KRONOS FUSION ENERGY

PAPER 2.14 · THE KRONOS 2026 SERIES

The Greenwald-Fraction Density Limit and the Radiative Boundary: Mapping the Safe Operating Envelope of a Compact Spherical-Tokamak Breeder

A fusion power plant lives inside a bounded operating window, and the density limit is one of its walls. We pose the operating-window boundary of the Hyperion breeder — a compact...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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 Greenwald-Fraction Density Limit and the Radiative Boundary: Mapping the Safe Operating Envelope of a Compact Spherical-Tokamak Breeder

A fusion power plant lives inside a bounded operating window, and the density limit is one of its walls. We pose the operating-window boundary of the Hyperion breeder — a compact negative-triangularity spherical tokamak ($I_p = 9.66$ MA, $R_0 = 1.20$ m, aspect ratio $A = 2.5$, minor radius $a = 0.48$ m, elongation $\kappa = 2.0$, triangularity $\delta = -0.30$, $B_0 = 8$ T) — as a formal feasibility problem and map its density axis. From the geometry the Greenwald density is $n_{\rm GW} = I_p/(\pi a^2) = 1.33 \times 10^{21}\ {\rm m}^{-3}$; the design point operates at a Greenwald fraction $f_G \approx 0.35$, a line-averaged density $\bar{n}_e \approx 4.7 \times 10^{20}\ {\rm m}^{-3}$ and a density headroom of a factor ~ 2.9 below the empirical limit. We derive the four walls of the window — the Greenwald/radiative density boundary, the ideal-MHD β limit, the scrape-off-layer detachment boundary and the radiative-collapse floor — as explicit inequalities, and we solve the enclosed feasible set on a nine-point corner grid: the fusion gain holds a window $Q = 2.735$–3.460 ($P_{\rm fus} = 58.6$–118.9 MW) with every physics gate satisfied, enclosing the frozen design point $Q = 3.076$, $P_{\rm fus} = 85.04$ MW. Because the geometry is fixed the neutron source strength and the first-wall loading are a strict one-to-one invariant, so density is a controllable lever inside the window ($Q = 3.865$ at $0.967\times$ the reference wall loading) rather than a cliff. Edge-transport solves place the high-density boundary at divertor detachment access near $0.6\,n_{\rm GW}$ — reached with a sub-one-percent nitrogen seed that dissipates $\sim 94\%$ of the exhaust power and holds the target heat flux within the plasma-facing-component limit — well above the operating point, so the practical ceiling is a controlled radiative boundary and not a disruptive one. All quantitative results are frozen anchors reproduced on an NVIDIA GPU HPC campaign spanning a reduced-order transport engine, the UEDGE and SD1D/BOUT++ edge codes and the SOLPS-ITER reference edge code. The result is a quantified, gate-by-gate safe-operating envelope with a wide density margin on the design point.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645947 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: Greenwald density limit · operating window · negative triangularity · detachment · radiative boundary · 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)

Greenwald density limit; operating window; negative triangularity; detachment; radiative boundary; spherical-tokamak breeder

§ 1

Introduction

THE OPERATING WINDOW AND WHY THE DENSITY AXIS MATTERS

Every magnetically confined fusion device operates inside a bounded region of parameter space. Push the density too high and the discharge terminates against the empirical density limit; push the pressure too high and it crosses the ideal-magnetohydrodynamic (MHD) β limit; let the edge radiate too little and the divertor overheats; let it radiate too much and the plasma suffers a radiative collapse. For a power plant the engineering question is not where these walls lie in the abstract but where the *design point* sits relative to them, and how much of the window is available as controllable headroom. This paper poses the operating-window boundary of the Hyperion breeder as a formal feasibility problem, derives its four walls as explicit inequalities, and maps the design point′s position along the density axis — the axis governed by the density limit, the radiative boundary and divertor detachment.

The density limit is the most robust of the empirical operational boundaries. Since Greenwald′s compilation [1,2] the maximum line-averaged density accessible to a tokamak has been found to scale, to within the scatter of the multi-machine database, as the *Greenwald density*

$$n_{\mathrm{GW}}\,[10^{20}\,\mathrm{m}^{-3}] = \frac{I_p\,[\mathrm{MA}]}{\pi\,a^2\,[\mathrm{m}^2]},$$

depending only on the plasma current I_p and the minor radius a and, strikingly, not on the heating power, the magnetic field, or the isotope. Exceeding n_{GW} in an inductive discharge drives a chain of edge cooling, current contraction and a $2/1$ tearing mode to disruption. The empirical robustness of 1 across four decades of experiments makes the Greenwald fraction $f_G = \bar{n}_e/n_{\mathrm{GW}}$ the natural coordinate for the density axis of any new design′s operating window.

PRIOR ART: WHAT SETS THE DENSITY LIMIT

The physics behind 1 has been the subject of a long line of work, and the modern picture is that the limit is a radiative–transport boundary rather than a hard cliff. Radiative-condensation models trace the limit to a loss of thermal equilibrium at the plasma edge: when impurity line radiation, which scales as n_e^2, can no longer be balanced by the heating power delivered to a cooling edge layer, the edge collapses into a MARFE (multifaceted asymmetric radiation from the edge) [7] and the current channel contracts. Gates and Delgado-Aparicio [3] showed that the onset of radiation-driven magnetic islands, when combined with a cylindrical model of the current channel, reproduces the empirical I_p/a^2 scaling. Zanca *et al.* [4] formalised the boundary as the disappearance of a magnetic- and thermal-equilibrium solution of a power-balance model, and Giacomin *et al.* [5] derived a first-principles scaling from the increase of boundary turbulent transport with edge collisionality, predicting for reactor-scale plasmas a larger safety margin than the raw Greenwald value in the event of a confinement back-transition. The common thread is that the limit is a *power-balance* boundary: it moves with heating power and edge radiation, and the fractional margin f_G to 1 is a conservative lower bound on the true accessible margin [6,2].

The high-density boundary is coupled to the exhaust problem. As the upstream density rises the divertor target cools; below a few electron-volts the target enters detachment, in which volumetric recombination and momentum loss detach the plasma from the plate and cap the target heat flux [12,13]. Detachment is the *intended* high-density operating regime of a reactor divertor, accessed with seeded impurity radiation [11]; the scrape-off-layer (SOL) heat-flux width that sets the unmitigated load follows the drift-based and multi-machine scalings of Goldston [8] and Eich *et al.* [9]. For a compact, high-power-density spherical tokamak the exhaust channel is narrow and the parallel heat flux is severe [10,15], so where the detachment boundary sits relative to the density limit — and relative to the operating point — is a first-order design question.

CONTRIBUTION

This paper contributes a quantified, gate-by-gate map of the Hyperion breeder's operating window along its density axis, cast as a formal feasibility problem and solved on a reproducible GPU HPC campaign. Specifically we: (i) derive the four walls of the window — the Greenwald/radiative density boundary, the ideal-MHD β limit, the SOL detachment boundary and the radiative-collapse floor — as explicit inequalities from a zero-dimensional power balance, a radiative-thermal-equilibrium criterion, and a two-point SOL model (§2); (ii) solve the enclosed feasible set on a nine-point corner grid and show the design point sits inside a fusion-gain window of $Q = 2.735$–3.460 rather than at its edge (§4); (iii) establish that source strength and first-wall loading are a strict one-to-one invariant at fixed geometry, so density is a controllable lever; and (iv) place the high-density boundary at detachment access near $0.6\, n_{\mathrm{GW}}$, reached with a sub-one-percent seed, well above the $f_G \approx 0.35$ operating point. The design point operates at a factor ~ 2.9 below the density limit with every wall bracketed. The reduced-order engine, the corner grid and the detachment onset are frozen anchors in the Kronos Physics De-Risking Register [29] and are regenerable from the archived scripts. The result complements the negative-triangularity confinement and stability results of the Kronos series [31,32,33] and the divertor-exhaust and detachment-control results [34,35,36], and situates the density-limit question within the platform mission of Hyperion as a compact strategic-materials source [30].

§ 2

Governing equations and the formal operating-window problem

We treat the operating window as a feasible set in the two-dimensional control plane (f_G, T_{i0}) of Greenwald fraction and on-axis ion temperature, at fixed geometry and triangularity. Each physics limit is an inequality $g_k(f_G, T_{i0}) \leq 0$; the window is the intersection

$$\Omega \;=\; \{(f_G, T_{i0}) : g_k(f_G, T_{i0}) \leq 0,\; k = 1, \ldots, K\},$$

and the design task is to certify that the frozen design point $x_\star = (f_G^\star, T_{i0}^\star)$ lies in the interior of Ω with a stated margin $\min_k |g_k(x_\star)|$. Figure 1 sketches the window and its four walls. The remainder of this section derives the g_k.

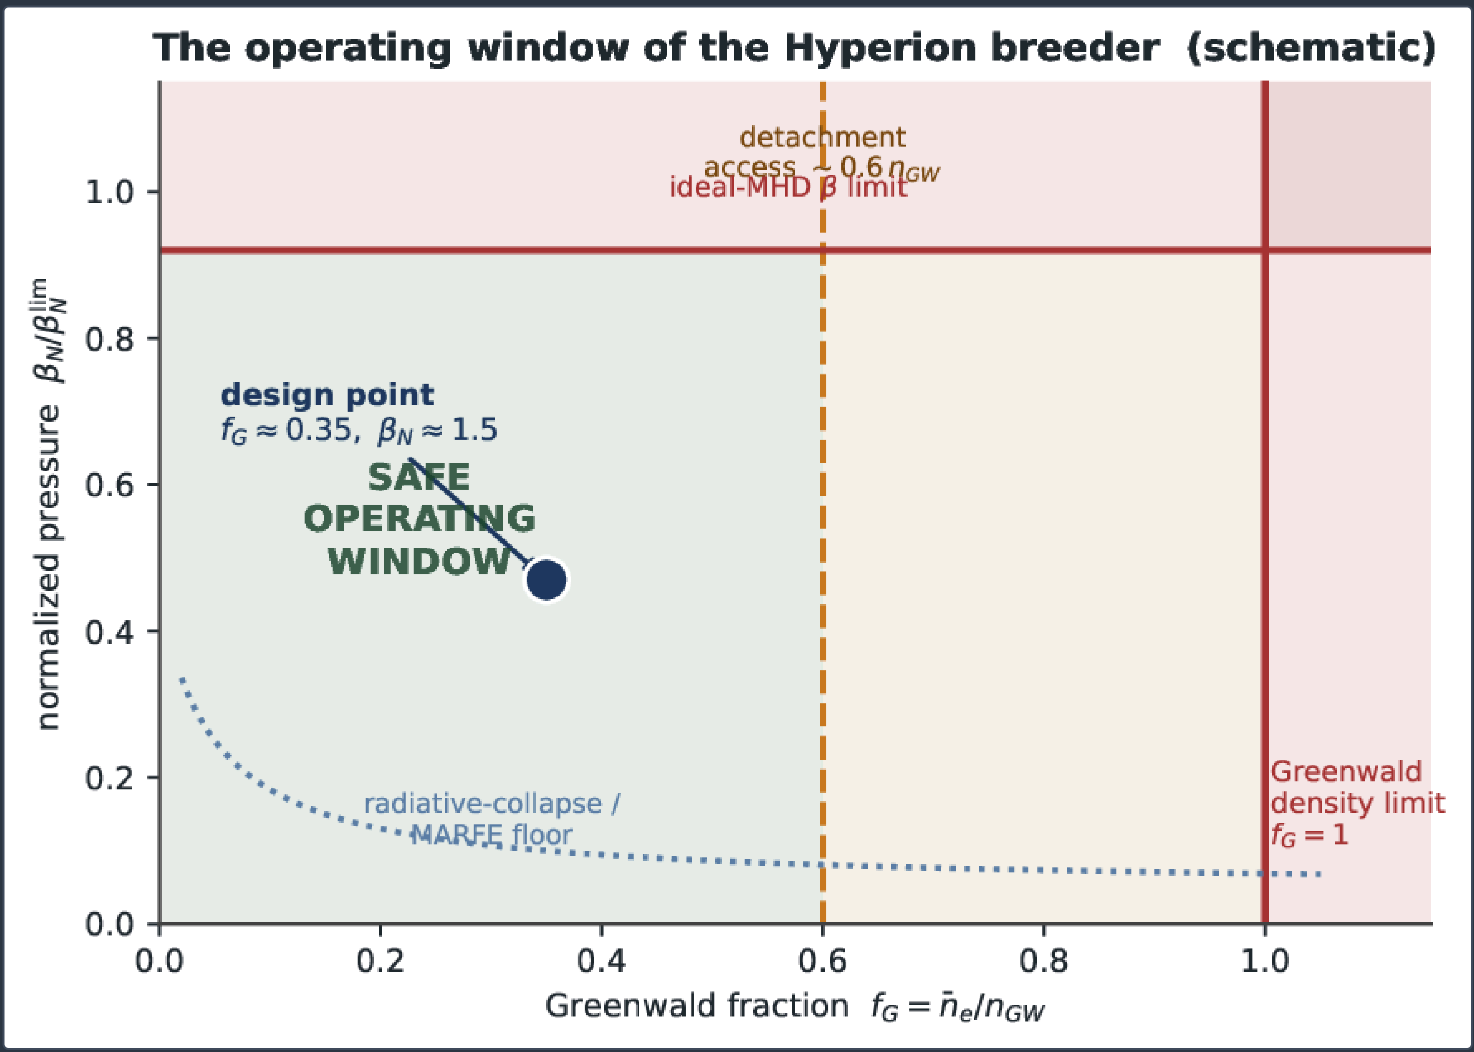

(Schematic.) The operating window of the Hyperion breeder in the density–pressure plane. The four walls are the Greenwald/radiative density limit ($f_G = 1$), the detachment-access boundary near $0.6\,n_{\mathrm{GW}}$, the ideal-MHD β limit, and the radiative-collapse (MARFE) floor. The design point sits at $f_G \approx 0.35$, $\beta_N \approx 1.5$, in the interior of the feasible set defined by equation 2.

THE GREENWALD DENSITY AND THE GREENWALD FRACTION

The Greenwald density 1 follows from the geometry alone. With the Hyperion parameters $I_p = 9.66\,\mathrm{MA}$ and $a = R_0/A = 1.20/2.5 = 0.48\,\mathrm{m}$,

$$n_{\mathrm{GW}} = \frac{I_p}{\pi a^2} = \frac{9.66}{\pi\,(0.48)^2} \times 10^{20} = 13.35 \times 10^{20}\,\mathrm{m}^{-3} = 1.33 \times 10^{21}\,\mathrm{m}^{-3}.$$

The plasma current itself is set by the field, aspect ratio and safety factor through the cylindrical estimate $I_p \simeq 2\pi a^2 B_0 (1+\kappa^2)/(2\mu_0 R_0 q_{95})$, and is sustained without a central solenoid by an aligned bootstrap fraction and off-axis current drive; here we take I_p as the frozen anchor. The density wall is the inequality

$$g_1: \quad f_G \equiv \frac{\bar{n}_e}{n_{\mathrm{GW}}} - 1 \;\le\; 0.$$

The design point operates at $f_G^\star \approx 0.35$, i.e. $\bar{n}_e \approx 4.7 \times 10^{20}\,\mathrm{m}^{-3}$, so $g_1(\quad x_\star) \approx -0.65$: a factor $n_{\mathrm{GW}}/\bar{n}_e = 1/f_G \approx 2.9$ of density headroom.

ZERO-DIMENSIONAL POWER BALANCE AND THE FUSION-GAIN SURFACE

The operating point is the steady state of the volume-averaged energy balance

$$\frac{3}{2}\frac{\mathrm{d}(\bar{n}\bar{T})}{\mathrm{d}t} = P_{\mathrm{aux}} + P_\alpha - \frac{W}{\tau_E} - P_{\mathrm{rad}}, \qquad W = \tfrac{3}{2}\int n_e (T_e + T_i)\,\mathrm{d}V,$$

where $P_\alpha = \frac{1}{5}P_{\mathrm{fus}}$ is the charged-particle heating for a deuterium–tritium plasma, τ_E is the energy confinement time taken from an H-factor-scaled confinement law, and P_{rad} is the total radiated power. The fusion power density is quadratic in density,

$$p_{\mathrm{fus}} = \tfrac{1}{4}\, n_e^2\, \langle\sigma v\rangle(T_i)\, E_{\mathrm{fus}}, \qquad P_{\mathrm{fus}} = \int p_{\mathrm{fus}}\,\mathrm{d}V,$$

for a 50:50 fuel mix with $E_{\mathrm{fus}} = 17.6\,\mathrm{MeV}$, and the fusion gain and Lawson condition are

$$Q = \frac{P_{\mathrm{fus}}}{P_{\mathrm{heat}}}, \qquad n_e T_i \tau_E \;\gtrsim\; \left(n T \tau_E\right)_{\mathrm{ign}}.$$

Solving 5–7 in steady state on the control plane (f_G, T_{i0}) defines the gain surface $Q(f_G, T_{i0})$. Because $\langle\sigma v\rangle$ is a steeply increasing function of T_i up to $\sim 70\,\mathrm{keV}$ and $P_{\mathrm{fus}} \propto n_e^2$, the surface is smooth and monotone in the neighbourhood of the design point; the operating-window constraint is the requirement that the machine remains a net gain source,

$$g_2: \quad Q_{\min} - Q(f_G, T_{i0}) \le 0, \qquad Q_{\min} = 1.$$

At the frozen anchor the balance closes at $Q^\star = 3.076$, $P_{\mathrm{fus}}^\star = 85.04\,\mathrm{MW}$, neutron power $P_n^\star = 67.79\,\mathrm{MW}$ and neutron source strength $S_n^\star = 3.009 \times 10^{19}\,\mathrm{s}^{-1}$.

THE RADIATIVE BOUNDARY: THE DENSITY LIMIT AS A LOSS OF THERMAL EQUILIBRIUM

The physical density limit is not 4 directly but the loss of a stable thermal-equilibrium solution at the plasma edge. Consider a cooling edge layer of volume V_{edge} heated by a power P_{edge} and cooled by impurity line radiation with cooling function $L_z(T_e)$ and impurity concentration c_z. The local power balance is

$$P_{\mathrm{edge}}(T_e) \;=\; c_z\, \bar{n}_e^2\, L_z(T_e)\, V_{\mathrm{edge}},$$

whose right-hand side scales as $\bar{n}_e^2$. As $\bar{n}_e$ rises, the radiative sink grows quadratically while the conducted heating grows more slowly; above a critical density the two curves no longer intersect at a stable (high-temperature) root and the edge undergoes radiative collapse. Linearising 9 about the edge temperature, the stable branch is lost when

$$\frac{\partial}{\partial T_e}\left[P_{\mathrm{edge}} - c_z \bar{n}_e^2 L_z V_{\mathrm{edge}}\right] = 0 \implies \bar{n}_{e,\mathrm{crit}} \;\propto\; \sqrt{\frac{P_{\mathrm{edge}}}{c_z \langle L_z\rangle V_{\mathrm{edge}}}},$$

recovering the power-balance density limit of [3,4,5]: a boundary that increases with heating power and decreases with impurity content. The radiative-collapse floor is the complementary constraint that the core not be over-radiated,

$$g_3: \quad f_{\rm rad}(f_G, T_{i0}) - f_{\rm rad}^{\max} \le 0, \qquad f_{\rm rad} \equiv \frac{P_{\rm rad}}{P_{\rm heat}},$$

with $f_{\rm rad}^{\max}$ below unity to preserve the burn. In the negative-triangularity edge the impurity screening is favourable and the collisional baseline is benign, so 11 is comfortably satisfied at the operating point; the full collisional-radiative closure of $L_z(T_e)$ is treated in the companion impurity-radiation study [36].

THE SCRAPE-OFF-LAYER TWO-POINT MODEL AND DETACHMENT ONSET

The high-density boundary that matters operationally is the divertor exhaust. The parallel heat flux entering the SOL is

$$q_\parallel = \frac{P_{\rm SOL}}{2\pi R\,\lambda_q \sin\theta},$$

where λ_q is the heat-flux width from the drift/multi-machine scalings [8,9] and θ is the field-line pitch at the target. The upstream and target conditions are related by the two-point model [12,13]. Conduction along the field gives the upstream temperature

$$T_u = \left(T_t^{7/2} + \frac{7}{2}\,\frac{q_\parallel L_\parallel}{\kappa_{0e}}\right)^{2/7} \approx \left(\frac{7\,q_\parallel L_\parallel}{2\,\kappa_{0e}}\right)^{2/7} \quad (T_t \ll T_u),$$

with κ_{0e} the Spitzer–H\"arm conductivity coefficient and $L_\parallel$ the connection length; pressure balance with a momentum-loss factor $f_{\rm mom} \in [0,1]$ and total-pressure conservation gives

$$2\,n_t T_t = (1 - f_{\rm mom})\,n_u T_u, \qquad T_t = \frac{m_i}{2e\,\gamma^2}\left(\frac{q_\parallel}{e\,n_u T_u}\right)^2 (1 - f_{\rm mom})^2,$$

with $\gamma \simeq 7$ the sheath transmission coefficient. Equation 14 shows the target temperature falls as $T_t \propto q_\parallel^2/(n_u^2 T_u) \propto n_u^{-2}$ at fixed exhaust power — the collapse observed in the edge-code scan of §4. Detachment sets in when T_t falls to a few electron-volts, where volumetric recombination and charge-exchange momentum loss detach the target ($n_t/n_u \gg 1$). Solving the target-collapse condition $T_t = T_{\rm det}$ for the upstream density yields the detachment-access scaling

$$n_{u,\rm det} \;\propto\; q_\parallel^{5/7}\,L_\parallel^{2/7},$$

which for the Hyperion exhaust places density-only detachment above the Greenwald limit — inaccessible without seeding — but brings the boundary down to $\sim 0.6\,n_{\rm GW}$ once a modest impurity seed is added (below). The detachment-access constraint is

$$g_4: \quad f_G - f_G^{\rm det}(c_z) \le 0, \qquad f_G^{\rm det} \approx 0.6 \text{ at } c_z \lesssim 1\%.$$

IMPURITY-SEEDED DETACHMENT: THE LENGYEL REQUIREMENT

The seed concentration required to detach follows from integrating the radiated power along the field. With $n_e T_e \simeq \text{const}$ across the radiating layer, the Lengyel integral gives the parallel heat flux consumed by radiation,

$$q_{\parallel,u}^2 \;=\; q_{\parallel,t}^2 + 2\,\kappa_{0e}\,c_z\,(n_u T_u)^2 \int_{T_t}^{T_u} L_z(T)\,\sqrt{T}\,{\rm d}T,$$

so full dissipation ($q_{\parallel,t} \to 0$, target detached) requires the seed fraction

$$c_z \;\ge\; \frac{q_{\parallel,u}^2}{2\,\kappa_{0e}\,(n_u T_u)^2 \displaystyle\int_{T_t}^{T_u} L_z \sqrt{T}\,{\rm d}T}.$$

Evaluated for the Hyperion exhaust with a nitrogen seed and an ADAS/`radas` cooling curve, 18 gives $c_z \lesssim 1\%$ for the $\sim 94\%$ radiated fraction that holds the target heat flux within the plasma-facing- component limit, consistent with the collisional-radiative closure [36] and the stacked detachment portfolio [34]. Detachment is thus a controllable operating choice, actuated in real time by the strike-point and seeding controller [35], not a passive limit.

THE WINDOW AS A FEASIBLE SET

Assembling 4, 8, 11 and 16 gives the four walls of Ω. The β wall is a fifth constraint, the ideal-MHD Troyon-type limit $\beta_N \leq \beta_N^{\text{lim}}$ certified in the free-boundary equilibrium and stability study [32]; it is bracketed at the operating point and is not the binding density-axis constraint, so we carry it as a passive gate. The certification task is to show $x_\star \in \textbf{int}\,\Omega$ with margin — which §4 does numerically.

§ 3

Computational methodology: the NVIDIA GPU HPC campaign

Every quantitative result in this paper was produced or reproduced on the Kronos NVIDIA GPU HPC campaign, a verification-first computing programme that re-runs each design gate in an independent code and stamps the result against a frozen anchor. The campaign spans three latency tiers of silicon — microsecond field-programmable-gate- array edge logic, millisecond local GPU twins, and an offline GPU HPC back end on NVIDIA A100/H100-class accelerators with an A10 pool for lighter jobs — with the heavy first-principles work run on the A100/H100 platform and the reduced-order and analytic gates on an in-house 56-thread node. No cost or economic quantity enters this paper; the campaign is described purely as the compute provenance of the numbers.

THE REDUCED-ORDER TRANSPORT ENGINE AND THE CORNER-GRID SOLVE

The operating-window map is produced by the KRONOS reduced-order engine, a deterministic zero/one-dimensional solver that closes 5–7 at a fixed point by Newton iteration on $(\bar{n}_e, \bar{T}_e, \bar{T}_i)$ with an H-factor-scaled confinement law and a collisional-radiative radiation term. The engine evaluates, for each control point (f_G, T_{i0}), the gain Q, the fusion and neutron powers, the wall loading, and the physics-gate flags (coil, density, β, duty). Two campaigns drive the present results:

the design-closure corner map (register serial KX-L1-A1): a 3×3 grid in $(f_G, T_{i0}) = (0.27, 0.30, 0.33) \times (14, 15, 16)\,\mathrm{keV}$ at fixed $\delta = -0.30$, evaluated with a deterministic seed for byte-reproducibility, returning the gain window $Q = 2.735$–3.460 and $P_{\mathrm{fus}} = 58.6$–$118.9\,\mathrm{MW}$;
the mission-point re-optimisation (register serial BR-L1-A2): a 99-point scan of $f_G \in [0.25, 0.35], T_{i0} \in [13, 17]\,\mathrm{keV}$ around the anchor, returning the source-strength/wall-loading invariant. itemize
Uncertainty is propagated by quasi-/Monte-Carlo sampling over the input ranges with the standard-error estimator $\mathrm{SE} = \sigma/\sqrt{N}$ and a Sobol variance decomposition $S_i = \mathrm{Var}_{x_i}(E_{x_{\sim i}}[g\,|\,x_i])/\mathrm{Var}(g)$ to rank the dominant drivers; the global sensitivity of the breeder gain is dominated by the confinement quality [30], which the corner grid brackets explicitly.

EDGE-TRANSPORT SOLVES FOR THE DETACHMENT BOUNDARY

The high-density boundary 16 is set by edge-transport codes run on the GPU HPC campaign and cross-checked against one another:

UEDGE (register serial BR-L2-A21): a finite-volume one-dimensional SOL solve on a 6 radial $\times$ 16 poloidal slab at fixed slab power, scanning upstream density to map the target-temperature collapse and the detachment onset; the first-SOL-tube values are quoted.
SD1D/Hermes-3 (BOUT++; register serials HX-14b/15b): a one-dimensional self-consistent SOL solve with neutral recycling and nitrogen seeding, integrated with a CVODE stiff solver, confirming detachment onset near $0.6\,n_{\mathrm{GW}}$ with a sub-one-percent seed.
SOLPS-ITER (register serials HX-14/BR-L2-A11; the ITER reference edge code): the two-dimensional B2.5–EIRENE deck for config-22021, used as the detachment arbiter and to set the $\sim 94\%$ radiated-fraction requirement. itemize
The equilibrium geometry underpinning the SOL connection length and the corner grid is taken from a free-boundary Grad–Shafranov solve (FreeGS) at the frozen shape, and the neutron source strength that sets the wall loading is cross-checked against Monte-Carlo neutronics (OpenMC) [23]; the negative-triangularity confinement that places the design point is grounded in nonlinear gyrokinetics (CGY-RO) [22,31].

GPU ACCELERATION, PROBLEM SCALE AND CONVERGENCE

The heavy first-principles verifications — nonlinear CGYRO at the real electron mass ratio, the two-dimensional SOLPS-ITER B2.5–EIRENE solve, and the three-dimensional OpenMC neutronics — are the GPU-bound components of the campaign and run on the A100/H100 platform; the reduced-order corner grid and mission-point scans, being deterministic fixed-point solves, run at interactive speed on the in-house node and are embarrassingly parallel over the grid. Convergence of the edge solves is monitored by the residual of the discretised transport equations and by cross-method agreement: the UEDGE divertor/upstream density ratio at onset $(n_t/n_u = 31)$ agrees within $\sim 15\%$ with the one-dimensional arbiter (~ 27), and the SD1D, self-written and Lengyel estimates of the seed fraction agree that a sub-one-percent nitrogen seed detaches. One honest fidelity caveat is carried: the SD1D integrator was CVODE-stiff at the full design parallel heat flux and was run at an ITER-class throat, with the full-flux case confirmed by the self-written solver and the analytic Lengyel model (§6).

VERIFICATION, FROZEN ANCHORS AND REPRODUCIBILITY

The campaign is verification-first [30]: each result is a frozen anchor with a serial identifier, a deterministic seed (`20260726`), and an independent-code reproduction. The design point, the corner-grid window and the detachment onset are the anchors BR-L1-A2, KX-L1-A1 and HX-14b/BR-L2-A21; a regression suite re-verifies the full anchor set on every build. All figures in this paper are regenerated from a single archived script that reads the anchor values, so the map is reproducible byte-for-byte from the deposit.

§ 4

Results

THE DESIGN POINT AND ITS PARAMETERS

Table 1 collects the frozen design-point parameters. The geometry fixes the Greenwald density through 3, and the power balance fixes the gain and the neutron output. The line-averaged density at $f_G^\star \approx 0.35$ is $\bar{n}_e \approx 4.7 \times 10^{20}\ \mathrm{m}^{-3}$.

QUANTITY	SYMBOL	VALUE
plasma current	I_p	9.66 MA
major radius	R_0	1.20 m
aspect ratio	A	2.5
minor radius	$a = R_0/A$	0.48 m
elongation	κ	2.0
triangularity	δ	−0.30
on-axis field	B_0	8 T
centrepost peak field	B_{cp}	16.84 T
Greenwald density	n_{GW}	$1.33 \times 10^{21}\ \mathrm{m}^{-3}$
Greenwald fraction (operating)	f_G	≈ 0.35
line-averaged density	$\bar{n}_e$	$\approx 4.7 \times 10^{20}\ \mathrm{m}^{-3}$
fusion gain	Q	3.076
fusion power	P_{fus}	85.04 MW
neutron power	P_n	67.79 MW
neutron source strength	S_n	$3.009 \times 10^{19}\ \mathrm{s}^{-1}$
first-wall neutron loading	Γ_{wall}	$1.885\ \mathrm{MW\,m}^{-2}$

Frozen design-point parameters of the Hyperion breeder (register anchors BR-L1-A2, KX-L1-A1). Two distinct magnets are modelled — the breeder centrepost at 16.84 T and the separate burner plug at 26.49 T — and only the breeder centrepost is relevant here.

THE DENSITY MARGIN

Figure 2 places the operating point on the Greenwald-fraction axis. The design point at $f_G \approx 0.35$ sits a factor $1/f_G \approx 2.9$ below the Greenwald limit at $f_G = 1.0$, with the detachment onset near $0.6\,n_{\mathrm{GW}}$ in between. Operating at roughly a third of the Greenwald density is a deliberate design choice: it keeps the machine far from the radiative/disruptive density boundary while leaving density available as a control lever for source strength (§4.4). The geometric origin of the limit is shown in Figure 3, which maps $n_{\mathrm{GW}} = I_p/(\pi a^2)$ over the (I_p, a) plane with the design point marked; the limit is set by current and minor radius alone, so the wide margin is robust to the physics scatter in the empirical scaling [2,5].

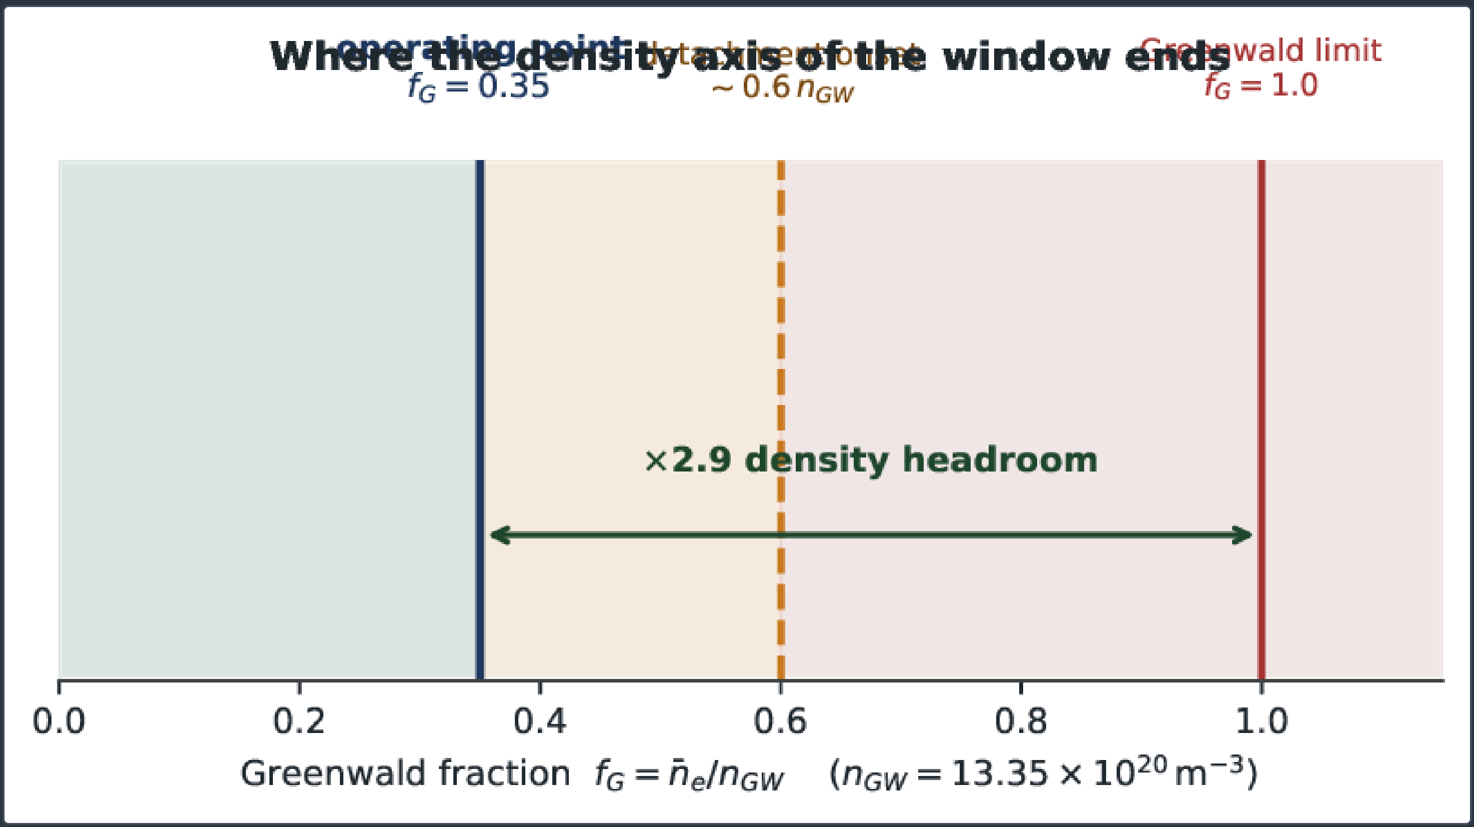

The density axis of the operating window. The operating point at $f_G \approx 0.35$ (navy), the detachment onset near $0.6\,n_{\mathrm{GW}}$ (amber), and the Greenwald limit at $f_G = 1.0$ (red), with a density headroom of a factor ~ 2.9. Data anchor: BR-L1-A2/HX-14b.

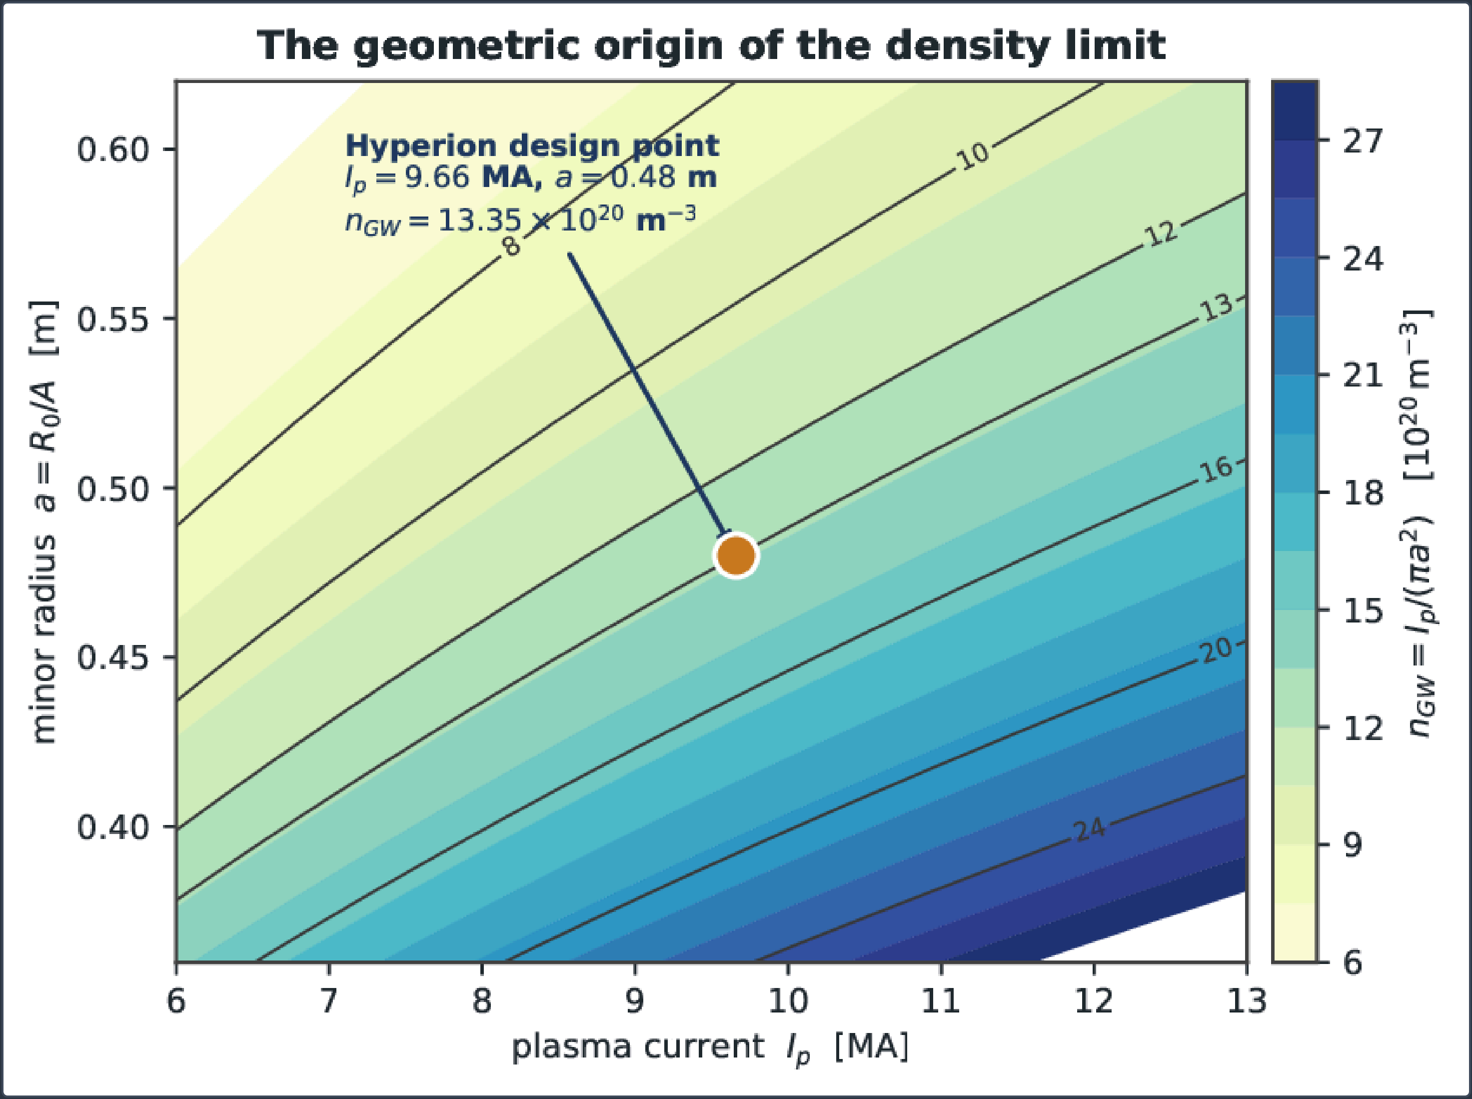

The geometric origin of the density limit. Contours of $n_{\rm GW} = I_p/(\pi a^2)$ over plasma current and minor radius; the Hyperion design point ($I_p = 9.66$ MA, $a = 0.48$ m) gives $n_{\rm GW} = 1.33 \times 10^{21}$ m^{-3}. The limit depends on current and minor radius alone, not on power or field.

THE FUSION-GAIN WINDOW

Figure 4 and Table 2 give the gain window solved on the nine-point corner grid. The fusion gain holds a window $Q = 2.735$–3.460 with fusion power $P_{\text{fus}} = 58.6$–$118.9\,\text{MW}$, and every physics gate — coil, density, β — is green at every corner. The design point ($Q = 3.076$, $P_{\text{fus}} = 85.04\,\text{MW}$) sits inside this window, not at its edge: it lies at 47% of the way through the Q range and 44% through the P_{fus} range. Three of the nine corners (the hot, dense corners) tripped the tritium duty band, an over-breeding exceedance rather than a physics failure, confirming that the binding constraint at high density is a mission-duty choice and not a stability wall.

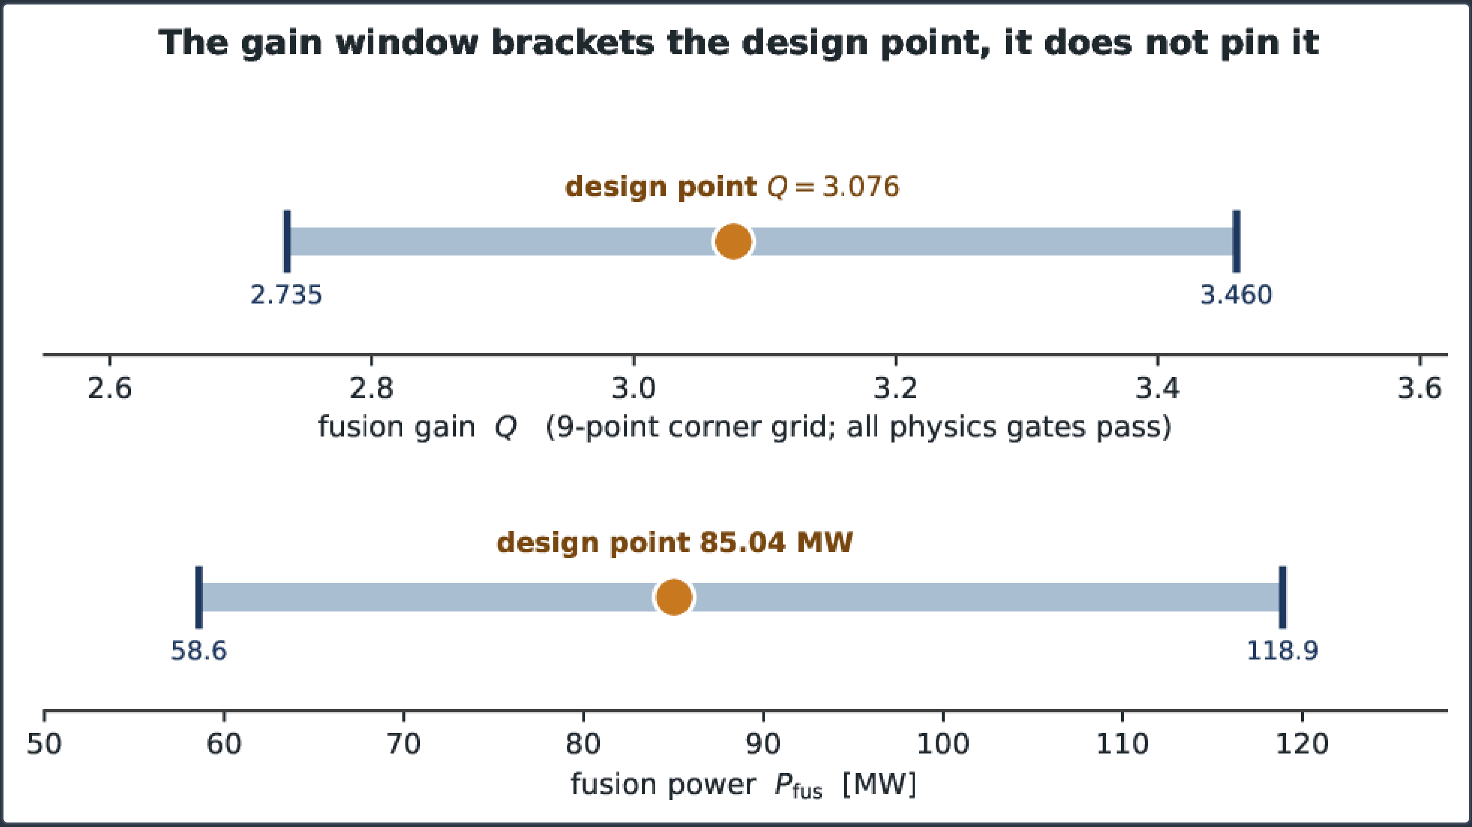

The fusion-gain window over the nine-point corner grid (register KX-L1-A1). Top: gain window $Q = 2.735$–3.460; bottom: fusion-power window $P_{\text{fus}} = 58.6$–$118.9\,\text{MW}$. The design point (amber) is enclosed in both, not pinned to an edge.

BOUNDARY	VALUE	DESIGN-POINT MARGIN
Greenwald density limit	$f_G = 1.0\,(1.33 \times 10^{21}\,\text{m}^{-3})$	operates at $f_G \approx 0.35$ ($\sim 2.9\times$)
detachment onset (seeded)	$\sim 0.6\,n_{\text{GW}}$	below onset at the operating point
radiative-collapse floor	$f_{\text{rad}} < f_{\text{rad}}^{\text{max}}$	satisfied (favourable NT screening)
ideal-MHD β limit	$\beta_N \leq \beta_N^{\text{lim}}$	bracketed (passive gate)
fusion-gain window (Q)	2.735–3.460	design point 3.076 enclosed
fusion power (P_{fus})	58.6–$118.9\,\text{MW}$	design point $85.04\,\text{MW}$ enclosed

The operating-window boundary along the density axis and the design-point margins.

SOURCE STRENGTH AND FIRST-WALL LOADING: THE ONE-TO-ONE INVARIANT

Because the geometry is fixed, both the neutron source strength S_n and the first-wall neutron loading Γ_{wall} scale directly with the neutron power P_n, so their ratio is invariant across the whole (f_G, T_{i0}) plane. Figure 5 and Table 3 show the invariant and three representative operating points. Raising density and temperature to the hot–dense corner ($f_G = 0.35$, $T_{i0} = 17\,\mathrm{keV}$) reaches a peak source strength of $5.388 \times 10^{19}\ \mathrm{s}^{-1}$, a factor 1.790 above the design point, at exactly $1.790\times$ the reference wall loading — off the duty band. The Pareto-interesting neighbour ($f_G = 0.35$, $T_{i0} = 13\,\mathrm{keV}$) delivers $0.967\times$ the design source strength at a slightly lower wall loading with a higher gain $Q = 3.865$, on-duty. There is no free operating-point gain: source strength and centrepost life are a strict one-to-one trade, and density is a smooth lever inside the window rather than a cliff. The lifetime-constrained operating-window synthesis [37] closes this trade against the consumable- centrepost replacement cadence.

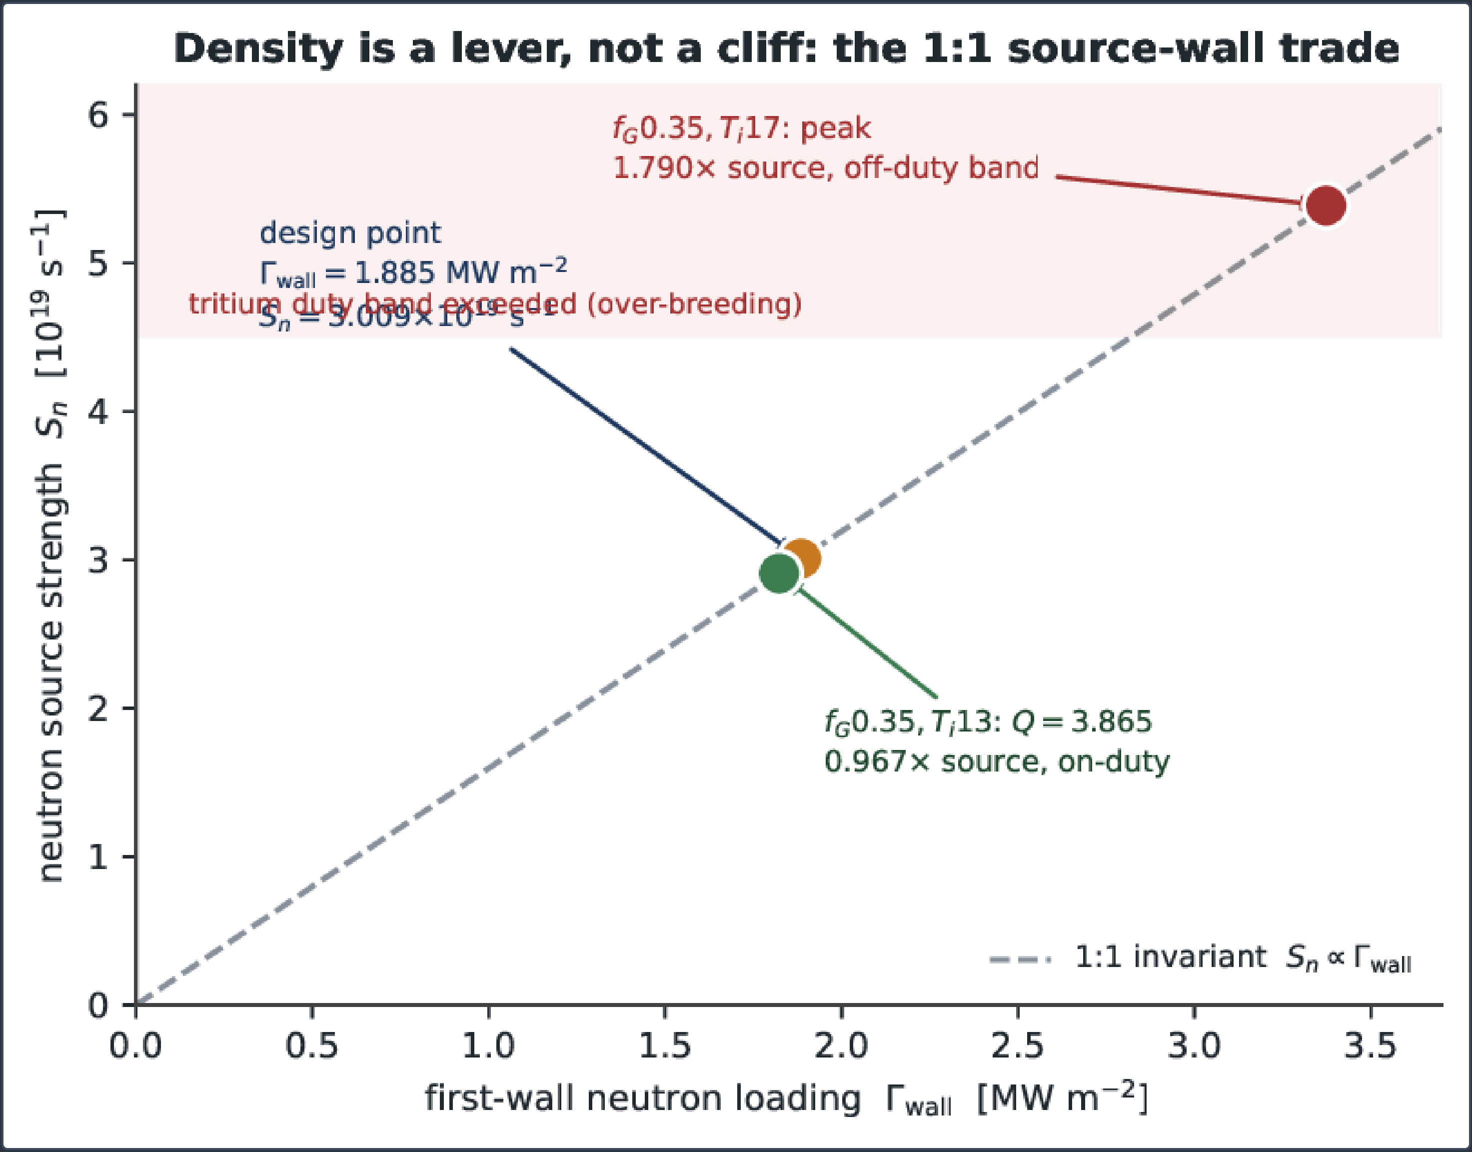

The source-strength/first-wall-loading invariant (register BR-L1-A2). Both quantities scale with the neutron power, so they trade one-for-one along the dashed invariant. Design point (amber); the $T_{i0} = 13\,\mathrm{keV}$ neighbour (green, $Q = 3.865$, on-duty); the $T_{i0} = 17\,\mathrm{keV}$ peak (red, $1.790\times$, off-duty). Above $\sim 1.5\times$ the reference loading the tritium duty band is exceeded by over-breeding.

OPERATING POINT	f_G	T_{i0} [KEV]	S_n [10^{19} S^{-1}]	Γ_{wall} [MW M^{-2}]	Q
design point	≈ 0.35	—	3.009 (1.000×)	1.885	3.076
low-T neighbour	0.35	13	2.909 (0.967×)	1.823	3.865 (on-duty)
hot–dense peak	0.35	17	5.388 (1.790×)	3.373	off-duty

Representative operating points from the mission-point scan (register BR-L1-A2). S_n and Γ_{wall} are a strict one-to-one invariant at fixed geometry.

THE HIGH-DENSITY BOUNDARY: DETACHMENT, NOT DISRUPTION

The practical high-density edge of the window is set by divertor exhaust, not by the disruptive density limit. Figure 6 shows the UEDGE one-dimensional SOL scan: as the upstream Greenwald fraction rises the target temperature collapses from 14.4 eV (attached) through 5.28 eV (marginal at the 5 eV detachment threshold) to 1.05 eV (detachment onset), following the two-point-model scaling $T_{e,t} \propto \bar{n}_e^{-1.7}$ steepening to $\bar{n}_e^{-2.8}$ near onset, with the divertor/upstream density ratio rising to $n_t/n_u = 31$. Mapped to the Greenwald-fraction axis the onset lies near $0.6\,n_{\mathrm{GW}}$, well above the $f_G \approx 0.35$ operating point, so the machine can raise density toward detachment access without approaching the disruptive limit.

Figure 7 makes the density-limit/detachment relation explicit through the detachment-access scaling 15. Density-only detachment at the design parallel heat flux would require an upstream density of $(1\text{–}2) \times n_{\mathrm{GW}}$ — above the disruptive limit and therefore inaccessible without seeding. Adding a sub-one-percent nitrogen seed, which supplies the $\sim 94\%$ radiated fraction required to hold the target flux within the plasma-facing-component limit, brings detachment access down to $\sim 0.6\,n_{\mathrm{GW}}$. Both the operational ceiling (seeded detachment access at $\sim 0.6\,n_{\mathrm{GW}}$) and the physics ceiling (the Greenwald limit at $1.0\,n_{\mathrm{GW}}$) therefore lie above the operating point, and the radiative boundary is a controllable operating choice inside the window (Figure 9), actuated by the real-time detachment controller [35] and backed by the collisional-radiative cooling curves that set the seed requirement [36].

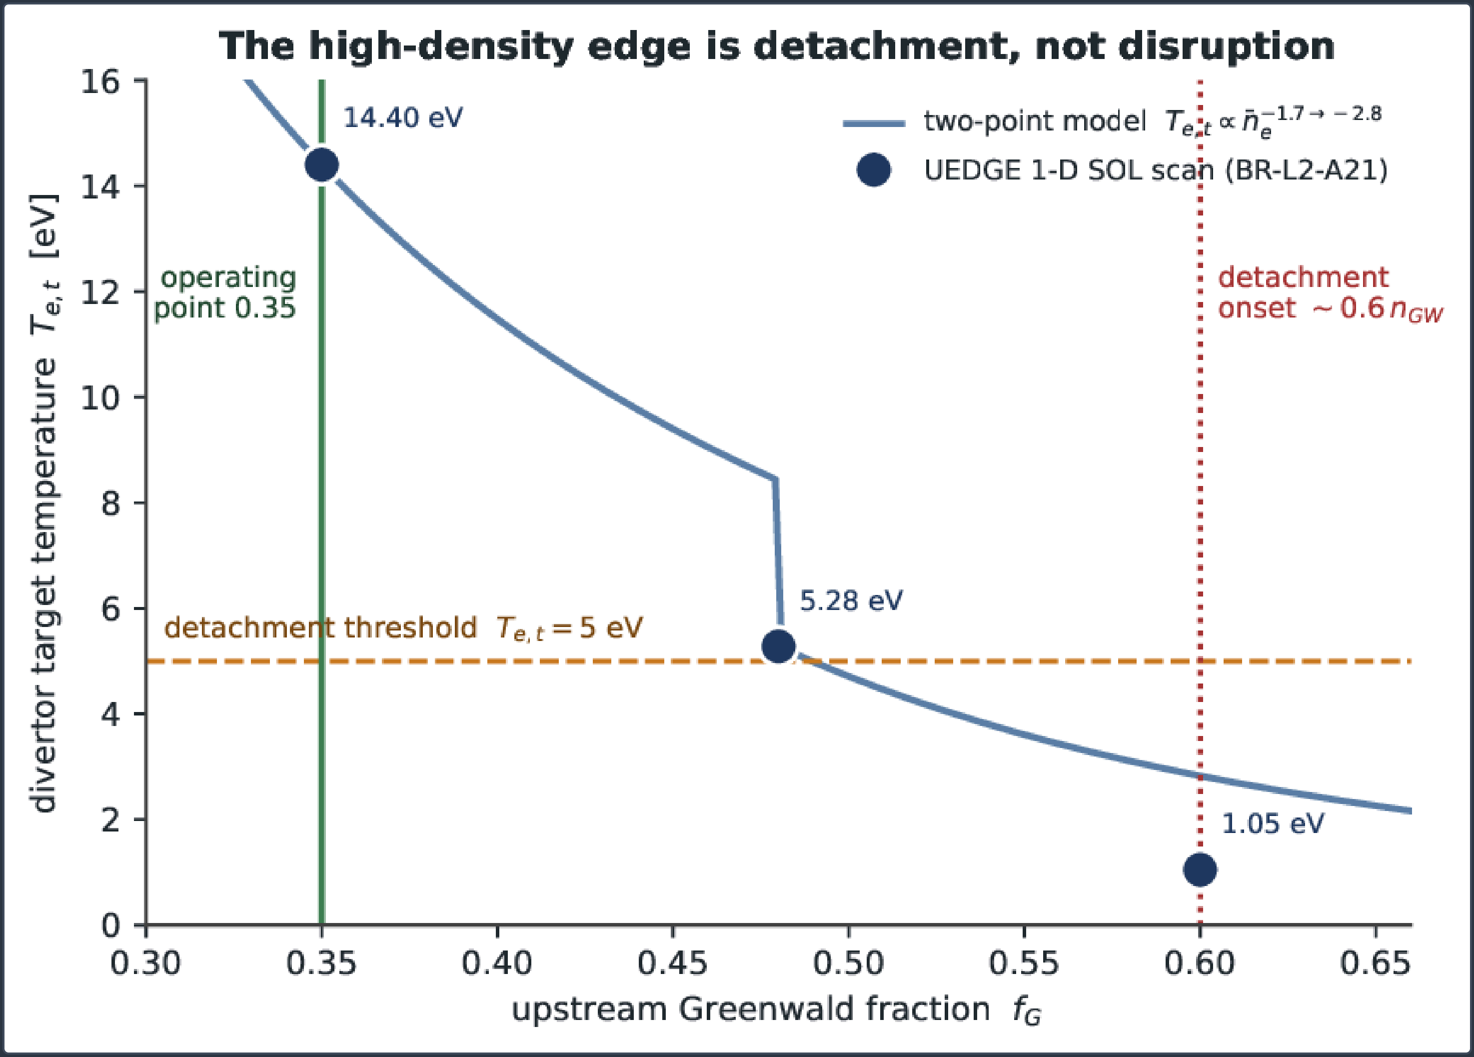

The divertor target temperature versus upstream Greenwald fraction (register BR-L2-A21, UEDGE). The measured attached/marginal/detached points follow the two-point-model collapse $T_{e,t} \propto \bar{n}_e^{-1.7\to-2.8}$; detachment onset (target 5 eV) lies near $0.6\, n_{\mathrm{GW}}$, above the operating point at 0.35.

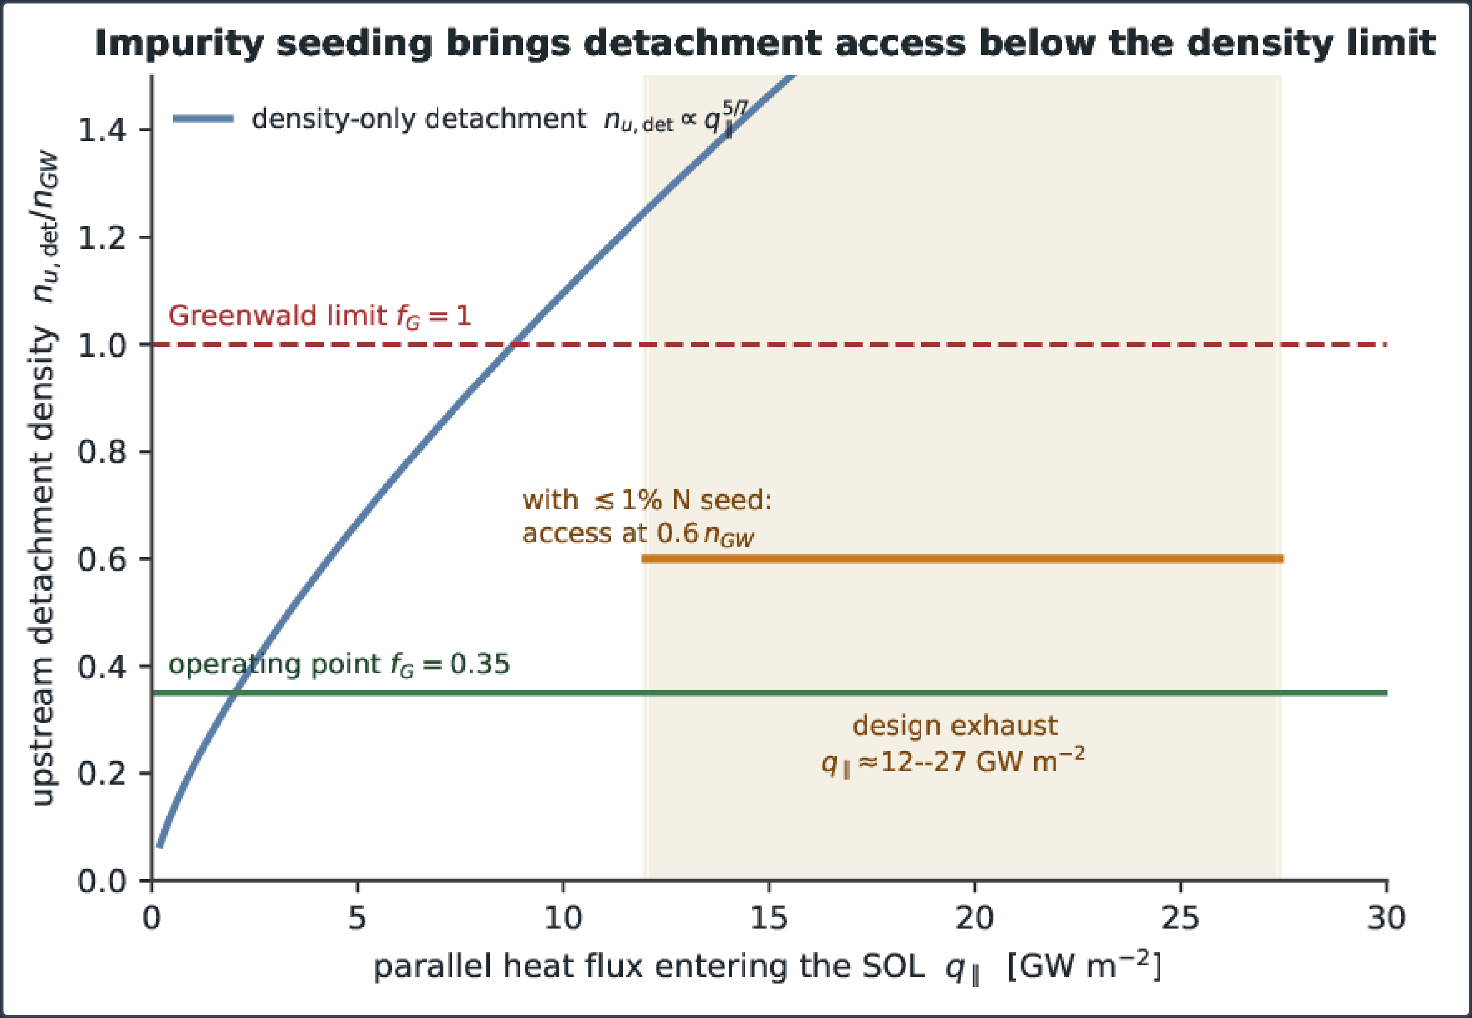

Detachment access versus parallel heat flux, from the two-point scaling $n_{u,\mathrm{det}} \propto q_\parallel^{5/7}$. Density-only detachment at the design exhaust ($q_\parallel \approx 12$–27 GW m^{-2}) would require $(1–2) \times n_{\mathrm{GW}}$; a sub-one-percent nitrogen seed brings access to $0.6\, n_{\mathrm{GW}}$, still above the operating point.

THE MARGIN SUMMARY AND THE RADIATIVE BOUNDARY

Figure 8 summarises the position of the design point within every wall of the window as the fraction of available margin used. The density axis is the most generous: the design point uses only 35% of the distance to the Greenwald limit and sits below the detachment onset. The gain and power windows bracket the design point near their centres, and the wall-loading duty band binds at 1.5× the reference loading — a mission choice, not a physics wall. Figure 9 shows the radiative boundary as the controlled loop it is: a seed of $c_z \lesssim 1\%$ produces the $\sim 94\%$ SOL radiated fraction that drives the target to $\sim 10\,\mathrm{MW\,m^{-2}}$ and $T_{e,t} \to 1.9\,\mathrm{eV}$, a detached, plasma-facing-component-safe divertor held by real-time feedback.

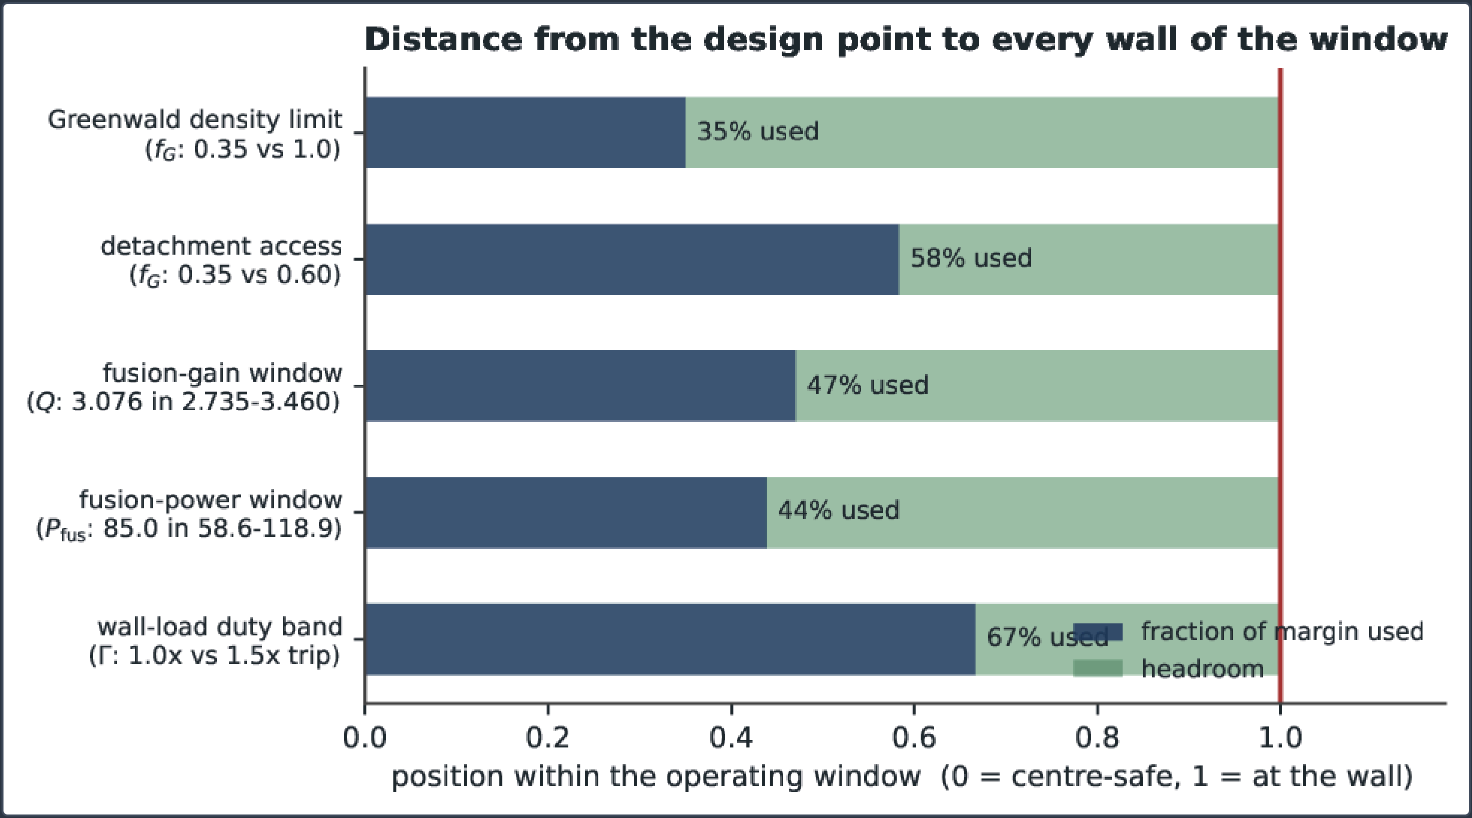

Distance from the design point to every wall of the operating window, as the fraction of available margin used (navy) versus headroom (green). The density axis is the most generous; the wall-load duty band is a mission choice rather than a physics limit.

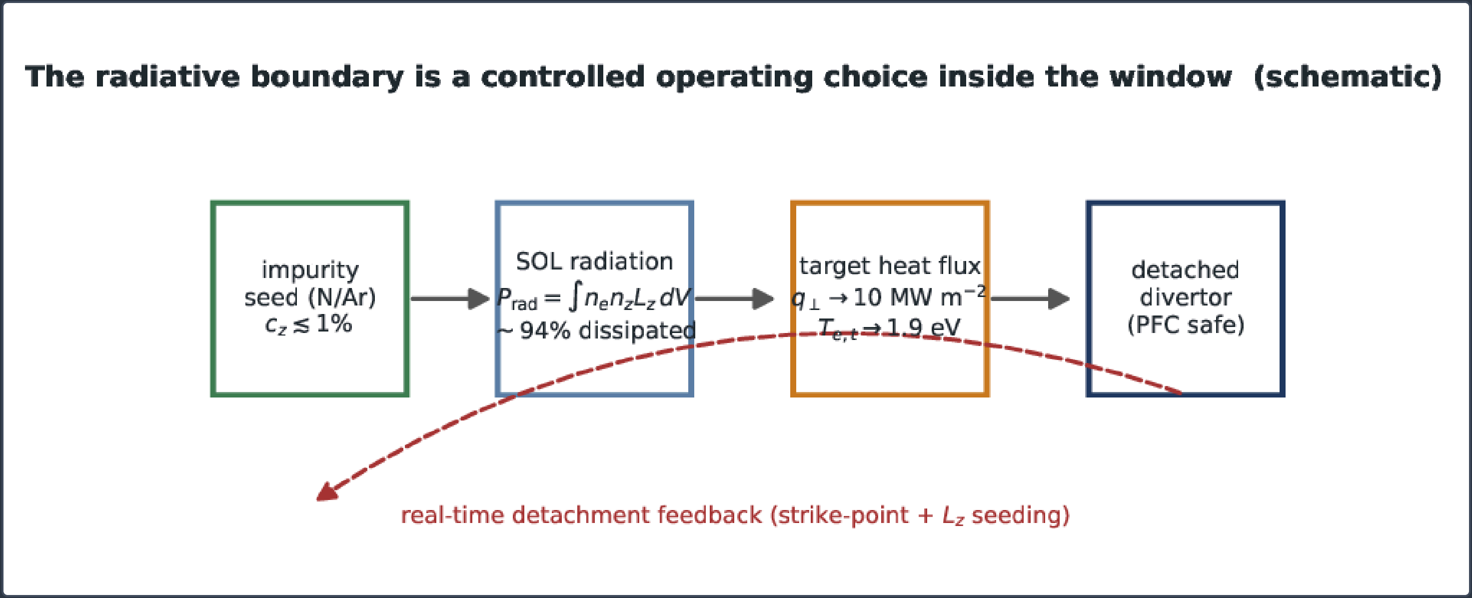

(Schematic.) The radiative boundary as a controlled operating choice. A sub-one-percent impurity seed produces the SOL radiated fraction ($\sim 94\%$) that caps the target heat flux and detaches the divertor, held inside the window by real-time detachment feedback.

THE BURNER DESIGN-SPACE, FOR CONTRAST

The same discipline — engineer the operating point to sit inside the window with margin — governs the separate burner line, a distinct D–^{3}He tandem-mirror machine with its own window and its own magnets. In the burner design-space database every one of the scanned corner points clears net electrical gain ($Q_E > 1$), with Q_E rising with density along a temperature ridge near $85\,\mathrm{keV}$; the burner design point sits at $Q_E = 1.318$ with neutron fraction $f_n = 5.44\%$, direct-conversion efficiency $\eta_{\mathrm{DEC}} = 0.49$ and a plug field of 26.49 T. We cite it here only to show that the design-inside-the-window discipline is common to both product lines; the burner′s magnets (26.49 T plug) and the breeder′s magnets (16.84 T centrepost) are two distinct devices and are never conflated.

§ 5

Discussion

THE DENSITY MARGIN AGAINST THE PUBLISHED DENSITY-LIMIT LITERATURE

The design point's factor $\sim \mathbf{2.9}$ density margin should be read against the modern understanding that the Greenwald value is a conservative, power-dependent boundary rather than a hard cliff. The empirical scaling 3 carries a scatter of tens of percent across the multi-machine database [1,2]; the first-principles turbulent-transport scaling of Giacomin *et al.* [5] predicts, for reactor-scale plasmas at full heating power, a limit *above* the raw Greenwald value, and the power-balance model of Zanca *et al.* [4] and the radiation-driven-island criterion of Gates and Delgado-Aparicio [3] both make the limit increase with heating power. A design that operates at $f_G \approx \mathbf{0.35}$ therefore sits safely below even the most pessimistic reading of the boundary, and the large fractional margin is robust precisely because it does not depend on which physics model of the limit is correct. In H-mode devices the density limit is reached through a sequence of confinement degradation and back-transition [6]; operating at a third of the Greenwald density keeps the machine far from the first rung of that sequence. The confinement quality that places the operating point on the gain surface is itself the dominant driver of the gain [30], and is grounded in the negative-triangularity turbulence suppression that gives the edge its quiet, ELM-free character [31,33].

THE EXHAUST BOUNDARY AGAINST THE DIVERTOR LITERATURE

The finding that density-only detachment is inaccessible below the Greenwald limit, while seeded detachment is reached at $\sim \mathbf{0.6}\, n_{\mathrm{GW}}$, is consistent with the divertor physics of high-power-density compact tokamaks. The parallel heat flux 12 for a compact spherical tokamak is severe because the heat-flux width [8,9] is narrow at high current; the SPARC divertor analysis reaches the same conclusion that the unmitigated parallel flux is an order of magnitude above the target limit and that seeded, detached operation is mandatory [10]. The Lengyel requirement 18 of a sub-one-percent nitrogen seed for a $\sim \mathbf{94\%}$ radiated fraction is in line with the analytic partially-detached-divertor scalings of Kallenbach *et al.* [11] and the ultimate-detachment physics of Krasheninnikov and Kukushkin [12]. Because the detachment onset lies at $\mathbf{0.6}\, n_{\mathrm{GW}}$, above the operating point, the divertor solution is decoupled from the density limit: the machine can seed and detach without approaching the disruptive boundary, and the two-dimensional SOLPS-ITER arbiter and the stacked Super-X/double-null portfolio close the exhaust gate at a tolerable seed [34,35]. The radiative-collapse floor 11 is the complementary constraint — the core must not be over-radiated — and is comfortably satisfied because the negative-triangularity edge screens impurities favourably [13].

DENSITY AS A CONTROL LEVER

The one-to-one source/wall invariant (§4.4) reframes the density axis as a control lever rather than a constraint. Within the window the operator can raise density to increase the neutron source strength (at a proportional cost in centrepost life), lower it to extend life, or push toward the detachment boundary to protect the divertor, all with a factor $\sim \mathbf{2.9}$ of headroom to the density limit throughout. This is the operating philosophy that the real-time control suite and the certified safe-envelope controller are built to exploit [35,30]: density is a smooth, bidirectional actuator inside a certified feasible set, not a boundary the controller must avoid. The disruption and runaway-electron loads that a density-limit excursion would drive are quantified and mitigated in the companion disruption study [38], so even the failure mode that the density limit guards against is bounded.

§ 6

Limitations

The claims rest on one honest fidelity gate. The Greenwald limit is an empirical scaling boundary rather than a first-principles one, and the operating-window claim rests on the large fractional margin ($f_G \approx 0.35$) to that empirical boundary and on the detachment onset lying above the operating point — both robust to the scatter in the scaling and to the choice of density-limit physics model. The one open computational item is the edge-code fidelity at the full design parallel heat flux: the SD1D/Hermes-3 integrator was CVODE-stiff at the full exhaust flux and was run at an ITER-class throat, with the full-flux detachment confirmed by a self-written solver and by the analytic Lengyel model, and the two-dimensional SOLPS-ITER arbiter for config-22021 remains the next fidelity step. The sub-one-percent seed requirement and the $0.6\,n_{\mathrm{GW}}$ onset are cross-method-agreed across UEDGE, SD1D and Lengyel, so the qualitative conclusion — a wide density margin and a controllable, seeded high-density boundary — does not depend on closing that item, which is a refinement rather than a gate.

§ 7

Conclusion

The Hyperion breeder operates inside a quantified safe envelope with a wide density margin. From the geometry the Greenwald density is $n_{\mathrm{GW}} = 1.33 \times 10^{21}\ \mathrm{m}^{-3}$, and the design point operates at a Greenwald fraction $f_G \approx 0.35$ — a density headroom of a factor ~ 2.9 below the empirical limit. The four walls of the operating window — the Greenwald/radiative density boundary, the ideal-MHD β limit, the SOL detachment boundary and the radiative-collapse floor — were derived as explicit inequalities and the enclosed feasible set was solved on a nine-point corner grid: the fusion gain holds a window $Q = 2.735$–3.460 ($P_{\mathrm{fus}} = 58.6$–$118.9\,\mathrm{MW}$) with every physics gate satisfied, enclosing the frozen design point $Q = 3.076$, $P_{\mathrm{fus}} = 85.04\,\mathrm{MW}$. Because the geometry is fixed, source strength and first-wall loading are a strict one-to-one invariant, so density is a controllable lever ($Q = 3.865$ at $0.967\times$ the reference wall loading) rather than a cliff. The high-density boundary is set by seeded divertor detachment near $0.6\,n_{\mathrm{GW}}$ — reached with a sub-one-percent nitrogen seed that dissipates $\sim 94\%$ of the exhaust power and holds the target heat flux within the plasma-facing-component limit — well above the operating point, so the density limit is a distant wall and the radiative boundary is a controlled operating choice, not an operating constraint.

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 on which the operating-window map and the edge-transport solves were produced and independently reproduced.

Reproducibility. The operating point (BR-L1-A2), the corner-grid window (KX-L1-A1) and the detachment onset (HX-14b/BR-L2-A21) are frozen anchors in the Kronos Physics De-Risking Register [29] (doi:10.5281/zenodo.22133057) and are regenerable from the reduced-order engine and the single figure-generation script deposited with this paper; this paper is archived at doi:10.5281/zenodo.22645947; official version: https://www.kronosfusionenergy.com/publications/the-greenwald-fraction-density-limit-and-the-radiative-.html.

Data availability tocsectionData availability The Greenwald-fraction map, the corner-grid gain window, the mission-point invariant and the edge-code detachment scan used for the figures are archived with the deposit at doi:10.5281/zenodo.22645947, which references the Kronos 2026 series and the de-risking register [29]. All quantitative values trace to the register by the serial identifiers cited inline. Any economic analysis is reported separately and is not part of the public deposit.

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

tocsectionReferences

APPARATUS

Portfolio, People & Record

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

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

INTELLECTUAL PROPERTY

The patent portfolio

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

GRANTED & PENDING UTILITY

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

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

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

TRADEMARK & ORIGIN

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

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

THE TEAM

Founding partners, board, advisors & auditors

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

FOUNDING PARTNERS — SCIENCE

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

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

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

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

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

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

SCIENTIFIC ADVISORS

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

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

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

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

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

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

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

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

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

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

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

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

SCIENTIFIC AUDITORS

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

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

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

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

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

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

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

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

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

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

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

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

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

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

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

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

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

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

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

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

Andrea Romero
Marketing Lead
2022–Present

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

SOURCES

References

1 M Greenwald, J L Terry, S M Wolfe, S Ejima, M G Bell, S M Kaye and G H Neilson, A new look at density limits in tokamaks, *Nucl. Fusion* **28** 2199 (1988). doi:10.1088/0029-5515/28/12/009

2 M Greenwald, Density limits in toroidal plasmas, *Plasma Phys. Control. Fusion* **44** R27 (2002). doi:10.1088/0741-3335/44/8/201

3 D A Gates and L Delgado-Aparicio, Origin of tokamak density limit scalings, *Phys. Rev. Lett.* **108** 165004 (2012). doi:10.1103/PhysRevLett.108.165004

4 P Zanca, F Sattin, D F Escande, G Pucella and O Tudisco, A unified model of density limit in fusion plasmas, *Nucl. Fusion* **57** 056010 (2017). doi:10.1088/1741-4326/aa6230

5 M Giacomin, A Pau, P Ricci, O Sauter and T Eich, First-principles density limit scaling in tokamaks based on edge turbulent transport and implications for ITER, *Phys. Rev. Lett.* **128** 185003 (2022). doi:10.1103/PhysRevLett.128.185003

6 M Bernert *et al*, The H-mode density limit in the full tungsten ASDEX Upgrade tokamak, *Plasma Phys. Control. Fusion* **57** 014038 (2015). doi:10.1088/0741-3335/57/1/014038

7 B Lipschultz, B LaBombard, E S Marmar, M M Pickrell, J L Terry, R Watterson and S M Wolfe, MARFE: an edge plasma phenomenon, *Nucl. Fusion* **24** 977 (1984). doi:10.1088/0029-5515/24/8/002

8 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

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 A Q Kuang *et al*, Divertor heat flux challenge and mitigation in SPARC, *J. Plasma Phys.* **86** 865860505 (2020). doi:10.1017/S0022377820001117

11 A Kallenbach, M Bernert, R Dux, F Reimold and M Wischmeier, Analytical calculations for impurity seeded partially detached divertor conditions, *Plasma Phys. Control. Fusion* **58** 045013 (2016). doi:10.1088/0741-3335/58/4/045013

12 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

13 A Loarte *et al*, Chapter 4: Power and particle control (Progress in the ITER Physics Basis), *Nucl. Fusion* **47** S203 (2007). doi:10.1088/0029-5515/47/6/S04

14 E J Doyle *et al*, Chapter 2: Plasma confinement and transport (Progress in the ITER Physics Basis), *Nucl. Fusion* **47** S18 (2007). doi:10.1088/0029-5515/47/6/S02

15 H Zohm *et al*, On the physics guidelines for a tokamak DEMO, *Nucl. Fusion* **53** 073019 (2013). doi:10.1088/0029-5515/53/7/073019

16 G Verdoolaege *et al*, The updated ITPA global H-mode confinement database: description and analysis, *Nucl. Fusion* **61** 076006 (2021). doi:10.1088/1741-4326/abdb91

17 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

18 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

19 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

20 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

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

22 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

23 P K Romano, N E Horelik, B R Herman, A G Nelson, B Forget and K Smith, OpenMC: a state-of-the-art Monte Carlo code for research and development, *Ann. Nucl. Energy* **82** 90 (2015). doi:10.1016/j.anucene.2014.07.048

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

25 J Wesson, *Tokamaks*, 4th edn, Oxford University Press, Oxford (2011).

26 S Wiesen *et al*, The new SOLPS-ITER code package, *J. Nucl. Mater.* **463** 480 (2015).

27 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).

28 B D Dudson, M V Umansky, X Q Xu, P B Snyder and H R Wilson, BOUT++: a framework for parallel plasma fluid simulations, *Comput. Phys. Commun.* **180** 1467 (2009).

29 P I Ford, *The Physics De-Risking Register: Gate-by-Gate Evidence, Independently Re-Run and Stamped*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22133057

30 P I Ford, *Hyperion as a Strategic-Materials Platform: Gain, Tritium Breeding, He-3 Co-Production and 14 MeV Neutron Output*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645691

31 P I Ford and G L Kulcinski, *Negative-Triangularity Turbulence Suppression in a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645722

32 P I Ford, *Free-Boundary Equilibrium and Ideal-MHD Stability of a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645740

33 P I Ford, *The Negative-Triangularity Edge: Heat-Flux Width, the ELM-Free Pedestal, and the Scrape-Off Layer*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645830

34 P I Ford, *A Stacked Detachment Portfolio for a Compact Spherical-Tokamak Breeder: Super-X, Double-Null and Impurity Seeding*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645742

35 P I Ford, *Real-Time Detachment Control for a Super-X Divertor in a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645874

36 P I Ford, *Collisional-Radiative Impurity-Radiation Modeling for a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645883

37 P I Ford, *Lifetime-Constrained Operating-Window Synthesis for a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645799

38 P I Ford, *Disruptions, Runaway-Electron Avalanche and Electromagnetic Loads in a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645724

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 Greenwald-Fraction Density Limit and the Radiative Boundary: Mapping the Safe Operating Envelope of a Compact Spherical-Tokamak Breeder

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

Zenodo · doi:10.5281/zenodo.22645947

Read on kronosfusionenergy.com

The Physics De-Risking Register