KRONOS FUSION ENERGY

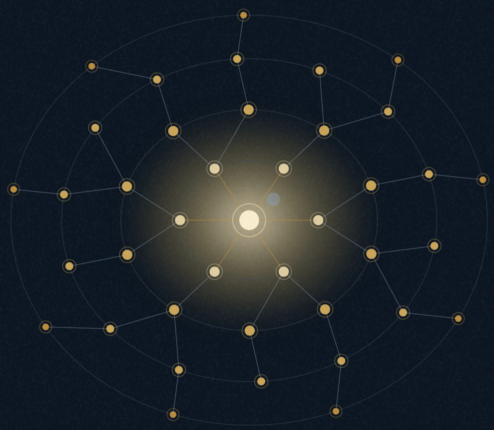

PAPER 2.93 · THE KRONOS 2026 SERIES

Quantum Algorithms for the Real-Time Control Loop: HHL Equilibrium Reconstruction, Disruption State-Estimation, and Edge-Turbulence Simulation with Resource-Honest Advantage Accounting

The real-time control loop of a compact spherical-tokamak breeder must reconstruct the equilibrium, watch for disruptions, and shape the plasma edge faster than a plasma at can…

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

BY P I FORDOFFICIAL VERSION: HTTPS://WWW.KRONOSFUSIONENERGY.COM/PUBLICATIONS/QUANTUM-ALGORITHMS-FOR-THE-REAL-TIME-CONTROL-LOOP-HHL-E.HTML

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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 Algorithms for the Real-Time Control Loop: HHL Equilibrium Reconstruction, Disruption State-Estimation, and Edge-Turbulence Simulation with Resource-Honest Advantage Accounting

The real-time control loop of a compact spherical-tokamak breeder must reconstruct the equilibrium, watch for disruptions, and shape the plasma edge faster than a $9.66\,MA$ plasma at $R_0 = 1.20\,m$ can move — a loop served today by classical GPU surrogates and a demonstrated in-winding quantum sensor. This paper asks, and answers with equal weight on both sides, where a quantum *computer* (as distinct from a quantum sensor) could accelerate that loop. We formalise a stack of four control-loop quantum algorithms as rigorous problems: an HHL-class linear solver for live Grad–Shafranov equilibrium reconstruction; a quantum-kernel state-estimator for disruption and runaway-electron precursor detection inside the shattered-pellet mitigation window; a Hamiltonian simulation of edge and scrape-off-layer turbulence to feed a predictive detachment controller; and a quantum-machine-learning surrogate for live transport prediction. For each we derive the governing operator, its discretization, and its fault-tolerant complexity, and we place them against the classical frontier through an explicit *advantage inequality*. Alongside the designs we report a completed characterization: the quantum-ML classifier *loses* discriminating power as it scales — area under the receiver-operating-characteristic curve falls from 0.817 at four qubits to 0.691 at eight, the calibrated signature of exponential kernel concentration — and a fault-tolerant kernel evaluation is 9–11 orders of magnitude slower than the classical route for this problem. Fault-tolerant resource estimates for the solver-class algorithms (edge/scrape-off-layer kinetics: ~ 44 logical qubits and ~ 19 days of single-instance runtime; linear kinetics: 11–19 logical qubits) confirm that none clears the advantage bar today. The honest conclusion is stated once: **there is no measured quantum advantage on these control problems today**. The deliverable is a validated set of algorithm designs, a resource-honest map of where a fault-tolerant machine will and provably will not help, and a dated roadmap that keeps the classical loop authoritative until a quantum route clears an explicit latency-and-fidelity bar. All numbers trace to frozen entries in the Kronos Physics De-Risking Register by their serial identifiers.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645925 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: quantum algorithms · HHL linear solver · quantum state estimation · Hamiltonian simulation · quantum machine learning · kernel concentration · fault-tolerant resource estimation · real-time control · spherical-tokamak breeder

CONTENTS

Table of Contents

Quantum Algorithms for the Real-Time Control Loop: HHL Equilibrium Reconstruction, Disruption State-Estimation, and Edge-Turbulence Simulation with Resource-Honest Advantage Accounting 3
Introduction 6
The designed control-loop quantum stack 7
Governing equations and the formal problem 8
Computational methodology: the NVIDIA GPU HPC campaign 11
Results 12
Discussion 19
Limitations 20
Conclusion 21
References 26
One programme, many papers 31

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 COMPACT-ST CONTROL PROBLEM AND ITS TIMESCALES

A spherical tokamak (ST) confines a large plasma current in a small volume, which is exactly what makes it an attractive neutron source and breeder platform [1,2] and, at the same time, what makes it demanding to control. The Hyperion breeder holds a plasma current $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$, aspect ratio $A = 2.5$, elongation $\kappa = 2.0$, on-axis field $B_0 = 8\,T$ behind a centre-post magnet at a peak field of $16.84\,T$, in a negative-triangularity ($\delta = -0.30$) regime [3], at a fusion power of $P_{\text{fus}} = 85.04\,MW$ and gain $Q = 3.076$. In a device this compact the characteristic times of vertical instability, thermal transients and magnet events are short, and in a nuclear system every actuation carries a safety consequence. The control loop must therefore reconstruct the equilibrium, forecast disruptions, and hold the quiet negative-triangularity edge inside a narrow window, all faster than the plasma evolves.

Feedback control of tokamak equilibria and profiles is a mature discipline: magnetic diagnostics feed real-time equilibrium reconstruction [4] and physics-based profile simulators run inside the control cycle [5]; more recently, deep reinforcement learning has commanded coil voltages directly from magnetics [6] and deep networks have forecast disruptions across machines [7]. The common thread is that control quality is bounded by the speed and fidelity of the underlying plasma model. That model runs today on the classical GPU high-performance-computing (HPC) campaign — learned flux and equilibrium operators [8,9] inside a four-twin shadow architecture [39] — and on a demonstrated in-winding quantum sensor [42].

QUANTUM COMPUTING VERSUS QUANTUM SENSING IN THE LOOP

The productive near-term quantum contribution to this loop is already identified and demonstrated, and it is a *sensor*, not a solver. An in-winding nitrogen-vacancy (NV) diamond magnetometer resolves a magnet-quench precursor at a sensitivity of $0.189\,\mathrm{pT\,Hz^{-1/2}}$ over a 125 Hertz$bandwidth, catchingthe$\sim100 mT$conductor - scale fieldstepatapaper - shotsignal - to - noiseratiooforder$10^9\$ (register serial KX-L1-A22, reproduced) [42,12,13]. This paper concerns the distinct and harder question: where could a quantum *computer* accelerate the loop, and where could it not? We answer it in the only way useful to a plant designer — by presenting the algorithm designs specified for the loop as rigorous problems, and by reporting what a direct characterization of the most control-relevant of them actually found. The result is deliberately two-sided: the designs are real and the sensing technology is demonstrated, while the general-purpose quantum-machine-learning route, when measured, shows no advantage on this class of problem. Reporting both is the point.

PRIOR ART

The four algorithms rest on established quantum primitives. Linear-system solving descends from the Harrow–Hassidim–Lloyd (HHL) algorithm [14] and its precision-improved successors [15,16], whose caveats were sharpened by Aaronson's "fine print" [17]. Hamiltonian simulation descends from Lloyd [18] through qubitization [19] and truncated-Taylor / linear-combination-of-unitaries methods [20]; a broad survey of end-to-end algorithm complexities is given in [21]. Quantum machine learning in feature Hilbert spaces [22] and the quantum-kernel classifier [23,24] are now understood to be limited by two structural obstacles: the *barren-plateau* vanishing of trainable gradients [27,28] and, for kernels specifically, *exponential concentration* of kernel values [26], against a backdrop in which classical models with data can match many claimed quantum advantages [25]. On the plasma-physics side, quantum algorithms for the Vlasov equation [31], resource studies for plasma simulation [32], and the community review of quantum computing for fusion [30] frame the Hamiltonian-simulation route. Finally, the fault-tolerant cost of any of these is set by the surface code [33] and the magic-state and T-count overheads that make even quadratic speedups insufficient in the near term [34,35,36]. The Kronos AI-and-quantum programme records the underlying methods [38].

CONTRIBUTION

Prior art treats the primitives and the plasma problems separately. What is missing for a compact nuclear ST is a single, control-loop-scoped treatment that (i) writes each of the four hardest real-time tasks as a formal quantum problem with its operator, discretization and complexity; (ii) places each against the classical frontier through an explicit advantage inequality rather than an asymptotic claim; (iii) reports the one component that can be measured today, honestly, including its failure; and (iv) reduces the whole to a dated roadmap gated on acceptance. This paper contributes exactly that. We give the designed stack and its gates (§2); the governing equations and formal problems for all four algorithms, with heavy use of the Grad–Shafranov, kernel-concentration, Hamiltonian-simulation and surface-code mathematics (§3); the GPU-HPC methodology that produced every number (§4); the results, including the measured no-advantage characterization and the fault-tolerant resource estimates (§5); a quantitative comparison to the published quantum-algorithm and quantum-ML literature (§6); one honest limitation (§7); and the conclusion (§8). Every quantitative claim traces to a frozen result in the Kronos Physics De-Risking Register [37]; the serial identifiers (e.g. BR-SX-08, KX-L3-A1) are given inline so each number is auditable.

§ 2

The designed control-loop quantum stack

Four algorithms target the four hardest real-time tasks in the loop (Table 1, Figure 1). Three are solver-class and require fault-tolerant hardware; the fourth is characterizable today and is characterized in §5.

HHL equilibrium reconstruction (BR-QX-01). The live equilibrium is the solution of a sparse linear system obtained by discretizing the Grad–Shafranov operator, with the measured magnetics and coil currents as right-hand side. A quantum linear-system algorithm (QLSA) is designed to return the flux state ψ for the real-time loop.
Disruption state-estimation (BR-QX-03). A quantum-kernel state-estimator is designed to flag disruption and runaway-electron (RE) precursors within the shattered-pellet-injection (SPI) lead-time window, feeding the always-armed mitigation floor [41] rather than replacing it.
Edge-turbulence Hamiltonian simulation (BR-QX-05). A Hamiltonian simulation of edge and scrape-off-layer (SOL) turbulence is designed to feed a predictive detachment / edge-localised-mode controller, resource-scoped for fault-tolerant hardware.
Quantum-ML transport surrogate (BR-QX-06). A quantum-ML surrogate is designed for live core transport and profile prediction, with an explicit go/no-go against the characterization of §5. itemize

ALGORITHM (SERIAL)	REAL-TIME TASK	PRESENT GATE
HHL / QLSA (BR-QX-01)	live Grad–Shafranov equilibrium	fault-tolerant hardware
state estimation (BR-QX-03)	disruption / RE precursor flag	fault-tolerant hardware
Hamiltonian sim (BR-QX-05)	edge/SOL turbulence for detachment	fault-tolerant hardware
quantum-ML surrogate (BR-QX-06)	live transport / profile prediction	**characterized: no advantage**

The designed control-loop quantum stack, its target real-time task, and its present gate. Serials refer to the Kronos Physics De-Risking Register [37].

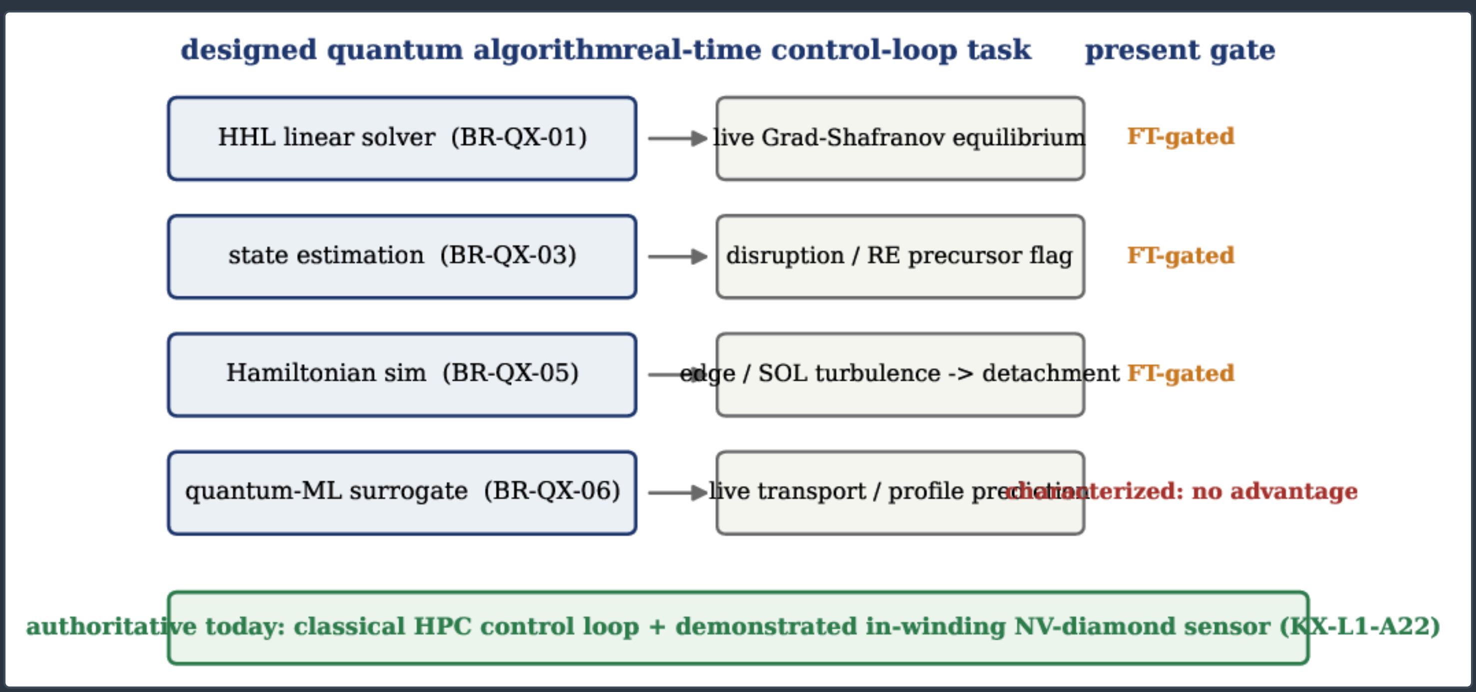

(Schematic.) The designed control-loop quantum stack. Each algorithm (left) targets one real-time task (centre) and carries a present gate (right); three are fault-tolerant-gated designs and one is characterized and shown to have no advantage. The authoritative element today is the classical HPC control loop together with the demonstrated in-winding NV-diamond sensor (KX-L1-A22). No data values are shown; the figure encodes structure only.

§ 3

Governing equations and the formal problem

We now state each algorithm as a formal problem: the continuous operator, its discretization, and its complexity against the classical frontier. Throughout, n denotes the number of logical (data) qubits, $N = 2^n$ the dimension of the encoded Hilbert space, κ the condition number of a matrix, s its sparsity (nonzeros per row), and ε a target accuracy.

HHL EQUILIBRIUM RECONSTRUCTION AS A QUANTUM LINEAR SYSTEM

Continuous operator. The axisymmetric plasma equilibrium is governed by the Grad–Shafranov (GS) equation for the poloidal flux $\psi(R, Z)$,

$$\Delta^*\psi \;\equiv\; R\,\frac{\partial}{\partial R}\!\left(\frac{1}{R}\frac{\partial\psi}{\partial R}\right) + \frac{\partial^2\psi}{\partial Z^2} \;=\; -\,\mu_0 R^2\, p'(\psi) \;-\; F(\psi)\,F'(\psi),$$

with kinetic pressure $p(\psi)$ and poloidal-current function $F(\psi) = RB_\phi$. In real-time reconstruction the right-hand-side profiles are treated by a Picard/Newton iteration; each iterate solves a linear elliptic problem for ψ with the measured magnetics and coil currents entering through the boundary condition and the free-boundary Green's-function coupling to the poloidal-field coil set [4,10,46].

Discretization. On an $N_R \times N_Z$ grid with uniform spacing h, the operator Δ^* discretizes to a sparse matrix using centred differences,

$$(\Delta^*\psi)_{i,j} \approx \frac{\psi_{i+1,j} - 2\psi_{i,j} + \psi_{i-1,j}}{h^2} - \frac{1}{R_i}\,\frac{\psi_{i+1,j} - \psi_{i-1,j}}{2h} + \frac{\psi_{i,j+1} - 2\psi_{i,j} + \psi_{i,j-1}}{h^2},$$

a five-point stencil (sparsity $s = 5$). Collecting the interior unknowns into a vector ψ and the source-plus-boundary data into b, each Picard iterate is the linear system

$$A \quad \psi = \quad b, \qquad A \in \mathbb{R}^{N_{\text{dof}}\times N_{\text{dof}}}, \quad N_{\text{dof}} = N_R N_Z.$$

The eigenvalues of the leading (symmetric) part of A inherit those of the discrete Laplacian, $\lambda_{jk} = \frac{4}{h^2}\Big[\sin^2\frac{j\pi h}{2} + \sin^2\frac{k\pi h}{2}\Big]$, so that the condition number scales as

$$\kappa(A) = \frac{\lambda_{\max}}{\lambda_{\min}} \;\sim\; \frac{1}{\sin^2(\pi h/2)} \;=\; \mathcal{O}(h^{-2}) \;=\; \mathcal{O}(N_\perp^2),$$

with $N_\perp$ the number of points per dimension (Figure 2). This $N_\perp^2$ stiffness is the single most important number for the quantum route, because HHL cost is quadratic (original) or linear (improved) in κ.

Quantum linear-system algorithm and its cost. The HHL algorithm [14] prepares the normalized state $|b\rangle = \sum_i b_i\,|i\rangle/\| \quad b\|$, applies quantum phase estimation on e^{iAt} to write the eigenphases of A into an ancilla register, performs a controlled rotation implementing $\lambda \mapsto \lambda^{-1}$, and uncomputes, yielding a state proportional to $|x\rangle = A^{-1}\,|b\rangle/\|A^{-1}\,|b\rangle\|$. Its gate complexity is

$$T_{\text{HHL}} \;=\; \tilde{\mathcal{O}}\!\left(\frac{\kappa^2 s^2 \log N_{\text{dof}}}{\varepsilon}\right),$$

improved by Childs–Kothari–Somma [15] and quantum singular-value transformation [16] to

$$T_{\text{QLSA}} \;=\; \tilde{\mathcal{O}}(\kappa\, s\, \mathrm{polylog}(N_{\text{dof}}/\varepsilon)).$$

The comparison is against the best classical route for a sparse elliptic operator: geometric multigrid solves Eq. (3) in $\mathcal{O}(N_{\text{dof}})$ work [5], and a control-latency neural-operator surrogate reconstructs ψ in sub-millisecond time [40].

The fine print. Three caveats [17] determine whether Eqs. (5)–(6) translate into a loop advantage. (i) *Readout.* HHL returns $|x\rangle$ as a quantum state; extracting the full flux map $\psi(R, Z)$ requires $\mathcal{O}(N_{\text{dof}})$ tomographic measurements, which erases the $\log N_{\text{dof}}$ advantage. The only readout that preserves it is a scalar functional $\langle x|M|x\rangle$ — e.g. a stored-energy or a shape moment — not the field the controller needs. (ii) *State preparation.* Loading $|b\rangle$ is cheap here because the magnetics vector is small, but a QRAM assumption elsewhere would dominate. (iii) *Conditioning.* The $\kappa \sim N_\perp^2$ stiffness of Eq. (4) enters quadratically in Eq. (5); even in the improved form it multiplies the polylog prefactor. The honest formal conclusion is that HHL has *no* advantage for full-field GS reconstruction, and a possible advantage only for scalar-functional queries of a well-conditioned reduced operator — a regime we scope but do not claim.

DISRUPTION STATE-ESTIMATION AS A QUANTUM-KERNEL PROBLEM

Feature map and kernel. A quantum feature map encodes a classical precursor vector x (magnetics, locked-mode amplitude, radiated-power fraction, q_{95}, internal inductance) into an n-qubit state $|\phi(x)\rangle = U_\phi(x)|0\rangle^{\otimes n}$ [22,23]. The induced kernel is the squared fidelity

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

and the disruption/RE decision rule is a support-vector machine over the Gram matrix, $f(x) = \mathrm{sgn}\big(\sum_i \alpha_i y_i\, k(x_i, x) + b\big)$, trained to fire inside the SPI mitigation lead-time [41,7].

Exponential concentration. The obstruction is structural. For expressive feature maps the kernel values concentrate exponentially in n about a fixed value μ [26],

$$\mathrm{Var}_{x,x'}\big[k(x, x')\big] \;\le\; V_0\, e^{-\alpha n}, \qquad \alpha > 0,$$

so that the off-diagonal Gram entries become statistically indistinguishable from μ and the Gram matrix tends to the identity. Resolving an informative off-diagonal element to relative precision then costs $M_{\text{shots}} \sim 1/\,\mathrm{Var} \sim e^{\alpha n}$ measurements, which destroys trainability at the very scale where the quantum feature space would be interesting. This is the same untrainability that barren plateaus impose on variational circuits [27,28], specialised to kernels, and it is precisely what the measurement of §5 exhibits: calibrating Eq. (8) to the two measured operating points gives $\alpha \simeq \mathbf{0.14}$ per qubit (Figure 4).

EDGE-TURBULENCE HAMILTONIAN SIMULATION

From kinetics to a Hamiltonian. The edge/SOL turbulence that sets the detachment window is described by a gyrokinetic/Vlasov evolution of the perturbed distribution f,

$$\frac{\partial f}{\partial t} = \mathcal{L}\, f,$$

with $\mathcal{L}$ the (linearised) Vlasov operator built from streaming, the $E \times B$ drift and the self-consistent field response [11]. Following Engel–Smith–Parker [31], the linearised system is cast into Schr\"odinger form by splitting $\mathcal{L} = \mathcal{L}_{\mathrm{H}} + \mathcal{L}_{\mathrm{A}}$ into anti-Hermitian and Hermitian parts and simulating the unitary generated by the Hermitised operator $H = i\mathcal{L}_{\mathrm{H}}$,

$$i\,\frac{\partial}{\partial t}|\Psi\rangle = H\,|\Psi\rangle, \qquad |\Psi(t)\rangle = e^{-iHt}\,|\Psi(0)\rangle,$$

where a phase-space grid of M points is amplitude-encoded into $n = \lceil \log_2 M \rceil$ qubits — the source of the exponential state compression.

Trotter error and query complexity. Writing $H = \sum_{j=1}^{L} H_j$ as a sum of efficiently simulable terms, the first-order product formula obeys the operator bound

$$\Big\| e^{-iHt} - \big(\textstyle\prod_j e^{-iH_j t/r}\big)^r \Big\| \;\le\; \tfrac{t^2}{2r} \textstyle\sum_{j<k} \big\|[H_j, H_k]\big\| \;\equiv\; \tfrac{\Lambda t^2}{2r},$$

so that $r = \mathcal{O}(\Lambda t^2/\varepsilon)$ Trotter steps reach accuracy ε; the commutator sum Λ is the physically meaningful cost. Qubitization [19] and truncated-Taylor methods [20] improve this to near-optimal query complexity

$$Q \;=\; \mathcal{O}\!\left(t\,\|H\|_1 + \frac{\log(1/\varepsilon)}{\log\log(1/\varepsilon)} \right),$$

linear in the simulated time and additive in precision. The potential advantage is real in principle: classical gyrokinetics scales polynomially in the phase-space resolution M, while Eq. (12) scales in $n = \log_2 M$. The obstruction is that Eq. (9) is only the *linear* operator; the turbulent saturation that fixes the transport is nonlinear, and the readout of a transport coefficient again reintroduces a measurement cost. The register's resource estimate is correspondingly staged: a linearised kinetic instance needs **11–19** logical qubits (KX-L3-A6), the six-dimensional kinetic encoding $\sim$**44** logical qubits and $\sim$**19** days of single-instance fault-tolerant runtime (KX-L3-A1), and a full nonlinear five-dimensional gyrokinetic simulation is placed beyond **2035** (KX-L3-A7).

RESOURCE-HONEST ADVANTAGE ACCOUNTING

Fault-tolerant overhead. None of Eqs. (5), (6) or (12) is a wall-clock time; each must be multiplied by the cost of realising a logical operation on error-corrected hardware. Under the surface code [33], a logical qubit at code distance d occupies

$$n_{\text{phys}} = n_{\log} \cdot 2d^2$$

physical qubits (KX-L3-A6), and the logical error rate per code cycle is suppressed as

$$p_L = A\left(\frac{p}{p_{\text{th}}}\right)^{\lfloor (d+1)/2 \rfloor},$$

with physical error rate p, threshold p_{th} and $\mathcal{O}(1)$ constant A. Requiring p_L below the inverse of the total (Clifford $+\,T$) gate count fixes d; the T/Toffoli gates then dominate the wall-clock time because each consumes a distilled magic state over $\mathcal{O}(d)$ code cycles [34]. A computation of N_T Toffoli gates at code-cycle time τ_c therefore costs of order $T_{\text{wall}} \sim N_T\, d\, \tau_c$ times the distillation factory throughput — which is how an $\mathcal{O}(10^9)$-Toffoli chemistry instance reaches days of runtime, and how the amplitude-amplification iteration count $N_{\text{iter}} = 1.6\sqrt{2^n}$ (KX-L3-A6) becomes prohibitive at loop scale.

The advantage inequality. A quantum route is admitted to the control loop only if it clears an explicit two-part bar, in the same spirit as the twin-admission criteria of the shadow architecture [39]: for a given task at problem size N, conditioning κ and accuracy ε,

$$\boxed{T_{\text{Q}}(N,\kappa,\varepsilon)\cdot C_{\text{FT}}(d) \;<\; T_{\text{C}}(N,\kappa,\varepsilon) \quad \wedge \quad \mathcal{F}_{\text{Q}} \;\geq\; \mathcal{F}_{\text{bar}}}$$

where $T_{\text{Q}}, T_{\text{C}}$ are the quantum and classical time-to-solution, $C_{\text{FT}}(d)$ the fault-tolerant constant-factor overhead from Eqs. (13)–(14), and $\mathcal{F}_{\text{Q}}, \mathcal{F}_{\text{bar}}$ the achieved and required fidelities. The whole of §5 is an evaluation of Eq. (15) for the four algorithms; not one of them satisfies it today, and for the quantum-ML surrogate the left inequality fails by 9–11 orders of magnitude.

§ 4

Computational methodology: the NVIDIA GPU HPC campaign

Every number in this paper was produced on the classical NVIDIA GPU HPC campaign; there is no quantum hardware in the loop and none is claimed. The campaign runs the four coupled digital twins [39] on an NVIDIA GPU fleet (A100/H100/A10-class accelerators), and the quantum-algorithm characterization is one workstream inside it.

Codes run for this paper. (i) The quantum-ML classifier characterization (BR-SX-08) was executed as a GPU state-vector emulation of the feature-map circuits of Eq. (7): the encoding unitaries $U_\phi(\quad x)$ were simulated exactly, the fidelity kernel of Eq. (7) assembled into the Gram matrix over the trial set, and a classical support-vector machine trained and scored against a matched classical kernel baseline. (ii) The fault-tolerant resource estimator (KX-L3-A6) recomputes the logical-qubit and gate budgets from Eqs. (13)–(14) and $N_{\text{iter}} = 1.6\sqrt{2^n}$; all 50 of its register rows were recomputed with 0 failures. (iii) The classical control-loop baselines against which the quantum routes are judged are the free-boundary Grad–Shafranov solver of the equilibrium twin [10,40], the CGYRO gyrokinetic operator [11] whose linearisation Eq. (9) encodes, and the quantum-inspired tensor-network kinetic solver (KX-L1-A16), whose singular-value-rank error decays exponentially (4.7×10^{-3} at rank 4, 2.9×10^{-6} at rank 8) and defines the classical frontier the quantum route must beat. (iv) The NV-diamond sensor model (KX-L1-A22) is a NumPy/SciPy shot-noise calculation.

GPU acceleration and problem scale. State-vector emulation of an n-qubit feature map costs $\mathcal{O}(2^n)$ memory and $\mathcal{O}(g\,2^n)$ time for g gates; the GPU throughput sets the practical ceiling of the characterization at $n = 8$ qubits, which is itself why the measured scaling series stops at eight (the $\mathcal{O}(2^n)$ Gram-matrix assembly over the trial set becomes the binding cost). The fault-tolerant estimates are analytic and inexpensive; their value is in being bounded and reproducible rather than compute-heavy.

Verification and reproducibility. The characterization is a frozen, negative-preserving result: the $0.817 \to 0.691$ AUC series, the 9–11-order slowdown, and the resource budgets are register entries (BR-SX-08, KX-L3-A1/A2/A6/A7) reproduced from archived inputs, and the $50/50$ resource-row recomputation with 0 failures is the internal cross-check. No result is permitted to over-claim advantage; negatives are kept verbatim. No compute-cost or economic figures are reported here or in the public deposit.

§ 5

Results

GRAD–SHAFRANOV STIFFNESS AND THE HHL SCALING

Figure 2 plots the condition number of the discretized GS operator, computed from the five-point stencil eigenvalues, against grid resolution; it follows the $\kappa \propto N_\perp^2$ law of Eq. (4) over the whole tested range (16 to 512 points per dimension). Feeding this κ into the cost models of Eqs. (5)–(6) with a representative fault-tolerant constant-factor overhead gives Figure 2: against an $\mathcal{O}(N_{\mathrm{dof}})$ classical multigrid solve, neither the original HHL nor the precision-improved QLSA crosses below the classical curve for full-field reconstruction across the equilibrium degrees of freedom relevant to the loop. The scaling exponents are the honest content; the constant factor only worsens the quantum curves.

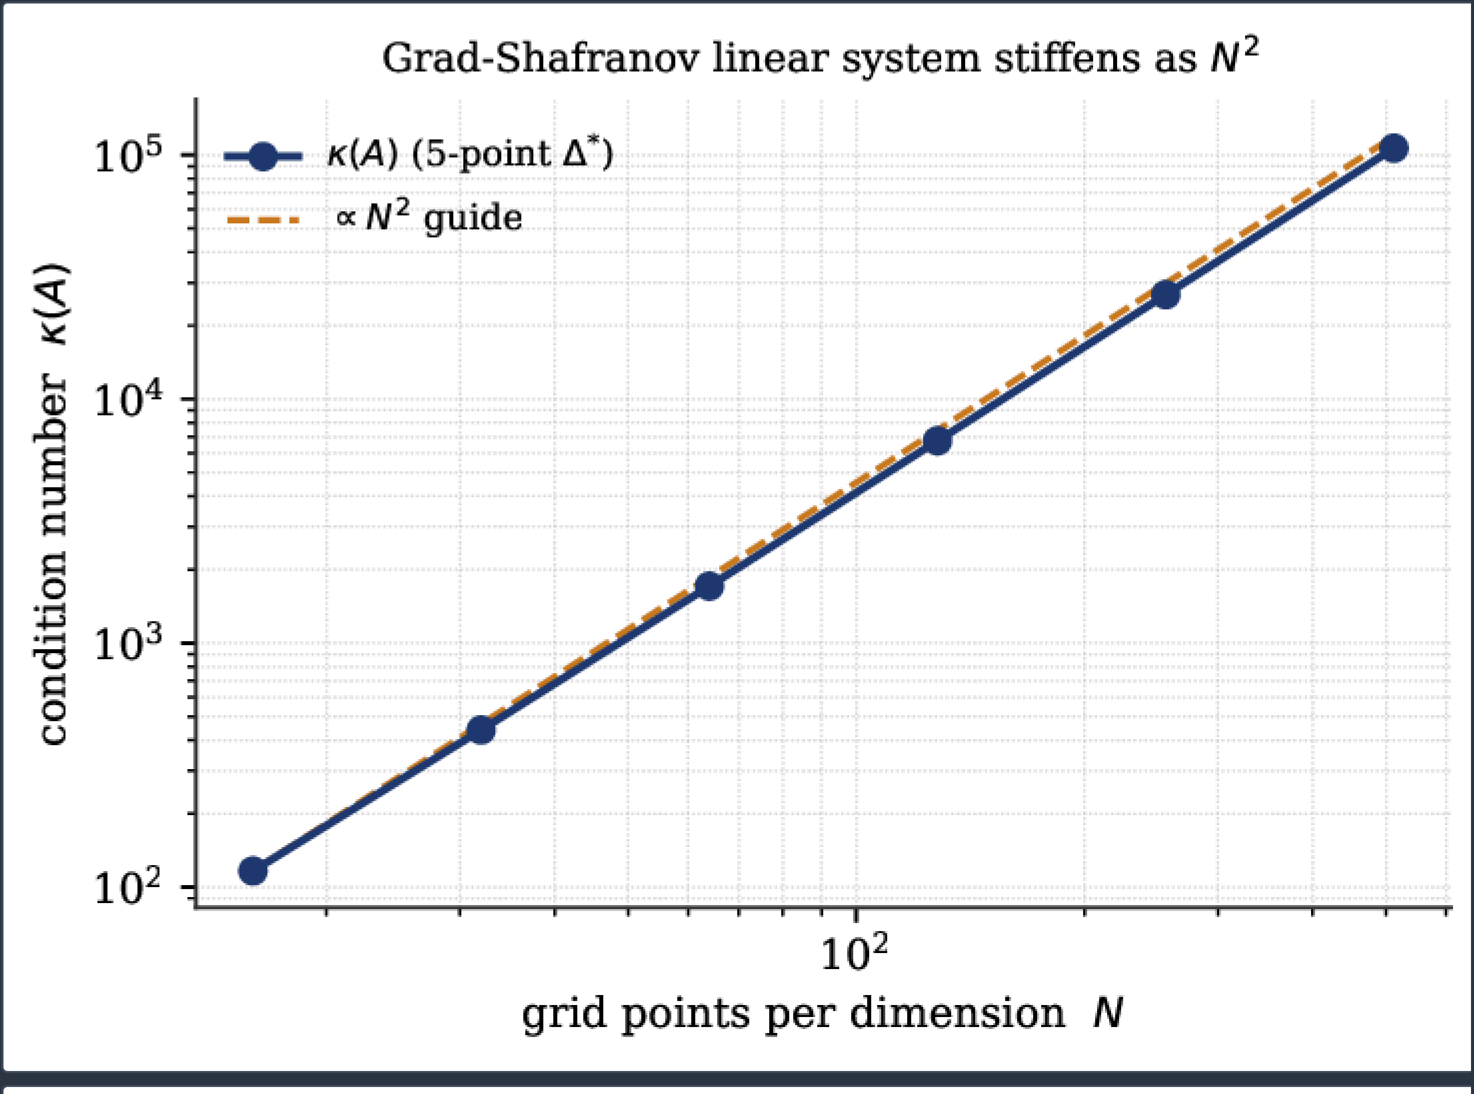
Grad-Shafranov linear system stiffens as N^2
$\kappa(A)$ (5-point Δ^*)
$\propto N^2$ guide
condition number $\kappa(A)$
10^5
10^4
10^3
10^2
10^2
grid points per dimension N

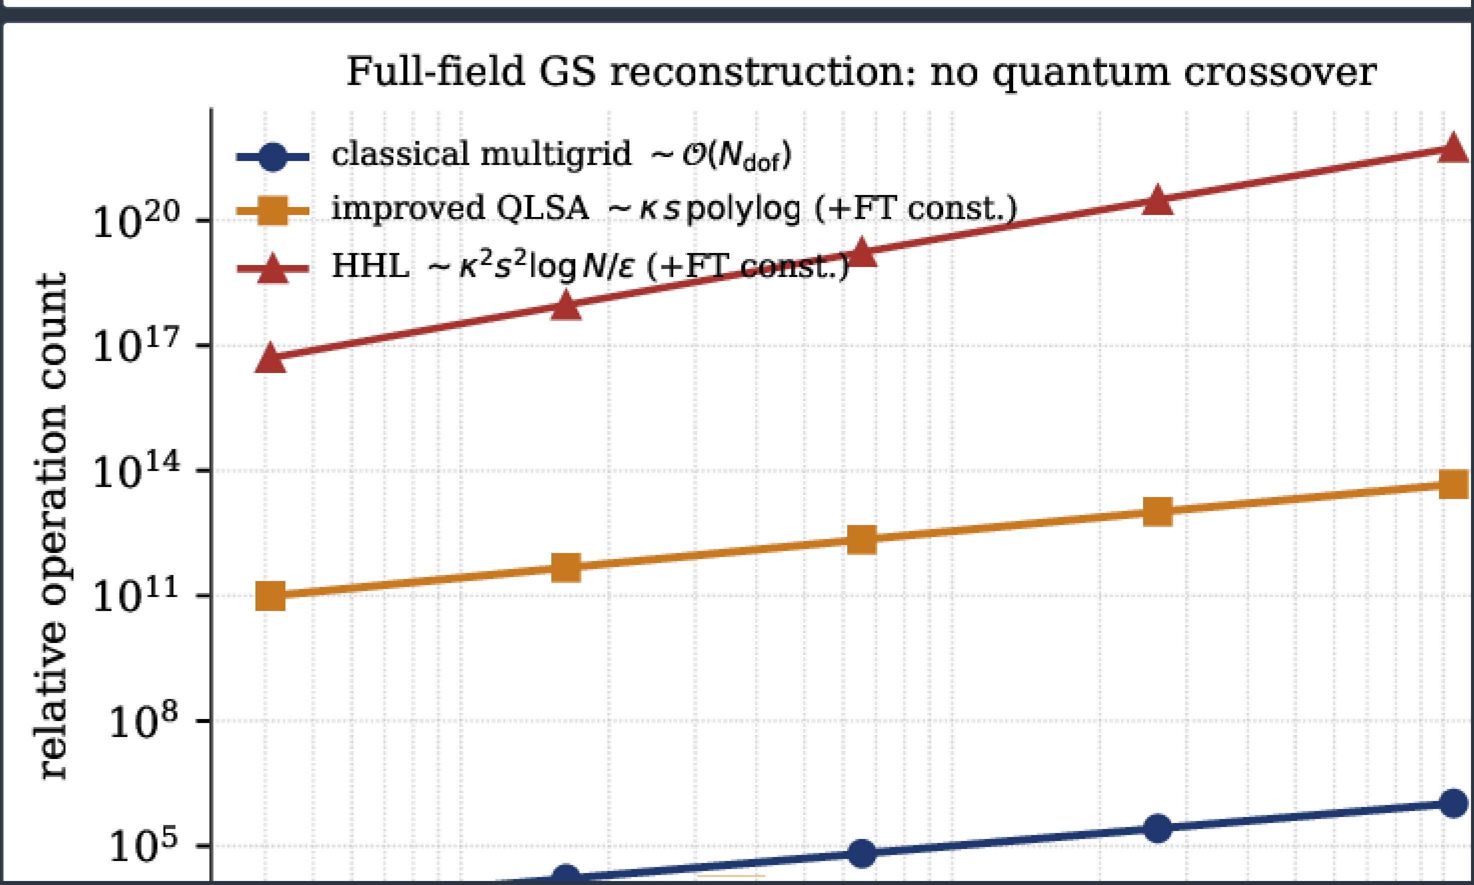
Full-field GS reconstruction: no quantum crossover
classical multigrid $\sim \mathcal{O}(N_{\mathrm{dof}})$
improved QLSA $\sim \kappa\, s\,\mathrm{polylog}$ (+FT const.)
HHL $\sim \kappa^2 s^2 \log N/\varepsilon$ (+FT const.)
relative operation count
10^{20}
10^{17}
10^{14}
10^{11}
10^8
10^5

THE MEASURED NO-ADVANTAGE RESULT

The quantum-ML surrogate is the one component of the stack that can be characterized on classical and emulated hardware today, and it was characterized directly (BR-SX-08). Figure 3(a) shows that the classifier's AUC *degrades* with scale, falling from 0.817 at four qubits to 0.691 at eight; panel (b) gives the resource-honest time-to-solution, a fault-tolerant kernel evaluation being 9–11 orders of magnitude slower than the classical route. Figure 4 shows that this is the calibrated signature of exponential kernel concentration: fitting Eq. (8) to the two measured operating points gives a per-qubit variance decay $\alpha \simeq 0.14$, and the induced shot cost $\sim 1/\operatorname{Var}$ grows exponentially over the same band. Table 2 collects the measured series and the derived concentration parameters. For this task the verdict is not merely "un-advantaged" but architecturally wrong: a quantum kernel is the wrong tool. The same knee is what governs the entanglement-enhanced sensing companion result — entanglement helps only above a noise threshold $\sigma \approx 0.14$ that all Kronos diagnostics sit below [43].

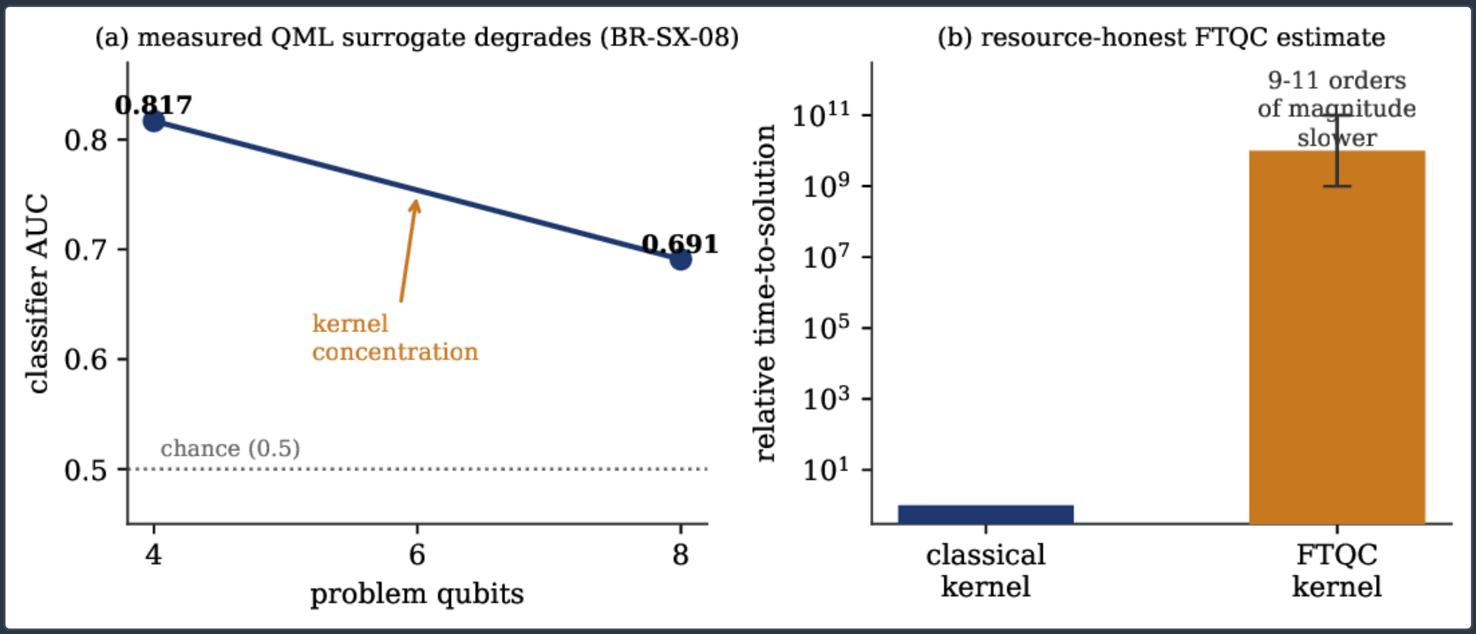

The measured honest-gate result (BR-SX-08). **(a)** Quantum-ML classifier AUC versus problem size: 0.817 at four qubits falling to 0.691 at eight — the signature of exponential kernel concentration. **(b)** Resource-honest time-to-solution: a fault-tolerant kernel evaluation is 9–11 orders of magnitude slower than the classical route. Measured (a) and resource-estimated (b) data.

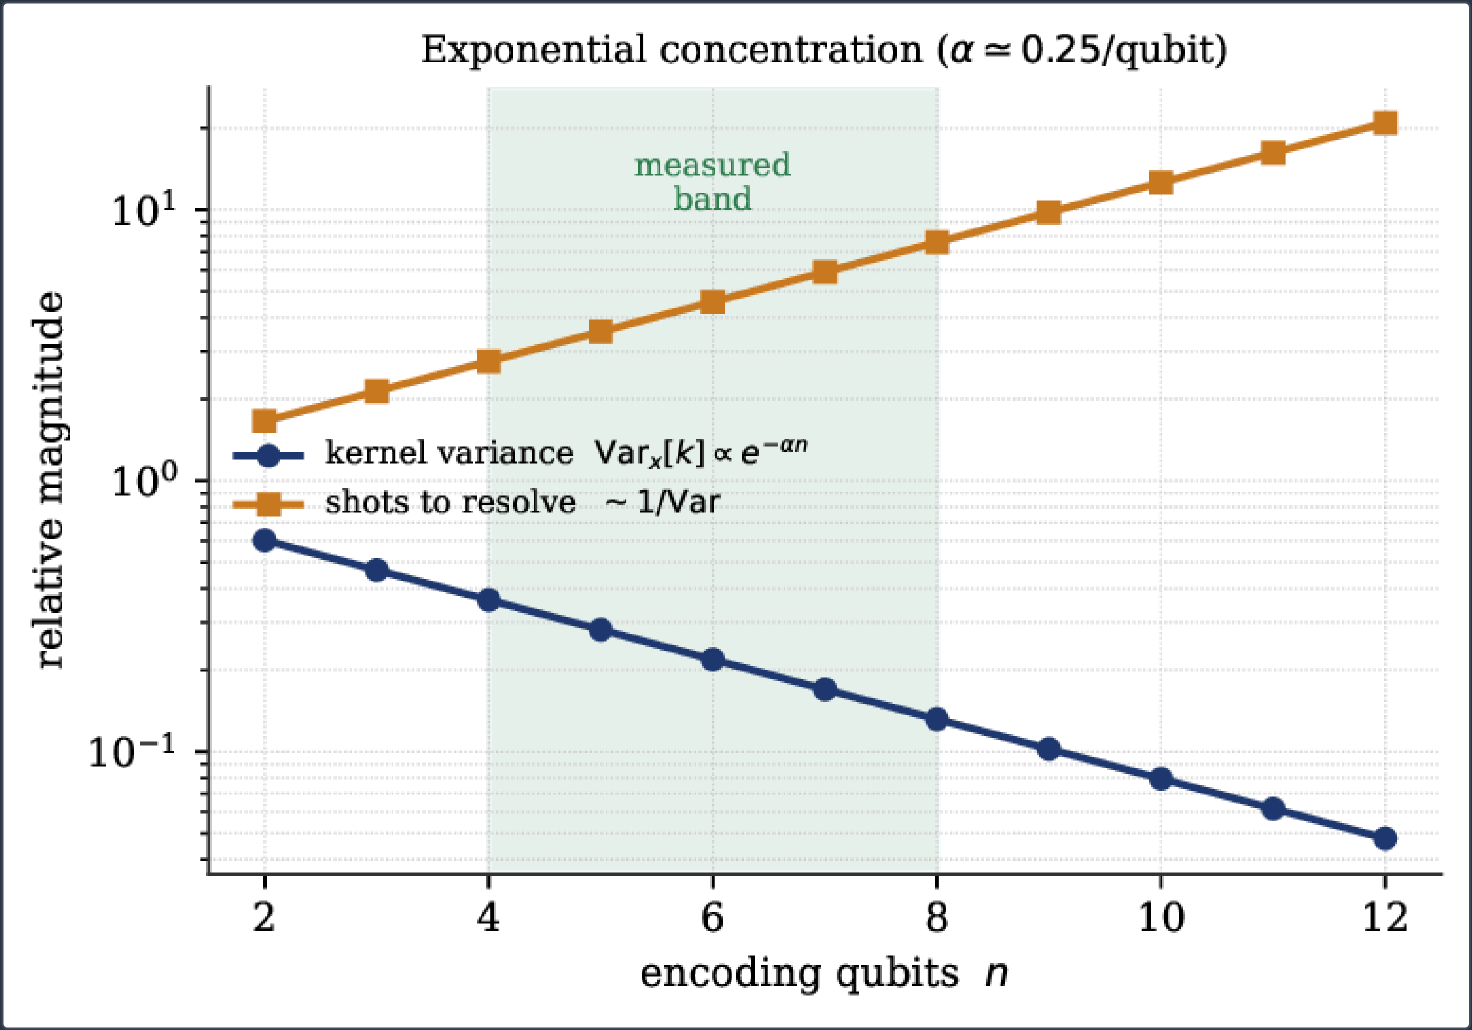

Exponential kernel concentration (Eq. 8) calibrated to the two measured AUC points of Figure 3(a): kernel variance $\mathrm{Var}_x[k] \propto e^{-\alpha n}$ with $\alpha \simeq 0.14$ per qubit, and the shot count $\sim 1/\mathrm{Var}$ needed to resolve an informative off-diagonal Gram entry. The shaded band marks the four-to-eight-qubit measured range. Derived model consistent with the measurement.

QUBITS n	MEASURED AUC	REL. KERNEL VARIANCE	REL. SHOTS TO RESOLVE
4	0.817	1.00	1.0
8	0.691	0.57	1.8
4l FTQC kernel evaluation: 9–11 orders of magnitude slower than classical (BR-SX-08).			

Quantum-ML surrogate characterization (BR-SX-08). Measured AUC series and the derived concentration parameters of Eq. (8) (α calibrated to the two measured points; relative variance $e^{-\alpha n}$ normalised to the four-qubit value; shots $\propto 1/\mathrm{Var}$).

FAULT-TOLERANT RESOURCE ESTIMATES FOR THE SOLVER-CLASS ALGORITHMS

The three solver-class algorithms are not characterizable on today's hardware; what is owned is a bounded resource estimate. Figure 5 gives the logical-qubit counts and single-instance runtimes from the register: a linearised kinetic instance at the core of the edge-turbulence simulation needs 11–19 logical qubits (KX-L3-A6), the full six-dimensional kinetic encoding ~ 44 logical qubits and ~ 19 days (KX-L3-A1), with the alloy-chemistry quantum-phase-estimation instance (~ 238 logical qubits, ~ 38 days, KX-L3-A2) shown for scale. Figure 5 translates these into physical resources through Eqs. (13)–(14): the physical-qubit count is $n_{\log}\, 2d^2$ at code distance d, and the code distance needed to reach a 10^{-9} logical error rate is read from the surface-code suppression curve. Table 3 collects the estimates with their register serials.

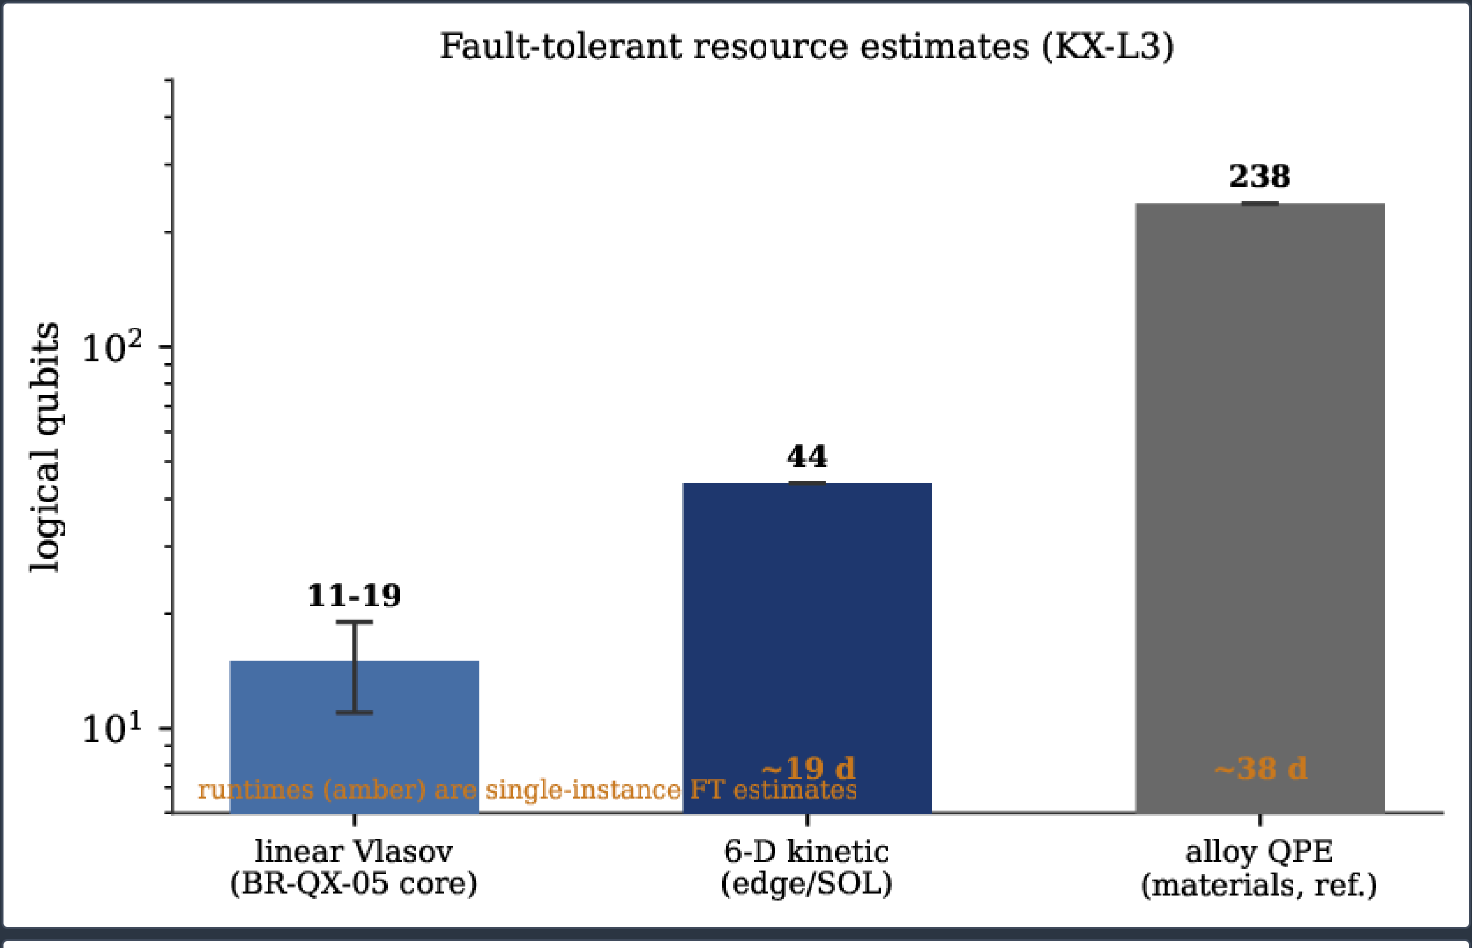

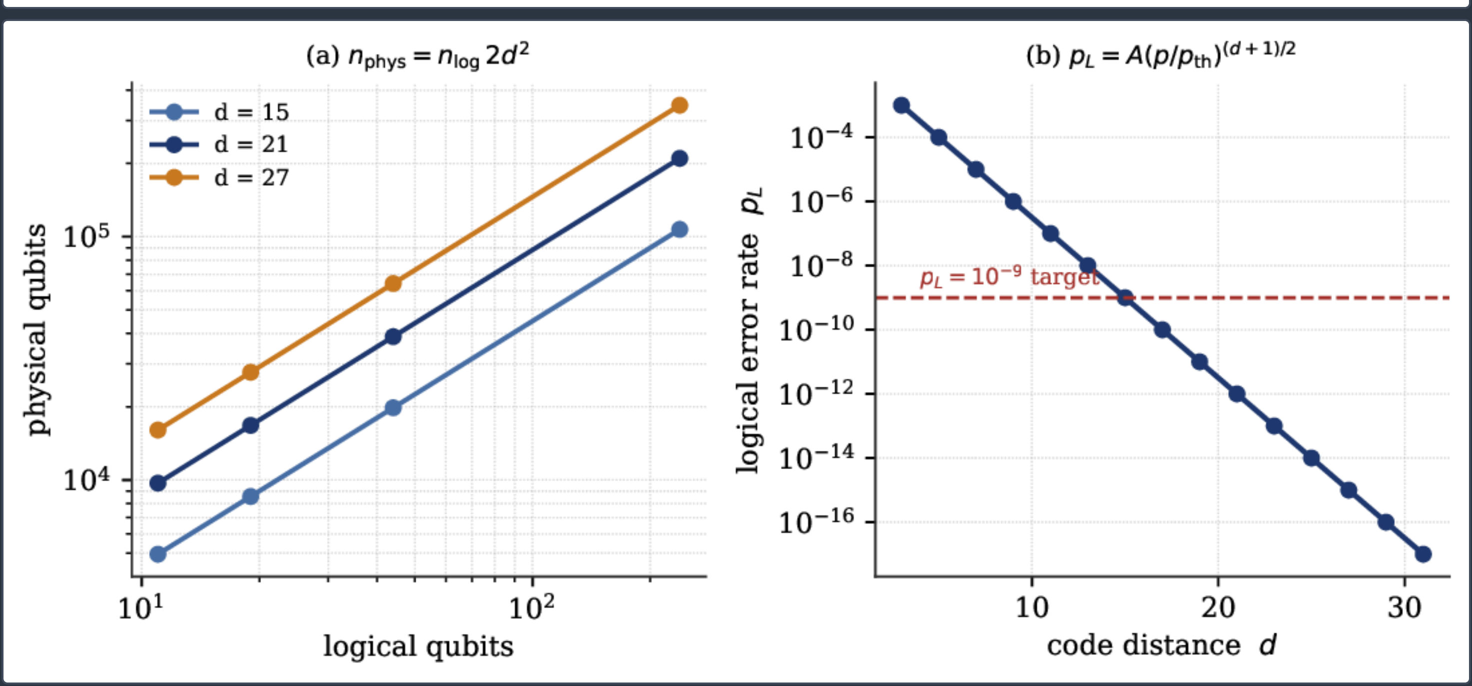

Fault-tolerant resource estimates (KX-L3). Logical-qubit counts (bars) and single-instance runtimes (amber) for the control-relevant kinetic instances, with the alloy-chemistry QPE case shown for scale. Estimated data from the register.

TASK	LOGICAL QUBITS	RUNTIME	SERIAL
linear kinetics (BR-QX-05 core)	11–19	—	KX-L3-A6
6-D kinetic (edge/SOL)	~ 44	~ 19 d	KX-L3-A1
alloy QPE (materials, reference)	~ 238	~ 38 d	KX-L3-A2
4l Amplitude-amplification iteration count $N_{\text{iter}} = 1.6\sqrt{2^n}$; fleet Grover-type search $\sim 1.8 \times 10^{18}$ iterations (KX-L3-A7): never.			

Fault-tolerant resource estimates for the control-loop and reference quantum tasks. Logical qubits and single-instance runtimes are frozen register entries; physical qubits are $n_{\log}\, 2d^2$ at $d = 27$ (Eq. 13). Serials refer to the register [37].

THE ADVANTAGE MAP AND THE STAGED ROADMAP

Figure 6 evaluates the advantage inequality Eq. (15) for the four algorithms as the ratio of quantum to classical time-to-solution: the quantum-ML surrogate is 9–11 orders slower (measured), and the three solver-class designs sit well above the advantage bar on their fault-tolerant estimates. No element crosses the bar today. Figure 7 places the corresponding schedule against the Kronos programme milestones (KX-L3-A7): the demonstrated NV sensor and the completed no-advantage characterization are the present quantum content; early fault-tolerant linear kinetics is placed at $\sim 2029 \pm 2$; the classical loop stays authoritative through first plasma ($\sim$2030–31); and nonlinear five-dimensional gyrokinetics is placed beyond 2035. Solver-class algorithms are admitted to the loop only when they clear Eq. (15).

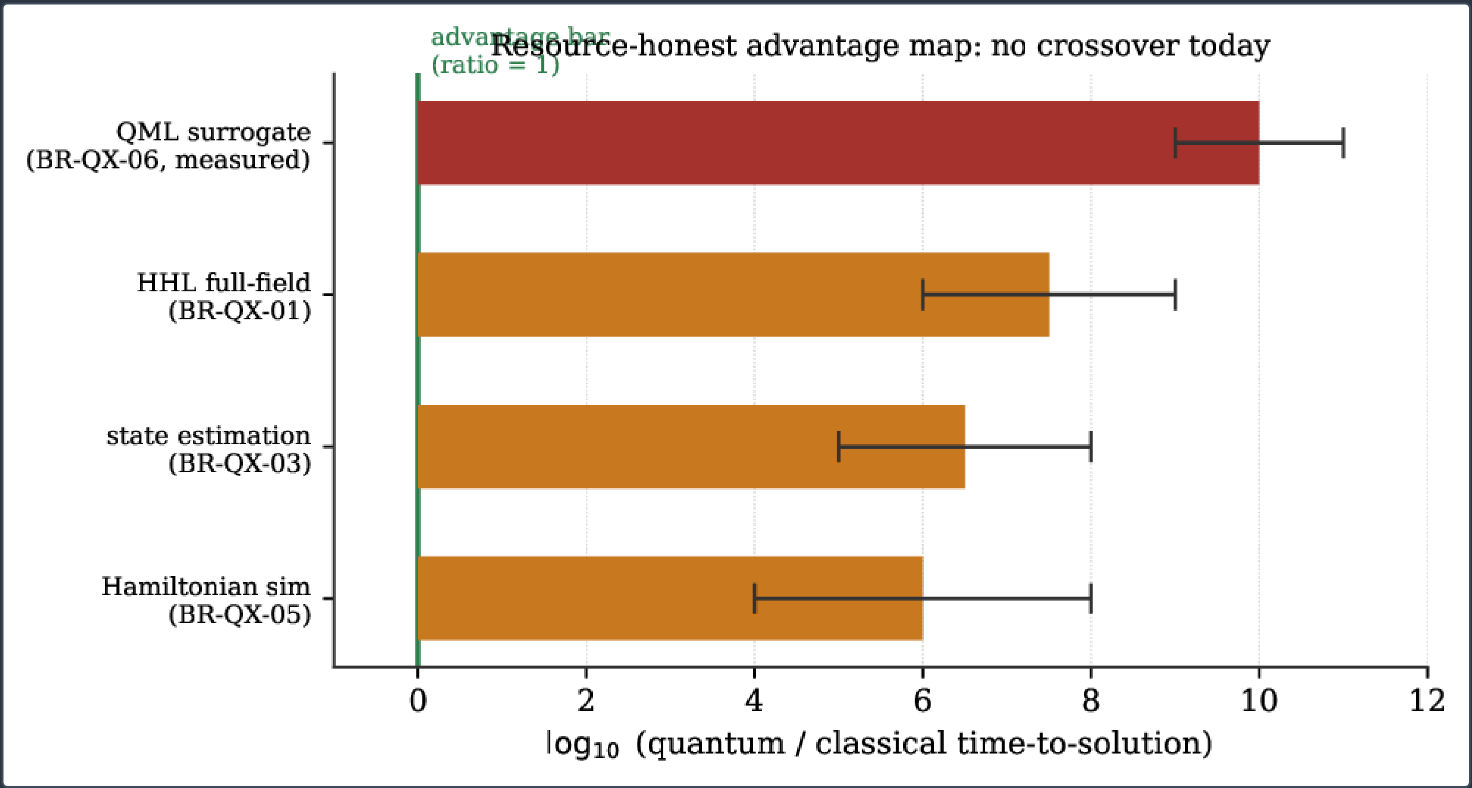

Resource-honest advantage map: $\log_{10}$ of the quantum-to-classical time-to-solution ratio for each stack element (Eq. 15). The quantum-ML bar is the measured 9–11 orders (BR-SX-08); the three solver-class bars are fault-tolerant design estimates. None crosses the advantage bar at ratio 1. Derived from the register estimates.

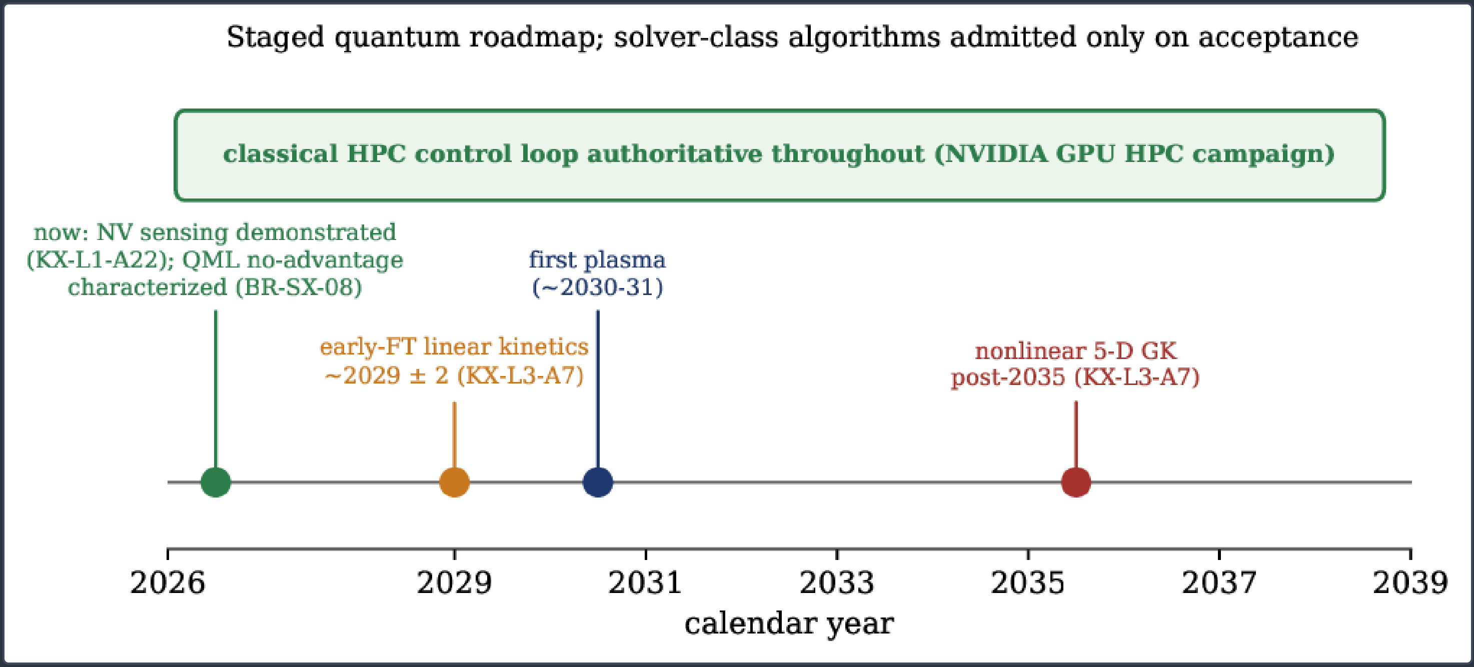

(Schematic.) Staged quantum roadmap aligned to Kronos milestones (KX-L3-A7). The classical HPC control loop is authoritative throughout; the demonstrated NV sensor and the no-advantage characterization are the present quantum content; solver-class algorithms are admitted only on the acceptance test of Eq. (15). Dates are programme placements, not measured values.

§ 6

Discussion

Against the quantum-algorithm literature. The HHL scaling of §3.1 reproduces, for the specific GS operator, the general caution of Aaronson's "fine print" [17]: an exponential speedup in the linear-algebra core (Eqs. 5–6) is neutralised by state preparation, readout and conditioning. The $\kappa \propto N_\perp^2$ stiffness we compute (Figure 2) is the concrete conditioning obstacle, and the improved QLSA [15,16] softens but does not remove it. The Hamiltonian-simulation route rests on the strongest theory of the four — qubitization [19] and truncated-Taylor methods [20] give near-optimal query complexity (Eq. 12) — and the Vlasov encoding of Engel *et al.* [31] and the plasma resource studies of Dodin & Startsev [32] and Joseph *et al.* [30] frame it precisely; our contribution is to attach a control-loop-scoped resource estimate (~ 44 logical qubits, ~ 19 days for the six-dimensional kinetic instance) and to note that the loop needs the *nonlinear* saturation and a transport-coefficient readout, both of which reintroduce cost.

Against the quantum-ML literature. The measured $0.817 \rightarrow 0.691$ degradation is not an implementation artefact; it is the exponential-concentration phenomenon proven by Thanasilp *et al.* [26] for expressive quantum kernels, the kernel-method counterpart of the barren-plateau result [27,28], and it is consistent with the broader finding of Huang *et al.* [25] that classical models with data match many claimed quantum-ML advantages. Our calibrated $\alpha \simeq 0.14$ per qubit (Figure 4) quantifies the effect on this specific control task. The lesson for a plant designer is that the quantum-kernel classifier [22,23] is the wrong tool for disruption state-estimation at scale, and the classical learned reconstructor and disruption model [40,41,6] remain authoritative.

Against the fault-tolerant-cost literature. The surface-code overhead of §3.4 (Eqs. 13–14) and the T-count-dominated wall-clock argument follow Fowler *et al.* [33] and the factoring cost model of Gidney & Eker [34]; the conclusion that these overheads defeat modest speedups in the near term is exactly that of Babbush *et al.* [35] and Beverland *et al.* [36]. What we add is the explicit advantage inequality Eq. (15) as the admission gate: a quantum route enters the loop only when it beats the classical surrogate on both latency and fidelity, and today none does. This is the same discipline applied in the Kronos resource-honest quantum-chemistry and quantum-optimization companions, where ADAPT-VQE and QAOA likewise show no measured advantage [44,45], and it is what keeps the programme's quantum investment aligned to the demonstrated sensor now and the solver-class designs later.

§ 7

Limitations

One honest gate bounds the solver-class claims: the HHL equilibrium reconstruction, the disruption state-estimation, and the edge-turbulence Hamiltonian simulation all require fault-tolerant quantum hardware at a logical-qubit scale (tens for the kinetic instances, Table 3) that does not yet exist, so they remain resource-scoped designs rather than running components. The classical HPC loop and the demonstrated NV sensor carry the control function in the interim, and each quantum route is admitted only on the explicit latency-and-fidelity acceptance test of Eq. (15). The one component measured today, the quantum-ML surrogate, is shown not to help; the resource estimates for the others are bounded but are estimates, not measurements.

§ 8

Conclusion

Kronos has a designed quantum-algorithm stack for the real-time control loop — HHL equilibrium reconstruction (BR-QX-01), disruption state-estimation (BR-QX-03), and edge-turbulence Hamiltonian simulation (BR-QX-05) — each stated here as a formal problem with its governing operator, discretization and fault-tolerant complexity, and each placed against the classical frontier through the explicit advantage inequality of Eq. (15). Alongside the designs it has a demonstrated in-winding quantum sensor and a completed, resource-honest characterization of the fourth, quantum-ML, route. That characterization is decisive and is stated once: there is no measured quantum advantage on these control problems today, the quantum-ML classifier degrades with scale ($\mathbf{0.817} \rightarrow \mathbf{0.691}$ AUC from four to eight qubits, the calibrated signature of exponential concentration at $\alpha \simeq \mathbf{0.14}$ per qubit), and a fault-tolerant kernel is $\mathbf{9\text{–}11}$ orders of magnitude slower for this problem. The control loop stays classical and authoritative on the NVIDIA GPU HPC campaign, the NV sensor is the working quantum element, and the solver-class designs are staged behind an explicit acceptance bar. A program that reports only its designs overstates its position; one that reports only the negative understates a real staged asset. Reporting all three is what lets the control architecture spend its quantum budget where it can pay.

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 GPU HPC campaign that produced them.

Reproducibility. The measured quantum-ML scaling ($\mathbf{0.817} \rightarrow \mathbf{0.691}$ AUC) and the $\mathbf{9\text{–}11}$-order fault-tolerant slowdown (BR-SX-08), the fault-tolerant resource estimates (KX-L3-A1/A2/A6/A7), the four control-algorithm designs (BR-QX-01/03/05/06), and the NV-diamond sensing metrics (KX-L1-A22) are frozen entries in the Kronos Physics De-Risking Register (doi:10.5281/zenodo.22133057). This paper is archived at doi:10.5281/zenodo.22645925 and its official version is hosted at https://www.kronosfusionenergy.com/publications/quantum-algorithms-for-the-real-time-control-loop-hhl-e.html; the figure-generation script (`make_figs_2p93_qalgo_control.py`) is deposited with the paper and regenerates every figure from the serials cited inline.

Data availability tocsectionData availability The quantum-ML scaling data, the resource estimates, and the algorithm specifications are archived with the deposit at doi:10.5281/zenodo.22645925 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 [37] by the serial identifiers cited inline. Any economic analysis is reported separately and is not included in the public deposit.

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

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 E Menard *et al*, Fusion nuclear science facilities and pilot plants based on the spherical tokamak, *Nucl. Fusion* **56** 106023 (2016). doi:10.1088/0029-5515/56/10/106023

2 B N Sorbom *et al*, ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* **100** 378 (2015). doi:10.1016/j.fusengdes.2015.07.008

3 A Marinoni, O Sauter and S Coda, A brief history of negative triangularity tokamak plasmas, *Rev. Mod. Plasma Phys.* **5** 6 (2021). doi:10.1007/s41614-021-00054-0

4 L L Lao, H St John, R D Stambaugh, A G Kellman and W Pfeiffer, Reconstruction of current profile parameters and plasma shapes in tokamaks, *Nucl. Fusion* **25** 1611 (1985). doi:10.1088/0029-5515/25/11/007

5 F Felici *et al*, Real-time physics-model-based simulation of the current density profile in tokamak plasmas, *Nucl. Fusion* **51** 083052 (2011). doi:10.1088/0029-5515/51/8/083052

6 J Degrave *et al*, Magnetic control of tokamak plasmas through deep reinforcement learning, *Nature* **602** 414 (2022). doi:10.1038/s41586-021-04301-9

7 J Kates-Harbeck, A Svyatkovskiy and W Tang, Predicting disruptive instabilities in controlled fusion plasmas through deep learning, *Nature* **568** 526 (2019). doi:10.1038/s41586-019-1116-4

8 L Lu, P Jin, G Pang, Z Zhang and G E Karniadakis, Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, *Nat. Mach. Intell.* **3** 218 (2021). doi:10.1038/s42256-021-00302-5

9 Z Li *et al*, Fourier neural operator for parametric partial differential equations, *Int. Conf. on Learning Representations (ICLR)* (2021). arXiv:2010.08895

10 N C Amorisco *et al*, FreeGSNKE: a Python-based dynamic free-boundary toroidal plasma equilibrium solver, *Phys. Plasmas* **31** 042517 (2024). doi:10.1063/5.0188467

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

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

13 J F Barry *et al*, Sensitivity optimization for NV-diamond magnetometry, *Rev. Mod. Phys.* **92** 015004 (2020). doi:10.1103/RevModPhys.92.015004

14 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

15 A M Childs, R Kothari and R D Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, *SIAM J. Comput.* **46** 1920 (2017). doi:10.1137/16M1087072

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

17 S Aaronson, Read the fine print, *Nat. Phys.* **11** 291 (2015). doi:10.1038/nphys3272

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

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

20 D W Berry, A M Childs, R Cleve, R Kothari and R D Somma, Simulating Hamiltonian dynamics with a truncated Taylor series, *Phys. Rev. Lett.* **114** 090502 (2015). doi:10.1103/PhysRevLett.114.090502

21 A Montanaro, Quantum algorithms: an overview, *npj Quantum Inf.* **2** 15023 (2016). doi:10.1038/npjqi.2015.23

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

23 V Havl\'icek *et al*, Supervised learning with quantum-enhanced feature spaces, *Nature* **567** 209 (2019). doi:10.1038/s41586-019-0980-2

24 J Biamonte *et al*, Quantum machine learning, *Nature* **549** 195 (2017). doi:10.1038/nature23474

25 H-Y Huang *et al*, Power of data in quantum machine learning, *Nat. Commun.* **12** 2631 (2021). doi:10.1038/s41467-021-22539-9

26 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

27 J R McClean *et al*, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9** 4812 (2018). doi:10.1038/s41467-018-07090-w

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

29 A Peruzzo *et al*, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5** 4213 (2014). doi:10.1038/ncomms5213

30 I Joseph *et al*, Quantum computing for fusion energy science applications, *Phys. Plasmas* **30** 010501 (2023). doi:10.1063/5.0123765

31 A Engel, G Smith and S E Parker, Quantum algorithm for the Vlasov equation, *Phys. Rev. A* **100** 062315 (2019). doi:10.1103/PhysRevA.100.062315

32 I Y Dodin and E A Startsev, On applications of quantum computing to plasma simulations, *Phys. Plasmas* **28** 092101 (2021). doi:10.1063/5.0056974

33 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

34 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

35 R Babbush *et al*, Focus beyond quadratic speedups for error-corrected quantum advantage, *PRX Quantum* **2** 010103 (2021). doi:10.1103/PRXQuantum.2.010103

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

37 P I Ford, *Kronos Physics De-Risking Register*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22133057

38 P I Ford, *AI and Quantum Computing for Fusion: ML Surrogates, Real-Time Control, and a Resource-Honest Quantum Frontier*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645702. https://www.kronosfusionenergy.com/publications/ai-quantum-control.html

39 P I Ford, *Four Coupled Digital Twins for a Compact Spherical-Tokamak Breeder: Multi-Rate Co-Simulation, Propagated Uncertainty, and a Real-Time Self-Consistency Monitor*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645805. https://www.kronosfusionenergy.com/publications/four-coupled-digital-twins-for-a-compact-spherical-toka.html

40 P I Ford, *Sub-Millisecond Neural Equilibrium Reconstruction Robust to Sensor Dropout: A Virtual-Sensor Mesh for Radiation-Degraded Diagnostics in a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645801. https://www.kronosfusionenergy.com/publications/sub-millisecond-neural-equilibrium-reconstruction-robus.html

41 P I Ford, *Disruptions, Runaway-Electron Avalanche and Electromagnetic Loads in a Compact Spherical-Tokamak Breeder: The Design Case for Shattered-Pellet Mitigation*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645724. https://www.kronosfusionenergy.com/publications/disruptions-and-runaway-electrons.html

42 P I Ford, *Fused Strain-Rate and Magnetization Quench Detection with In-Winding NV-Diamond Sensing for No-Insulation REBCO Breeder Magnets*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645752. https://www.kronosfusionenergy.com/publications/quench-detection-and-nv-sensing.html

43 P I Ford, *Quantum Machine-Learning Surrogates, Entanglement-Enhanced Magnetometry and Fault-Tolerant Resource Estimates for Fusion: A Resource-Honest Roadmap*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645905. https://www.kronosfusionenergy.com/publications/quantum-machine-learning-surrogates-entanglement-enhanc.html

44 P I Ford, *Resource-Honest Quantum Computing for Fusion: Variational Landau-Damping Emulation and ADAPT-VQE Low-Activation Alloy Chemistry Against the Classical Frontier*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645718. https://www.kronosfusionenergy.com/publications/quantum-computing-for-fusion.html

45 P I Ford, *Quantum Optimization and Reinforcement Learning for Real-Time Plasma Control: QUBO/QAOA Actuator Scheduling and Quantum RL for Shape Control, without Claimed Advantage*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645903. https://www.kronosfusionenergy.com/publications/quantum-optimization-and-reinforcement-learning-for-rea.html

46 P I Ford, *Poloidal-Field Coil Set and Equilibrium-Current Optimization for a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645945. https://www.kronosfusionenergy.com/publications/poloidal-field-coil-set-and-equilibrium-current-optimiz.html

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 Algorithms for the Real-Time Control Loop: HHL Equilibrium Reconstruction, Disruption State-Estimation, and Edge-Turbulence Simulation with Resource-Honest Advantage Accounting

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

Zenodo · doi:10.5281/zenodo.22645925

Read on kronosfusionenergy.com

The Physics De-Risking Register