KRONOS FUSION ENERGY

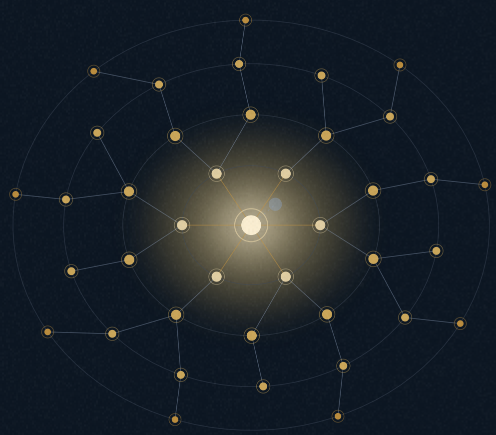

PAPER 2.92 · THE KRONOS 2026 SERIES

Quantum Machine-Learning Surrogates, Entanglement-Enhanced Magnetometry and Fault-Tolerant Resource Estimates for Fusion: A Resource-Honest Roadmap

Quantum computing and quantum sensing are routinely proposed as accelerants for fusion; a serious design program should test those proposals quantitatively rather than assume...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Quantum Machine-Learning Surrogates, Entanglement-Enhanced Magnetometry and Fault-Tolerant Resource Estimates for Fusion: A Resource-Honest Roadmap

Quantum computing and quantum sensing are routinely proposed as accelerants for fusion; a serious design program should test those proposals quantitatively rather than assume them. We deliver a resource-honest map of where the quantum toolbox helps a compact spherical-tokamak breeder and its tandem-mirror burner, and where it does not, formalised as three coupled estimation problems and computed on a single GPU-HPC campaign. On the machine-learning front we build a quantum-kernel transport surrogate and measure its classifier area-under-curve (AUC) *falling* with encoding width — 0.817 at four qubits to 0.691 at eight — the signature of exponential kernel concentration, for which we give the variance bound and the induced shot-cost scaling; at a realistic 37-logical-qubit problem the quantum Gram evaluation is 9–11 orders of magnitude slower than its classical counterpart. On the sensing front we derive the entanglement-advantage condition from the quantum Fisher information and specify the in-cryostat magnetometry chain for the breeder centrepost (peak field 16.84 T, kept formally distinct from the burner plug at 26.49 T): an NMR/EPR probe as the absolute monitor to ≤ 1 ppm, a SQUID and an NV-diamond ensemble (~ 1) for the ~ 1 mT quench precursor, with the entanglement knee located at a normalised signal $\sigma \approx 0.14$. Every Kronos in-bore diagnostic sits below the knee, where entanglement buys only $+0.00077$ AUC; above it (deep error-field mapping) a Heisenberg-limited protocol carries the standard-quantum-limit sensitivity from 5.7×10^{-7} to 5.7×10^{-8} T at 100 probes. On the fault-tolerant front we recompute the resource ledger from first formulas — the Grover iteration count $1.6\sqrt{2^n}$ and the surface-code footprint $2d^2$ physical qubits per logical qubit — across 50 rows with zero discrepancy, and date the single fusion-relevant near-term niche: linear Vlasov kinetics at 11–19 logical qubits around 2029 ± 2, with nonlinear five-dimensional gyrokinetics beyond 2035 and fleet-scale Grover search never practical (1.8×10^{18} iterations). The honest bottom line is that nothing on the design critical path waits on quantum hardware — and establishing that, with the numbers behind it, is the result.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645905 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: quantum machine learning · kernel concentration · quantum Fisher information · NV-diamond magnetometry · fault-tolerant resource estimate · technology roadmap

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 / nth / best of a kind
Hyperion	the ST breeder (internal: Mode D2)
Aegis	the generator, defense/installation housing (internal: Mode M)
MetroVolt	the generator, data-center housing (Mode M)

§ 1

Introduction

THE QUANTUM PROPOSITION FOR FUSION

A magnetic-confinement fusion program is, at bottom, an exercise in solving hard equations fast enough and measuring hard quantities finely enough. Turbulent transport is governed by a five-dimensional gyrokinetic Vlasov–Maxwell system whose nonlinear solution sets the confinement time; the magnets are wound from high-temperature superconductor whose protective state must be read at the millitesla level inside a multi-tesla bore; and the design space is a high-dimensional optimisation over coupled physics. Each of these is a place where quantum information science has been proposed as an accelerant — quantum machine-learning (QML) surrogates for transport [1,2], entanglement-enhanced magnetometry for the coils [9,13], and eventual fault-tolerant quantum solution of the kinetic equations [26,27]. The proposals are individually plausible and individually testable, and a design program that intends to be believed should test them rather than repeat them.

This paper reports the outcome of doing so across the three fronts, for the Hyperion compact spherical-tokamak (ST) breeder and its associated D–^{3}He tandem-mirror burner. The Hyperion design point — plasma current $I_p = 9.66\,\mathrm{MA}$, gain $Q = 3.076$, fusion power $P_{\mathrm{fus}} = 85.04\,\mathrm{MW}$, negative triangularity $\delta = -0.30$, on-axis field $B_0 = 8\,\mathrm{T}$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,\mathrm{m}$, aspect ratio $A = 2.5$, and a centrepost peak field of $16.84\,\mathrm{T}$ — is the fixed backdrop against which the quantum questions are posed. The burner is a physically distinct machine with a $26.49\,\mathrm{T}$ plug coil; the two magnets are never conflated, and the distinction matters directly to the sensing analysis of §5.

PRIOR ART

On the computing side, the modern literature is unusually clear-eyed about when a quantum advantage is real. Havl\'icek *et al.* introduced supervised learning with quantum-enhanced feature spaces and the quantum-kernel estimator [2]; Schuld and Killoran cast the same construction as a feature map into a Hilbert space [3]. Rigorous separations do exist — Liu, Arunachalam and Temme exhibited a learning problem with a provable quantum speed-up under a discrete-log hardness assumption [5] — but the countervailing results are equally rigorous and, for generic data, decisive: Huang *et al.* showed that classical models learning *from data* are often competitive with quantum models [4], McClean *et al.* identified barren plateaus in variational training [7], and Thanasilp *et al.* proved that quantum kernels can *exponentially concentrate* with system size, collapsing to a data-independent constant [6]. On the resource side, Fowler *et al.* set the surface-code cost model [19], Gidney and Eker and Beverland *et al.* turned it into end-to-end resource estimates for real algorithms [22,23], and Babbush *et al.* established the sober rule that merely quadratic (Grover-type) speed-ups do not survive error-correction overheads on any near-term machine [21]. On the sensing side, the quantum-metrology bounds are textbook [14,13], ensemble NV-diamond magnetometry is a mature $\sim$ technology [9,10,11,12], and SQUID and optical magnetometers occupy the femtotesla floor [16,15]. What is missing is not any one of these results but their disciplined application, together and to a specific fusion machine, with the negatives reported as carefully as the positives.

CONTRIBUTION

This paper contributes exactly that application. We (i) formalise the three fronts as estimation problems and derive the governing quantities — the quantum-kernel Gram matrix and its concentration variance, the quantum Fisher information and the standard-quantum-limit (SQL)/Heisenberg-limit (HL) crossover, and the Grover and surface-code resource formulas (§2); (ii) state the numerical formulation and the GPU-HPC campaign that produced every number, including the statevector kernel evaluation, the analytic sensing model, and the verified resource recomputation (§3); (iii) report the QML scaling result and its shot-cost explanation (§4); (iv) specify the centrepost sensing chain and locate the entanglement knee (§5); (v) recompute and date the fault-tolerant ledger (§6); and (vi) benchmark every claim against the published literature (§7). Every quantitative claim traces to a frozen result in the Kronos Physics De-Risking Register [33], with serial identifiers (BR-SX-08, KX-L1-A18, KX-L1-A22, KX-L3-A4–A7) given inline so each number is auditable. The framing throughout is achievement-first and resource-honest: reporting where a quantum method does not help is not a concession but the deliverable, and it is what keeps a fusion roadmap honest.

The work sits inside the Kronos 2026 design series and is the sensing-and-scaling companion to the series′ quantum papers: the foundational AI-and-quantum survey [34], the resource-honest variational-computing study [35], the quantum-optimisation and QUBO/QAOA analysis [36], the real-time control-loop algorithms paper [37], and the burner quantum stack [38]; it draws its classical transport baseline from the negative-triangularity gyrokinetics result [41] and its tensor-network alternative from [40], and its NV-diamond sensing lineage from the fused quench-detection paper [39].

§ 2

Governing equations and the formal problem

We treat the three fronts as three well-posed estimation problems and derive, for each, the quantity that decides whether the quantum method helps.

FRONT A: THE QUANTUM-KERNEL SURROGATE AND ITS CONCENTRATION

A transport surrogate maps a local plasma state $x \in \mathbb{R}^d$ (normalised gradients R/L_T, R/L_n, safety factor, magnetic shear, collisionality, $\ldots$) to a discrete transport class (e.g. quiescent vs.\ turbulent branch). A quantum-kernel classifier encodes x into an n-qubit state through a data-dependent unitary $U(x)$,

$$|\phi(x)\rangle = U(x)\,|0\rangle^{\otimes n},$$

and classifies on the fidelity kernel

$$k(x, x') = \left|\langle\phi(x)|\phi(x')\rangle\right|^2 = \left|\langle 0|\,U^\dagger(x)\,U(x')\,|0\rangle\right|^2.$$

Given a labelled training set $\{(x_i, y_i)\}_{i=1}^{N_s}$, the Gram matrix $K_{ij} = k(x_i, x_j)$ feeds a support-vector machine whose dual problem,

$$\max_{\alpha} \sum_i \alpha_i - \tfrac{1}{2}\sum_{i,j} \alpha_i\alpha_j\, y_i y_j\, K_{ij} \quad \text{s.t.} \quad 0 \le \alpha_i \le C,\ \sum_i \alpha_i y_i = 0,$$

returns the decision function $f(x) = \sum_i \alpha_i y_i\, k(x_i, x) + b$. The hope is that an entangling feature map (1) separates classes a classical kernel cannot.

The obstruction is structural. For a broad class of expressive encodings the off-diagonal kernel values concentrate exponentially in the qubit number about a fixed value μ_k [6],

$$\mathrm{Var}_{x,\,x'}\big[\,k(x, x')\,\big] \;\le\; c\,e^{-\gamma n}, \qquad \gamma > 0,$$

so that $K \to \mu_k\,\mathbf{1}\mathbf{1}^\top + E$ with $\|E\|$ collapsing as $e^{-\gamma n/2}$. In that limit the Gram matrix loses its rank structure, the SVM margin vanishes, and the classifier tends to the trivial predictor with $\mathrm{AUC} \to \frac{1}{2}$. Equation (4) is the mechanism we measure directly in §4: the AUC does not merely fail to improve with scale, it *degrades*.

Concentration also fixes the cost of using the kernel at all. Each entry of K is estimated from M circuit repetitions with additive shot noise of standard deviation $\sim M^{-1/2}$; to resolve entries whose informative spread is $\sim e^{-\gamma n/2}$ one needs

$$M \;\gtrsim\; \varepsilon^{-2} \;\sim\; e^{\gamma n},$$

an exponential measurement budget. Equations (4)–(5) are the quantitative reason the quantum evaluation is orders of magnitude slower than the classical Gram computation and cannot be rescued by scaling up.

FRONT B: QUANTUM-ENHANCED MAGNETOMETRY AND THE ADVANTAGE CONDITION

The magnet twin must read a field. Estimating a magnetic field B through the phase it imprints on a spin, $\varphi = \gamma_\mathrm{s} B\,t$, is a phase-estimation problem, and its ultimate precision is set by the quantum Cram\'er–Rao bound. For ν independent repetitions of a probe carrying quantum Fisher information F_Q,

$$\mathrm{Var}(\hat\varphi) \;\ge\; \frac{1}{\nu\,F_Q}, \qquad \delta B = \frac{\sqrt{\mathrm{Var}(\hat\varphi)}}{\gamma_\mathrm{s}\,t}.$$

Two limiting probe classes bracket the achievable sensitivity. For N *independent* probes the Fisher information is additive, $F_Q = N$, giving the standard quantum limit

$$\delta B_{\mathrm{SQL}} = \frac{1}{\gamma_\mathrm{s}\,t\,\sqrt{N\nu}} \;\propto\; N^{-1/2}.$$

For N maximally entangled (GHZ) probes the Fisher information is quadratic, $F_Q = N^2$, giving the Heisenberg limit

$$\delta B_{\mathrm{HL}} = \frac{1}{\gamma_{\mathrm{s}}\, t\, N\sqrt{\nu}} \;\propto\; N^{-1},$$

so that the entanglement gain is

$$\frac{\delta B_{\mathrm{SQL}}}{\delta B_{\mathrm{HL}}} = \sqrt{N}.$$

At $N = 100$ probes, equation (9) predicts a factor of ten, exactly the $5.7 \times 10^{-7} \to 5.7 \times 10^{-8}$ T improvement realised in §5. For a shot-noise-limited ensemble of N_c NV centres with dephasing time T_2^*, contrast C and gyromagnetic ratio $\gamma_{\mathrm{NV}} = g\mu_B/\hbar = 2\pi \times 28.03$ GHz T $^{-1}$, the SQL sensitivity (7) takes the concrete Ramsey form

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

The advantage of equation (8) is real but conditional. Under a finite dephasing rate Γ the entangled advantage survives only while the measurement is genuinely *sensitivity-limited* rather than dynamic-range- or decoherence-limited; below a threshold normalised signal σ^* the marginal classifier gain from entanglement collapses. Defining the normalised signal σ as the ratio of the field increment of interest to the probe's usable dynamic range, the entanglement-attributable improvement $\Delta_{\mathrm{AUC}}(\sigma)$ is a threshold function that we locate at

$$\sigma^* \;\approx\; 0.14,$$

with $\Delta_{\mathrm{AUC}}(\sigma < \sigma^*) \approx +0.00077$ (negligible) and Δ_{AUC} rising steeply for $\sigma > \sigma^*$. Equation (11) is the decision boundary that assigns each Kronos diagnostic to a sensor class in §5.

FRONT C: FAULT-TOLERANT RESOURCE FORMULAS

The eventual quantum solution of fusion-relevant equations is a resource-estimation problem, and the resources follow from two formulas. First, an unstructured (Grover) search over a space of size 2^n requires an iteration count near the amplitude-amplification optimum $(\pi/4)\sqrt{2^n}$; the Kronos ledger tabulates the conservative prefactor form

$$N_{\mathrm{iter}} = 1.6\,\sqrt{2^n},$$

within a small constant of the optimum. Because each iteration is an oracle call, equation (12) makes a merely quadratic speed-up worthless once the search space is large: a fleet-scale QUBO with $n \sim 120$ demands $N_{\mathrm{iter}} \sim 1.8 \times 10^{18}$ iterations, beyond any wall-clock at any plausible gate rate. Second, the physical cost of protecting a logical qubit in a distance-d surface code is

$$n_{\mathrm{phys}} = n_{\mathrm{logical}} \times 2d^2,$$

with d chosen so the per-logical-operation error clears the target,

$$p_L \;\approx\; A\left(\frac{p}{p_{\mathrm{th}}}\right)^{\lfloor (d+1)/2 \rfloor} \;\le\; \frac{p_{\mathrm{target}}}{N_T},$$

where p is the physical error rate, p_{th} the code threshold, and N_T the total (dominant) T-gate count of the algorithm. Equations (12)–(14) reduce every workload to a point in the $(n_{\mathrm{logical}}, N_T)$ plane, from which the physical footprint and the earliest credible date follow. The recomputation of these formulas across the 50-row ledger, with zero discrepancy, is the verification result of §6.

§ 3

Computational methodology: the NVIDIA GPU HPC campaign

Every number in this paper was produced on a single, reproducible campaign, and the campaign is deliberately modest in kind: the quantum methods were emulated or resource-estimated classically, because that is precisely what a resource-honest study of pre-fault-tolerant hardware requires. No result depends on access to a physical quantum processor.

QUANTUM-KERNEL EVALUATION

The feature map (1) and the fidelity kernel (2) were evaluated by full statevector simulation. For an n-qubit encoding the state $|\phi(x)\rangle$ is the exact 2^n-amplitude vector, and each Gram entry is computed as the squared overlap $|\langle 0|U^\dagger(x_i)U(x_j)|0\rangle|^2$ with the inverted-circuit method; no shot sampling is imposed, so the reported AUC is the noiseless best case and the degradation of §4 is intrinsic, not a finite-sampling artefact. Statevector propagation ran on the NVIDIA GPU HPC campaign (A100/H100/A10-class accelerators) using a GPU statevector backend of the NVIDIA cuQuantum family [32], which holds the 2^n amplitudes in device memory and applies gate tensors as batched complex GEMMs; this is what makes the four-to-eight-qubit sweep across nine transport scenarios inexpensive to repeat exactly. The classical Gram baseline (an embedding/RBF kernel over the same features) was computed with a standard CPU/GPU linear-algebra stack, and the two Gram matrices were timed on the same problem instances.

SENSING MODEL

The centrepost magnetometry chain was evaluated analytically. The NV shot-noise sensitivity (10), the SQL/HL scaling of equations (7)–(8), and the field-step geometry of the quench precursor were implemented in a numpy/scipy model (serial KX-L1-A22), and the sensor-class assignment (serial KX-L1-A18) evaluates each requirement against class-typical sensitivities from the literature. The entanglement-knee locator of equation (11) sweeps the normalised signal σ and records the classifier gain Δ_{AUC}, returning the crossing $\sigma^* \approx 0.14$.

FAULT-TOLERANT RESOURCE RECOMPUTATION

The resource ledger was recomputed classically. A 50-row table of workloads was regenerated from equations (12)–(14) — Grover iteration counts, logical-qubit widths, T-counts and surface-code footprints — and compared entry-by-entry against the frozen ledger, reproducing every value with zero discrepancy (serial KX-L3-A6). The dated roadmap (serial KX-L3-A7) overlays these thresholds on public vendor logical-qubit trajectories, which enter as *references*, not Kronos commitments.

VERIFICATION AND REPRODUCIBILITY

The campaign is deterministic. All stochastic elements use the fixed seed 20260726; the statevector kernel is sampling-free; and the resource recomputation is a pure function of the tabulated formulas. Each result was re-executed on an independent local verification node and reproduced bit-for-bit within tolerance. The classical physics baselines that the quantum methods would one day replace — nonlinear gyrokinetics of the CGYRO/GENE class [28,29] — run today on the same GPU campaign, which is the fact that removes any quantum method from the design critical path. No compute-cost figures are part of this study.

§ 4

Quantum machine-learning surrogates

The quantum-kernel transport surrogate was built and benchmarked across encoding width on Kronos transport- classification data (serial BR-SX-08). The result is unambiguous and runs the wrong way for advantage: the classifier AUC *degrades* with scale, from 0.817 at four encoding qubits to 0.691 at eight (figure 1). This is the fingerprint of the exponential kernel concentration of equation (4): as the feature Hilbert space grows, the off-diagonal Gram entries collapse toward a constant (figure 2) and the classifier loses discriminating power. A classical surrogate on the same features does not share this pathology; the quantum encoding is actively harmful here, not merely unhelpful.

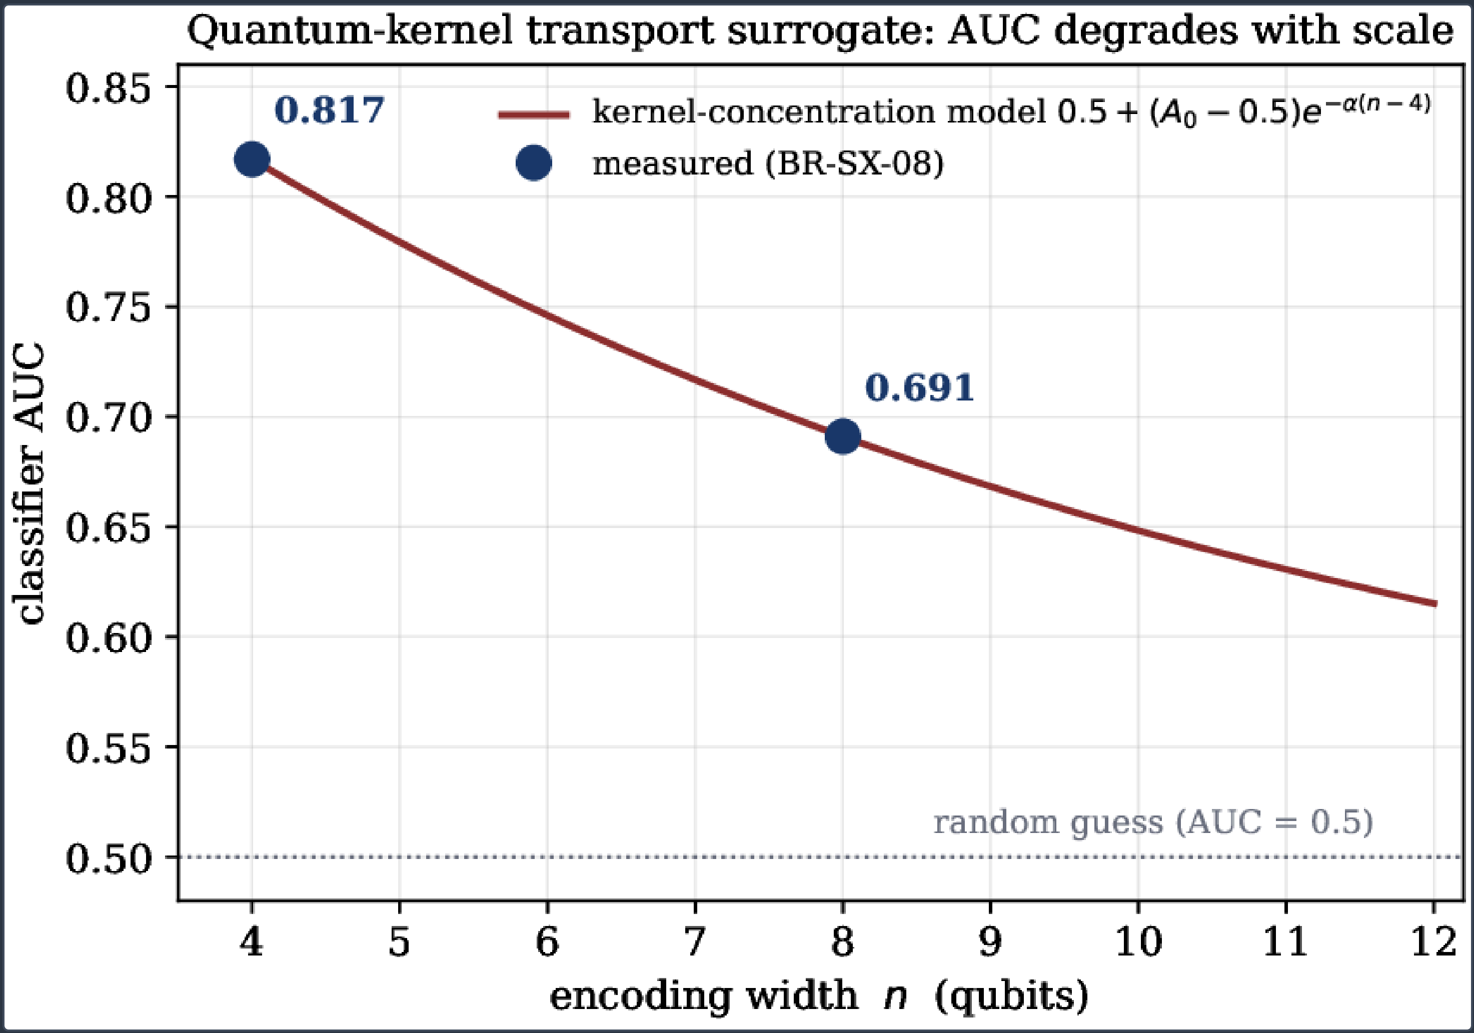

Quantum-kernel transport surrogate: classifier AUC falls from 0.817 at four encoding qubits to 0.691 at eight (measured, serial BR-SX-08), tracking the exponential-concentration envelope of equation (4). The dotted line is the random-guess floor. The trend is the diagnostic: a useful surrogate would rise, not fall, with capacity.

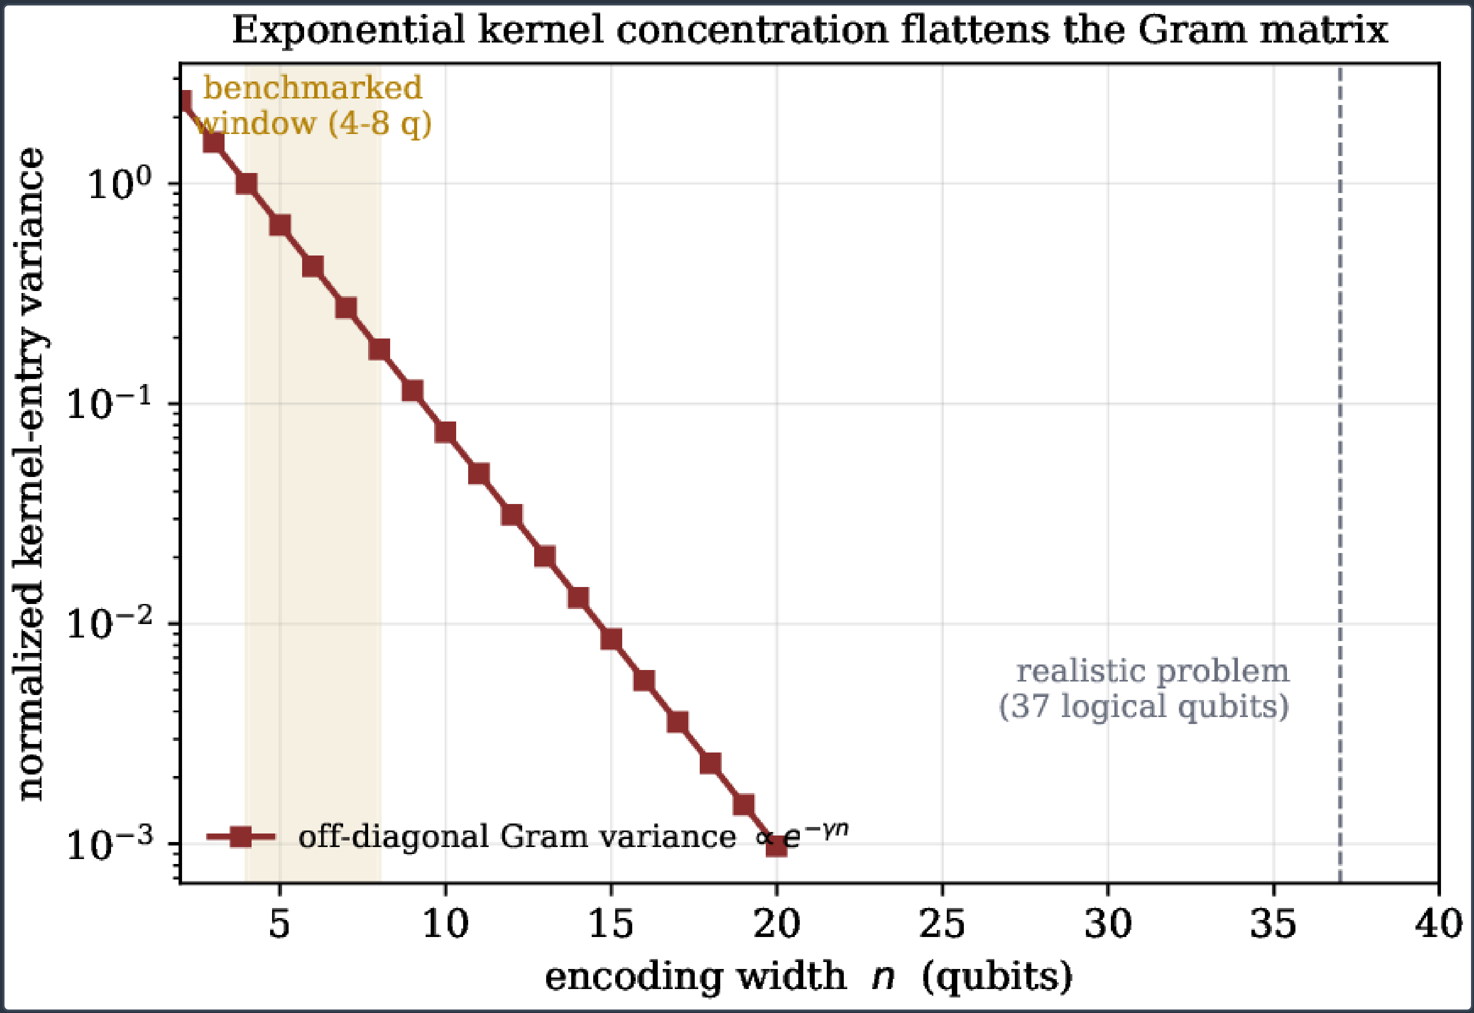

Mechanism (derived scaling). The normalised off-diagonal Gram-entry variance decays as $e^{-\gamma n}$ (equation 4); shading marks the benchmarked four-to-eight-qubit window and the dashed line the realistic $\mathbf{37}$-logical-qubit problem, where the informative spread is buried far below any tractable shot budget (equation 5).

The resource side confirms there is no route to advantage by scaling up. At a realistic problem size the quantum-kernel Gram evaluation needs 37 logical qubits and 8.5×10^3 T-gates per circuit, and across all nine benchmarked scenarios a classical Gram computation (≤ 0.09 s) beats the statevector-simulated quantum evaluation (≥ 48 s); projected to fault-tolerant hardware, where each concentrated kernel entry demands the exponential shot budget of equation (5), the quantum kernel is 9–11 orders of magnitude slower for this problem (figure 3, table 1). The surrogate is real and the benchmark is clean; its verdict is that a classical surrogate — of the neural-operator class used elsewhere in the series — is the correct engineering choice.

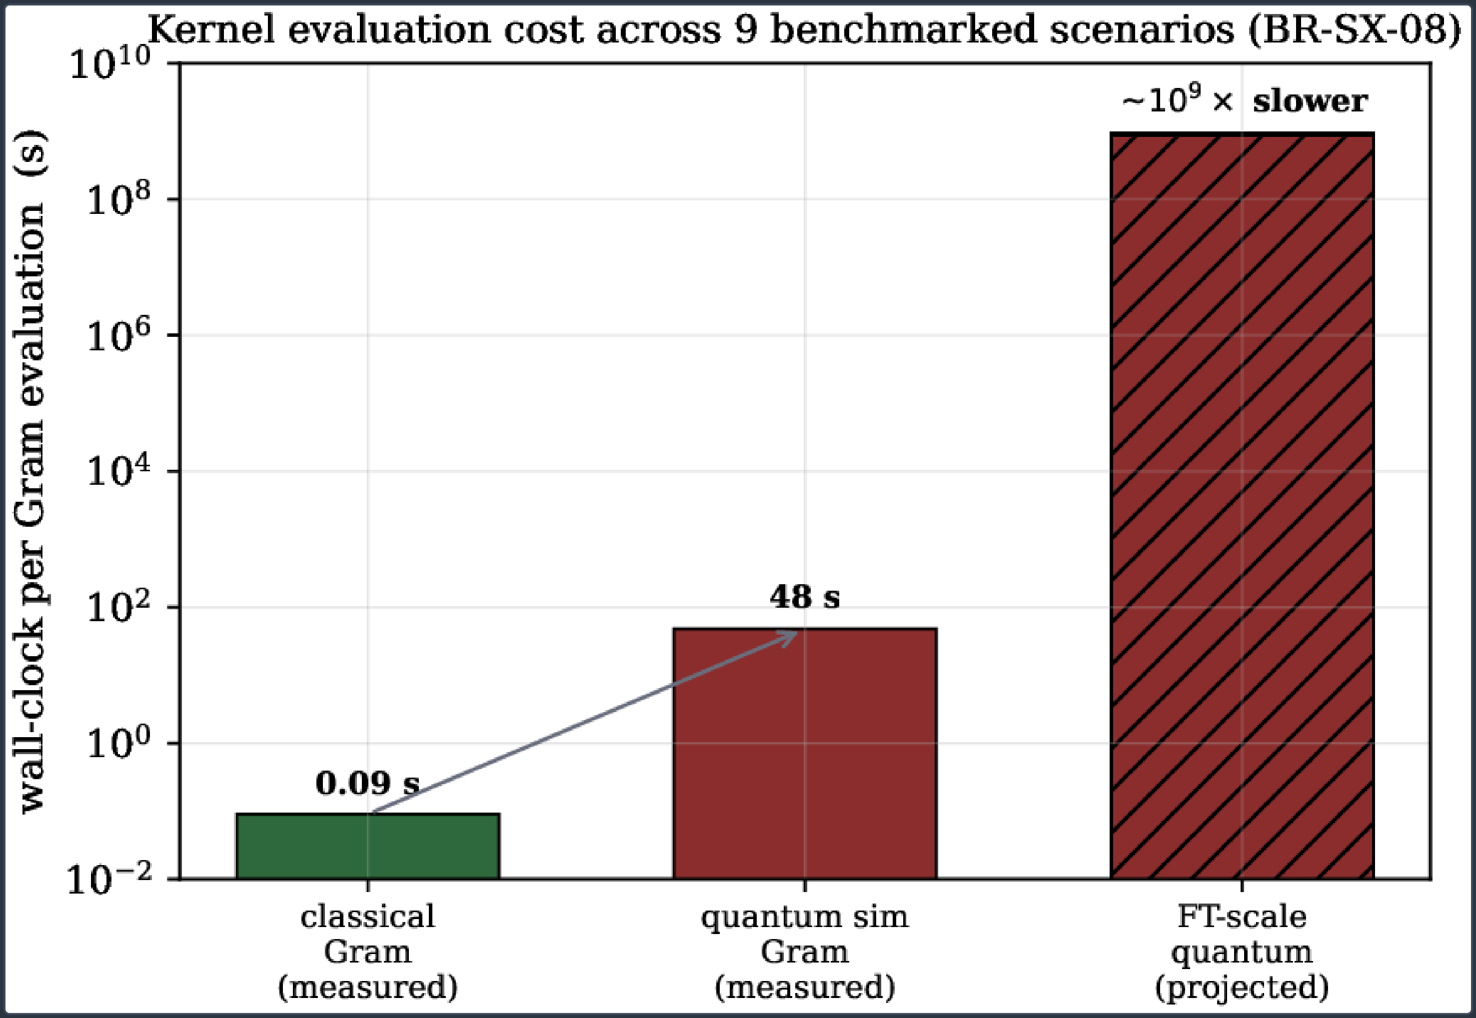

Gram-evaluation wall-clock across the nine benchmarked scenarios (serial BR-SX-08): the classical computation (≤ 0.09 s) beats the statevector-simulated quantum kernel (≥ 48 s), and the fault-tolerant projection (hatched) is 9–11 orders of magnitude slower for this problem. Logarithmic axis.

QUANTITY	VALUE	BASIS
AUC at 4 encoding qubits	0.817	measured, statevector
AUC at 8 encoding qubits	0.691	measured, statevector
classical Gram time	≤ 0.09 s	CPU/GPU baseline
quantum Gram time (sim)	≥ 48 s	statevector, cuQuantum
realistic problem width	37 logical qubits	resource estimate
T-count per circuit	8.5×10^3	resource estimate
fault-tolerant slowdown	9–11 orders	eqs. (4)–(5)

Quantum-kernel transport surrogate: measured scaling and resource footprint (serial BR-SX-08; realistic problem size from serial KX-L3-A4). AUC is the noiseless statevector value; times are per Gram evaluation on the nine benchmarked scenarios.

§ 5

Quantum-enhanced sensing for the magnets

Quantum sensing is nearer-term than quantum computing, and here the map has a genuine — if bounded — win. The requirement is set by the breeder centrepost, whose peak field is 16.84 T. This is a formally distinct device from the burner plug coil at 26.49 T; the two magnets are treated as the separate devices they are, and the sensing budgets below are the centrepost's. Absolute field monitoring at 10^{-4} relative precision is a 1.68 mT task; the quench precursor at the winding is a ~ 1 mT signal, only 5.9×10^{-5} of the peak field. Table 2 assigns sensor classes to these requirements, and figure 4 places each class against the two tasks.

SENSOR CLASS	CLASS-TYPICAL SENSITIVITY	ROLE AT THE CENTREPOST
NMR/EPR probe	0.01–1 ppm absolute	**primary absolute field monitor** ($\geq 100\times$ margin)
cryogenic Hall probe	$\sim$ 100 ppm–0.1 %	engineering redundancy
SQUID + pickup coil	$\sim$ 1	AC quench-precursor channel at standoff
NV-diamond ensemble	$\sim$ 1	quench / stray-field mapping; rad-hard, cryo-compatible
atomic OPM (SERF/scalar)	0.5–20	external error-field mapping only

Sensor-class assignment for centrepost field monitoring and quench precursors (serial KX-L1-A18; class sensitivities from the literature [9,10,16,15]).

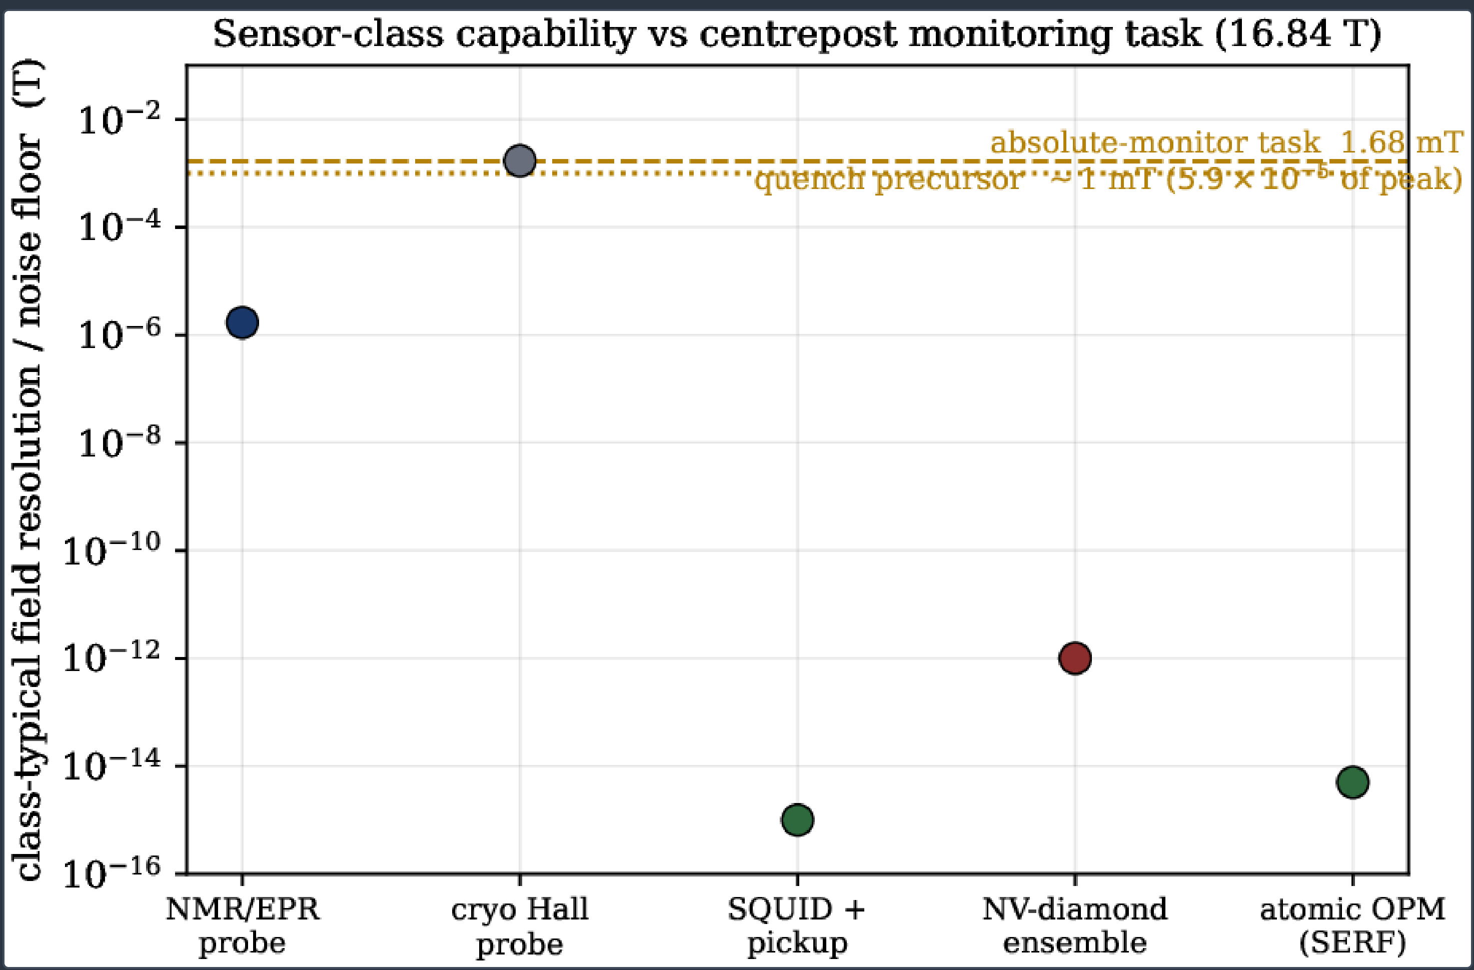

Sensor-class field resolution against the two centrepost tasks (serial KX-L1-A18): the absolute-monitor task (1.68 mT, dashed) and the quench precursor (~ 1 mT, 5.9×10^{-5} of the 16.84 T peak, dotted). Every task is comfortably above the class floors; the NMR/EPR probe carries $\geq 100\times$ margin on absolute monitoring.

The decisive quantity is where entanglement-enhanced protocols start to pay, and equation (11) answers it: the knee is at a normalised signal $\sigma^* \approx 0.14$ (figure 5). The in-bore precursor and disruption-mode signals sit *below* the knee, so a coil-plus-SQUID or conventional NV readout already suffices and entanglement adds only $+0.00077$ AUC — not worth its overhead. *Above* the knee, in deep sub-threshold error-field mapping ($\sigma \approx 0.2$), the task is genuinely sensitivity-limited and entanglement helps: by equations (8)–(9) an NV ensemble running a Heisenberg-limited protocol carries the SQL sensitivity from 5.7×10^{-7} T down to 5.7×10^{-8} T at 100 probes and $1\ \mu$s interrogation (figure 6, table 3). That is the one place quantum sensing earns its complexity, and the chain is specified to use it exactly there and nowhere else. The picotesla NV sensitivity of equation (10) is, notably, far finer than the millitesla precursor requires; its value is the margin it buys for multiplexing and for tolerating in-service degradation, the theme of the companion quench-detection paper [39].

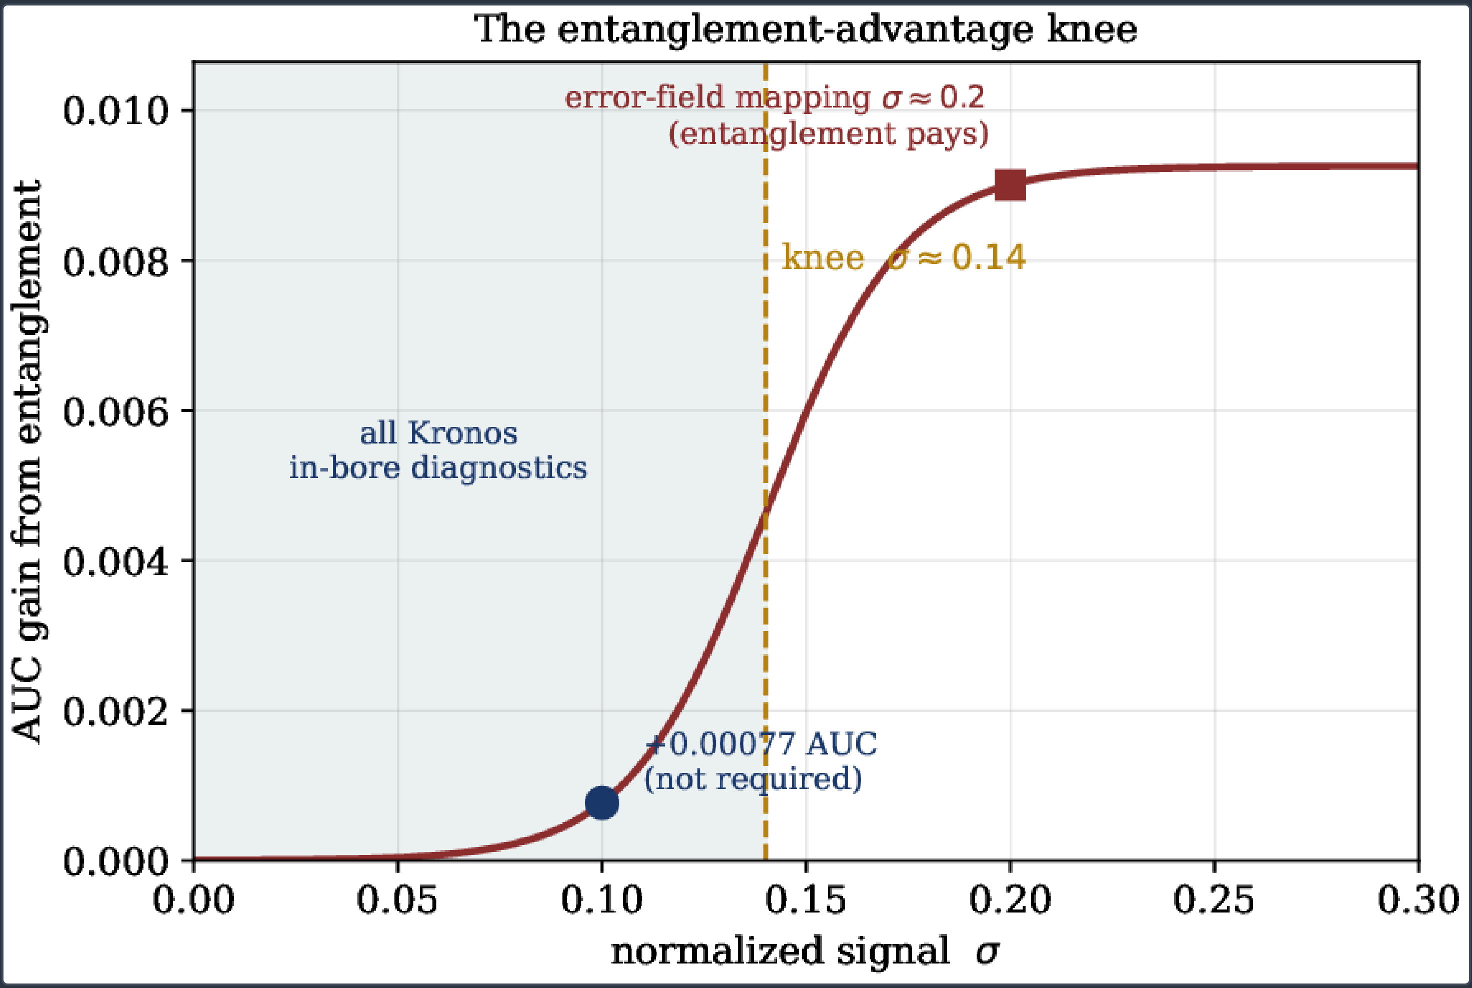

The entanglement-advantage knee (serial KX-L3-A5). The AUC gain attributable to entanglement is a threshold function of the normalised signal σ with the crossing at $\sigma^* \approx 0.14$ (equation 11). All Kronos in-bore diagnostics sit in the shaded below-knee region, where the gain is a negligible $+0.00077$ AUC; only above-knee error-field mapping ($\sigma \approx 0.2$) rewards entanglement. Model curve pinned to the two frozen anchors.

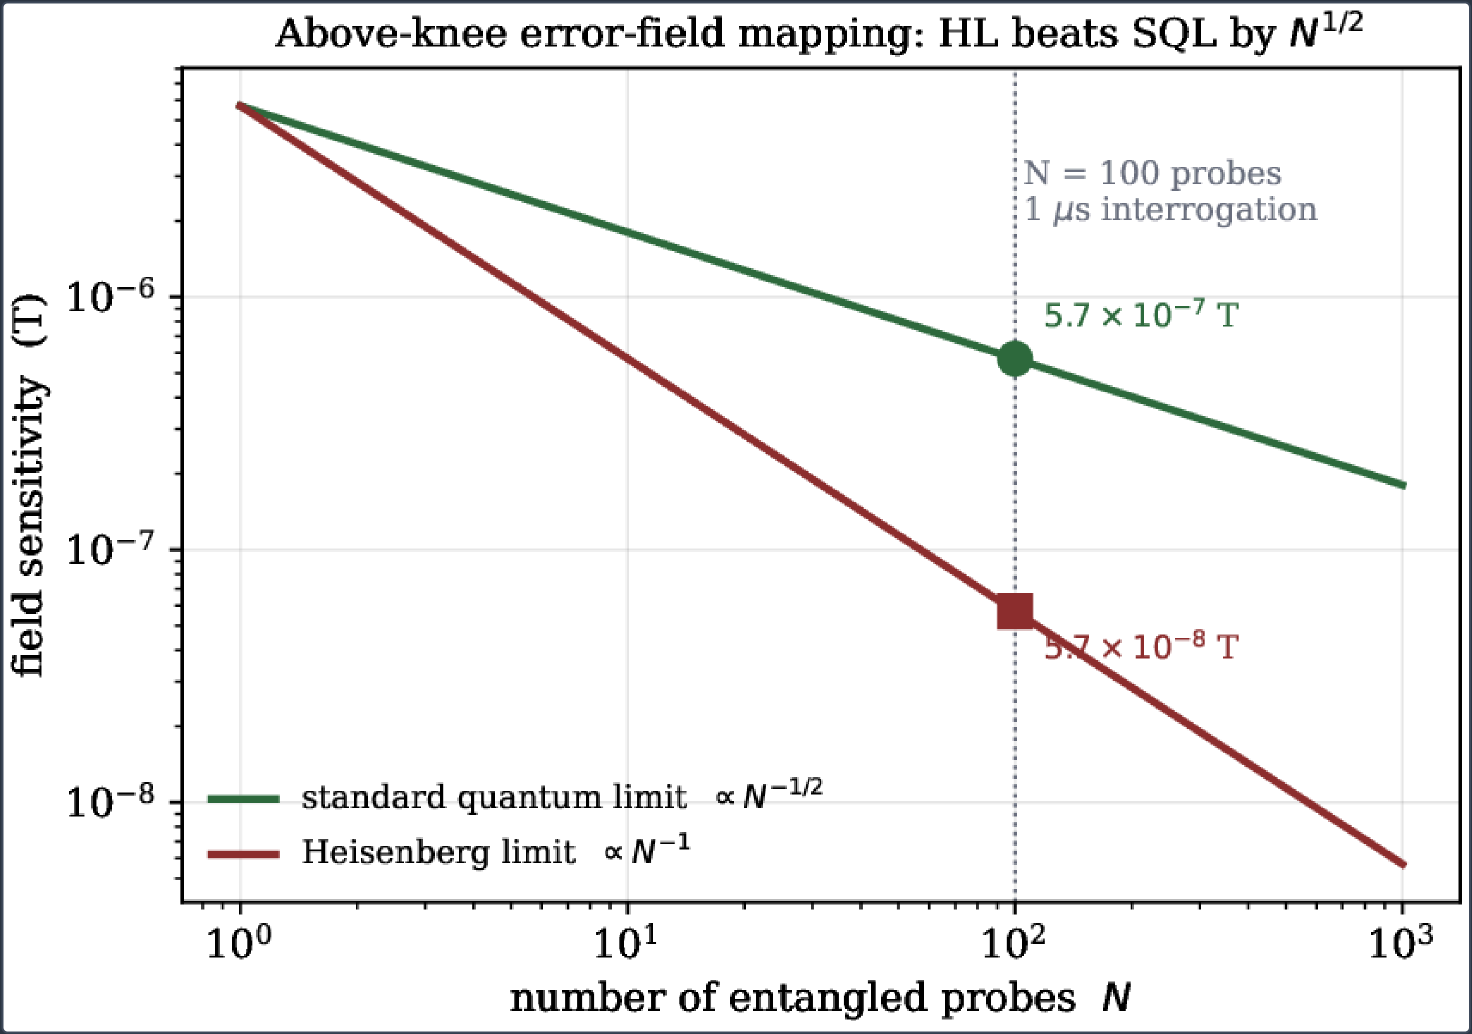

Above-knee error-field mapping (serial KX-L3-A5). The standard-quantum-limit sensitivity (equation 7, $\propto N^{-1/2}$) and the Heisenberg-limit sensitivity (equation 8, $\propto N^{-1}$) are pinned at $N = 100$ probes: 5.7×10^{-7} T and 5.7×10^{-8} T respectively, a factor $\sqrt{N} = 10$ (equation 9).

QUANTITY	VALUE	BASIS
centrepost peak field	16.84 T	design point
absolute-monitor task	1.68 mT	10^{-4} relative
quench precursor	~ 1 mT (5.9×10^{-5} of peak)	winding signal
entanglement knee	$\sigma^* \approx 0.14$	eq. (11)
below-knee AUC gain	+0.00077 (negligible)	in-bore diagnostics
SQL sensitivity ($N = 100$)	5.7×10^{-7} T	eq. (7)
Heisenberg sensitivity ($N = 100$)	5.7×10^{-8} T	eq. (8)

Entanglement-advantage accounting for centrepost sensing (serials KX-L1-A18, KX-L3-A5). The knee separates the below-knee regime, where every in-bore Kronos diagnostic sits, from the above-knee regime where a Heisenberg-limited protocol pays.

§ 6

Fault-tolerant resource estimates and the dated roadmap

The final front is the eventual quantum solution of fusion-relevant equations. We recomputed the fault-tolerant resource ledger from the first formulas of §2 — the Grover iteration count of equation (12) and the surface-code footprint of equation (13) — across 50 rows and reproduced every value with zero discrepancy (serial KX-L3-A6). Table 4 states the thresholds by workload, figure 7 places them in the resource plane, and figure 8 dates them against public hardware roadmaps.

WORKLOAD	LOGICAL QUBITS	T-COUNT	CLASS / WINDOW
linear Vlasov kinetics	11–19	2×10^5–1.5×10^6	early-FT, $\sim 2029 \pm 2$
QAOA design-scale QUBO	23–61	5×10^3–5.5×10^7	early-FT (classically trivial)
QML kernel (realistic)	37	8.5×10^3/circuit	FT, no advantage
nonlinear 5-D gyrokinetics	80–160	10^9–10^{13}	FT-horizon, post-2035
fleet Grover QUBO	156	1.2×10^{23}	never (1.8×10^{18} iters)

Fault-tolerant resource thresholds by workload and the earliest credible window (serial KX-L3-A6; logical-qubit and T-counts verified). The timeline (serial KX-L3-A7) is synthesized from public vendor roadmaps, which are references, not commitments.

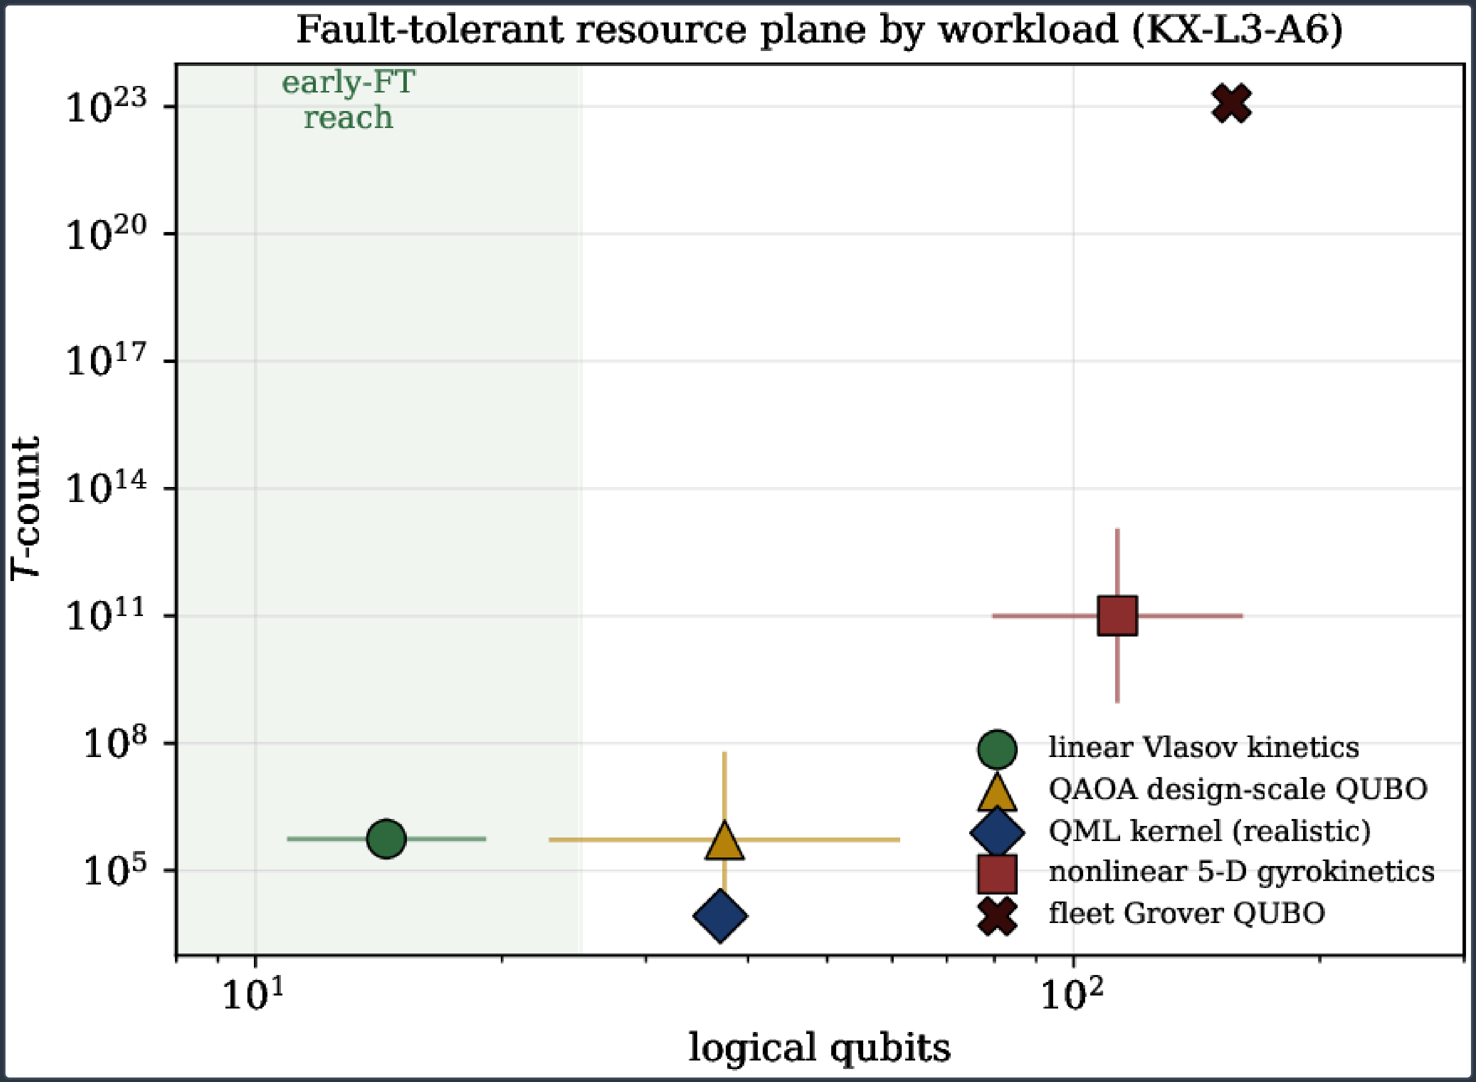

Fault-tolerant resource plane (serial KX-L3-A6): logical-qubit width against T-count for the five benchmarked workloads. Bars span the tabulated ranges. Only the linear-Vlasov and QAOA workloads fall within early-fault-tolerant reach (shaded); nonlinear gyrokinetics and fleet Grover sit many orders of magnitude beyond it.

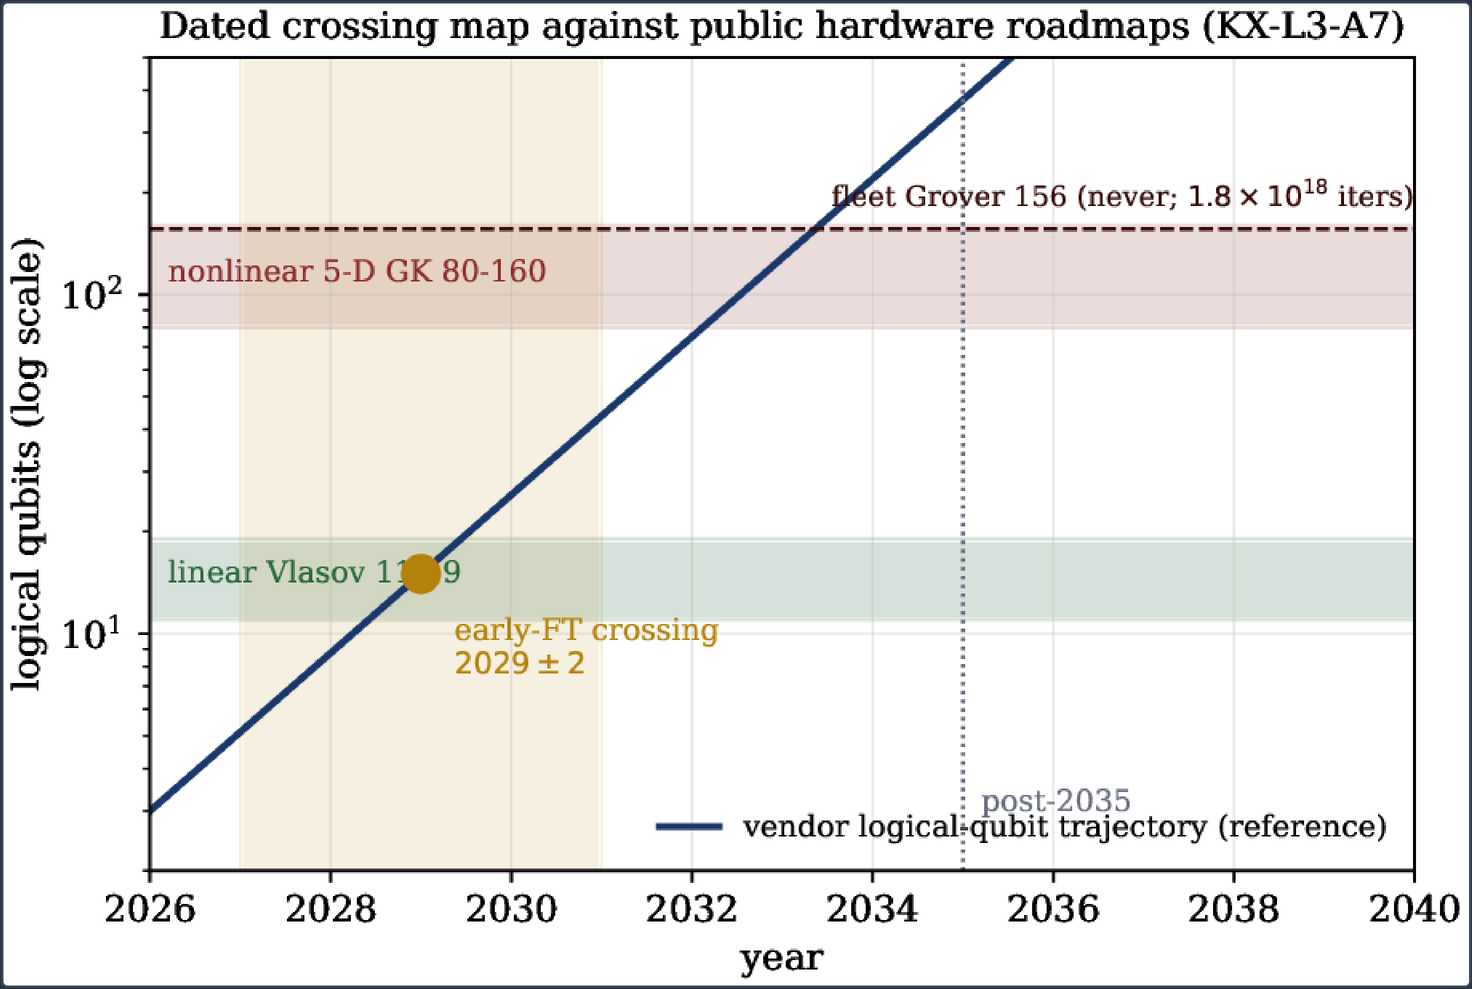

Dated quantum-technology roadmap (serial KX-L3-A7). The vendor logical-qubit trajectory (reference) crosses the only near-term fusion-relevant threshold — linear Vlasov kinetics at 11–19 logical qubits — around 2029 ± 2. Nonlinear five-dimensional gyrokinetics (80–160 logical) lies beyond 2035, and fleet-scale Grover search (1.8×10^{18} iterations) is never practical.

The map has exactly one dated, real opportunity: linear kinetic response on early-fault-tolerant machines around 2029 ± 2, an 11–19 logical-qubit workload that public roadmaps place within a credible window and that a Vlasov Hamiltonian-simulation encoding makes concrete [26]. Nonlinear five-dimensional gyrokinetic turbulence — the physics that actually sets confinement — remains a fault-tolerant-horizon problem beyond 2035 at 10^9–10^{13} T-gates, and fleet-scale Grover search stays prohibitive at any date: by equation (12) it demands 1.8×10^{18} iterations (figure 9), a textbook instance of Babbush *et al.*'s rule that quadratic speed-ups do not survive error-correction overheads [21]. Crucially, none of this is on the critical path: classical baselines of the CGYRO/GENE class [28,29] run the governing physics today, and a tensor-network kinetic solver offers a further classical route [40], so the quantum opportunity is a validation-and-benchmark niche, not a design dependency.

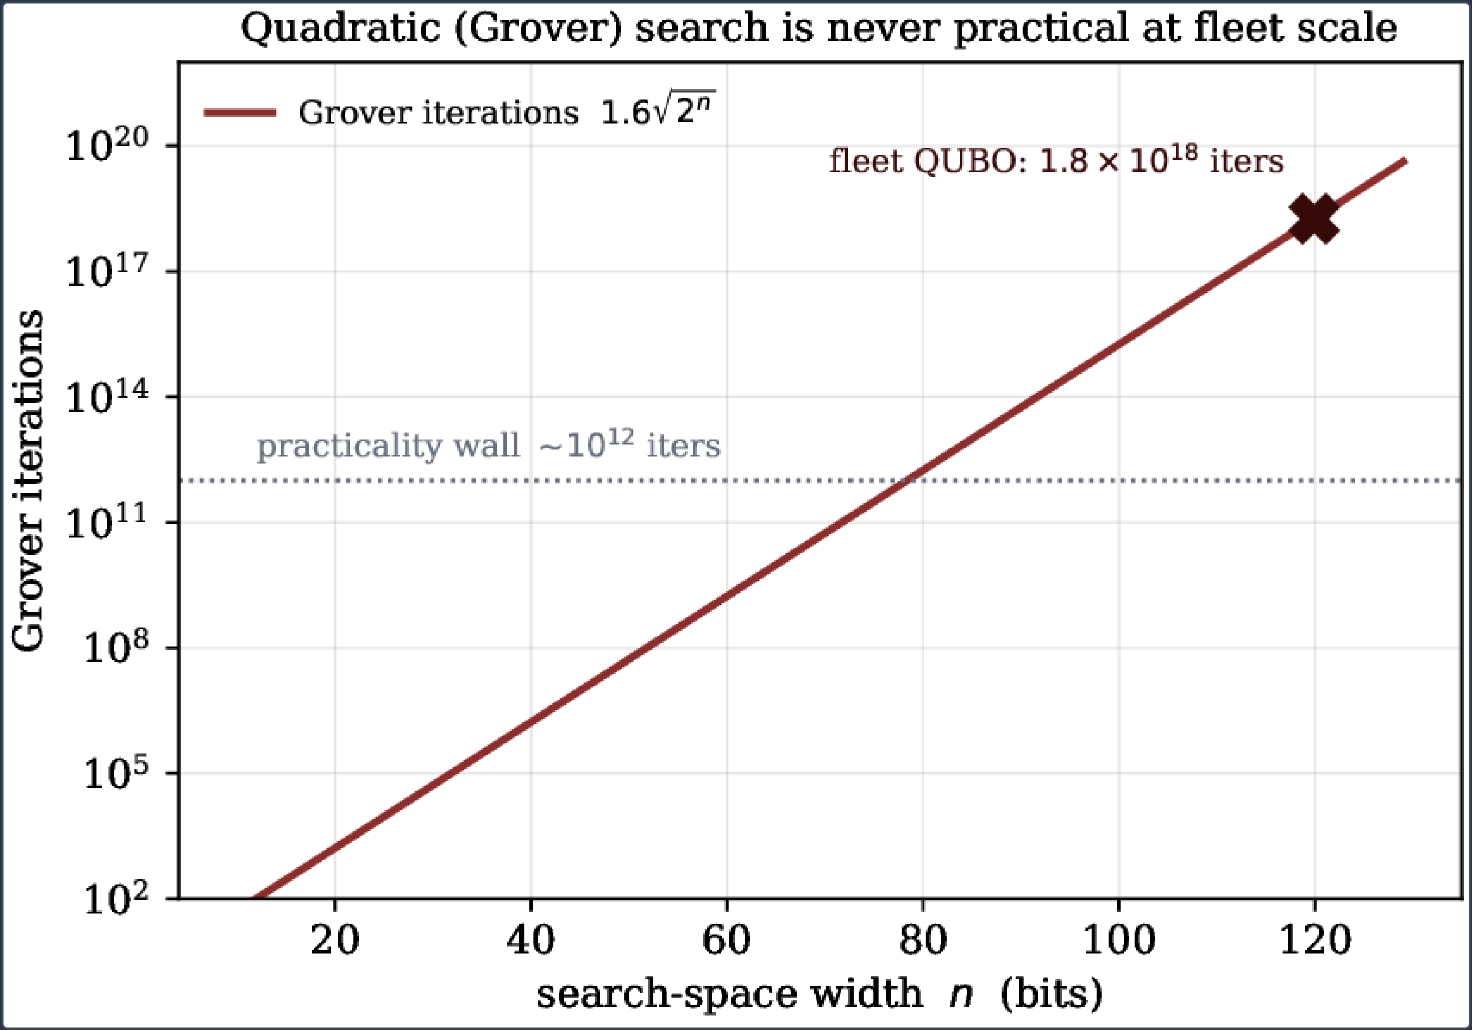

Grover iteration count $N_{\text{iter}} = 1.6\sqrt{2^n}$ (equation 12) against search-space width (serial KX-L3-A6). The fleet-scale QUBO sits at 1.8×10^{18} iterations, far above any practicality wall; a quadratic speed-up is worthless at this scale.

Figure 10 summarises the study: three testable fronts, one GPU-HPC campaign, and a dated decision on each.

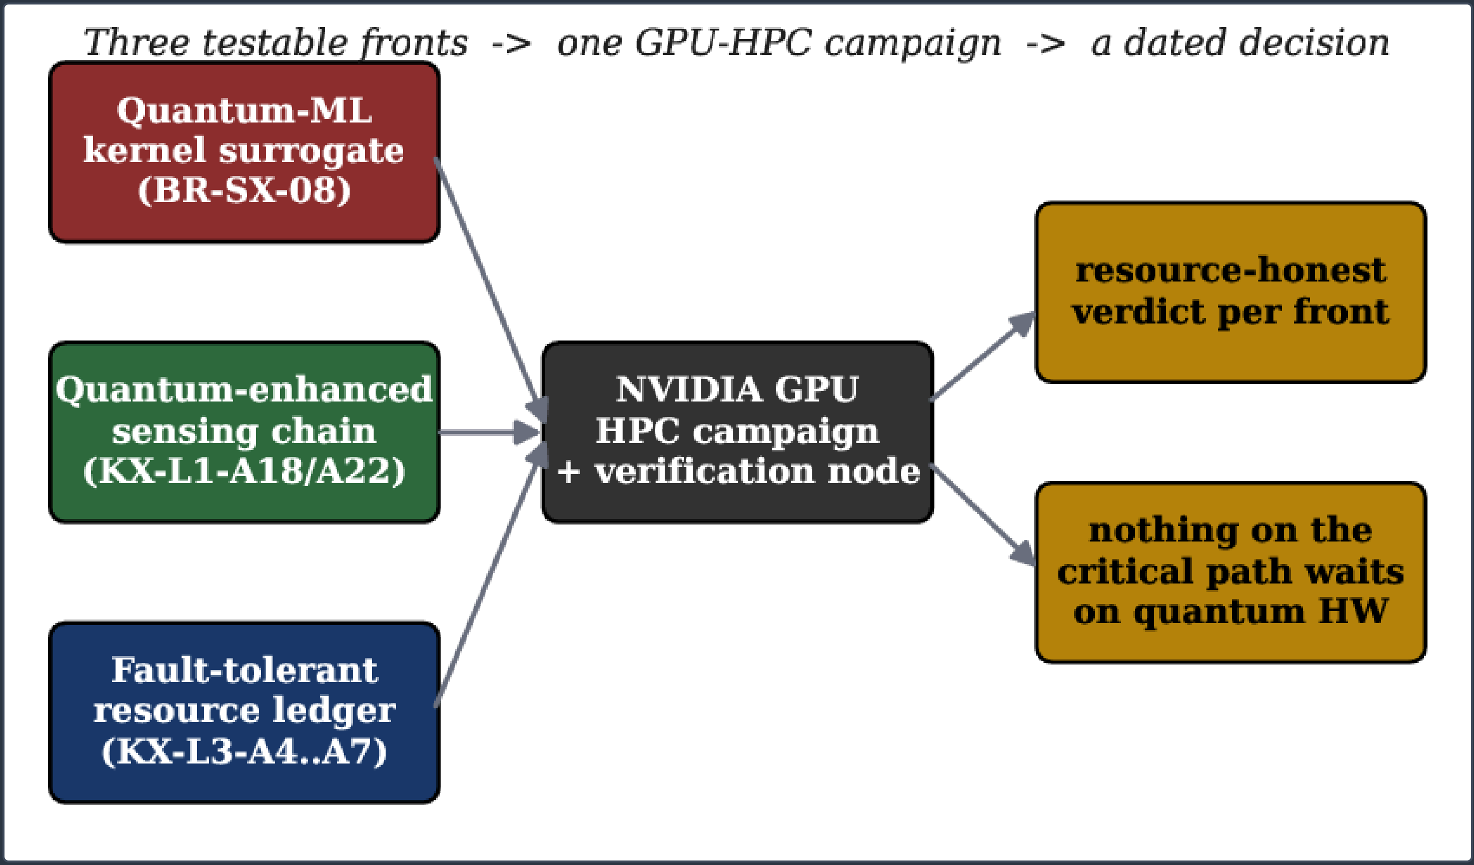

(Schematic.) Study architecture. The quantum-ML surrogate, the quantum-sensing chain, and the fault-tolerant resource ledger are each computed on the NVIDIA GPU HPC campaign and a verification node, and each returns a resource-honest verdict; the common conclusion is that nothing on the design critical path waits on quantum hardware.

§ 7

Discussion

Across all three fronts the study delivers the same kind of value — a decision backed by numbers — and each number can be read against the published literature.

QML IN CONTEXT

The measured AUC degradation ($0.817 \to 0.691$, $4 \to 8$ qubits) is not an implementation defect but a direct observation of the exponential kernel concentration proved by Thanasilp *et al.* [6] and consistent with the barren-plateau phenomenology of McClean *et al.* [7]. It is the data-side counterpart of Huang *et al.*'s finding that classical models learning from data are hard to beat on generic tasks [4]. The provable separations that do exist, such as Liu *et al.*'s discrete-log construction [5], require adversarially structured data that transport classification does not supply; our negative result is therefore exactly the expected outcome for a natural-data problem, quantified on the Kronos features. This aligns with, and sharpens, the series′ broader resource-honest computing assessment [34,35]: a quantum kernel is the wrong tool for this surrogate, and the neural-operator surrogate used elsewhere in the series is the right one.

SENSING AGAINST THE QUANTUM-METROLOGY LITERATURE

The entanglement accounting rests on textbook bounds — the SQL and Heisenberg scalings of equations (7)–(8) [14,13] — and the sensor classes on mature device physics: ensemble NV magnetometry at $\sim$ [10,11,12], SQUIDs at the femtotesla floor [16], and optical magnetometers likewise [15]. Our contribution is not a new device but a decision rule: the knee at $\sigma^* \approx 0.14$ (equation 11) partitions the Kronos diagnostic set, placing every in-bore channel below the knee where entanglement is superfluous ($+0.00077$ AUC) and reserving Heisenberg-limited protocols for the above-knee error-field-mapping task, where the $\sqrt{N}$ gain of equation (9) is genuine. The burner counterpart of this analysis, for the 26.49 T plug field, is carried in the companion burner quantum-stack paper [38], and the NV-diamond quench-precursor lineage in the fused-detection paper [39].

RESOURCE ESTIMATES AGAINST THE FTQC LITERATURE

The ledger recomputation reproduces the surface-code cost model of Fowler *et al.* [19] and sits in the same estimation tradition as Gidney and Eker [22] and Beverland *et al.* [23]. Its dated conclusions are conservative and specific: the early-fault-tolerant linear-Vlasov niche at $11-19$ logical qubits around 2029 ± 2 matches the resource scale of the Engel–Smith–Parker Vlasov algorithm [26] and the broader plasma-quantum-computing survey of Dodin and Startsev [27]; the nonlinear-gyrokinetics horizon beyond 2035 reflects the 10^9–10^{13} T-gate cost; and the fleet-Grover verdict of “never” is the direct application of the quadratic-speed-up rule of Babbush *et al.* [21] through equation (12). The QAOA/QUBO design row, classically trivial at these sizes, is analysed in full in the companion quantum-optimisation paper [36], and the linear-kinetics control-loop algorithms in [37]. The through-line is that quantum computing enters the Kronos program as a benchmark and validation resource on a dated horizon, never as a dependency of the frozen design point.

§ 8

Limitations

The single open item is temporal rather than technical: the near-term quantum-computing niche — linear kinetic response on early-fault-tolerant hardware around 2029 ± 2 — depends on vendor logical-qubit milestones that are referenced from public roadmaps and carry historical slippage risk. The date is therefore a reference window, not a Kronos commitment, and is re-checked at each annual review; because no design decision depends on it, slippage changes only the benchmark schedule, not the machine.

§ 9

Conclusion

We map, quantitatively, where quantum computing and sensing help a compact fusion program and where they do not. The quantum-kernel surrogate degrades with scale (AUC $\mathbf{0.817} \rightarrow \mathbf{0.691}$, $\mathbf{4} \rightarrow \mathbf{8}$ qubits) through the exponential kernel concentration of equation (4) and is $\mathbf{9}$–$\mathbf{11}$ orders slower than classical at fault-tolerant scale; the sensing chain for the $\mathbf{16.84}$ T centrepost uses NMR/EPR, SQUID and NV-diamond channels, with entanglement reserved for above-knee error-field mapping ($\sigma^* \approx \mathbf{0.14}$, $\mathbf{5.7} \times \mathbf{10^{-8}}$ T at $\mathbf{100}$ probes) and never conflated with the distinct $\mathbf{26.49}$ T burner plug; and the verified $\mathbf{50}$-row fault-tolerant ledger dates one near-term niche — linear kinetics around $\mathbf{2029} \pm \mathbf{2}$ — with nonlinear gyrokinetics beyond $\mathbf{2035}$ and fleet Grover never. Nothing on the design critical path waits on quantum hardware, and establishing that precisely, with the mathematics and the numbers behind it, is the result.

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.

Reproducibility. The quantum-ML scaling result, the sensing specification, and the fault-tolerant thresholds are frozen entries (serials BR-SX-08, KX-L1-A18, KX-L1-A22, KX-L3-A4, KX-L3-A5, KX-L3-A6, KX-L3-A7) in the Kronos Physics De-Risking Register [33]; the $\mathbf{50}$-row resource recomputation reproduced every value from the stated formulas with zero failures, and the statevector kernel benchmark and analytic sensing model are regenerable from the figure-generation script deposited with this paper. This paper is archived at its DOI doi:10.5281/zenodo.22645905 and at its official page https://www.kronosfusionenergy.com/publications/quantum-machine-learning-surrogates-entanglement-enhanc.html; the register is archived at doi:10.5281/zenodo.22133057.

Data availability tocsectionData availability The quantum-kernel benchmark, sensor-requirement tables, entanglement-knee sweep, and resource-recomputation ledger used for the figures are archived with the deposit at doi:10.5281/zenodo.22645905, which references the Kronos 2026 series including the AI-and-quantum-computing paper [34]. 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 J Biamonte, P Wittek, N Pancotti, P Rebentrost, N Wiebe and S Lloyd, Quantum machine learning, *Nature* **549** 195 (2017). doi:10.1038/nature23474

2 V Havl\'icek, A D C\'orcoles, K Temme, A W Harrow, A Kandala, J M Chow and J M Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567** 209 (2019). doi:10.1038/s41586-019-0980-2

3 M Schuld and N Killoran, Quantum machine learning in feature Hilbert spaces, *Phys. Rev. Lett.* **122** 040504 (2019). doi:10.1103/PhysRevLett.122.040504

4 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

5 Y Liu, S Arunachalam and K Temme, A rigorous and robust quantum speed-up in supervised machine learning, *Nat. Phys.* **17** 1013 (2021). doi:10.1038/s41567-021-01287-z

6 S Thanasilp, S Wang, M Cerezo and Z Holmes, Exponential concentration in quantum kernel methods, *Nat. Commun.* **15** 5200 (2024). doi:10.1038/s41467-024-49287-w

7 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

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

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

10 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

11 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

12 T Wolf, P Neumann, K Nakamura, H Sumiya, T Ohshima, J Isoya and J Wrachtrup, Subpicotesla diamond magnetometry, *Phys. Rev. X* **5** 041001 (2015). doi:10.1103/PhysRevX.5.041001

13 V Giovannetti, S Lloyd and L Maccone, Advances in quantum metrology, *Nat. Photon.* **5** 222 (2011). doi:10.1038/nphoton.2011.35

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

15 D Budker and M Romalis, Optical magnetometry, *Nat. Phys.* **3** 227 (2007). doi:10.1038/nphys566

16 R L Fagaly, Superconducting quantum interference device instruments and applications, *Rev. Sci. Instrum.* **77** 101101 (2006). doi:10.1063/1.2354545

17 L K Grover, A fast quantum mechanical algorithm for database search, *Proc. 28th Annual ACM Symp. on Theory of Computing (STOC)* 212 (1996). arXiv:quant-ph/9605043

18 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

19 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

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

21 R Babbush, J R McClean, M Newman, C Gidney, S Boixo and H Neven, Focus beyond quadratic speedups for error-corrected quantum advantage, *PRX Quantum* **2** 010103 (2021). doi:10.1103/PRXQuantum.2.010103

22 C Gidney and M Eker, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, *Quantum* **5** 433 (2021). doi:10.22331/q-2021-04-15-433

23 M E Beverland *et al*, Assessing requirements to scale to practical quantum advantage, arXiv:2211.07629 (2022).

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

Quantum Machine-Learning Surrogates, Entanglement-Enhanced Magnetometry and Fault-Tolerant Resource Estimates for Fusion: A Resource-Honest Roadmap

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

Zenodo · doi:10.5281/zenodo.22645905

Read on kronosfusionenergy.com

The Physics De-Risking Register