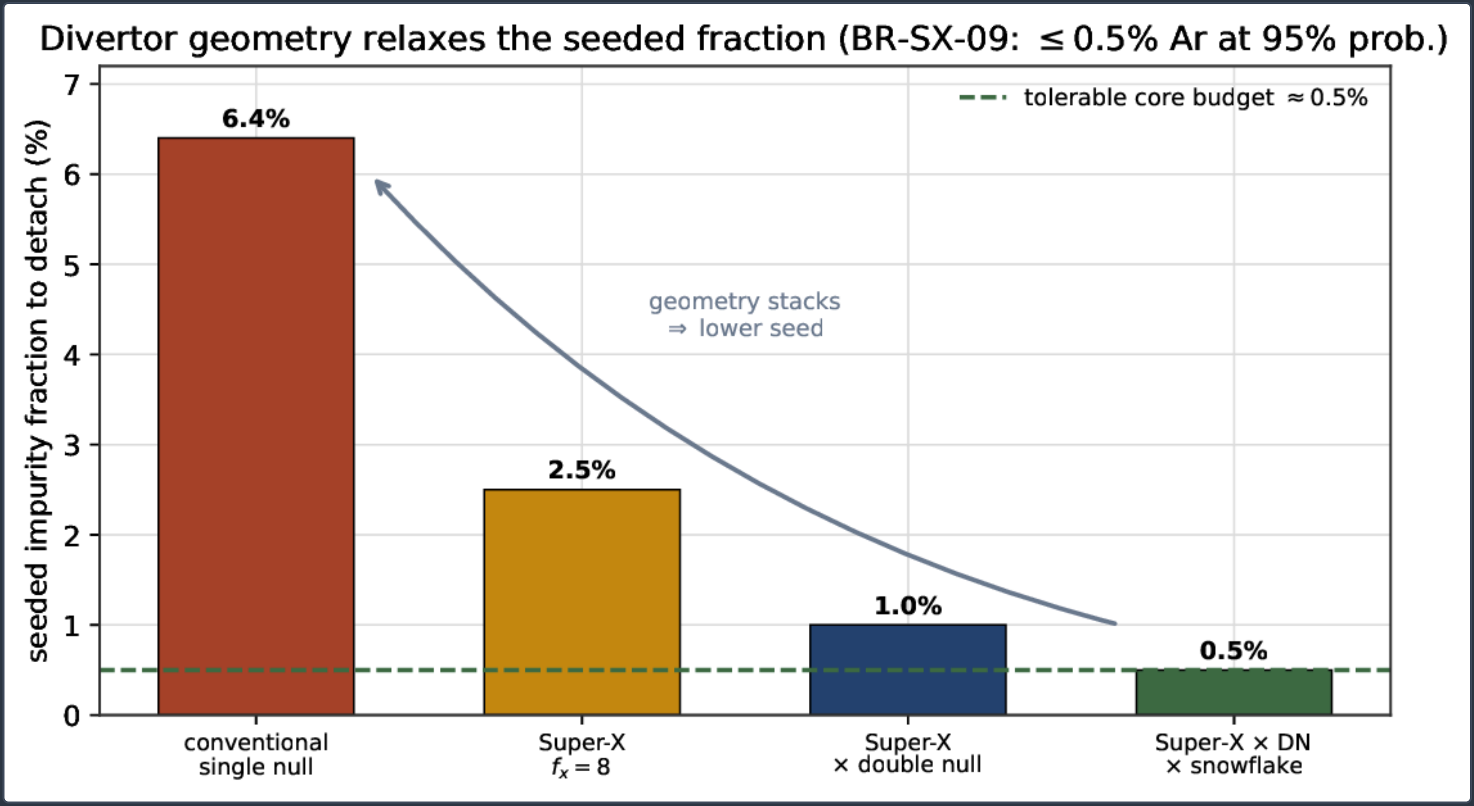

Divertor-geometry levers relax the seeded fraction required to detach, from $\sim 6.4\%$ in a seeded single null to $\sim 0.5\%$ with a stacked advanced divertor (registers BR-L2-A11-SUM, BR-SX-09).

PLASMA–MATERIAL INTERACTION AT THE DETACHED TARGET

The cold, low-temperature target that detachment secures is a benign sputtering environment, confirmed by a first-principles binary-collision calculation (figure 9). Across the RustBCA scan, deuterium on tungsten reaches a peak physical-sputtering yield of 1.35×10^{-2} atoms ion^{-1} — consistent with the established literature value of order 10^{-2} [25] — at the highest scanned energy and grazing angle, with sub-monolayer-per-pulse gross erosion at the design-point target flux. The physically decisive result follows from the sheath boundary condition of §2: at the detached target ($T_{e,t} \approx 1.9\,eV$) the sheath energy $E_i \approx 5T_{e,t} \approx 10\,eV$ is far below the $E_{\text{th}} \approx 209\,eV$ deuterium-on-tungsten threshold, so the operating-point yield collapses many orders of magnitude below the scan peak (figure 9b). Even the marginally attached target ($T_{e,t} \approx 5$–14 eV) sits at or below threshold. Chemical sputtering and prompt redeposition are refinements that lower the net erosion further; the full plasma-material-interaction budget is developed in the companion PMI study [37]. Target erosion is not a life-limiter for the detached NT-edge divertor.

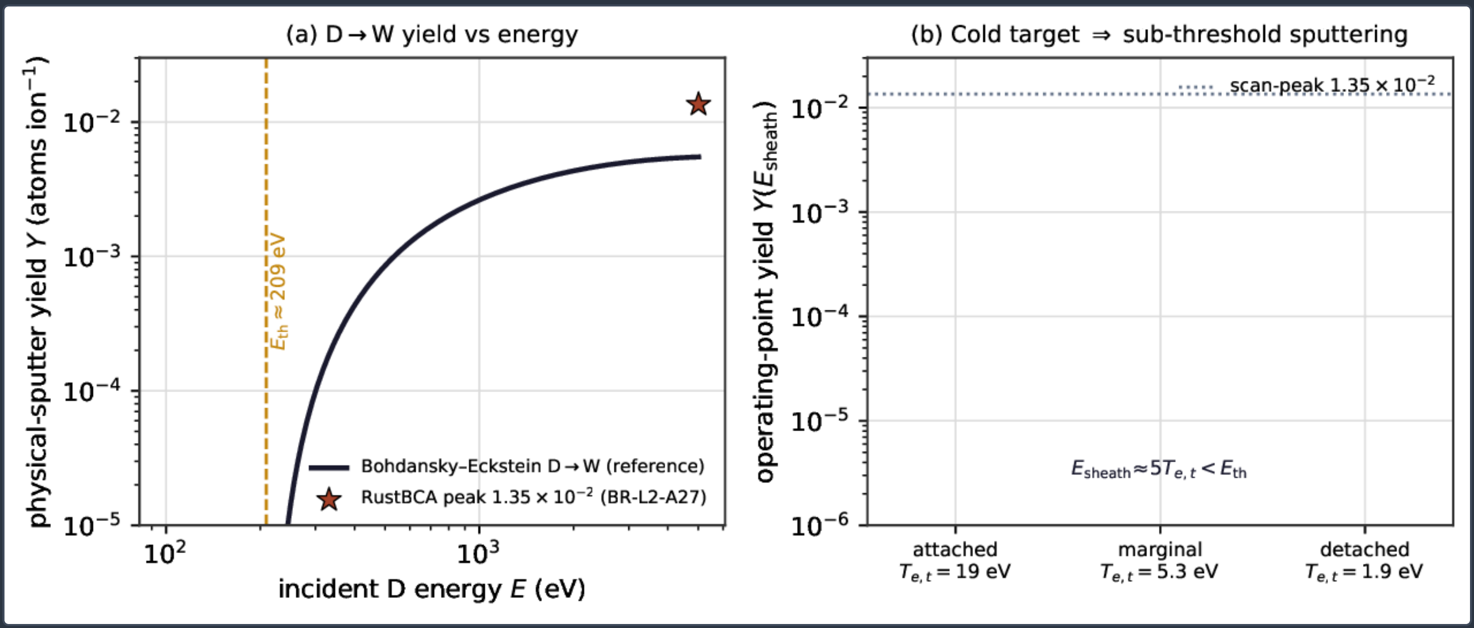

Target sputtering (register BR-L2-A27). (a) Deuterium-on-tungsten physical-sputter yield versus energy: the Bohdansky–Eckstein reference curve anchored to the RustBCA scan peak 1.35×10^{-2} atoms ion^{-1}, with the $E_{\text{th}} \approx 209\,eV$ threshold. (b) Operating-point yield at the sheath energy $E_i \approx 5T_{e,t}$: the cold detached target sits below threshold, collapsing the yield.

THE ELM-FREE EDGE IN THE STABILITY PLANE

Figure 10 shows why the whole analysis is steady-state. In the $(\alpha, \hat{J})$ plane of equation (8), a conventional H-mode pedestal rides the peeling–ballooning boundary and periodically crosses it, driving the ELM cycle. The Hyperion NT operating point sits well inside the stable interior: the inverted shape caps the edge pressure gradient and raises the shear, so there is no pedestal to collapse and no transient to size the divertor. The steady exhaust resolved above is therefore the whole power-handling problem.

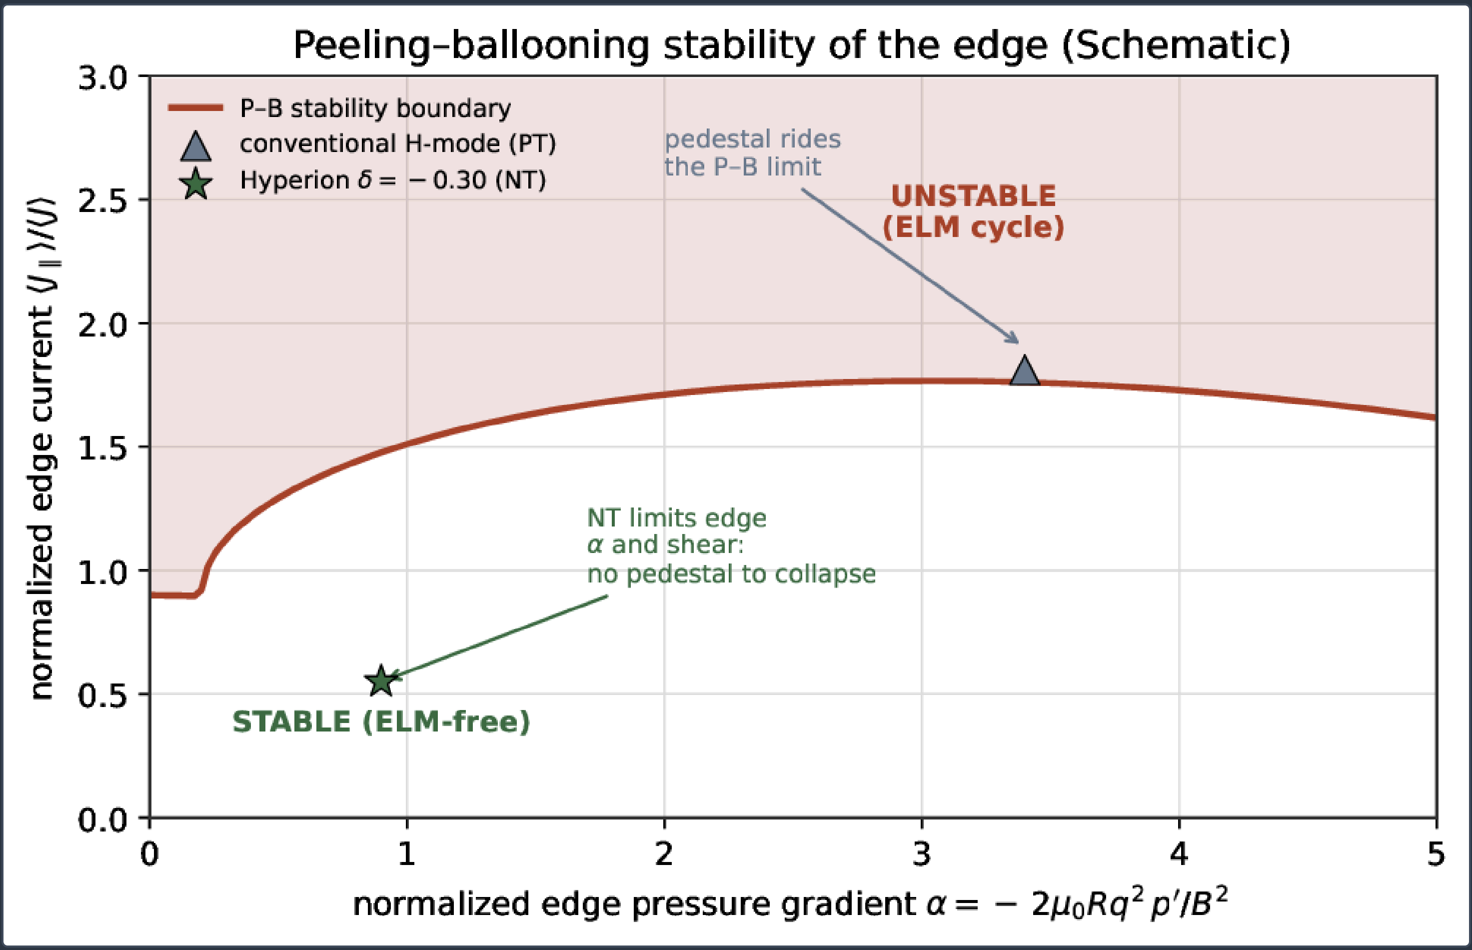

(Schematic.) Peeling–ballooning stability in the normalised edge-pressure-gradient/edge-current plane. A conventional PT pedestal rides the boundary (ELM cycle); the Hyperion $\delta = -0.30$ point sits in the stable interior — ELM-free by construction. The axes and boundary are illustrative; the placement of the two operating points is the content.

§ 5

Discussion

HEAT-FLUX WIDTH AND PARALLEL LOAD AGAINST THE LITERATURE

The computed $\lambda_q = 0.207\,mm$ is a direct evaluation of the Eich regression [9] at the breeder's outboard poloidal field and is bracketed by the Goldston HD model [10]; both are foundational, multi-machine-validated predictors, and the inclusion of spherical tokamaks does not shift the Eich fit, so the extrapolation is well-grounded. The narrow width and the resulting $q_\parallel = 27.4\,GW/squarem$ are not anomalies of the design but the generic signature of a compact, high-current device: they are of the same order as, and consistent with, the values reported for other compact high-field concepts [2] and roughly $30\times$ the ITER divertor value [3]. The design implication — that essentially all of P_{SOL} must be radiated — is shared by every reactor-class tokamak and is the reason the divertor, not the confinement, is the binding edge constraint.

DETACHMENT THRESHOLD AGAINST EXPERIMENT AND MODELLING

The detachment picture reproduces the established phenomenology quantitatively. The $T_{e,t} \simeq 5\,eV$ criterion, the $\sim n_u^{-2}$ collapse steepening at onset, and the target ion-flux rollover are the textbook two-point/detachment signatures [11,12,13]. The UEDGE divertor/upstream compression $n_t/n_u = 31$ at onset agrees with the independent 1-D arbiter (≈ 27) to within $\sim 15\%$ and is in the range measured on detaching tokamak divertors. The seed requirement — sub-1% nitrogen at the divertor-throat flux, $\sim 1\%$ at the full OMP flux — sits inside the impurity budget and is consistent with the impurity-seeding exhaust strategy established from present devices through ITER to DEMO [14,15]. Most directly, the highly-radiative NT pilot-plant analysis of Kuang *et al.* [19] finds that a standard divertor at $P_{\text{SOL}} = 25\,MW$ reaches $\sim 7.8\,MW/squarem$ peak target load with only $\sim 0.13\%$ neon — the same conclusion, on a comparable NT machine, that a modest low-Z seed suffices once the ELM transient is removed. Our Super-X result likewise aligns with the experimental demonstration of detachment access and control on the MAST-U Super-X divertor [18], and the snowflake and Super-X flux-expansion gains with their original analyses [16,17].

THE ROLE OF NEGATIVE TRIANGULARITY

The load-bearing simplification — that the divertor is sized by steady exhaust, not by an ELM — rests on the ELM-free NT edge, which is now an experimentally robust result: strong negative triangularity suppresses all ELM activity across a wide operating range on DIII-D, consistent with the P–B prediction that the NT edge cannot access the ELMing H-mode [6,8,7]. This paper takes that established edge and carries its consequence all the way to the target: because there is no transient, the entire power-handling case reduces to the steady boundary-value problem of §2, which we have closed. The companion peeling–ballooning study [31] establishes the ELM-free edge itself; the present work is its power-handling closure.

§ 6

Limitations

Two honest caveats bound the claims. First, the detachment physics is resolved here in converged 1-D and quasi-1-D solves (UEDGE slab, SD1D, and the independent 1-D solver); the 2-D SOLPS-ITER divertor solve [20], whose deck is built on the config-22021 equilibrium (register BR-L2-A11b-s3, toolchain staged), is the one open confirmation and is the measurement that fixes the *absolute* detachment threshold and the stacked-geometry seed fractions, which are presently reduced-order superpositions rather than 2-D-solved values. Second, the converged SD1D solve was run at a divertor-throat parallel flux of 5 GW/squarem rather than the full OMP peak of 27.4 GW/squarem, because the community code is numerically stiff on this grid at the extreme flux; the absolute seed fraction scales with $q_{\|}$, so the SD1D percentage is not directly comparable in magnitude to the full-flux estimate — the cross-method comparison is of the *physics and trend*, which is robust, not of the absolute percentage. Neither caveat gates the design: the divertor closes with impurity seeding as intended, on multiple independent paths, and the 2-D solve is a fidelity upgrade of an already-consistent result.

§ 7

Conclusion

Under negative triangularity the Hyperion edge is ELM-free by construction, and the power-handling problem reduces to the steady scrape-off layer, which we have posed formally and closed on the design equilibrium. A first-principles heat-flux width $\lambda_q = 0.207\, mm$ (Eich, bracketed by Goldston HD) sets a parallel heat flux $q_{\|} = 27.4\, GW/squarem$ and an unmitigated attached target load of 1370 (conventional) and 171 MW/squarem (Super-X), both far above the material limit. A converged UEDGE solve resolves the density-driven target collapse from 14.4 through 5.3 to 1.05 eV with $n_t/n_u = 31$ at onset (within $\sim 15\%$ of the 1-D arbiter); a converged SD1D solve reproduces the physics and detaches the target with $< 0.5\%$ nitrogen; and four independent methods agree that density alone cannot detach the leg while a sub-1%-to-\ 1% seed can. Divertor geometry walks the required seed fraction from 6.4% in a single null to $\sim 0.5\%$ with a stacked advanced divertor, with a flowing-lithium leg as a seed-free second path. At the cold detached target the deuterium-on-tungsten sputtering yield peaks at 1.35×10^{-2} atoms ion^{-1} and the sheath energy falls below threshold, so target erosion is not a life-limiter. The negative-triangularity breeder edge is quantified, bounded, and controllable, with a 2-D SOLPS-ITER divertor solve as the confirming step.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The authors thank the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and the NVIDIA GPU HPC campaign for the compute that produced the edge-code and BCA solves.

Reproducibility. The heat-flux-width, two-point, detachment, seeding-portfolio and sputtering results are frozen anchors in the Kronos Physics De-Risking Register [30]: the heat-flux width and two-point model (HX-14), the neutral/impurity detachment window (HX-15), the converged UEDGE solve (BR-L2-A21), the converged SD1D solve (HX-14b), the seeding portfolio (BR-L2-A11-SUM, BR-SX-09), and the D →W sputtering (BR-L2-A27), computed with UEDGE 8.1.2, SD1D on BOUT++, and RustBCA on the config-22021 geometry and regenerable from the archived run inputs and the deposited figure-generation script. The register is deposited at doi:10.5281/zenodo.22133057 and this paper is archived at doi:10.5281/zenodo.22645830; official version: https://www.kronosfusionenergy.com/publications/the-negative-triangularity-edge-heat-flux-width-the-elm.html.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645830. The scrape-off-layer scan inputs, the equilibrium the edge is built on, and the figure-generation script are archived with the deposit and reference the Kronos 2026 series. Any economic analysis is reported separately and is not included in 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 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

3 R A Pitts *et al*, Physics basis for the first ITER tungsten divertor, *Nucl. Mater. Energy* **20** 100696 (2019). doi:10.1016/j.nme.2019.100696

4 P B Snyder, H R Wilson, J R Ferron, L L Lao, A W Leonard, T H Osborne, A D Turnbull, D Mossessian, M Murakami and X Q Xu, Edge localized modes and the pedestal: a model based on coupled peeling–ballooning modes, *Phys. Plasmas* **9** 2037 (2002). doi:10.1063/1.1449463

5 A W Leonard, Edge-localized-modes in tokamaks, *Phys. Plasmas* **21** 090501 (2014). doi:10.1063/1.4894742

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 A O Nelson, L Schmitz, C Paz-Soldan, K E Thome, T B Cote, N Leuthold, F Scotti, M E Austin, A Hyatt and T Osborne, Robust avoidance of edge-localized modes alongside gradient formation in the negative triangularity tokamak edge, *Phys. Rev. Lett.* **131** 195101 (2023). doi:10.1103/PhysRevLett.131.195101

9 T Eich *et al*, Scaling of the tokamak near the scrape-off layer H-mode power width and implications for ITER, *Nucl. Fusion* **53** 093031 (2013). doi:10.1088/0029-5515/53/9/093031

10 R J Goldston, Heuristic drift-based model of the power scrape-off width in low-gas-puff H-mode tokamaks, *Nucl. Fusion* **52** 013009 (2012). doi:10.1088/0029-5515/52/1/013009

11 P C Stangeby, *The Plasma Boundary of Magnetic Fusion Devices* (Institute of Physics Publishing, Bristol, 2000).

12 A Loarte *et al*, Chapter 4: Power and particle control, *Nucl. Fusion* **47** S203 (2007). doi:10.1088/0029-5515/47/6/S04

13 S I Krasheninnikov and A S Kukushkin, Physics of ultimate detachment of a tokamak divertor plasma, *J. Plasma Phys.* **83** 155830501 (2017). doi:10.1017/S0022377817000654

14 A Kallenbach *et al*, Impurity seeding for tokamak power exhaust: from present devices via ITER to DEMO, *Plasma Phys. Control. Fusion* **55** 124041 (2013). doi:10.1088/0741-3335/55/12/124041

15 R J Goldston, M L Reinke and J A Schwartz, A new scaling for divertor detachment, *Plasma Phys. Control. Fusion* **59** 055015 (2017). doi:10.1088/1361-6587/aa5e6e

16 P M Valanju, M Kotschenreuther, S M Mahajan and J Canik, Super-X divertors and high power density fusion devices, *Phys. Plasmas* **16** 056110 (2009). doi:10.1063/1.3110984

17 D D Ryutov, Geometrical properties of a “snowflake” divertor, *Phys. Plasmas* **14** 064502 (2007). doi:10.1063/1.2738399

18 K Verhaegh *et al*, Spectroscopic investigations of detachment on the MAST Upgrade Super-X divertor, *Nucl. Fusion* **63** 016014 (2023). doi:10.1088/1741-4326/aca10a

19 A Q Kuang *et al*, Power handling in a highly-radiative negative triangularity pilot plant, *Plasma Phys. Control. Fusion* **66** 125001 (2024). doi:10.1088/1361-6587/ad867a

20 S Wiesen *et al*, The new SOLPS-ITER code package, *J. Nucl. Mater.* **463** 480 (2015). doi:10.1016/j.jnucmat.2014.10.012

21 T D Rognlien, J L Milovich, M E Rensink and G D Porter, A fully implicit, time-dependent 2-D fluid code for modeling tokamak edge plasmas, *J. Nucl. Mater.* **196–198** 347 (1992).

22 B D Dudson, J Allen, T Body, B Chapman, C Lau, L Townley, D Moulton, J Harrison and B Lipschultz, The role of particle, energy and momentum losses in 1D simulations of divertor detachment, *Plasma Phys. Control. Fusion* **61** 065008 (2019). doi:10.1088/1361-

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

The Negative-Triangularity Edge: Heat-Flux Width, the ELM-Free Pedestal, and the Scrape-Off-Layer Boundary Condition for Power Handling in a Compact Spherical-Tokamak Breeder

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

Zenodo · doi:10.5281/zenodo.22645830

Read on kronosfusionenergy.com

The Physics De-Risking Register