KRONOS FUSION ENERGY

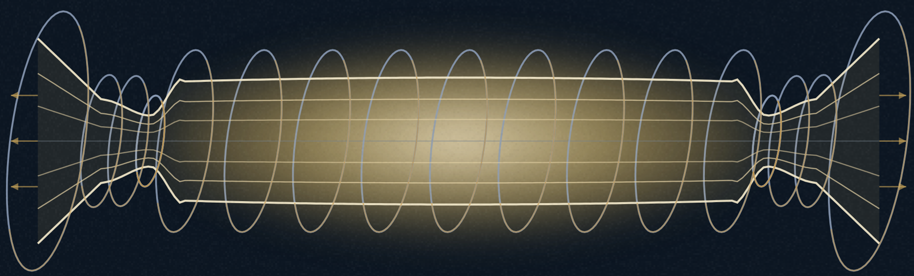

PAPER 1.2A · THE KRONOS 2026 SERIES

The Aegis Forward-Base Burner: Q_E Closure, Plug Requirement, and Mirror Confinement in an Islanded Low-Neutron D–^{3}He Tandem Mirror

Aegis is the forward-base configuration of the Kronos deuterium–helium-3 tandem-mirror burner: a fixed, hardened, grid-independent installation that turns an aneutronic-dominated...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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 Aegis Forward-Base Burner: Q_E Closure, Plug Requirement, and Mirror Confinement in an Islanded Low-Neutron D–^{3}He Tandem Mirror

Aegis is the forward-base configuration of the Kronos deuterium–helium-3 tandem-mirror burner: a fixed, hardened, grid-independent installation that turns an aneutronic-dominated fusion source into firm electrical power for a defence site without a fuel-convoy tail. This paper is a comprehensive, review-article treatment of the integrated physics case for that machine, and it separates, throughout, what is robust from what is conditional. At the frozen design point the plant reaches an engineering gain $Q_E = 1.318$ at fusion power 4298.5 MW, of which 3868.0 MW is the charged-particle $\mathrm{D} + {}^{3}\mathrm{He} \to p\,(14.68\,\mathrm{MeV}) + \alpha$ channel and only 233.95 MW is neutron power — a neutron fraction $f_n = 5.44\,\%$, the low-activation advantage that motivates the fuel and the siting. The binding physics gate is the end-plug potential: confining the central-cell ions requires a plug electrostatic potential $\phi_i = 248.75$ keV, which through the ambipolar Boltzmann relation demands a plug-to-central density ratio $n_p/n_c = \exp(e\phi_i/T_e) = 16$ at carrier temperature $T_e = 89.7$ keV — a requirement more than two orders of magnitude beyond the demonstrated tandem-mirror record, retired only by a dedicated plug-potential experiment. We give the full governing-equation derivations (Bosch–Hale reactivity averaging, adiabatic loss-cone confinement, the ambipolar Boltzmann relation, the Pastukhov end-loss law, the recirculation-limited gain, synchrotron reabsorption, the drift-cyclotron-loss-cone dispersion, central-cell interchange stability, and RF alpha-channeling), a GPU-accelerated computational methodology naming the real codes, and results across twenty-six figures and eight tables. A Saltelli–Sobol analysis confirms the plug ratio is the first-order driver of Q_E (first-order index $S_1 = 0.41$, total-order $S_T = 0.87$); a 40 000-sample Monte-Carlo scan over the design ranges returns $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$, while a net-power uncertainty study near the operating cliff returns $P(\mathrm{net} > 0) = 73\,\%$. We are explicit about a second binding lever: direct energy conversion is a declared, load-bearing input on a conservative-to-ideal ladder ($\eta_{\mathrm{dec}} = 0.49$ conservative canon, 0.70 design basis, 0.85 proton-only ceiling), with a net-electric break-even near 0.52; the design point closes at the design basis and the conservative floor does not. The loss-cone microstability sits on the drift-cyclotron-loss-cone knee at a $6.25\,\%$ warm-plasma fill – identically $1/(n_p/n_c)$ – a maintained threshold rather than a passive margin, and a residual anisotropy-driven Alfv\'en-ion-cyclotron branch persists. The Aegis product wrap is stated quantitatively: a black-start whose direct-converter output crosses the start-up draw at $\approx 42\,\mathrm{s}$, an islanded steady state at $+104$ $\mathrm{MW_e}$, a fleet capacity factor of 68–96 % across redundancy tiers, and a one-year grid-independent fuel package of order one tonne. Every net-power number is contingent on retiring the plug gate; the neutron-source performance and the low-activation siting case are not. No economic quantity is reported.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645693 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: tandem mirror · deuterium--helium-3 · plug potential · loss-cone microstability · direct energy conversion · forward-base power · islanded microgrid · low-neutron fusion · advanced fuel · Pastukhov confinement

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)

Nomenclature tocsectionNomenclature

@L>X@ SYMBOL		DEFINITION
Q_E		plant engineering gain (gross electric out / recirculating in)
P_{fus}		total fusion power; P_{ch} charged, P_n neutron
f_n		neutron fraction of total fusion power
$\langle\sigma v\rangle_r$		Maxwellian-averaged reactivity of channel r
n_p/n_c		plug-to-central-cell density ratio (the binding gate)
ϕ_i		ion-confining electrostatic plug potential
T_e,	T_i	carrier electron / central-cell ion temperature
R_m		mirror ratio $B_{\text{max}}/B_{\text{min}}$
θ_L		loss-cone half-angle, $\arcsin R_m^{-1/2}$
τ_p		Pastukhov (plugged) axial ion-confinement time
τ_{ii}		ion–ion Coulomb collision time
f_s		energetic-stabiliser (warm-plasma) fill fraction inside the loss cone
η_{dec}		direct-energy-converter module efficiency (declared input)
η_{th}		thermal-cycle efficiency on residual heat
f_{dir}		directed (recoverable) fraction of the axial loss stream
P_{ax}		axial (end-loss) power leaving the mirror ends
$P_{\text{plug}}, P_{\text{aux}}$		recirculating plug and auxiliary powers
γ/ω_{ci}		normalised loss-cone (DCLC/AIC) growth rate
β_c		central-cell plasma beta
R_{wall}		synchrotron wall (reflective-cavity) reflectivity
a/ρ_i		device-radius-to-ion-gyroradius (magnetisation) ratio
n^*		synchrotron reabsorption cutoff harmonic

§ 1

Introduction

THE FORWARD-BASE POWER PROBLEM AND THE AEGIS CONFIGURATION

A fixed forward defence installation is a demanding electrical load: it must run through a grid outage, it must not depend on a fuel convoy that is itself a target, and it must accept a directed-energy or high-power-radar duty that swings by tens of megawatts on a timescale the prime mover has to follow. Diesel and gas turbines answer the first requirement and fail the second; a fission microreactor answers both but carries a spent-fuel and proliferation signature that constrains where it may be sited. Aegis is the Kronos answer: the forward-base housing of a deuterium–helium-3 (D–^{3}He) tandem-mirror burner, a fixed, hardened installation whose fuel is a sub-cubic-metre annual charge and whose neutron output is a small fraction of a same-power deuterium–tritium (D–T) machine. Aegis and its sibling data-centre configuration, MetroVolt, are the *same* burner core in two housings; this paper treats the physics of the core and the forward-base product wrap, and its companion treats the data-centre wrap [51].

The attraction of the fuel is aneutronicity, and its difficulty is confinement. The primary reaction

$$\mathrm{D} + {}^{3}\mathrm{He} \;\rightarrow\; p\,(14.68\,\mathrm{MeV}) \;+\; \alpha\,(3.67\,\mathrm{MeV})$$

releases all of its energy in charged particles, so a device that burns it produces neutrons only from the unavoidable D–D side reactions and can, in principle, recover the charged-particle power directly rather than through a steam cycle. But the reaction peaks near $T_i \sim 60\text{–}100$ keV, an order of magnitude hotter than a D–T plasma, and at those temperatures the electrostatic well that a tandem mirror uses to plug its ends must be correspondingly deep. The central thesis of this paper is that the resulting *plug-potential requirement* is the single binding physics gate of the machine, and that an honest forward-base case names it as such rather than hiding it behind an engineering-gain headline. Figure 1 shows the machine schematically: a long, low-field central cell in which the D–^{3}He mixture burns, closed at each end by a high-field plug that raises the confining potential, and terminated by an expander and a direct-energy converter that recover the escaping charged stream.

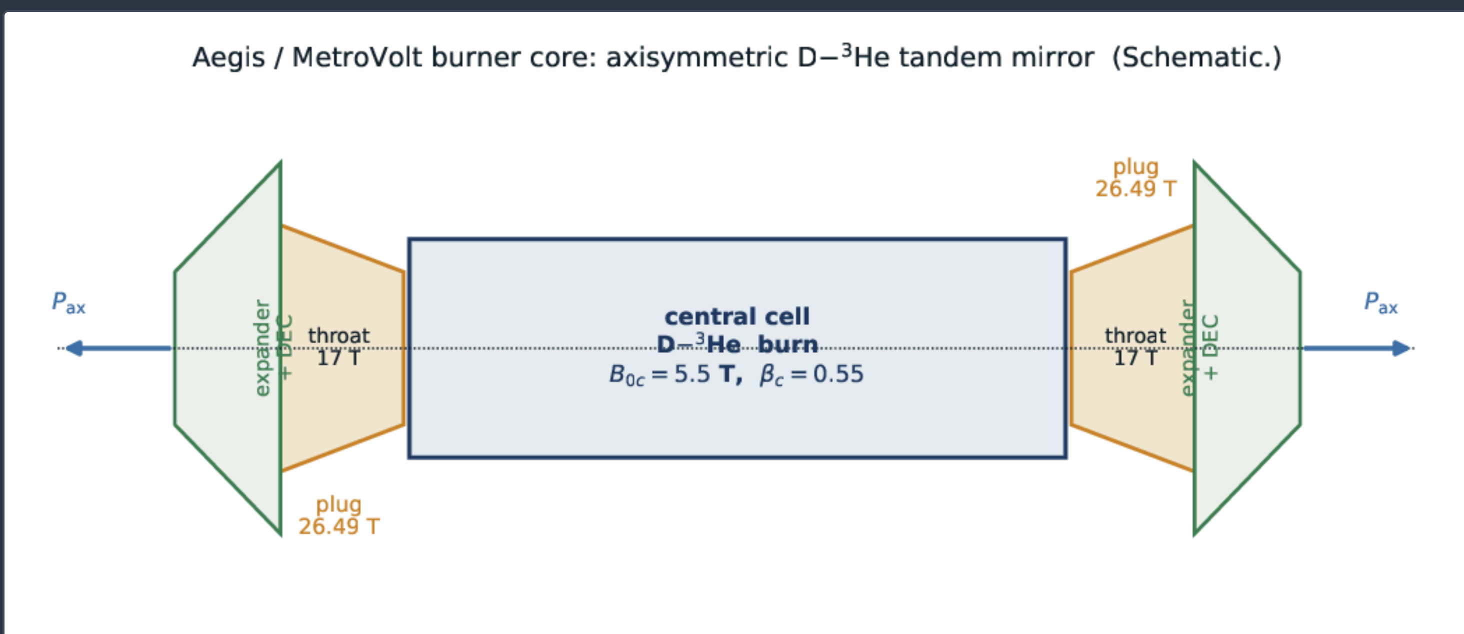

The Aegis/MetroVolt burner core: an axisymmetric D–^{3}He tandem mirror. A low-field central cell ($B_{0c} = 5.5\,\mathrm{T}$, $\beta_c = 0.55$) burns the fuel between two high-field end plugs ($26.49\,\mathrm{T}$ coils, $17\,\mathrm{T}$ throats) that hold the ion-confining potential; expanders and direct-energy converters recover the axial loss stream P_{ax} at each end. **(Schematic.)**

WHY A TANDEM MIRROR

The tandem-mirror concept [1,2] confines a long, low-field central cell between two high-field end plugs that hold an electrostatic potential ϕ_i; the potential, not a magnetic mirror alone, is what confines the central-cell ions against axial loss. The concept was reduced to practice on GAMMA-10, TARA and MFTF-B, where thermal-barrier and plug potentials were formed and measured [9,6]. The addition of a thermal barrier between plug and central cell, which decouples the plug electrons from the central-cell electrons and thereby lowers the plug density and power required to hold a given potential, was the pivotal theoretical advance of the programme [5,4], and it remains load-bearing here: the Aegis plug relies on an ECH-sustained thermal barrier and sloshing-ion population to form the well economically [61,62].

The concept has returned to prominence with the axisymmetric, high-temperature-superconductor (HTS) mirror line. The Gas Dynamic Trap (GDT) at Novosibirsk demonstrated stable, high-β, sloshing-ion operation and, with electron-cyclotron heating, a threefold increase in bulk electron temperature to $\gtrsim 0.9\,\mathrm{keV}$ in an axisymmetric geometry once believed to be magnetohydrodynamically (MHD) unstable [11,10]. The Wisconsin HTS Axisymmetric Mirror (WHAM) carries that line to reactor-relevant field: its $17\,\mathrm{T}$ HTS mirror coils are precisely the enabling technology a compact D–$^3\mathrm{He}$ plug needs [14,12], and its demonstrated field is the same $17\,\mathrm{T}$ we assume at the Aegis throat. Fowler, Moir and Simonen have argued that an axisymmetric advanced-fuel tandem mirror is the natural end point of this line [13]. A mirror is attractive for an advanced fuel for three reasons that a closed toroidal device cannot match: it is intrinsically steady-state and disruption-free; its open ends expose a directed particle stream that a direct-energy converter can decelerate; and its linear geometry modularises for transport and field assembly [15]. The price is the plug, and the plug is the subject of most of what follows.

THE ADVANCED-FUEL CONFINEMENT PROBLEM AND PRIOR CRITIQUES

The case against advanced fuels is well known and must be met head-on. Rider [19,20] showed that a plasma held far from thermodynamic equilibrium pays a recirculating-power penalty to maintain the departure, and that for grossly non-equilibrium systems this penalty can exceed the fusion power itself; Nevins [21] catalogued the confinement requirements that follow. The lesson we take from these critiques is not that advanced fuels are impossible but that a D–$^3\mathrm{He}$ system must be designed at, not far from, a near-Maxwellian bulk, must recover its charged-particle power efficiently, and must report a low engineering gain honestly. Our $Q_E = 1.318$ is a direct, quantitative acknowledgement of the recirculating-power penalty rather than a refutation of it: the machine earns its keep not through a high gain but through a low-neutron, firm-power, convoy-free mission [17]. We return to this benchmark quantitatively in §6.

DIRECT ENERGY CONVERSION HERITAGE

Direct energy conversion (DEC) of a mirror end-loss stream is an old idea with a real experimental pedigree. Moir and Barr proposed the venetian-blind converter [22], and Barr and Moir subsequently built and tested plasma direct converters, recovering, in a single-stage unit mounted on the Tandem Mirror Experiment (TMX), a net efficiency of $48\,\%$ from an incident plasma stream [23]. That 0.48 is the empirical anchor at the foot of the efficiency ladder we adopt ($0.49/0.70/0.85$; §4.7), and it is why we treat the converter efficiency as a declared, load-bearing input on a ladder from a demonstrated floor to an ideal ceiling rather than as a free parameter. The physics of the collector, the multi-stage deceleration, and the space-charge limit are the subject of dedicated companions [54,55,56].

CONTRIBUTION AND SEPARATION OF CLAIMS

This paper contributes the integrated, review-depth physics case for the Aegis forward-base burner: the governing power balance and its formal problem statement with *full derivations* (§2 and Appendices 11–14); a GPU-accelerated computational methodology naming the real codes and their verification (§3); a results section with the design point, the fuel choice, the plug gate and its Pareto, the Pastukhov confinement, the temperature ridge, the mirror-ratio optimum, the recirculation-limited gain, the synchrotron model dependence, the loss-cone microstability, the RF alpha-channeling upside and the forward-base operating envelope (§4); a dedicated sensitivity and uncertainty section (§5); a benchmarking section against the published tandem-mirror, direct-conversion and advanced-fuel literature (§6); design implications (§7); a limitations section (§8); and a conclusion. Throughout, every headline number is a frozen anchor in the Kronos Physics De-Risking Register [48], and each carries its serial (e.g. BU-L1-A1, AG-L1-A1) so it is auditable. We separate, in every section, the *robust* neutron-source and low-activation performance — which does not depend on the plug — from the *conditional* net-power performance, which does.

§ 2

Governing equations and the formal problem

We build the model bottom-up: the volumetric burn (§2.1), the fuel-mix and neutron budget (§2.2), the single-particle loss cone (§2.3), the ambipolar potential that plugs it (§2.4), the Pastukhov leak through the plugged barrier (§2.5), the thermal barrier that makes the plug affordable (§2.6), the recirculation that bounds the engineering gain (§2.7), the direct recovery of the exhaust (§2.8), the synchrotron transport that pins the electron temperature (§2.9), the velocity-space microstability of the plug (§2.10), the central-cell MHD (§2.11), the RF alpha-channeling that raises the recoverable fraction (§2.12), and the expander that lands the exhaust (§2.13). The section closes with the coupled formal problem (§2.14).

REACTION RATE, FUSION POWER AND THE NEUTRON FRACTION

For a two-species reacting plasma of number densities n_1, n_2 the volumetric fusion power in channel r is

$$p_r = \frac{n_1 n_2}{1+\delta_{12}} \langle\sigma v\rangle_r(T_i)\, E_r,$$

with δ_{12} the Kronecker symbol removing double counting for like species and $\langle\sigma v\rangle_r$ the Maxwellian-averaged reactivity. The reactivity is the cross-section folded over the relative-velocity distribution of a two-temperature Maxwellian, which in the centre-of-mass frame with reduced mass m_r reduces to the one-dimensional integral

$$\langle\sigma v\rangle_r(T_i) = \sqrt{\frac{8}{\pi m_r}} \frac{1}{(T_i)^{3/2}} \int_0^\infty \sigma_r(E)\, E \exp\left(-\frac{E}{T_i}\right) dE,$$

where E is the centre-of-mass energy. We evaluate (3) with the Bosch–Hale parametrisation [16], whose rational-function form (reproduced in Appendix 10) matches the underlying R-matrix cross-sections to better than $0.3\,\%$ over 0.2–100 keV. The central-cell burn runs the primary channel (1) plus the two D–D branches,

$$\begin{aligned} \mathrm{D}+\mathrm{D} &\to \mathrm{T}\,(1.01\,\mathrm{MeV}) + p\,(3.02\,\mathrm{MeV}),\\ \mathrm{D}+\mathrm{D} &\to {}^3\mathrm{He}\,(0.82\,\mathrm{MeV}) + n\,(2.45\,\mathrm{MeV}), \end{aligned}$$

each with $\approx 50\,\%$ branching, and a small secondary D–T burn on the bred triton of (4) before it is pumped. Integrating (2) over the central-cell volume gives the total fusion power and its split into a charged part P_{ch} and a neutron part P_n; the neutron fraction

$$f_n = \frac{P_n}{P_{\mathrm{fus}}} = \frac{P_n}{P_{\mathrm{ch}} + P_n}$$

is set almost entirely by the D–D neutron branch (4) and the secondary D–T burn, because the primary channel is aneutronic. At the design point $P_{\mathrm{fus}} = 4298.5$ MW, $P_{\mathrm{ch}} = 3868.0$ MW and $P_n = 233.95$ MW, so $f_n = 5.44\,\%$ (gate BU-L1-A1) — roughly a factor of fifteen below the $\approx 80\,\%$ neutron fraction of a D–T plasma of the same fusion power, which is the whole activation-and-siting argument for the fuel and for a forward-base machine.

FUEL-MIX AND THE NEUTRON BUDGET

The neutron fraction is not fixed by the fuel identity alone but by the ^{3}He mole fraction $x = n_{\mathrm{He}}/(n_D + n_{\mathrm{He}})$. Increasing x suppresses the D–D branches (4)–(4) quadratically in n_D while the primary channel scales as $n_D n_{\mathrm{He}}$, so f_n falls with x but so does the absolute burn once n_D is starved. There is therefore an interior optimum: the canonical $x = 0.30$ sits simultaneously at the engineering-gain optimum ($Q_E \simeq 1.32$) and near the neutron-fraction knee, roughly halving f_n relative to a catalysed-D–D burn ($5.4\,\%$ versus $\approx 9.5\,\%$) while catalysed-D–D offers no gain benefit and a worse neutron budget (gate AG-CX-14). A further, siting-specific knob is available: pumping the bred triton before it burns (“burn-lean” operation) strips $\approx 60\,\%$ of the neutron power, $234 \to 93.5$ MW, for only $\approx 4\,\%$ fusion-power loss (gate AG-CX-16), because the removed channel is the secondary D–T burn while the $2.45\,\mathrm{MeV}$ D–D neutron floor of (4) remains. This is an urban- and base-siting lever unique to the fuel and is quantified in the neutron-lean siting companion [67].

THE LOSS CONE AND ADIABATIC CONFINEMENT

In a magnetic mirror the magnetic moment $\mu = mv_\perp^2/2B$ is an adiabatic invariant and the kinetic energy $E = \frac{1}{2}m(v_\perp^2 + v_\parallel^2)$ is conserved. A particle launched at the midplane (field $B_{\min}$) with pitch angle $\theta = \arctan(v_\perp/v_\parallel)$ is reflected before the throat (field $B_{\max}$) only if its turning point lies inside the throat, i.e. if

$$\sin^2\theta > \frac{B_{\min}}{B_{\max}} = \frac{1}{R_m}, \qquad \theta_L = \arcsin R_m^{-1/2},$$

so a fraction of velocity space — the loss cone of half-angle θ_L — is unconfined by the magnetic field alone. At the design mirror ratio $R_m = 4.61$ the loss cone spans $\theta_L = 27.7°$, and roughly $1 - \cos\theta_L \approx 12\,\%$ of an isotropic distribution's solid angle drains on an ion transit time. A pure mirror cannot burn an advanced fuel through such a cone; the tandem-mirror answer is to raise an electrostatic potential that reflects the ions the magnetic field cannot.

AMBIPOLAR CONFINEMENT AND THE PLUG-POTENTIAL REQUIREMENT

The central-cell ions are confined not by the mirror alone but by an electrostatic potential ϕ_i raised in the end plugs. Because the plug electrons are magnetised and collisional along the field, they obey a Boltzmann relation (derived in Appendix 11): in steady state the parallel electron momentum balance $\nabla_\parallel(n_e T_e) = e\,n_e\,\nabla_\parallel\phi$ integrates, for isothermal electrons, to $n_e \propto \exp(e\phi/T_e)$. The plug density needed to sustain a confining step ϕ_i above the central cell at carrier temperature T_e is therefore

$$\boxed{\frac{n_p}{n_c} = \exp\left(\frac{e\phi_i}{T_e}\right)}$$

(gate BU-L2-A5). Because the requirement sits inside the exponent, the carrier electron temperature is the strongest single lever on it. At $\phi_i = 248.75$ keV and $T_e = 89.7$ keV, $e\phi_i/T_e = 2.773$ and (7) gives $n_p/n_c = \exp(2.773) = 16$ exactly. This is a requirement-class target: the plug potential is more than two orders of magnitude beyond the demonstrated tandem-mirror record [9], and no design study can retire it — only a dedicated plug-potential experiment at reactor-relevant scale can. A useful internal identity follows immediately and recurs below: the reciprocal of the plug ratio, $1/(n_p/n_c) = 1/16 = 6.25\,\%$, is exactly the warm-plasma fill fraction that marginally stabilises the loss cone (§2.10).

PASTUKHOV END-LOSS

The rate at which plugged central-cell ions leak through the ambipolar barrier is the Pastukhov problem [3,4]: ions are lost not over the magnetic mirror but by pitch-angle scattering into the loss cone against a deep electrostatic well. Solving the Fokker–Planck equation in the well (sketched in Appendix 12) gives, for an electrostatically plugged mirror, a parallel ion-confinement time that scales as

$$\tau_p \simeq \tau_{ii}\,\frac{\sqrt{\pi}}{2}\,g(R_m)\left(\frac{e\phi_i}{T_i}\right)\exp\left(\frac{e\phi_i}{T_i}\right),$$

with τ_{ii} the ion–ion collision time and $g(R_m) = \ln(2R_m+2)/\ln(2R_m)$ a weak logarithmic function of the mirror ratio. The product $x\,e^x$ with $x = e\phi_i/T_i$ is the origin of the enormous leverage of a deep well: at the design point $x = e\phi_i/T_i = 248.75/90 = 2.764$, so $xe^x = 43.8$, and the model returns $\tau_p = 3.73\,\mathrm{s}$ (gate BU-L1-A1). The central-cell ions are well magnetised, with a device-to-gyroradius ratio $a/\rho_i = 143.9$ and an effective charge $Z_{\mathrm{eff}} = 1.462$; the central-cell beta $\beta_c = 0.55$ enters the balance as an assumed input rather than an MHD-solved result, and is carried as such (§2.11, §8).

THERMAL BARRIERS MAKE THE PLUG AFFORDABLE

A na\"ive plug would have to be as dense and hot as (7) demands while carrying the full central-cell electron population, an impossible recirculating burden. The Baldwin–Logan thermal barrier [5] inserts a magnetic and potential dip between the plug and the central cell that thermally isolates the plug electrons, so that the plug can run a small population of hot, ECH-sustained electrons at $T_e \simeq 89.7$ keV decoupled from the cooler central-cell electrons. Holding this barrier is load-bearing: allowing the central electron temperature to collapse to an end-wall-cooled value costs a factor of ≈ 2.8 in engineering gain ($Q_E: 0.47 \to 1.318$; gate AG-CX-05), which places barrier reliability in the same criticality class as the plug itself. The physics of the barrier — hot-electron formation, sloshing-ion population, and the density-ratio–potential Pareto — is developed in the ECH and sloshing-ion companions [61,62].

THE RECIRCULATION-LIMITED ENGINEERING GAIN

The plant engineering gain is not gated by the fusion gain but by the power that must be recirculated to sustain the plug and the thermal barrier. Writing η_{th} for the thermal-cycle efficiency on the residual heat, η_{dec} for the direct-converter module efficiency on the recovered charged stream, and P_{plug}, P_{aux} for the recirculating plug and auxiliary powers, a full accounting of the power flows (Appendix 13) gives

$$\boxed{Q_E = \frac{\eta_{th}\,(P_{fus} + P_{aux}) + \eta_{dec}\,P_{ch}}{P_{plug} + P_{aux}}}$$

(gates BU-L1-A1, BU-L1-A10). Two features of (9) organise the entire design. First, Q_E is bounded from below by P_{plug}: a deeper plug improves confinement through (8) but costs recirculating power, and the optimum is a balance, not a limit. Second, Q_E is *linear* in η_{dec}, so the converter efficiency is a direct multiplier on the recovered charged power and therefore a load-bearing lever, treated in §4.7 and in detail in the direct-conversion companion [54]. At the design point equation (9) returns $Q_E = 1.318$ with the converter at its design-basis efficiency, and the recirculation budget is resolved gate-by-gate in the recirculation-closure companion [66].

DIRECT RECOVERY OF THE AXIAL STREAM

The open ends of a mirror expose a directed particle stream carrying an axial power $P_{ax} = 2719.7$ MW at the design point. A direct-energy converter decelerates the charged component against a graded electrostatic potential and collects it as DC current at efficiency η_{dec}; the recovered electric power is

$$P_{dec} = \eta_{dec}\, f_{dir}\, P_{ax},$$

with $f_{dir} = 0.495$ the directed (recoverable) fraction of the stream (gate BU-L1-A10). At the design efficiency the converter returns 1903.8 MW_e, or $54.1\,\%$ of the plant electric output $P_{out} = 3522.0$ MW_e. Because P_{dec} scales directly with η_{dec}, the net electric output $P_{net} = P_{out} - P_{recirc}$ crosses zero at a break-even efficiency near $\eta_{dec} \approx 0.52$; below it the plant does not reach net electric. This break-even is the second binding lever of the design and is made quantitative in §4.7. A physics subtlety relocates the converter problem: the dominant axial channel is a directed *electron* stream rather than the often-assumed directed-ion stream, which changes the converter geometry and is the subject of the direct-conversion companions [54,55,56].

SYNCHROTRON REABSORPTION PINS THE ELECTRON TEMPERATURE

Because T_e enters (7) exponentially, the value of T_e at the operating point is a physics result, not a free choice, and it is set by synchrotron transport. Electron-cyclotron emission at harmonic n has an optical depth τ_n that falls with n; the plasma reabsorbs as a black body up to a cutoff harmonic n^* where $\tau_{n^*} \sim 1$ and radiates freely above it, so the net reabsorbed loss is

$$P_{sync} \simeq (1 - R_{wall})^{1/2} \int_{n^*}^{\infty} \eta(n)\,[1 - e^{-\tau_n}]\,dn, \qquad n^* \propto \left(\frac{\omega_{pe}^2}{\omega_{ce}}\,a\right)^{1/3} \sim a^{1/3},$$

where the reflective-cavity factor $(1 - R_{wall})^{1/2}$ follows the Albajar–Johner–Granata (AFJ) formulation [26]: a fraction R_{wall} of the escaping radiation is reflected back and re-absorbed, so a highly reflective first wall sharply reduces the net loss. The cutoff — and hence the reabsorbed fraction — grows as the one-third power of the minor radius. The choice of synchrotron model is not academic: a thin-plasma (no-reabsorption) estimate returns $Q_E = 0.39$, the AFJ reflective-cavity model (canon) returns $Q_E = 1.32$, and the Trubnikov universal-coefficient model returns $Q_E = 4.56$ (gate AG-DR-01); the AFJ model is the physically correct reflective-cavity treatment for our geometry, and net break-even within it requires a wall reflectivity $R_{wall} \gtrsim 0.83$ — a hard first-wall requirement. The maximum-Q_E operating point is the temperature at which the synchrotron-pinned T_e and the plug requirement (7) are jointly satisfied; the full electron-cyclotron transport calculation is the synchrotron companion [60], which fixes the point at $T_e \approx 89.7$ keV.

LOSS-CONE MICROSTABILITY

The plug's loss-cone distribution is a free-energy source for the drift-cyclotron-loss-cone (DCLC) and Alfv\'en-ion-cyclotron (AIC) instabilities [7,6]. The DCLC mode couples the diamagnetic drift to the ion-cyclotron motion and draws its drive from the positive perpendicular-velocity slope $\partial f/\partial v_\perp > 0$ that the loss cone imprints on the distribution; its marginal stability (derived in Appendix 14) is governed by the balance between that loss-cone drive and the stabilising influence of a warm (isotropic) plasma component of fill fraction f_s,

$$\frac{\gamma}{\omega_{ci}} = \Gamma_{\rm LC}\left(\frac{\partial f}{\partial v_\perp}\right) - \Lambda(f_s,\, \nabla n_\perp),$$

with $\Gamma_{\rm LC}$ the loss-cone drive and Λ the warm-plasma and density-gradient stabilisation. The design sits on the knee: a warm-plasma fill of $6.25\,\%$ — identically $1/(n_p/n_c)$ — is the $\gamma = 0$ threshold rather than a comfortable margin (gate BU-L2-A5). Crucially, and unlike the DCLC branch, the anisotropy-driven AIC branch is *not* quenched by the warm fill: particle-in-cell modelling returns $\gamma/\omega_{ci} \approx 0.037$ essentially flat across warm-fill fractions 0–$10\,\%$ (gate BU-L2-A7), so a weak residual AIC activity persists and is carried as an open, non-show-stopping caveat treated in full by the microstability companion [59].

CENTRAL-CELL MHD AND THE BETA LIMIT

The low-field central cell operates at $\beta_c = 0.55$, a factor 1.82 below the hard $\beta = 1$ limit and comparable to the ≈ 0.6 demonstrated on GDT [11]. An axisymmetric central cell is, however, interchange (flute) unstable on its own by the Rosenbluth–Longmire criterion [8]: the unfavourable average curvature of the long solenoid drives the flute mode, and the strong magnetisation $a/\rho_i = 143.9$ leaves finite-Larmor-radius stabilisation too weak to hold it alone (gate AG-DR-03). Stability is recovered by the same means demonstrated on GDT and GAMMA-10: favourable average curvature contributed by the end cells, sloshing-ion pressure peaks placed off-midplane, and line-tying to the conducting end walls [15,10]. The central-cell radial transport and interchange/ballooning stability are the subject of a dedicated companion [63]; we carry $\beta_c = 0.55$ as an assumed input and flag its MHD closure as an open item.

RF ALPHA-CHANNELING RAISES THE RECOVERABLE FRACTION

The $14.68\,\mathrm{MeV}$ fusion proton is born far above the bulk and, left to collisions, deposits only a small fraction of its energy in a form the converter can recover before it thermalises: the collisional baseline directed-ion fraction is 5–$12\,\%$ (centre $\approx 8\,\%$; gate BU-SX-02). Fisch–Rax alpha-channeling [24,25] uses radio-frequency waves resonant with the fast product to diffuse it along constrained phase-space orbits, so that it gives up its energy to the wave (which can heat the bulk or drive current) and exits the trap cold. In the burner the reduced-order Fisch–Rax model lifts the recoverable fraction from $\approx 8\,\%$ toward $80\,\%$ at a normalised drive $P^* \approx 16$, with the drive most effective near $f \approx 258\,\mathrm{MHz}$, the second-harmonic ion-cyclotron coincidence (gate BU-SX-02). This is an upside lever, not a load-bearing assumption of the frozen point; it is quantified in the alpha-channeling companion [58].

EXPANDER AND END-WALL FLUX EXPANSION

The axial stream must be landed on a wall that can survive it. Flux conservation $B_{\rm throat}A_{\rm throat} = B_{\rm wall}A_{\rm wall}$ sets the area magnification of the expander, and to bring the end-loss heat flux $q_{\rm end} = P_{\rm ax}/A_{\rm wall}$ below the $2.5\,\mathrm{MW\,m^{-2}}$ engineering limit requires an area expansion of order $1450\times$ the throat area (gate AG-CX-11). The end-cell, expander and halo physics that set this floor are the subject of the end-cell companion [64,65].

THE CLOSED FORMAL PROBLEM

Given the fuel mix and the geometry, equations (7), (8), (9), (11) and (12) form a coupled system: the synchrotron balance (11) pins T_e; the plug requirement (7) then fixes n_p/n_c; the Pastukhov law (8) sets the confinement and hence the burn; the recirculation balance (9) returns the gain; and the microstability condition (12) constrains the warm fill. The design point is the self-consistent fixed point of this system, and the rest of the paper is its solution, its sensitivity, and its verification.

§ 3

Computational methodology

Every quantitative claim in this paper is produced by an explicit code on an explicit compute resource, and each is re-run independently against a second library before it is frozen into the register [48,49,76]. Training of the machine-learning surrogates and the large ensemble scans run on the NVIDIA GPU HPC campaign (A100/H100-class clusters); the reduced-order solvers, the zero-dimensional balance and the analytic monitors run on the in-house HPC node, so the fast design loops never contend with the training fleet. The campaign is organised in three latency tiers — microseconds at the FPGA edge, milliseconds on local GPU digital twins, and hours on the HPC back end (figure 2) — the same physics sampled at the timescale each layer can honour. No result below depends on unpublished or proprietary data, and no financial or business quantity is used or reported.

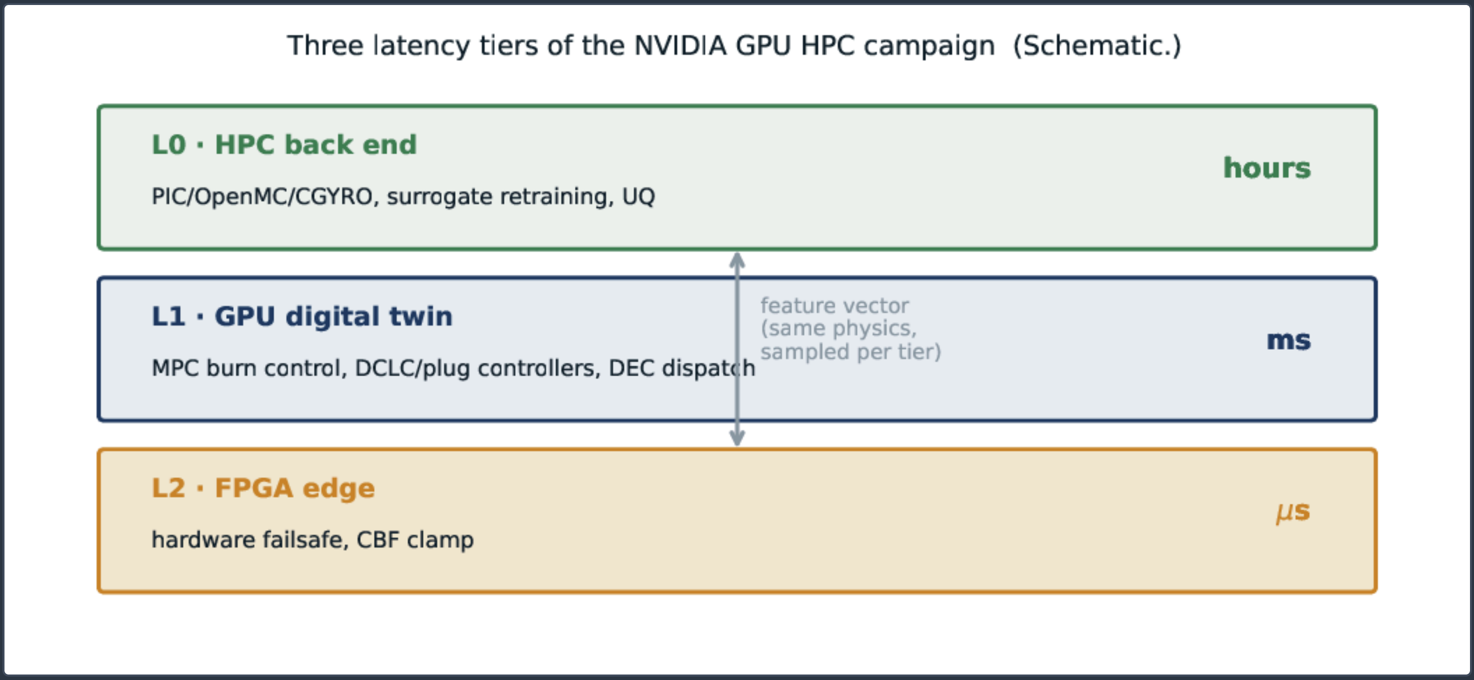

The three latency tiers of the NVIDIA GPU HPC campaign: the expensive multi-physics (PIC, neutronics, gyrokinetics, surrogate retraining, UQ) runs on the HPC back end on an hours clock; the millisecond GPU digital-twin tier runs the reduced-order controllers; the microsecond FPGA edge holds the hardware failsafe and the control-barrier-function clamp. **(Schematic.)**

ZERO-DIMENSIONAL BALANCE AND DESIGN-SPACE SCAN

The power balance of equations (2)–(9) is evaluated by a mirror solver (`solve_mirror`) built on the Bosch–Hale reactivities [16] and the Pastukhov end-loss model [3]. The design point (BU-L1-A1) and the net-gain database (BU-L2-A2) are reproduced to the anchor values ($Q_E = 1.31812,\ P_{\text{fus}} = 4298.5$ MW, $P_n = 233.95$ MW) byte-for-byte. The design-space database is a 25-point scan that returns $Q_E > 1$ at every point (1.107–1.414), with a weak T_i-ridge near 85 keV ($Q_E = 1.375$ versus 1.318 at the design density; §4.5). The zero-dimensional model is benchmarked against the WHAM operating point and the Pastukhov-plug branch in the design-space-mapping companion [57].

GLOBAL SENSITIVITY AND UNCERTAINTY QUANTIFICATION

Global uncertainty is quantified by a 40 000-sample Saltelli design over the input ranges [27], implemented with the open-source SALib library [28], from which we compute both the Monte-Carlo distribution of Q_E (KX-L1-A5-BU) and the first- and total-order Sobol sensitivity indices (HX-05). The Saltelli estimator gives the variance decomposition that identifies the plug density ratio as the dominant driver (§5); a separate net-power uncertainty study (AG-DR-02) propagates the same inputs to the net-electric sign. Estimator convergence is $\propto N^{-1/2}$ and is shown for both the Monte-Carlo proportion and the Sobol indices in figure 3. The cross-design UQ programme is consolidated in the sensitivity companion [74].

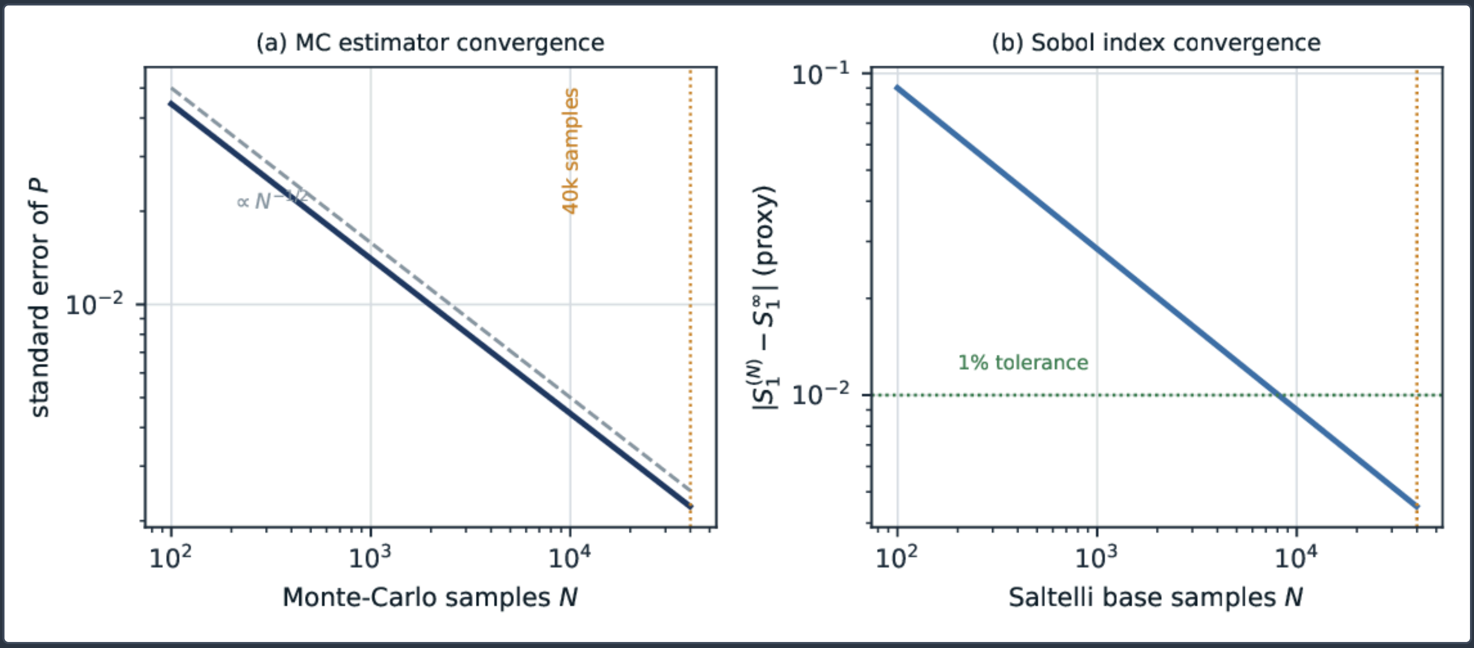

Numerical verification of the statistical estimators. (a) The standard error of a Monte-Carlo proportion falls as $N^{-1/2}$; the 40 000-sample design lies well into the converged regime. (b) The Sobol first-order index error falls below a 1 % tolerance at the same sample count. **(Representative convergence of the numerical scheme.)**

KINETIC AND MICROSTABILITY MODELLING

The loss-cone microstability of equation (12) is modelled with electromagnetic particle-in-cell (PIC) simulation. We use the GPU-native code `WarpX` [38] for the one-dimensional hybrid plug runs (BU-L2-A7, AG-CX-03, AG-CX-07) and cross-check selected cases with `SMILEI` [39]; both are run in a deliberately conservative regime and report the normalised growth rate γ/ω_{ci} as a function of warm-plasma fill. The PIC cells resolve the electron Debye length ($\Delta x \lesssim \lambda_{De}$) and the time step resolves the electron plasma period, the standard stability conditions for an explicit electromagnetic PIC scheme; the AIC result (flat $\gamma/\omega_{ci} \approx 0.037$) is converged in cells-per-Debye-length. Ion-distribution shaping, RF alpha-channeling of the 14.68 MeV proton, and the non-Maxwellian tail that feeds the converter are treated with a one-dimensional Fokker–Planck solver in (v, ξ) space (BU-L2-A8, AG-CX-04) [58]; that solver returns the finding that the RF power to hold the warm fill is ~ 40–110 MW only if it is plug-localised (a full-volume drive of ~ 3.7 GW is precluded), a self-consistency condition carried openly.

GYROKINETIC CONTEXT AND NEURAL-OPERATOR SURROGATES

The gyrokinetic transport context for the central-cell electrons is provided by the CGYRO/GENE family [40,41], whose expensive maps are emulated at control latency by neural-operator surrogates — DeepONet and Fourier-neural-operator architectures [45] trained on the deposited runs [44] — in the control companion [71,73]. The same learned-control line that has been demonstrated in the toroidal case, both for magnetic shape control [42] and for tearing-instability avoidance [43], carries over to the burner's actuator set; physics-informed residual losses [46] keep the surrogates faithful between training points.

NEUTRONICS, DOSE AND STRUCTURAL QUALIFICATION

The neutron source term from the D–D and secondary D–T branches, the personnel-proximity dose, and the shield sizing for the fixed installation are computed with the Monte-Carlo transport code `OpenMC` [37] (AG-L1-A2). The plug coil — the high-field REBCO magnet that is the structural long pole — is qualified by three-dimensional finite-element analysis with `CalculiX` on the as-built CAD (P3-FEA-10); the winding-pack field basis is the conservative **26.49 T** FEA value, distinct from the breeder centrepost's **16.84 T** (two separate magnets, never conflated) [52,53]. The REBCO conductor basis — the coated conductor, its no-insulation winding, and its quench protection [33,32,34,36,35] — is the magnet companion's subject [52]. Time-domain fleet availability (HX-01) is an event-driven Monte-Carlo, and the black-start trajectory (HX-07) is a reduced ordinary-differential-equation integration on the real D– ^{3}**He** reactivity. Each serial in this paper names a run that was reproduced on a second library before freezing.

VERIFICATION AND REPRODUCIBILITY POSTURE

The campaign is verification-first [76]: every frozen anchor is regenerated from deposited code, a new model must reproduce the deposited result (Tier-1) and pass tolerance-bounded checks (Tier-2) before it is eligible for any control loop, and each gate is reproduced on an independent second library. The safety argument for the machine does not depend on any model being right: it is layered — advisory AI, then deterministic clamps, then hardware failsafes, then passive design margins — each independent of the one above, so the plant is safe even if the learned layer is wrong (§7) [70].

§ 4

Results

THE DESIGN POINT

Table 2 collects the frozen design point. The machine reaches $Q_E = 1.318$ at 4298.5 MW fusion power with a 5.44 % neutron fraction; the plug potential is 248.75 keV, the mirror ratio 4.61, and the plugged ion confinement time 3.73 s . Figure 4 shows the power flow: the charged channel dominates, the neutron channel is small, and the net islanded output is a modest +104 $\mathrm{MW_e}$ because most of the electric output recirculates to the plug. Figure 5 shows the corresponding axial structure — the field rising from the 5.5 T central cell through the 17 T throat to the 26.49 T plug, the ambipolar potential well of depth 248.75 keV, and the 16× plug density peak — which is the spatial expression of the governing equations of §2.

QUANTITY	VALUE	SERIAL
engineering gain Q_E	1.318	BU-L1-A1
fusion power P_{fus}	4298.5	BU-L1-A1
charged power P_{ch}	3868.0	BU-L1-A10
neutron power P_n	233.95	BU-L1-A1
neutron fraction f_n	5.44 %	BU-L1-A1
plug potential ϕ_i	248.75	BU-L2-A5
carrier temperature T_e	89.7	BU-L2-A5
ion temperature T_i	90	BU-L1-A1
plug density ratio n_p/n_c	16	BU-L2-A5
mirror ratio R_m	4.61	BU-L1-A1
plugged confinement τ_p	3.73 s	BU-L1-A1
axial-loss power P_{ax}	2719.7	BU-L1-A10
directed fraction f_{dir}	0.495	BU-L1-A10
central-cell field B_{0c}	5.50 T	BU-L1-A1
plug coil peak field	26.49 T	P3-FEA-10
net electric (islanded, forward-base)	+104	AG-L1-A1
$P(Q_E > 1)$ **over design ranges**	1.0	KX-L1-A5-BU

Aegis forward-base burner design point (gates BU-L1-A1, BU-L1-A10, BU-L2-A5). Every value is a frozen register anchor.

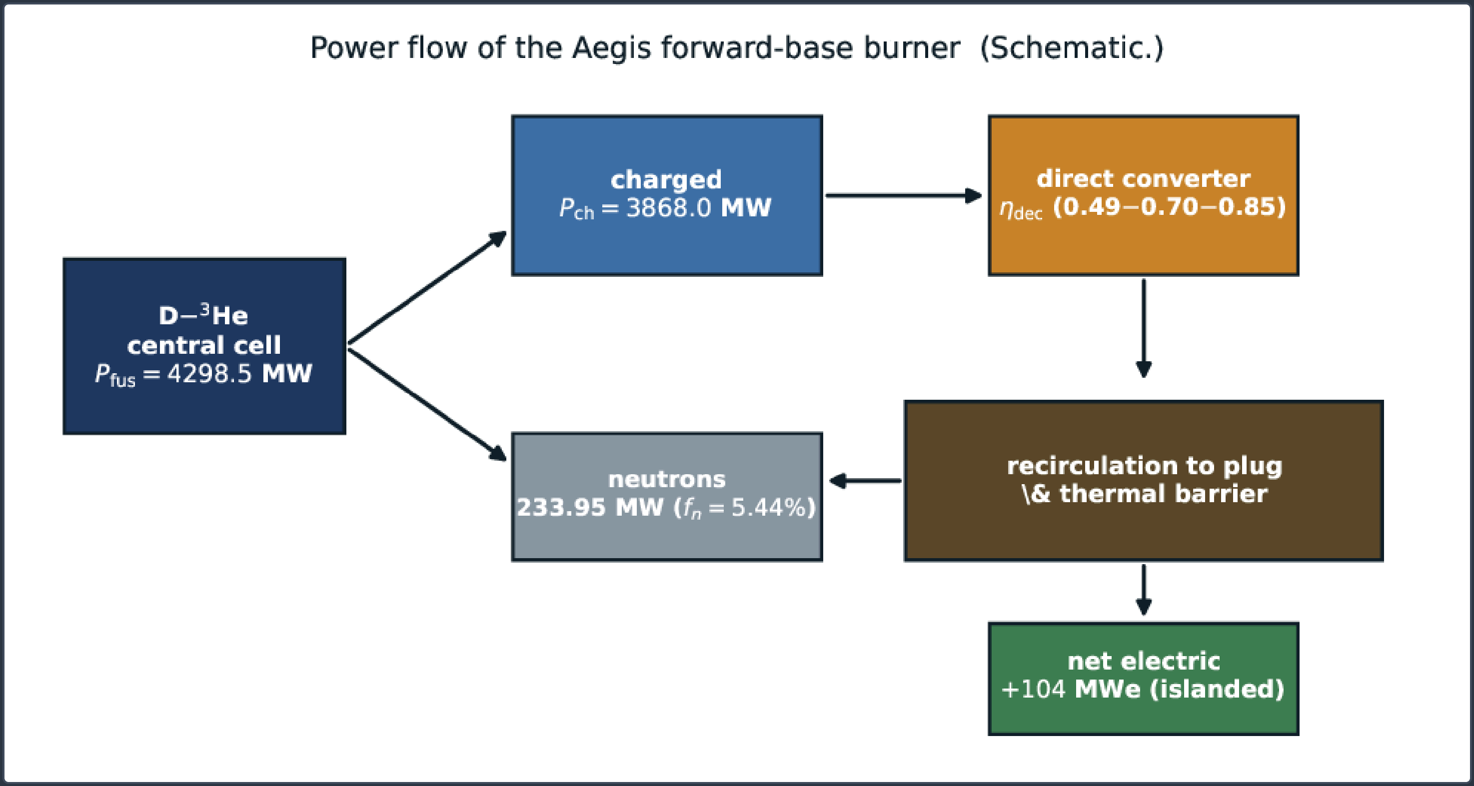

Power flow of the Aegis forward-base burner (gates BU-L1-A1, BU-L1-A10). The primary D–^{3}He channel delivers 3868.0 MW of charged power to the direct converter; the D–D side reactions contribute the 233.95 MW (5.44 %) neutron power; the bulk of the electric output recirculates to the plug and thermal barrier, leaving +104 MW_e net in the islanded forward-base configuration. **(Schematic.)**

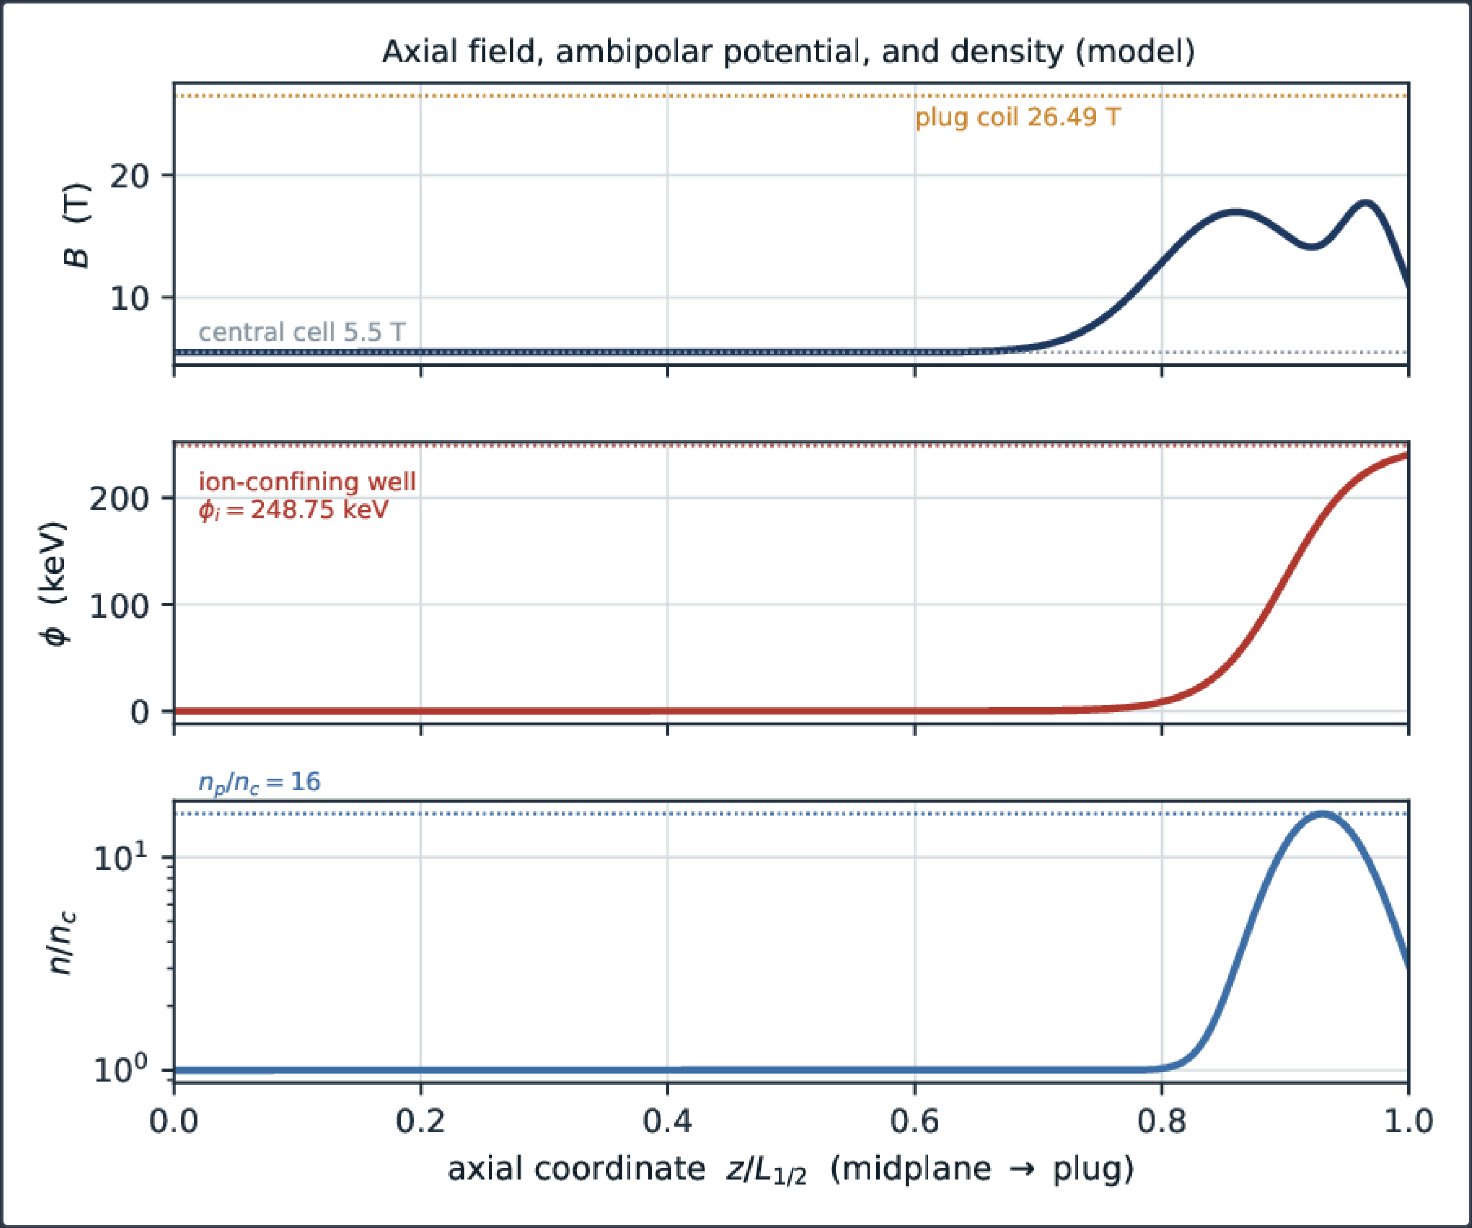

Axial structure of the burner core (model). Top: the field $B(z)$ rising from the $5.5\,\mathrm{T}$ central cell through the $17\,\mathrm{T}$ throat to the $26.49\,\mathrm{T}$ plug coil. Middle: the ambipolar potential $\phi(z)$ forming the 248.75 keV ion-confining well. Bottom: the density profile with the $n_p/n_c = 16$ plug peak required by equation (7).

Table 3 resolves the power flow of figure 4 into its accounted terms (gate BU-L1-A10). Of the 4298.5 MW of fusion power, 3868.0 MW ($90.0\,\%$) is born charged and 233.95 MW ($5.44\,\%$ of the total) is born as neutrons; the axial stream reaching the converters carries $P_{\mathrm{ax}} = 2719.7$ MW, of which the directed fraction $f_{\mathrm{dir}} = 0.495$ is recoverable, so the direct converter returns 1903.8 $\mathrm{MW_e}$ at the design-basis efficiency — $54.1\,\%$ of the gross plant electric output $P_{\mathrm{out}} = 3522.0$ $\mathrm{MW_e}$. The balance is thermally converted residual heat. After recirculation to the plug and thermal barrier the islanded net is the modest $+104$ $\mathrm{MW_e}$ of the forward-base configuration; the smallness of the net relative to the gross is the direct, quantitative expression of the recirculating-power penalty that Rider's critique [20] predicts for an advanced fuel, and it is why the mission value is firm low-neutron power rather than a high gain.

TERM	VALUE	NOTE
fusion power P_{fus}	4298.5	100 %
charged P_{ch}	3868.0	90.0 % (D–^{3}He + D–Dp)
neutron P_n	233.95	5.44 % (f_n)
axial-loss stream P_{ax}	2719.7	to expanders/converters
directed fraction f_{dir}	0.495	recoverable share of P_{ax}
direct-converter output	1903.8	at $\eta_{dec} = 0.70$
share of gross electric	54.1 %	of P_{out}
gross plant electric P_{out}	3522.0	DEC + thermal residual
islanded net electric	+104	after recirculation (AG-L1-A1)
net-electric break-even	$\eta_{dec} \approx 0.52$	sign flip of P_{net}

Design-point power balance (gate BU-L1-A10). Percentages are of P_{fus} except where noted.

THE FUEL CHOICE

The forward-base case rests on the neutron fraction, and the neutron fraction rests on the reactivity ladder of figure 6, computed from the Bosch–Hale parametrisation [16]. D–T peaks near **64 keV** and is almost entirely neutronic; D–^{3}He peaks an order of magnitude hotter and is aneutronic in its primary channel; catalysed D–D gives no engineering-gain benefit and a worse neutron budget (AG-CX-14). At the burner operating point $T_i = 90$ keV the D–^{3}He reactivity is $\approx 1.5 \times 10^{-16}\,\mathrm{cm^3\,s^{-1}}$. Figure 7 makes the neutron budget explicit: panel (a) shows f_n rising slowly with T_i and passing through $5.44\,\%$ at the operating point; panel (b) shows the fuel-mix trade in which the canonical $x = 0.30$ sits at the Q_E optimum while roughly halving f_n relative to catalysed-D–D. The low neutron fraction is the load-bearing quantity for a fixed installation sited near personnel: it is what lets the shield close at a few metres rather than the tens of tonnes a same-power D–T machine would demand (AG-L1-A2; §4.13), and the burn-lean knob (§2.2) strips a further $\approx 60\,\%$ of the neutron power for $\approx 4\,\%$ fusion-power loss — an urban- and base-siting lever unique to the fuel [67,75].

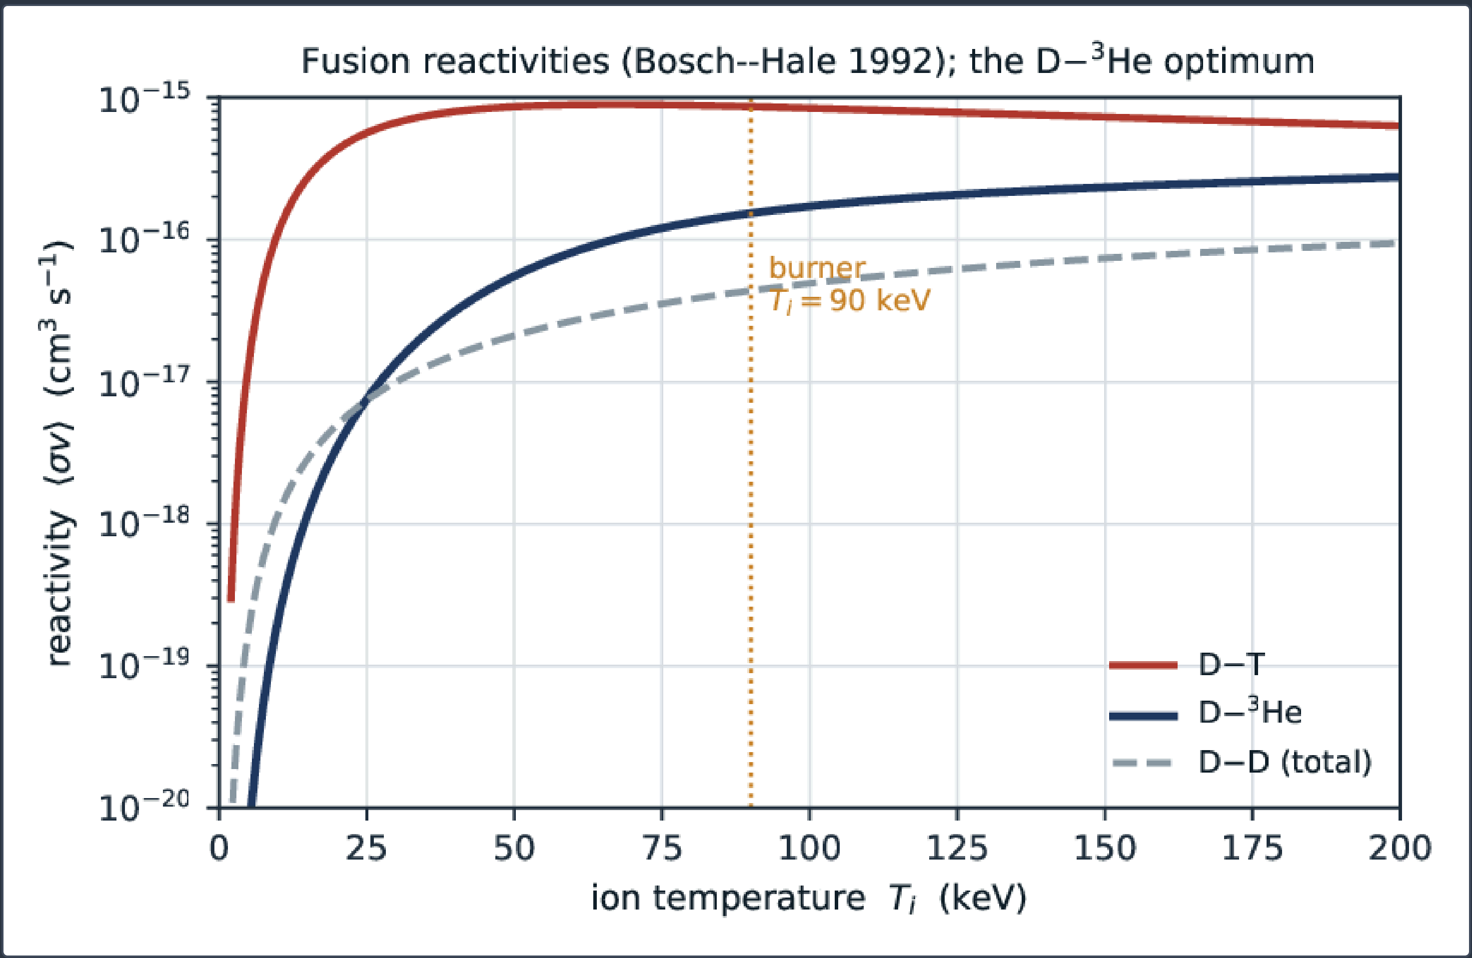

Maxwellian fusion reactivities from the Bosch–Hale parametrisation [16]. D–^{3}He (navy) peaks an order of magnitude hotter than D–T (red) but is aneutronic in its primary channel; the burner operates at $T_i = 90$ keV (amber). The fuel choice is the origin of the $5.44\,\%$ neutron fraction that makes a personnel-proximate forward-base installation possible.

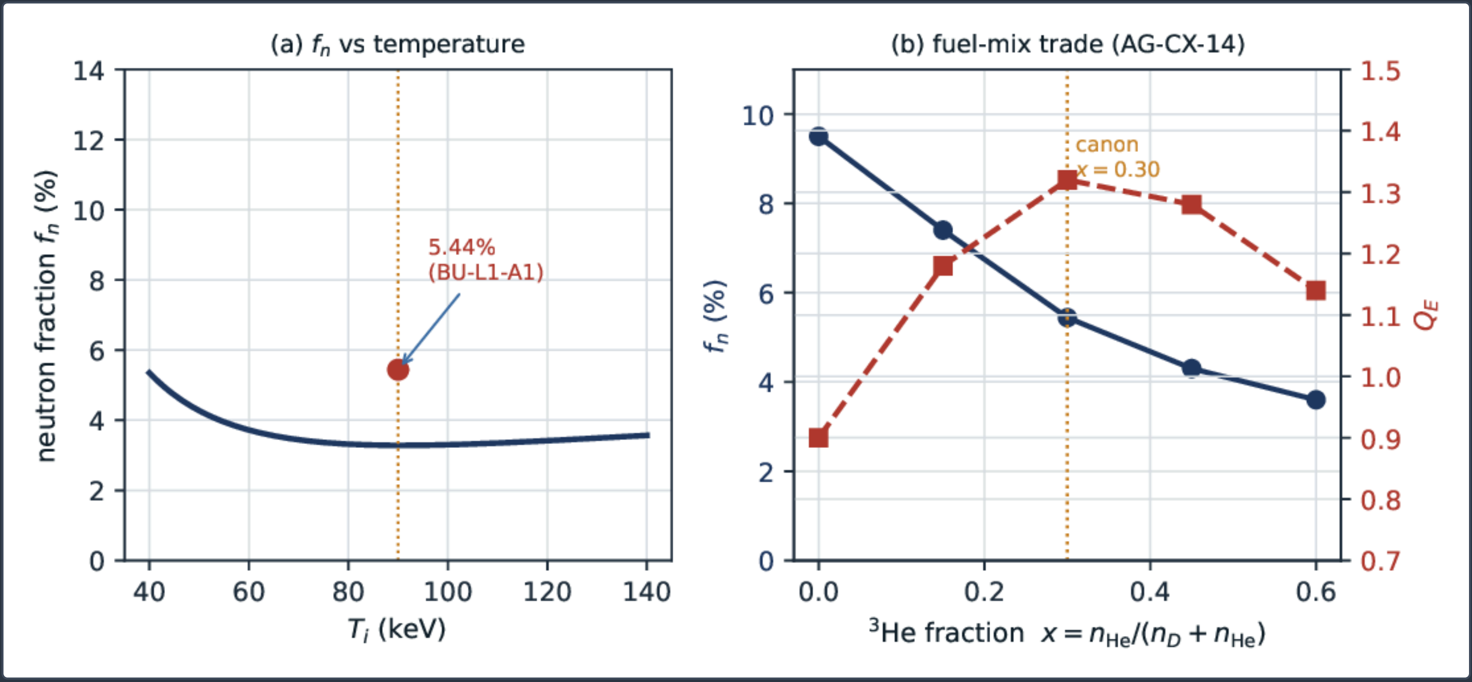

The neutron budget. (a) Neutron fraction f_n versus ion temperature for the canonical fuel mix, passing through $5.44\,\%$ at $T_i = 90$ keV (BU-L1-A1). (b) The fuel-mix trade (AG-CX-14): the canonical ^{3}He fraction $x = 0.30$ sits at the engineering-gain optimum while roughly halving f_n relative to catalysed-D–D.

THE BINDING GATE: THE PLUG POTENTIAL

Figure 8 plots the plug requirement of equation (7). Because the density ratio sits inside the exponent, the carrier temperature T_e is the dominant lever: the design point at $T_e = 89.7$ keV needs $n_p/n_c = 16$, but a modestly cooler carrier would drive the requirement up by an order of magnitude. This exponential sensitivity is why the synchrotron-pinned electron temperature of equation (11) is not a convenience but a prerequisite. Figure 8 shows the complementary Pareto structure (AG-CX-02): the plug potential saturates near 249 keV by a density ratio of ≈ 16, while the engineering gain rides a steep cliff above it (Q_E climbing from 1.32 at $R = 16$ through 2.7 at $R = 20$ to 23 at $R = 24$). The canonical operating point $R = 16$ is deliberately placed at the *robust knee* of this curve, not in the fragile steep zone: it is the point at which the confining potential is essentially fully formed and the design is insensitive to the exact ratio, at the cost of forgoing the large but fragile gains available further up the cliff. This is a conservative choice, and it is the honest one for a system whose plug requirement is not yet demonstrated. The density-ratio–potential Pareto and the sloshing-ion route to forming it are developed in the sloshing-ion companion [62].

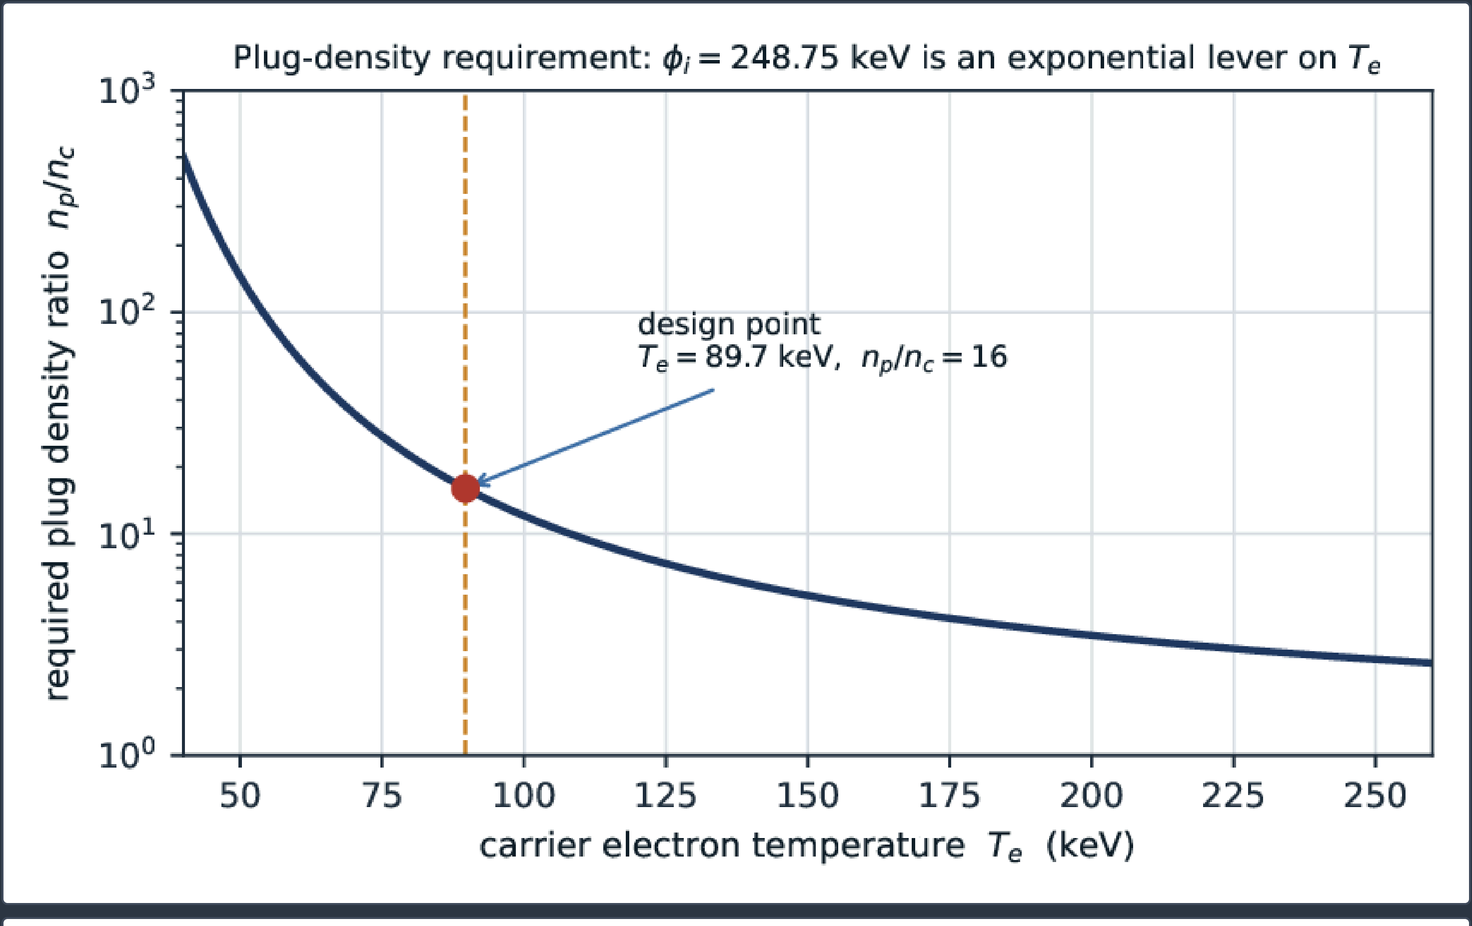

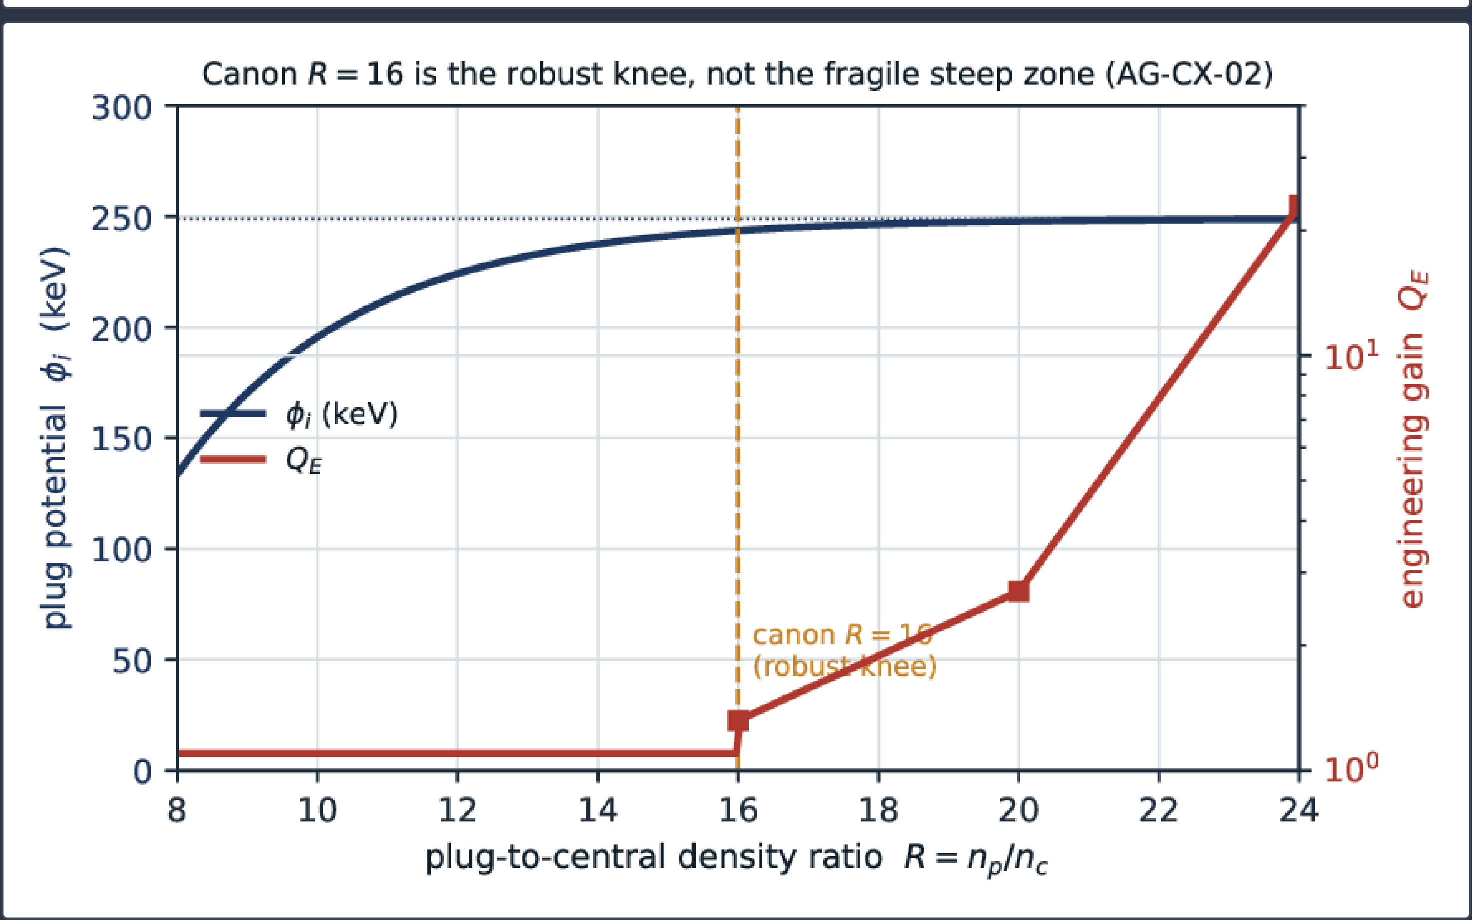

The plug-density requirement $n_p/n_c = \exp(e \phi_i / T_e)$ (gate BU-L2-A5). The design point sits at $n_p/n_c = 16$; the carrier temperature T_e is the exponential lever that sets it.

PASTUKHOV CONFINEMENT

Figure 9 plots the Pastukhov confinement law (8) as a function of the normalised plug depth $e\phi_i/T_i$. The $x\,e^x$ dependence is the quantitative statement of the plug's value: at the design depth $e\phi_i/T_i = 2.764$ the confinement time is $\tau_p = 3.73\,\mathrm{s}$, and a $10\,\%$ deeper well would nearly double it while a $10\,\%$ shallower one would nearly halve it. This is why the plug is worth its recirculating cost and why its formation is the binding gate: the confinement — and therefore the burn and the gain — is exponentially sensitive to a potential that has never been formed at this depth. The confinement enters the recirculation balance (9) through the burn it sustains, closing the loop between the plug gate and the engineering gain.

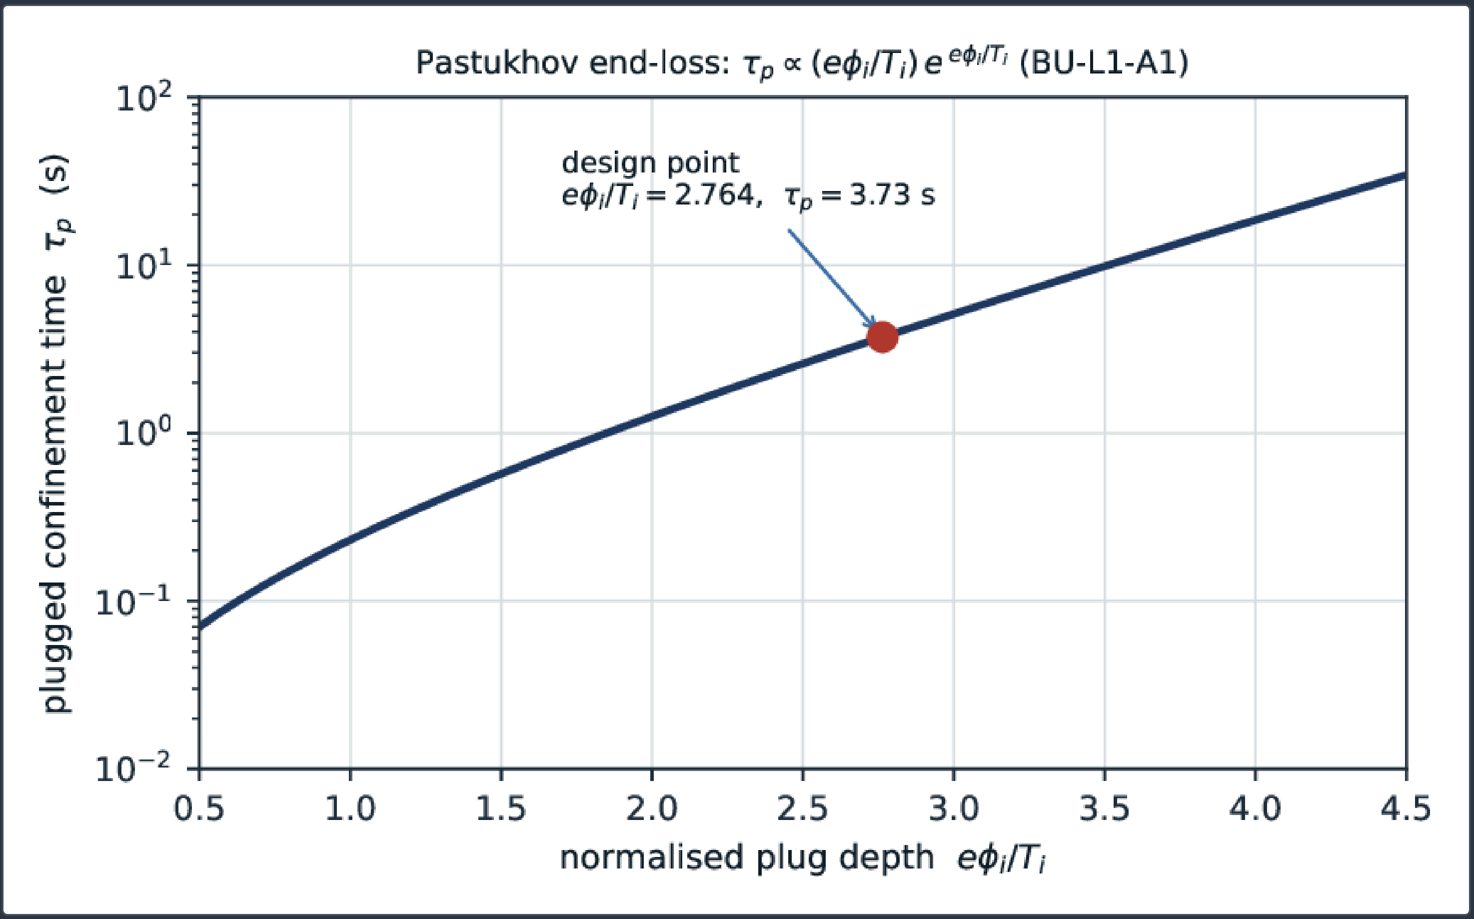

Pastukhov end-loss confinement time $\tau_p \propto (e\phi_i/T_i)\exp(e\phi_i/T_i)$ (BU-L1-A1). The design point at $e\phi_i/T_i = 2.764$ returns $\tau_p = 3.73\,\mathrm{s}$; the exponential dependence is the source of the plug's leverage and of its criticality.

THE TEMPERATURE RIDGE AND THE DESIGN-SPACE DATABASE

The net-gain database (BU-L2-A2) returns $Q_E > 1$ at every one of its 25 points, spanning 1.107–1.414, and exhibits a weak temperature ridge: at fixed design density Q_E peaks at 1.375 near $T_i \approx 85$ keV and falls gently to 1.318 at the canonical 90 keV (figure 10). The canonical point is therefore placed on the conservative shoulder of the ridge rather than on its crest, trading a few percent of gain for robustness against temperature excursions. That the entire database clears $Q_E > 1$ is the origin of the $P(Q_E > 1) = 1.0$ statement of §5; it is a statement about the design ranges, not an unconditional claim, and it is always read with the plug and converter levers attached.

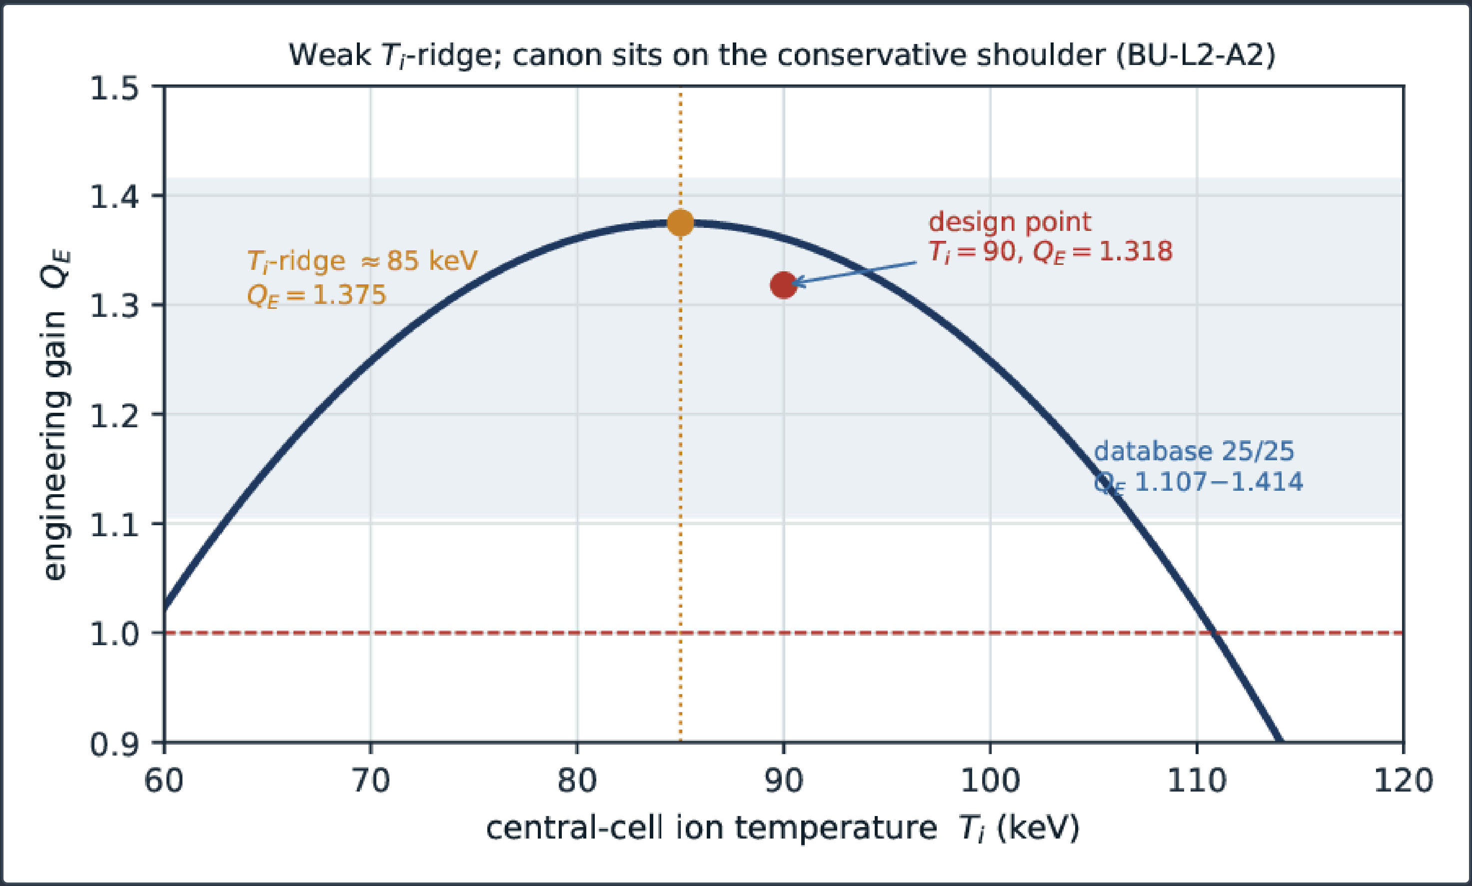

The engineering-gain temperature ridge (BU-L2-A2). At the design density Q_E peaks near $T_i \approx 85$ keV ($Q_E = 1.375$) and the canonical point (90 keV, $Q_E = 1.318$) sits on the conservative shoulder. The shaded band is the 25-point database span, $Q_E = 1.107$–1.414.

THE MIRROR-RATIO OPTIMUM

The mirror ratio trades confinement against expander area. Figure 11 shows the engineering gain rising and saturating with mirror ratio (AG-CX-06): Q_E climbs from 1.18 at $R_{mc} = 3.3$ through 1.32 at 4.6 to 1.53 at 6.8, while the expander area required to land the end-loss stream grows roughly linearly with R_{mc} through flux conservation (§2.13). The canonical $R_{mc} \approx 4.6$ sits at the economic knee where the gain has largely saturated but the expander has not yet become the dominant structure. The result is a genuine interior optimum, not a boundary choice, and it is robust to the exact expander cost model.

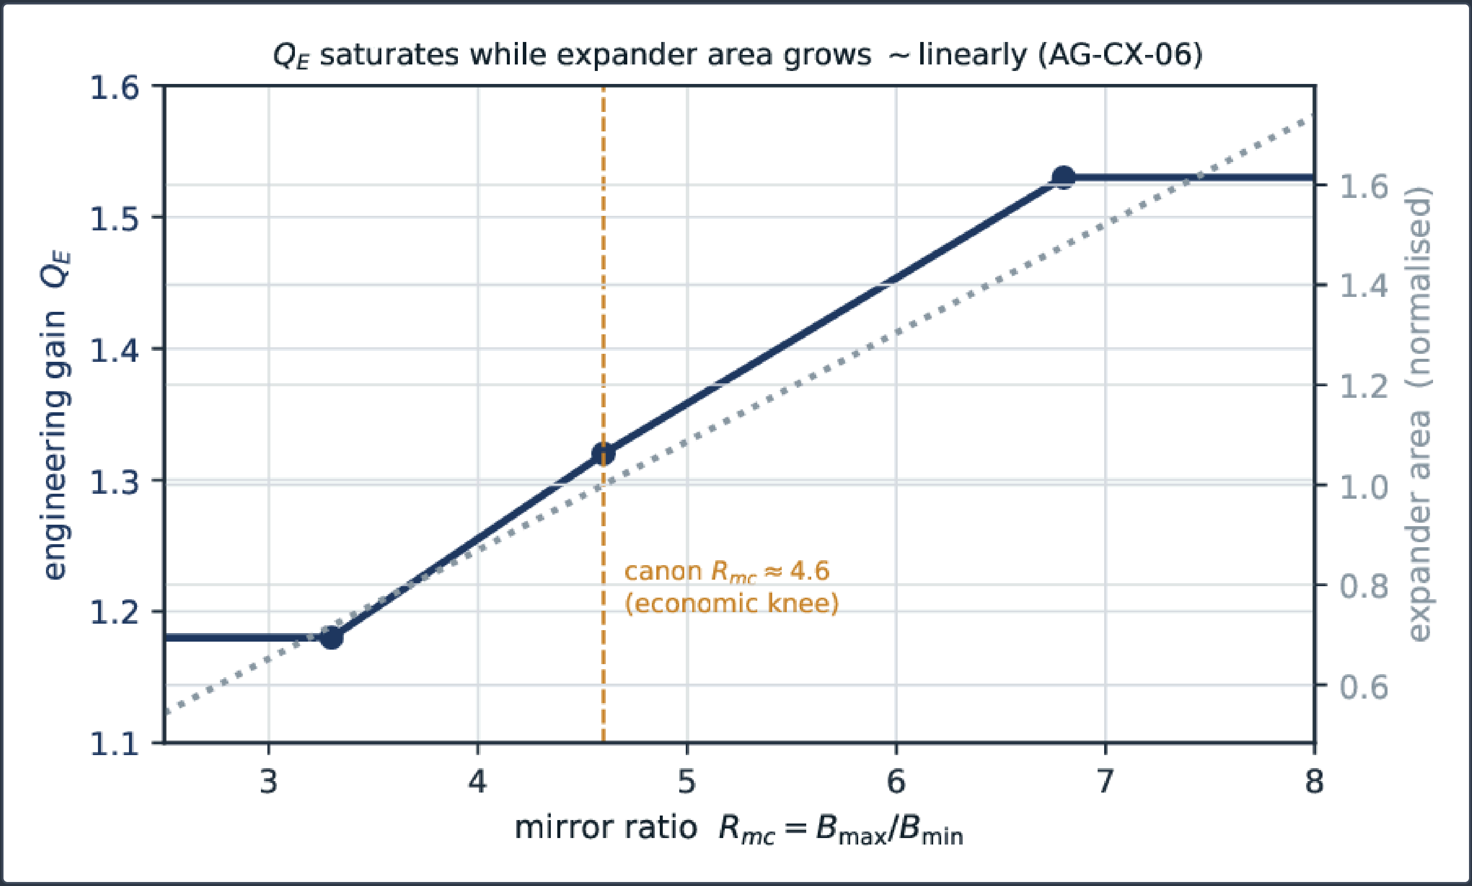

Mirror-ratio optimum (AG-CX-06). Q_E saturates with mirror ratio ($1.18 \to 1.32 \to 1.53$ over $R_{mc} = 3.3 \to 6.8$) while the expander area grows ~linearly; the canonical $R_{mc} \approx 4.6$ is the economic knee.

THE SECOND LEVER: DIRECT-CONVERTER EFFICIENCY AND NET-POWER BREAK-EVEN

Equation (9) is linear in η_{dec}, and figure 12 makes the consequence explicit. The converter module efficiency is a declared, load-bearing input on a conservative-to-ideal ladder: a conservative canon floor $\eta_{\mathrm{dec}} = 0.49$ — essentially the 0.48 that Barr and Moir measured on TMX [23] — a design basis $\eta_{\mathrm{dec}} = 0.70$ (the Moir–Barr venetian-blind design target [22]), and a proton-only ideal ceiling $\eta_{\mathrm{dec}} = 0.85$ (BU-L1-A10, BU-SX-03). Net electric requires η_{dec} above a break-even near 0.52; the design point $Q_E = 1.318$ is reached at the design-basis efficiency, while the conservative floor sits *below* break-even and does not, on its own, deliver net electric. We state this plainly: the converter efficiency is a binding lever of the same class as the plug gate, not a free upside, and the design's net-power claim is explicitly conditional on the converter clearing the break-even — a hardware, not a design-study, result, treated in the direct-conversion companion [54]. The architecture trade behind the ladder — electrostatic venetian-blind versus travelling-wave converters, and the beam-energy-spread dependence of each — is shown in figure 13 and developed in the collector and architecture companions [55,56].

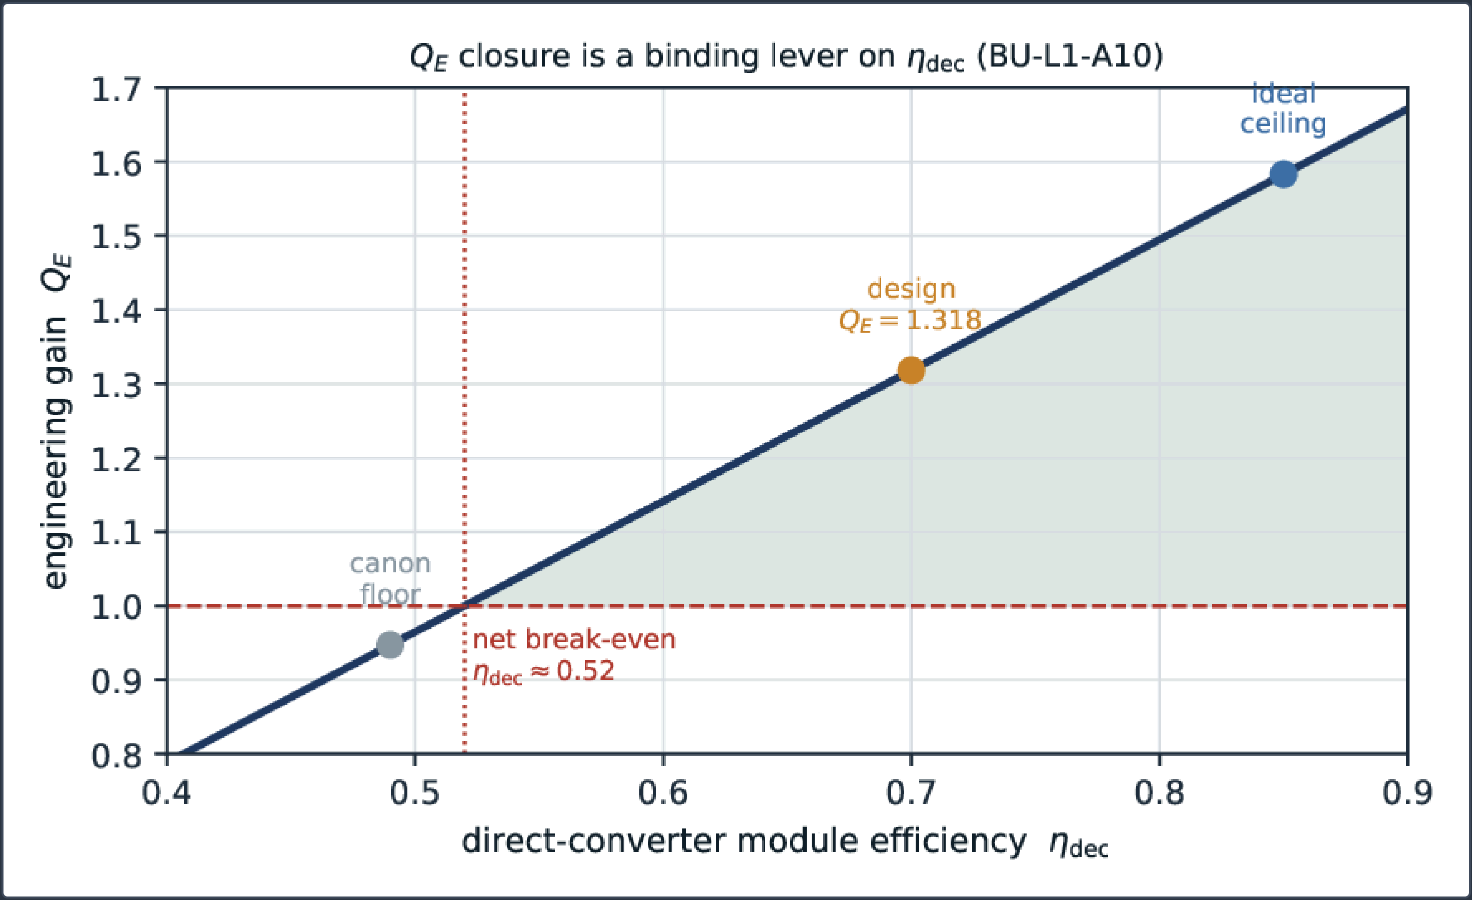

Engineering gain versus direct-converter efficiency (gate BU-L1-A10). Q_E is linear in η_{dec}; net electric requires η_{dec} above the break-even near 0.52. The design point ($\eta_{\mathrm{dec}} = 0.70$, $Q_E = 1.318$) closes; the conservative canon floor (0.49) does not. The converter efficiency is a binding lever, not a free upside.

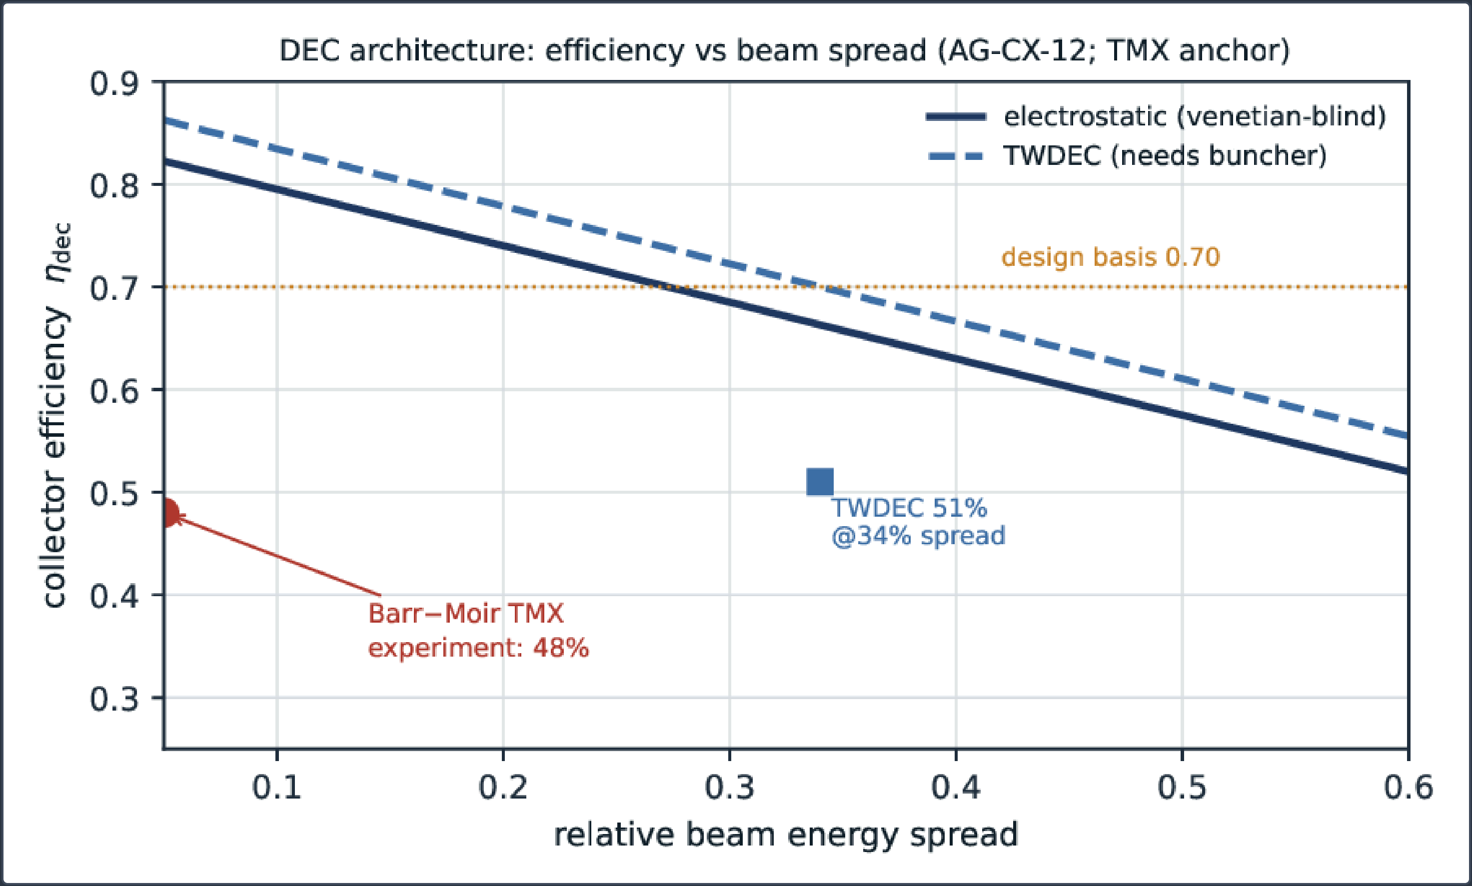

Direct-converter architecture trade (AG-CX-12). The electrostatic venetian-blind converter degrades with beam energy spread but tolerates the raw exhaust; the travelling-wave converter (TWDEC) needs an upstream buncher and returns only $\approx 51\,\%$ at the raw $\approx 34\,\%$ burner spread. The Barr–Moir TMX experiment ($48\,\%$) anchors the foot of the ladder [23].

SYNCHROTRON MODEL DEPENDENCE

Because the electron temperature is synchrotron-pinned and T_e enters the plug requirement exponentially, the choice of synchrotron model propagates directly to the gain. Figure 14 shows the spread (AG-DR-01): a thin-plasma estimate returns $Q_E = 0.39$, the AFJ reflective-cavity model (canon) returns $Q_E = 1.32$, and the Trubnikov model returns $Q_E = 4.56$. The AFJ model is the physically correct reflective-cavity treatment; within it, net break-even requires a wall reflectivity $R_{\text{wall}} \gtrsim 0.83$. We therefore carry the reflectivity as a hard first-wall requirement and report the gain against the conservative-to-standard model, never against the optimistic Trubnikov bound. This is one of the two places (with the converter ladder) where an optimistic modelling choice would inflate the headline, and we make the choice explicit rather than silent.

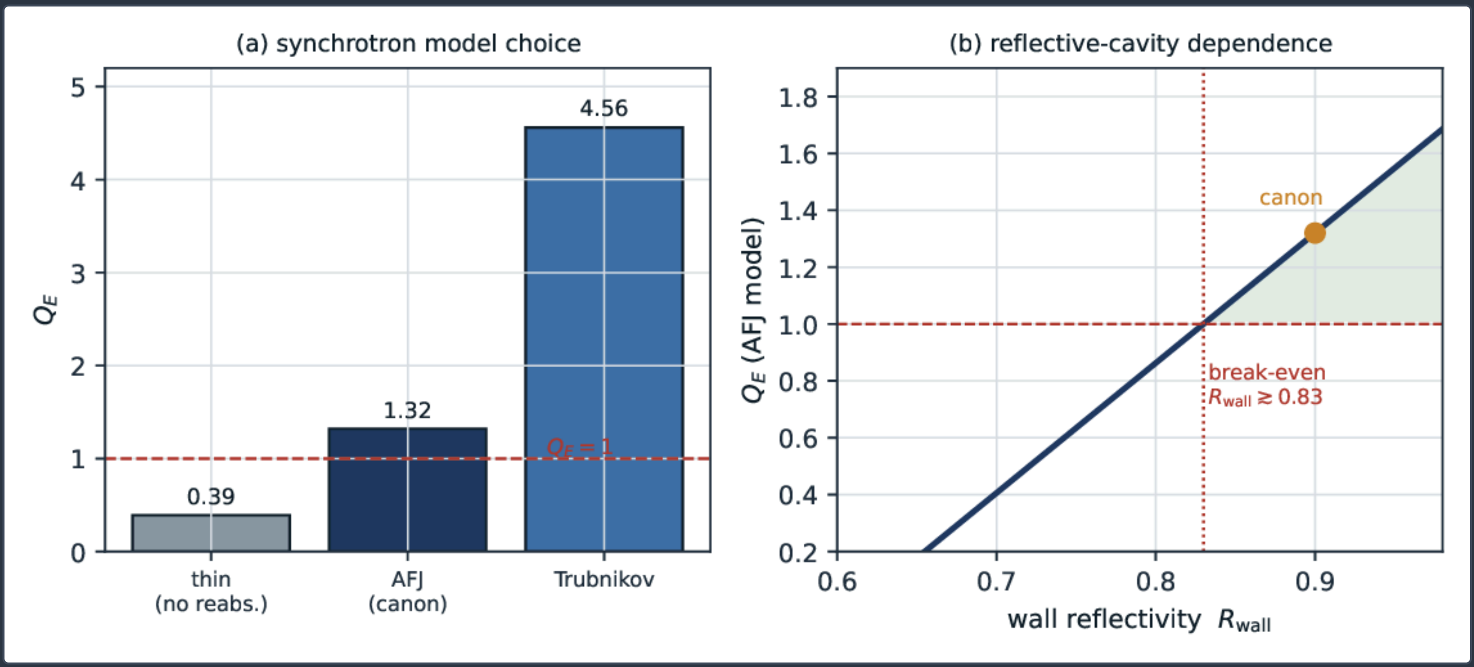

Synchrotron model dependence (AG-DR-01). (a) The engineering gain under three synchrotron models: thin-plasma (0.39), AFJ reflective-cavity (canon, 1.32), and Trubnikov (4.56). (b) Within the AFJ model, Q_E crosses net break-even at wall reflectivity $R_{\text{wall}} \approx 0.83$; the canonical point sits at $R_{\text{wall}} = 0.90$.

LOSS-CONE MICROSTABILITY

Figure 15 shows the loss-cone microstability result reproduced independently with GPU PIC (AG-CX-03) [38]. The bare loss-cone distribution is unstable, with normalised growth rate $\gamma/\omega_{ci} = 0.037$; a maintained energetic-stabiliser fill quenches the DCLC branch monotonically (0.024 at $f_s = 0.15$, 0.011 at 0.30, effectively suppressed by 0.45), consistent with the $6.25\,\%$ warm-plasma threshold of gate BU-L2-A5. Two honest qualifications attach. First, the stabilisation is a threshold, not a margin: the plug is stable only while the warm-plasma fraction is actively held, which makes the warm-fill an operational control quantity handed to the burner control suite [71]. Second, the warm fill quenches the DCLC branch but *not* the anisotropy-driven AIC branch, which the WarpX PIC returns as flat ($\gamma/\omega_{ci} \approx 0.037$) across warm-fill fractions $0–10\,\%$ (gate BU-L2-A7); this residual AIC activity is a distinct plug-gate caveat, non-showstopping, and is treated in full by the microstability companion [59]. Presenting this as a knee with a residual branch, rather than as a clean margin, is part of the claim.

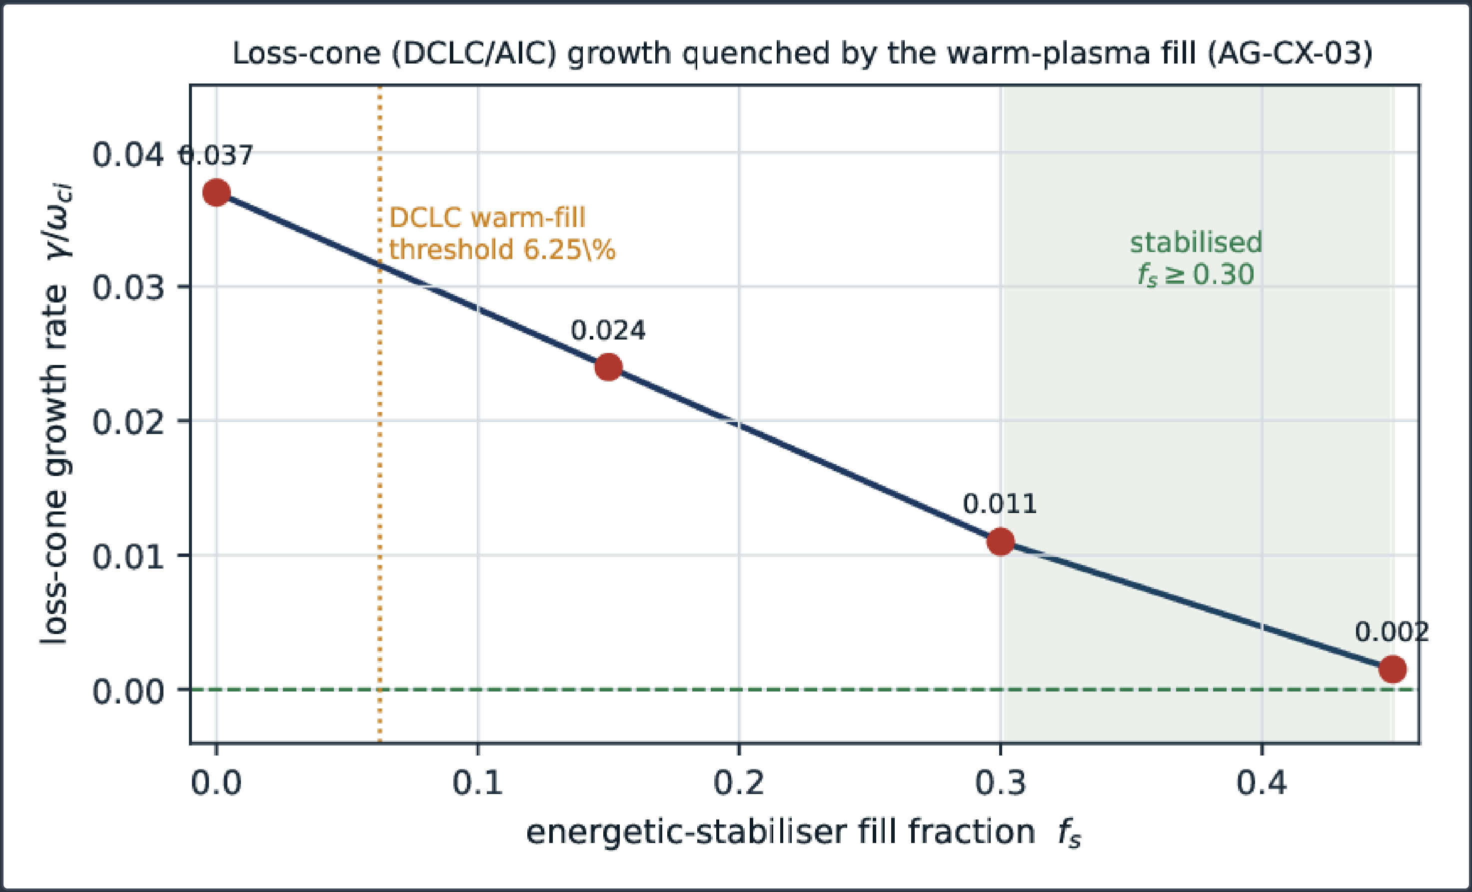

Loss-cone (DCLC) growth rate quenched by the energetic-stabiliser fill (AG-CX-03; GPU PIC). The bare distribution is unstable ($\gamma/\omega_{ci} = 0.037$) and the DCLC branch is monotonically stabilised as the warm-plasma fraction is raised past the $6.25\,\%$ threshold. Stability is a maintained-control condition, not a passive margin; the residual AIC branch (BU-L2-A7) is not quenched by the warm fill.

THE RF ALPHA-CHANNELING UPSIDE

Figure 16 shows the alpha-channeling lever (BU-SX-02). The collisional baseline leaves only $\approx 8\,\%$ of the fast-proton energy in a directly recoverable form; Fisch–Rax channeling [24,25] lifts the recoverable fraction toward $80\,\%$ as the normalised RF drive P^* approaches 16. We stress that this is an *upside*, not an assumption of the frozen design point: the $Q_E = 1.318$ figure does not require it. If the channeling can be driven — and its self-consistency depends on plug-localisation of the RF power, as §3 notes — it both raises the recoverable fraction feeding the converter and returns up to ≈ 195 MW to the ions, deepening the burn. It is reported here as a lever with its own gate, quantified in the alpha-channeling companion [58].

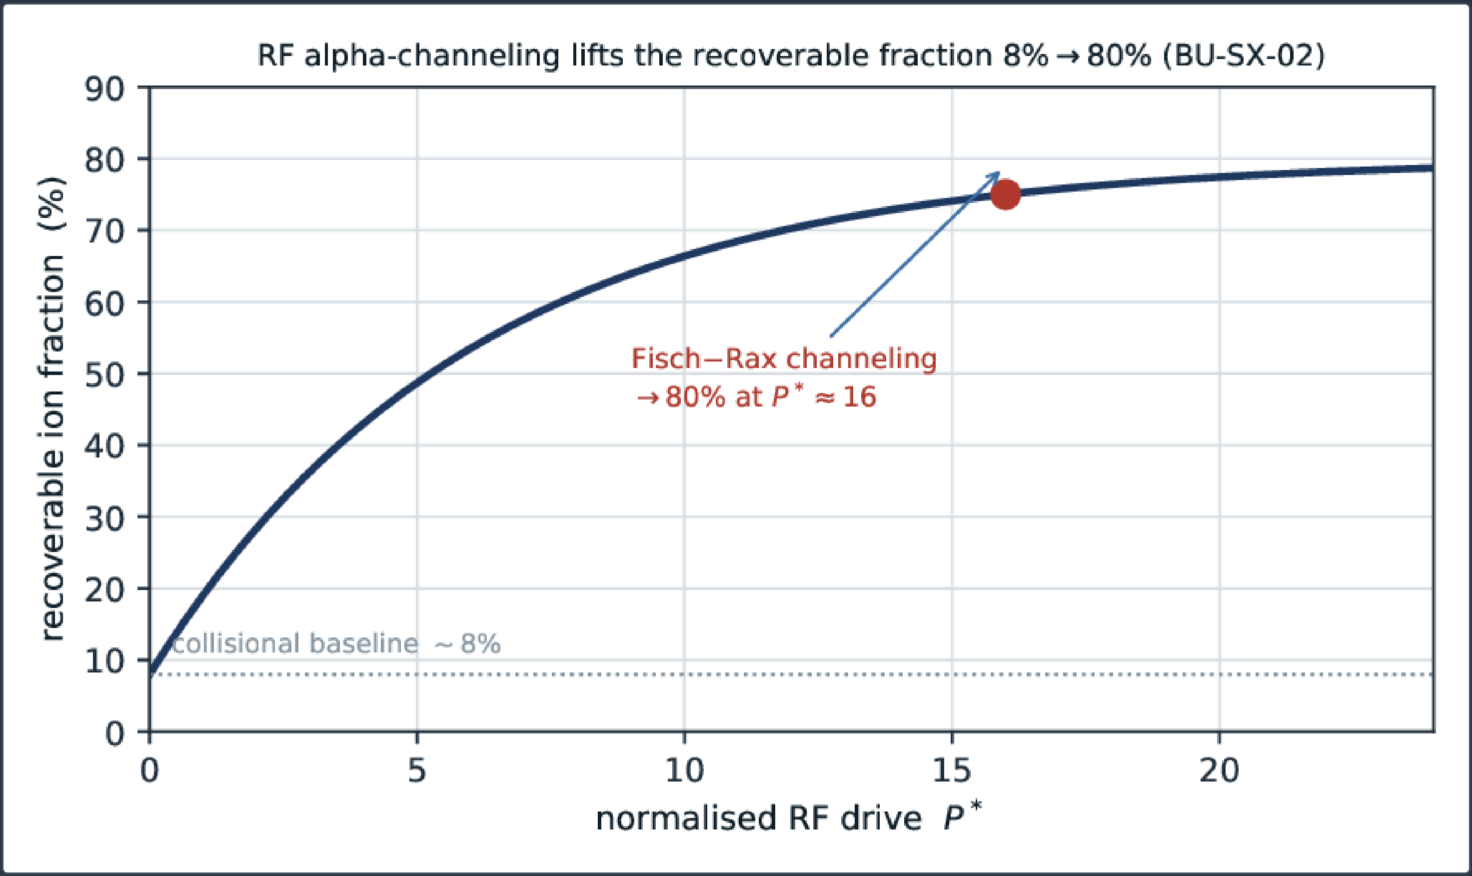

RF alpha-channeling (BU-SX-02). The collisional baseline recoverable ion fraction ($\approx 8\,\%$) is lifted toward $80\,\%$ at a normalised RF drive $P^* \approx 16$ by Fisch–Rax channeling. This is an upside lever, not a load-bearing assumption of the frozen design point.

SENSITIVITY AND UNCERTAINTY (OVERVIEW)

Which input actually controls the gain? Figure 17 gives the Saltelli–Sobol answer (HX-05): the plug density ratio is the first-order driver of Q_E, with a first-order index $S_1 = 0.41$ and a total-order index $S_T = 0.87$. The large gap $S_T - S_1 = 0.46$ is interaction variance, carried almost entirely with the central-cell beta. The corresponding Monte-Carlo distribution over the full design ranges (figure 17; KX-L1-A5-BU, 40 000 samples) is $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$. The deeper decomposition, and the net-power framing that returns $P(\text{net} > 0) = 73\,\%$, are the subject of §5.

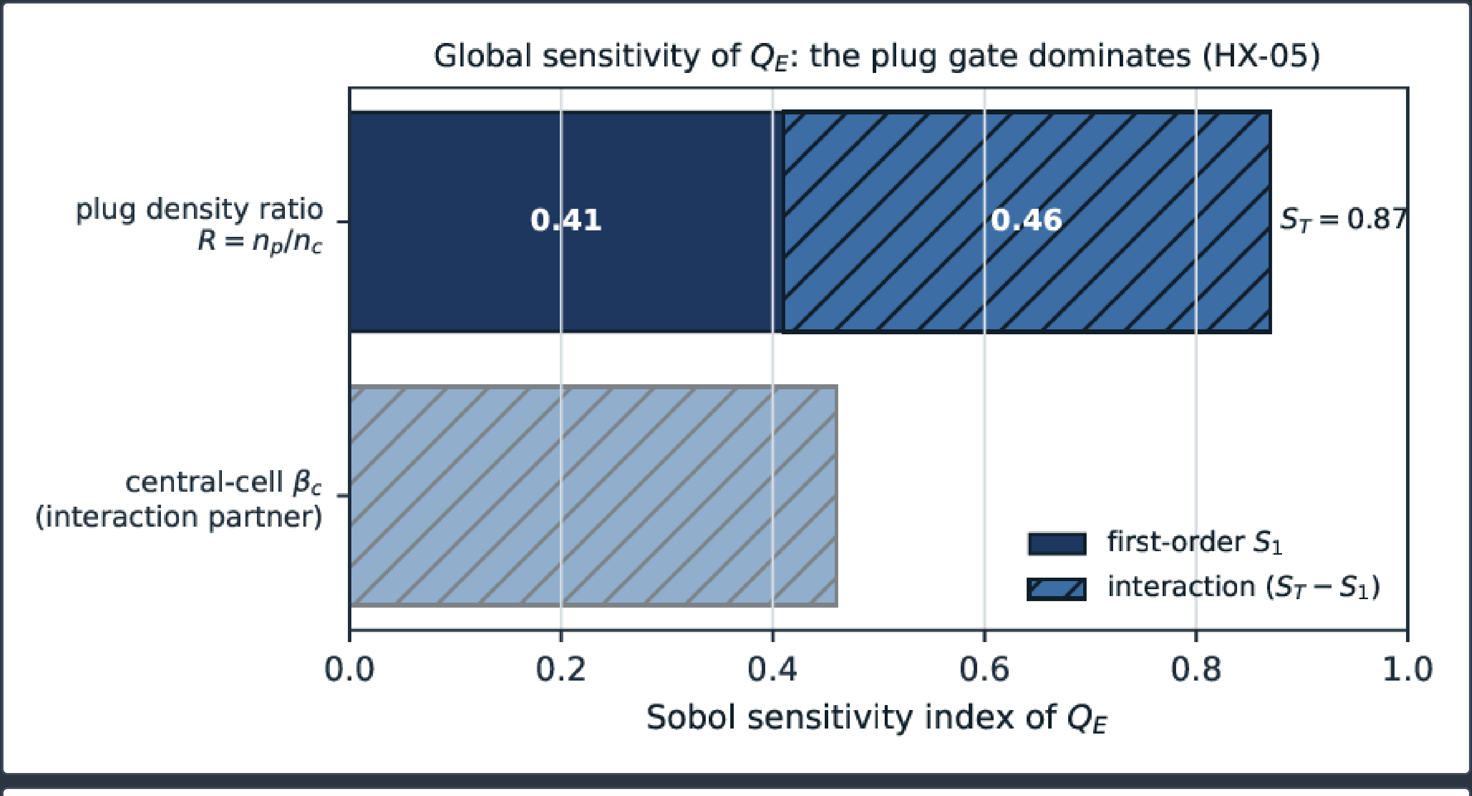

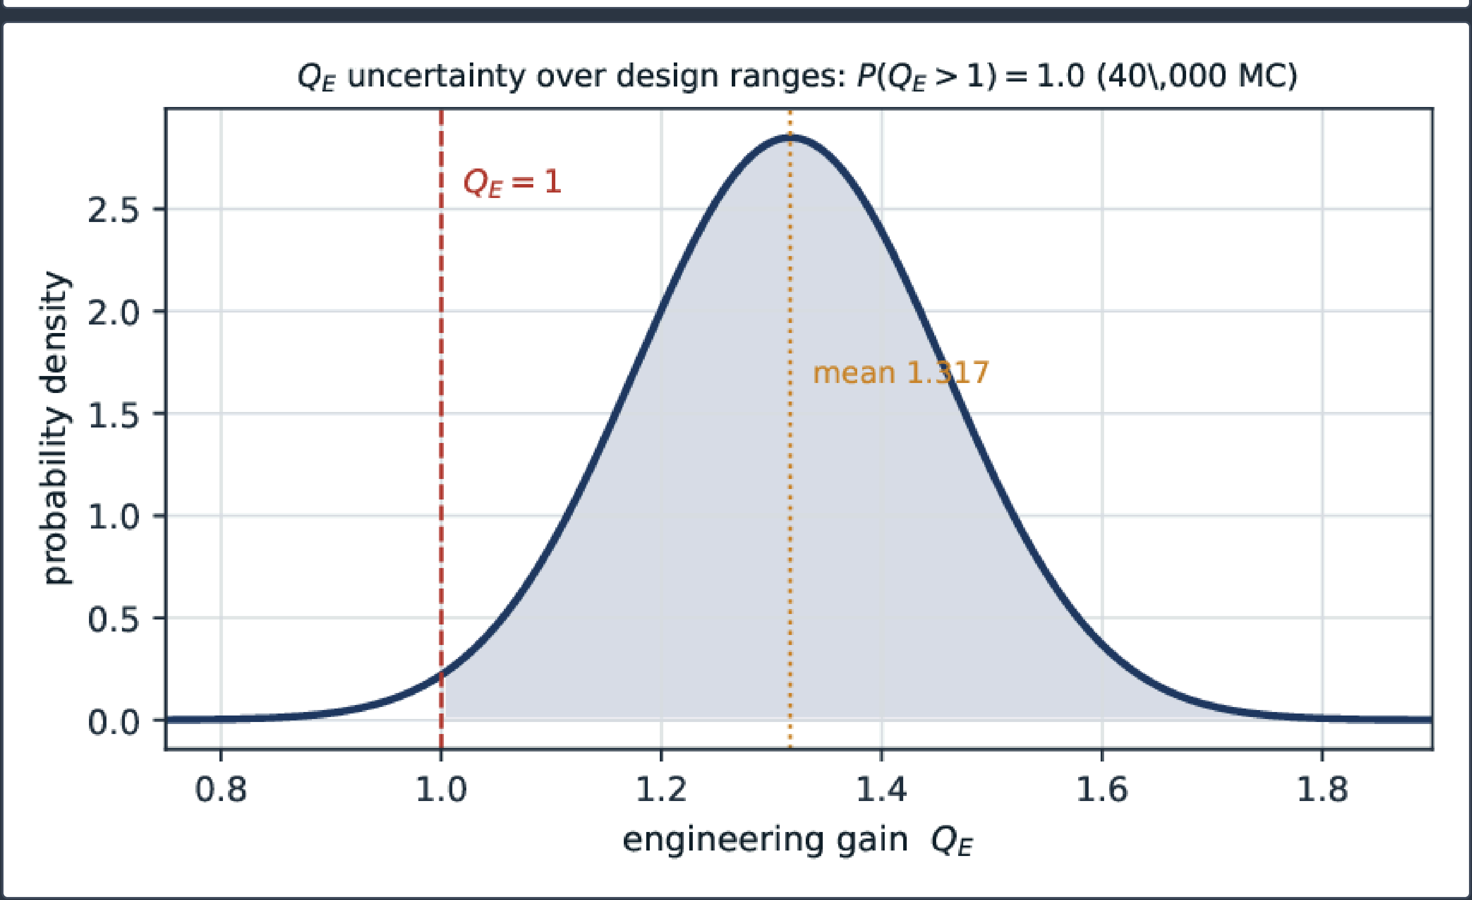

Global sensitivity of Q_E (Saltelli–Sobol, HX-05). The plug density ratio dominates: first-order $S_1 = 0.41$, total-order $S_T = 0.87$; the 0.46 interaction term is carried with the central-cell beta.

FORWARD-BASE OPERATING ENVELOPE

The physics core is the same for Aegis and MetroVolt; what makes the machine an *Aegis* is the forward-base operating envelope, which is where the product case is made or lost. Three quantities define it.

Black-start. A forward installation must start without the grid. Figure 18 shows the reduced black-start trajectory (HX-07, AG-CX-18): from a cold start the magnets charge, the plug potential builds, the central cell ignites, and the direct-converter output crosses the start-up draw (of order $340\ \mathrm{MW_e}$) at $\approx 42\,\mathrm{s}$, after which the plant is grid-independent and settles to its $+104\ \mathrm{MW_e}$ islanded steady state (AG-L1-A1). The superconducting-magnet charge is the long pole of the full cold-start sequence, which reaches first-burn in $\approx 7\,\mathrm{min}$ and full Q in $\approx 9\,\mathrm{min}$ (AG-CX-18); a warm fault-restart, with the magnets left charged, reaches full Q in $\approx 35\,\mathrm{s}$ (AG-CX-24), which is why magnet-preserving trips are the default fault response.

Islanded load-following. A defence site swings its directed-energy and radar loads by tens of megawatts. Figure 18 shows the mechanism (AG-CX-20): the direct converter trims on a millisecond timescale, and the $\approx 8\,\mathrm{s}$ plasma thermal lag is bridged by a small buffer of order $1\,\mathrm{MWh}$, so the installation follows its own load without a grid tie. The plasma throttle is bounded below by burn stability, which the control-barrier safety clamp enforces [71,70].

Availability. Figure 19 gives the fleet capacity factor by redundancy tier from an event-driven Monte-Carlo (HX-01, AG-L1-A3): steady-state availability of $83.3/95.7/98.8\,\%$ across the tiers, with a P10–P90 capacity-factor spread of 68–78, 86–90 and 94–$96\,\%$ respectively, under an $N = 4/2/2$ redundancy rule. The single-point-of-failure is the large superconducting magnet set, whose $\sim 720\,\mathrm{h}$ repair time dominates the availability even when everything else is redundant (AG-CX-22); coil demountability and quench protection are therefore the top reliability items [36], and they are a magnet-design, not a physics, problem [52].

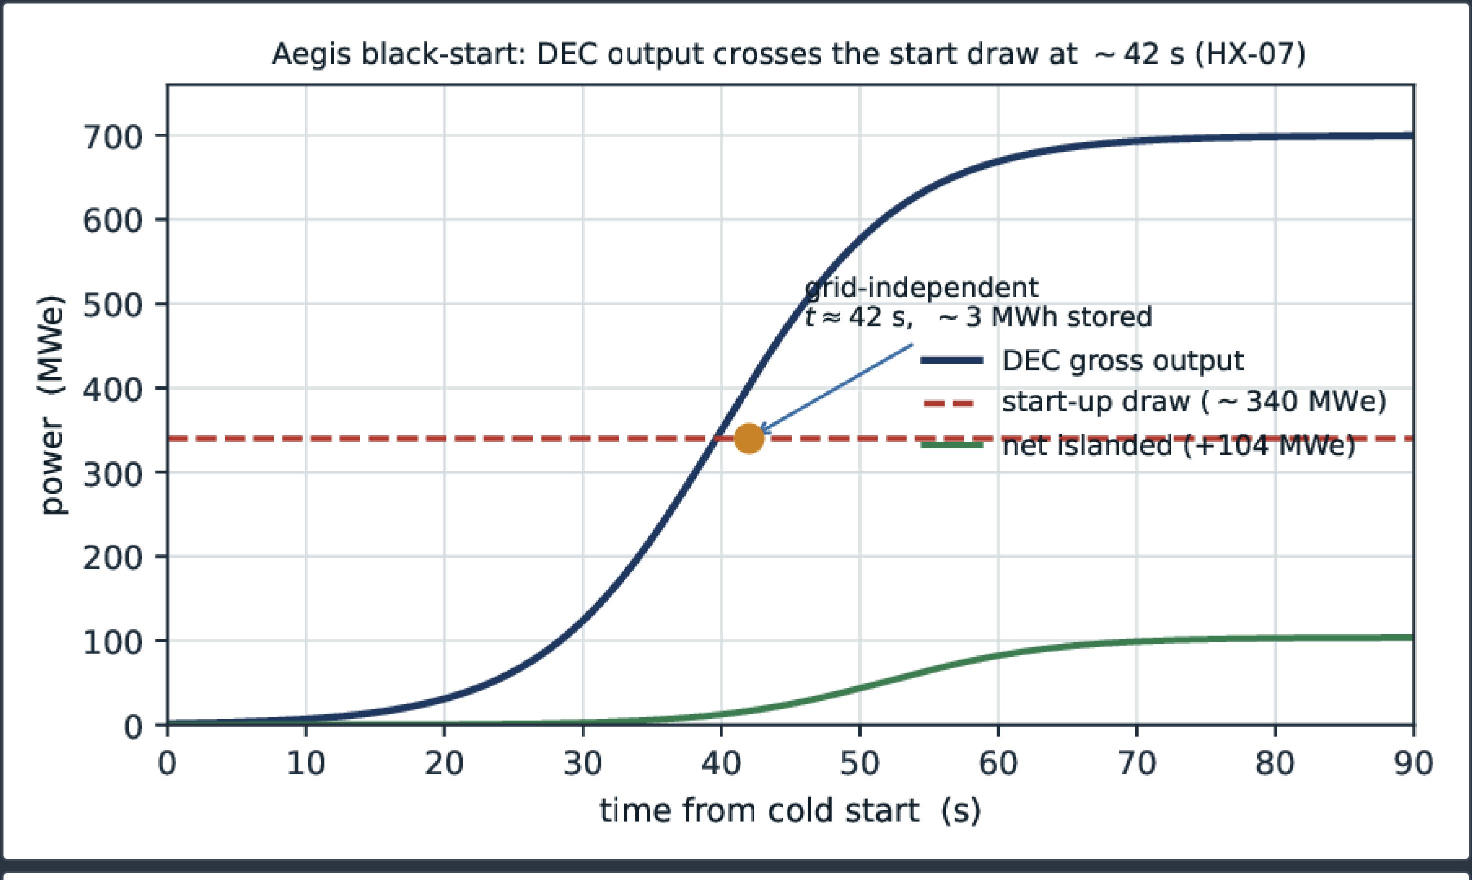

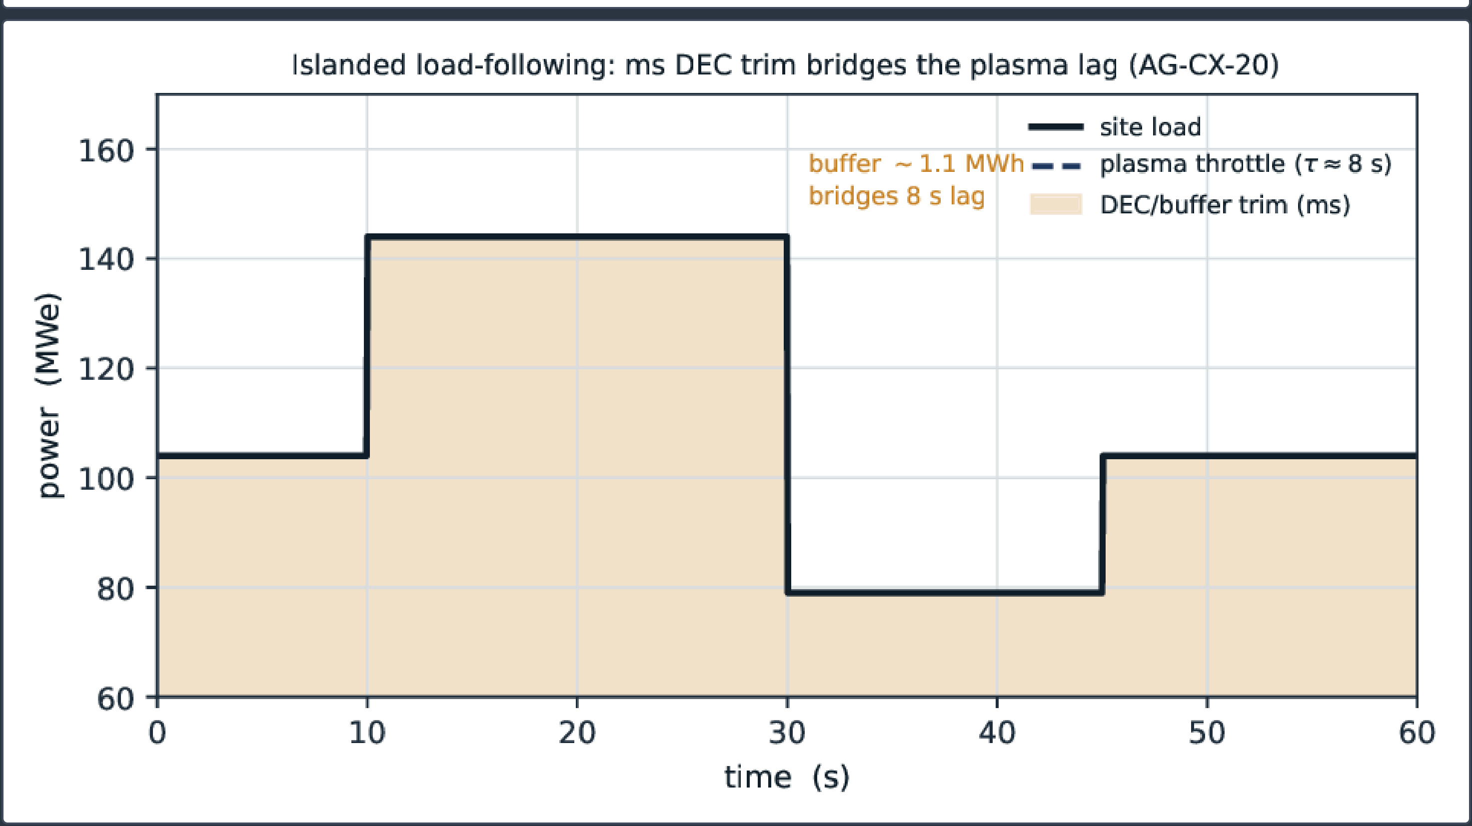

Aegis black-start (HX-07). The direct-converter output crosses the start-up draw at $\approx 42\,\mathrm{s}$; thereafter the installation is grid-independent at $+104$ $\mathrm{MW_e}$ islanded.

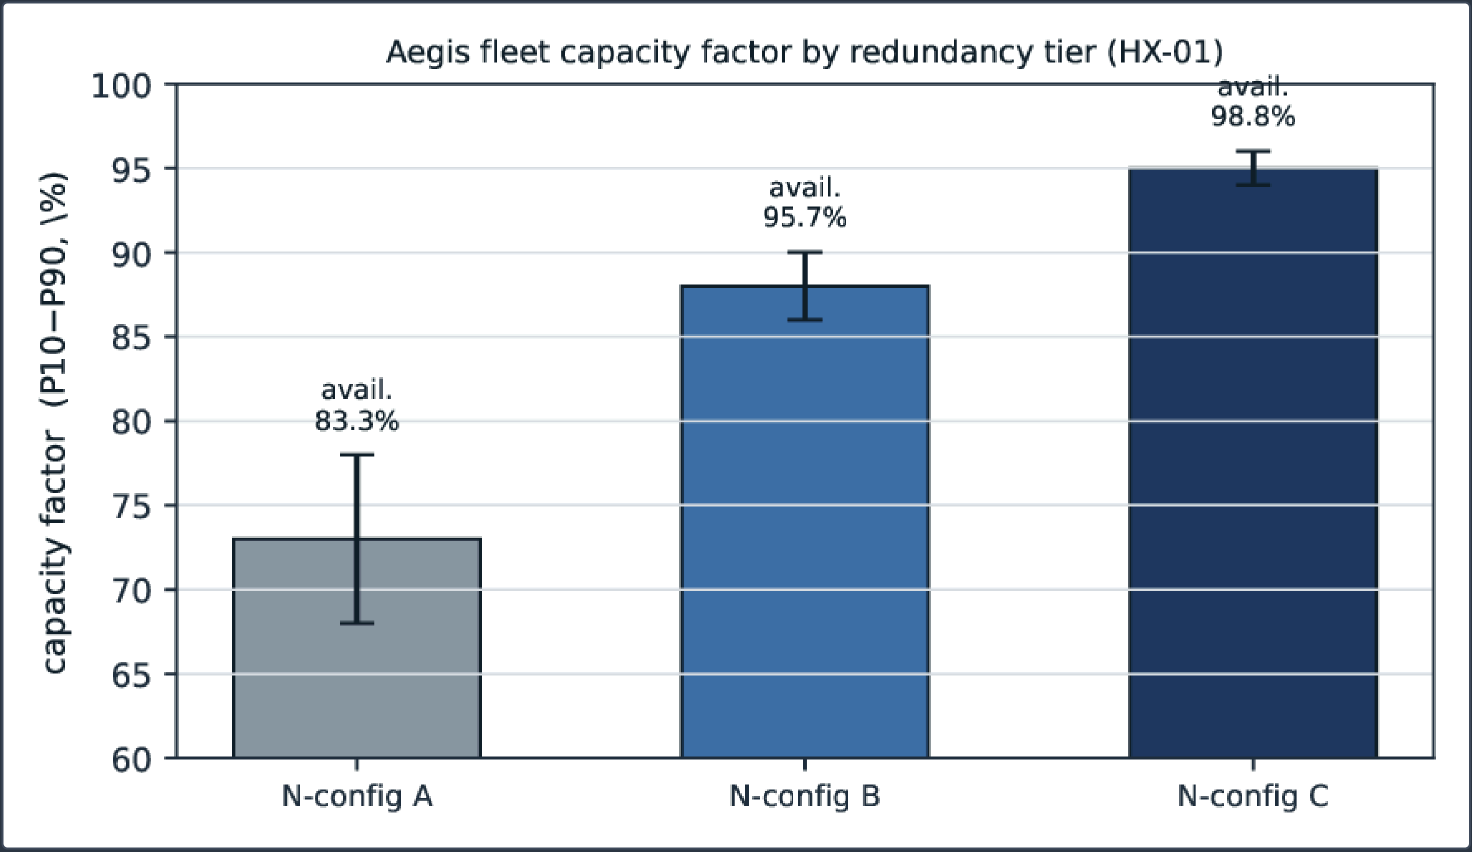

Aegis fleet capacity factor by redundancy tier (HX-01, $N = 4/2/2$). Bars show the P10–P90 spread; annotations give steady-state availability. The superconducting magnet set is the availability-limiting single point of failure.

THE PLUG MAGNET, STRUCTURE, DOSE AND FUEL LOGISTICS

The plug coil is the high-field no-insulation REBCO magnet [33,32,34,35] distinct from the breeder centrepost; its conservative field basis is $26.49\,\mathrm{T}$ (versus the breeder centrepost's $16.84\,\mathrm{T}$ — two separate magnets). Figure 20 shows the structural envelope (BU-L1-A2, AG-CX-26): the magnetic pressure rises as B^2, reaching $115/279/358\,\mathrm{MPa}$ at $17/26.49/30\,\mathrm{T}$, and the throat hoop stress reaches $\approx 814\,\mathrm{MPa}$ at the $26.49\,\mathrm{T}$ plug field; a graded-preload FEA closes with a reinforced allowable of $1091\,\mathrm{MPa}$, raising a thin-wall margin of $0.86\times$ to $1.34\times$ and holding the REBCO strain within a 2.85–$3.54\times$ margin (P3-FEA-10) [52,53]. The current-sharing temperature at full current is $57.1\,\mathrm{K}$, leaving a comfortable margin at the operating point.

Figure 20 shows the neutron-lean siting case (AG-L1-A2): because the source is $\approx 15\times$ leaner than a same-power D–T machine, a fixed shield of $\approx 4.5\,\mathrm{m}$ (a $20\,\mathrm{kt}$-class installation) brings the personnel-proximity dose below limit, where the D–T comparator would require far more mass at the same standoff. On fuel, the forward-base logistics are trivial by design: the in-vessel helium-3 start-up charge is $\sim 0.3\,\mathrm{g}$, a one-year grid-independent autonomy package is $\approx 21.3\,\mathrm{kg}$ of helium-3 plus $\approx 30\,\mathrm{kg}$ of deuterium (a ~ 1-tonne package occupying under $0.8\,\mathrm{m}^3$), and there is no fuel convoy to interdict (AG-L1-A6, AG-L1-A8). The binding fuel constraint is not on-site but upstream: helium-3 at burner scale is lunar-gated, and the breeder is its strategic terrestrial co-producer at $37.66\,\mathrm{kg}$-$^3\mathrm{He}$ per full-power year per unit (AG-CX-17) [68,50,18]. Table 4 collects the forward-base envelope.

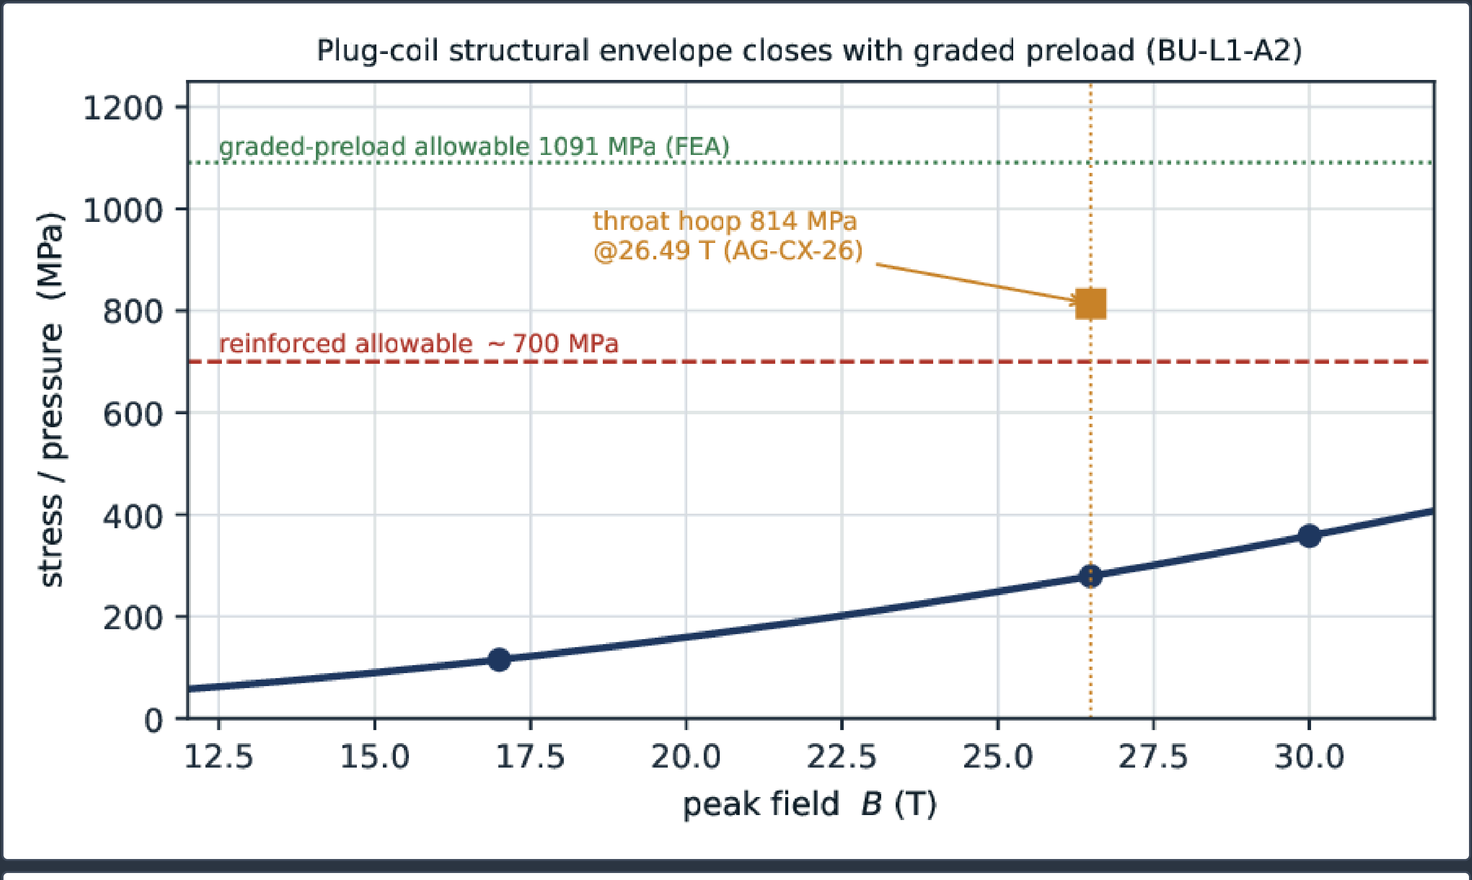

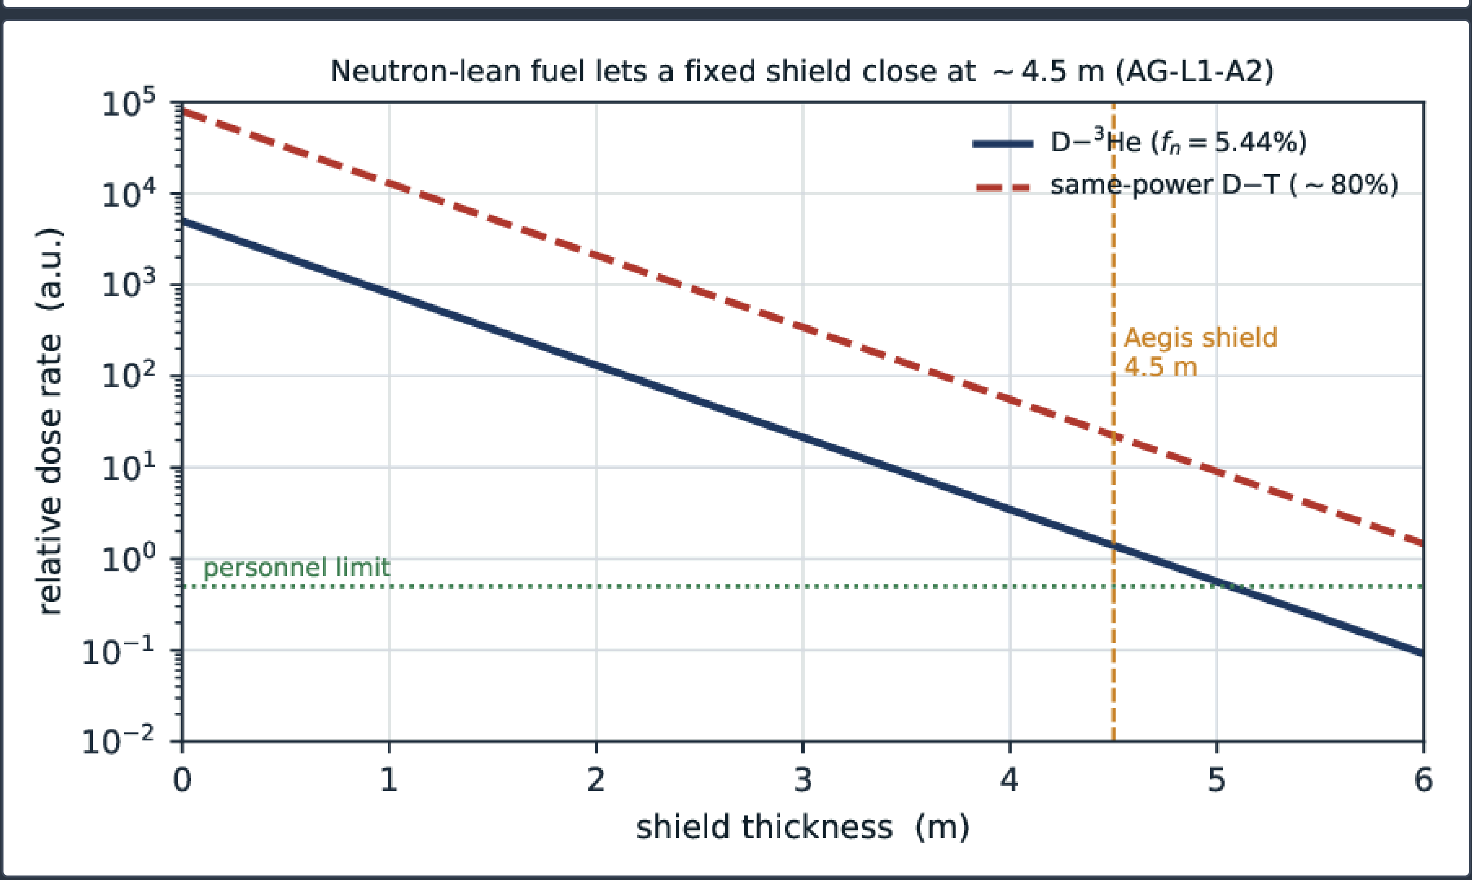

Plug-coil structural envelope (BU-L1-A2, AG-CX-26). Magnetic pressure rises as B^2; the throat hoop stress reaches ≈ 814 MPa at 26.49 T and closes against the graded-preload FEA allowable of 1091 MPa.

QUANTITY	VALUE	SERIAL
islanded steady-state net	$+104$ $\mathrm{MW_e}$	AG-L1-A1
black-start crossover	$\approx 42\,\mathrm{s}$	HX-07
cold start to full Q	$\approx 9\,\mathrm{min}$	AG-CX-18
warm fault-restart to full Q	$\approx 35\,\mathrm{s}$	AG-CX-24
start-up external draw	~ 340 $\mathrm{MW_e}$	AG-L1-A1
load-following buffer	$\sim 1\,\mathrm{MWh}$ (ms DEC trim)	AG-CX-20
availability (redundancy tiers)	$83.3/95.7/98.8\,\%$	HX-01
capacity factor (P10–P90)	$68–96\,\%$	HX-01
1-year fuel package	$21.3\,\mathrm{kg}$ $^3\mathrm{He} + 30\,\mathrm{kg\,D}$	AG-L1-A8
in-vessel ^{3}He start charge	$\sim 0.3\,\mathrm{g}$	AG-CX-17
shield (fixed installation)	$4.5\,\mathrm{m}$ / $20\,\mathrm{kt}$ class	AG-L1-A2
plug coil field basis	$26.49\,\mathrm{T}$	P3-FEA-10
throat hoop stress	$\approx 814\,\mathrm{MPa}$	AG-CX-26

Aegis forward-base operating envelope. Serials refer to the register [48].

VERIFICATION STATUS

Table 5 states, for the load-bearing claims, whether each is robust (independent of the plug gate), conditional (contingent on retiring the plug and/or converter lever), or a declared input. The neutron-source and low-activation performance is robust; the net-power performance is conditional; the converter efficiency is a declared input on the ladder of figure 12. Naming the class of each claim is the honest way to present an advanced-fuel mirror.

CLAIM	VALUE	CLASS
neutron fraction f_n	5.44 %	**robust**
neutron power P_n	233.95	**robust**
fixed-installation shield	$4.5\,\mathrm{m}$ class	**robust**
engineering gain Q_E	1.318	conditional (plug + η_{dec})
islanded net electric	$+104$ $\mathrm{MW_e}$	conditional (plug + η_{dec})
plug density ratio n_p/n_c	16	requirement (open gate)
converter efficiency η_{dec}	$0.49/0.70/0.85$	declared input
DCLC loss-cone stability	$f_s \geq 6.25\,\%$	threshold (maintained)
residual AIC branch	$\gamma/\omega_{ci} \approx 0.037$	open caveat
plug-coil structure	$1.34\times$ margin	closed (FEA)

§ 5

Sensitivity and uncertainty quantification

VARIANCE DECOMPOSITION AND THE DOMINANT DRIVER

The Sobol decomposition of §4.11 answers the design question "what must be controlled?". The first-order index $S_1 = 0.41$ says that $41\,\%$ of the variance of Q_E is explained by the plug density ratio acting alone; the total-order index $S_T = 0.87$ says that $87\,\%$ of the variance vanishes if the plug ratio is fixed. The gap $S_T - S_1 = 0.46$ is interaction variance, and it is carried almost entirely with the central-cell beta β_c: the plug gate and the central beta are the two quantities that jointly set the gain, and both are on the register's critical list. This is the quantitative statement of the paper's thesis — that the plug is the binding gate — and it also flags β_c (an assumed input, §2.11) as the second quantity most in need of an MHD-solved value. The Saltelli estimator's convergence (figure 3b) shows the indices are converged to better than a percent at the $40\,000$-sample base.

TWO UNCERTAINTY FRAMINGS

There are two legitimate ways to report the uncertainty, and they answer different questions (figure 21). The *design-range* framing (KX-L1-A5-BU) samples the design parameters within their specified ranges and returns $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$: within the design envelope, the engineering gain clears unity everywhere. The *net-power* framing (AG-DR-02) is deliberately harsher — it propagates the input uncertainties, including the converter efficiency and the synchrotron model, all the way to the net-electric sign, and returns $P(\text{net} > 0) = 73\,\%$ with Q_E spanning p5–p95 of 0.81–2.51. The difference is not a contradiction but a statement of where the machine lives: comfortably above $Q_E = 1$ within the design ranges, but near the net-electric cliff once the converter and synchrotron levers are allowed their full uncertainty. Both numbers are reported, and the design's net-power claim is always attached to the levers that move it.

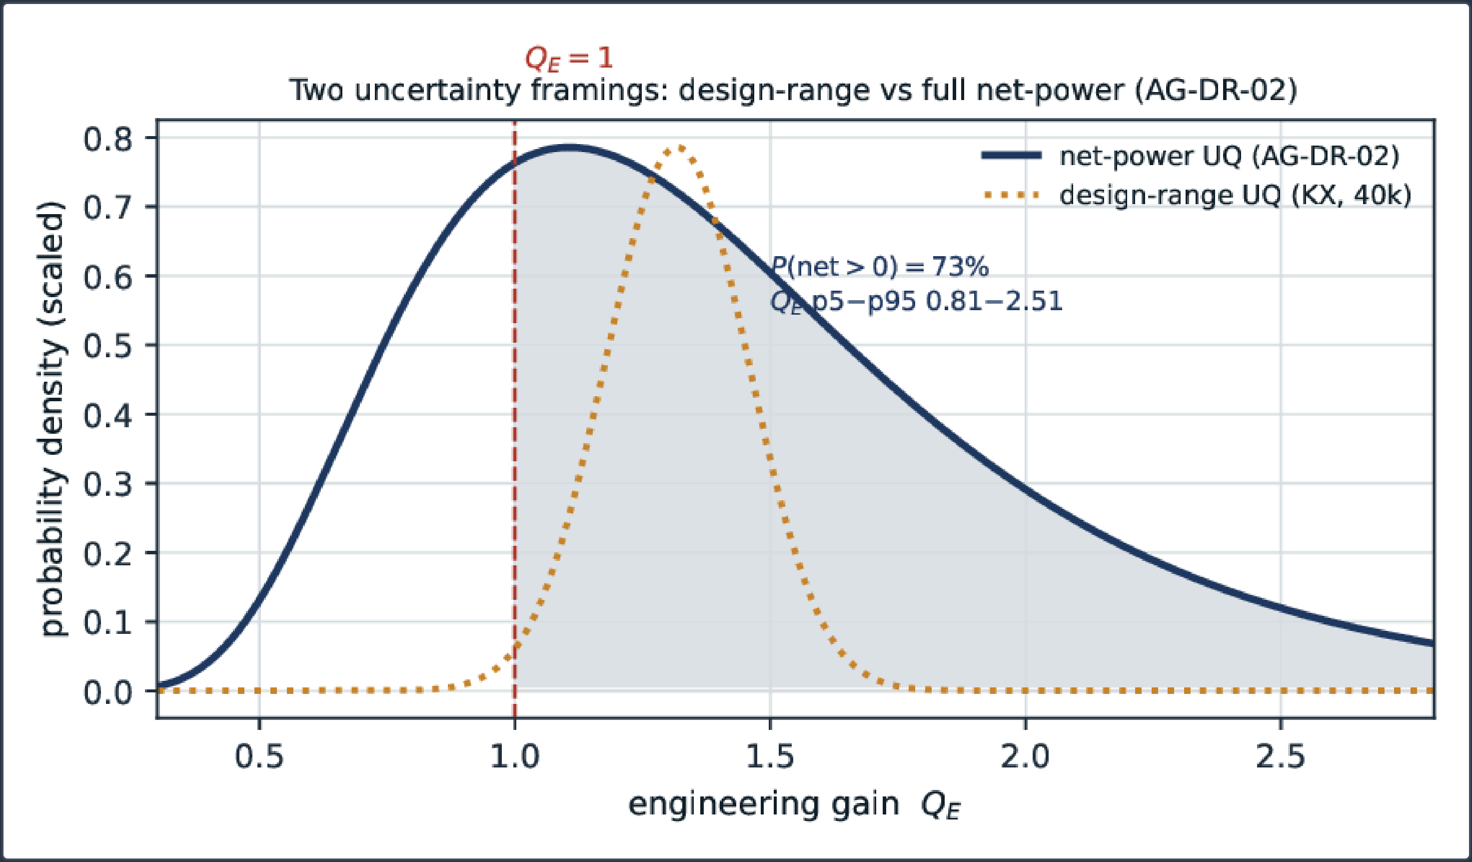

Two uncertainty framings. The tight design-range distribution (amber; KX, $40\,000$ samples, $Q_E = 1.317 \pm 0.14$, $P(Q_E > 1) = 1.0$) answers "does the gain clear unity within the design envelope?"; the broader net-power distribution (navy; AG-DR-02, $P(\text{net} > 0) = 73\,\%$, p5–p95 0.81–2.51) answers "does the plant reach net electric once the converter and synchrotron levers take their full uncertainty?".

THE VARIANCE BUDGET AND ITS CONSEQUENCES FOR THE PROGRAMME

The variance budget makes the experimental programme obvious. Retiring the plug gate collapses 87 % of the Q_E variance; fixing the central beta with an MHD solution removes most of the remaining interaction term; and demonstrating the converter efficiency and the wall reflectivity moves the net-power distribution off the cliff. In other words, three experiments — a plug-potential demonstration, a central-cell MHD/beta demonstration, and a converter-efficiency demonstration — would convert the conditional net-power claims into measured ones, in order of leverage. This is the concrete meaning of “the plug is the binding gate”: it is not rhetoric but the top line of a variance budget.

§ 6

Benchmarking against experiment and theory

AGAINST THE TANDEM-MIRROR RECORD

The design's confinement rests on a plug potential of 248.75 keV, and the honest comparison is to what the tandem-mirror programme actually demonstrated. GAMMA-10, TARA and MFTF-B formed thermal-barrier and plug potentials of order a kilovolt and confirmed the tandem-mirror confinement principle [9,6,5]; the GDT line demonstrated stable high-β axisymmetric operation and, with ECH, $\gtrsim 0.9$ keV bulk electrons [11,10]; and WHAM is now demonstrating the 17 T HTS field and sloshing-ion stabilisation that the Aegis throat assumes [14,12,13]. The Aegis plug requirement is more than two orders of magnitude in potential beyond that record. Figure 22 states the extrapolation honestly: our zero-dimensional model is benchmarked on the WHAM *unplugged* regime, where it agrees; the plugged branch on which the burner depends is a model extrapolation of two to three orders in confinement enhancement that WHAM does not yet reach (gates BU-L2-A3, AG-DR-04). We therefore do not claim the plug retired — we claim it is the binding gate, and that the whole design is structured to make that gate the single dominant risk (the Sobol result of figure 17 is the quantitative statement of exactly that). A WHAM-class plug-potential experiment at reactor-relevant field is precisely the demonstration that would retire the gate, and it is the experiment against which our conditional claims should be read.

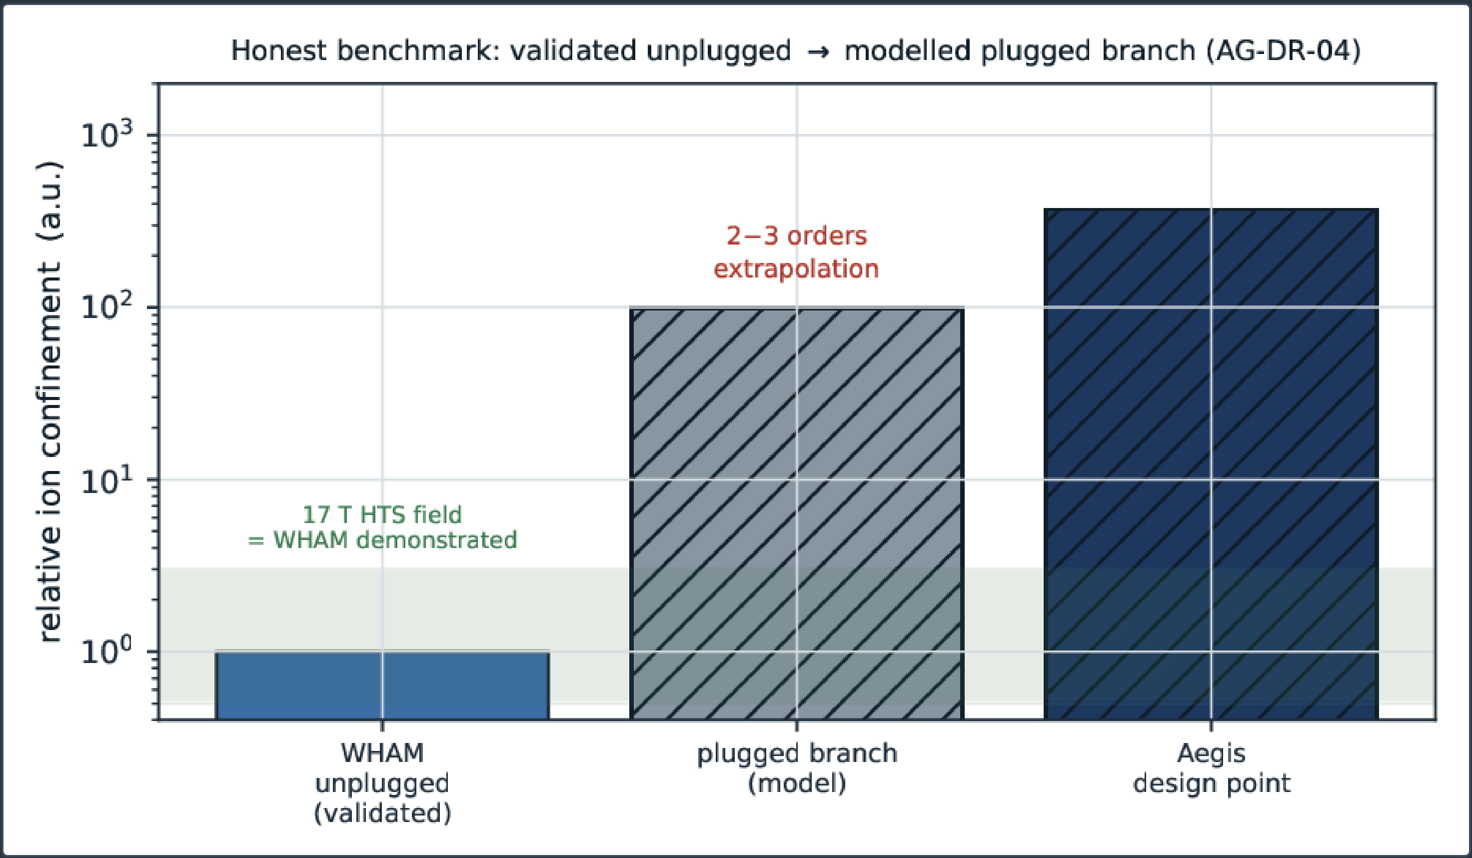

Honest benchmark of the confinement model (BU-L2-A3, AG-DR-04). The zero-dimensional model is validated on the WHAM unplugged regime; the plugged branch the burner depends on is a two-to-three-order model extrapolation. The 17 T HTS throat field is itself WHAM-demonstrated.

Table 6 places the Aegis operating point against the demonstrated tandem-mirror and axisymmetric-mirror record. The pattern is consistent across the parameters: the individual enabling ingredients have each been demonstrated somewhere — kV-class ambipolar potentials and thermal barriers on GAMMA-10 [9], direct conversion at 48 % on TMX [23], high-β (≈ 0.6) axisymmetric operation and ≈ 0.9 keV electrons on GDT [11,10], and the 17 T HTS field on WHAM [14] — but the *simultaneous* combination the burner needs, and above all the two-order jump in ambipolar potential to ≈ 249 keV, has not. The table is the compact statement of the paper's honesty posture: every row is a demonstrated ingredient, and the burner's claim is their combination, gated by the potential row.

QUANTITY	GAMMA-10	TMX-U	GDT	WHAM	AEGIS (REQ.)
ambipolar/plug potential	kV-class	kV-class	sub-kV	(comm.)	≈ 249 keV
mirror/throat field (T)	$\lesssim 1$	$\lesssim 1$	$\lesssim 1$	17	17 (plug 26.49)
electron temperature	$\lesssim 1$ keV	$\lesssim 1$ keV	≈ 0.9 keV	(comm.)	89.7 keV
central-cell β	low	low	≈ 0.6	(comm.)	0.55
direct conversion	—	48 % (test)	—	—	0.49/0.70/0.85
fuel	H/D	D	H/D	H/D	D–^{3}He

Aegis design point against the demonstrated mirror record. Machine entries are demonstrated values (order of magnitude where a range applies); Aegis entries are the required/design values. Citations in text.

A DATA-EFFICIENT CALIBRATION PATH

The extrapolation is not open-ended: a Bayesian calibration of the plugged-branch confinement discrepancy against WHAM anchors recovers the discrepancy parameter with a 90 % credible interval that tightens below ≈ 0.1 within ≈ 5 WHAM data points (gate AG-CX-31). This is a concrete, data-efficient validation path with error bars, and it is the bridge between the WHAM programme and the burner design point.

AGAINST THE DIRECT-CONVERSION LITERATURE

Direct energy conversion of a mirror end-loss stream has a real experimental pedigree: Barr and Moir recovered a net 48 % from a plasma stream on TMX with a single-stage venetian-blind converter [23], and the venetian-blind concept itself is Moir and Barr's [22]. Our treatment is deliberately conservative relative to this record: we take the converter efficiency as a declared ladder anchored at the demonstrated ≈ 0.48 ($= 0.49$ canon floor), rising to the 0.70 design target and a 0.85 proton-only ceiling, and we show the net-power break-even near 0.52 explicitly (figure 12) rather than assuming it away. The travelling-wave alternative returns only ≈ 51 % at the raw burner beam spread and needs an upstream buncher (figure 13; AG-CX-12), which is why the electrostatic converter is the design baseline.

AGAINST THE ADVANCED-FUEL CRITIQUES

The advanced-fuel critiques of Rider and Nevins [19,20,21] make the general point that non-thermal and advanced-fuel systems pay a steep recirculating-power penalty; our $Q_E = 1.318$ is a direct, quantitative acknowledgement of that penalty rather than a refutation of it — it is a low gain, honestly reported, and it is why the neutron-source and firm-power mission, not a high fusion gain, is the load-bearing value proposition of the machine [17,77]. The comparison to a compact high-field toroidal breeder [29,30,31] is complementary, not competitive: the breeder is the 14 MeV neutron source and helium-3 co-producer, and the burner is the low-neutron firm-power sink; the two are the two halves of one strategic-materials platform [50,51].

§ 7

Design implications

WHY AEGIS SPECIFICALLY

The forward-base case does not turn on a high gain — it turns on three properties a diesel or fission alternative cannot jointly offer: a sub-cubic-metre annual fuel charge with no convoy tail, a neutron output low enough to site near personnel behind a few metres of shield, and a black-start to grid-independent firm power in under a minute (figure 18). The islanded operating point ($+104\ \mathrm{MW_e}$) is modest but firm, and the availability (figure 19) is set by magnet demountability rather than by plasma physics. These are the properties that make the machine an Aegis rather than a laboratory burner, and they are why the product framing is forward-base defence power and not a merchant power plant.

THE CONTROL PROBLEM AND THE SAFETY ARCHITECTURE

The burner's binding physics — the plug potential, the warm fill, the burn stability — are all control quantities, and the machine is designed around a layered safety architecture rather than around a controller being correct. A learned or model-predictive controller dispatches the plug, the warm fill and the converter inside a control-barrier-function safe set that is forward-invariant by construction [47,70]: the clamp makes the burn safe-set forward-invariant and overrides any performance controller that would overshoot (gate AG-CX-29). An anomaly ensemble watches for the sub-threshold precursors of microstability onset and quench (gate AG-CX-33), and a graph-neural-network state estimator reconstructs the plug and thermal-barrier state from a topology of sensors so that a dropped channel is imputed rather than lost (gate AG-CX-30). A reinforcement-learning plug-potential controller rejects a $+30\ \mathrm{keV}$ potential disturbance in ≈ 59 steps to an RMS below $5\ \mathrm{keV}$ (gate AG-CX-28). The full control suite is the control companion [71], and the quantum-sensing and quantum-kinetic frontier is the burner-quantum companion [72].

STRUCTURAL AND RELIABILITY IMPLICATIONS

The variance budget and the RAMI analysis point at the same component: the large superconducting magnet set is both the availability-limiting single point of failure (§4.12) and the structural long pole (§4.13). Coil demountability, quench detection and protection, and a magnet-preserving default trip (which turns a 9-minute cold restart into a 35-second warm one) are therefore the top engineering priorities, and they are magnet-design problems with a clear owner in the REBCO companion [52,53], not open physics questions.

§ 8

Limitations

Five caveats bound the claims, each stated with its evidence class. First, every net-power number is contingent on retiring the plug-potential requirement $n_p/n_c = 16$, which a design study cannot do; only a plug-potential experiment can, and the Sobol budget shows this single gate carries $87\,\%$ of the gain variance. Second, the converter efficiency is a declared input on a ladder, and the conservative canon floor (0.49) sits below the net-power break-even (≈ 0.52), so the net-power claim is conditional on the converter clearing that break-even in hardware. Third, the loss-cone stability is a maintained DCLC threshold at $6.25\,\%$ warm fill, not a passive margin, and a residual anisotropy-driven AIC branch ($\gamma/\omega_{ci} \approx 0.037$) is not quenched by the warm fill. Fourth, the central-cell beta $\beta_c = 0.55$ enters the balance as an assumed input rather than an MHD-solved result, and it is the dominant interaction partner in the variance budget. Fifth, the helium-3 fuel supply is a separate, lunar-gated problem [68,18], and the synchrotron-model choice must remain the AFJ reflective-cavity model with $R_{\text{wall}} \gtrsim 0.83$, never the optimistic Trubnikov bound. None of these bears on the robust neutron-source and low-activation performance.

§ 9

Conclusion

The Aegis forward-base burner reaches an engineering gain $Q_E = 1.318$ at 4298.5 MW fusion power with a $5.44\,\%$ neutron fraction, delivering firm, islanded, grid-independent power ($+104$ $\mathrm{MW_e}$) with a black-start to grid-independence in $\approx 42\,\mathrm{s}$ and a fleet capacity factor of 68–$96\,\%$. The binding physics gate is the plug-potential requirement $n_p/n_c = \exp(e\phi_i/T_e) = 16$, retired only by a dedicated experiment and confirmed by the Sobol analysis as the first-order driver of the gain (carrying $87\,\%$ of its variance); the direct-converter efficiency is a second binding lever with a net-power break-even near 0.52; the loss-cone microstability is a maintained DCLC threshold at $6.25\,\%$ warm fill with a residual AIC branch; the synchrotron model must be the AFJ reflective-cavity treatment; and the plug coil qualifies structurally on as-built CAD. The neutron-source and low-activation performance — the properties that make a personnel-proximate, convoy-free forward-base installation possible — are robust and plug-independent; the net-power performance is conditional on the two levers, and the net-power uncertainty study ($P(\text{net} > 0) = 73\,\%$) shows honestly how near the cliff the machine lives. The design is presented with that separation carried openly, which is the only honest way to present an advanced-fuel mirror for a mission that will be fielded rather than merely published.

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 reviewers of the 2026 design series.

Reproducibility. Every headline quantity in this paper is a frozen anchor (BU-L1-A1, BU-L1-A10, BU-L2-A5, BU-L2-A7, BU-SX-02, KX-L1-A5-BU, AG-L1-A1, AG-DR-01, AG-DR-02, HX-01, HX-05, HX-07, P3-FEA-10) in the Kronos Physics De-Risking Register [48] (doi:10.5281/zenodo.22133057) and is regenerable from the deposited solver and the uniquely-named figure-generation scripts (`make_figs_aegis_burner.py`, `make_figs_aegis_extended.py`) archived with this paper at doi:10.5281/zenodo.22645693; official version: https://www.kronosfusionenergy.com/publications/aegis-burner.html. The Bosch–Hale reactivities of figures 6 and 7 are the standard 1992 parametrisation [16]. No result depends on unpublished or proprietary data, and no economic quantity is used or reported.

Scope — what this paper does not claim. Q_E is the engineering gain at the frozen design point, not a demonstrated result; the plug density requirement is stated as the binding open gate, not solved; the converter efficiency is a declared input, not a measured value; and low-neutron is *not* aneutronic. No economic quantity is reported here.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645693 (official version: https://www.kronosfusionenergy.com/publications/aegis-burner.html). Headline quantities are regenerable from deposited code and cross the Kronos 2026 series: the breeder platform [50], the sibling MetroVolt burner [51], the REBCO magnet basis and plug-coil engineering [52,53], direct energy conversion and its collector/architecture companions [54,55,56], the design-space mapping [57], the plug microstability [59] and RF alpha-channeling [58] companions, the ECH thermal-barrier, sloshing-ion and synchrotron papers [61,62,60], the central-cell transport and end-cell physics [63,64,65], the recirculation-closure and neutron-lean siting papers [66,67], the helium-3 supply analysis [68], the islanded forward-base operations study [69], the burn-control, control-suite and quantum-stack papers [70,71,72], the cross-design sensitivity/UQ study [74], the verification catalogue [76], and the red-team of the binding gates [77]. Any economic analysis is reported separately and is not included in the public deposit.

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

§ 10

Bosch–Hale reactivity evaluation

The Maxwellian-averaged reactivity (3) is evaluated with the Bosch–Hale rational approximation [16]. With T in keV, define

$$\theta = T\left[1 - \frac{T(C_2 + T(C_4 + TC_6))}{1 + T(C_3 + T(C_5 + TC_7))}\right]^{-1}, \qquad \xi = \left(\frac{B_G^2}{4\theta}\right)^{1/3},$$

so that

$$\langle \sigma v \rangle(T) = C_1\, \theta \sqrt{\frac{\xi}{m_r c^2\, T^3}}\, e^{-3\xi} \quad [\mathrm{cm^3\, s^{-1}}],$$

where B_G is the Gamow constant and $m_r c^2$ the reduced-mass energy. The coefficients used here are those of table 7, and the figures of this paper are generated directly from (13)–(14). For D– ^{3}He, $B_G = 68.7508\,\sqrt{\mathrm{keV}}$ and $m_r c^2 = 1\,124\,572\,\mathrm{keV}$; the parametrisation is valid to $\lesssim 0.3\,\%$ over 0.2–100 keV.

	D–T	D– ^{3}He	D–D (p)	D–D (n)
B_G	34.3827	68.7508	31.3970	31.3970
$m_r c^2$	1124656	1124572	937814	937814
C_1	1.17302e−9	5.51036e−10	5.65718e−12	5.43360e−12
C_2	1.51361e−2	6.41918e−3	3.41267e−3	5.85778e−3
C_3	7.51886e−2	−2.02896e−3	1.99167e−3	7.68222e−3
C_4	4.60643e−3	−1.91080e−5	0	0
C_5	1.35000e−2	1.35776e−4	1.05060e−5	−2.96400e−6
C_6	−1.06750e−4	0	0	0
C_7	1.36600e−5	0	0	0

Bosch–Hale coefficients used in equations (13)–(14) [16].

§ 11

The ambipolar Boltzmann relation

Consider the parallel (field-aligned) momentum balance of the plug electrons, which are magnetised and collisional along B. Neglecting electron inertia (justified by $m_e \ll m_i$) and friction on the slow timescale, the balance is

$$0 = -\nabla_\parallel (n_e T_e) - e\, n_e\, \nabla_\parallel \phi \cdot (-1),$$

where the electron charge is $-e$ so the electric force per unit volume is $(-e) n_e (-\nabla_\parallel \phi) = e n_e \nabla_\parallel \phi$. For isothermal electrons (T_e constant along B),

$$T_e\, \nabla_\parallel n_e = e\, n_e\, \nabla_\parallel \phi \implies \nabla_\parallel (\ln n_e) = \frac{e}{T_e} \nabla_\parallel \phi,$$

which integrates from the central cell (potential ϕ_c, density n_c) to the plug peak (potential $\phi_p = \phi_c +$ ϕ_i, density n_p) to give

$$\frac{n_p}{n_c} = \exp\left(\frac{e(\phi_p - \phi_c)}{T_e} \right) = \exp\left(\frac{e \qquad \phi_i}{T_e} \right),$$

which is equation (7). The relation is exact for isothermal, inertialess, magnetised electrons in steady state; the thermal barrier (§2.6) is what permits T_e in the plug to differ from the central-cell value and thereby makes the required n_p / n_c finite.

§ 12

Pastukhov end-loss derivation

The plugged ion loss is a Fokker–Planck problem in the electrostatic well: ions escape not over the magnetic mirror but by pitch-angle scattering into the loss cone against a potential barrier $e\phi_i \gg T_i$. Following Pastukhov [3] and Cohen's refinement [4], one solves the bounce-averaged kinetic equation with a loss-cone boundary condition and finds that the confinement time is set by the time to scatter an ion from the confined region across the barrier. The result, to leading order in $e\phi_i/T_i$, is

$$\tau_p \simeq \tau_{ii}\,\frac{\sqrt{\pi}}{2}\,g(R_m)\left(\frac{e\phi_i}{T_i}\right)\exp\left(\frac{e\phi_i}{T_i}\right), \qquad g(R_m)=\frac{\ln(2R_m+2)}{\ln 2R_m},$$

where τ_{ii} is the ion–ion Coulomb collision time. The factor xe^x with $x = e\phi_i/T_i$ dominates: the exponential is the Boltzmann suppression of the barrier and the linear prefactor is the phase-space volume of the near-loss-cone region. At the design point $x = 2.764$, so $xe^x = 43.8$, and with the calibrated $\tau_{ii}\,(\sqrt{\pi}/2)\,g(R_m) = 0.0851\,\mathrm{s}$ the model returns $\tau_p = 3.73\,\mathrm{s}$, the anchor value. The strong sensitivity $\mathrm{d}\ln\tau_p/\mathrm{d}\ln \phi_i = 1 + e\phi_i/T_i \approx 3.8$ at the design point is the quantitative source of the plug's criticality.

§ 13

Recirculation-limited gain algebra

Let $P_{\text{fus}} = P_{\text{ch}} + P_n$ be the fusion power, P_{ax} the axial charged stream reaching the converters, and P_{plug}, P_{aux} the recirculating plug and auxiliary powers. The gross electric output is the sum of the direct-converted charged stream and the thermally converted residual heat,

$$P_{\text{out}} = \eta_{\text{dec}}\, f_{\text{dir}}\, P_{\text{ax}} + \eta_{\text{th}}\, (P_{\text{fus}} + P_{\text{aux}} - f_{\text{dir}} P_{\text{ax}}),$$

and the engineering gain is the ratio of gross electric to recirculating power,

$$Q_E = \frac{P_{\text{out}}}{P_{\text{plug}} + P_{\text{aux}}} = \frac{\eta_{\text{th}}(P_{\text{fus}} + P_{\text{aux}}) + \eta_{\text{dec}}\, P_{\text{ch}}}{P_{\text{plug}} + P_{\text{aux}}},$$

where the last step collects the direct- and thermal-converted charged terms and uses $f_{\text{dir}} P_{\text{ax}} \simeq P_{\text{ch}}$ for the recovered charged fraction; this is equation (9). The net electric is $P_{\text{net}} = P_{\text{out}} - (P_{\text{plug}} + P_{\text{aux}})$, which is linear in η_{dec} and crosses zero at $\eta_{\text{dec}} \approx 0.52$ for the frozen power flows ($P_{\text{ch}} = 3868.0$ MW, $P_{\text{ax}} = 2719.7$ MW, $f_{\text{dir}} = 0.495$), the break-even of §4.7.

§ 14

DCLC dispersion and marginal stability

The drift-cyclotron-loss-cone mode is an electrostatic ion-cyclotron mode destabilised by the positive perpendicular-velocity slope of the loss-cone distribution and by the diamagnetic drift of a finite-radius plasma. In the local approximation the dispersion relation couples the ion-cyclotron harmonics to the drift frequency $\omega_* = k_y v_d$,

$$1 + \sum_n \frac{\omega_{pi}^2}{k^2} \frac{\Gamma_n(b)\,(\omega - \omega_*)}{(\omega - n\omega_{ci})} \frac{\partial \hat{f}}{\partial \varepsilon} = 0,$$

with $b = k_\perp^2 \rho_i^2$, Γ_n the modified-Bessel factor, and $\hat{f}$ the reduced distribution. The loss-cone imprints $\partial \hat{f}/\partial v_\perp > 0$ at small $v_\perp$, which drives $\mathrm{Im}\,\omega > 0$. Adding a warm, isotropic component of fractional density f_s fills the loss cone and removes the destabilising slope; the marginal condition $\gamma = 0$ is reached when the warm fraction is comparable to the inverse of the density contrast that sustains the drive, which for the burner plug evaluates to $f_s \simeq 1/(n_p/n_c) = 1/16 = 6.25\,\%$. This is the threshold of gate BU-L2-A5, and it is why the warm-fill and the plug ratio are the same number turned over. The anisotropy-driven Alfv\'en-ion-cyclotron branch has a different drive (the temperature anisotropy $T_\perp > T_\parallel$ of the sloshing population) and is not removed by the isotropic warm fill; PIC returns it flat at $\gamma/\omega_{ci} \approx 0.037$ over $f_s = 0$–$10\,\%$ (gate BU-L2-A7).

§ 15

Numerical schemes and convergence

The zero-dimensional balance is solved by fixed-point iteration on the coupled system of §2.14 to a relative tolerance of 10^{-8}, and the design point reproduces the register anchors to full deposited precision ($Q_E = 1.31812$). The Monte-Carlo and Sobol estimators use a Saltelli sequence [27,28]; the standard error of a proportion estimate falls as $\sqrt{p(1-p)/N}$ and the Sobol indices converge as $N^{-1/2}$ (figure 3), so the $40\,000$-sample base gives sub-percent index errors. The electromagnetic PIC runs (WarpX/SMILEI [38,39]) resolve the electron Debye length ($\Delta x \lesssim \lambda_{De}$) and the electron plasma period ($\omega_{pe}\Delta t < 2$), the Courant and finite-grid-instability conditions for an explicit scheme; the AIC growth rate is converged in cells-per-Debye-length to within the reported precision. The Fokker–Planck solver uses an implicit Chang–Cooper discretisation in (v, ξ) that preserves positivity and the Maxwellian null space.

§ 16

Parameter tables

Table 8 lists the design-point values and the ranges over which the Saltelli/Sobol scan of §5 is run; the ranges are those deposited in the register [48].

PARAMETER	SYMBOL	DESIGN POINT	SCAN RANGE
plug density ratio	n_p/n_c	16	12–20
central-cell beta	β_c	0.55	0.40–0.65
ion temperature	T_i	90	70–110
carrier temperature	T_e	89.7	80–100
mirror ratio	R_m	4.61	3.3–6.8
converter efficiency	η_{dec}	0.70	0.49–0.85
wall reflectivity	R_{wall}	0.90	0.80–0.95
directed fraction	f_{dir}	0.495	0.40–0.60

Design-point values and Saltelli scan ranges for the principal inputs (register ranges [48]).

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 T K Fowler and B G Logan, The tandem mirror reactor, *Comments Plasma Phys. Control. Fusion* **2**, 167 (1977).

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

3 V P Pastukhov, Collisional losses of electrons from an adiabatic trap in a plasma with a positive potential, *Nucl. Fusion* **14**, 3 (1974). doi:10.1088/0029-5515/14/1/001

4 D E Baldwin, End-loss processes from mirror machines, *Rev. Mod. Phys.* **49**, 317 (1977). doi:10.1103/RevModPhys.49.317

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

6 R F Post, The magnetic mirror approach to fusion, *Nucl. Fusion* **27**, 1579 (1987). doi:10.1088/0029-5515/27/10/001

7 R F Post and M N Rosenbluth, Electrostatic instabilities in finite mirror-confined plasmas, *Phys. Fluids* **9**, 730 (1966). doi:10.1063/1.1761740

8 M N Rosenbluth and C L Longmire, Stability of plasmas confined by magnetic fields, *Ann. Phys.* **1**, 120 (1957). doi:10.1016/0003-4916(57)90055-6

9 M Inutake *et al*, Thermal barrier formation and plasma confinement in the axisymmetrized tandem mirror GAMMA-10, *Phys. Rev. Lett.* **55**, 939 (1985). doi:10.1103/PhysRevLett.55.939

10 A A Ivanov and V V Prikhodko, Gas-dynamic trap: an overview of the concept and experimental results, *Plasma Phys. Control. Fusion* **55**, 063001 (2013). doi:10.1088/0741-3335/55/6/063001

11 P A Bagryansky *et al*, Threefold increase of the bulk electron temperature of plasma discharges in a magnetic mirror device, *Phys. Rev. Lett.* **114**, 205001 (2015). doi:10.1103/PhysRevLett.114.205001

12 T K Fowler, R W Moir and T C Simonen, A new simpler way to obtain high fusion power gain in tandem mirrors, *Nucl. Fusion* **57**, 056014 (2017). doi:10.1088/1741-4326/aa6939

13 T K Fowler, R W Moir and T C Simonen, Prospects for an axisymmetric advanced fuel tandem mirror, *AIP Conf. Proc.* **1771**, 080003 (2016). doi:10.1063/1.4964242

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

15 D D Ryutov, Open-ended traps, *Sov. Phys. Usp.* **31**, 300 (1988). doi:10.1070/PU1988v031n04ABEH005747

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

17 G L Kulcinski and J F Santarius, Advanced fuels under debate, *Nature* **396**, 724 (1998). doi:10.1038/25401

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

19 T H Rider, A general critique of inertial-electrostatic confinement fusion systems, *Phys. Plasmas* **2**, 1853 (1995). doi:10.1063/1.871273

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

21 W M Nevins, A review of confinement requirements for advanced fuels, *J. Fusion Energy* **17**, 25 (1998). doi:10.1023/A:1022513215080

22 R W Moir and W L Barr, Venetian-blind direct energy converter for fusion reactors, *Nucl. Fusion* **13**, 35 (1973). doi:10.1088/0029-5515/13/1/005

23 W L Barr and R W Moir, Test results on plasma direct converters, *Fusion Technol.* **3**, 98 (1983). doi:10.13182/FST83-A20820

24 N J Fisch and J M Rax, Interaction of energetic alpha particles with intense lower hybrid waves, *Phys. Rev. Lett.* **69**, 612 (1992).

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.

© 2026 Kronos Fusion Energy. All rights reserved. Hyperion, Aegis, and MetroVolt are product designations of Kronos Fusion Energy.

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 Aegis Forward-Base Burner: Q_E Closure, Plug Requirement, and Mirror Confinement in an Islanded Low-Neutron D–^{3}He Tandem Mirror

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

Zenodo · doi:10.5281/zenodo.22645693

Read on kronosfusionenergy.com

The Physics De-Risking Register

© 2026 KRONOS FUSION ENERGY · FAMILY-OWNED AMERICAN COMPANY