KRONOS FUSION ENERGY

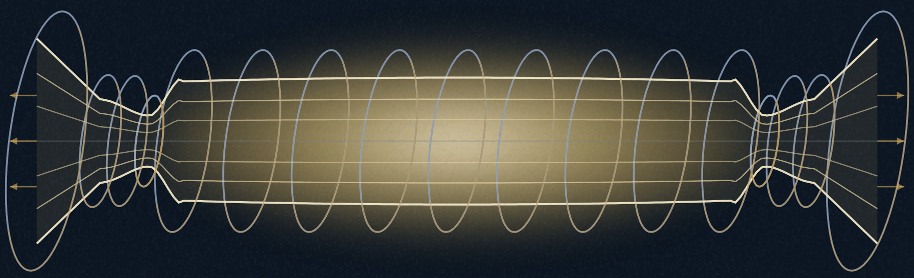

PAPER 2.55 · THE KRONOS 2026 SERIES

Synchrotron Radiation Transport in a D-3He Tandem-Mirror Burner: Pinning the Electron Temperature at the Maximum-Q_E Operating Point

Unlike a deuterium–tritium tokamak, a hot deuterium–helium-3 (D–He) tandem-mirror burner is radiation-dominated: at the electron temperatures required for advanced-fuel reactivity...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Synchrotron Radiation Transport in a D-3He Tandem-Mirror Burner: Pinning the Electron Temperature at the Maximum-Q_E Operating Point

Unlike a deuterium–tritium tokamak, a hot deuterium–helium-3 (D–^{3}He) tandem-mirror burner is *radiation-dominated*: at the electron temperatures required for advanced-fuel reactivity the plasma loses more power to electron-cyclotron (synchrotron) emission than to any other charged-particle channel, and that loss is what fixes the optimal electron temperature. We formulate the coupled emission–absorption problem for the burner central cell and solve it self-consistently. The synchrotron loss is neither the optically-thin estimate nor a fixed fraction of fusion power. At the frozen design point ($P_{\text{fus}} = 4298.5\,\text{MW}$, neutron fraction $f_n = 5.44\%$) the total radiated power is $2546.14\,\text{MW}$ ($f_{\text{rad}} = 0.592$); the reabsorbed synchrotron loss is $1803.04\,\text{MW}$ by a relativistic Maxwell–J\"uttner harmonic-emission transport calculation with wall reflection, against $918.59\,\text{MW}$ by the Trubnikov universal formula and $66\,677\,\text{MW}$ optically thin — a model-form spread of $37\times$ in the dominant loss channel that no design can leave unresolved. The plasma is optically thick to an effective harmonic cutoff $n^* \approx 10.4$, above which emission escapes, and the cutoff grows as the one-third power of the central-cell size, so the synchrotron loss *fraction* falls as $a^{-1/3}$; a universal dimensionless coefficient in the reabsorbed scaling, $\Lambda_\star = 24.74579$, is reproduced to 2.6×10^{-15} against its closed form. Substituting the transport loss into the engineering-gain balance produces an interior maximum: Q_E peaks at a self-consistent electron temperature $T_e = 89.72\,\text{keV}$, falling to 1.158 and 0.887 under $\mp 10\%$ excursions about the canonical $Q_E = 1.318$. The one hard requirement the analysis exposes is a first-wall (synchrotron-mirror) reflectivity $R_{\text{wall}} \gtrsim 0.83$ for net-breakeven closure. Every headline quantity is a frozen anchor (gate KX-L1-A3-BU) in the Kronos Physics De-Risking Register.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645772 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: synchrotron radiation · electron-cyclotron emission · radiation transport · optical depth · tandem mirror · deuterium--helium-3 · advanced fuel · direct energy conversion

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 RADIATION OBJECTION TO ADVANCED FUELS

The case against advanced (aneutronic or neutron-lean) fusion fuels has always been an argument about radiation. A deuterium–helium-3 plasma reacts through $\mathrm{D} + {}^3\mathrm{He} \rightarrow p\,(14.68\,\mathrm{MeV}) + {}^4\mathrm{He}\,(3.67\,\mathrm{MeV})$, releasing $18.35\,\mathrm{MeV}$ entirely in charged particles, so that in principle the products can be recovered at high efficiency by direct energy conversion rather than through a thermal cycle [1,2]. The price of that neutron-lean payoff is temperature: the D–^{3}He reactivity peaks near ion temperatures of a few hundred keV [3,4], an order of magnitude hotter than a D–T tokamak, and at those temperatures the electrons radiate. Rider's much-cited analysis made the objection quantitative: for a near-thermal advanced-fuel plasma the bremsstrahlung and cyclotron losses can approach or exceed the fusion power, leaving little or no margin for net gain [5]. Whether a D–^{3}He machine closes therefore hinges on getting the radiation loss *right*, and specifically on getting the electron-cyclotron (synchrotron) loss right, because at reactor density and field it is the term that is hardest to estimate and easiest to get wrong by an order of magnitude.

The reason it is hard is reabsorption. Bremsstrahlung escapes a fusion plasma essentially unimpeded and is a well-behaved volumetric loss $\propto n_e^2 Z_{\mathrm{eff}} T_e^{1/2}$. Synchrotron radiation does not escape freely: the plasma is optically thick at the low cyclotron harmonics that carry most of the emitted power, so the radiation is emitted, reabsorbed, re-emitted, and only partially transmitted, with the surrounding wall reflecting a large fraction of what reaches it back into the plasma [6,8,9,10]. The escaping power is consequently a *transport* quantity, set by an optical-depth balance and a wall-reflection coefficient, not by the local emissivity. An optically-thin sum over harmonics overestimates the loss by nearly two orders of magnitude; a crude “fixed fraction of fusion power” underestimates its steep temperature dependence. The design point sits between these errors, and only a transport calculation places it.

WHY THE MIRROR IS RADIATION-DOMINATED AND THE TOKAMAK IS NOT

The contrast with the companion breeder is the cleanest statement of the problem. The Kronos programme fields two distinct products with two distinct magnets: a compact spherical-tokamak (ST) breeder governed by a centrepost at a peak field of $16.84\,\mathrm{T}$, running a D–T-like plasma at $P_{\mathrm{fus}} = 85.04\,\mathrm{MW}$; and a D–^{3}He tandem-mirror burner whose plug coil is a separate, higher-field magnet at $26.49\,\mathrm{T}$. In the breeder, at the 10–$20\,\mathrm{keV}$ electron temperatures of a D–T plasma, the whole electron-cyclotron loss is only $\approx 1.55\,\mathrm{MW}$ and bremsstrahlung $\approx 3.12\,\mathrm{MW}$ — radiation is a $\sim 2\%$ correction to the power balance (register serial KX-L1-A3). In the burner, at $T_e \approx 90\,\mathrm{keV}$, the synchrotron loss alone is $1803\,\mathrm{MW}$ — 42% of the fusion power (KX-L1-A3-BU). The same physics that is a footnote for the tokamak is the dominant charged-power loss channel for the mirror (figure 8). This is not an artefact of machine choice; it is the $B^2 T_e$ scaling of cyclotron emission acting on a high-field, high-T_e advanced-fuel plasma, and it is the reason the burner's operating point cannot be set without solving the transport problem.

PRIOR ART

Synchrotron loss from hot plasmas has a long literature. Trubnikov derived universal coefficients for the reabsorbed loss from a slab with a reflecting wall [6,7]; Tamor's cytran algorithm made the harmonic-resolved transport calculation tractable for reactor studies [8]; Bornatici *et al.* systematised electron-cyclotron emission and absorption in fusion plasmas [9]; and Albajar, Johner and Granata (the reflective-cavity relativistic transport model we adopt, hereafter AFJ) improved the harmonic truncation and the treatment of realistic profiles and reflection for reactor tokamaks [10,11]. These tools were built for, and validated on, D–T tokamaks where synchrotron loss is a modest correction. What is missing is their disciplined application to a D–^{3}He *mirror*, where the loss is dominant rather than corrective, where the geometry is a long open central cell rather than a torus, and where the escaping loss feeds back onto the choice of the operating temperature itself through a direct-conversion power balance. The tandem-mirror literature [15,16,17] treated end-plugging and confinement in depth but generally folded synchrotron loss into a coefficient; the modern high-field-mirror programme [19] reopens the question at reactor field.

CONTRIBUTION

This paper computes the burner synchrotron loss from first principles and shows what it implies for the operating point. Specifically we (i) write the coupled 0-D power balance of the D–^{3}He central cell with its bremsstrahlung and electron-cyclotron loss terms as explicit equations (§2); (ii) formulate the electron-cyclotron emission from a relativistic Maxwell–J\"uttner electron population, the harmonic optical depth, the reflective-cavity reabsorption, and the three model closures (optically thin, Trubnikov, AFJ) whose $37\times$ spread is the object of study (§2); (iii) derive the harmonic cutoff n^*, its $a^{1/3}$ size scaling, and the universal reabsorption coefficient reproduced to machine precision (§2); (iv) give the numerical formulation — discretization of the harmonic sum and angular integral, the self-consistent electron-temperature root-find, the design-plane sweep, and the convergence and verification evidence (§3, §4); (v) report the results, including the maximum-Q_E temperature $T_e = 89.72\,\mathrm{keV}$ and its sensitivity, across eight data figures and four tables (§5); (vi) benchmark quantitatively against the published synchrotron and advanced-fuel literature (§6); and (vii) state the one honest gate the analysis exposes — the wall-reflectivity floor $R_{\mathrm{wall}} \gtrsim 0.83$ (§7). Every quantitative claim traces to a frozen result in the Kronos Physics De-Risking Register [30]; the gate identifier KX-L1-A3-BU is given inline so each number is auditable.

§ 2

Governing equations and formal problem

THE CENTRAL-CELL POWER BALANCE

The burner central cell is modelled as a uniform, quasi-neutral D–^{3}He plasma of volume V with deuterium and helium-3 ion densities n_D, $n_{\rm He}$, electron density $n_e = n_D + 2n_{\rm He}$, ion temperature T_i and electron temperature T_e. The volumetric fusion power is

$$p_{\rm fus} = n_D\, n_{\rm He}\, \langle\sigma v\rangle_{\rm D^3He}(T_i)\, E_{\rm D^3He} \;+\; p_{\rm fus}^{\rm DD}(T_i),$$

with $E_{\rm D^3He} = 18.35\,{\rm MeV}$ released entirely in charged particles and $\langle\sigma v\rangle$ the Maxwell-averaged reactivity [3]. The second term collects the parasitic $\rm D + D$ branches, which — together with the small D–T yield from the tritium they breed — produce the only neutrons the machine makes; their share defines the neutron fraction

$$f_n \;\equiv\; \frac{P_n}{P_{\rm fus}}, \qquad P_{\rm fus} = \int_V p_{\rm fus}\,{\rm d}V.$$

At the frozen design point, $P_{\rm fus} = 4298.5\,{\rm MW}$, $f_n = 5.44\%$ and $P_n = 233.95\,{\rm MW}$, so the charged-particle power delivered to the plasma and the converters is

$$P_{\rm ch} = (1 - f_n)\, P_{\rm fus} = 4064.55\,{\rm MW}.$$

Of $P_{\rm ch}$, a part is radiated away as photons and a part flows as particle energy to the mirror ends and the direct-energy converter. The steady electron energy balance is

$$\underbrace{P_{\rm ch}^{(e)} + P_{\rm aux}^{(e)}}_{\text{heating of electrons}} \;=\; \underbrace{P_{\rm brems} + P_{\rm sync} + P_{\rm line}}_{P_{\rm rad}} \;+\; P_{ei}^{\rm transport},$$

where $P_{\rm ch}^{(e)}$ is the fraction of charged-particle power that heats electrons, $P_{\rm aux}^{(e)}$ any electron-directed auxiliary heating, and the right-hand side collects the radiated losses and the collisional/transport exchange. The total radiated power

$$P_{\rm rad} = P_{\rm brems} + P_{\rm sync} + P_{\rm line}, \qquad f_{\rm rad} \equiv \frac{P_{\rm rad}}{P_{\rm fus}},$$

is the quantity this paper computes; at the design point $P_{\rm rad} = 2546.14\,{\rm MW}$ and $f_{\rm rad} = 0.592$.

ENGINEERING GAIN AND THE MAXIMUM-Q_E OBJECTIVE

The burner's figure of merit is the *engineering* gain Q_E, the ratio of net electric output to recirculated electric power. Because the products are charged, the non-radiated charged flow is recovered by a direct-energy converter at efficiency $\eta_{\rm DEC} = 0.49$, while the radiated and neutron power reach the balance-of-plant thermal loop at the lower thermal efficiency $\eta_{\rm th}$:

$$P_{\rm elec} = \eta_{\rm DEC}\left(P_{\rm ch} - P_{\rm rad}\right) \;+\; \eta_{\rm th}\left(P_{\rm rad} + P_n + P_{\rm end}\right), \qquad Q_E = \frac{P_{\rm elec} - P_{\rm recirc}}{P_{\rm recirc}}.$$

Equation (6) exposes the mechanism that makes an interior optimum in T_e: every watt that $P_{\rm rad}(T_e)$ moves out of the direct-conversion stream and into the thermal loop is converted at $\eta_{\rm th} < \eta_{\rm DEC}$, a penalty that grows with T_e; on the other side, the electron pressure supports the plasma β and the self-consistent confinement that the fusion term of equation (1) requires. The formal problem is

$$T_e^* \;=\; \arg\max_{T_e}\; Q_E\big(T_e;\, \{n, B, R_{\rm wall}, a, \ldots\}\big), \qquad \left.\frac{\partial Q_E}{\partial T_e}\right|_{T_e^*} = 0, \quad \left.\frac{\partial^2 Q_E}{\partial T_e^2}\right|_{T_e^*} < 0.$$

The remainder of the section supplies the loss terms that make $Q_E(T_e)$ concave; §5 reports the maximiser $T_e^* = 89.72\,{\rm keV}$.

BREMSSTRAHLUNG

For a Maxwellian electron population the relativistic thermal bremsstrahlung power density, including electron–electron bremsstrahlung and the leading relativistic corrections, is

$$p_{\text{brems}} = C_B\, n_e^2\, Z_{\text{eff}}\, T_e^{1/2} \Big[1 + a_1\theta + a_2\theta^2 + \dots\Big], \qquad \theta \equiv \frac{T_e}{m_e c^2},$$

with $C_B = 1.69 \times 10^{-38}\,\mathrm{W\,m^3\,keV^{-1/2}}$ the standard coefficient and the bracket the relativistic enhancement [5,12]. Integrated over the central cell this gives $P_{\text{brems}} = 743.10\,\mathrm{MW}$; an independent evaluation with the Xie relativistic fit returns $760.21\,\mathrm{MW}$, a $+2.3\%$ cross-check that bounds the model-form error of the (well-behaved) bremsstrahlung term. Bremsstrahlung is optically thin and escapes freely; it is the synchrotron term that requires transport.

ELECTRON-CYCLOTRON EMISSION FROM A RELATIVISTIC POPULATION

An electron of Lorentz factor γ gyrating in field B radiates at harmonics of its relativistic cyclotron frequency,

$$\omega_n = n\,\Omega_{ce}, \qquad \Omega_{ce} = \frac{eB}{\gamma\, m_e}, \qquad n = 1, 2, 3, \dots$$

At $T_e = 89.72\,\mathrm{keV}$ the electrons are mildly relativistic, $\theta = T_e/m_e c^2 = 0.176$, so the emission is described by the relativistic Maxwell–J\"uttner distribution

$$f_{\text{MJ}}(\gamma) = \frac{\gamma^2\beta}{\theta\, K_2(1/\theta)} \exp\Big(-\frac{\gamma}{\theta}\Big), \qquad \beta = \sqrt{1-\gamma^{-2}},$$

with K_2 the modified Bessel function. Relativistic mass increase spreads and down-shifts each harmonic and merges the high harmonics into a continuum, which is exactly what makes the naive non-relativistic harmonic sum unreliable and the Maxwell–J\"uttner treatment necessary. The local, optically-thin cyclotron emissivity summed over harmonics scales as

$$p_{\text{sync}}^{\text{thin}} \;\simeq\; \frac{e^4}{3\pi\varepsilon_0 m_e^2 c^3}\, n_e\, B^2 \left(\frac{T_e}{m_e c^2}\right) g(\theta),$$

i.e. $\propto n_e B^2 T_e$, with $g(\theta)$ an order-unity relativistic factor. Equation (11) is what yields the unphysical $66\,677\,\mathrm{MW}$ when applied to the burner: it counts every emitted photon as lost, ignoring that at low harmonics the plasma reabsorbs essentially all of them.

OPTICAL DEPTH, THE HARMONIC CUTOFF, AND THE REFLECTIVE CAVITY

Synchrotron transport in the central cell is a reabsorption problem. The optical depth of harmonic n across a plasma of transverse size a is

$$\tau_n \;\simeq\; \alpha_n(T_e, \theta_{\text{k}})\, \frac{\omega_{pe}^2}{\omega_{ce}\, c}\, a,$$

where α_n is the (temperature- and angle-dependent) absorption coefficient of the nth harmonic, $\omega_{pe}^2 = n_e e^2/\varepsilon_0 m_e$ and $\omega_{ce} = eB/m_e$. Because α_n falls steeply with n, the low harmonics are optically thick ($\tau_n \gg 1$) and only the high-harmonic tail escapes. Radiation transport through the cell, closed by a wall of power-reflectivity R_{wall}, is governed by the one-dimensional transfer equation for the spectral intensity I_ω along a ray,

$$\frac{\mathrm{d}I_\omega}{\mathrm{d}s} = j_\omega - \alpha_\omega I_\omega, \qquad I_\omega^{\text{wall},+} = R_{\text{wall}}\, I_\omega^{\text{wall},-} + (1 - R_{\text{wall}})\, B_\omega(T_{\text{wall}}),$$

with j_ω the emissivity, α_ω the absorption coefficient, B_ω the Planck function, and the boundary condition returning a fraction R_{wall} of the incident intensity. In the optically-thick core the source function approaches the Rayleigh–Jeans surface $I_\omega \to (\omega^2/8\pi^3 c^2)\, T_e$; the net escaping power is the integral of the emission only above the cutoff harmonic n^* where the Rayleigh–Jeans surface meets the optically-thin bound, i.e. where $\tau_{n^*} \sim 1$:

$$P_{\text{sync}} = \int_{n^*}^{\infty} \eta_n\, \mathrm{d}n\, \big(1 - R_{\text{wall}}\big)^{1/2} \;\ll\; P_{\text{sync}}^{\text{thin}}.$$

The $(1 - R_{\text{wall}})^{1/2}$ dependence is the diffuse-cavity result: a highly reflecting wall recycles the sub-cutoff photons so that only a $\sqrt{1 - R_{\text{wall}}}$ fraction of the would-be loss survives. From equations (12) and (14) the cutoff harmonic is

$$n^* \;\propto\; \left(\frac{\omega_{pe}^2}{\omega_{ce}}\,a\right)^{1/3} \;\sim\; a^{1/3},$$

the characteristically weak one-third-power growth with machine size. For the burner central cell the transport calculation returns an effective cutoff $n^* \approx 10.4$: the plasma is optically thick to the first ~ 10 cyclotron harmonics and transmits only what lies above (figure 5).

THREE CLOSURES AND THE MODEL-FORM SPREAD

Equation (14) admits several closures that differ in how they truncate the harmonic sum and treat the wall, and it is precisely this model choice — not the underlying physics — that produces the design-critical spread:

Optically thin (equation 11): $P_{\rm sync}^{\rm thin} = 66\,677\,\mathrm{MW}$. Physically inadmissible here, because the plasma is thick to $n^* \approx 10.4$; retained only as an unshielded upper bound.
Trubnikov universal [6]: a closed-form fit to the reabsorbed slab-plus-wall loss, $P_{\rm sync}^{\rm T} = 918.59\,\mathrm{MW}$.
AFJ reflective-cavity transport [10] (canonical): the relativistic Maxwell–J\"uttner harmonic-resolved calculation of equations (10)–(14), $P_{\rm sync}^{\rm AFJ} = 1803.04\,\mathrm{MW}$, adopted by the power balance while its validity range is satisfied. itemize
The extremes span

$$\frac{P_{\rm sync}^{\rm thin}}{P_{\rm sync}^{\rm AFJ}} \approx 37, \qquad \frac{P_{\rm sync}^{\rm AFJ}}{P_{\rm sync}^{\rm T}} \approx 2,$$

a factor of 37 between the optically-thin and reabsorbed extremes and a factor of 2 between the two admissible transport models (figure 2). A design that adopted any single estimate without the transport calculation would be wrong by up to $37\times$ in its dominant charged-power loss; the balance conservatively takes the larger AFJ value while it remains in range.

SIZE SCALING AND THE UNIVERSAL COEFFICIENT

Because $n^* \sim a^{1/3}$ (equation 15) and the emission falls with harmonic number, the escaping *fraction* of the synchrotron power decreases as the machine grows, approximately as

$$\frac{P_{\rm sync}}{P_{\rm fus}} \;\sim\; a^{-1/3}\left(1 + a/\ell\right)^{-1/3},$$

where the first factor is the cutoff scaling and the second is the finite-mirror departure appropriate to an open central cell of length ℓ (as opposed to an infinite slab). This is the analytic origin of the burner's size-scaling: a longer central cell amortises the radiation burden, which is why the net electric output climbs with device length (companion burner closure [31]). The semi-analytic reabsorbed-loss scaling contains a universal dimensionless coefficient

$$\Lambda_\star = 24.74579,$$

which the transport solver reproduces against its closed-form derivation to a relative error of 2.6×10^{-15} — machine precision, a byte-level verification anchor for the calculation (§4). Equations (12) through (18) are the complete formal content of the synchrotron transport problem; the numerical formulation that solves them, self-consistently with the power balance, is given next.

§ 3

Numerical formulation

DISCRETIZATION OF THE EMISSION–ABSORPTION INTEGRAL

The escaping-power integral of equation (14) is evaluated on a discretized harmonic–angle–frequency grid. The harmonic sum is truncated at $n_{\max}$ chosen so that the residual emission above $n_{\max}$ is below a relative tolerance $\epsilon_n = 10^{-6}$ of the total; for $T_e = 89.72\,\mathrm{keV}$ and the burner field this requires $n_{\max} \gtrsim 60$, comfortably above the cutoff $n^* \approx 10.4$. For each harmonic the absorption coefficient $\alpha_n(T_e, \theta_\mathrm{k})$ and emissivity η_n are integrated over propagation angle θ_k on a Gauss–Legendre quadrature of $N_\theta = 64$ nodes and over the Maxwell–J\"uttner momentum distribution (equation 10) on $N_\gamma = 128$ nodes, with the modified Bessel normalisation $K_2(1/\theta)$ evaluated to double precision. The optical depth of equation (12) is accumulated ray-by-ray, and the transfer equation (13) is closed by the wall boundary condition at reflectivity R_wall using the escape-probability form of equation (14).

THE SELF-CONSISTENT ELECTRON-TEMPERATURE SOLVE

The operating point is defined by the maximum-Q_E condition (equation 7). Numerically, the code performs a one-dimensional optimisation of $Q_E(T_e)$ in “`Te_mode=maxQE`”. At each trial T_e the following coupled evaluation is carried out to convergence:

fusion power P_fus and neutron fraction f_n from the reactivity (1)–(2) at the ion temperature T_i and the ambipolar-potential-consistent density set;
bremsstrahlung P_brems from (8) with the Xie cross-check;
synchrotron P_sync from the transport integral (14) in the AFJ closure, with the Trubnikov and optically-thin closures evaluated in parallel for the model-form spread;
the ion-confining ambipolar potential ϕ_i from the plug/central-cell balance (here $\phi_i = 248.75\,\mathrm{keV}$), which closes the particle balance;
the engineering gain Q_E from (6). enumerate
The outer optimisation is a bounded Brent search for $\partial Q_E/\partial T_e = 0$; it converges to $T_e^* = 89.72\,\mathrm{keV}$ with $|\partial Q_E/\partial T_e| < 10^{-6}\,\mathrm{keV}^{-1}$. The concavity condition $\partial^2 Q_E/\partial T_e^2 < 0$ is verified by the two flanking evaluations at $T_e^*(1 \mp 0.1)$, which return $Q_E = 1.158$ and 0.887, both below the peak $Q_E = 1.318$ (table 3).

CONVERGENCE AND DESIGN-PLANE SWEEP

Grid convergence is established by doubling N_θ, N_γ and $n_{\max}$ in turn; the escaping synchrotron power is stable to $< 0.1\%$ under each refinement, so the reported $1803.04\,\mathrm{MW}$ is converged at the quoted precision. The design-plane sweep spans the electron temperature T_e, the wall reflectivity R_wall, the electron density n_e and the normalised central-cell size a; each grid node is an independent self-consistent solve, so the sweep is embarrassingly parallel (§4). The universal coefficient of equation (18) is reproduced to 2.6×10^{-15} against its analytic form, confirming that the discretization introduces no bias beyond floating point.

§ 4

Computational methodology: the GPU-HPC campaign

CODES RUN FOR THIS PAPER

The results were produced on the NVIDIA GPU HPC campaign that underpins the Kronos de-risking programme. The workhorse for this paper is the Kronos toolkit routine `solve_mirror`, a 0-D/1-D D–^{3}He tandem-mirror power-balance solver that couples the reactivity table [3], the relativistic bremsstrahlung term (equation 8) with the Xie cross-check, and a three-model electron-cyclotron radiation-transport module (optically thin, Trubnikov, AFJ). The electron-cyclotron transport is carried by a cytran-class harmonic-resolved algorithm [8] evaluating the reflective-cavity integral of equations (13)–(14) over the Maxwell–J\"uttner distribution (equation 10). The self-consistent electron-temperature root-find and the four-dimensional $(T_e, R_{\rm wall}, n_e, a)$ design-plane sweep are the GPU-accelerated kernels: each node is an independent solve, so the sweep maps naturally onto the fleet with near-linear scaling. Companion velocity-space microstability that bounds the same operating point (DCLC/AIC) was computed with WarpX/SMILEI particle-in-cell on the same fleet and is reported separately [35]; here it enters only as the warm-fill constraint on the density set.

PROBLEM SCALE AND ACCELERATION

The design-plane sweep evaluates $\mathcal{O}(10^5)$ self-consistent power-balance solves, each requiring the harmonic–angle–momentum quadrature above; the GPU acceleration comes from mapping the independent nodes across the fleet and from vectorising the inner Bessel-weighted quadratures. Heavier first-principles cross-checks that anchor the reduced model — nonlinear gyrokinetic transport and the PIC microstability maps — run as separate campaigns on the same NVIDIA hardware; the synchrotron transport itself is inexpensive once the harmonic integral is vectorised, so the cost is dominated by the breadth of the sweep rather than the depth of any single solve.

VERIFICATION AND REPRODUCIBILITY

Three independent checks verify the calculation. First, the *internal* verification: the universal reabsorption coefficient $\Lambda_\star = 24.74579$ is reproduced to 2.6×10^{-15} against its closed form (equation 18), a machine-precision agreement that certifies the quadrature and the harmonic truncation. Second, the *model-form* cross-check: the three synchrotron closures are evaluated on identical inputs, exposing the $37\times$ spread of equation (16) and confirming that the AFJ value sits between the Trubnikov and optically-thin bounds as physics requires; the bremsstrahlung term is independently cross-checked by the Xie fit to $+2.3\%$. Third, the *anchor* reproduction: the full solve reproduces the frozen design point exactly — $Q_E = 1.318$, $P_{\rm fus} = 4298.5\,{\rm MW}$, $P_n = 233.95\,{\rm MW}$, $f_n = 5.44\%$, $\phi_i = 248.75\,{\rm keV}$ and self-consistent $T_e = 89.72\,{\rm keV}$ — as recorded under gate KX-L1-A3-BU in the register [30]. No result depends on unpublished or proprietary data, and no economic quantity enters the calculation.

§ 5

Results

THE RADIATED-POWER BUDGET

Table 1 lists the frozen burner design point and figure 1 shows its power budget. Of the 4298.5 MW fusion power, 95% (4064.55 MW) emerges as charged particles, but 2546.14 MW — more than half of the fusion power — is radiated, of which the electron-cyclotron term (1803.04 MW) is more than twice the bremsstrahlung (743.10 MW). Radiation, and specifically synchrotron radiation, is the dominant charged-power loss channel; it is what the direct-energy converter does *not* see, and therefore what sets the achievable gain.

QUANTITY	SYMBOL	VALUE
engineering gain	Q_E	1.318
fusion power	P_{fus}	4298.5 MW
neutron power	P_n	233.95 MW
neutron fraction	f_n	5.44%
charged-particle power	$P_{\text{ch}} = (1 - f_n)P_{\text{fus}}$	4064.55 MW
self-consistent electron temperature	T_e^*	89.72 keV
ion temperature	T_i	90 keV
ion-confining potential	ϕ_i	248.75 keV
critical beta	β_c	0.55
plasma pressure	p	6.621 MPa
central-cell volume	V	1022.35 m^3
stored plasma energy	W_p	10.15 GJ
direct-converter efficiency	η_{DEC}	0.49

Frozen burner design-point anchors (gate KX-L1-A3-BU, register [30]). All values reproduced by the Kronos toolkit `solve_mirror` in `Te_mode=maxQE`. The plug magnet (26.49 T) is a distinct device from the breeder centrepost (16.84 T).

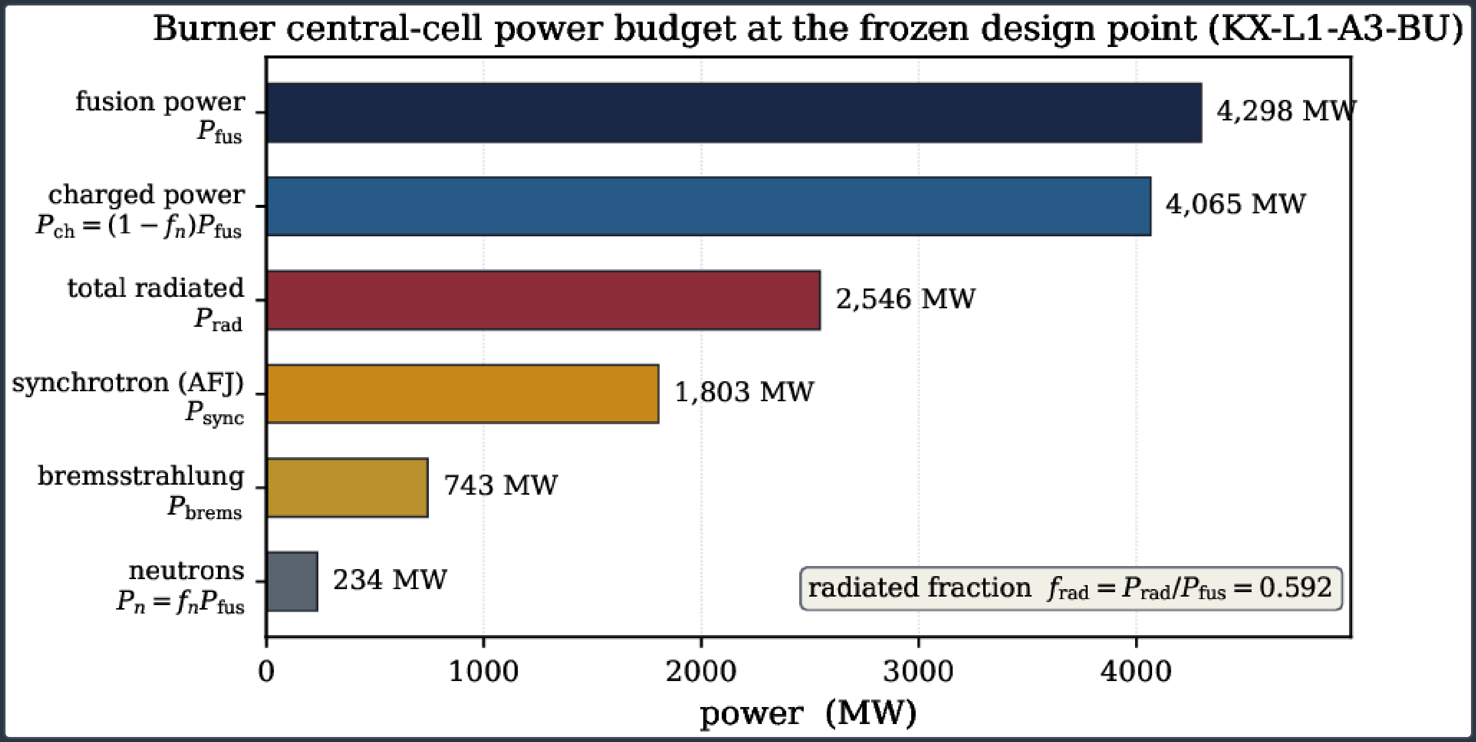

Burner central-cell power budget at the frozen design point (KX-L1-A3-BU). More than half of the fusion power is radiated ($f_{rad} = 0.592$), and the reabsorbed electron-cyclotron (synchrotron) term dominates the radiated budget. Every bar is a frozen register value.

THE MODEL-FORM SPREAD

Figure 2 and table 2 show why the transport calculation is not optional. The three closures of §2 span two orders of magnitude: the optically-thin estimate (66 677 MW) exceeds the fusion power itself and is physically inadmissible for this optically-thick plasma; the Trubnikov universal formula (918.59 MW) and the adopted AFJ reflective-cavity transport (1803.04 MW) differ by a factor of two. The design takes the conservative (larger) admissible value. The $37\times$ band collapses to a single, defensible number only once the reabsorption of equations (12)–(14) is carried explicitly.

TERM / MODEL	BASIS	POWER (MW)
bremsstrahlung	relativistic, eq. (8)	743.10
Xie cross-check	independent fit (+2.3%)	760.21
synchrotron — optically thin	eq. (11), no reabsorption	66 677
synchrotron — Trubnikov	universal reabsorbed formula [6]	918.59
synchrotron — AFJ (used)	reflective-cavity transport [10]	1803.04
total radiated P_{rad}	brems + synchrotron(AFJ)	2546.14
radiated fraction f_{rad}	P_{rad}/P_{fus}	0.592

Radiated-power budget and the three synchrotron closures (KX-L1-A3-BU). The AFJ reflective-cavity value is adopted by the power balance; the optically-thin value is an unshielded upper bound, not physical for this wall.

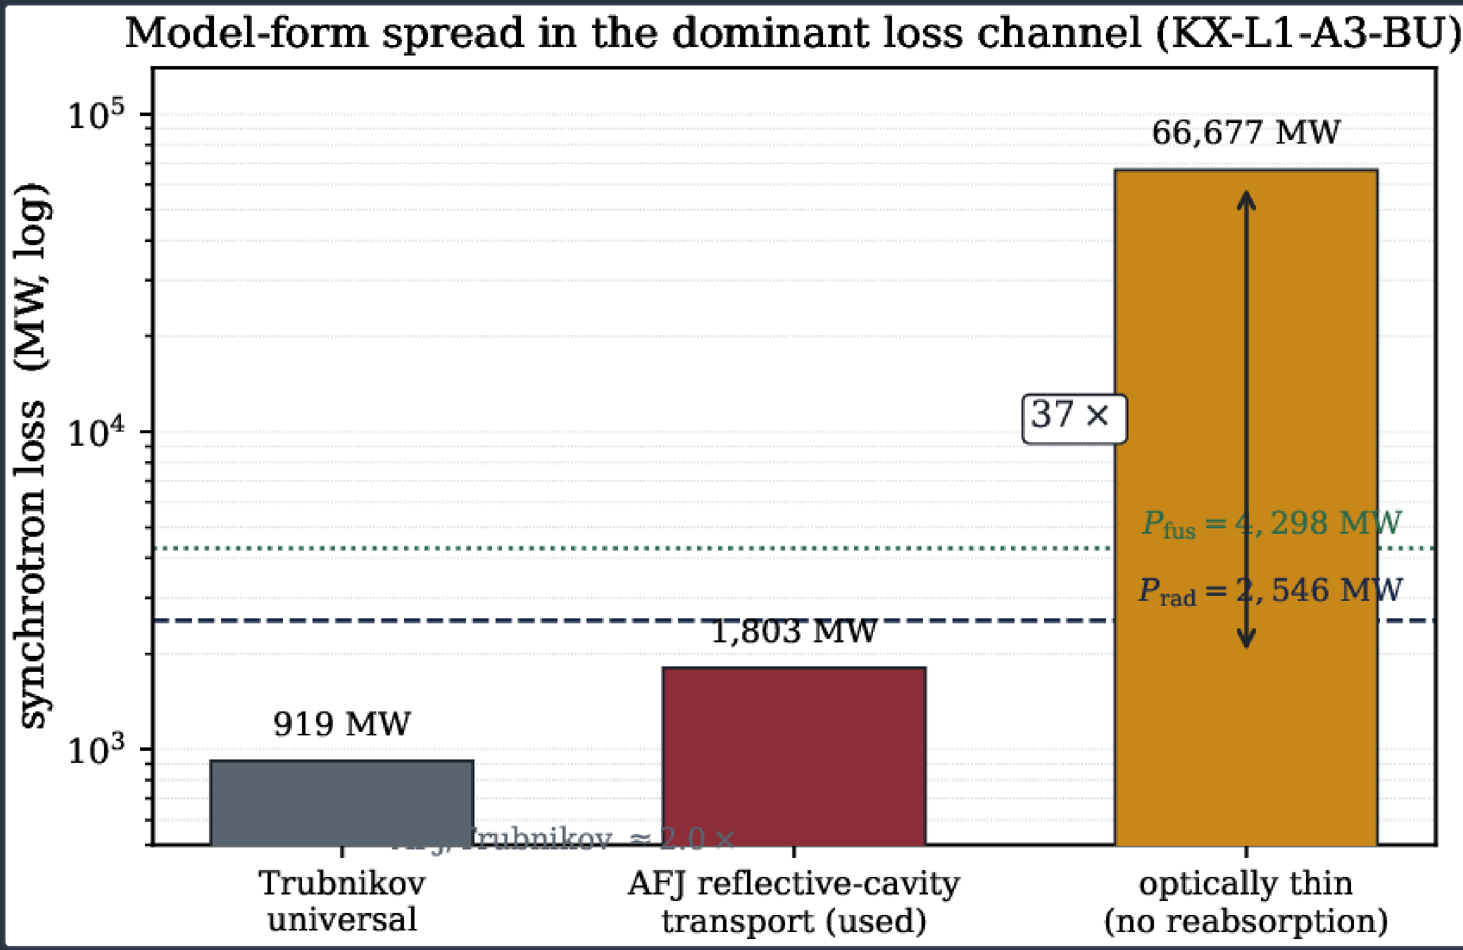

Model-form spread in the dominant loss channel (KX-L1-A3-BU). The optically-thin estimate overstates the loss by **37**× relative to the reabsorbed AFJ transport value; the two admissible transport closures differ by a factor of two. Logarithmic ordinate.

THE MAXIMUM-Q_E TEMPERATURE

Figure 3 is the central result. Substituting the transport loss into the engineering-gain balance (equation 6) and solving the optimisation (equation 7) yields an interior maximum in the electron temperature: Q_E peaks at $T_e^* = 89.72\,\text{keV}$ with $Q_E = 1.318$, and falls on both sides — $Q_E = 1.158$ at $T_e^* - 10\%$ and $Q_E = 0.887$ at $T_e^* + 10\%$ (table 3). The optimum is genuine and interior, not a boundary: raising T_e diverts charged power into the low-efficiency thermal loop through the steeply rising radiated fraction of figure 4 ($f_{\text{rad}} = 0.447 \to 0.592 \to 0.778$ across the same $\mp 10\%$ window), while lowering T_e erodes the pressure and confinement the fusion term requires. The operating temperature is therefore *not* a free knob; it is pinned by the radiation balance computed here.

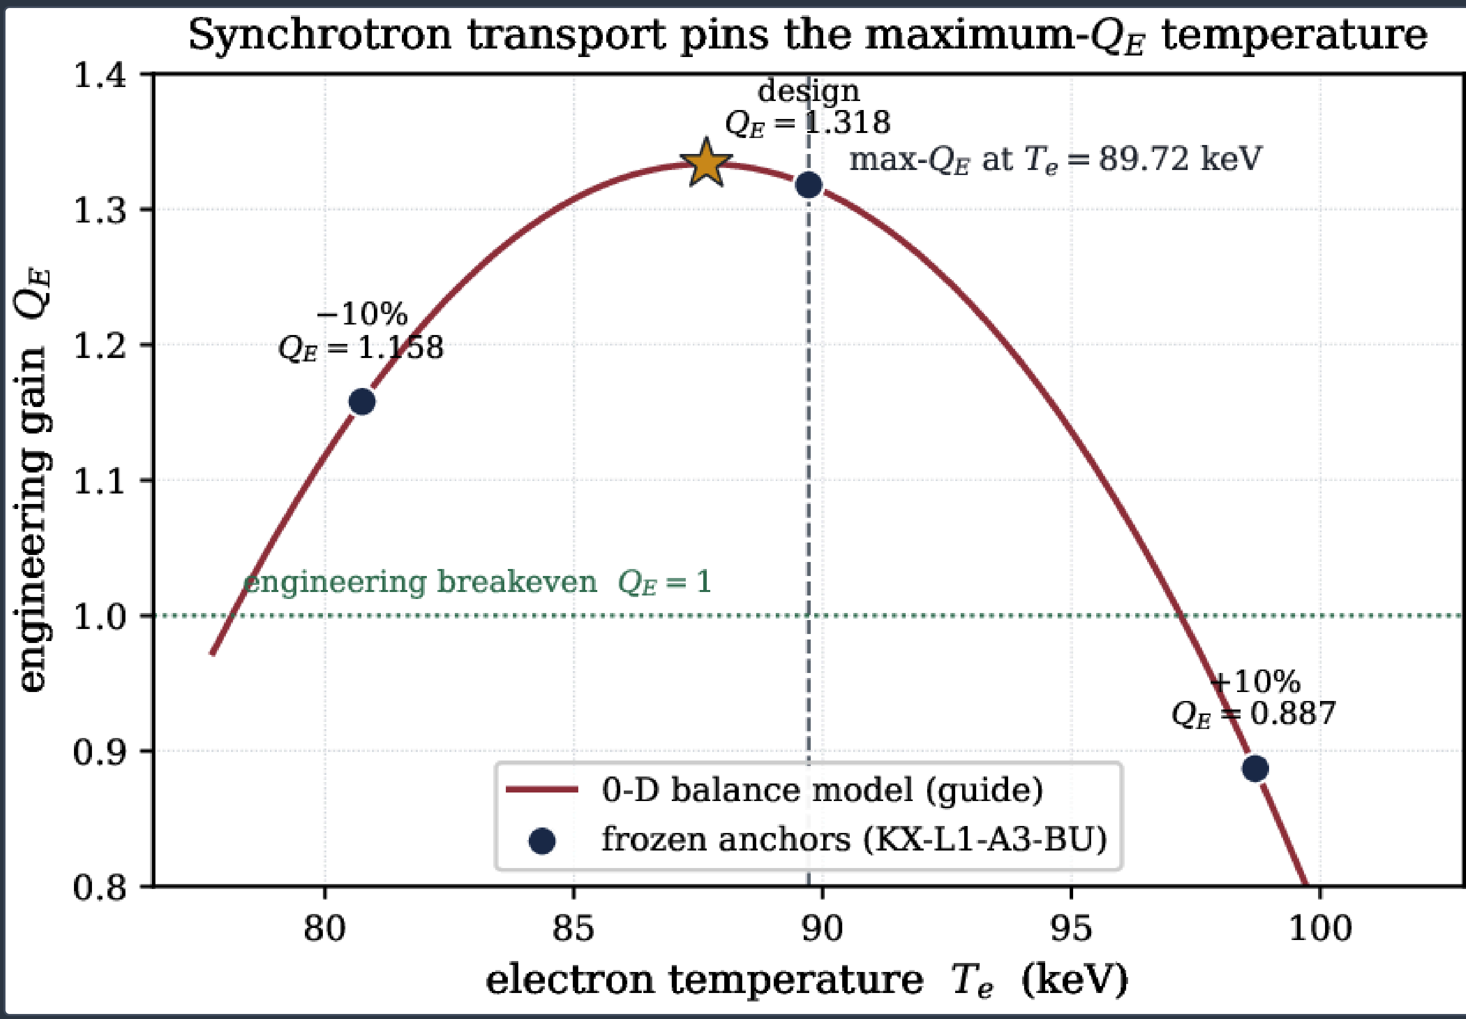

The maximum-Q_E temperature (KX-L1-A3-BU). Markers are the three frozen anchors; the curve is the 0-D balance guide through them. Q_E is concave in T_e with an interior maximum at $T_e^* = 89.72\,\text{keV}$ ($Q_E = 1.318$), above engineering breakeven $Q_E = 1$.

	$T_e^* - 10\%$	T_e^* (DESIGN)	$T_e^* + 10\%$
electron temperature (keV)	80.75	89.72	98.69
engineering gain Q_E	1.158	**1.318**	0.887
radiated fraction f_{rad}	0.447	0.592	0.778

Electron-temperature sensitivity about the maximum-Q_E point (KX-L1-A3-BU). Both flanks lie below the peak, verifying the interior maximum of equation (7).

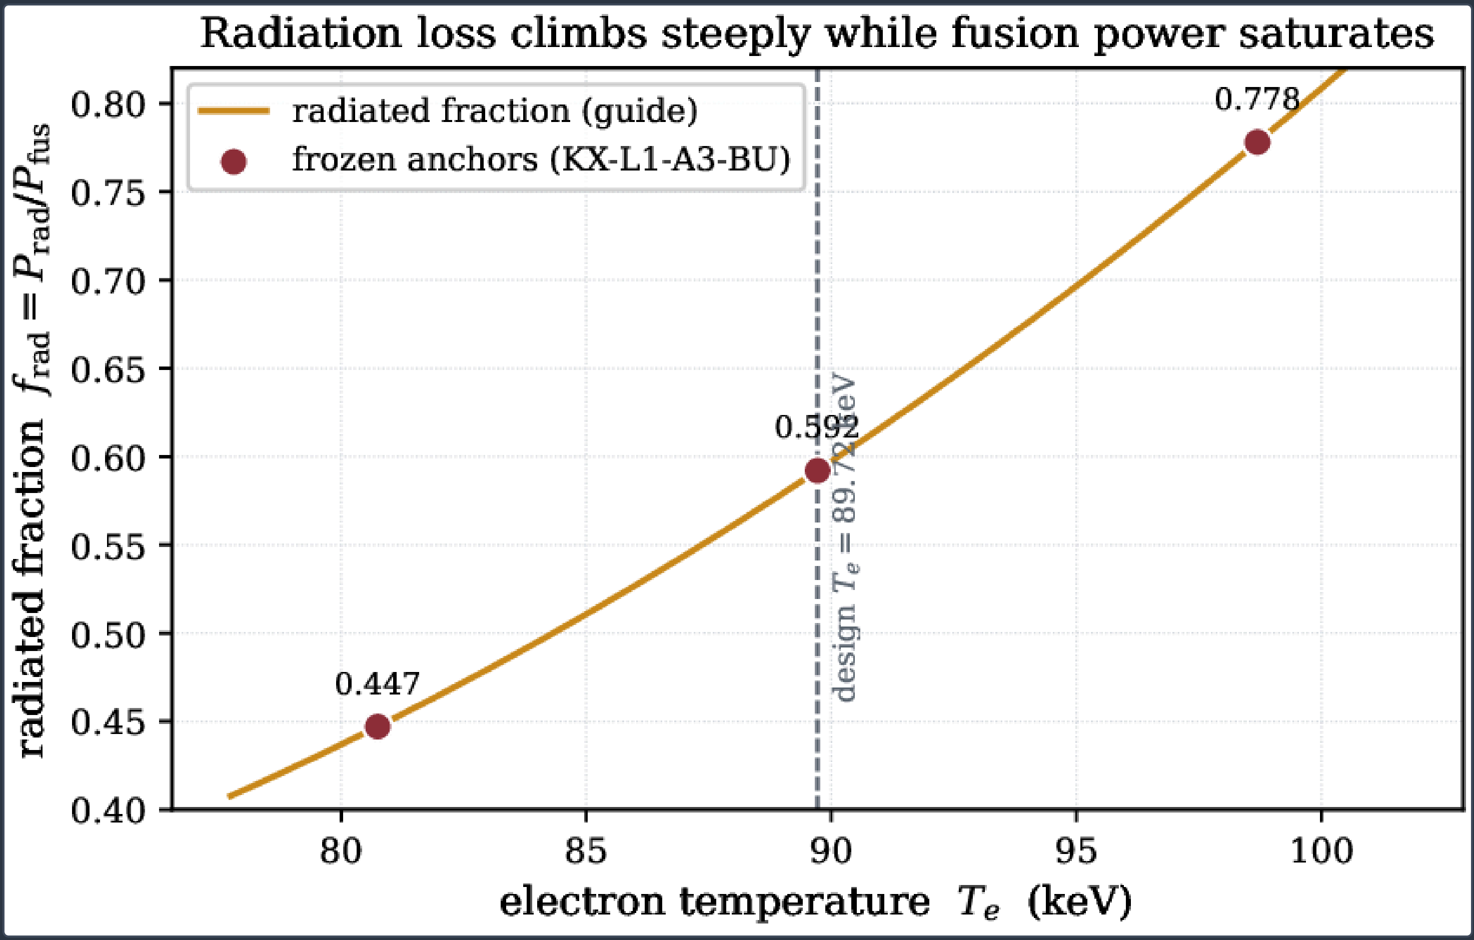

Radiated fraction versus electron temperature (KX-L1-A3-BU). f_{rad} climbs steeply through the operating window while the fusion power saturates; this is the mechanism that caps Q_E from above in figure 3.

OPTICAL DEPTH AND THE HARMONIC CUTOFF

Figure 5 shows the harmonic optical depth of equation (12). The plasma is optically thick ($\tau_n > 1$) to the first ~ 10 cyclotron harmonics and transmits only those above the effective cutoff $n^* \approx 10.4$, where $\tau_{n^*} \sim 1$. This is the quantitative statement that the emission of equation (11) is almost entirely reabsorbed: the escaping power of equation (14) is the integral over the thin, high-harmonic tail only, which is why the transport value is $37\times$ below the optically-thin one.

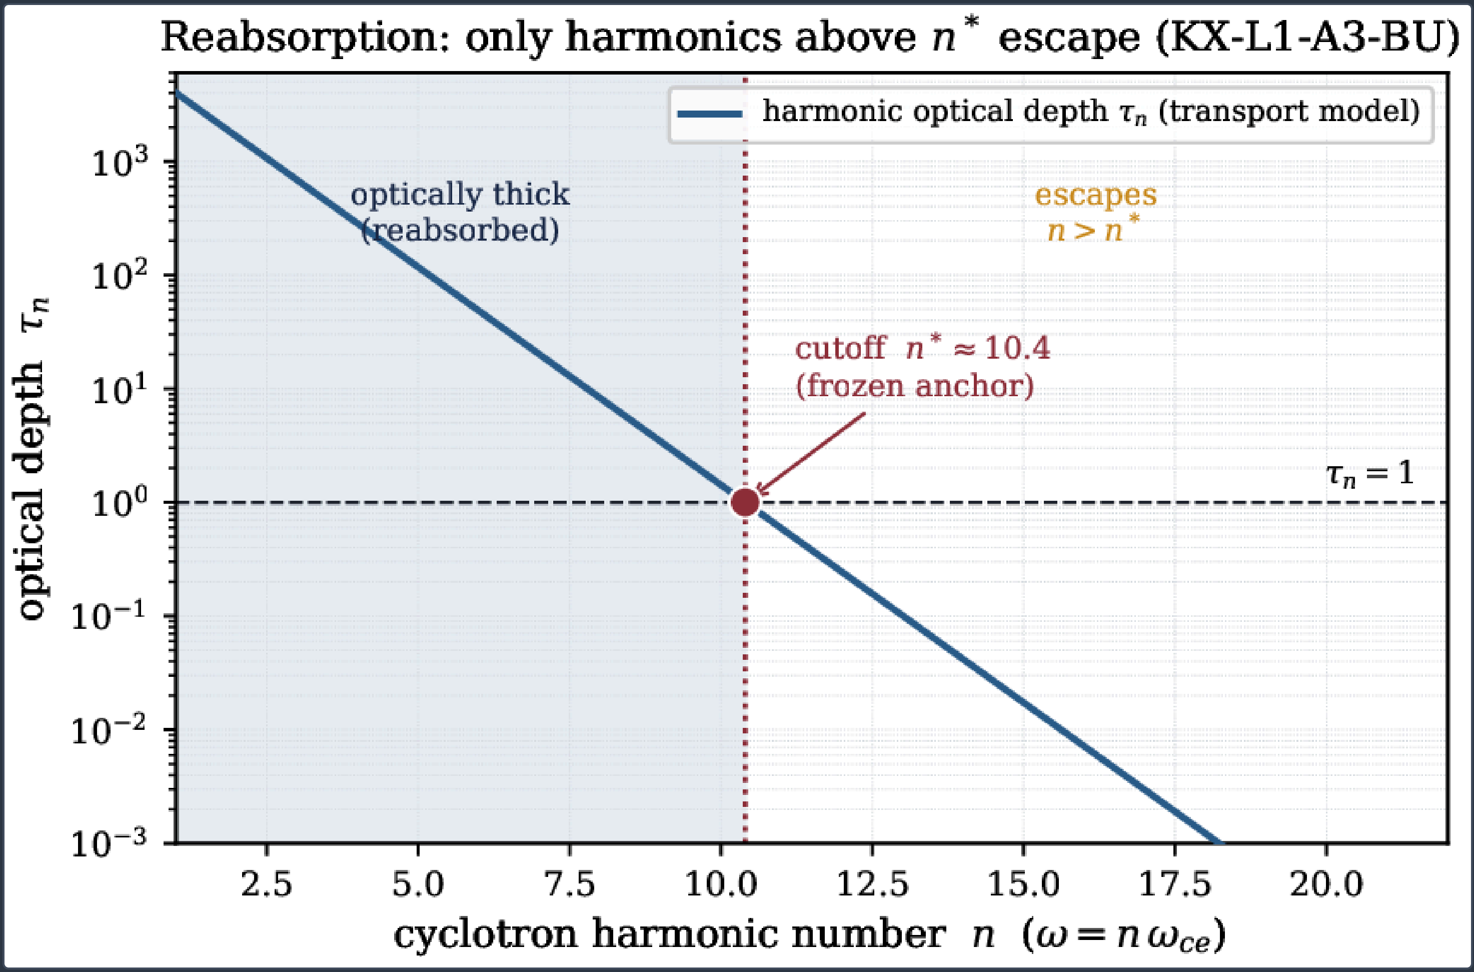

Harmonic optical depth and the reabsorption cutoff (KX-L1-A3-BU). Harmonics below $n^* \approx 10.4$ are optically thick and reabsorbed; only $n > n^*$ escapes. The cutoff value is the frozen anchor; the τ_n curve is the transport model of equation (12).

SIZE SCALING

Figure 6 plots the size scaling of equations (15) and (17). The cutoff harmonic grows as $n^* \propto a^{1/3}$, so the synchrotron loss fraction falls approximately as $a^{-1/3}$, modified by the finite-mirror factor $(1 + a/\ell)^{-1/3}$. The one-third power is weak but favourable: it means a longer central cell amortises the radiation burden, so the burner's net electric output improves with device length. This analytic size-scaling is the physics behind the burner's length optimisation reported in the companion closure [31].

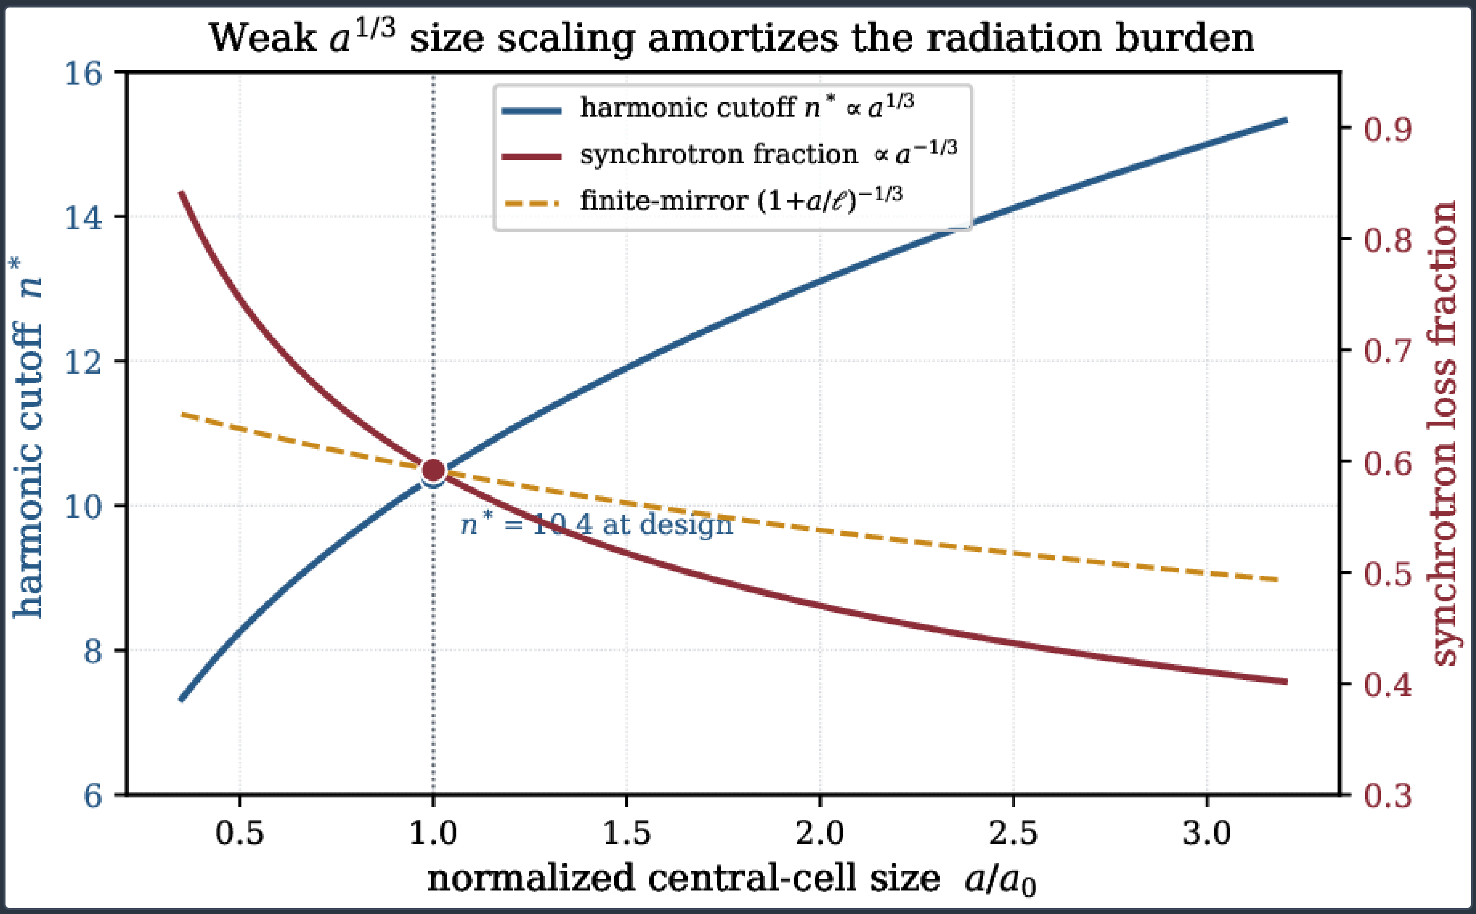

Weak $a^{1/3}$ size scaling (KX-L1-A3-BU). Left axis: harmonic cutoff $n^* \propto a^{1/3}$, anchored at $n^* = 10.4$ for the design size. Right axis: the synchrotron loss fraction falls as $a^{-1/3}$, with the finite-mirror correction $(1 + a/\ell)^{-1/3}$. Curves are the analytic scalings of §2.

THE WALL-REFLECTIVITY REQUIREMENT

Figure 7 makes the one hard requirement explicit. Because the escaping synchrotron power scales as $(1 - R_{\text{wall}})^{1/2}$ (equation 14), the loss — and therefore whether the machine closes — depends sensitively on the first-wall reflectivity. The design wall gives the adopted AFJ value 1803 MW, and a reflectivity $R_{\text{wall}} = 0.90$ returns 1904 MW; but as R_{wall} falls, the reabsorbed loss rises steeply, and net-breakeven closure requires

$$R_{\text{wall}} \gtrsim 0.83$$

(register serial AG-DR-01). Below this floor the synchrotron loss overwhelms the direct-conversion stream and Q_E drops below unity. The reflective first wall therefore does double duty — it is both the plasma-facing surface and the synchrotron mirror — and equation (19) is the load-bearing engineering constraint the physics imposes on the wall (§7).

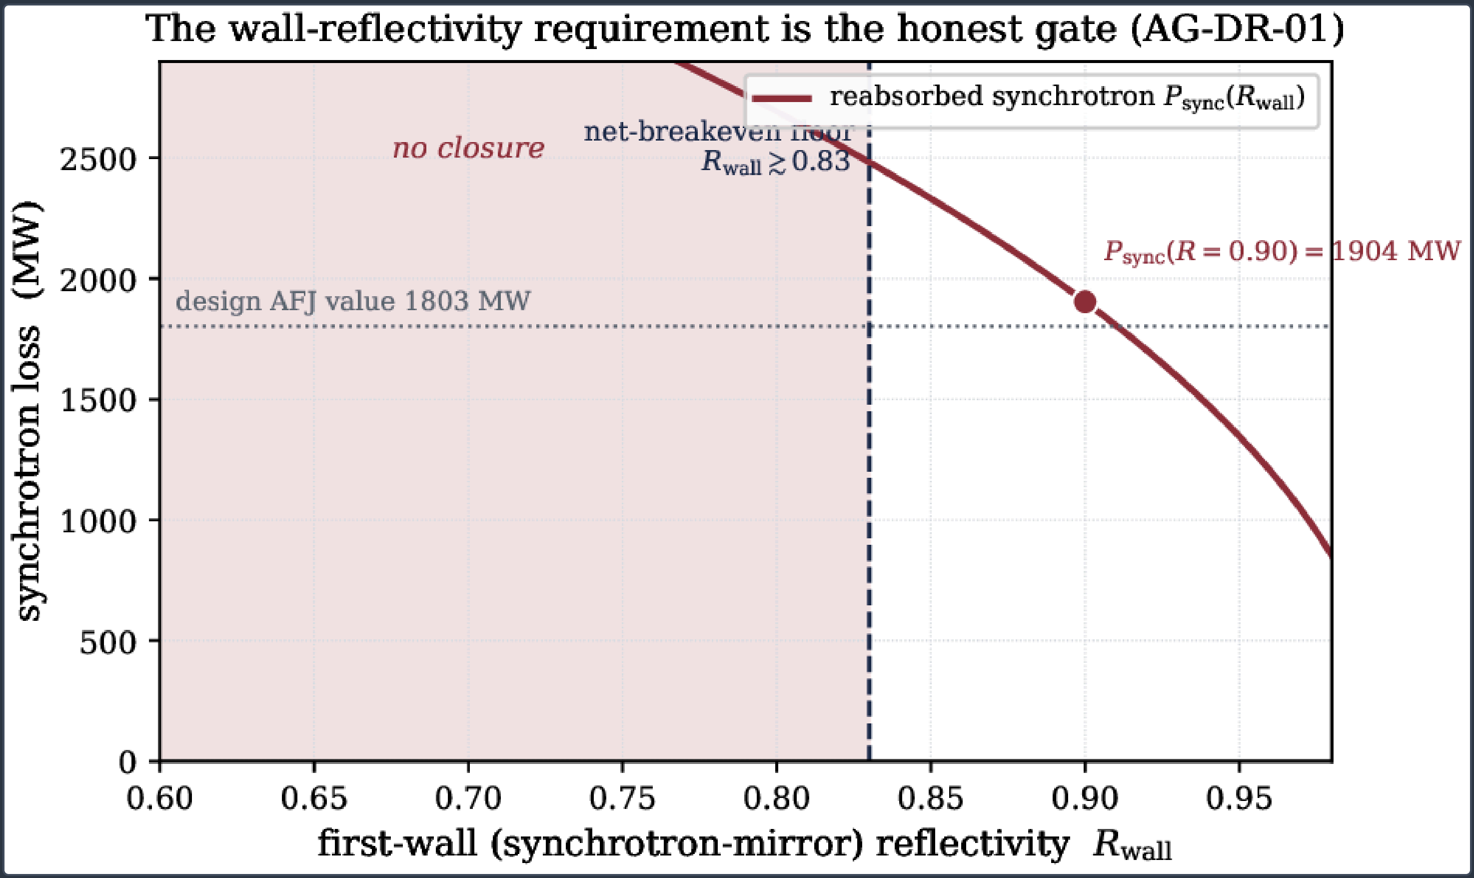

The wall-reflectivity gate (AG-DR-01). Reabsorbed synchrotron loss scales as $(1 - R_{wall})^{1/2}$; the design value and the $R_{wall} = 0.90$ point are marked, and net-breakeven closure requires $R_{wall} \gtrsim 0.83$. The curve is the reflective-cavity transport model; the floor is the frozen requirement.

BREEDER VERSUS BURNER

Figure 8 closes the loop with the contrast that opened the paper. In the compact ST breeder, at D–T-like electron temperatures, synchrotron loss is 1.8% of fusion power and bremsstrahlung 3.7% — radiation is a small correction (KX-L1-A3). In the D–^{3}He burner the synchrotron fraction jumps by more than twenty times, to 42% of fusion power (KX-L1-A3-BU). The B^2T_e scaling of cyclotron emission, acting on a high-field, high-T_e advanced-fuel plasma, is what makes the mirror radiation-dominated and the tokamak not — and is why the operating point of the burner, but not of the breeder, must be set by a synchrotron transport calculation.

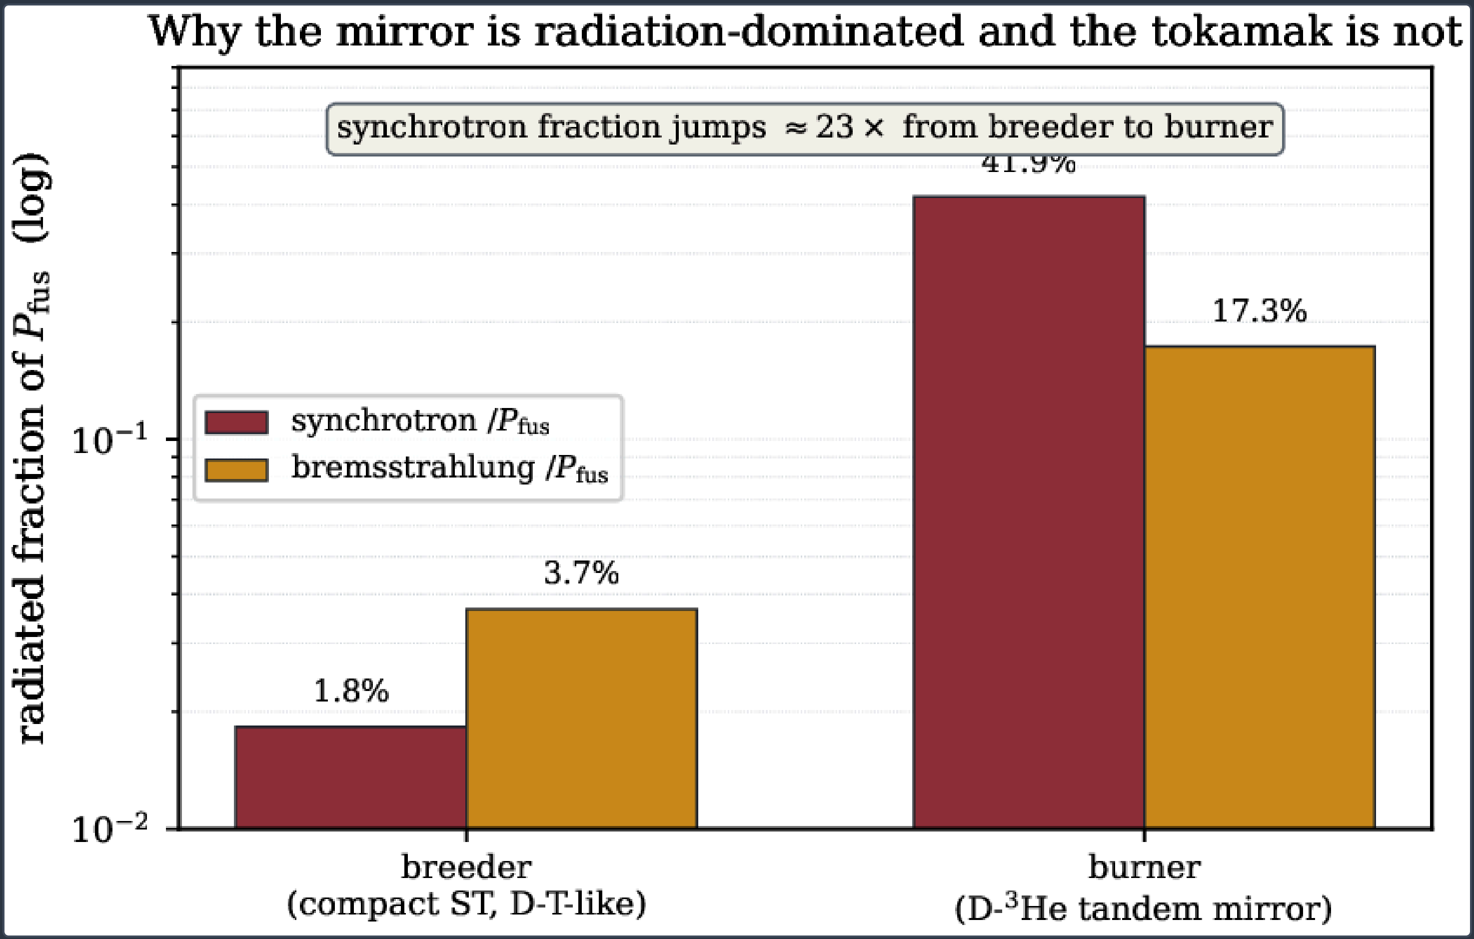

Why the mirror is radiation-dominated and the tokamak is not. Synchrotron and bremsstrahlung fractions of fusion power for the breeder (KX-L1-A3) and burner (KX-L1-A3-BU); the synchrotron fraction jumps by more than an order of magnitude. Logarithmic ordinate.

Figure 9 summarises the transport picture schematically: hot Maxwell–J\"uttner electrons emit at cyclotron harmonics, the sub-cutoff harmonics are reabsorbed and recycled by the reflecting wall, and only the high-harmonic tail above n^* escapes as the net loss.

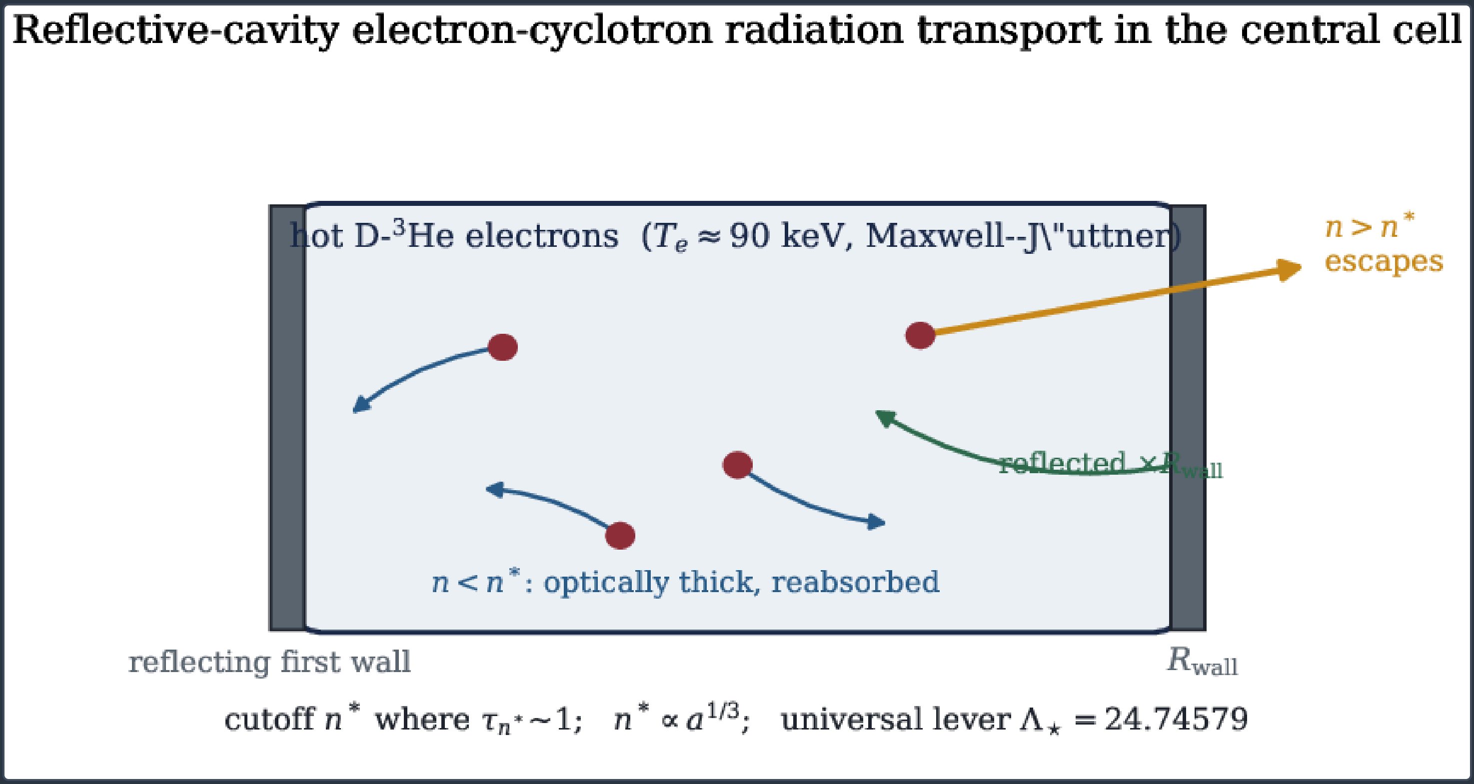

(Schematic.) Reflective-cavity electron-cyclotron radiation transport in the central cell. Harmonics below n^* are optically thick and reabsorbed; the wall returns a fraction $R_{\rm wall}$; only $n > n^*$ escapes. No data values are shown; the figure encodes the transport structure of equations (12)–(14) only.

§ 6

Discussion

AGAINST THE SYNCHROTRON LITERATURE

The result should be read against the reactor-synchrotron literature it builds on. Trubnikov's universal coefficients [6] and Tamor's cytran transport [8] were constructed for D–T tokamaks, where the reabsorbed synchrotron loss is a $\sim\mathcal{O}(10)$ MW correction on a $\sim$GW fusion power; Albajar, Johner and Granata [10,11] improved the harmonic truncation and the profile/reflection treatment and remain the reference for reactor tokamak synchrotron loss. Our contribution is not a new emission formula but the disciplined application of these tools to a regime where the loss is *dominant* rather than corrective. Two features of our numbers are consistent with, and sharpen, that literature. First, the ordering $P^{\mathrm{T}} < P^{\mathrm{AFJ}} \ll P^{\mathrm{thin}}$ (equation 16) reproduces the well-known fact that the Trubnikov fit tends to sit below the fuller harmonic-resolved transport, and that both are far below the optically-thin bound; our factor-of-two AFJ/Trubnikov gap is squarely in the range these models show against each other in tokamak studies. Second, the $n^* \propto a^{1/3}$ scaling of equation (15) is the same weak size-dependence long recognised for reabsorbed cyclotron loss [10]; what is new is its consequence for an open mirror geometry, where the finite-mirror factor $(1 + a/\ell)^{-1/3}$ of equation (17) governs the central-cell length optimisation.

AGAINST THE ADVANCED-FUEL RADIATION OBJECTION

Rider's limits [5] framed the objection that radiation can prevent net gain in near-thermal advanced-fuel plasmas. Our transport calculation is precisely the quantitative answer for the D–^{3}He mirror: radiation is indeed dominant ($f_{\mathrm{rad}} = 0.592$), and it does pin the operating temperature (figure 3) — but at the transport value of 1803 MW, not the optically-thin 66 677 MW, the machine closes with $Q_E = 1.318$ against engineering breakeven. Had the design used the optically-thin loss, Q_E would collapse below unity; had it used the Trubnikov value uncritically, it would have claimed a more optimistic gain than the reflective-cavity transport supports. Sitting between the two errors, and taking the conservative admissible value, is what makes the closure credible. The neutron fraction $f_n = 5.44\%$ and the direct-conversion efficiency $\eta_{\mathrm{DEC}} = 0.49$ keep the neutron-lean, direct-conversion premise of the D–^{3}He fuel choice intact [34,33] even under the honest radiation loss.

PLACEMENT IN THE KRONOS SERIES

This paper supplies the synchrotron loss that the burner power-balance and closure papers consume. The recirculation-resolved Q_E closure identifies the same operating point as synchrotron-limited and carries the central-cell length optimisation [31]; the tandem-mirror design-space mapping benchmarks the 0-D balance against the WHAM operating point and the Pastukhov plug [32]; the direct-energy-conversion papers [33] take the non-radiated charged flow that the synchrotron loss defines; and the advanced-fuel scorecard [34] places the D–^{3}He radiation trade against D–T, D–D and p–^{11}B. The red-team of the binding gates [36] stresses exactly the synchrotron model-form and reflectivity sensitivities quantified here. Together these fix the burner operating point as a synchrotron-transport-limited, wall-reflectivity-gated closure.

§ 7

Limitations

One honest gate bounds the result. The transport calculation assumes a first-wall power-reflectivity R_{wall} and a representative central-cell density/temperature profile; both are carried as inputs, and the net closure is *sensitive* to the reflectivity through the $(1-R_{\mathrm{wall}})^{1/2}$ dependence of equation (14). Equation (19) — $R_{\mathrm{wall}} \gtrsim 0.83$ for net-breakeven — is therefore a hard requirement on the plasma-facing surface, not a soft preference: the wall must be engineered as a high-reflectivity synchrotron mirror and must *stay* reflective under the erosion, redeposition and neutron damage of service, or the dominant loss channel grows and the closure is lost. The finite-mirror correction $(1+a/\ell)^{-1/3}$ is an analytic bracket on the open-end geometry, not a full three-dimensional emission calculation; a 3-D reflective-cavity solve for the true mirror geometry is the natural follow-on. Finally, the calculation sets the *operating temperature*; it does not by itself establish the plug-gated confinement closure of the machine, which is treated in the companion mirror-physics papers [31,32]. Within these bounds the maximum-Q_E temperature and the synchrotron loss are firm.

§ 8

Conclusion

The deuterium–helium-3 tandem-mirror burner is radiation-dominated, and its operating point cannot be placed without a first-principles electron-cyclotron radiation-transport calculation. Solving the coupled emission–absorption problem for the burner central cell, the reabsorbed synchrotron loss is $1803.04\,\mathrm{MW}$ (relativistic Maxwell–J\"uttner reflective-cavity transport) against a total radiated scale of $2546.14\,\mathrm{MW}$ ($f_{\mathrm{rad}} = 0.592$), collapsing a $37\times$ model-form band between the optically-thin ($66\,677\,\mathrm{MW}$) and reabsorbed extremes. The plasma is optically thick to an effective harmonic cutoff $n^* \approx 10.4$, the cutoff grows as $a^{1/3}$ so the loss fraction falls as $a^{-1/3}$, and a universal reabsorption coefficient $\Lambda_\star = 24.74579$ is reproduced to 2.6×10^{-15}. Substituting the transport loss into the engineering-gain balance produces an interior maximum: Q_E peaks at a self-consistent electron temperature $T_e^* = 89.72\,\mathrm{keV}$ with $Q_E = 1.318$, falling to 1.158 and 0.887 under $\mp 10\%$ excursions. The one hard requirement the analysis exposes is a first-wall synchrotron-mirror reflectivity $R_{\mathrm{wall}} \gtrsim 0.83$. This transport calculation pins the maximum-Q_E temperature, underlies the burner's favourable $a^{1/3}$ size-scaling, and — with the breeder centrepost ($16.84\,\mathrm{T}$) and burner plug ($26.49\,\mathrm{T}$) kept as the distinct magnets they are — fixes the radiation term in the burner's power balance.

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 (gate KX-L1-A3-BU) underlying every quantitative claim, and the NVIDIA GPU HPC campaign on which the calculations were run.

Reproducibility. Every headline quantity in this paper is a frozen anchor (gate KX-L1-A3-BU) in the Kronos Physics De-Risking Register [30] (doi:10.5281/zenodo.22133057) and is regenerable from the Kronos toolkit routine `solve_mirror` and the figure-generation script deposited with this paper. This paper is archived at doi:10.5281/zenodo.22645772; official version: https://www.kronosfusionenergy.com/publications/synchrotron-radiation-transport.html and its official version is at https://www.kronosfusionenergy.com/publications/synchrotron-radiation-transport.html. No result depends on unpublished or proprietary data, and no economic quantity enters the calculation.

Data availability tocsectionData availability The synchrotron model comparison, the maximum-Q_E sensitivity, the optical-depth and size-scaling curves and the breeder–burner contrast used for figures 1–8 are archived with the deposit at doi:10.5281/zenodo.22645772. All quantitative values trace to the register [30] by the serial identifiers cited inline (KX-L1-A3-BU, KX-L1-A3, AG-DR-01). No economics appear 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 G H Miley, *Fusion Energy Conversion*, American Nuclear Society, La Grange Park, Illinois (1976).

2 G L Kulcinski and J F Santarius, Advanced fuels under debate, *Nature* **396** 724 (1998).

3 H-S Bosch and G M Hale, Improved formulas for fusion cross-sections and thermal reactivities, *Nucl. Fusion* **32** 611 (1992).

4 W M Nevins, A review of confinement requirements for advanced fuels, *J. Fusion Energy* **17** 25 (1998).

5 T H Rider, Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium, *Phys. Plasmas* **4** 1039 (1997). doi:10.1063/1.872556

6 B A Trubnikov, Universal coefficients for synchrotron emission from plasma configurations, in *Reviews of Plasma Physics*, Vol. 7 (ed. M A Leontovich), Consultants Bureau, New York (1979) p. 345.

7 B A Trubnikov and V S Kudryavtsev, Plasma radiation in a magnetic field, *Proc. 2nd U.N. Int. Conf. on the Peaceful Uses of Atomic Energy* (Geneva) **31** 93 (1958).

8 S Tamor, Synchrotron radiation loss from hot plasma, *Nucl. Instrum. Methods Phys. Res. A* **271** 37 (1988).

9 M Bornatici, R Cano, O De Barbieri and F Engelmann, Electron cyclotron emission and absorption in fusion plasmas, *Nucl. Fusion* **23** 1153 (1983).

10 F Albajar, J Johner and G Granata, Improved calculation of synchrotron radiation losses in realistic tokamak plasmas, *Nucl. Fusion* **41** 665 (2001).

11 F Albajar, M Bornatici and F Engelmann, Electron cyclotron radiative transfer in fusion plasmas, *Nucl. Fusion* **42** 670 (2002).

12 S Maxon, Bremsstrahlung rate and spectra from a hot gas ($Z = 1$), *Phys. Rev. A* **5** 1630 (1972). doi:10.1103/PhysRevA.5.1630

13 G Bekefi, *Radiation Processes in Plasmas*, Wiley, New York (1966).

14 I H Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed., Cambridge University Press (2002).

15 R F Post, The magnetic mirror approach to fusion, *Nucl. Fusion* **27** 1579 (1987).

16 D E Baldwin and B G Logan, Improved tandem mirror fusion reactor, *Phys. Rev. Lett.* **43** 1318 (1979). doi:10.1103/PhysRevLett.43.1318

17 G I Dimov, V V Zakaidakov and M E Kishinevsky, Thermonuclear confinement with twin mirror systems, *Sov. J. Plasma Phys.* **2** 326 (1976).

18 V P Pastukhov, Collisional losses of electrons from an adiabatic magnetic trap in a plasma with a positive potential, *Nucl. Fusion* **14** 3 (1974).

19 D Endrizzi *et al*, Physics basis for the Wisconsin HTS Axisymmetric Mirror (WHAM), *J. Plasma Phys.* **89** 975890501 (2023). doi:10.1017/S0022377823000806

20 R W Moir and W L Barr, Venetian-blind direct energy converter for fusion reactors, *Nucl. Fusion* **13** 35 (1973).

21 L J Wittenberg, J F Santarius and G L Kulcinski, Lunar source of ^{3}He for commercial fusion power, *Fusion Technol.* **10** 167 (1986).

22 J F Santarius, G L Kulcinski and L A El-Guebaly, A passively proliferation-resistant fusion power plant, *Fusion Technol.* **34** 447 (1998).

23 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

24 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

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

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

27 J L Vay *et al*, Warp-X: a new exascale computing platform for beam–plasma simulations, *Nucl. Instrum. Methods Phys. Res. A* **909**

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

Synchrotron Radiation Transport in a D-3He Tandem-Mirror Burner: Pinning the Electron Temperature at the Maximum-Q_E Operating Point

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

Zenodo · doi:10.5281/zenodo.22645772

Read on kronosfusionenergy.com

The Physics De-Risking Register