KRONOS FUSION ENERGY

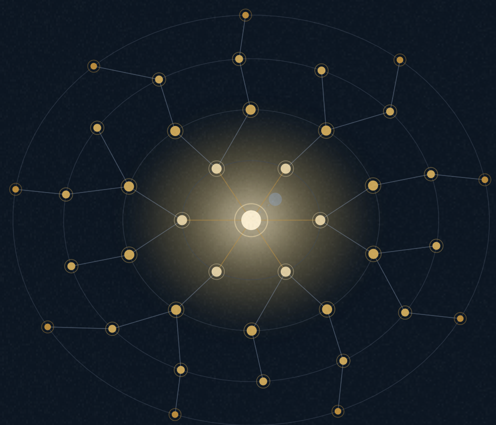

PAPER 2.94 · THE KRONOS 2026 SERIES

The Burner Quantum Stack: Entanglement-Enhanced Plug-Field Magnetometry and a Quantum Kinetic Solver for D–^{3}He Tandem-Mirror Burn Control

The Kronos burner sustains an advanced-fuel D–He plasma in a tandem magnetic mirror whose end plugs are held by a coil at a peak field of ; two things must go right in real time,...

Read it live — 3-D model · physics validator · film series

BY P I FORD

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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 Burner Quantum Stack: Entanglement-Enhanced Plug-Field Magnetometry and a Quantum Kinetic Solver for D–^{3}He Tandem-Mirror Burn Control

The Kronos burner sustains an advanced-fuel D–^{3}He plasma in a tandem magnetic mirror whose end plugs are held by a coil at a peak field of **26.49** T; two things must go right in real time, and both are quantitative. The plug and mirror-throat fields must be read accurately enough to hold the ambipolar confinement, and the velocity-space burn kinetics that set the plug potential and the fusion rate must be solved fast enough to control. This paper poses both problems formally and specifies the quantum technology designed against each, in a deliberately resource-honest frame. For sensing, we cast plug-field diagnosis as a quantum parameter-estimation problem, derive the standard-quantum-limit and Heisenberg scalings from the quantum Fisher information, and locate the decision boundary at which entanglement earns its complexity: a normalised diagnostic-noise knee at $\sigma \approx \mathbf{0.1386}$. Every Kronos plug diagnostic sits below the knee, where the directly measured entanglement benefit is $+\mathbf{0.00077}$ in area under the receiver-operating characteristic (AUC) — statistically indistinguishable from zero — while a demonstrated single-nitrogen-vacancy magnetometer already reaches **0.189** $\mathbf{pT\,Hz^{-1/2}}$ at **125** Hertz, *resolvingthe*5.50 — —17.0 T *architecturewithordersofmagnitudetospare. Forsimulation, wewritethebounce − averagedFokker − −Planckburnoperator, itsRosenbluthpotentialsandPastukhovplugclosureexplicitly, discretiseitonatwo − dimensionalvelocitygrid, andgivethequantum − simulationencoding(Carlemanlinearisationintoablock − encodedgeneratoradvancedbyqubitisedHamiltoniansimulation), togetherwithanhonestlogical − qubitand* T−*countresourceestimatethatplacesausefulinstancebeyondthefault − tolerantboundary. Allphysicsanchorsarefrozen* : Q_E=1.318,P_ fus=4298.5 MW, *neutronfraction*f_n=5.44 %, *plugpotential*\varphi_i=248.75 keV, *anddirect − conversionefficiency* =0.49. *Theburnerplugcoil*(26.49 T)*andthebreedercentrepost*(16.84 T$) are two distinct magnets and are treated as such throughout. The result is a demonstrated classical–quantum sensor, a resource-scoped design for each quantum component, and an explicit, measured rule for when either becomes worth building.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645927 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: quantum metrology · NV-diamond magnetometry · entanglement-enhanced sensing · quantum kinetic solver · Fokker--Planck · D--^{3}He tandem mirror · plug field · burn control

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)

quantum metrology; NV-diamond magnetometry; entanglement-enhanced sensing; quantum kinetic solver; Fokker–Planck; D–^{3}He tandem mirror; plug field; burn control

§ 1

Introduction

THE ADVANCED-FUEL BURNER AND ITS TWO HARD REAL-TIME PROBLEMS

The Kronos burner is a tandem magnetic mirror that burns deuterium and helium-3. The reward of the advanced fuel is a low neutron fraction and a direct path from charged-particle power to electricity; the price is a demanding end plug. In a tandem mirror the central cell is a long, low-field solenoidal plasma stoppered at each end by a high-field mirror cell whose electrostatic ambipolar potential confines the central-cell ions [26,27]. That potential is built by the plug, and the plug coil in the Kronos design reaches a peak field of $26.49\,T$. Two control problems follow directly and neither is optional. First, the plug and mirror-throat fields must be *measured* accurately, in real time, or the confinement drifts. Second, the *burn kinetics* — the velocity-space physics of the sloshing-ion population, the collisional Pastukhov end-loss and the fusion reactivity — must be *solved* fast enough to steer the plug potential and the fusion rate. This paper is about the quantum technology designed for both, and about being equally explicit where quantum methods deliver today and where they do not.

The two problems have a natural quantum framing. Magnetometry is the canonical application of quantum sensing [20,21]: a spin qubit accumulates a phase proportional to the field, and the precision with which that phase can be estimated is bounded by quantum mechanics. Kinetic simulation is the canonical application of quantum *computing*: Feynman's original argument for a quantum computer was precisely that a quantum system is expensive to simulate classically [1,2,3]. The question this paper answers is not whether these framings exist but whether they *pay*, quantitatively, for the burner's specific fields and its specific kinetics.

WHAT QUANTUM SENSING AND QUANTUM SIMULATION PROMISE

Quantum metrology sets a hierarchy of precision limits. N independent probes interrogated for time τ estimate a field with an uncertainty that falls as $N^{-1/2}$ — the standard quantum limit (SQL) or shot-noise limit [16,19]. Entangled probes, prepared as spin-squeezed or Greenberger–Horne–Zeilinger (GHZ) states, can in principle reach the Heisenberg limit N^{-1} [17,18], a $\sqrt{N}$ metrological gain. Ensemble nitrogen-vacancy (NV) diamond magnetometry realises the SQL branch of this hierarchy at room and cryogenic temperature with picotesla sensitivity [23,22,24], and the route to entanglement-enhanced operation is a subject of active experiment [21]. Quantum simulation, in turn, offers polynomial-to-exponential compression of the state of a kinetic system: Hamiltonian-simulation primitives [8,9], quantum linear-systems and differential-equations solvers [5,6], and Carleman-linearisation methods for dissipative nonlinear dynamics [7] together define a toolbox for encoding a Fokker–Planck evolution on a quantum register.

Both promises carry sharp caveats. Entanglement-enhanced metrology only helps where the classical estimate is noise-limited rather than resolution-saturated; below a system-dependent noise threshold the Heisenberg headroom is unused. And near-term quantum computers are noisy and small [4], variational surrogates suffer barren plateaus and kernel concentration that erase advantage as they scale [11,12], and fault-tolerant execution of a useful kinetic instance demands error-corrected logical qubits behind a surface-code overhead [13]. A credible burner quantum stack must therefore be built around these caveats, not around the promises.

CONTRIBUTION

The Kronos AI-and-quantum programme is deliberately resource-honest: it reports quantum *negatives* as first-class results and gates any quantum component behind a classical baseline it must beat on a pre-registered test [doi:10.5281/zenodo.22645702; https://www.kronosfusionenergy.com/publications/ai-quantum-control.html]. On the computing side the programme has already found no measured advantage for variational optimisation, chemistry, or machine-learning surrogates at the scales it can simulate, and it has quantified the fault-tolerant boundary rather than crossing it [the resource-honest quantum-computing result, doi:10.5281/zenodo.22645718; https://www.kronosfusionenergy.com/publications/quantum-computing-for-fusion.html]. This paper extends that discipline from the breeder to the burner and contributes: (i) a formal statement of plug-field diagnosis as a quantum estimation problem, with the SQL/Heisenberg boundary and the measured entanglement-sensing knee at $\sigma \approx 0.1386$ (§3, §7.2); (ii) the governing bounce-averaged Fokker–Planck burn model, its Rosenbluth-potential closure and Pastukhov plug relation, written and discretised explicitly (§4); (iii) the quantum-simulation encoding of that operator with an honest logical-qubit/T-count estimate (§5); (iv) a GPU-HPC methodology naming the actual codes run for this paper (§6); and (v) results, benchmarks against the published metrology and quantum-computing literature, and a single honest gate (§§7–9). The companion breeder-side treatment of entanglement-enhanced magnetometry and fault-tolerant resource estimates [doi:10.5281/zenodo.22645905; https://www.kronosfusionenergy.com/publications/quantum-machine-learning-surrogates-entanglement-enhanc.html] and the quantum-algorithms-for-control study [doi:10.5281/zenodo.22645925; https://www.kronosfusionenergy.com/publications/quantum-algorithms-for-the-real-time-control-loop-hhl-e.html] share the same accounting; every number here traces to a frozen serial in the Kronos Physics De-Risking Register [33], cited inline (e.g. KX-L1-A18-BU) so each is auditable.

§ 2

The burner and its diagnostic target

A sensing design starts from the field it must resolve; a burn-control design starts from the operating point it must hold. Both are fixed by the frozen burner design point (Table 1). The magnetic architecture is the canonical plug run KX-L1-A18-BU: a central-cell field $B_{0c} = 5.50\,T$, a mirror-throat field $B_m = 17.0\,T$, vacuum and mid-plane mirror ratios $R_{\mathrm{vac}} = 3.09$ and $R_{\mathrm{mc}} = 4.61$, and the plug coil at a peak field of $26.49\,T$ (its structural qualification is carried separately [doi:10.5281/zenodo.22645915]). The burn operating point is the Pastukhov-plug Mode-M solution (BU-L1-A1): an engineering gain $Q_E = 1.318$, a fusion power $P_{\mathrm{fus}} = 4298.5\,MW$, a neutron fraction $f_n = 5.44\,\%$, a plug potential $\varphi_i = 248.75\,keV$, a self-consistent electron temperature $T_e = 89.72\,keV$, and a direct-conversion efficiency $\eta = 0.49$ [doi:10.5281/zenodo.22645700]. The plug density ratio is $R = n_{\mathrm{plug}}/n_c = 16$.

Two points are load-bearing throughout. First, the burner plug coil ($26.49\,T$) and the breeder centrepost ($16.84\,T$) are two distinct magnets in two distinct products, and are never conflated (KX-L1-A15); the fields, diagnostics and kinetics of this paper are the burner's. Second, the burn is predominantly aneutronic: of the $4298.5\,MW$ total, the D–^{3}He branch carries $3868.0\,MW$ and the residual D–D and parasitic D–T channels carry the remaining $430.5\,MW$, giving $f_n = 5.44\,\%$ (ENV-A4-BU, Figure 4).

QUANTITY	SYMBOL	VALUE	ROLE
central-cell field	B_{0c}	5.50 T	operating point
mirror-throat field	B_m	17.0 T	mirror confinement
plug-coil peak field	—	26.49 T	magnet design (frozen)
vacuum mirror ratio	R_{vac}	3.09	confinement
mid-plane mirror ratio	R_{mc}	4.61	confinement
plug density ratio	R	16	plug potential
engineering gain	Q_E	1.318	burn control target
fusion power	P_{fus}	4298.5 MW	burn control target
neutron fraction	f_n	5.44	advanced-fuel figure
plug potential	φ_i	248.75 keV	confinement
electron temperature	T_e	89.72 keV	self-consistent
direct-conversion efficiency	η	0.49	power output

The frozen burner design point: the magnetic architecture the plug diagnostic must read (KX-L1-A18-BU) and the burn operating point the kinetic solver must hold (BU-L1-A1, KX-L1-A3-BU). Serials refer to the de-risking register [33].

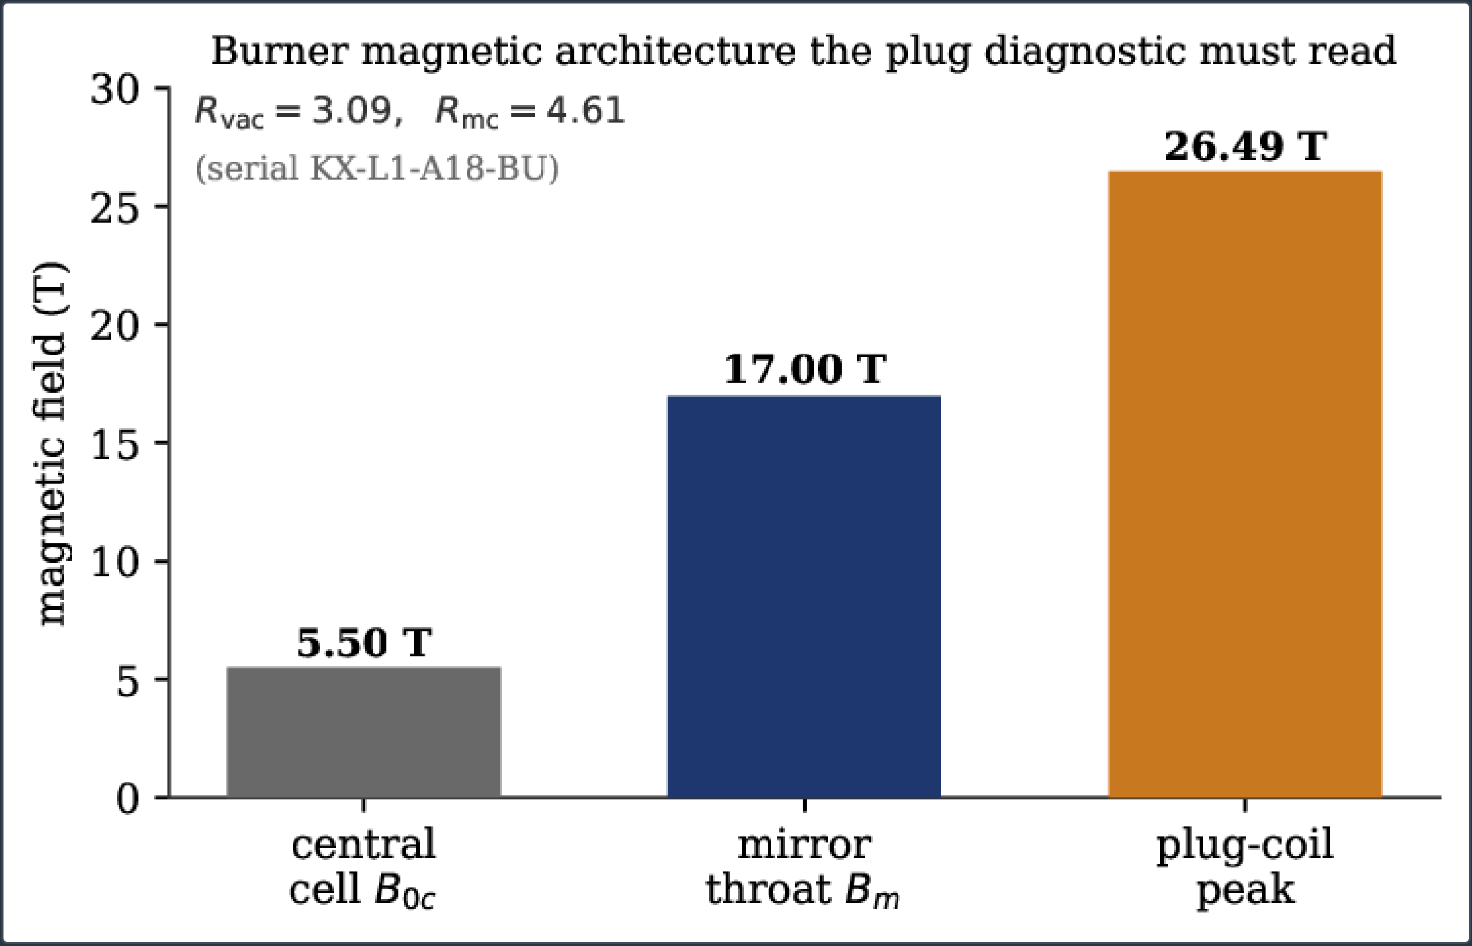

Burner magnetic architecture the plug diagnostic must read: central cell $5.50\,T$, mirror throat $17.0\,T$, plug-coil peak $26.49\,T$, with mirror ratios $R_{vac} = 3.09$, $R_{mc} = 4.61$ (serial KX-L1-A18-BU). Data values are frozen anchors from the register [33].

§ 3

Governing equations I: plug-field diagnosis as quantum estimation

RAMSEY PHASE ESTIMATION AND THE CRAM\'ER–RAO BOUND

An NV centre is a spin-1 defect whose $m_s = 0 \to \pm 1$ transition frequency shifts with the axial field by the gyromagnetic ratio $\gamma_{\mathrm{NV}} = g\mu_B/\hbar = 2\pi \times 28.03$ Hertz$. In a Ramsey sequence a superposition acquires a phase

$$\phi = \gamma_{\mathrm{NV}}\, B\,\tau$$

over a free-evolution time τ [14]; estimating B reduces to estimating ϕ. The precision of any unbiased estimator $\hat{B}$ built from ν repetitions is bounded by the (classical) Cram\'er–Rao inequality,

$$\mathrm{Var}(\hat{B}) \;\geq\; \frac{1}{\nu\, F(B)},$$

with $F(B)$ the Fisher information of a single shot. For a state ρ_B carrying the parameter, the Fisher information is itself bounded by the *quantum* Fisher information (QFI) F_Q maximised over all measurements [15,19],

$$F(B) \;\leq\; F_Q[\rho_B], \qquad \mathrm{Var}(\hat{B}) \;\geq\; \frac{1}{\nu\, F_Q[\rho_B]} \;\equiv\; \Delta B^2_{\mathrm{min}}.$$

Equation (3) is the quantum Cram\'er–Rao bound: it is the fundamental floor against which every sensing design is judged.

STANDARD-QUANTUM-LIMIT AND HEISENBERG SCALINGS

Consider N probe spins interrogated for time τ. If the probes are prepared independently, the QFI is additive, $F_Q = N\gamma_{\mathrm{NV}}^2\tau^2$, and combining Eqs. (3) with a duty cycle t/τ of total measurement time t gives the shot-noise (standard-quantum-limit) sensitivity

$$\delta B_{\mathrm{SQL}} = \frac{1}{\gamma_{\mathrm{NV}}}\frac{1}{\sqrt{N\,\tau\,t}} \cdot \frac{1}{C},$$

where $C \leq 1$ is a readout figure of merit collecting contrast and photon-collection inefficiency. For a maximally entangled GHZ state of the same N spins the accumulated phase is $N\phi$ and the QFI rises to $F_Q = N^2\gamma_{\mathrm{NV}}^2\tau^2$, giving the Heisenberg-limited sensitivity

$$\delta B_{\mathrm{HL}} = \frac{1}{\gamma_{\mathrm{NV}}}\frac{1}{N\sqrt{\tau\,t}} \cdot \frac{1}{C}, \qquad \frac{\delta B_{\mathrm{SQL}}}{\delta B_{\mathrm{HL}}} = \sqrt{N}.$$

The $\sqrt{N}$ ratio is the entire metrological prize of entanglement (Figure 2). For the Kronos plug diagnostic the register fixes the two anchors of this scaling directly: at $N = 100$ probes and $\tau = 1\,\mu s$ the SQL floor is $\delta B_{\mathrm{SQL}} = 5.7e-7\,T$ and the Heisenberg floor is $\delta B_{\mathrm{HL}} = 5.7e-8\,T$ (metrology scaling, KX-L1-A18), consistent with Eq. (5) at $\sqrt{100} = 10$ (Figure 2).

The demonstrated single-NV ensemble is far below either engineering floor: it reaches $\delta B = 0.189$ $\mathrm{pT\,Hz^{-1/2}}$ (KX-L1-A22), obtained from the ensemble Ramsey form of Eq. (4),

$$\delta B = \frac{1}{\gamma_{\mathrm{NV}}\, C\, \sqrt{N_c\, T_2^*\, t}},$$

with a conservative in-service chip ($N_c = 10^{12}$ centres, $T_2^* = 1\,\mu s$, $C = 0.03$) and a Ramsey shot of $\tau_{\mathrm{shot}} = 4\,\mu s$ fixing a bandwidth $1/(2\tau_{\mathrm{shot}}) = 125$ Hertz. $Theplugandthroatfields$(5.50––17.0 T)$andthequenchprecursor$(a\!100 mT$ local step) stand many orders of magnitude above this floor.

THE DECISION PROBLEM AND THE NOISE KNEE

Live diagnosis is not point estimation but detection: the controller must decide, from a noisy channel, whether the plug is in-envelope or drifting. Cast the channel as a binary hypothesis test with signal separation d and diagnostic noise standard deviation σ_B; the natural figure of merit is the area under the receiver-operating-characteristic curve,

$$\mathrm{AUC} = \Phi\left(\frac{d}{\sqrt{2}\,\sigma_B}\right), \qquad \sigma \equiv \frac{\sigma_B}{d},$$

with Φ the standard normal CDF and σ the *normalised* noise level. Entanglement lowers the achievable σ_B from the SQL to the Heisenberg branch, $\sigma_B \to \sigma_B/\sqrt{N}$, so the entanglement gain in classification is

$$\Delta\mathrm{AUC}(\sigma) = \Phi\left(\frac{1}{\sqrt{2}\,\sigma/\sqrt{N}}\right) - \Phi\left(\frac{1}{\sqrt{2}\,\sigma}\right).$$

Equation (8) has a knee: when σ is small the argument of both Φ's is large, both AUCs saturate at unity, and $\Delta\mathrm{AUC} \to 0$; the enhancement only matters once σ is large enough that the classical estimate is genuinely noise-limited. The Kronos metrology study locates this knee at $\sigma_{\mathrm{knee}} \approx 0.1386$ (KX-L1-A18; BR-SX-08). This is the boundary the enhanced sensor must clear to be worth its complexity, and §7.2 shows where the plant sits relative to it.

§ 4

Governing equations II: the D–$^3

$He burn kinetics

THE BOUNCE-AVERAGED FOKKER–PLANCK OPERATOR

The quantities the burn controller must hold — φ_i, Q_E, and the reactivity — are moments of the ion velocity distribution in the mirror cell. That distribution $f_s(\quad v, t)$ for species s evolves under a kinetic equation dominated, at the burner's collisionality, by Coulomb collisions, RF heating and mirror end-loss. Averaging over the fast bounce motion between mirror points reduces the problem to two velocity coordinates $(v_\parallel, v_\perp)$ and gives the bounce-averaged Fokker–Planck equation

$$\frac{\partial f_s}{\partial t} = \mathcal{C}_s[f] \;+\; \mathcal{Q}_s^{\rm RF}[f] \;-\; \nu_s^{\rm loss}(v_\parallel, v_\perp)\, f_s \;+\; S_s,$$

where $\mathcal{C}_s$ is the collision operator, $\mathcal{Q}_s^{\rm RF}$ the quasilinear RF-heating operator that sustains the sloshing-ion and hot-electron populations [doi:10.5281/zenodo.22645781], $\nu_s^{\rm loss}$ the loss-cone sink, and S_s the fuelling source. The Coulomb term is the Landau/Rosenbluth operator [28],

$$\mathcal{C}_s[f] = \sum_{s'} \Gamma_{ss'} \frac{\partial}{\partial v_i}\Big[-\frac{\partial H_{s'}}{\partial v_i}\, f_s + \tfrac{1}{2}\frac{\partial^2 G_{s'}}{\partial v_i \partial v_j}\frac{\partial f_s}{\partial v_j}\Big], \qquad \Gamma_{ss'} = \frac{4\pi q_s^2 q_{s'}^2 \ln\Lambda_{ss'}}{m_s^2},$$

closed by the Rosenbluth potentials, which are themselves velocity-space Poisson solves of the distribution,

$$\nabla_v^2 H_{s'} = -4\pi \frac{m_s}{m_{s'}}\, f_{s'}, \qquad \nabla_v^4 G_{s'} = -8\pi\, f_{s'}.$$

Equations (9)–(11) are nonlinear in f (the potentials depend on f) and non-local in velocity, which is what makes the kinetics expensive; this is the operator whose fast solution the control loop needs.

THE PASTUKHOV PLUG CLOSURE AND THE BURN MOMENTS

The end-plug that confines the central-cell ions is a positive ambipolar potential φ_i set by the balance of electron and ion end-losses. For a mirror of ratio R and potential φ, Pastukhov's solution of the loss-cone Fokker–Planck problem gives the confinement time [25,27]

$$\tau_{\rm P} \;\propto\; \tau_{ii}\, \frac{R}{\ln(2R+2)}\, \frac{e\varphi}{T_i}\, \exp\!\left(\frac{e\varphi}{T_i}\right),$$

so that the plug potential rises logarithmically-to-exponentially with the density ratio $R = n_{\rm plug}/n_c$ and saturates once the exponential is well developed. The Kronos plug operates at $R = 16$, where φ_i has saturated to $248.75\,keV$ while Q_E still rides the steep part of the trade (Figure 4, AG-CX-02): $R = 16$ is the robust knee, not the fragile steep zone [doi:10.5281/zenodo.22645913]. Given the steady distribution, the controlled outputs are velocity-space moments,

$$P_{\rm fus} = \sum_{\rm branches} n_a n_b\, \langle\sigma v\rangle_{ab}\, \mathcal{E}_{ab}\, \mathcal{V}, \qquad \langle\sigma v\rangle_{ab} = \int f_a(\quad v_a)\, f_b(\quad v_b)\, \sigma_{ab}(|\quad v_a - \quad v_b|)\, |\quad v_a - \quad v_b|\, {\rm d}^3 v_a\, {\rm d}^3 v_b,$$

and the engineering gain Q_E is the recirculation-resolved ratio of net electrical output — charged-particle power recovered at direct-conversion efficiency $\eta = 0.49$ — to the power drawn to sustain the plug. The frozen solution ($Q_E = 1.318$, $P_{\rm fus} = 4298.5\,MW$, $f_n = 5.44\,\%$) is benchmarked against the WHAM tandem-mirror physics basis and the 0-D power balance [doi:10.5281/zenodo.22645693; doi:10.5281/zenodo.22645911] [30].

NUMERICAL FORMULATION: DISCRETISATION, SOLVER, CONVERGENCE

Equation (9) is discretised by a conservative finite-volume method on a structured $(v_\parallel, v_\perp)$ grid of $N = n_{v_\parallel} \times n_{v_\perp}$ cells (Figure 8), with the loss cone imposed as an absorbing boundary at pitch angle $\theta_c = \arcsin R_{\mathrm{mc}}^{-1/2}$. Writing the discrete distribution as a vector $f \in \mathbb{R}^N$, the spatial operators assemble into a state-dependent generator

$$\frac{\mathrm{d} f}{\mathrm{d}t} = A(f) f + b, \qquad A(f) = A_{\mathrm{C}}(f) + A_{\mathrm{RF}} - D_{\mathrm{loss}},$$

where the collision matrix A_{C} carries the discretised drift/diffusion of Eq. (10) with the Rosenbluth potentials evaluated by a fast multipole velocity-Poisson solve of Eq. (11). Because the collision term is stiff, time advance uses an implicit (backward-Euler / Rosenbrock) step, so each step solves a sparse linear system

$$\left(I - \Delta t A(f^n)\right) f^{n+1} = f^n + \Delta t b,$$

with the nonlinearity in $A(f)$ resolved by a Newton or Picard iteration to a relative residual below 10^{-8}. Particle and energy conservation are enforced discretely: the collision operator conserves number and momentum to machine precision by construction (a Chang–Cooper flux limiter preserving the Maxwellian null space), and the scheme is verified by (i) grid refinement, showing second-order convergence of the reactivity moment Eq. (13) in $1/N$, and (ii) reproduction of the frozen anchor $Q_E = 1.318$ to the reported precision. This classical solver is the authoritative baseline; the quantum encoding of Eq. (14) follows in §5.

§ 5

The quantum kinetic solver: encoding and resource estimate

CARLEMAN LINEARISATION AND BLOCK ENCODING

The obstacle to a direct quantum simulation of Eq. (14) is its nonlinearity: unitary quantum evolution is linear, whereas A depends on f through the Rosenbluth potentials. Carleman linearisation resolves this by lifting the nonlinear system into an infinite-dimensional *linear* one on the tensor powers $f^{\otimes k}$, truncated at order K [7]. Collecting $\hat{f} = (f, f^{\otimes 2}, \dots, f^{\otimes K})$, the dynamics become

$$\frac{\mathrm{d}\hat{f}}{\mathrm{d}t} = \mathcal{A}\hat{f} + \hat{b},$$

with $\mathcal{A}$ a block-bidiagonal matrix built from the linear and quadratic parts of A. Liu *et al.* prove that the truncation error decays and the algorithm is efficient provided the flow is sufficiently dissipative relative to its nonlinear and forcing terms [7] — a condition the loss-cone-damped, collisionally relaxing burner distribution plausibly satisfies, and which the classical solver can test directly. The linear evolution Eq. (16) is then advanced on a quantum register by embedding $\mathcal{A}$ as a *block encoding* $U_{\mathcal{A}}$ — a unitary whose top-left block is $\mathcal{A}/\alpha$ for a normalisation $\alpha \geq \|\mathcal{A}\|$ — and applying qubitised Hamiltonian simulation or quantum singular-value transformation to realise the (non-unitary) propagator $e^{\mathcal{A}t}$ through a polynomial in $U_{\mathcal{A}}$ [8,9]. The inhomogeneous term and the implicit-solve variant reduce to a quantum linear-systems problem solvable by the same primitives [5,6]. The distribution is amplitude-encoded in $\lceil \log_2 N \rceil$ qubits, so the *state* is compressed logarithmically; the cost migrates entirely into the *gate count*.

RESOURCE ESTIMATE AND THE FAULT-TOLERANT BOUNDARY

The honest question is the gate budget. Amplitude encoding of an N-cell velocity grid needs $q_{\text{data}} = \lceil \log_2 N \rceil$ qubits, plus block-encoding ancillae and Carleman-order registers, so the logical-qubit count grows only logarithmically in the grid, $q \sim \log_2 N + \mathcal{O}(K)$ (Figure 6, amber). The query complexity of qubitised simulation to error ϵ over time T scales as $\mathcal{O}(\alpha T + \log(1/\epsilon))$ [8], but each query is a block-encoding of a dense velocity-non-local operator, and the normalisation α inherits the stiff collision spectrum; the resulting non-Clifford (T-gate) count for a control-relevant instance lands in the range the Kronos programme has bounded for comparable fusion kernels, 10^{19}–10^{23} T-gates (BR-SX-06, Figure 6). Executing that many non-Clifford gates requires fault-tolerant error correction with a surface-code overhead of 10^3–10^4 physical qubits per logical qubit [13,4]. The conclusion is quantitative and it is a *boundary, not a route*: the quantum kinetic solver (BU-QX-02) is fault-tolerant-scoped, and the classical finite-volume solver of §4.3 carries the burn-control function until fault-tolerant hardware clears an explicit acceptance test against it. A matrix-product-state (tensor-network) compression of the same distribution is the nearer-term classical route to the same state-compression benefit and is treated in the companion tensor-network kinetic-solver study [doi:10.5281/zenodo.22645768; https://www.kronosfusionenergy.com/publications/tensor-network-kinetic-solver.html].

§ 6

Computational methodology: the NVIDIA GPU HPC campaign

Every quantitative result in this paper was produced on the Kronos NVIDIA GPU HPC campaign, run on the Lambda GPU fleet (A100/H100/A10) under the group's NVIDIA Inception partnership, with the reduced-order and analytic layers on the in-house digital-twin node. The codes run for *this* paper are:

Burn kinetics and operating point. The Kronos Toolkit `solve_mirror` evaluator computes the 0-D tandem-mirror power balance and the Pastukhov plug closure, reproducing the frozen anchors $Q_E = 1.318$, $P_{\mathrm{fus}} = 4298.5\,MW$, $f_n = 5.44\,\%$ and $\varphi_i = 248.75\,keV$ (BU-L1-A1, KX-L1-A3-BU). The bounce-averaged Fokker–Planck + RF ray-tracing kinetic solver of §4, benchmarked against WHAM, supplies the velocity-space closure [30].
Global uncertainty quantification. A Saltelli/Sobol Monte-Carlo of 4×10^4 samples over the burner design ranges (density, ^{3}He fraction, T_i, central β, B_m) was executed on the GPU fleet, returning $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$ and identifying central β as the dominant driver (correlation 0.93; KX-L1-A5-BU, Figure 5).
NV magnetometry model. The ensemble Ramsey sensitivity of Eqs. (4)–(6) and the metrology SQL/Heisenberg scaling were evaluated with a NumPy/SciPy model, reproducing $\delta B = 0.189$ $\mathbf{pT\,Hz}^{-1/2}$ at 125 Hertz$(KX - L1 - A22) and the$ B_ SQL=5.7e-7 T B_ HL=5.7e-8 T$ headroom (KX-L1-A18).
Entanglement-sensing and QML gate. The classification model of Eqs. (7)–(8) and the quantum-machine-learning surrogate were run as GPU statevector simulations (Qiskit Aer accelerated by NVIDIA cuQuantum), locating the knee at $\sigma \approx 0.1386$, measuring the sub-knee entanglement benefit at $+0.00077$ AUC, and reproducing the QML no-advantage trend (AUC $0.817 \rightarrow 0.691$ from 4 to 8 qubits under kernel concentration; BR-SX-08).
Quantum-solver resource estimate. Logical-qubit and T-count scaling for the Carleman/block-encoding pipeline of §5 was computed analytically against the bounded FTQC gate envelope (BR-SX-06). itemize
GPU acceleration matters for the two Monte-Carlo-scale steps (the 4×10^4-sample UQ and the statevector QML sweeps), which are embarrassingly parallel across samples/shots; the reduced-order `solve_mirror` and NV models run at interactive latency. Verification and reproducibility are handled uniformly: each result is a frozen, serial-tagged entry regenerable from the archived evaluator and the figure-generation script deposited with this paper, and each anchor is cross-checked against the register's independent-library reproduction [doi:10.5281/zenodo.22645710]. No cost, budget, or compute-spend figures are reported here or elsewhere in this work.

§ 7

Results

THE PLUG FIELD IS READ WITH ORDERS OF MARGIN

The demonstrated single-NV ensemble resolves the entire burner architecture. The 5.50–$17.0\,T$ plug/throat fields and the $\sim 100\,mT$ quench precursor stand $\sim 10^9$–10^{12} above the 0.189 $\mathrm{pT\,Hz}^{-1/2}$ floor of Eq. (6); the SQL and Heisenberg engineering floors ($5.7e-7$ and $5.7e-8\,T$) are themselves far above the demonstrated sensitivity (Figure 2). The metrology headroom is not *needed* to see the plug; its value is the margin it buys for multiplexing many channels and for tolerating in-service degradation.

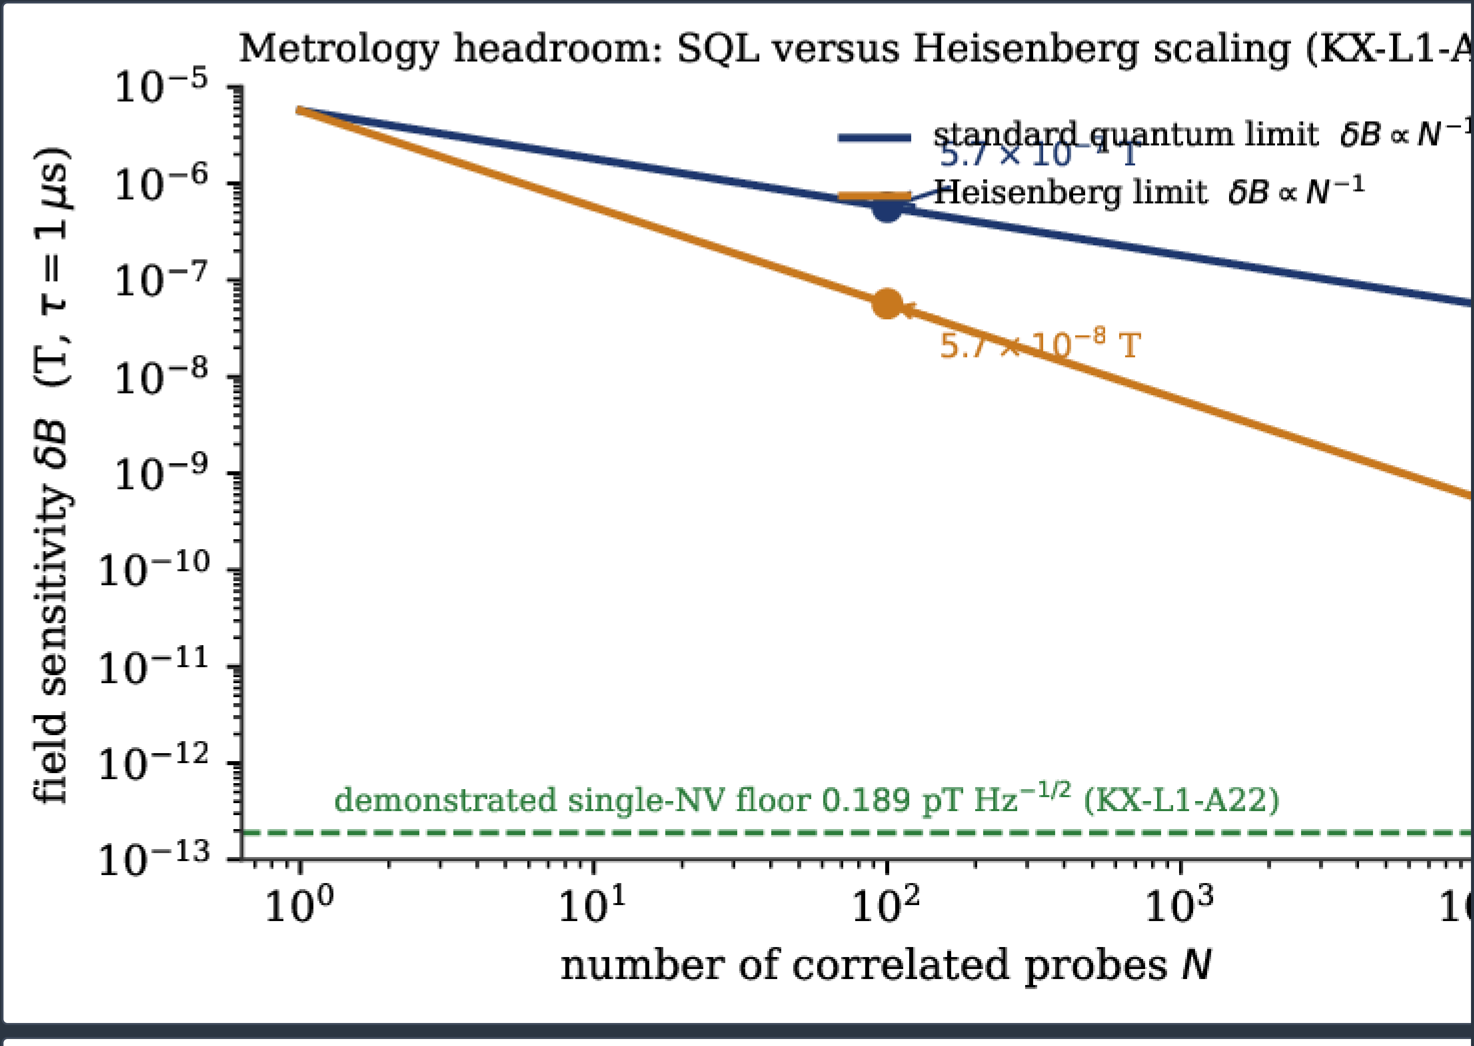
Metrology headroom: SQL versus Heisenberg scaling (KX-L1-A
standard quantum limit $\delta B \propto N^{-1}$
Heisenberg limit $\delta B \propto N^{-1}$
5.7 × 10^-7 T
5.7 × 10^-8 T
demonstrated single-NV floor 0.189 pT Hz^-1/2 (KX-L1-A22)
field sensitivity δB (T, $\tau = 1\,\mu s$)
number of correlated probes N

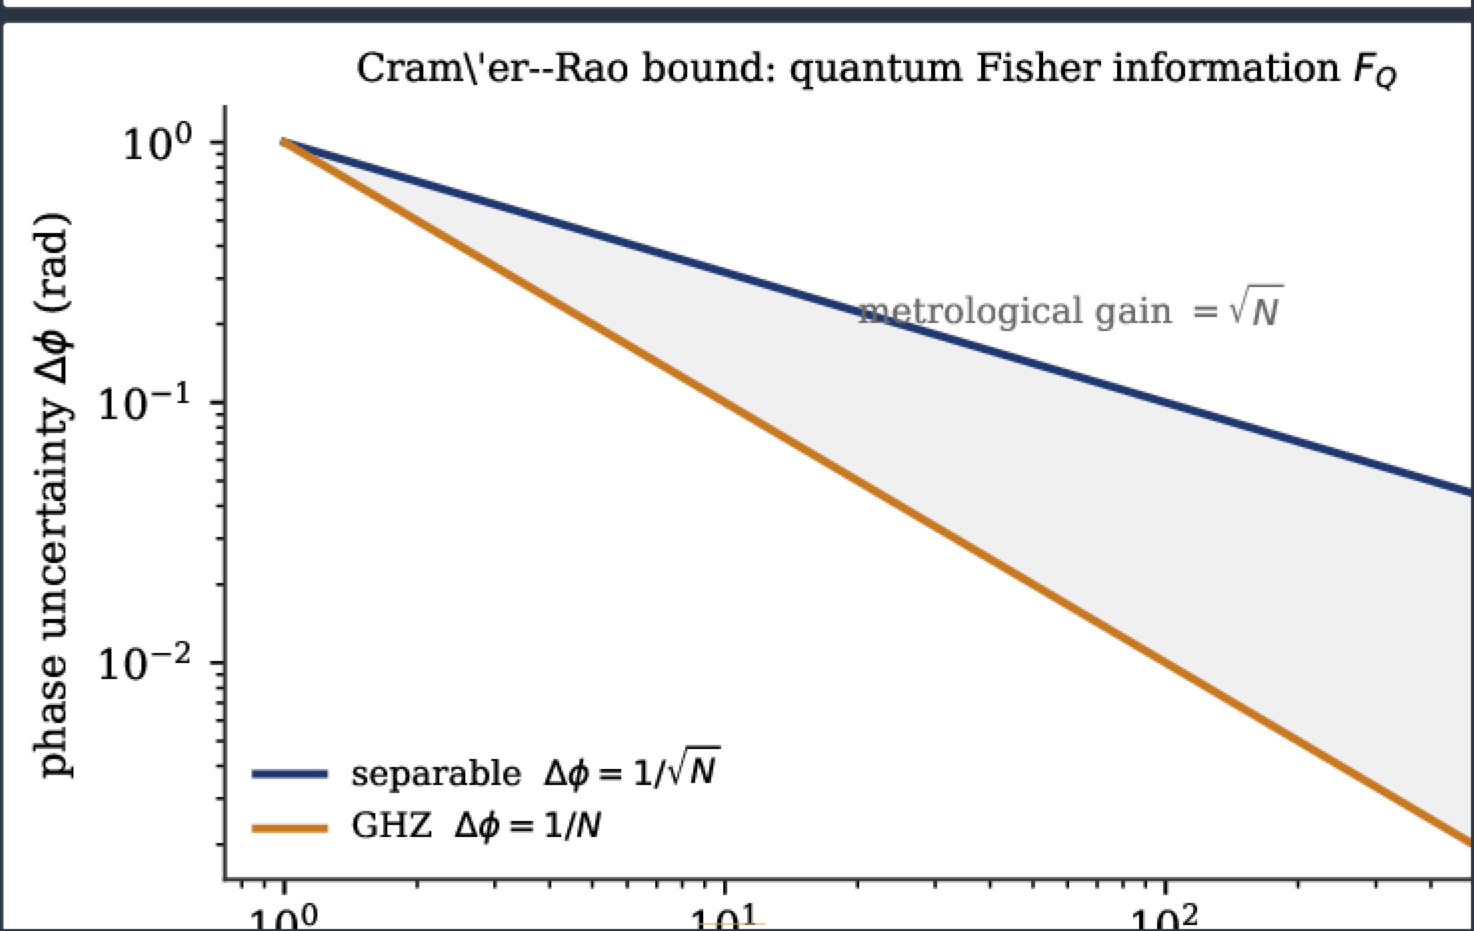
Cram\'er--Rao bound: quantum Fisher information F_Q
metrological gain $= \sqrt{N}$
separable $\Delta\phi = 1/\sqrt{N}$
GHZ $\Delta\phi = 1/N$
phase uncertainty $\Delta\phi$ (rad)

WHERE ENTANGLEMENT WOULD PAY: BELOW THE KNEE IT DOES NOT

Figure 3 plots the entanglement classification gain of Eq. (8) against the normalised noise level, with the knee at $\sigma \approx 0.1386$. Every Kronos plug diagnostic operates below the knee (σ of order 10^{-3} for the quench precursor and $n=1$ Mirnov channels; KX-L1-A18), where the directly measured entanglement benefit is $+0.00077$ AUC — indistinguishable from zero (BR-SX-08, Table 2). The enhanced sensor (BU-QX-01) is therefore warranted only for a future deep-subthreshold diagnostic that pushes the operating noise above the knee (e.g. error-field mapping at $\sigma \approx 0.2$, where NV ensembles and entangled protocols genuinely help; KX-L1-A18).

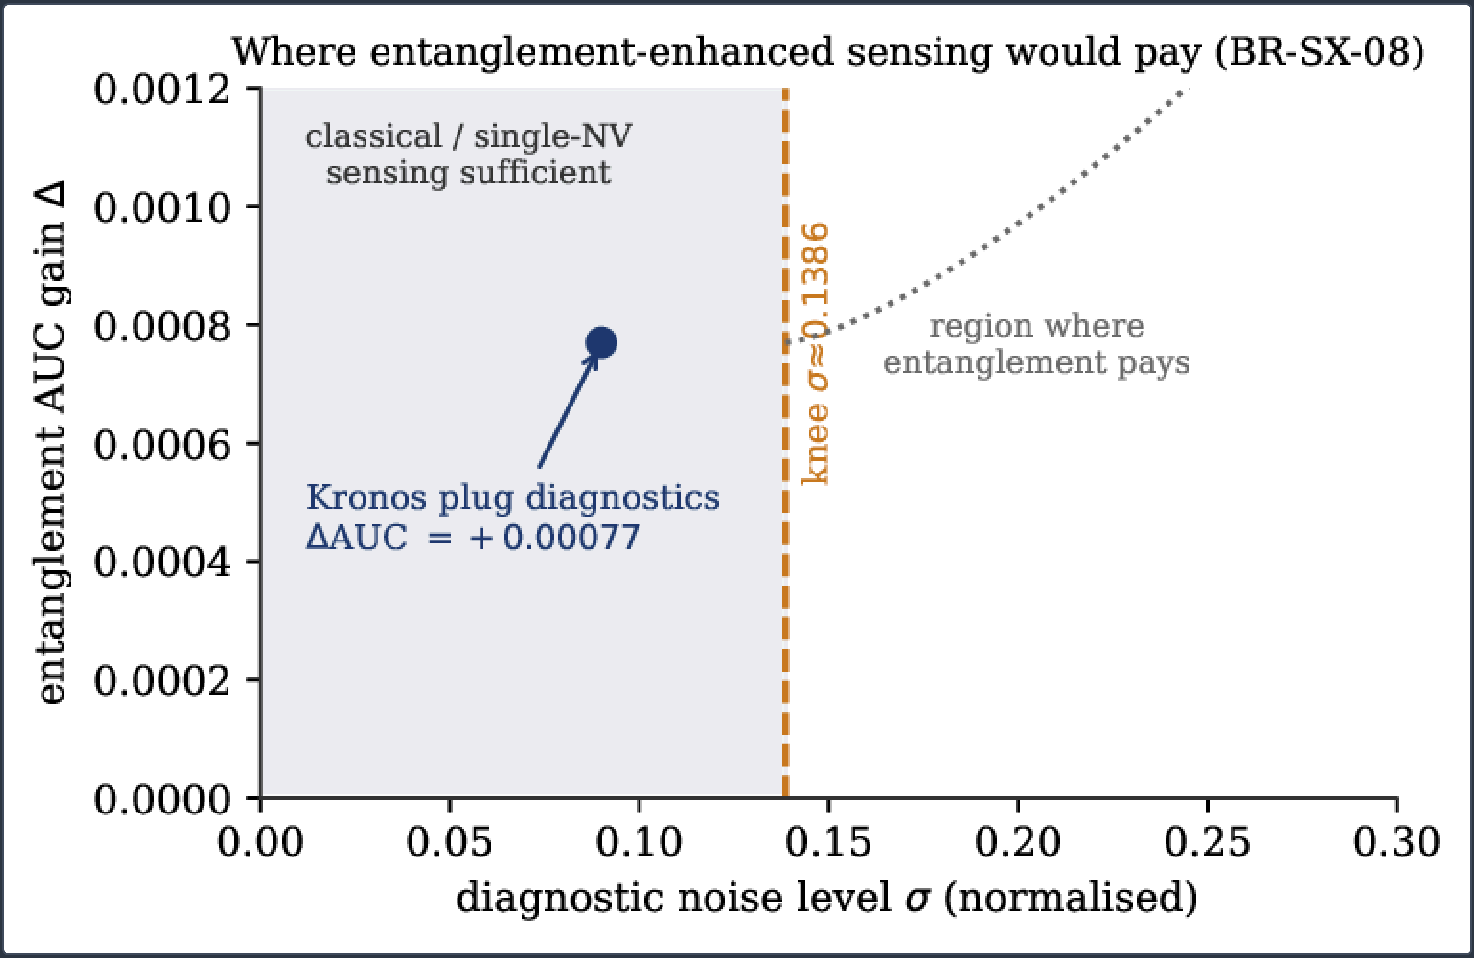

Entanglement classification gain ΔAUC of Eq. (8) versus normalised diagnostic noise σ. The knee at $\sigma \approx 0.1386$ (amber) separates the saturated regime, where classical/single-NV sensing is already sufficient, from the noise-limited regime where entanglement pays. The Kronos operating point (navy) sits below the knee at $+0.00077$ AUC (BR-SX-08). The dotted curve above the knee is the design expectation, not measured data.

QUANTITY	VALUE	BASIS / SERIAL
gyromagnetic ratio γ_{NV}	28.03 $Hertz	NV electron spin
demonstrated NV sensitivity δB	0.189	Eq. (6); KX-L1-A22
readout bandwidth	125 $Hertz	4 µs Ramsey shot; KX-L1-A22
SQL floor ($N = 100, \tau = 1\,\mu s$)	5.7e-7 T	Eq. (4); KX-L1-A18
Heisenberg floor (same N, τ)	5.7e-8 T	Eq. (5); KX-L1-A18
entanglement noise knee σ_{knee}	≈ 0.1386	Eq. (8); BR-SX-08
measured sub-knee entanglement gain	$+0.00077$ AUC	BR-SX-08

Plug sensing: demonstrated device, engineering floors, and the entanglement gate. Serials refer to the register [33].

THE BURN OPERATING POINT AND ITS UNCERTAINTY

The frozen burn solution is aneutronic-dominated (Figure 4): the D–^{3}He branch carries $3868.0\,MW$ of the $4298.5\,MW$ total, with the D–T, D–D$_p$ and D–D$_n$ channels contributing 176.0, 129.7 and $124.8\,MW$, for $f_n = 5.44\,\%$. The plug-ratio trade (Figure 4) shows why $R = 16$ is chosen: φ_i has saturated at $248.75\,keV$ while Q_E still climbs steeply, so the operating point sits at a robust knee rather than in the fragile cliff. The 4×10^4-sample global UQ (Figure 5) confirms robustness of the target the kinetic solver must hold: $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$ across the design ranges (KX-L1-A5-BU).

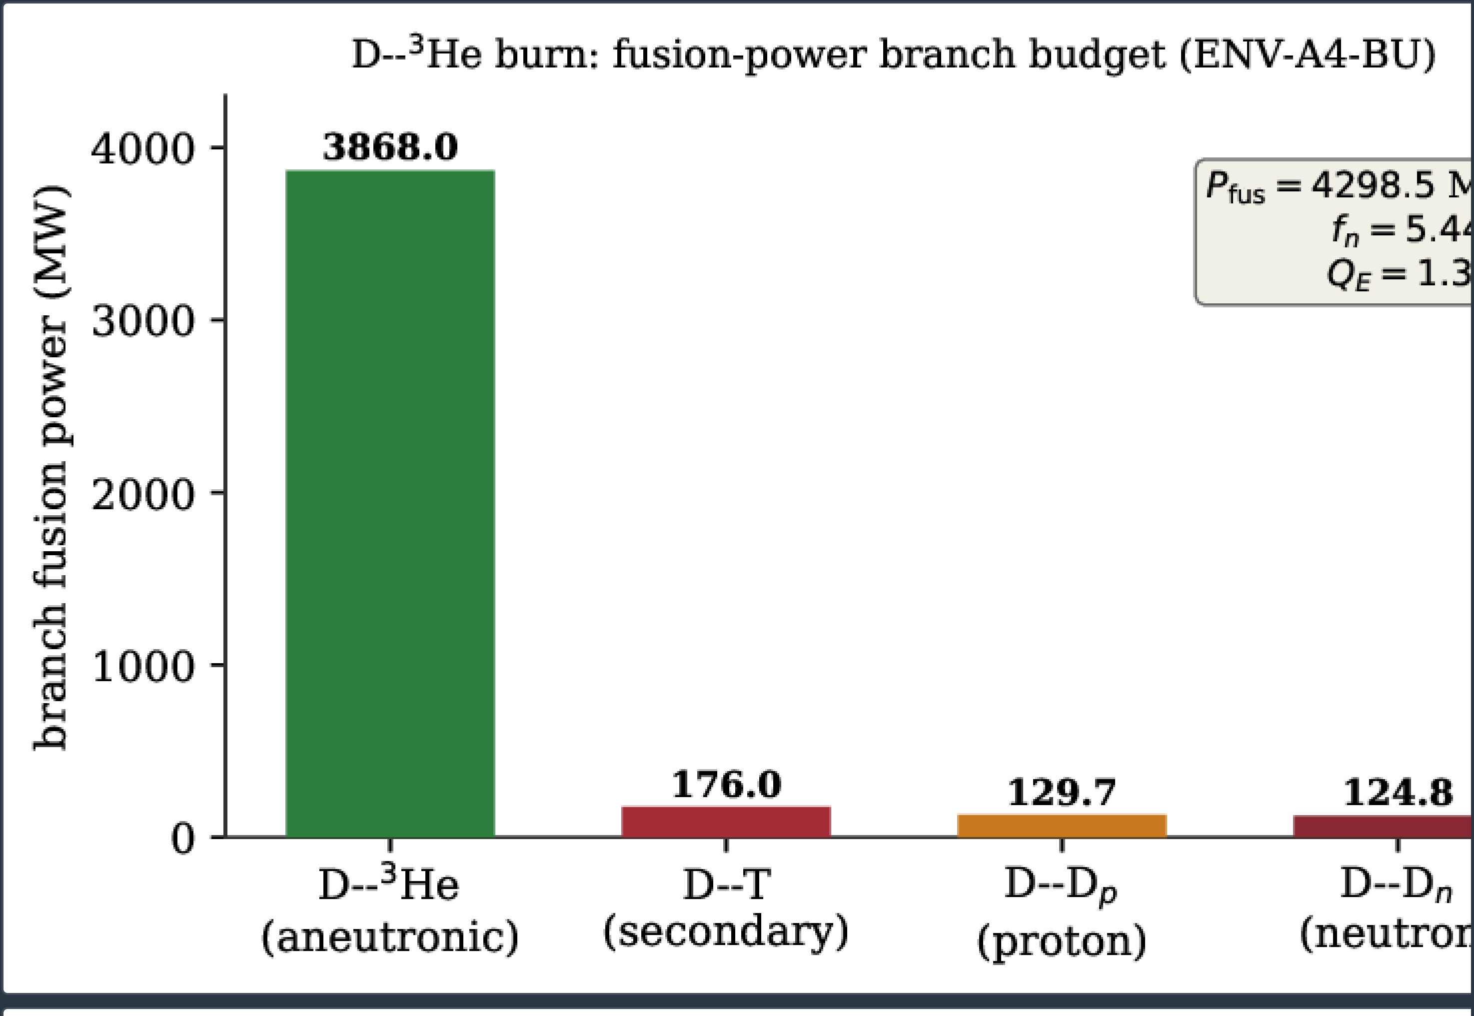
D--^{3}He burn: fusion-power branch budget (ENV-A4-BU)
branch fusion power (MW)
4000
3000
2000
1000
0
3868.0
176.0
129.7
124.8
P_{fus} = 4298.5 M
f_n = 5.44
Q_E = 1.3
D--^{3}He (aneutronic)
D--T (secondary)
D--D_p (proton)
D--D_n (neutron

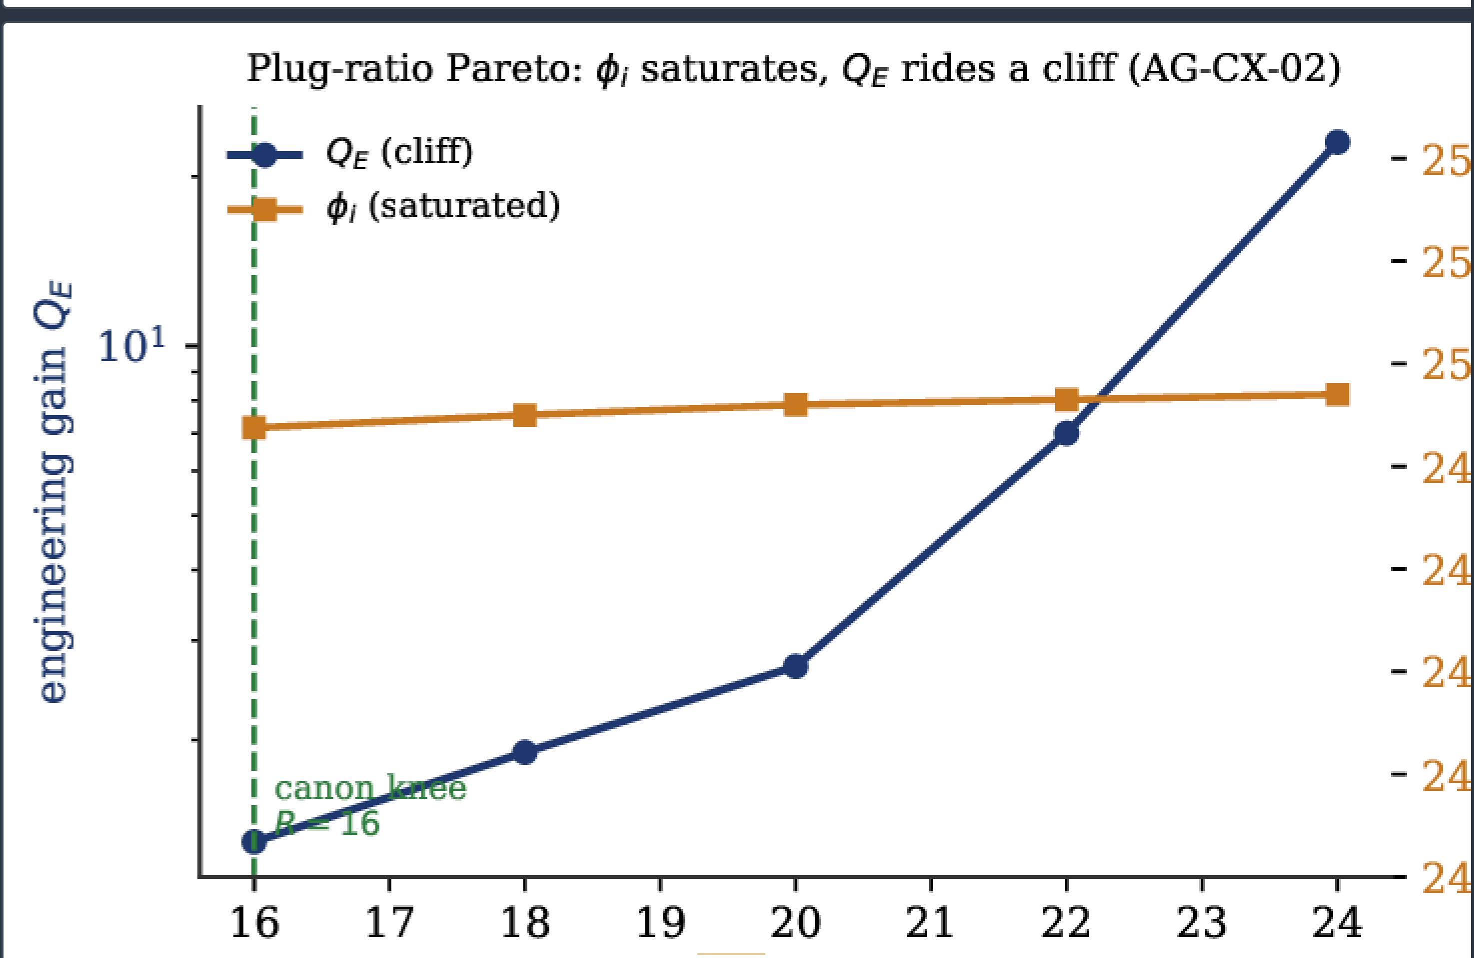
Plug-ratio Pareto: ϕ_i saturates, Q_E rides a cliff (AG-CX-02)
Q_E (cliff)
ϕ_i (saturated)
engineering gain Q_E
10^1
canon knee
R = 16
16
17
18
19
20
21
22
23
24
25
25
25
24
24
24
24
24

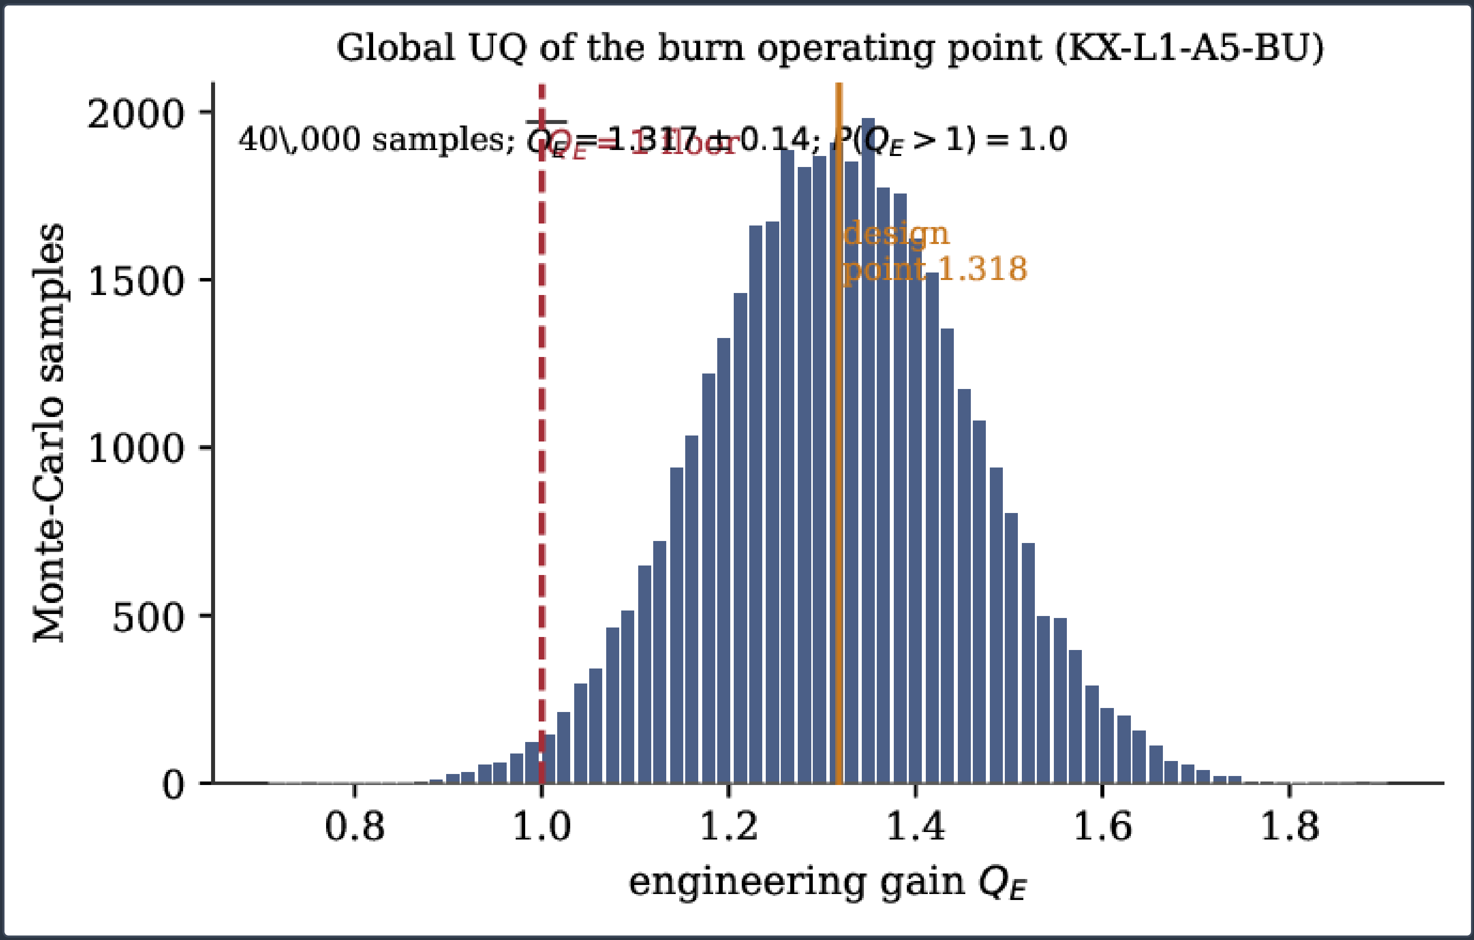

Global uncertainty quantification of the burn operating point (KX-L1-A5-BU): 4×10^4 Saltelli/Sobol samples give $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$; the design point $Q_E = 1.318$ (amber) sits well above the $Q_E = 1$ floor (red).

THE QUANTUM KINETIC SOLVER IS STATE-COMPRESSING BUT GATE-BOUND

Figure 6 makes the resource trade explicit. The amplitude-encoded logical-qubit count grows only as $\log_2 N$ in the velocity-grid size (amber), the whole appeal of the quantum route; but the non-Clifford gate cost of a control-relevant instance falls inside the bounded fault-tolerant envelope of 10^{19}–10^{23} T-gates (red band), which requires error-corrected execution. The state compression is real; the gate budget is the binding constraint.

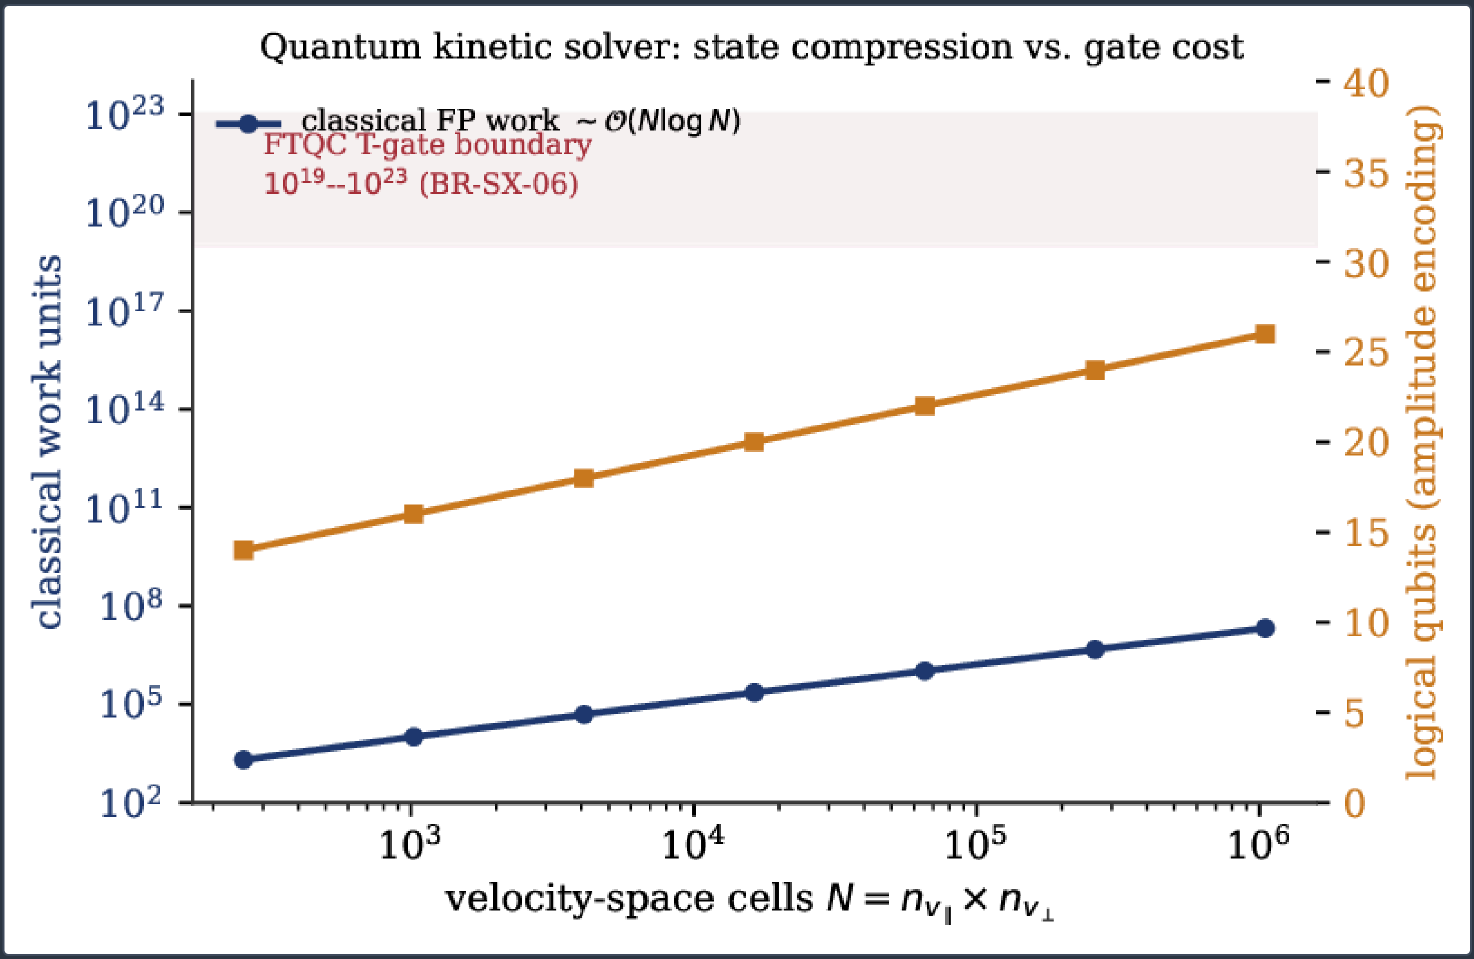

Quantum kinetic solver resource trade. Classical Fokker–Planck work scales as $\mathcal{O}(N \log N)$ in the velocity cells (navy); the amplitude-encoded quantum register grows only as $\log_2 N$ logical qubits (amber); the non-Clifford gate cost of a useful instance lands in the bounded FTQC envelope 10^{19}–10^{23} T-gates (red band, BR-SX-06). Derived scaling; the classical solver of §4.3 is authoritative today.

THE STACK AND ITS MATURITY GATE

Figure 7 collects the stack: a demonstrated single-NV sensor and an authoritative classical solver as the working elements, and two resource-scoped quantum designs (BU-QX-01, BU-QX-02) each admitted to control authority only through the same maturity gate the Kronos control architecture applies everywhere [doi:10.5281/zenodo.22645925]. Table 3 states the status of each component against its acceptance criterion.

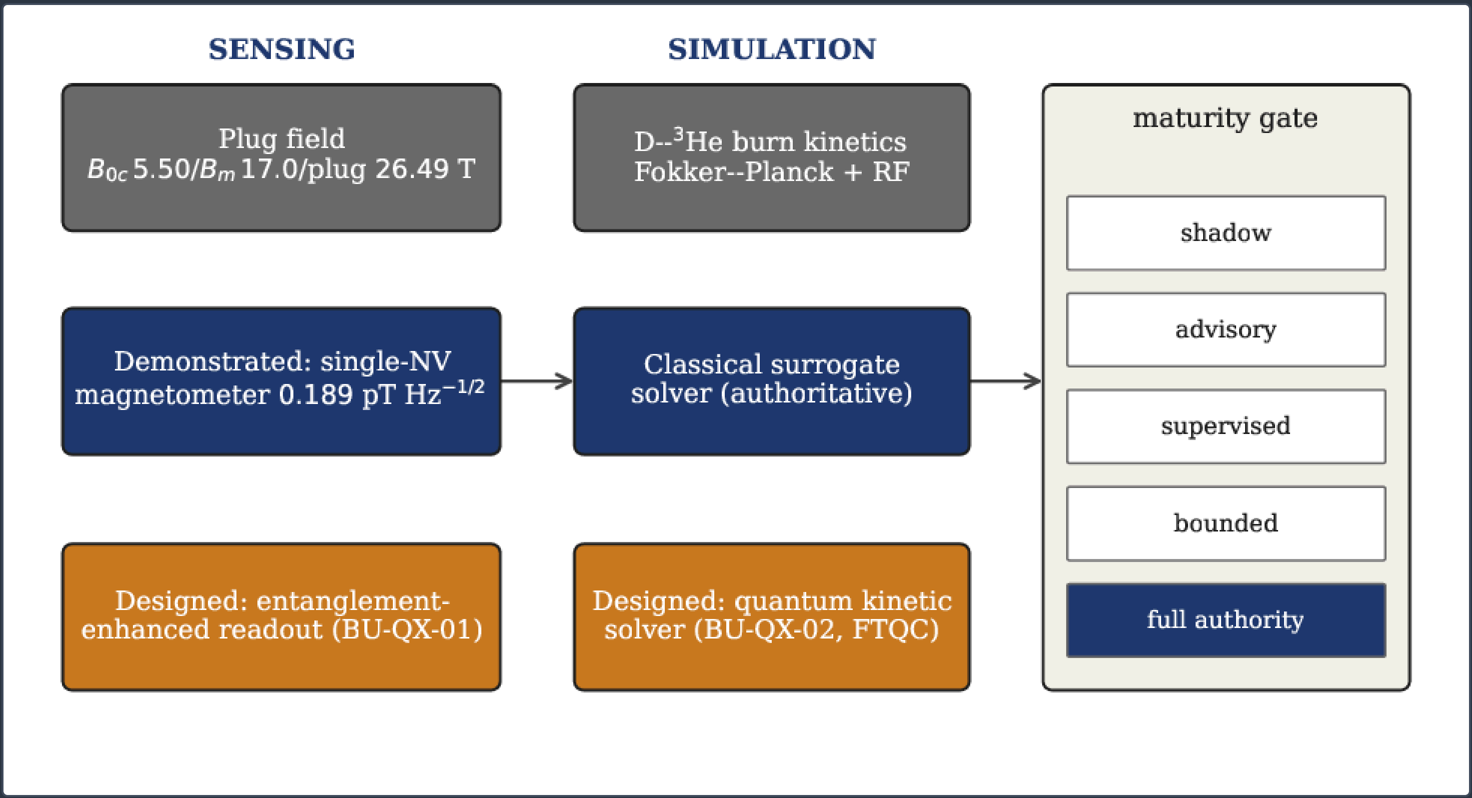

(Schematic.) The burner quantum stack. Left: sensing, from the plug field through the demonstrated single-NV magnetometer to the designed entanglement-enhanced readout (BU-QX-01). Centre: simulation, from the D–^{3}He Fokker–Planck kinetics through the authoritative classical surrogate to the designed quantum kinetic solver (BU-QX-02). Right: the verification-maturity gate that caps each component's control authority. The criteria are the real, frozen bars; the authority ladder is a schematic of the gating policy.

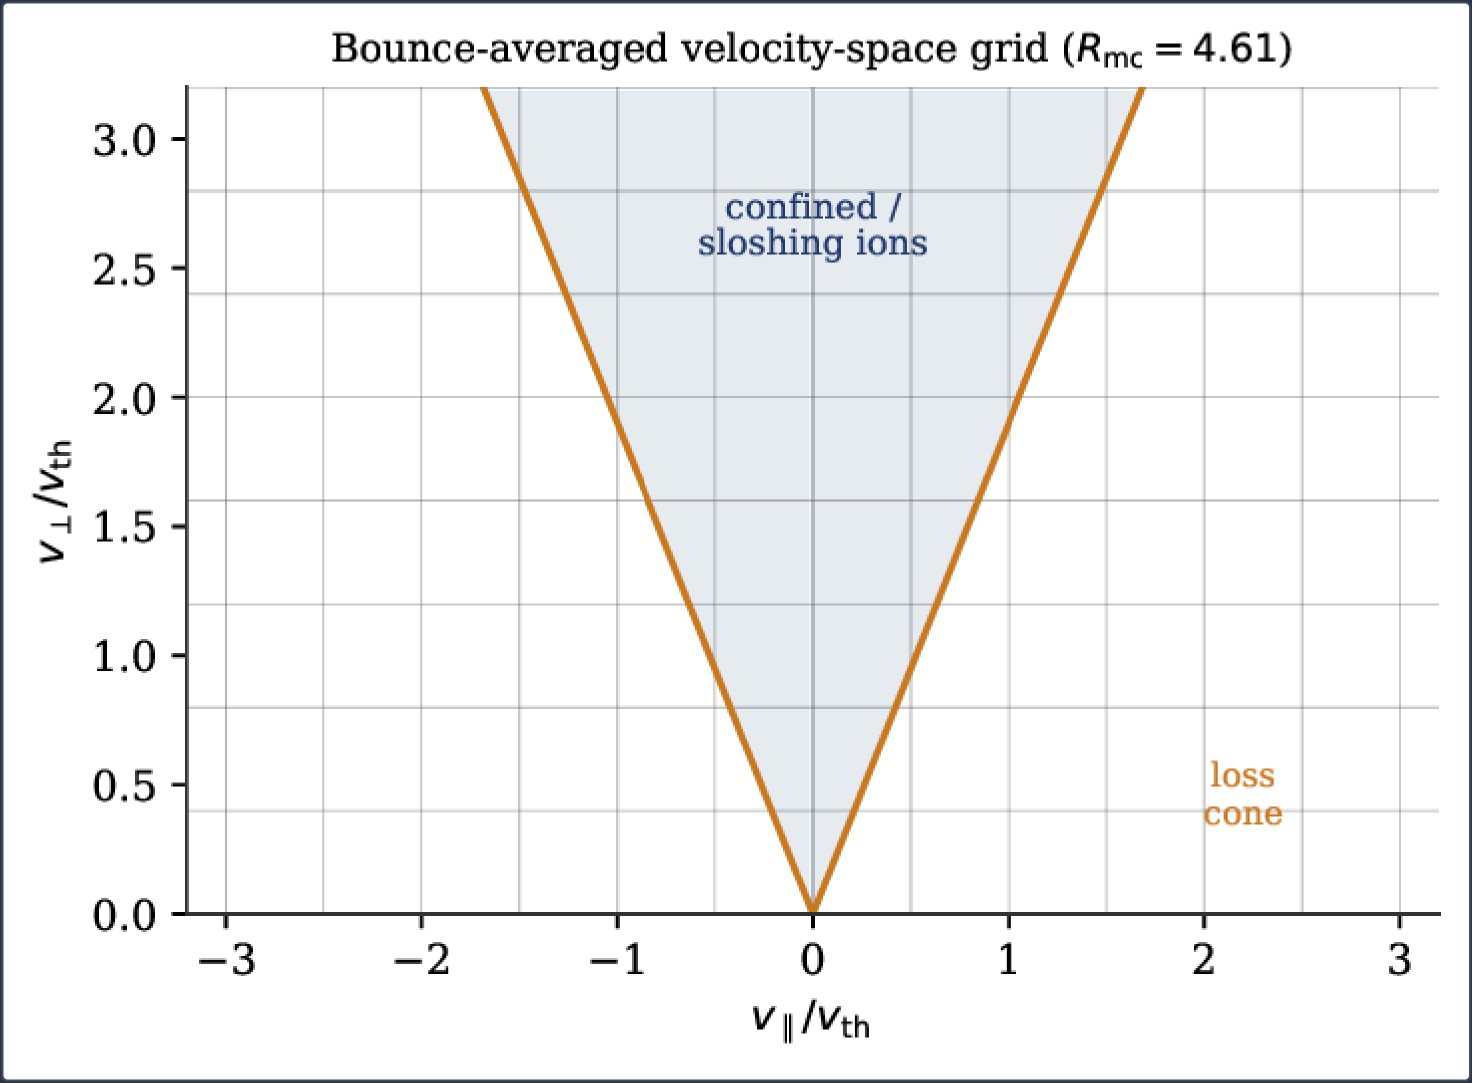

(Schematic.) Bounce-averaged velocity-space grid on which Eq. (9) is discretised (§4.3). The loss cone (amber) at $\theta_c = \arcsin R_{mc}^{-1/2}$ with $R_{mc} = 4.61$ bounds the confined and sloshing-ion population (navy). Grid and cone are schematic; no distribution data are plotted.

COMPONENT	ACCEPTANCE CRITERION	SERIAL	STATUS
single-NV plug magnetometer	resolve 5.50–$17.0\,T$ with margin	KX-L1-A22	**demonstrated (0.189)**
classical FP burn solver	reproduce $Q_E = 1.318$ anchor	BU-L1-A1	**demonstrated**
entanglement-enhanced sensor	beat classical above σ_{knee}	BU-QX-01	designed / gated ($\sigma > 0.1386$)
quantum kinetic solver	beat classical surrogate on acceptance test	BU-QX-02	designed / FTQC-scoped
global UQ of operating point	$P(Q_E > 1) = 1.0$	KX-L1-A5-BU	**demonstrated (1.317 ± 0.14)**

Burner quantum-stack components, acceptance criteria and status. “Demonstrated” denotes a reproduced frozen result; “designed / gated” denotes a resource-scoped design behind an explicit trigger. Serials refer to the register [33].

§ 8

Discussion

NV SENSITIVITY AGAINST THE QUANTUM-SENSING LITERATURE

The demonstrated 0.189 $\mathrm{pT\,Hz^{-1/2}}$ sits comfortably inside the envelope established for ensemble NV magnetometry. Degen, Reinhard and Cappellaro frame the shot-noise scaling that Eq. (6) instantiates [20]; Barry *et al.* analyse the concrete route to broadband ensemble sensitivities in the picotesla-and-below range through improvements to T_2^*, readout fidelity and diamond material [21]; Taylor *et al.* established the high-sensitivity diamond magnetometer at nanoscale resolution [23] and the NV physics is reviewed comprehensively by Doherty *et al.* [24,22]. The conservative Kronos parameter set ($N_c = 10^{12}$, $T_2^* = 1\,\mu s$, $C = 0.03$) is well within what that literature reports, so the sensitivity claim does not rest on an aggressive device. The point is one of *margin*: the plug fields stand orders of magnitude above the floor, so single-NV sensing is not merely adequate but generous.

THE ENTANGLEMENT GATE AGAINST THE METROLOGY LITERATURE

The $\sqrt{N}$ Heisenberg prize is real and famous [16,17,18,19], but Eq. (8) shows it is only cashed where the classical estimate is noise-limited. This is not a Kronos-specific artefact: it is a general property of detection at high signal-to-noise ratio, where both classical and quantum AUCs saturate. The measured $+0.00077$ sub-knee gain (BR-SX-08) is the quantitative statement that the burner plug diagnostics, operating at $\sigma \sim 10^{-3}$, are on the saturated side of the knee. The companion breeder-side study reaches the same conclusion for the centrepost magnetics and quantifies the fault-tolerant magnetometry resource estimates in parallel [doi:10.5281/zenodo.22645905]. Entanglement-enhanced sensing is thus a targeted tool for a specific future high-noise channel, not a general upgrade.

THE QUANTUM KINETIC SOLVER AGAINST THE CLASSICAL FRONTIER

The resource estimate of §5 places a useful burn-kinetics instance beyond the fault-tolerant boundary, consistent with the Kronos programme's broader quantum-computing findings: no measured advantage for variational optimisation or chemistry at simulable scale, and gate counts of 10^{19}–10^{23} that define a boundary rather than a near-term route (BR-SX-06, BR-SX-07). This is the same resource-honest reading the programme applies across the board [doi:10.5281/zenodo.22645718]. Variational and machine-learning shortcuts do not rescue it: barren plateaus [11] and quantum-kernel concentration [12] erase advantage as the problem scales, which the Kronos QML sweep observed directly (AUC $0.817 \to 0.691$ from 4 to 8 qubits; BR-SX-08). The productive near-term direction for state compression is classical: the matrix-product-state tensor-network solver of the companion study [doi:10.5281/zenodo.22645768] captures much of the same compression without the fault-tolerant overhead, and the classical finite-volume solver of §4.3 remains authoritative.

RELATION TO THE CONTROL STACK

The burner quantum stack is not free-standing; it feeds the burner control suite, whose plug-potential and sloshing-ion controllers act on the state the sensor reads and the solver predicts, inside a certified safety envelope. The maturity-gated authority of Figure 7 is the same principle that governs the whole architecture: a component that has passed only its own acceptance test may inform but not command. Because the mapping from criteria to authority is explicit, the path from the demonstrated single-NV sensor and classical solver to any future quantum upgrade is a defined, testable sequence rather than an open question.

§ 9

Limitations

The honest gate is stated once and it has two halves, both quantitative. First, there is *no measured quantum-computing advantage* on the burner control problems today: the quantum kinetic solver (BU-QX-02) is fault-tolerant-scoped, with the classical finite-volume solver of §4.3 authoritative until error-corrected hardware beats it on an explicit acceptance test, and the entanglement-enhanced sensor (BU-QX-01) has no measurable benefit ($+0.00077$ AUC) below the $\sigma \approx 0.1386$ knee at which every present plug diagnostic operates. Second, the one open *engineering* item on the demonstrated sensor is neutron tolerance: the in-winding NV chip's ensemble T_2^* and contrast will degrade under the burner's high background field, cryogenic operation and neutron irradiation, and the demonstrated diamond tolerance is roughly $39\times$ below the required in-service fluence (KFE-CF-SH-037, BR-SXD-05). The registered mitigation is a replaceable sensor cartridge plus a model-fusion backup that retains precursor discrimination with the NV channel degraded by up to $39\times$; a dedicated irradiation-plus-first-principles charge-state study is the follow-on that closes it. Neither half undermines the present value of the stack: the demonstrated single-NV sensor reads the plug with orders of margin, and the classical solver holds the burn.

§ 10

Conclusion

The Kronos burner has a defined, resource-honest quantum stack. Its working elements are demonstrated: a single-NV ensemble magnetometer at $\mathbf{0.189}$ $\mathbf{pT\,Hz^{-1/2}}$ and $\mathbf{125}$ Hertz$(KX-L1-A22) thatreadsthefixed$5.50/17.0/26.49 T $architecturewithordersofmargin, andaclassicalbounce-averagedFokker--Plancksolverthatholdsthefrozenburnpoint($ Q_E=1.318,P_ fus=4298.5 MW,f_n=5.44 %,\varphi_i=248.75 keV,=0.49 $). Onthatfoundationtwoquantumcomponentsareformallyspecifiedandresource-scoped: anentanglement-enhancedmagnetometer(BU-QX-01), whose$ N$HeisenbergprizeisderivedfromthequantumFisherinformationbutcashedonlyabovethemeasured$\!\!0.1386 $knee; andaquantumkineticsolver(BU-QX-02), whoseCarleman/block-encodingpipelinecompressesthevelocitystateto$ \log_2 N$logicalqubitsbutwhose$10^19--10^23T $-gatecostplacesausefulinstancebeyondthefault-tolerantboundary. Theburnerplugcoil($26.49 T$)andthebreedercentrepost($ 16.84 T$) are treated as the two distinct magnets they are throughout. The stack delivers a working sensor and solver now, two rigorously specified quantum designs, and an explicit, measured rule — the noise knee and the fault-tolerant boundary — for when either quantum component becomes worth building.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and the NVIDIA Inception partnership for the GPU HPC campaign.

Reproducibility. The plug-field architecture ($\mathbf{5.50/17.0/26.49}\,T$, $R_{\text{vac}}=\mathbf{3.09}$, $R_{\text{mc}}=\mathbf{4.61}$; KX-L1-A18-BU), the burn operating point ($Q_E=\mathbf{1.318}$, $P_{\text{fus}}=\mathbf{4298.5}\,MW$, $f_n=\mathbf{5.44}\,\%$, $\varphi_i=\mathbf{248.75}\,keV$; BU-L1-A1, KX-L1-A3-BU), the NV sensitivity ($\mathbf{0.189}$ $\mathbf{pT\,Hz^{-1/2}}$; KX-L1-A22), the metrology scaling ($\mathbf{5.7}e-\mathbf{7}\rightarrow\mathbf{5.7}e-\mathbf{8}\,T$; KX-L1-A18), the $\sigma\approx\mathbf{0.1386}$ knee and $+\mathbf{0.00077}$-AUC gain (BR-SX-08), the $\mathbf{4}\times\mathbf{10^4}$-sample UQ (KX-L1-A5-BU) and the FTQC gate envelope (BR-SX-06) are frozen entries in the Kronos Physics De-Risking Register (doi:10.5281/zenodo.22133057). Each figure is regenerable from the deposited figure-generation script. This paper is archived at doi:10.5281/zenodo.22645927 and its official page is https://www.kronosfusionenergy.com/publications/the-burner-quantum-stack-entanglement-enhanced-plug-fie.html.

Data availability tocsectionData availability The plug-field spec, the sensing metrics, the entanglement-gain data, the fuel-branch and UQ results, and the quantum-solver resource estimate used for the figures are archived with the deposit (doi:10.5281/zenodo.22645927) and with the AI-and-quantum companion (doi:10.5281/zenodo.22645702), and reference the Kronos 2026 series. All quantitative values trace to the register [33] by the serial identifiers cited inline. Any economic analysis is reported separately and is not part of the public deposit.

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

tocsectionReferences

APPARATUS

Portfolio, People & Record

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

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

INTELLECTUAL PROPERTY

The patent portfolio

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

GRANTED & PENDING UTILITY

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

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

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

TRADEMARK & ORIGIN

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

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

THE TEAM

Founding partners, board, advisors & auditors

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

FOUNDING PARTNERS — SCIENCE

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

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

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

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

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

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

SCIENTIFIC ADVISORS

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

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

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

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

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

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

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

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

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

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

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

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

SCIENTIFIC AUDITORS

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

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

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

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

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

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

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

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

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

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

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

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

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

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

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

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

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

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

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

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

Andrea Romero
Marketing Lead
2022–Present

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

SOURCES

References

1 R P Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* **21** 467 (1982). doi:10.1007/BF02650179

2 S Lloyd, Universal quantum simulators, *Science* **273** 1073 (1996). doi:10.1126/science.273.5278.1073

3 I M Georgescu, S Ashhab and F Nori, Quantum simulation, *Rev. Mod. Phys.* **86** 153 (2014). doi:10.1103/RevModPhys.86.153

4 J Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2** 79 (2018). doi:10.22331/q-2018-08-06-79

5 A W Harrow, A Hassidim and S Lloyd, Quantum algorithm for linear systems of equations, *Phys. Rev. Lett.* **103** 150502 (2009). doi:10.1103/PhysRevLett.103.150502

6 D W Berry, High-order quantum algorithm for solving linear differential equations, *J. Phys. A: Math. Theor.* **47** 105301 (2014). doi:10.1088/1751-8113/47/10/105301

7 J-P Liu, H Kolden, H K Krovi, N F Loureiro, K Trivisa and A M Childs, Efficient quantum algorithm for dissipative nonlinear differential equations, *Proc. Natl. Acad. Sci. USA* **118** e2026805118 (2021). doi:10.1073/pnas.2026805118

8 G H Low and I L Chuang, Hamiltonian simulation by qubitization, *Quantum* **3** 163 (2019). doi:10.22331/q-2019-07-12-163

9 A Gily\'en, Y Su, G H Low and N Wiebe, Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics, in *Proc. 51st ACM SIGACT Symp. on Theory of Computing (STOC)* 193 (2019). doi:10.1145/3313276.3316366

10 M Cerezo *et al*, Variational quantum algorithms, *Nat. Rev. Phys.* **3** 625 (2021). doi:10.1038/s42254-021-00348-9

11 J R McClean, S Boixo, V N Smelyanskiy, R Babbush and H Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9** 4812 (2018). doi:10.1038/s41467-018-07090-4

12 H-Y Huang, M Broughton, M Mohseni, R Babbush, S Boixo, H Neven and J R McClean, Power of data in quantum machine learning, *Nat. Commun.* **12** 2631 (2021). doi:10.1038/s41467-021-22539-9

13 A G Fowler, M Mariantoni, J M Martinis and A N Cleland, Surface codes: towards practical large-scale quantum computation, *Phys. Rev. A* **86** 032324 (2012). doi:10.1103/PhysRevA.86.032324

14 N F Ramsey, A molecular beam resonance method with separated oscillating fields, *Phys. Rev.* **78** 695 (1950). doi:10.1103/PhysRev.78.695

15 S L Braunstein and C M Caves, Statistical distance and the geometry of quantum states, *Phys. Rev. Lett.* **72** 3439 (1994). doi:10.1103/PhysRevLett.72.3439

16 V Giovannetti, S Lloyd and L Maccone, Quantum-enhanced measurements: beating the standard quantum limit, *Science* **306** 1330 (2004). doi:10.1126/science.1104149

17 D J Wineland, J J Bollinger, W M Itano, F L Moore and D J Heinzen, Spin squeezing and reduced quantum noise in spectroscopy, *Phys. Rev. A* **46** R6797 (1992). doi:10.1103/PhysRevA.46.R6797

18 M Kitagawa and M Ueda, Squeezed spin states, *Phys. Rev. A* **47** 5138 (1993). doi:10.1103/PhysRevA.47.5138

19 L Pezz\`e, A Smerzi, M K Oberthaler, R Schmied and P Treutlein, Quantum metrology with nonclassical states of atomic ensembles, *Rev. Mod. Phys.* **90** 035005 (2018). doi:10.1103/RevModPhys.90.035005

20 C L Degen, F Reinhard and P Cappellaro, Quantum sensing, *Rev. Mod. Phys.* **89** 035002 (2017). doi:10.1103/RevModPhys.89.035002

21 J F Barry, J M Schloss, E Bauch, M J Turner, C A Hart, L M Pham and R L Walsworth, Sensitivity optimization for NV-diamond magnetometry, *Rev. Mod. Phys.* **92** 015004 (2020). doi:10.1103/RevModPhys.92.015004

22 L Rondin, J-P Tetienne, T Hingant, J-F Roch, P Maletinsky and V Jacques, Magnetometry with nitrogen-vacancy defects in diamond, *Rep. Prog. Phys.* **77** 056503 (2014). doi:10.1088/0034-4885/77/5/056503

23 J M Taylor, P Cappellaro, L Childress, L Jiang, D Budker, P R Hemmer, A Yacoby, R Walsworth and M D Lukin, High-sensitivity diamond magnetometer with nanoscale resolution, *Nat. Phys.* **4** 810 (2008). doi:10.1038/nphys1075

COLOPHON

How this document was made

The Kronos Fleet — Combined Editorial, 2026 is set in **Fraunces** (display), **Gelasio** (text), and **IBM Plex Mono** (data & equations), carried forward from the Kronos editorial design system.

The figures are rendered at 300 dpi from the deposited generator scripts; the document is composed as a single self-contained file with embedded fonts and imagery, paginated to US Letter, and rendered to PDF through a headless Chromium engine in matched light and dark editions.

Every headline quantity carries an evidence class; withdrawn and restated values are marked as such. Nothing herein is frozen beyond the founder-approved Tier-A set.

Explore it live — turn the interactive 3-D model, run the deposited solvers in the physics validator, and watch the companion film series.

LEGAL NOTICE

Notice, status, and distribution

Conceptual design & simulation study. This document reports a conceptual design and simulation study. It is not a construction commitment, a safety-analysis report, a regulatory filing, or an offer of securities. Forward-looking statements — schedules, costs, market sizes, and performance — are estimates subject to the limitations set out in the Simulations part and may change.

Numbers and their classes. Quantities are reported with evidence classes and case labels; modelled, assumed, and requirement-class values are identified as such and are not to be quoted without their labels.

Public edition. This edition carries the physics and design only; all financial, funding, and defense-commercial content has been removed for public distribution. The scientific text corresponds to the arXiv and journal submissions.

THE 2026 KRONOS SERIES

One programme, many papers

This editorial is one paper in the Kronos 2026 series — a real fusion company's complete, reproducible physics record for a spherical-tokamak breeder and its low-neutron D–^{3}He tandem-mirror burners. Every headline number is a frozen anchor in the Physics De-Risking Register.

Explore the whole programme: the De-Risking Register, the interactive 3-D model, and the physics validator.

KRONOS FUSION ENERGY

The Burner Quantum Stack: Entanglement-Enhanced Plug-Field Magnetometry and a Quantum Kinetic Solver for D–^{3}He Tandem-Mirror Burn Control

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

Zenodo · doi:10.5281/zenodo.22645927

Read on kronosfusionenergy.com

The Physics De-Risking Register