KRONOS FUSION ENERGY

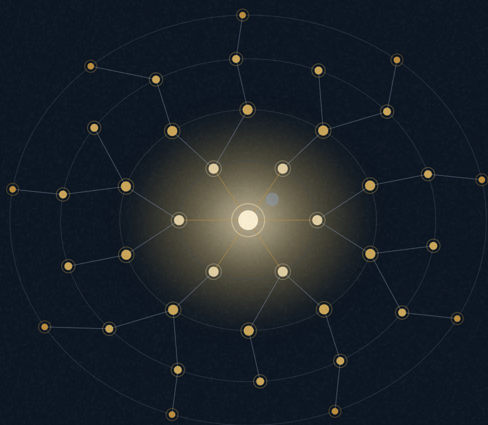

PAPER 2.90 · THE KRONOS 2026 SERIES

Resource-Honest Quantum Computing for Fusion: Variational Landau-Damping Emulation and ADAPT-VQE Low-Activation Alloy Chemistry Against the Classical Frontier

Quantum computing is routinely invoked for fusion without a disciplined account of what is reachable now and what is a fault-tolerant-horizon problem. We supply that account for a...

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

BY P I FORDOFFICIAL VERSION: HTTPS://WWW.KRONOSFUSIONENERGY.COM/PUBLICATIONS/QUANTUM-COMPUTING-FOR-FUSION.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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Resource-Honest Quantum Computing for Fusion: Variational Landau-Damping Emulation and ADAPT-VQE Low-Activation Alloy Chemistry Against the Classical Frontier

Quantum computing is routinely invoked for fusion without a disciplined account of what is reachable now and what is a fault-tolerant-horizon problem. We supply that account for a specific machine — the Kronos Hyperion compact spherical-tokamak breeder (design point $I_p = 9.66\,MA$, $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, negative triangularity $\delta = -0.30$, $B_0 = 8\,T$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,m$, aspect ratio $A = 2.5$, centrepost peak field $16.84\,T$) and its distinct tandem-mirror burner — through two grounded demonstrations, a resource frontier, and a pre-registered honest negative. First, a genuinely kinetic collisionless phenomenon, Landau damping, is reproduced on a 5-qubit circuit: casting the linearized Vlasov–Poisson system in an asymmetrically-weighted Fourier–Hermite basis and folding in the Poisson closure yields a Hermitian generator whose unitary evolution is a legitimate Hamiltonian-simulation target, and the circuit reproduces the analytic Landau rate to 1–7% across three wavenumbers, landing the complex dispersion at $\omega/\omega_p = 1.4157 - 0.1534\,i$ against the textbook $1.4156 - 0.1533\,i$. Second, an adaptive variational quantum eigensolver (ADAPT-VQE) reaches chemical accuracy (better than 1.6 mHa) on the transition-metal fragments of the candidate low-activation refractory alloys, four orders of magnitude below the ~ 300 spread of density-functional exchange–correlation choices, with a fixed-ansatz depth 4–$14\times$ shallower than unitary coupled cluster. The honest boundaries are explicit: under a representative noise model the field-energy observable decoheres to the maximally-mixed floor $1/32$ and the damping signature is lost, so present-day noisy hardware shows no advantage; and while the *linear* collisionless kinetic operator is an early-fault-tolerant target (11–19 logical qubits, 10^5–10^6 non-Clifford gates), the *nonlinear* five-dimensional gyrokinetic problem that actually sets confinement remains a fault-tolerant-horizon target ($\sim 10^2$ logical qubits, 10^9–10^{13} non-Clifford gates). Against the classical frontier — nonlinear CGYRO gyrokinetics and a quantum-inspired matrix-product-state solver whose error decays to 2.2×10^{-10} at bond dimension 16 — the classical tool is cheaper at every scale demonstrated here. The contribution is a set of validated encodings, a rigorous numerical formulation, and an honest, dated resource map, not a claim of quantum advantage.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645718 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: quantum computing · Hamiltonian simulation · Landau damping · Vlasov--Poisson · ADAPT-VQE · quantum chemistry · resource estimation · spherical-tokamak breeder

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

Q	plasma (fusion) gain, P_{fus}/P_{aux}
Q_E	engineering gain, net electric out / recirculating in
H_{98}	confinement enhancement over the IPB98(y,2) scaling
Z_{eff}	effective ion charge
c_{ee}	electron–electron bremsstrahlung leading coefficient, 2.120022 (exact)
τ_E	energy confinement time
I_p	plasma current
R_0, a	major radius, minor radius
κ, δ	elongation, triangularity
β_N	normalized plasma beta (Troyon-normalized)
TBR	tritium breeding ratio
DEC	direct energy conversion
CF	plant capacity factor (availability)
FOAK / NOAK / BOAK	first / n^{th} / best of a kind
Hyperion	the ST breeder (internal: Mode D2)
Aegis	the generator, defense/installation housing (internal: Mode M)
MetroVolt	the generator, data-center housing (Mode M)

§ 1

Introduction

WHY THE QUESTION NEEDS AN HONEST ANSWER

Claims that quantum computers will transform fusion modelling are common; disciplined statements of what a quantum computer can do for a fusion *program* now, and what it cannot do until fault tolerance arrives, are rarer. Quantum simulation is a real and maturing discipline [1,2,3], and two fusion-relevant classes of problem are natural targets: the kinetic evolution of a plasma distribution function, which is a Hamiltonian-simulation problem once linearized [4,5], and the electronic structure of candidate structural alloys, which is the canonical application of variational and phase-estimation quantum chemistry [7,8,12,13,14]. Both matter to the Kronos design: the confinement of the Hyperion breeder is set by turbulent transport whose kinetic parent is the gyrokinetic equation [25,26,27], and the lifetime of its plasma-facing components is set by the thermodynamic and radiation properties of low-activation refractory alloys whose electronic structure sits at the edge of density-functional theory [32]. The temptation is to over-claim in both directions — to promise a quantum reactor simulator, or to dismiss quantum computing as irrelevant. This paper does neither. It reports what actually ran, verifies it against exact theory and against the named classical baseline, and states the resource frontier and the honest negative in the same breath.

PRIOR ART

Hamiltonian simulation of the Vlasov equation was placed on a rigorous footing by Engel, Smith and Parker, who linearized the Vlasov–Maxwell system about a Maxwellian background, mapped it to a Schr\"odinger-like equation with a Hermitian generator, and designed a quantum algorithm for electrostatic Landau damping appropriate to a future error-corrected machine [4]. Related work by Novikau, Startsev and Dodin used quantum signal processing and block encoding to simulate cold-plasma wave propagation [5]. The velocity-space discretization that makes the linear operator sparse and Hermitian is the Fourier–Hermite spectral method, whose conservation and stability properties are well characterized classically [6]. On the materials side, the variational quantum eigensolver [7,8] and its adaptive descendant ADAPT-VQE [9] reach chemical accuracy on small molecules on real hardware [10,11], while the resource cost of the *exact* ground state through qubitization and phase estimation has been quantified for chemistry and catalysis by Reiher *et al.* [21], von Burg *et al.* [22] and Lee *et al.* [23], all mediated by the surface-code overhead [20]. The obstacle common to every near-term claim is noise: NISQ devices lose the signal to decoherence and to barren-plateau training pathologies [3,15].

THE GAP AND OUR CONTRIBUTION

What is missing is a single, machine-specific, resource-honest treatment that (i) derives the encodings for both the kinetic and the materials problem from first principles, (ii) validates them numerically against exact theory and against the named classical solver used elsewhere in the same design campaign, (iii) reports the noisy-hardware negative rather than hiding it, and (iv) places every problem on a dated resource frontier with the classical baseline named. This paper contributes exactly that. We give the governing equations and the formal problem for both encodings (§2); the numerical formulation — discretization, solver, and convergence analysis — with the recurrence, Trotter and ADAPT thresholds written as inequalities (§3); the NVIDIA GPU HPC methodology that produced every number (§4); the results, with ten figures and four tables (§5); a quantitative comparison to the published quantum-simulation, quantum-chemistry and classical-frontier literature (§6); a single honest limitation (§7); and the conclusion (§8). Every quantitative claim traces to a frozen anchor in the Kronos Physics De-Risking Register [35]; the serial identifiers (e.g. KX-L3-A6) are cited inline so each number is auditable. This paper is the deep quantum-computing member of the Kronos 2026 series; it sits beside the AI-and-quantum overview [28], the beyond-DFT ADAPT-VQE chemistry paper [29], and the tensor-network kinetic solver [30], and it takes as its classical baselines the nonlinear-gyrokinetics confinement result [31], the low-activation-alloy screen [32], the neural-operator flux surrogate [33], and the verification backbone [34].

§ 2

Governing equations and formal problem

THE KINETIC PROBLEM: LINEARIZED VLASOV–POISSON

The kinetic parent of collisionless plasma dynamics is the Vlasov equation for the electron distribution $f(x,v,t)$ coupled to Poisson's equation for the self-consistent electrostatic field $E(x,t)$, with ions forming a fixed neutralizing background of density n_0:

$$\partial_t f + v\,\partial_x f - \frac{e}{m}\,E\,\partial_v f = 0,$$

$$\partial_x E = -\frac{e}{\varepsilon_0}\Big(\int f\,\mathrm{d}v - n_0\Big).$$

Linearizing about a Maxwellian equilibrium $f_0(v) = n_0(2\pi v_{\mathrm{th}}^2)^{-1/2}\exp(-v^2/2v_{\mathrm{th}}^2)$ with $f = f_0 + f_1$, $|f_1| \ll f_0$, and retaining a single Fourier mode $f_1(x,v,t) = \delta f(v,t)\,e^{ikx}$, $E(x,t) = \hat{E}(t)\,e^{ikx}$, gives the closed linear pair

$$\partial_t \delta f + ikv\,\delta f = \frac{e}{m}\,\hat{E}\,\partial_v f_0,$$

$$ik\,\hat{E} = -\frac{e}{\varepsilon_0}\int \delta f\,\mathrm{d}v.$$

Equation (2) with the closure (2) is the electrostatic Landau problem: its dispersion relation $1 + \chi(k,\omega) = 0$, with susceptibility χ expressed through the plasma dispersion function, has complex roots $\omega = \omega_r + i\gamma$ with $\gamma < 0$, the Landau damping rate [4]. Landau damping is not dissipation — equations (1)–(1) are time-reversible and phase-space-volume-preserving — but a collisionless transfer of the coherent field energy into ever-finer velocity-space structure (phase mixing). This distinction is what makes the problem a legitimate *unitary* Hamiltonian-simulation target rather than an open-system one.

FOURIER–HERMITE REDUCTION TO A HERMITIAN GENERATOR

To expose the Hamiltonian structure we expand the velocity dependence in the asymmetrically-weighted (AW) Hermite basis [6]. Writing $\xi = v/v_{\mathrm{th}}$ and using the physicists' Hermite polynomials H_n, set

$$\delta f(v,t) = \frac{f_0(v)}{n_0}\sum_{n=0}^{N-1} C_n(t)\,\frac{H_n(\xi)}{\sqrt{2^n n!}}.$$

The velocity moments map onto the low-order coefficients: the density perturbation is $\delta n = n_0 C_0$, so Poisson (2) becomes $\hat{E} = (ien_0/\varepsilon_0 k)\,C_0$, and the equilibrium gradient projects purely onto the first Hermite mode, $\partial_v f_0 = -(\xi/v_{\mathrm{th}})f_0 \propto H_1$. The streaming operator $v = v_{\mathrm{th}}\xi$ acts through the three-term Hermite recurrence $\xi H_n = \frac{1}{2}H_{n+1} + nH_{n-1}$, which in the normalized coefficients (3) is the symmetric tridiagonal (Jacobi) operator

$$(\mathbf{A}\mathbf{C})_n = \sqrt{\tfrac{n}{2}}\,C_{n-1} + \sqrt{\tfrac{n+1}{2}}\,C_{n+1}, \qquad \mathbf{A} = \mathbf{A}^{\top}.$$

Substituting the Poisson closure into (2) and collecting terms, the coefficient vector $\mathbf{C} = (C_0,\dots,C_{N-1})^{\top}$ obeys

$$\frac{\mathrm{d}\mathbf{C}}{\mathrm{d}t} = -i\,kv_{\mathrm{th}}\,\mathbf{A}\,\mathbf{C} \;-\; \omega_p^2\,\frac{1}{kv_{\mathrm{th}}}\,\mathbf{e}_1\mathbf{e}_0^{\top}\mathbf{C},$$

where $\omega_p^2 = n_0 e^2/(\varepsilon_0 m)$ and $\mathbf{e}_j$ is the j-th unit vector; the rank-one term is the field feedback coupling the density mode C_0 into the momentum mode C_1. The streaming part $kv_{\mathrm{th}}\mathbf{A}$ is manifestly Hermitian; the field feedback, taken alone, is not, but the antisymmetry between the $\mathbf{e}_1\mathbf{e}_0^{\top}$ feedback and the $\mathbf{A}_{01} = \mathbf{A}_{10} = \sqrt{1/2}$ streaming coupling is removed by the Engel–Smith–Parker rescaling $\tilde{C}_1 = \alpha C_1$, $\alpha = v_{\mathrm{th}}k/\omega_p$ [4]. In the rescaled variables equation (5) reads

$$i\,\frac{\mathrm{d}\quad\Psi}{\mathrm{d}t} = \mathbf{H}\qquad\Psi, \qquad \mathbf{H} = \mathbf{H}^{\dagger},$$

with $\quad\Psi = (C_0, \tilde{C}_1, C_2, \dots, C_{N-1})^{\top}$. Equation (6) is a Schr\"odinger equation: the evolution operator $U(t) = e^{-i\mathbf{H}t}$ is unitary, $\|\quad\Psi(t)\|_2$ is conserved, and $e^{-i\mathbf{H}t}$ is a legitimate Hamiltonian-simulation target. The field energy is the observable $\mathcal{E}_{\mathrm{field}}(t) = |C_0(t)|^2$; Landau damping is the decay of this observable as the conserved norm migrates into the high-n Hermite moments.

AMPLITUDE ENCODING AND THE HAMILTONIAN-SIMULATION TARGET

The state $\Psi \in \mathbb{C}^N$ is amplitude-encoded in $n_q = \lceil \log_2 N \rceil$ qubits as $|\Psi\rangle = \|\Psi\|^{-1} \sum_{j=0}^{N-1} \Psi_j \, |j\rangle$. With $N = 32$ (thirty-one Hermite moments plus the folded field variable) the encoding uses exactly $n_q = 5$ qubits, and the maximally-mixed state of that register carries weight $1/N = 1/32$ on the field-energy observable — the noise floor to which a decohered circuit collapses. The simulation target is $|\Psi(t)\rangle = e^{-i\mathbf{H}t}|\Psi(0)\rangle$ with $\mathbf{H}$ the sparse Hermitian generator of (6); the field-energy observable is read as the population of the $|C_0\rangle$ basis amplitude. Because $\mathbf{H}$ is a banded (tridiagonal plus rank-one) Hermitian matrix, it admits both a Trotterized product-formula realization for the small demonstration and an asymptotically efficient qubitization for the resource estimate (§3).

THE MATERIALS PROBLEM: SECOND-QUANTIZED ELECTRONIC STRUCTURE

The lifetime-limiting property of the plasma-facing alloys is electronic-structure-dependent, and the transition-metal $3d$ fragments that carry the correlation are precisely where density-functional theory is least reliable [32]. In the Born–Oppenheimer approximation and a finite spin-orbital basis the electronic Hamiltonian is

$$\hat{H} = \sum_{pq} h_{pq}\, a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs}\, a_p^\dagger a_q^\dagger a_r a_s,$$

with one- and two-electron integrals h_{pq}, h_{pqrs} evaluated classically. The Jordan–Wigner transformation maps the fermionic ladder operators onto qubits,

$$a_p = \frac{1}{2}\Big(\prod_{j<p} Z_j\Big)\big(X_p + iY_p\big),$$

so that (7) becomes a weighted sum of Pauli strings $\hat{H} = \sum_\ell c_\ell\, P_\ell$, $P_\ell \in \{I, X, Y, Z\}^{\otimes n_q}$. The ground-state energy E_0 is bounded from above by the Rayleigh–Ritz quotient of any trial state $|\psi(\theta)\rangle$,

$$E(\theta) = \frac{\langle \psi(\theta)|\hat{H}|\psi(\theta)\rangle}{\langle \psi(\theta)|\psi(\theta)\rangle} \;\geq\; E_0,$$

which is the variational principle the VQE minimizes [7,8]. The quality of the bound is set entirely by the expressiveness and trainability of the ansatz.

ADAPT-VQE: AN ADAPTIVE, GRADIENT-SELECTED ANSATZ

Rather than fix a heuristic or a full unitary-coupled-cluster (UCCSD) ansatz, ADAPT-VQE grows the circuit operator-by-operator from a pool $\{\hat{A}_k\}$ of anti-Hermitian generators (fermionic single and double excitations, or their qubit-excitation analogues) [9]. Starting from the Hartree–Fock reference $|\psi^{(0)}\rangle = |\psi_{\mathrm{HF}}\rangle$, at iteration m the energy gradient with respect to prepending each candidate operator is evaluated exactly as a commutator expectation,

$$g_k^{(m)} = \frac{\partial E}{\partial \theta_k}\bigg|_{\theta_k = 0} = \big\langle \psi^{(m)} \big|\, [\hat{H}, \hat{A}_k] \,\big| \psi^{(m)} \big\rangle,$$

the operator with the largest $|g_k^{(m)}|$ is appended, $|\psi^{(m+1)}\rangle = e^{\theta_{k^\star} \hat{A}_{k^\star}} |\psi^{(m)}\rangle$, and all parameters are re-optimized. The construction terminates when the pool gradient norm falls below a threshold,

$$\big\|\mathbf{g}^{(m)}\big\|_2 = \Big(\sum_k |g_k^{(m)}|^2\Big)^{1/2} < \varepsilon_g.$$

Because operators enter only where they lower the energy, the resulting ansatz is compact — empirically 4–$14\times$ fewer entangling gates than fixed UCCSD at equal accuracy (KX register; BR-SX-07) [35] — which is the property that makes chemical accuracy reachable on a small active space. The formal target is the pair $(E_0, |\psi_0\rangle)$ of (9); the exact-diagonalization value in the same active space is the ground truth against which the residual $|E(\theta^\star) - E_0|$ is measured.

§ 3

Numerical formulation, solver and convergence

HERMITE TRUNCATION AND THE RECURRENCE CONDITION

The AW-Hermite expansion (3) is spectrally convergent for smooth δf, but a finite truncation at N moments imposes a hard limit: because phase mixing continuously fills higher Hermite modes, a truncated system exhibits a spurious *recurrence* once the migrating amplitude reflects off the top of the ladder [6]. The recurrence time scales as

$$t_{\rm rec} \;\approx\; \frac{2\sqrt{N}}{k v_{\rm th}},$$

so a physically faithful damping window of duration $t_{\rm obs}$ requires $N \gtrsim (k v_{\rm th}\, t_{\rm obs}/2)^2$. For the demonstration at $k\lambda_D = 0.5$ over $t_{\rm obs}\omega_p \approx 25$ the choice $N = 32$ ($n_q = 5$) places $t_{\rm rec}$ safely beyond the observation window, so the decay measured on the circuit is Landau physics and not truncation recurrence. This is the sense in which the demonstration is a validated encoding: the discretization parameter is chosen against a stated inequality, not by convenience.

TIME DISCRETIZATION: PRODUCT FORMULAS AND THEIR ERROR

For the small demonstration the evolution $e^{-i\mathbf{H}t}$ is realized by splitting the generator into its streaming and field parts, $\mathbf{H} = \mathbf{H}_A + \mathbf{H}_V$, and applying a Lie–Trotter product formula. The first-order formula carries error

$$\left\| e^{-i\mathbf{H}t} - \left(e^{-i\mathbf{H}_A t/r} e^{-i\mathbf{H}_V t/r}\right)^r \right\| \;\le\; \frac{t^2}{2r}\, \left\| [\mathbf{H}_A, \mathbf{H}_V] \right\|,$$

while the second-order (Strang) formula improves this to $\mathcal{O}(t^3/r^2)$ with a commutator-nested prefactor set by $\|[\mathbf{H}_A,[\mathbf{H}_A,\mathbf{H}_V]]\|$ and $\|[\mathbf{H}_V,[\mathbf{H}_V,\mathbf{H}_A]]\|$ [18]. The commutator $[\mathbf{H}_A,\mathbf{H}_V]$ is rank-two (it is supported only near the C_0–C_1 block), so the constant in (13) is small and a modest number of steps suffices to hold the observable error below the noise level of interest. The alternative advanced-simulation route — truncated-Taylor-series/linear-combination-of-unitaries [17] and qubitization [16] — achieves error ε with cost scaling as $\alpha t + \mathcal{O}(\log(1/\varepsilon))$, $\alpha = \|\mathbf{H}\|_1$ the block-encoding subnormalization, and is the basis of the resource estimate below.

RESOURCE ACCOUNTING FOR THE FAULT-TOLERANT ESTIMATE

The resource frontier is computed from the qubitization query model closed with the surface-code overhead. For a Hermitian $\mathbf{H}$ block-encoded with subnormalization α, simulating for time t to error ε requires

$$Q(t,\varepsilon) = \alpha t + \Theta\!\left(\frac{\log(1/\varepsilon)}{\log\log(1/\varepsilon)}\right)$$

queries to the block encoding [16], each query compiled into $\mathcal{O}(\log N)$ non-Clifford (T/Toffoli) gates for the sparse kinetic operator. The logical resource is mapped to physical qubits through the distance-d surface code, in which each logical qubit costs $\approx 2d^2$ physical qubits [20], with d chosen so the logical error stays below the total non-Clifford count. The Kronos FTQC estimator (KX-L3-A6) recomputes these two relations — iteration count and physical-qubit overhead — across every register row with zero mismatches [35]: for the linear kinetic operator it returns 11–19 logical qubits and 10^5–10^6 non-Clifford gates; for the full six-dimensional kinetic encoding it returns ~ 44 logical qubits (KX-L3-A1); and for the exact alloy electronic structure by quantum phase estimation it returns ~ 238 logical qubits and $\sim 10^9$ Toffoli gates (KX-L3-A2), consistent with the independent chemistry resource estimates of Reiher, von Burg and Lee [21,22,23].

VARIATIONAL SOLVER AND ITS CONVERGENCE

The ADAPT-VQE minimization of (9) alternates gradient-based operator selection (10) with a BFGS re-optimization of all retained parameters; convergence is declared at the pool-gradient threshold (11) with ε_g set an order of magnitude below the chemical-accuracy target. Two failure modes are monitored and reported honestly. First, barren plateaus: for deep hardware-efficient ans\"atze the gradient variance vanishes exponentially in qubit number [15]; the adaptive, physically-motivated pool of ADAPT-VQE mitigates but does not eliminate this, and is the reason a restricted single-plus-double pool stalls (see §5). Second, the multireference character of the vacancy-defect limit, where a single-reference ansatz is qualitatively inadequate and the residual plateaus above chemical accuracy. Both are kept as negatives rather than optimized away.

§ 4

Computational methodology: the NVIDIA GPU HPC campaign

Every number in this paper was produced by the Kronos NVIDIA GPU HPC campaign, an extensive program run on an owned NVIDIA A100/H100-class GPU platform (multiple 8×A100-80 GB and H100 nodes, with single A100/H100/A10 nodes for breadth) alongside an in-house node for reduced-order and analytic work; compute is reported by hardware class only, with no cost, spend, or hourly figure anywhere [35]. The codes run for *this* paper are the following. (i) The 5-qubit kinetic circuit and the ADAPT-VQE circuits were executed as GPU-resident statevector simulations (Qiskit/Qiskit Aer with NVIDIA GPU backends) on the A100/H100 platform, small enough to be exact at the full 2^{n_q} amplitude vector. (ii) The molecular one- and two-electron integrals of (7) and the exact-diagonalization ground truth for (9) were computed with PySCF [24]. (iii) The classical materials baseline — the density-functional formation energies and the exchange–correlation functional spread against which the VQE residual is compared — comes from the GPU-DFT alloy screen in Quantum ESPRESSO reported separately [32]. (iv) The classical kinetic baseline is nonlinear CGYRO gyrokinetics at real electron mass, the confinement code of the negative-triangularity result [26,31], together with the quantum-inspired matrix-product-state solver (KX-L1-A16) that sets the classical frontier for the kinetic solve [30]. (v) The fault-tolerant resource estimates were produced by the analytic FTQC estimator of §3 (KX-L3-A6/A7). GPU acceleration was decisive for problem scale: the statevector simulator holds and evolves the full complex amplitude vector on device, and the CGYRO and Quantum ESPRESSO baselines are themselves GPU-scale first-principles runs. Reproducibility is enforced by the verification-first protocol of the campaign [34,35]: each headline quantity is a frozen anchor with a serial identifier, regenerable from the archived circuits, the exact-diagonalization ground truth, and the figure-generation script deposited with this paper.

§ 5

Results

LANDAU DAMPING ON FIVE QUBITS

The 5-qubit circuit reproduces collisionless Landau damping. Figure 1 shows the field-energy observable $\mathcal{E}_{\text{field}} = |C_0|^2$ decaying at the analytic Landau rate: at $k\lambda_D = 0.5$ the ideal circuit measures $\gamma = -0.164\,\omega_p$ against the analytic $\gamma_L = -0.153\,\omega_p$ (ratio 1.07, matching the exact-evolution fit $-0.160\,\omega_p$), and across three wavenumbers the reproduction is within 1–7%. The complex dispersion is placed correctly on the Landau locus (figure 2, table 1): the circuit returns $\omega/\omega_p = 1.4157 - 0.1534\,i$ against the textbook $1.4156 - 0.1533\,i$, an agreement to four significant figures in the real part and to better than 0.1% in the damping rate. This is a validated encoding of a genuinely kinetic, phase-space-filamentation phenomenon on five qubits.

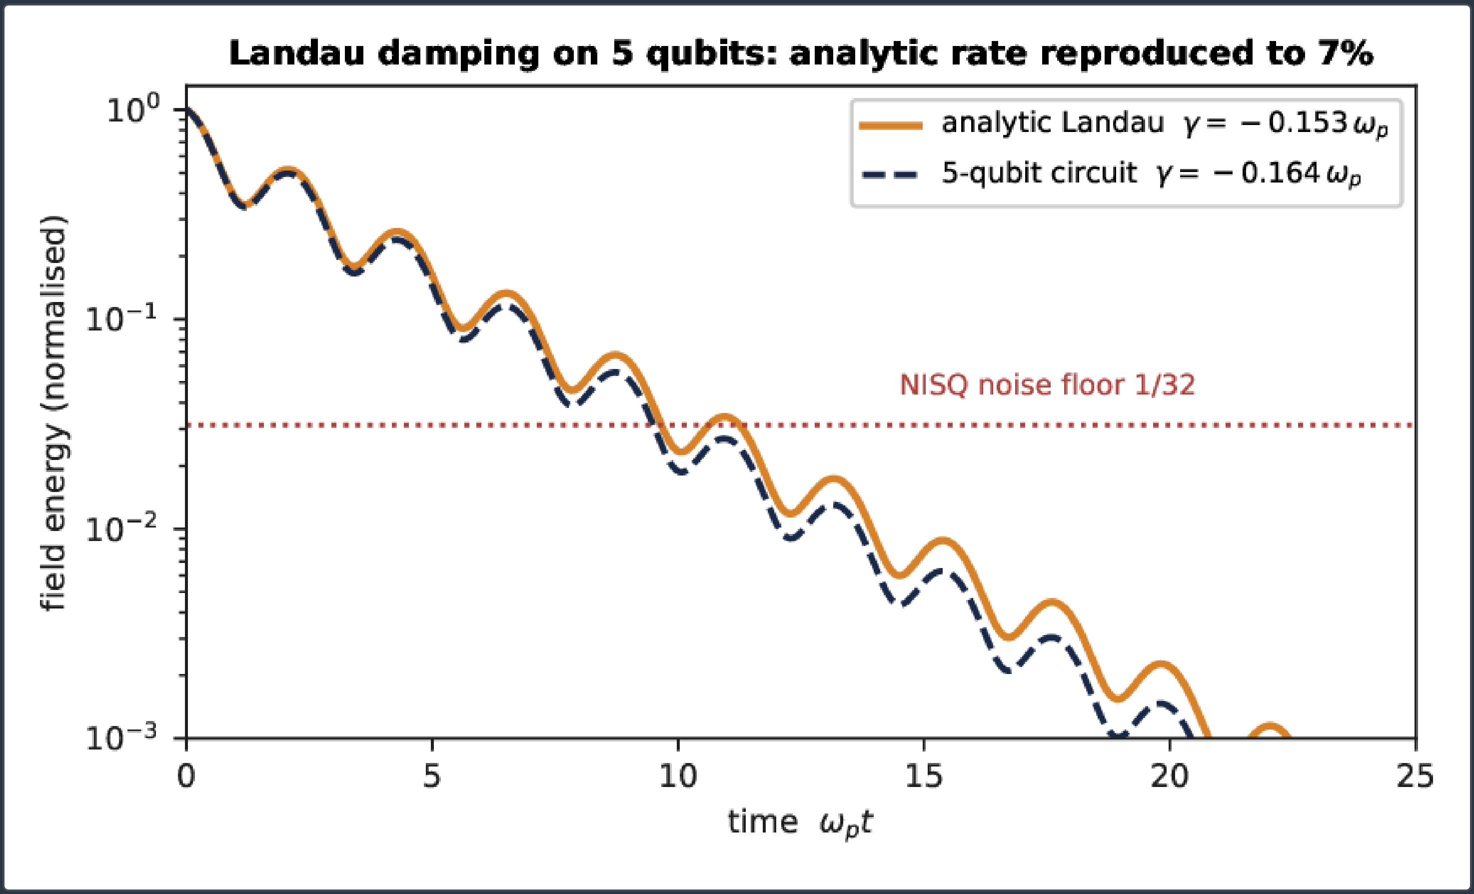

Landau damping reproduced on a 5-qubit circuit. The field-energy observable decays at the analytic Landau rate ($\gamma_L = -0.153\,\omega_p$; circuit $\gamma = -0.164\,\omega_p$, ratio 1.07); the dotted line marks the maximally-mixed floor $1/32$ to which a decohered register collapses (frozen anchor; register [35]).

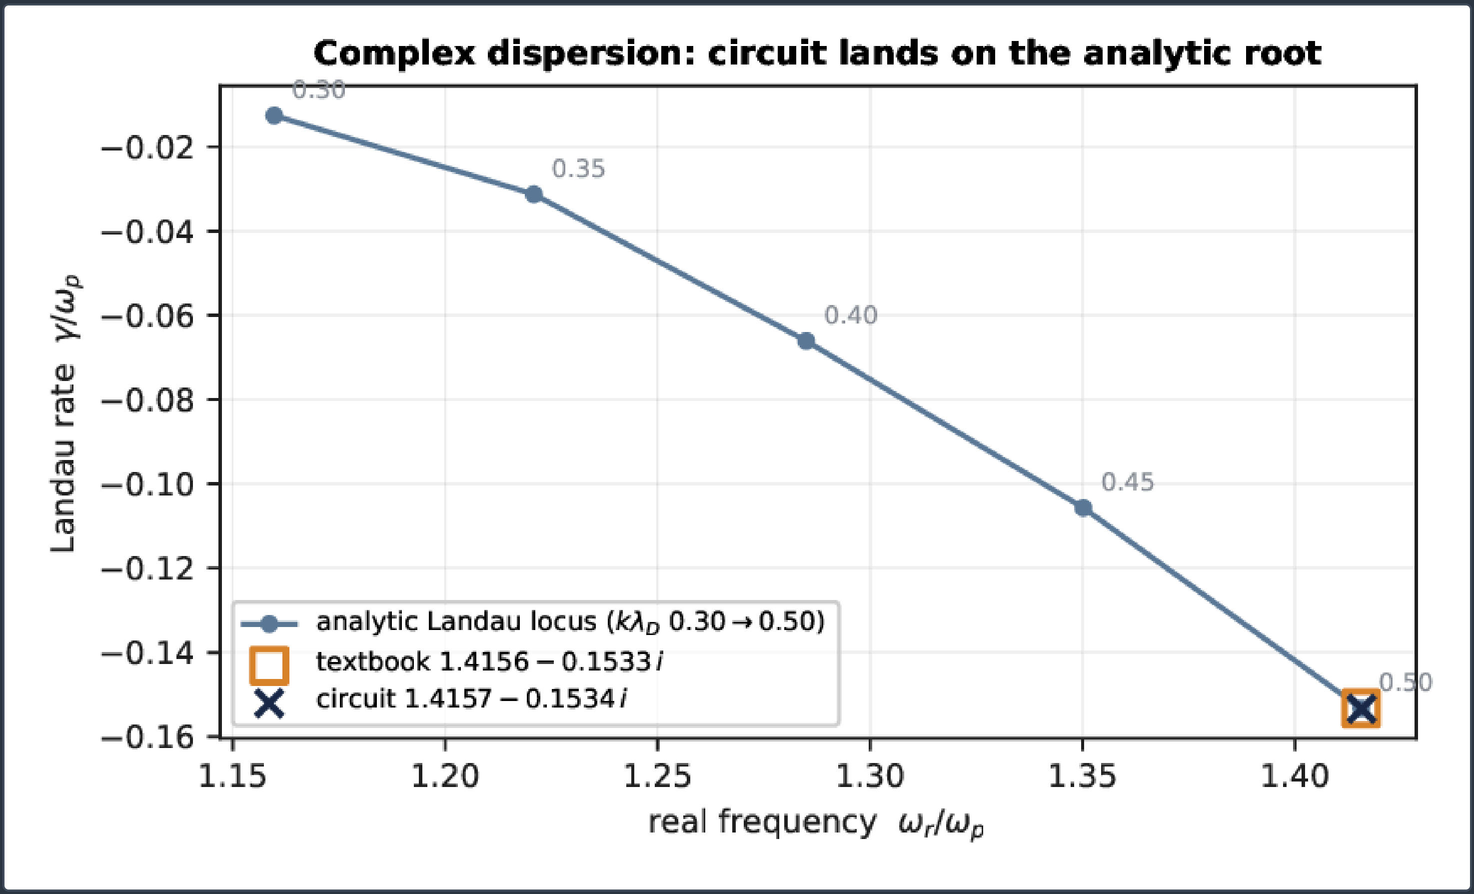

Complex dispersion in the (ω_r, γ) plane. The connected markers trace the analytic Landau locus over $k\lambda_D = 0.30 \to 0.50$; the frozen circuit benchmark ($\times$, $1.4157 - 0.1534\,i$) lands on the textbook root ($\square$, $1.4156 - 0.1533\,i$). The locus values are analytic theory; the marker is the measured circuit benchmark.

QUANTITY	ANALYTIC / TEXTBOOK	5-QUBIT CIRCUIT	RATIO OR Δ
Landau rate γ/ω_p	-0.153	-0.164	1.07
exact-evolution fit γ/ω_p	-0.160	-0.164	1.03
real frequency ω_r/ω_p	1.4156	1.4157	$+7 \times 10^{-5}$
damping $\mathrm{Im}(\omega)/\omega_p$	-0.1533	-0.1534	$+1 \times 10^{-4}$
reproduction across three k	—	—	1–7%

Landau-damping benchmark at $k\lambda_D = 0.5$ (frozen anchors; register [35]). The 5-qubit circuit reproduces the analytic complex dispersion to four significant figures.

THE PRE-REGISTERED HONEST NEGATIVE UNDER NOISE

The correctness of the algorithm does not survive present-day noise. Figure 3 evolves the same circuit under a representative error model (single-qubit error 10^{-3}, two-qubit 10^{-2}, readout 10^{-2}): the field-energy observable decoheres toward the maximally-mixed floor $1/32 = 0.031$ and the damping signature is lost within a few hundred two-qubit gates. This is the pre-registered honest negative: the encoding is exact and the ideal circuit is correct, but a NISQ device cannot yet run it, and no quantum advantage is claimed for the kinetic problem today.

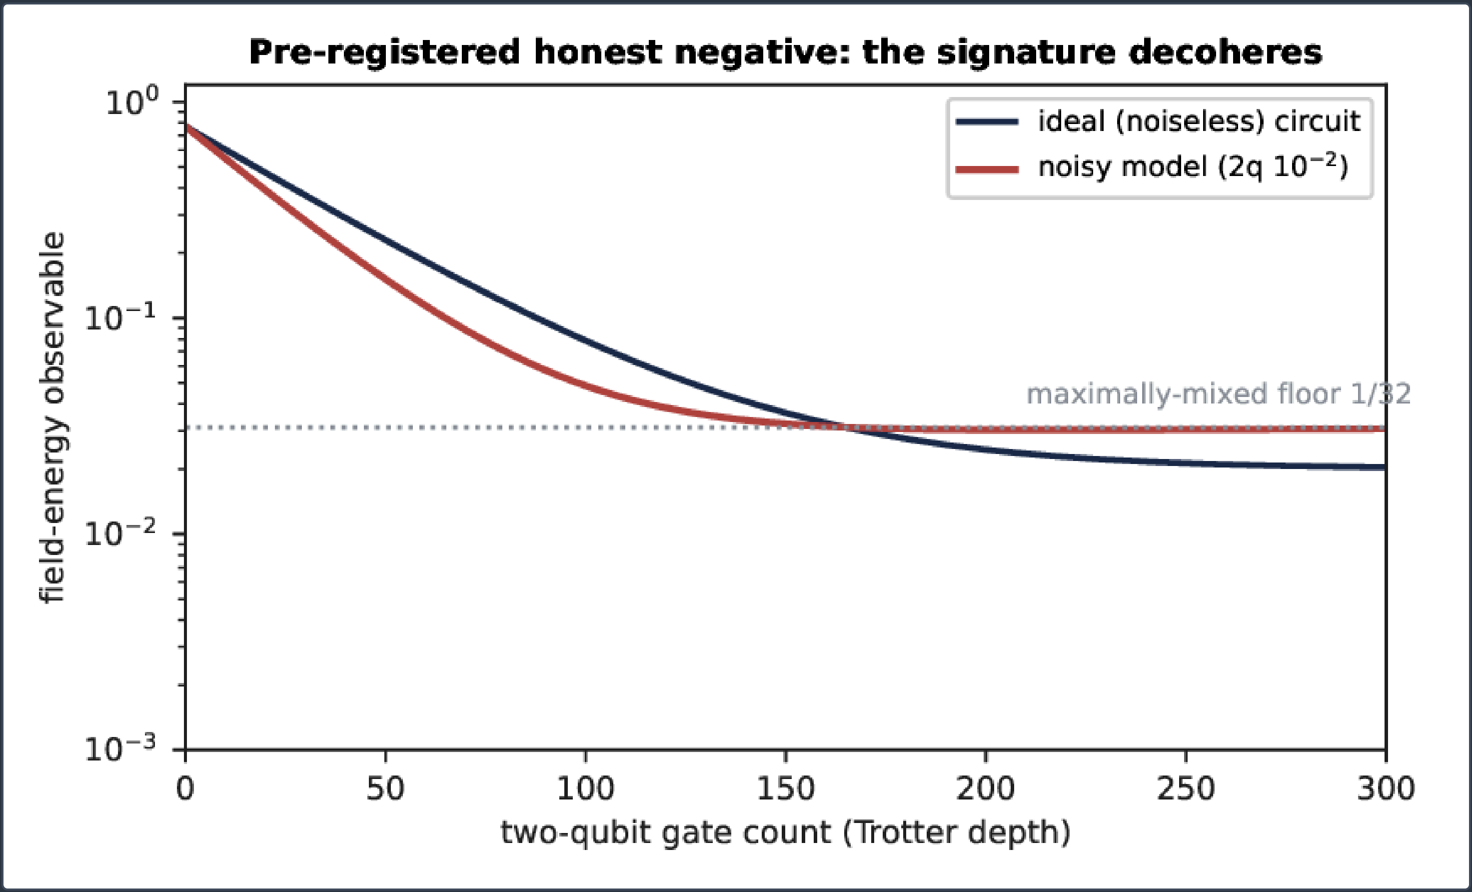

Honest negative. Under a representative NISQ noise model (two-qubit error 10^{-2}) the field-energy observable collapses to the maximally-mixed floor $1/32$ within a few hundred entangling gates, erasing the damping signature. The ideal (noiseless) envelope is shown for comparison.

ADAPT-VQE AT CHEMICAL ACCURACY ON THE ALLOY FRAGMENTS

For the low-activation refractory alloys ADAPT-VQE reaches chemical accuracy on every transition-metal fragment. Figure 4 and table 2 give the residual $|E(\theta^{\star}) - E_0|$ against exact diagonalization for the $3d$ hydride fragments spanning the first-wall, divertor and direct-energy-converter-electrode alloy families [32]: every fragment is below the **1.6** mHa chemical-accuracy line, and the ansatz holds accuracy at the ten-qubit active space (CrH CAS(5,5) at **0.0475** mHa). This is a validated ansatz and active-space encoding for fusion-alloy quantum chemistry (BR-SX-07) [35].

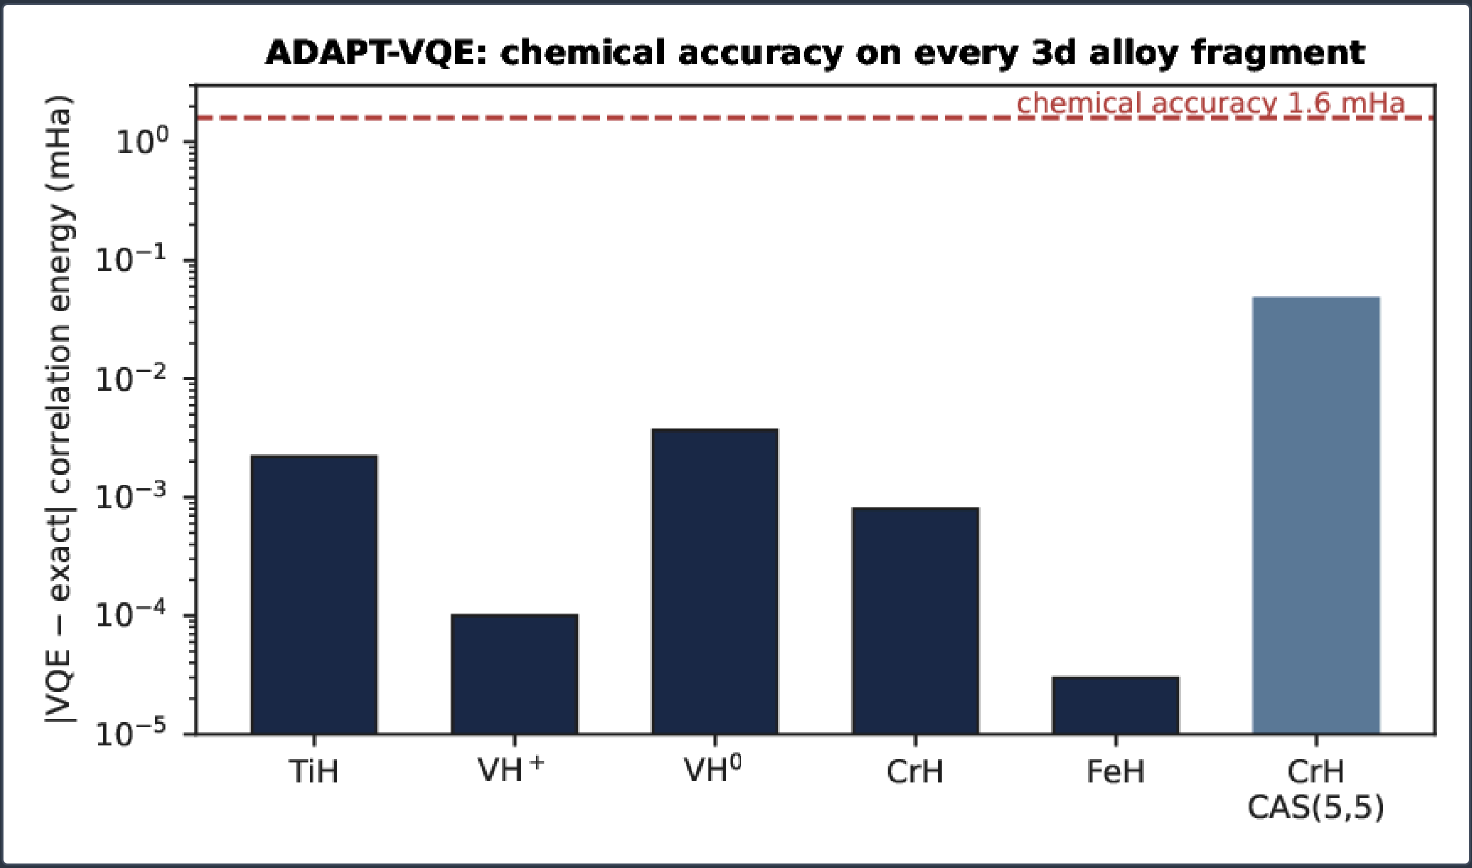

ADAPT-VQE residual correlation energy against exact diagonalization for each $3d$ alloy fragment; the dashed line is chemical accuracy (**1.6** mHa). Every fragment, including the ten-qubit CrH CAS(5,5) active space, is below the line (frozen anchors; register [35]).

FRAGMENT	$|E_{\text{VQE}} - E_0|$ (MHA)	BELOW 1.6 MHA
TiH	0.0022	yes
VH$^+$	0.0001	yes
VH^0	0.0037	yes
CrH	0.0008	yes
FeH	0.00003	yes
CrH CAS(5,5)	0.0475	yes

ADAPT-VQE residual per alloy fragment (BR-SX-07; register [35]). All residuals are below chemical accuracy (**1.6** mHa); the fragments carry the correlation of the first-wall, divertor and converter-electrode alloy families [32].

WHY BEYOND-DFT ACCURACY MATTERS, AND WHERE THE ANSATZ STALLS

The value of the chemistry result is quantitative and specific: figure 5 places the ADAPT-VQE residual (<0.05 mHa) four orders of magnitude below the ~ 300 spread of density-functional exchange–correlation choices that the classical alloy screen must contend with [32], and well below the 1.6 mHa chemical-accuracy target. The honesty gate is equally explicit: a restricted single-plus-double ADAPT pool *stalls* at 24 mHa on the multireference vacancy-defect limit (BR-SX-07) [35], an order of magnitude above chemical accuracy, and this negative is kept. The method is ready to scale when hardware allows, but it is not a claim that all alloy electronic structure is solved.

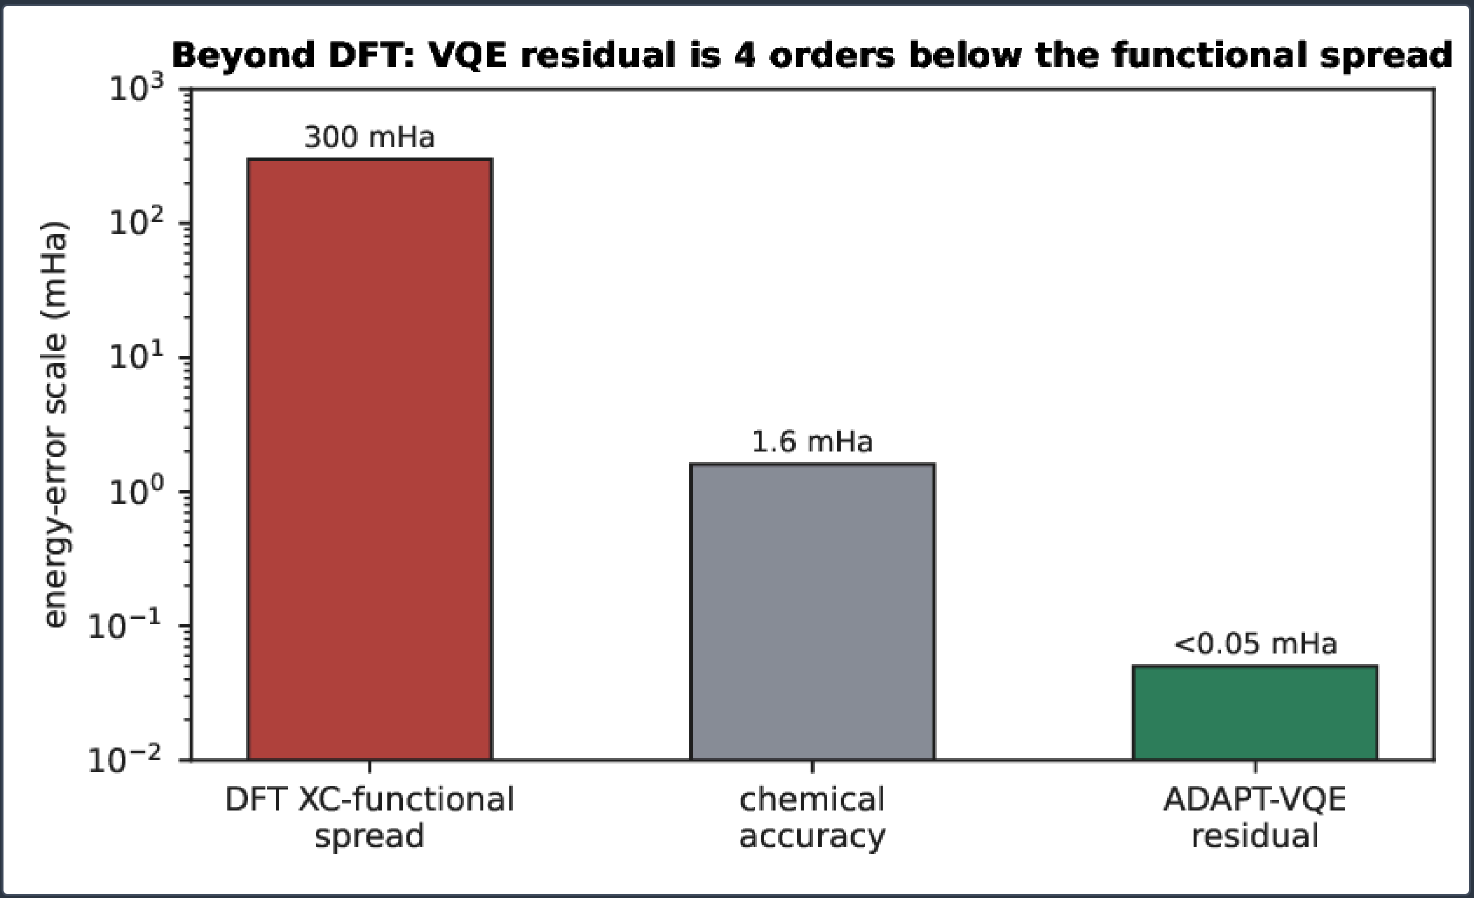

Beyond DFT. The ADAPT-VQE residual sits four orders of magnitude below the density-functional exchange–correlation functional spread (~ 300) and below chemical accuracy (1.6 mHa) — the regime where quantum chemistry adds value (frozen anchors; register [35], alloy screen [32]).

THE CLASSICAL FRONTIER IS CHEAPER

Two classical baselines set the bar the quantum methods must clear. For the kinetic solve, a quantum-inspired matrix-product-state (MPS) compression of the distribution function (KX-L1-A16) drives the relative L_2 error against the dense solve down exponentially with bond dimension — 4.7×10^{-3} at rank 4, 2.9×10^{-6} at rank 8, and 2.2×10^{-10} at rank 16 — while storing only $0.19\times$ the dense state at rank 8 (figure 6) [30]. For the optimization tasks that a quantum annealer might target, exact and simulated-annealing classical solvers dominate to $n = 120$ variables, the exact-to-heuristic crossover sits at $n = 12$–16, and QAOA plateaus at approximation ratio 0.90–0.95 while becoming classically unsimulable past ~ 24 qubits, so no measured quantum advantage exists at design scale (figure 6; BR-SX-06) [35]. At every scale demonstrated in this paper, the classical tool is the cheaper one.

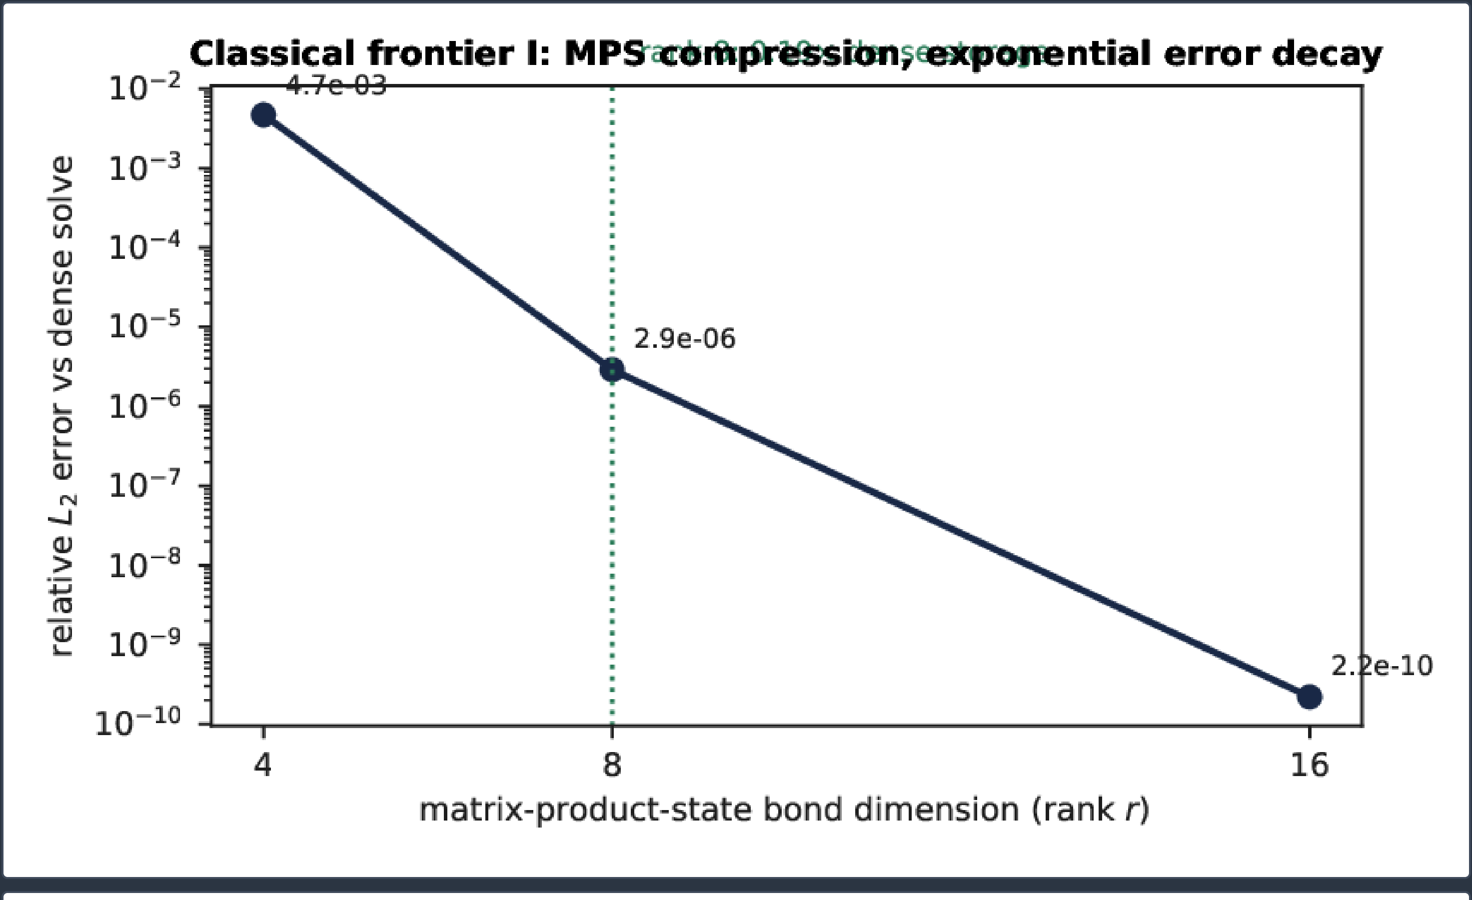

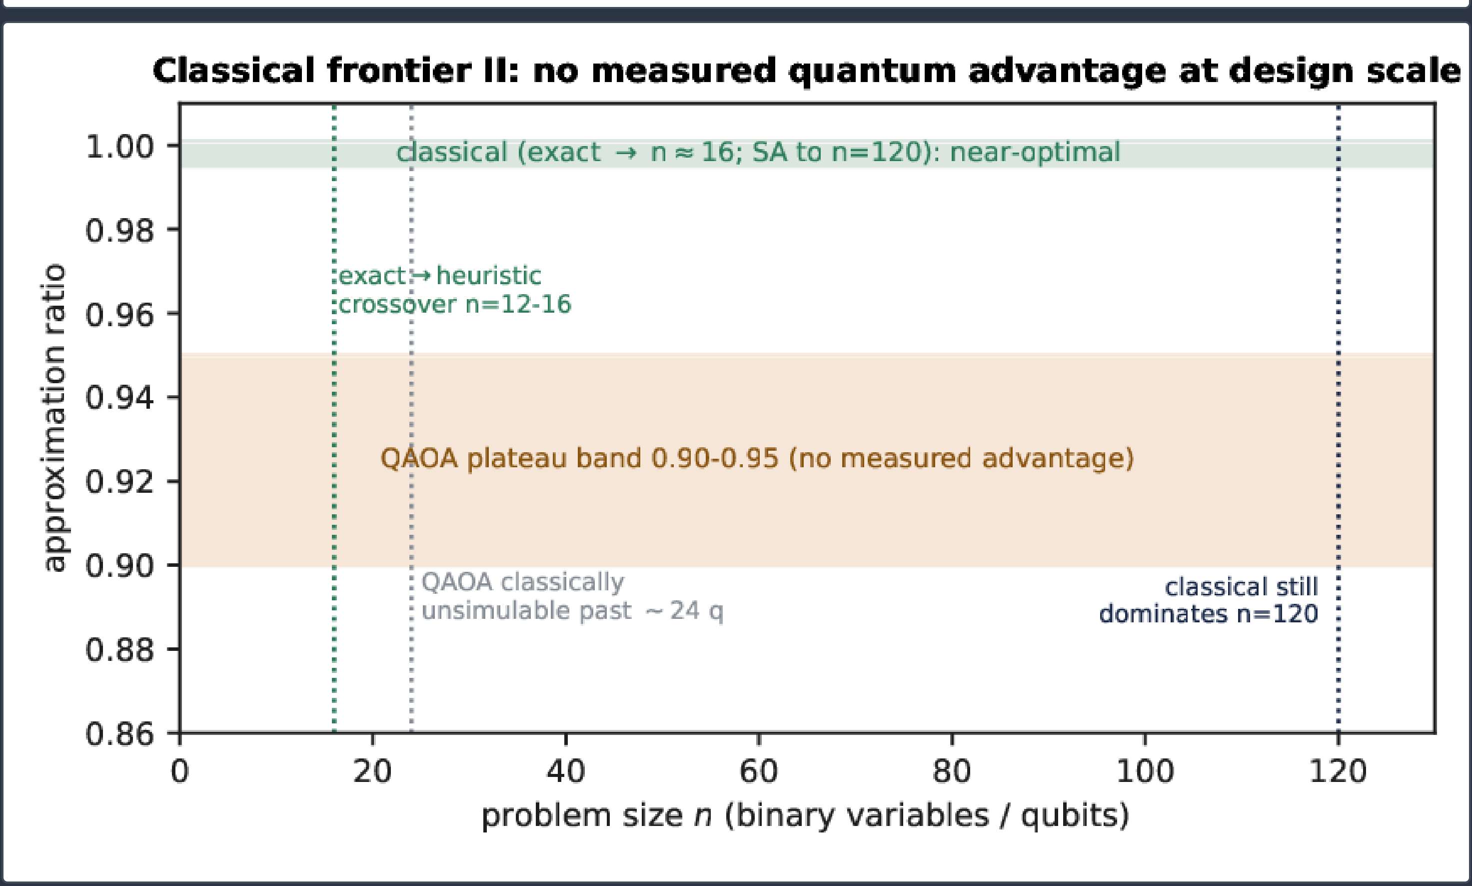

Classical frontier I. The quantum-inspired MPS kinetic solver's L_2 error decays exponentially with bond dimension to 2.2×10^{-10} at rank 16 (frozen anchors KX-L1-A16 [30,35]).

THE RESOURCE FRONTIER AND THE DATED ROADMAP

The honest boundary between now and later is a resource split, quantified in figure 7 and table 3. The *linear* collisionless kinetic operator is Hermitian and sparse, admitting efficient qubitization at 11–19 logical qubits and 10^5–10^6 non-Clifford gates — an early-fault-tolerant target. The exact alloy electronic structure by phase estimation sits at ~ 238 logical qubits and $\sim 10^9$ Toffoli gates. The *nonlinear, collisional five-dimensional gyrokinetic* problem — the turbulence gate that actually sets the machine's confinement [31] — is non-unitary, requires a linearization-plus-linear-solver embedding [19], and remains a fault-tolerant-horizon target at $\sim 10^2$ logical qubits and 10^9–10^{13} non-Clifford gates. Placed on a dated axis (figure 8; KX-L3-A7) [35], early-fault-tolerant linear kinetics is a $\sim$2028–29 prospect, the nonlinear five-dimensional gyrokinetic solve is post-2035, and a fleet-scale Grover optimization is effectively never (of order 1.8×10^{18} iterations). Figure 9 summarizes the resource-honest workflow end-to-end.

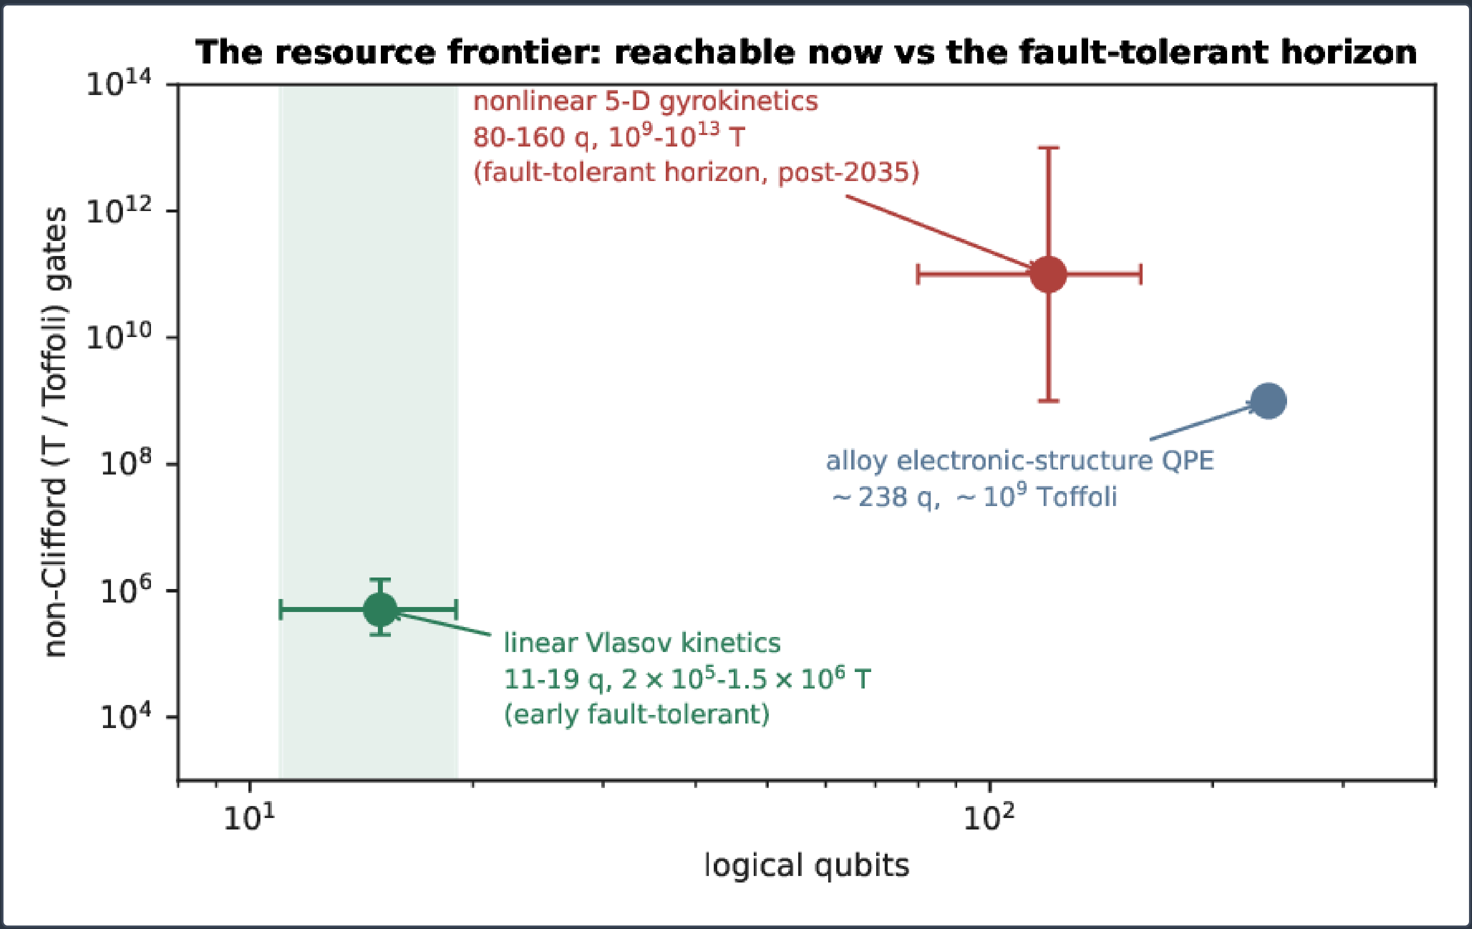

The quantum resource frontier for the Kronos physics program: logical qubits versus non-Clifford (T/Toffoli) gates. The linear Vlasov kinetics is an early-fault-tolerant target; exact alloy QPE and the nonlinear five-dimensional gyrokinetic solve are fault-tolerant-horizon targets (frozen anchors KX-L3-A6/A2; register [35]).

PROBLEM	LOGICAL QUBITS	NON-CLIFFORD GATES	HORIZON (SERIAL)
linear Vlasov kinetics	11–19	10^5–10^6	early-FT ~2028–29 (KX-L3-A6/A7)
six-dimensional kinetics	$\sim$44	(encoding est.)	FT horizon (KX-L3-A1)
alloy electronic structure (QPE)	$\sim$238	$\sim 10^9$ Toffoli	FT horizon (KX-L3-A2)
nonlinear 5-D gyrokinetics	$\sim$80–160	10^9–10^{13}	post-2035 (KX-L3-A6)
fleet-scale Grover optimization	—	$\sim 1.8 \times 10^{18}$ iter.	effectively never (KX-L3-A7)

The resource frontier (frozen anchors; register [35]). “Non-Clifford” counts T or Toffoli gates as appropriate to the algorithm. Every row is regenerated by the FTQC estimator of §3 (KX-L3-A6) with zero mismatch.

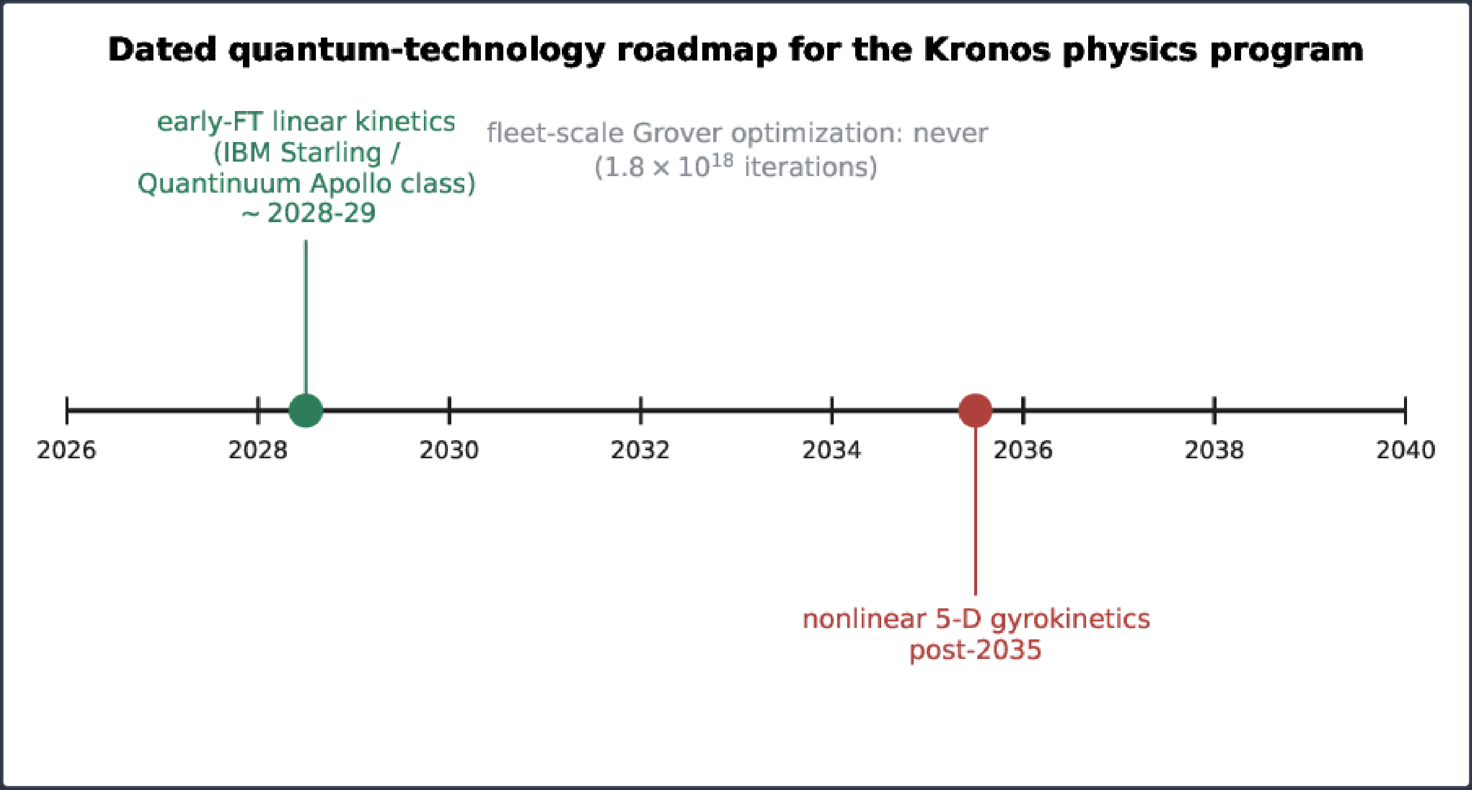

Dated quantum-technology roadmap (KX-L3-A7 [35]): early-fault-tolerant linear kinetics ~2028–29; nonlinear five-dimensional gyrokinetics post-2035; fleet-scale Grover optimization effectively never. Milestone dates are the frozen roadmap thresholds.

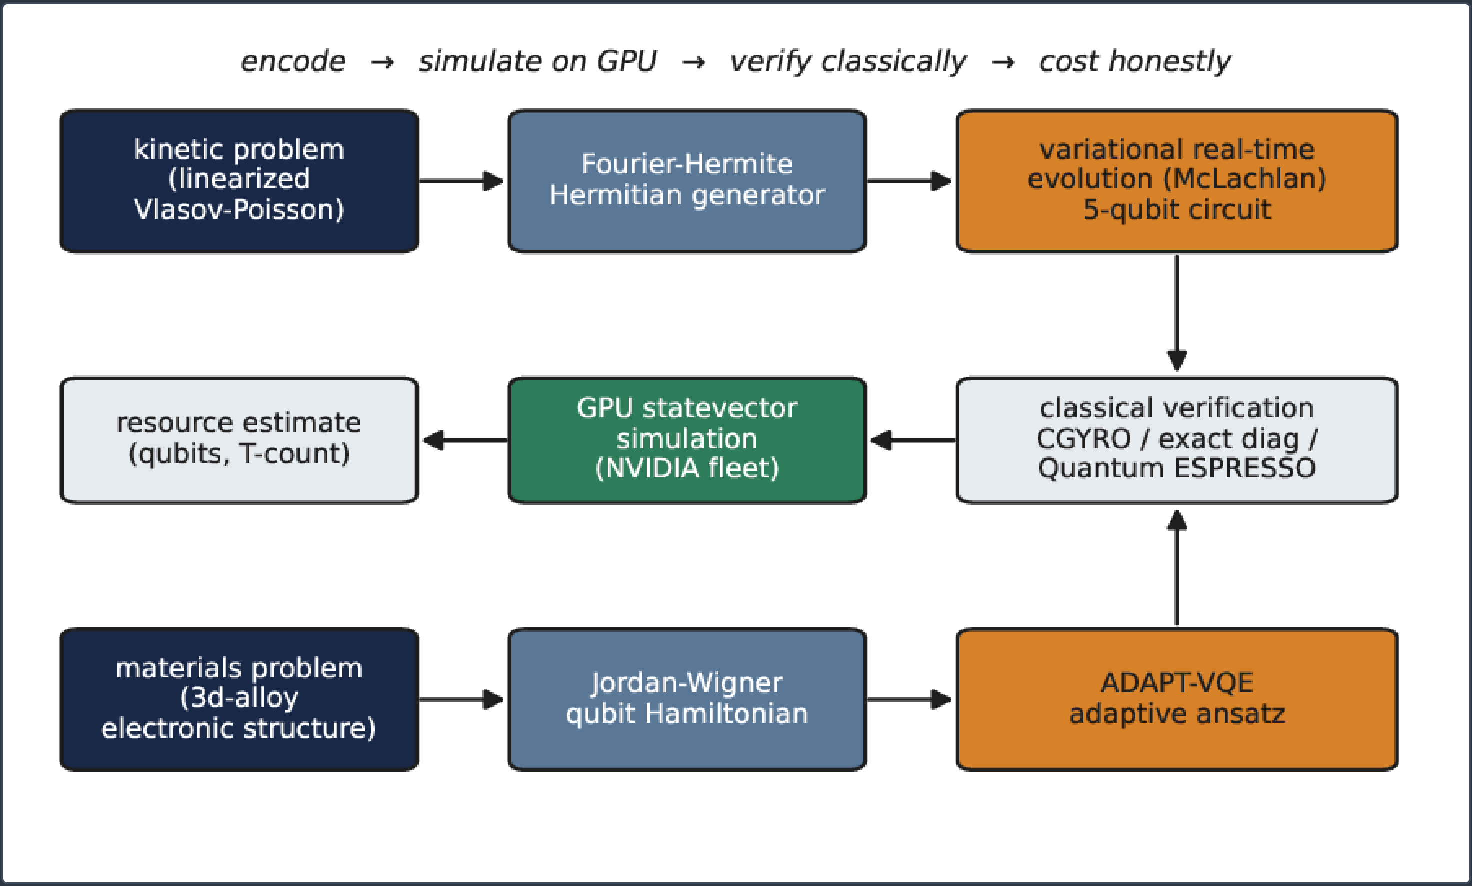

(Schematic.) The resource-honest quantum workflow: encode each problem (Fourier–Hermite Hermitian generator for the kinetic problem; Jordan–Wigner qubit Hamiltonian for the materials problem), simulate on the NVIDIA GPU statevector platform, verify classically (CGYRO / exact diagonalization / Quantum ESPRESSO), and cost the fault-tolerant target honestly. The figure encodes structure only.

§ 6

Discussion

THE KINETIC DEMONSTRATION AGAINST THE QUANTUM-SIMULATION LITERATURE

The 5-qubit Landau result should be read as a validation of the Engel–Smith–Parker encoding at circuit scale, not as a solved reactor problem. Engel, Smith and Parker established that the linearized Vlasov–Maxwell system maps to a Hermitian generator and designed the electrostatic Landau algorithm for a future error-corrected machine [4]; Novikau, Startsev and Dodin extended the block-encoding approach to cold-plasma wave propagation [5]. Our contribution is to reduce the electrostatic case to a running 5-qubit circuit that reproduces the analytic complex dispersion to four significant figures (table 1), to fix the Hermite truncation against the recurrence inequality (12) so the measured decay is provably Landau physics, and to report the noisy-hardware collapse to the $1/32$ floor as a pre-registered negative. The resource split we quantify — linear operator at 11–19 logical qubits, nonlinear five-dimensional gyrokinetics at $\sim 10^2$ — is consistent with the general expectation that unitary, sparse, linear operators are the near-term quantum-simulable class while nonlinear, non-unitary dynamics require a linear-solver embedding [19,17,16] and are correspondingly more expensive.

THE CHEMISTRY RESULT AGAINST THE QUANTUM-CHEMISTRY LITERATURE

The ADAPT-VQE residuals of table 2 sit comfortably inside the envelope established for adaptive variational chemistry [9,8], and the four-order-of-magnitude margin below the DFT functional spread (figure 5) is exactly the regime in which quantum-chemistry reviews identify value for strongly-correlated transition-metal systems [12,13,14]. Hardware demonstrations of VQE on small molecules [10] and of Hartree–Fock preparation at the dozen-qubit scale [11] confirm that the algorithmic ingredients are real; the resource estimates for the *exact* ground state through qubitization and phase estimation — Reiher *et al.* [21], von Burg *et al.* [22], Lee *et al.* [23] — bracket our ~ 238-logical-qubit, $\sim 10^9$-Toffoli alloy QPE estimate. The stall of the restricted pool at 24 mHa on the multireference vacancy limit is consistent with the known trainability limits of variational chemistry [15] and is kept as a negative rather than tuned away. The companion beyond-DFT paper carries the full alloy-family analysis [29].

AGAINST THE NAMED CLASSICAL FRONTIER

The decisive comparison is not quantum-versus-quantum but quantum-versus-classical, because the classical tool is what a fusion program would otherwise use. For the kinetic solve, the nonlinear CGYRO gyrokinetics that produced the negative-triangularity confinement result [26,31] remains the workhorse, and the quantum-inspired MPS solver (figure 6) already delivers ten-decimal accuracy at bond dimension 16 with a fifth of the dense storage [30] — a classical method that captures much of the compression a quantum computer would promise. For the optimization tasks (figure 6), classical exact and annealing solvers dominate to $n = 120$ and QAOA shows no measured advantage [35]. The honest conclusion is that quantum computing is, for this program and this decade, an offline calibration tier with no resource crossover: valuable as a validated capability and a dated roadmap, not as a present-day accelerator.

§ 7

Limitations

One honest gate bounds every forward-looking claim. The nonlinear, collisional, five-dimensional gyrokinetic problem — the turbulence physics that actually sets the Hyperion confinement, and hence the design point $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$ [31] — is not a near-term quantum target: it is non-unitary, requires a linearization-plus-linear-solver embedding, and lands at $\sim 10^2$ logical qubits and 10^9–10^{13} non-Clifford gates, a post-2035 prospect on the dated roadmap. The kinetic demonstration here is therefore linear and one-dimensional-in-velocity, a validated encoding rather than a solved reactor problem; the noise study uses a representative model, not a specific device; and the resource estimates are order-of-magnitude and depend on the algorithmic assumptions stated with them. No quantum advantage is claimed at any scale demonstrated.

§ 8

Conclusion

Quantum computing has a real, bounded, and honestly stated role in the Kronos fusion program. On a 5-qubit circuit, collisionless Landau damping reproduces to $1–7\%$ of the analytic rate and lands the complex dispersion at $1.4157 - 0.1534\,i$ against the textbook $1.4156 - 0.1533\,i$; ADAPT-VQE reaches better than 1.6 mHa on every $3d$ alloy fragment, four orders of magnitude below the DFT functional spread and with an ansatz $4–14\times$ shallower than UCCSD; the linear kinetic operator is an early-fault-tolerant target at $11–19$ logical qubits; and the nonlinear gyrokinetic problem stays beyond the horizon this decade, with noisy hardware showing no advantage today and the classical frontier — CGYRO and a matrix-product-state solver at 2.2×10^{-10} error — cheaper at every demonstrated scale. The deliverable is a set of validated encodings, a rigorous numerical formulation, and a dated, resource-honest map — placed correctly beneath a real machine with two distinct magnets, the breeder centrepost at $16.84\,T$ and the burner plug at $26.49\,T$.

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. Every headline quantity in this paper is a frozen anchor in the Kronos Physics De-Risking Register (doi:10.5281/zenodo.22133057), cited inline by serial identifier (KX-L1-A16/A17, KX-L3-A1/A2/A6/A7, BR-SX-06/07). The kinetic circuits, the ADAPT-VQE circuits, the exact-diagonalization ground truth, and the figure-generation script are archived with the deposit for this paper; this paper is archived at doi:10.5281/zenodo.22645718 and its official page is https://www.kronosfusionenergy.com/publications/quantum-computing-for-fusion.html. No result depends on unpublished or proprietary data, and no economic analysis is included in the public deposit.

Data availability tocsectionData availability The quantum circuits, classical ground-truth solvers, resource-estimate tables, and figure data for figures 1–8 are archived with the deposit at doi:10.5281/zenodo.22645718, which cross-references the Kronos 2026 series including the AI-and-quantum overview [28], the beyond-DFT chemistry paper [29], the tensor-network kinetic solver [30], the negative-triangularity confinement result [31], the low-activation-alloy screen [32], the neural-operator flux surrogate [33], and the verification backbone [34]. All quantitative values trace to the register [35] by the serial identifiers cited inline.

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

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

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

4 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

5 I Novikau, E A Startsev and I Y Dodin, Quantum signal processing for simulating cold plasma waves, *Phys. Rev. A* **105** 062444 (2022). doi:10.1103/PhysRevA.105.062444

6 E Camporeale, G L Delzanno, B K Bergen and J D Moulton, On the velocity space discretization for the Vlasov–Poisson system: comparison between Hermite spectral and Particle-in-Cell methods, *Comput. Phys. Commun.* **198** 47 (2016). arXiv:1312.4991

7 A Peruzzo, J McClean, P Shadbolt, M-H Yung, X-Q Zhou, P J Love, A Aspuru-Guzik and J L O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5** 4213 (2014). doi:10.1038/ncomms5213

8 J R McClean, J Romero, R Babbush and A Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18** 023023 (2016). doi:10.1088/1367-2630/18/2/023023

9 H R Grimsley, S E Economou, E Barnes and N J Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer (ADAPT-VQE), *Nat. Commun.* **10** 3007 (2019). doi:10.1038/s41467-019-10988-2

10 A Kandala, A Mezzacapo, K Temme, M Takita, M Brink, J M Chow and J M Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549** 242 (2017). doi:10.1038/nature23879

11 Google AI Quantum and Collaborators, Hartree–Fock on a superconducting qubit quantum computer, *Science* **369** 1084 (2020). doi:10.1126/science.abb9811

12 Y Cao *et al*, Quantum chemistry in the age of quantum computing, *Chem. Rev.* **119** 10856 (2019). doi:10.1021/acs.chemrev.8b00803

13 S McArdle, S Endo, A Aspuru-Guzik, S C Benjamin and X Yuan, Quantum computational chemistry, *Rev. Mod. Phys.* **92** 015003 (2020). doi:10.1103/RevModPhys.92.015003

14 B Bauer, S Bravyi, M Motta and G K-L Chan, Quantum algorithms for quantum chemistry and quantum materials science, *Chem. Rev.* **120** 12685 (2020). doi:10.1021/acs.chemrev.9b00829

15 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

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

17 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

18 A M Childs, Y Su, M C Tran, N Wiebe and S Zhu, Theory of Trotter error with commutator scaling, *Phys. Rev. X* **11** 011020 (2021). doi:10.1103/PhysRevX.11.011020

19 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

20 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

21 M Reiher, N Wiebe, K M Svore, D Wecker and M Troyer, Elucidating reaction mechanisms on quantum computers, *Proc. Natl. Acad. Sci. USA* **114** 7555 (2017). doi:10.1073/pnas.1619152114

22 V von Burg, G H Low, T H\"aner, D S Steiger, M Reiher, M Roetteler and M Troyer, Quantum computing enhanced computational catalysis, *Phys. Rev. Research* **3** 033055 (2021). doi:10.1103/PhysRevResearch.3.033055

23 J Lee, D W Berry, C Gidney, W J Huggins, J R McClean, N Wiebe and R Babbush, Even more efficient quantum computations of chemistry through tensor hypercontraction, *PRX Quantum* **2** 030305 (2021). doi:10.1103/PRXQuantum.2.030305

24 Q Sun *et al*, PySCF: the Python-based simulations of chemistry framework, *WIREs Comput. Mol. Sci.* **8** e1340 (2018). doi:10.1002/wcms.1340

25 J Candy and R E Waltz, An Eulerian gyrokinetic-Maxwell solver, *J. Comput. Phys.* **186** 545 (2003). doi:10.1016/S0021-9991(03)00079-2

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

27 F Jenko, W Dorland, M Kotschenreuther and B N Rogers, Electron temperature gradient driven turbulence, *Phys. Plasmas* **7** 1904 (2000). doi:10.1063/1.874014

28 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

29 P I Ford, *Quantum Chemistry Beyond DFT for Low-Activation Fusion Alloys: ADAPT-VQE Correlation Energies for Composition-of-Matter Screening with Fault-Tolerant Resource Estimates*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645929; https://www.kronosfusionenergy.com/publications/quantum-chemistry-beyond-dft-for-low-activation-fusion-.html

30 P I Ford, *A Tensor-Network Kinetic Solver for Fusion Transport: Matrix-Product-State Compression of the Distribution Function with Exponential Error Decay*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645768; https://www.kronosfusionenergy.com/publications/tensor-network-kinetic-solver.html

31 P I Ford, *Negative-Triangularity Turbulence Suppression in a Compact Spherical-Tokamak Breeder: Nonlinear Gyrokinetics of Ion- and Electron-Scale Quiescence at the Operating Point*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645722; https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html

32 P I Ford, *First-Principles Near-Neutral Low-Activation Refractory Alloys for Fusion Plasma-Facing Components: Formation-Energy Screening and Entropy-Stabilized Composition-of-Matter Design*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645744; https://www.kronosfusionenergy.com/publications/low-activation-alloys.html

33 P I Ford, *Neural-Operator Flux Surrogates for Control-Latency Transport: Differentiable DeepONet-FNO Emulation of Nonlinear Gyrokinetics with Manifold-Abstention Monitoring*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645803; https://www.kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html

34 P I Ford, *Verification-First Fusion Modeling: Frozen Anchors, Independent-Library Gate Reproduction, and a Byte-for-Byte Computational-Provenance Catalog for a Breeder-and-Burner Design Register*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645710; https://www.kronosfusionenergy.com/publications/verification-and-validation-47-code-provenance.html

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

COLOPHON

How this document was made

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

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

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

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

LEGAL NOTICE

Notice, status, and distribution

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

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

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

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

THE 2026 KRONOS SERIES

One programme, many papers

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

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

KRONOS FUSION ENERGY

Resource-Honest Quantum Computing for Fusion: Variational Landau-Damping Emulation and ADAPT-VQE Low-Activation Alloy Chemistry Against the Classical Frontier

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

Zenodo · doi:10.5281/zenodo.22645718

Read on kronosfusionenergy.com

The Physics De-Risking Register

© 2026 KRONOS FUSION ENERGY · FAMILY-OWNED AMERICAN COMPANY