KRONOS FUSION ENERGY

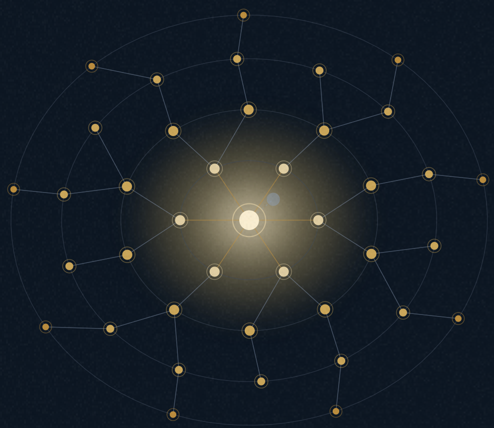

PAPER 2.91 · THE KRONOS 2026 SERIES

Quantum Optimization and Reinforcement Learning for Real-Time Plasma Control: QUBO/QAOA Actuator Scheduling and Quantum RL for Shape Control, without Claimed Advantage

Real-time control of a compact spherical-tokamak breeder poses, at several points, a combinatorial-optimization problem inside a control cycle: which actuators to schedule, how to...

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

BY P I FORD

2026 · EDITORIAL EDITION

THE OFFICIAL RECORD

What this document is, and how to read its numbers

This is a combined editorial. It gathers, in one volume, the two design studies Kronos Fusion Energy released in 2026 — a compact spherical-tokamak tritium/helium-3 breeder and a deuterium–helium-3 tandem-mirror generator — together with the intellectual-property estate, the people, and, in the internal edition, the full commercial case. The scientific text is the same text that appears in the arXiv preprints and the journal submissions; the editorial adds the apparatus a reader needs to hold the whole program at once, and nothing in the physics is altered to fit it.

THE HONESTY STANCE IS THE ASSET

Every headline quantity carries an **evidence class** — DERIVED RECOMPUTED MEASURED REQUIREMENT — and, where a number is a modelled extrapolation rather than a demonstrated value, it is labelled as such and its downside is shown next to it. Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly. This is deliberate: a diligence team will find the load-bearing assumptions, so the document surfaces them first.

TWO EDITIONS

This volume is issued in two editions. The **public edition** (light) carries the physics and the design only — no dollar, price, or per-kilogram figure appears anywhere in it. The **internal edition** (dark) adds the economics, the funding architecture, and the defense-commercial analysis, and is marked *Confidential* accordingly; it is not for public distribution. Where the commercial case rests on a modelled assumption rather than a demonstrated one, it is labelled as such and its downside is shown alongside it, on the same terms as the physics.

TITLE	The Kronos Fleet — Hyperion · Aegis · MetroVolt: A Combined Design & Commercialization Study
AUTHOR	Priyanca Ford, Founder & CEO, on behalf of Kronos Fusion Energy
EDITION	2026 Combined Editorial · first issue
COMPANION WORKS	The five-paper arXiv drop — (1) Breeder, (2) Burner, (3) REBCO magnet & tape, (4) Direct Energy Conversion, (5) AI/ML/Quantum control — with matching numbered Zenodo deposits, plus the IOP journal and IAEA-FEC submissions. Patents were filed before the drop.
NAMING	Hyperion = the breeder; Aegis (defense) and MetroVolt (data-center) = one generator in two housings. Internal mode letters (D1, L) do not appear on public surfaces.
STATUS	Conceptual design & simulation study. Nothing herein is a construction commitment.

HOW TO READ THIS VOLUME

The fleet at a glance, and the way its numbers are marked

This volume opens with *Orientation*, presents the five research papers as a bound sequence, and closes with the *Apparatus*. *Foundations* sets out the three general results both machines rest on. *Paper One* is the Hyperion breeder; *Paper Two* the Aegis / MetroVolt generator; *Papers Three, Four and Five* are the enabling science — high-field REBCO magnets and tape, direct energy conversion, and the AI / ML / quantum control stack. A *Digital Twin & 3-D Model* part follows, then the *Environmental* profile. The *Economics* and levelized cost and the *Business Case* and diligence record appear in the internal edition only; the *Simulations* proof record and the Apparatus — patent portfolio, team, limitations, and sources — close both editions.

THE FLEET AT A GLANCE

	HYPERION	AEGIS	METROVOLT
Role	Strategic-isotope foundry	Defense installation power	Campus power
Machine	**Spherical tokamak**	**Tandem mirror**	**Tandem mirror**
Fuel	Deuterium–tritium	Deuterium–helium-3	Deuterium–helium-3
Delivers	Tritium · ^{3}He · 14 MeV n	Resilient installation power	Firm campus power
Physics bar	**Fusion gain only**	Closure (gates named)	Closure + lunar ^{3}He
In the fleet	Breeds the fuel	Proves the generator	Commercial destination

HOW THE NUMBERS ARE MARKED

Every headline quantity carries an evidence class — DERIVED from first principles, RECOMPUTED from raw data, MEASURED in hardware, or a REQUIREMENT yet to be met. Financial figures carry a case label and are confined to the internal (confidential) edition. Values that moved after first publication are shown as withdrawn or restated rather than quietly changed. A capability you can evaluate is one whose limits are on the page.

ABSTRACT

Quantum Optimization and Reinforcement Learning for Real-Time Plasma Control: QUBO/QAOA Actuator Scheduling and Quantum RL for Shape Control, without Claimed Advantage

Real-time control of a compact spherical-tokamak breeder poses, at several points, a combinatorial-optimization problem inside a control cycle: which actuators to schedule, how to lay out a discrete blanket, and — the control-relevant case — how to sequence poloidal- and toroidal-field coil currents under hard constraints while tracking a shape and vertical-position target. Because these problems are naturally binary and quadratic, they are the canonical near-term target for quantum optimization, and it is a fair, testable question whether a quantum processor solves them faster or better than a classical one. This paper answers that question with a formal treatment and a pre-registered result. We derive compact quadratic-unconstrained-binary-optimization (QUBO) encodings for all three problems from first principles — pulse cadence as a number-partitioning Ising model, blanket layout as a penalty-embedded assignment, and coil-current/actuator scheduling as a fixed-point-encoded quadratic tracking objective built on the free-boundary equilibrium shape-response Jacobian — and map each to two fault-tolerant primitives, a depth-three quantum approximate optimization ansatz (QAOA, $p = 3$) and a D\"urr–Hyer Grover minimum-finder. At the frozen breeder design point ($I_p = 9.66\,\mathrm{MA}$, $Q = 3.076$, $P_{\mathrm{fus}} = 85.04\,\mathrm{MW}$, negative triangularity $\delta = -0.30$, $B_0 = 8\,\mathrm{T}$, $\kappa = 2.0$, $R_0 = 1.20\,\mathrm{m}$, aspect ratio $A = 2.5$, centrepost peak field $16.84\,\mathrm{T}$) the design-scale instances need 23–61 logical qubits and 5×10^3–5.5×10^7 T-gates, verified against the Fowler surface-code and D\"urr–Hyer formulas across 50 resource rows with zero discrepancies. On the question of whether quantum hardware *helps*, we report an explicit negative: classical simulated annealing matches or beats QAOA on all three breeder QUBOs out to $n = 120$ variables; the exact-to-heuristic crossover sits at $n = 12$–16 (exact optima in 0.03–292 s of wall time); the QAOA approximation ratio plateaus at 0.90–0.95 and the circuits become classically unverifiable past $n \approx 24$; and a fleet-scale ($n = 120$) Grover formulation would need 1.8×10^{18} oracle calls and 1.2×10^{23} T-gates — a resource boundary, not a route. We therefore place variational quantum reinforcement learning for shape and vertical-position control on a dated, benchmark-gated schedule: a quantum policy is admitted only if it meets or beats the classical controller at equal wall-clock with a control-barrier-certified safety filter held inviolate. All quantum and classical runs were carried out on the NVIDIA GPU HPC campaign. The deliverable is the honest, resource-grounded map of where quantum computing does — and does not — pay for real-time fusion control.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645903 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: QAOA · QUBO · quantum reinforcement learning · coil-current scheduling · control barrier function · fault-tolerant resource estimate · 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

REAL-TIME CONTROL OF A COMPACT SPHERICAL TOKAMAK LEAVES LITTLE TIME TO DECIDE

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 Kronos Hyperion breeder holds a plasma current $I_p = 9.66$ MA at a major radius of only $R_0 = 1.20$ m and aspect ratio $A = 2.5$, at toroidal field $B_0 = 8$ T behind a centrepost magnet whose peak field is 16.84 T, in a negative-triangularity ($\delta = -0.30$), highly elongated ($\kappa = 2.0$) regime, at fusion gain $Q = 3.076$ and fusion power $P_{\mathrm{fus}} = 85.04$ MW. Negative triangularity is chosen because it accesses good confinement without the edge pressure pedestal and the associated edge-localised-mode cycle of conventional H-mode [3,4,33]; the corresponding quiet edge is an operational asset only if the controller can actively hold the plasma inside the narrow window that keeps it quiet. In a device this compact the characteristic times of vertical instability, thermal transients, and magnet events are short, and every actuation carries a nuclear-safety consequence.

Feedback control of tokamak equilibria and profiles is a mature discipline: magnetic diagnostics feed real-time equilibrium reconstruction [5] and shape/position controllers [6], and physics-based profile simulators now run inside the control cycle [7]. More recently, learned controllers have entered the loop: a deep reinforcement-learning agent has been trained to command the TCV coil voltages and produce a variety of shapes — including negative triangularity — directly from magnetics [8]. Threaded through this progress is a family of problems that are, at their core, *combinatorial*: at a control tick the plant must choose among discrete actuator settings and coil-current increments subject to hard rate and saturation limits, a decision that is naturally posed as a constrained quadratic program over binary variables.

COMBINATORIAL CONTROL PROBLEMS ARE THE CANONICAL NEAR-TERM QUANTUM TARGET

Quadratic unconstrained binary optimization (QUBO), equivalently the Ising energy-minimization problem, is the lingua franca of near-term quantum optimization: nearly every NP-hard combinatorial problem admits a compact Ising encoding [9], quantum annealing was conceived precisely to minimize such an energy by transverse-field tunnelling [10], and the gate-model quantum approximate optimization algorithm (QAOA) targets the same objective with a shallow, variational, adiabatic-inspired ansatz [11,12]. Where a provable, if only quadratic, speedup exists it comes from Grover search [13] and its minimum-finding specialization [14]. It is therefore both fair and necessary to ask, for a fusion control stack, whether these methods can beat a classical solver on the control-critical combinatorial problems — and, if not now, on what fault-tolerant hardware and by when.

The honest backdrop to that question is that the near-term quantum-computing literature has become disciplined about advantage. Noisy intermediate-scale (NISQ) devices are powerful physics tools but not general accelerators [16]; variational quantum algorithms face trainability obstructions — barren plateaus in which gradients vanish exponentially in qubit number [17,18] — and fault-tolerant resource accounting for even celebrated algorithms runs to millions of physical qubits and astronomically many non-Clifford gates [15,19]. A responsible fusion program must therefore hold quantum optimization to the same evidentiary bar it holds any other subsystem.

CONTRIBUTION AND THE RESOURCE-HONEST POSTURE

This paper makes three contributions. First (§3), we give a from-first-principles derivation of the three governing breeder optimization problems as QUBOs, culminating in the control-relevant coil-current/actuator-scheduling QUBO built on the linearized free-boundary equilibrium shape-response Jacobian, and we state the QAOA, Grover/D\"urr–Hyer, and surface-code fault-tolerant primitives as explicit equations. Second (§4), we specify the GPU-accelerated computational methodology — the codes, the discretization and penalty embedding, the classical and quantum solvers, the statevector-simulation wall, and a 50-row resource recomputation that reproduces every fault-tolerant estimate from its stated formula with zero discrepancy. Third (§5–§6), we report a pre-registered negative on quantum advantage for these problems today and quantify it against strong classical baselines and against the published quantum-computing literature, then place variational quantum reinforcement learning (QRL) for shape and position control on a benchmark-gated schedule inside a control-barrier-certified safe set (§3.8).

The posture is deliberately resource-honest. The QUBO encodings are the deliverable: they are the interface through which any future quantum advantage would arrive, and their fault-tolerant footprints are now known to the logical qubit and the T-gate. The advantage claim is the pre-registered negative: at the scales that matter, classical methods win today, and establishing precisely where the crossover sits — and what hardware would move it — is itself the point, because it tells the program which parts of the control stack should, and should not, wait on quantum computing. This paper sits within the Kronos 2026 quantum series and shares its resource-honest frame with the umbrella AI-and-quantum study [30] and the sister variational-chemistry and kinetics work [31]; the companion quantum-algorithms paper treats the linear-algebra primitives (Harrow–Hassidim–Lloyd [20] equilibrium reconstruction, disruption state estimation, edge-turbulence simulation) that this paper does not [32]. Every quantitative claim traces to a frozen result in the Kronos Physics De-Risking Register [29] through the serial identifiers cited inline (e.g. BR-SX-06, KX-L3-A6).

§ 2

The frozen design point and the three optimization problems

The breeder operating point is a frozen anchor set held constant across the Kronos design campaign (Table 1); the burner is a physically distinct product with its own, higher-field magnet, and the two magnets are never conflated. Three optimization problems recur across breeder design and operation, and each maps to a QUBO of the standard form $\min_{x\in\{0,1\}^n} x^\top Q x$:

Pulse-cadence scheduling ("cadence") — assigning maintenance and duty windows across a fleet period to balance load, a number-partitioning / Ising problem.
Blanket layout ("blanket") — a discrete module-placement problem over the breeding blanket, a penalty-embedded assignment problem.
Coil-current / actuator scheduling ("control") — the control-relevant case: sequencing poloidal-field (PF) and toroidal-field (TF) coil currents and actuators as a constrained QUBO for live re-optimization inside a control cycle. itemize
The first two are design/operations problems solved offline; the third lives on the real-time control critical path and is the one whose classical-versus-quantum verdict actually governs the control architecture. We treat all three at both a design scale and a fleet/long-horizon scale.

SYSTEM	QUANTITY	VALUE
breeder	fusion gain Q	3.076
breeder	fusion power P_{fus}	85.04 MW
breeder	plasma current I_p	9.66 MA
breeder	triangularity δ / elongation κ	−0.30 / 2.0
breeder	major radius R_0 / aspect ratio A	1.20 m / 2.5
breeder	toroidal field B_0 / centrepost peak	8 T / 16.84 T
burner (distinct)	plug field / DEC efficiency	26.49 T / 0.49
burner (distinct)	engineering gain Q_E / neutron fraction f_n	1.318 / 5.44%

Frozen anchors used in this paper. Breeder values are the Tier-2 design point; the burner is a physically distinct product cited only to keep the two magnets separate. No economic quantity appears in this work.

§ 3

Governing equations and formal problem

THE QUBO / ISING CANONICAL FORM

A QUBO over n binary variables $x \in \{0,1\}^n$ minimizes a quadratic pseudo-Boolean objective

$$f(x) = x^{\top} Q x = \sum_{i \le j} Q_{ij}\, x_i x_j,$$

with $Q \in \mathbb{R}^{n\times n}$ upper-triangular. Because every variable is idempotent, $x_i^2 = x_i$, linear terms are carried on the diagonal of Q and no separate linear vector is needed. The change of variables $x_i = (1-s_i)/2$, $s_i \in \{-1,+1\}$, maps (1) to the classical Ising energy

$$H(s) = \sum_{i<j} J_{ij}\, s_i s_j + \sum_i h_i\, s_i + \text{const}, \qquad J_{ij} = \tfrac{1}{4} Q_{ij}, \quad h_i = -\tfrac{1}{2} Q_{ii} - \tfrac{1}{4} \sum_{j \ne i} (Q_{ij} + Q_{ji}),$$

so that minimizing f is finding the ground state of H [9]. Hard equality constraints $Ax = b$ are folded into the objective by a quadratic penalty,

$$f_\lambda(x) = x^{\top} Q x + \lambda\, \|Ax - b\|_2^2,$$

which remains a QUBO because $\|Ax-b\|_2^2$ is quadratic in x. The penalty weight λ must be large enough that violating any constraint costs more than the largest feasible objective improvement; writing $\Delta f_{\text{feas}} = \max_{x,x' \in \mathcal{F}} |f(x) - f(x')|$ for the feasible objective span and $g_{\min}$ for the smallest nonzero constraint residual, a sufficient condition is

$$\lambda \;>\; \frac{\Delta f_{\text{feas}}}{g_{\min}^2},$$

so that the global minimizer of f_λ is feasible and optimal for the original problem [9]. Equations (1)–(4) are the common substrate; the three problems differ only in how Q, A, and b are built.

PULSE-CADENCE SCHEDULING AS NUMBER PARTITIONING

Given a set of maintenance/duty tasks with loads $\{a_j\}_{j=1}^n$ to be split between two windows of a fleet period so as to balance the total load, assign $s_j \in \{-1,+1\}$ to windows and minimize the squared imbalance,

$$H_{\text{cad}}(s) = \Big(\sum_{j=1}^{n} a_j\, s_j\Big)^2 = \sum_{j \ne k} a_j a_k\, s_j s_k + \sum_j a_j^2,$$

a fully connected Ising model with couplings $J_{jk} = a_j a_k$ — the canonical number-partitioning problem [9]. The multi-window generalization uses one-hot indicators $x_{u,w} = 1$ iff unit u is serviced in window w, with objective $\sum_w \big(\sum_u L_u x_{u,w} - \bar{L}\big)^2$ balancing loads L_u across windows against the mean $\bar{L}$, subject to $\sum_w x_{u,w} = 1$ enforced through (3); this yields $n = UW$ binary variables for U units and W windows.

BLANKET LAYOUT AS PENALTY-EMBEDDED ASSIGNMENT

Let $x_{m,k} = 1$ iff breeding module type m occupies blanket segment k. Maximizing a coverage/breeding benefit $\sum_{m,k} b_{m,k}\, x_{m,k}$ with a quadratic adjacency interaction $\sum c_{m,k;m'k'} x_{m,k} x_{m',k'}$, subject to at most one module per segment ($\sum_m x_{m,k} \le 1$) and each module placed ($\sum_k x_{m,k} = 1$), is a constrained assignment problem. Converting the inequality with a binary slack and folding both constraints through (3) gives a QUBO in $n = MK$ variables for M module types and K segments. The blanket QUBO is an offline design problem; it appears here because it shares the encoding machinery and because its fault-tolerant footprint is representative.

COIL-CURRENT / ACTUATOR SCHEDULING: THE CONTROL-RELEVANT QUBO

The control-critical case is the live sequencing of PF/TF coil-current increments to track a plasma shape and vertical-position target. To first order the shape descriptors $g \in \mathbb{R}^p$ (control gaps, X-point coordinates, elongation, triangularity) respond linearly to coil-current increments $\delta I \in \mathbb{R}^C$ about the operating point through the free-boundary equilibrium shape-response Jacobian

$$g(\delta I) = g_0 + M\,\delta I, \qquad M = \frac{\partial g}{\partial I}\bigg|_{\text{op}} \in \mathbb{R}^{p\times C},$$

where M is computed from the axisymmetric free-boundary equilibrium (below) and g_0 is the present shape estimate. The controller minimizes a weighted tracking cost with actuator regularization,

$$J(\delta I) = \left\|W_g\big(g_0 + M\,\delta I - g^\star\big)\right\|_2^2 + \delta I^\top R\,\delta I, \qquad \text{s.t. } |\delta I_c| \le \Delta I_c^{\max},\ \ \delta I \in \mathcal{U},$$

with target $g^\star$, output weights W_g, and control weights $R \succ 0$. Equation (7) is the continuous quadratic program solved classically by the real-time optimization backbone [38] and, for the equilibrium currents themselves, by the PF-coil optimizer [37]. To pose the *same* objective as a QUBO for live quantum re-optimization, each increment is fixed-point encoded in B bits,

$$\delta I_c = \delta I_c^{\min} + \Delta_c \sum_{b=0}^{B-1} 2^b\, x_{c,b}, \qquad \Delta_c = \frac{\delta I_c^{\max} - \delta I_c^{\min}}{2^B - 1},$$

so the rate/saturation limits are enforced by the bit budget itself. Substituting the affine map (8) into (7) and collecting $x = \mathrm{vec}(x_{c,b}) \in \{0,1\}^n$, $n = C\,B$, gives

$$J(x) = x^\top \tilde{Q}\, x + \tilde{q}^\top x + \text{const}, \qquad \tilde{Q} = E^\top\big(M^\top W_g^\top W_g M + R\big)E,$$

where E is the linear encoding matrix implied by (8); absorbing $\tilde{q}$ onto the diagonal recovers the canonical form (1). Equation (9) is the object whose classical-versus-quantum verdict decides whether the control loop should ever wait on a quantum processor. The response matrix M derives from the Grad–Shafranov free-boundary equilibrium for the poloidal flux $\psi(R, Z)$,

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

with pressure $p(\psi)$ and poloidal-current function $F(\psi)$, solved with a free-boundary code [28]; differentiating the converged mapping $g(\textbf{coil currents})$ gives M.

QAOA FORMULATION AND THE APPROXIMATION RATIO

For a QUBO in Ising form the cost operator is diagonal in the computational basis, obtained by promoting each spin to a Pauli-Z,

$$H_C = \sum_{i<j} J_{ij}\, Z_i Z_j + \sum_i h_i\, Z_i,$$

and the QAOA mixing operator is the transverse field $H_B = \sum_i X_i$ [11]. A depth-p ansatz alternates p cost and mixer layers on the uniform superposition,

$$|\psi_p(\quad\gamma,\quad\beta)\rangle = \prod_{\ell=1}^{p} e^{-i\beta_\ell H_B}\, e^{-i\gamma_\ell H_C}\, |+\rangle^{\otimes n},$$

parameterized by $2p$ angles $(\quad\gamma,\quad\beta)$. A classical outer optimizer minimizes the energy expectation

$$F_p(\quad\gamma,\quad\beta) = \langle\psi_p(\quad\gamma,\quad\beta)|\, H_C\, |\psi_p(\quad\gamma,\quad\beta)\rangle.$$

With $C_{\min}$ the true minimum and $C_{\max}$ the maximum of the cost, the approximation ratio

$$r_p = \frac{C_{\max} - \min_{\quad\gamma,\quad\beta} F_p(\quad\gamma,\quad\beta)}{C_{\max} - C_{\min}} \in [0,1]$$

measures the fraction of the optimal gap captured; $r_p \to 1$ monotonically as $p \to \infty$ (the adiabatic limit), but at fixed shallow depth it is bounded away from unity, and the landscape suffers cost concentration and, for expressive ans\"atze, barren plateaus [17,18]. On the three breeder QUBOs at $p = 3$, r_p plateaus at 0.90–0.95 (§5).

GROVER / D\"URR–H

yer minimum finding The only route to a *provable* speedup for these problems is amplitude amplification. Grover search finds a marked element among $N = 2^n$ candidates in $O(\sqrt{N})$ oracle calls [13], and the D\"urr–Hyer procedure turns this into minimum finding by repeatedly lowering a threshold [14]. The register's frozen iteration count for the minimum-finder is

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

a quadratic improvement over the 2^n of exhaustive search. Each iteration applies the marking oracle U_f and the diffusion operator $D = 2|s\rangle\langle s| - \mathbb{I}$ (with $|s\rangle = |+\rangle^{\otimes n}$); both contain an n-controlled phase whose fault-tolerant decomposition costs $O(n)$ non-Clifford gates, so the total non-Clifford budget is

$$N_T = N_{\text{iter}} \cdot c_T(n),$$

with c_T the per-iteration T-count. At the fleet scale $n = 120$, equation (15) gives $N_{\text{iter}} = 1.6 \times 2^{60} \approx 1.8 \times 10^{18}$ and equation (16) gives $N_T \approx 1.2 \times 10^{23}$ (§5). No foreseeable hardware executes 10^{18} *sequential* oracle calls inside a control-relevant time; equation (15) is thus a boundary, not a route.

FAULT-TOLERANT FOOTPRINT: THE SURFACE-CODE OVERHEAD

Neither QAOA nor Grover runs on raw physical qubits; each logical qubit must be encoded. In the rotated planar surface code at code distance d a logical qubit occupies

$$n_{\text{phys/logical}} = 2d^2\,(-1) \approx 2d^2, \qquad N_{\text{phys}} = N_{\text{logical}} \cdot 2d^2,$$

physical qubits [15]. The required distance follows from the per-cycle logical error rate, which for physical error p below threshold $p_{\text{th}} \approx 10^{-2}$ scales as

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

so that, to keep the total logical error over a computation of space–time volume $V = N_{\text{logical}}\,N_{\text{cyc}}$ below a target ε, one selects the smallest d with $V\,p_L \le \varepsilon$. For the design-scale QUBOs this yields a small d and hence the tens-to-hundreds of physical qubits tabulated in §5; for the fleet Grover it is the T-count of equation (16), not the qubit count, that closes the case.

VARIATIONAL QUANTUM REINFORCEMENT LEARNING INSIDE A CERTIFIED SAFE SET

The second half of the quantum-control question is reinforcement learning: whether a variational quantum policy circuit can control plasma shape and vertical position at least as well as a classical RL controller. We treat this as a designed, benchmark-gated element, and give its formal object here. A parameterized-quantum-circuit (PQC) policy encodes the state s by a data-embedding unitary $U_{\text{enc}}(s)$, applies trainable layers $U(\theta)$, and reads out a set of observables $\{O_a\}$ indexed by control actions a; the policy is the softmax

$$\pi_{\theta}(a \mid s) = \frac{\exp\big(\beta\,\langle O_a\rangle_{s,\theta}\big)}{\sum_{a'}\exp\big(\beta\,\langle O_{a'}\rangle_{s,\theta}\big)}, \qquad \langle O_a\rangle_{s,\theta} = \langle 0|\,U_{\text{enc}}^{\dagger}(s)U^{\dagger}(\theta)\,O_a\,U(\theta)U_{\text{enc}}(s)\,|0\rangle,$$

with inverse temperature β [24,25]. Training uses the policy-gradient (REINFORCE) estimator

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\tau\sim\pi_{\theta}}\Big[\textstyle\sum_t \nabla_{\theta}\log\pi_{\theta}(a_t \mid s_t)\,G_t\Big], \qquad G_t = \textstyle\sum_{t'\ge t}\gamma^{\,t'-t}r_{t'},$$

where the expectation-value gradients are evaluated exactly on hardware by the parameter-shift rule [22,23]

$$\frac{\partial\langle O\rangle}{\partial\theta_j} = \tfrac{1}{2}\Big[\langle O\rangle_{\theta+\frac{\pi}{2}e_j} - \langle O\rangle_{\theta-\frac{\pi}{2}e_j}\Big].$$

Crucially, learning is confined to a certified safe set. With a control barrier function (CBF) $h(x)$ defining the safe set $\mathcal{C} = \{x : h(x) \ge 0\}$, forward invariance is guaranteed when the applied control keeps $\dot{h} \ge -\alpha(h)$ [26]; the safety filter projects the (quantum or classical) policy action onto the admissible set,

$$u^{\star}(s) = \arg\min_{u\in\mathcal{U}} \big\|u - \pi_{\theta}(\cdot \mid s)\big\|_2^2 \quad \text{s.t.} \quad \nabla h(x)\cdot f(x,u) \;\ge\; -\,\alpha\,h(x),$$

the same forward-invariant clamp that governs the Kronos control stack [35,36]. The acceptance test is fixed in advance: a quantum policy is admitted only if its expected return meets or beats the classical RL baseline at equal wall-clock while (22) holds for every step,

$$\text{admit } \pi^Q_\theta \iff J_{\text{ret}}(\pi^Q_\theta) \geq J_{\text{ret}}(\pi^{\text{cl}}) \wedge h(x_t) \geq 0\,\forall t \wedge \text{(equal wall-clock)}.$$

Because the classical RL controller is itself the operating baseline [34], equation (23) is a controlled A/B comparison with a clear, falsifiable bar — and the trainability caveats of variational circuits [17,18] are exactly why the bar, rather than a claim, is what we commit to.

§ 4

Computational methodology: the NVIDIA GPU HPC campaign

CODES, ROLES, AND GPU ACCELERATION

All quantum and classical runs supporting this paper were carried out on the NVIDIA GPU HPC campaign (A100/H100/A10-class accelerators), with a local node reserved for the resource-recomputation verification so the fastest checks never contend with the simulation fleet. Table 2 lists the codes and their roles. QAOA circuits (12) are simulated with a GPU statevector backend (Qiskit with the NVIDIA cuQuantum cuStateVec library), whose amplitude-parallel kernels give near-linear speedups over CPU statevector at the qubit counts of interest; tensor-network contraction (cuTensorNet) is used to cross-check shallow-circuit expectation values where the treewidth permits. The classical baselines — exact enumeration and simulated annealing [21] — are GPU-parallelized across independent restarts and spin-flip proposals. The equilibrium shape-response Jacobian M of equation (6) is assembled from a free-boundary Grad–Shafranov solve [28]; the negative-triangularity transport that fixes the design point derives from nonlinear gyrokinetics [27,33]. Resource estimation (15)–(17) is a deterministic recomputation harness with no stochastic component.

CODE / LIBRARY	ROLE IN THIS PAPER	ACCELERATION
Qiskit + cuQuantum cuStateVec	QAOA statevector simulation, eq. (12)–(13)	GPU statevector
cuQuantum cuTensorNet	shallow-circuit expectation cross-check	GPU tensor-network
simulated annealing (custom)	classical baseline, eq. (2)	GPU multi-restart
exact enumeration	ground-truth optimum for $n \leq 24$	GPU/CPU
free-boundary GS solver [28]	shape-response Jacobian M, eq. (6)	CPU/GPU
resource-recomputation harness	FT footprints, eq. (15)–(17)	CPU (verification)

Codes run for this paper and their roles. Compute was provided by the NVIDIA GPU HPC campaign; no cost, budget, or economic quantity is reported.

DISCRETIZATION, PENALTY EMBEDDING, AND PROBLEM SCALE

Each problem is reduced to a QUBO by the encodings of §3: one-hot indicators for cadence and blanket, fixed-point bit encoding (8) for the control increments, and the penalty embedding (3) with weight chosen by (4) for all constraints. The design-scale control QUBO uses C PF/TF channels at B bits each; at the design scale this places the instances at 23–61 logical qubits, and at the fleet/long-horizon scale at $n = 120$–125 (Table 3). Penalty weights were validated by confirming that the unconstrained minimizer of f_λ is feasible for every instance at the chosen λ.

SOLVERS AND THE STATEVECTOR WALL

The classical baseline is Metropolis simulated annealing with acceptance probability $P = \min\{1, \exp(-\Delta H/T_k)\}$ on a geometric cooling schedule $T_k = T_0\,\alpha^k$ $(0 < \alpha < 1)$, restarted many times in parallel; for $n \leq 16$ its output is checked against exact enumeration and agrees on the global optimum. The QAOA outer loop (13) is minimized with a derivative-free optimizer (COBYLA) and, where gradients are used, the parameter-shift rule (21). Noiseless statevector simulation stores 2^n complex amplitudes, i.e. $2^n \times 16$ bytes in double precision — ~ 0.27 GB at $n = 24$, ~ 17 GB at $n = 30$ — so full-fidelity QAOA simulation, and the brute-force verification of QAOA optimality against ground truth, both become impractical past $n \approx 24$. This is the "verification wall": above it, the regime where a quantum device could in principle add value is exactly the regime where its solution quality can no longer be checked against the true optimum.

VERIFICATION AND REPRODUCIBILITY

The fault-tolerant estimates are not asserted but recomputed. A 50-row resource table — 36 optimization rows plus companion kinetics and sensing rows — is regenerated from the stated formulas: D\"urr–Hyer iterations (15) and the Fowler surface-code footprint (17). All 50 rows reproduced their register values with *zero* discrepancies (KX-L3-A6), an independent-recomputation gate of the kind the Kronos verification-first methodology applies across the campaign [39]. The recomputation harness, the QUBO encoders, and the QAOA-versus-classical benchmark are deposited with the paper.

§ 5

Results

FAULT-TOLERANT FOOTPRINTS ARE REAL, CHECKED, AND SMALL

Table 3 and Figure 1 give the fault-tolerant footprints of the three breeder QUBOs. The design-scale instances need only **23–61** logical qubits and, at a practical code distance, a few tens to a bit over one hundred physical qubits; even the fleet/long-horizon cases stay under a few hundred physical qubits. The T-counts at design scale span 5×10^3–5.5×10^7. Every entry was reproduced from equations (15)–(17) with zero discrepancy (KX-L3-A6). The encodings are the deliverable: they are compact, auditable, and now fixed to the logical qubit and the T-gate.

PROBLEM	SCALE	QAOA LOGICAL	QAOA PHYS.	GROVER PHYS.	QAOA T-COUNT
cadence	design	23–61	60	122	$\sim 5 \times 10^3$–5.5×10^7
blanket	design	23–61	46	108	$\sim 5 \times 10^3$–5.5×10^7
control	design	23–61	46	108	$\sim 5 \times 10^3$–5.5×10^7
cadence	fleet	125	250	312	$\sim 1 \times 10^5$ (heuristic)
control	long	125	202	264	$\sim 1 \times 10^5$ (heuristic)

The three breeder QUBOs and their fault-tolerant footprint. Physical qubits from the surface-code formula (17) at a practical code distance; T-counts at design scale. Every row reproduced from the stated formulas with zero discrepancy (KX-L3-A6).

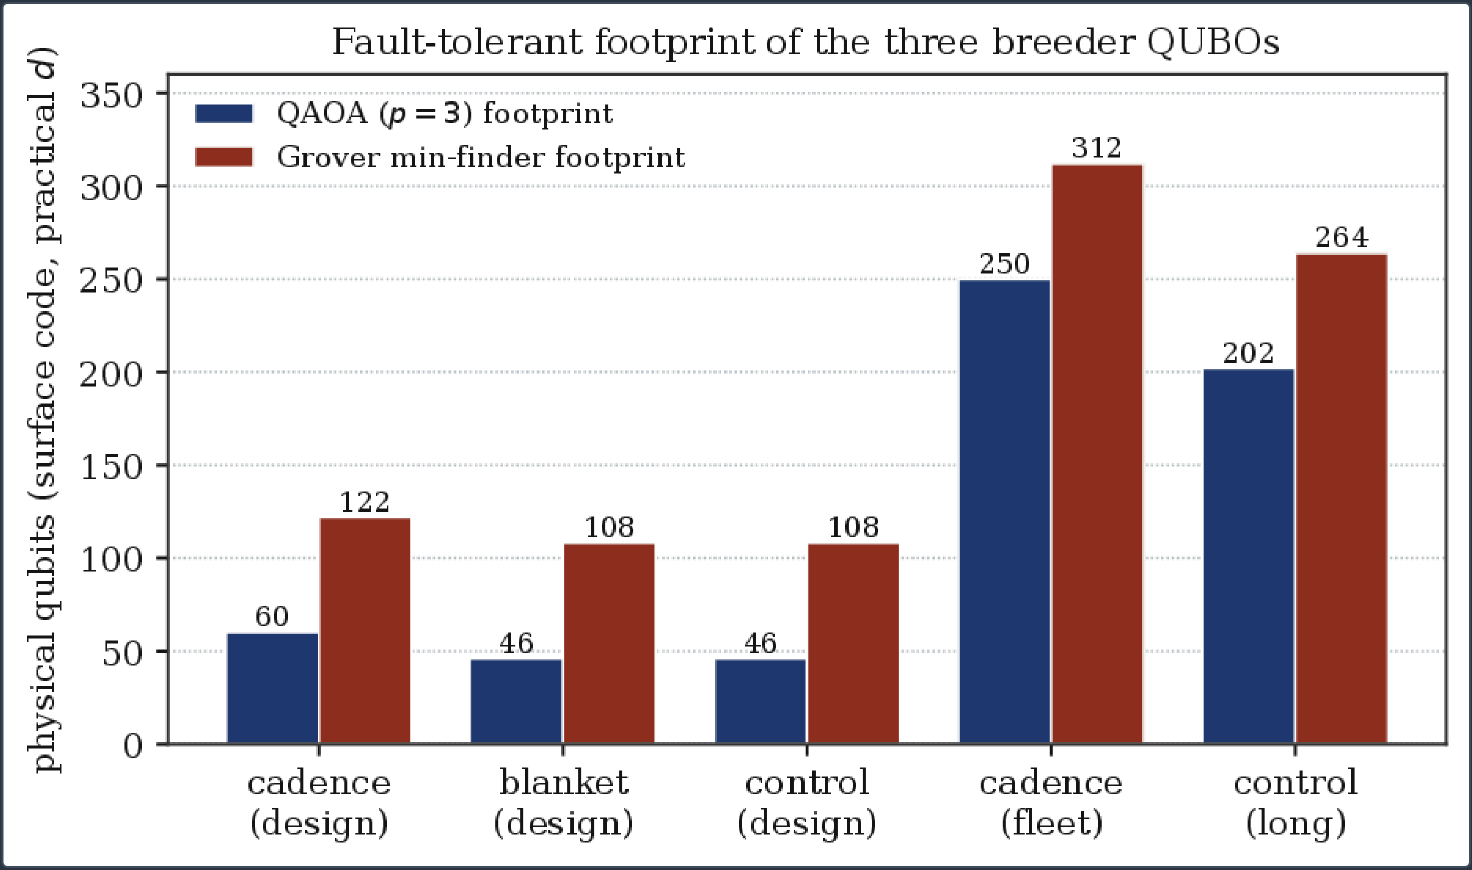

Physical-qubit footprint (surface code, practical code distance) for the three breeder QUBOs under a depth-three QAOA ansatz and a Grover minimum-finder, from equation (17). Design-scale instances fit in tens of physical qubits; even the fleet/long-horizon cases stay under a few hundred. Values reproduced from the resource formulas (KX-L3-A6).

THE RESOURCE BOUNDARY AT FLEET SCALE

Compactness at design scale does not extend to the fleet-scale provable-speedup route. Figure 2 places the T-budgets side by side: design-scale QAOA sits within early fault-tolerant reach, the fleet QAOA heuristic needs $\sim 10^5$ T-gates, and the fleet-scale Grover minimum-finder — the only formulation carrying a provable quadratic speedup — needs 1.2×10^{23} T-gates. Figure 3 shows why: the iteration count (15) reaches 1.8×10^{18} oracle calls at $n = 120$, sequential by construction. Figure 3 shows that the qubit overhead itself, growing as $2d^2$ per logical qubit (17), is modest at these logical counts — so it is the non-Clifford *time*, equation (16), and not the qubit *space*, that renders fleet Grover a boundary rather than a route.

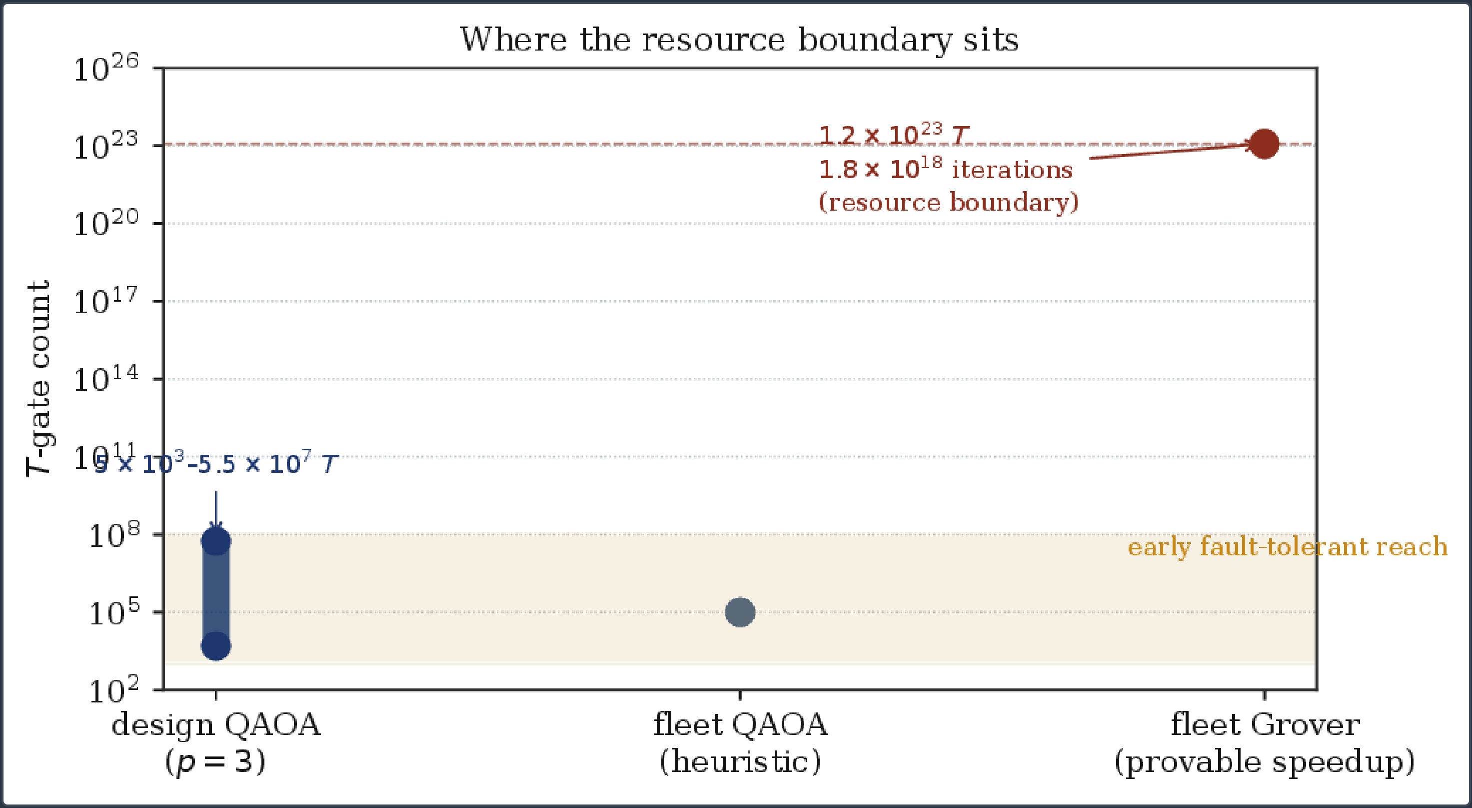

T-gate requirement across formulations. Design-scale QAOA sits within early fault-tolerant reach but is classically trivial to beat; the fleet-scale Grover search that would carry a provable speedup needs 1.2×10^{23} T-gates and 1.8×10^{18} iterations — a resource boundary. Values from the verified resource estimate (KX-L3-A6).

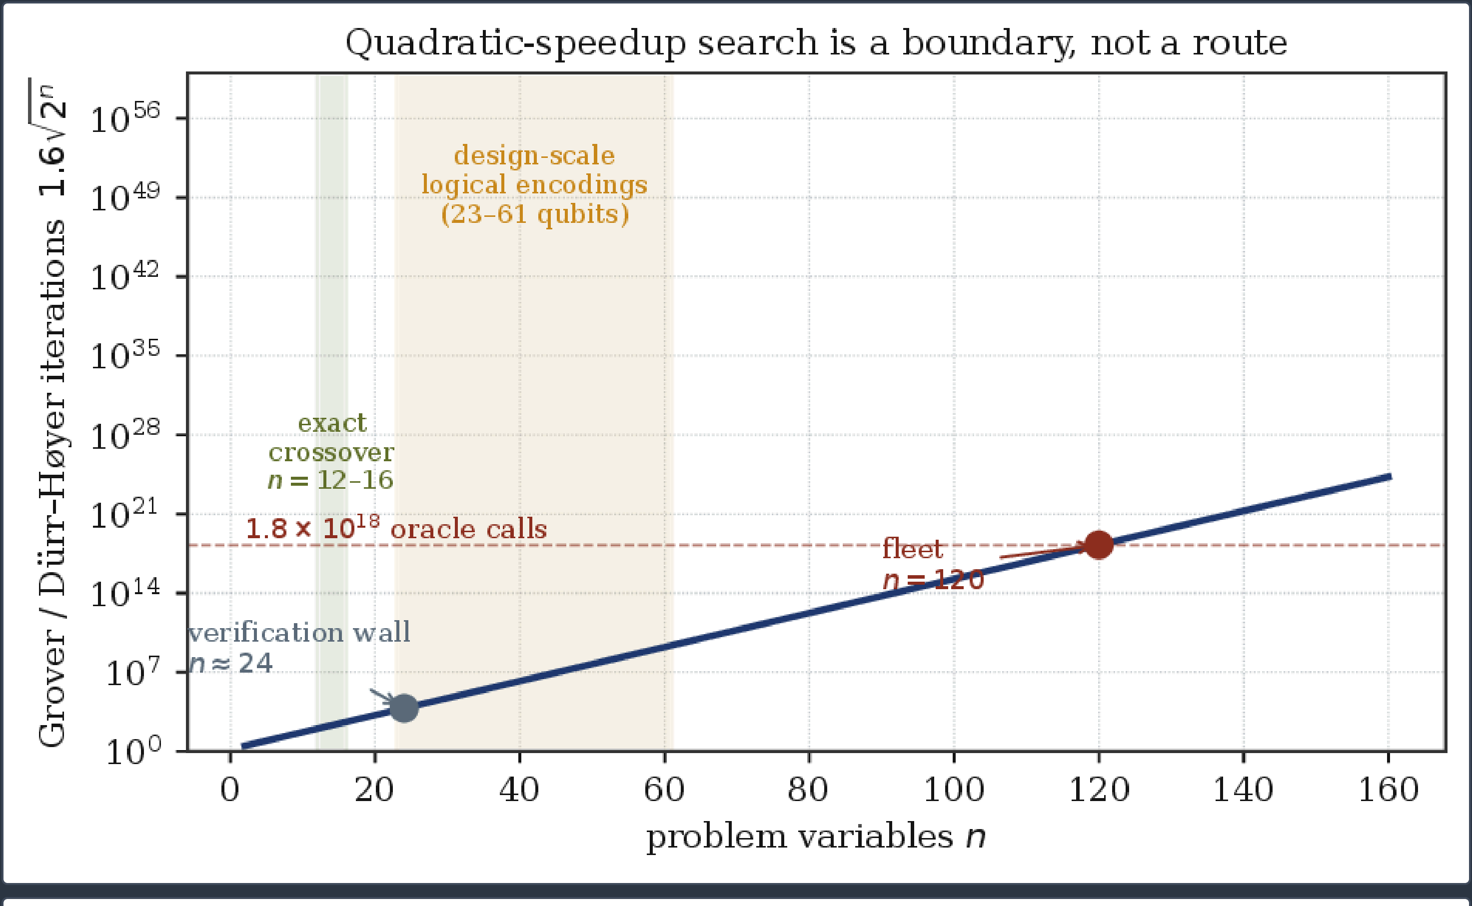

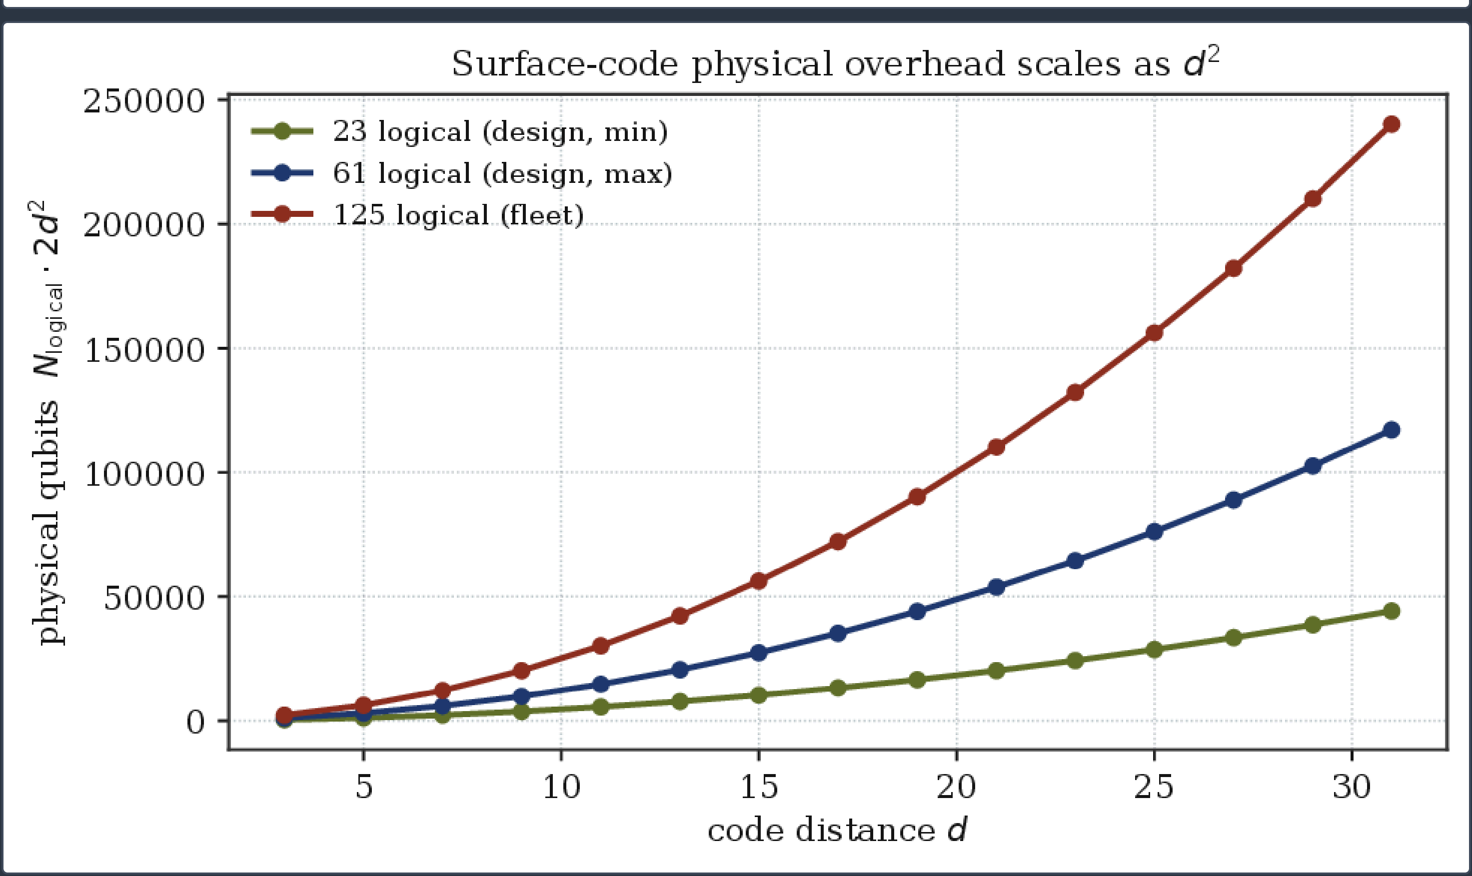

Grover / D\"urr–Hyer iteration count $1.6\sqrt{2^n}$ from equation (15) versus problem size n (log scale); design encodings (23–61) and the exact crossover band ($n = 12$–16) shaded, fleet point $n = 120$ at 1.8×10^{18} oracle calls (KX-L3-A6).

CLASSICAL DOMINANCE AND THE EXACT-SOLVE REGIME

The advantage question is settled by the classical baseline. On all three breeder QUBOs, classical simulated annealing matches or beats QAOA out to $n = 120$ variables, the largest instances we pose (BR-SX-06). Below the exact-to-heuristic crossover at $n = 12$–16 (Figure 4), exhaustive enumeration returns the certified global optimum in **0.03–292** s of wall time at design scale — so there is nothing for any heuristic, quantum or classical, to improve upon. Table 4 collects the three defining measurements.

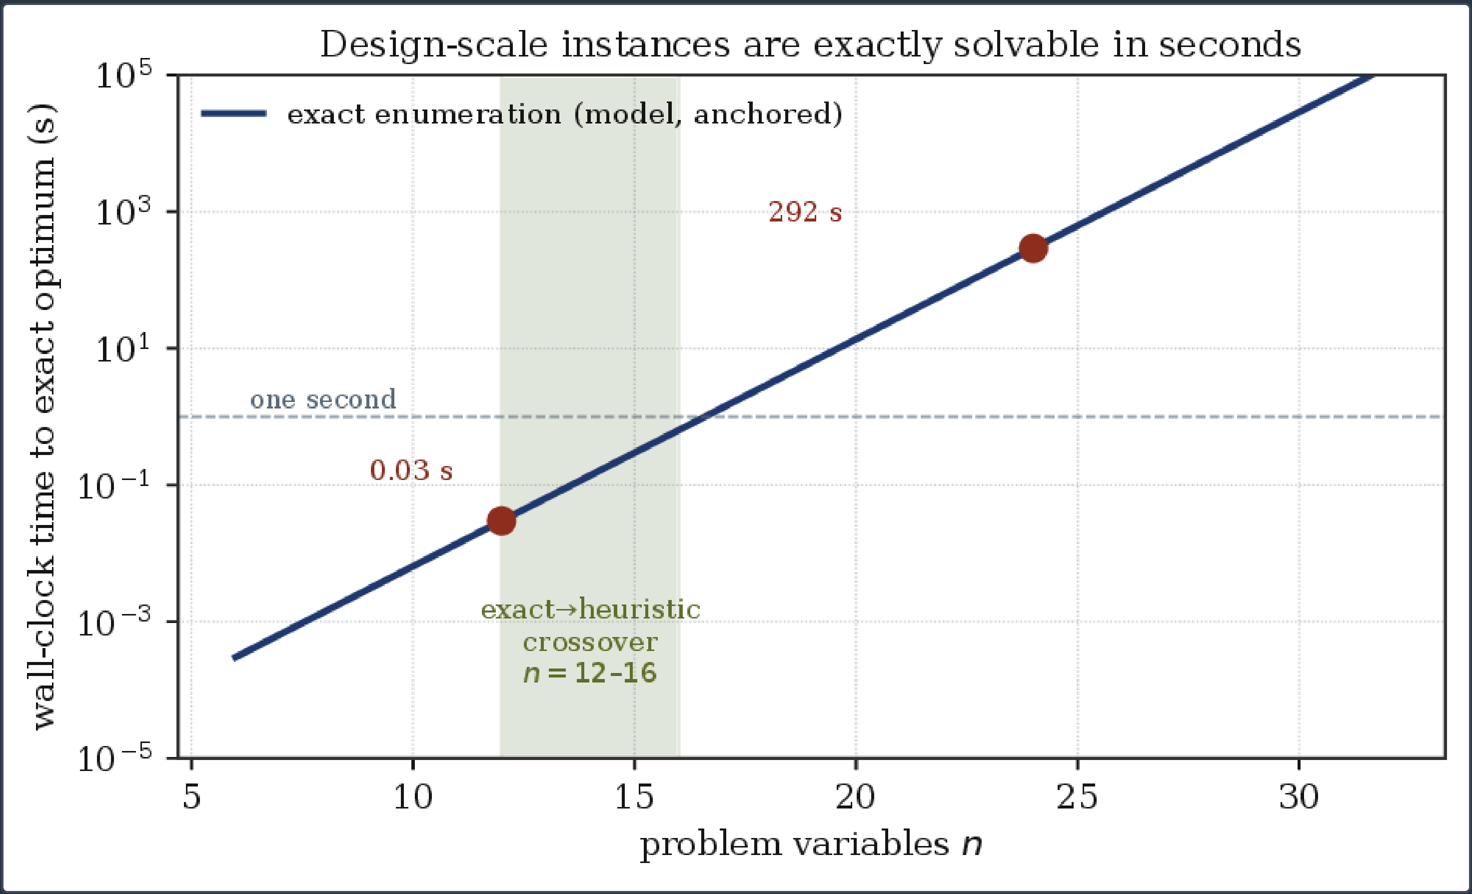

Exact-enumeration wall time (log scale): an 2^n cost model anchored to the frozen design-scale endpoints (**0.03** s and **292** s; BR-SX-06/KX-L3-A6), with the exact-to-heuristic crossover band $n = 12$–16 shaded. Design-scale instances are solved to certified optimality in seconds.

FINDING	VALUE	MEANING
classical SA vs QAOA	SA matches/beats to $n = 120$	no measured quantum advantage today
exact→heuristic crossover	$n = 12$–16 (0.03–292 s)	below it, exact optimum is available
QAOA approximation ratio	plateaus at 0.90–0.95	bounded quality at $p = 3$
verification wall	$n \approx 24$	above it, optimality unverifiable
fleet Grover	1.8×10^{18} iters, $1.2 \times 10^{23}\,T$	resource boundary, not a route

The pre-registered quantum-advantage verdict for the breeder optimization problems (BR-SX-06, KX-L3-A6). Each entry is a frozen measurement, not a projection.

THE HONEST MAP

Figure 5 draws the resulting regime map along problem size. Design-scale problems ($n \lesssim 16$) are exactly and verifiably solvable in seconds; a narrow band ($16 \lesssim n \lesssim 24$) is heuristic but still verifiable against ground truth; and past the verification wall ($n \gtrsim 24$) the problems are simultaneously classically unsimulable at full fidelity and, in their provable-speedup form, quantum-prohibitive. The QAOA approximation ratio plateaus at 0.90–0.95 throughout the accessible band. There is thus no window where quantum optimization is *needed* for breeder control: below the wall classical methods are exact and instantaneous, and above it the only provable-speedup formulation is a resource boundary.

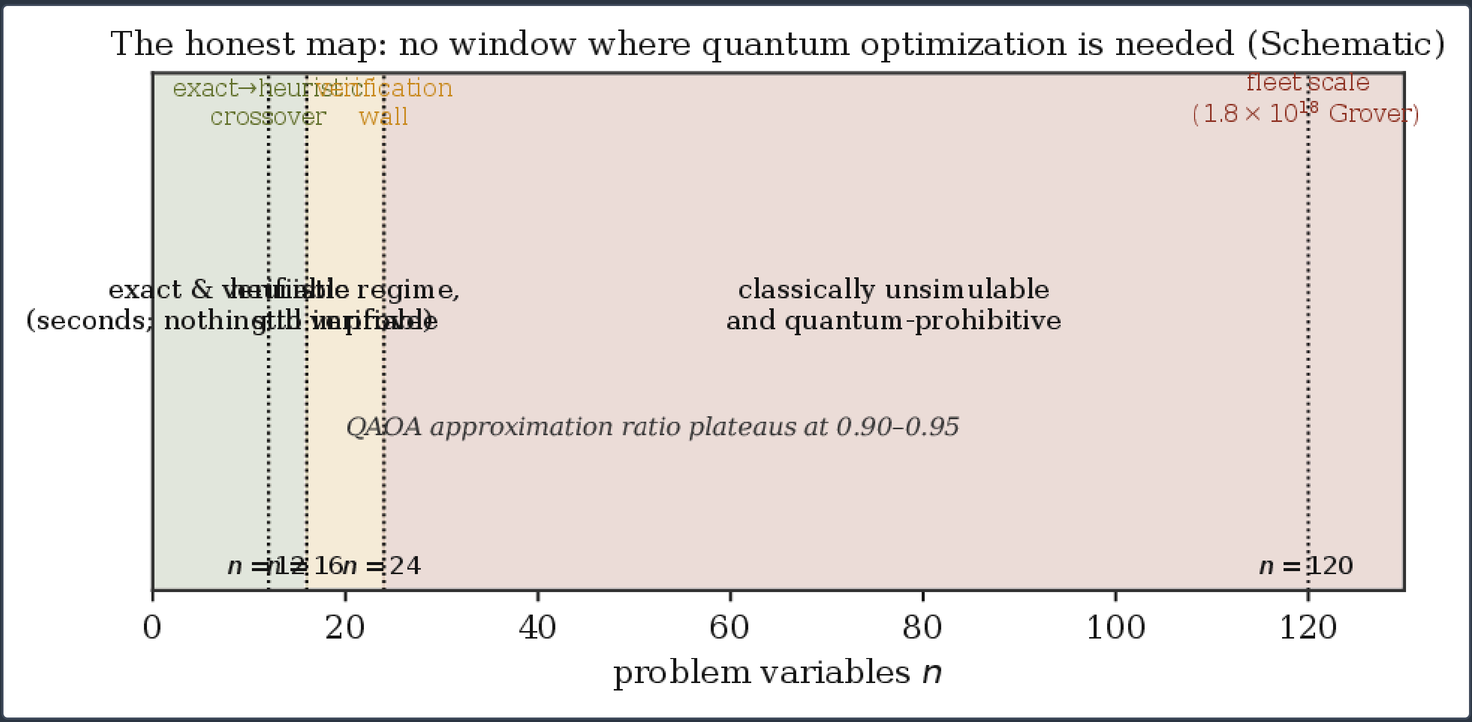

(Schematic.) The honest map along problem size n: an exact-and-verifiable regime ($n \lesssim 16$), a heuristic-but-verifiable band (16–24), and a classically-unsimulable, quantum-prohibitive regime ($n \gtrsim 24$). Boundaries ($n = 12$–16, $n \approx 24$, $n = 120$) are the frozen values of Table 4; the QAOA plateau annotation is the measured envelope. Structure only; no data are plotted.

QAOA STRUCTURE AND THE QUANTUM-RL GATE

Figure 6 shows the depth-three QAOA ansatz of equation (12) whose footprints populate Table 3. Figure 6 shows the quantum-RL element as a controlled gate: a variational quantum policy (19), trained by policy gradient (20) with parameter-shift derivatives (21), is admitted to command authority only if it clears the classical baseline at equal wall-clock (23) with the CBF safe-set filter (22) inviolate. This is a designed candidate with a falsifiable acceptance test, not a claim.

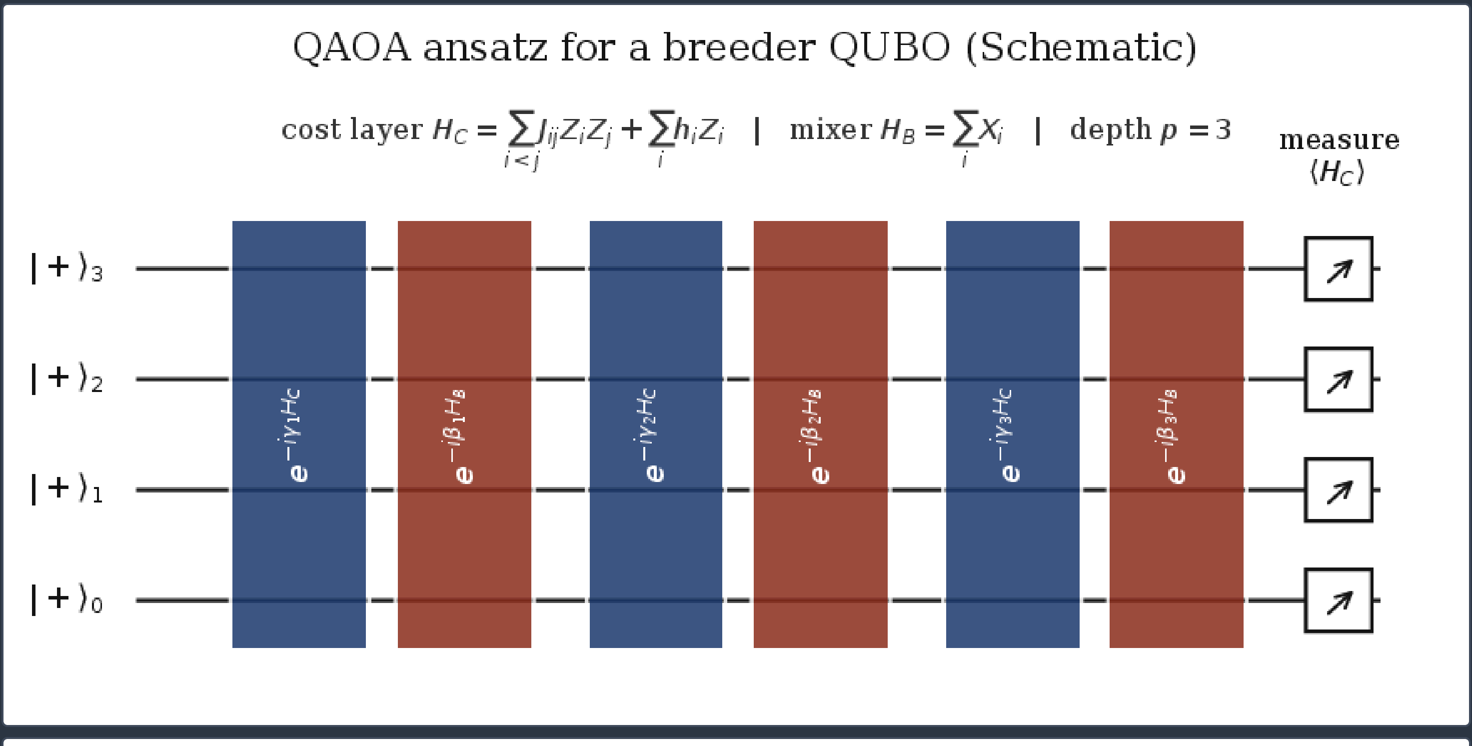

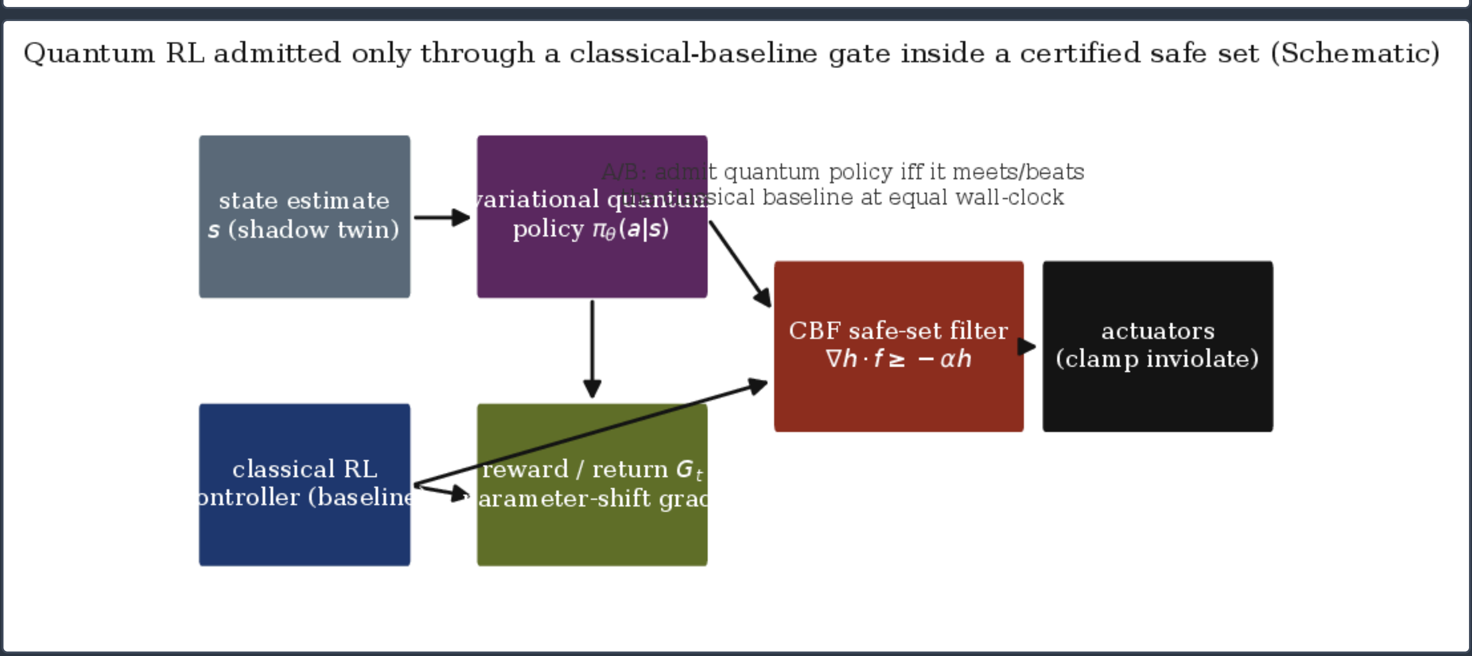

(Schematic.) The depth-three QAOA ansatz of equation (12): alternating cost layers $e^{-i\gamma_\ell H_C}$ built from the breeder QUBO (11) and transverse-field mixer layers $e^{-i\beta_\ell H_B}$, with a final energy measurement. Structure only.

THE DATED ROADMAP

Placing the same accounting on a calendar (Figure 7, KX-L3-A7) closes the picture. Design-scale QUBOs already encode on NISQ hardware but are classically dominated; the earliest fault-tolerant window relevant to fusion — linear-kinetics-scale problems at **11–19** logical qubits and 2×10^5–1.5×10^6 T-gates — opens around **2028–29** (IBM Starling / Quantinuum Apollo class), and it is against this milestone that the quantum-RL benchmark gate (23) is scheduled; larger, nonlinear-scale fault-tolerant computation stays on the far side of **2035**; and the fleet-scale Grover QUBO, at 1.8×10^{18} iterations, is never practical. Nothing on the breeder control critical path waits on quantum computing.

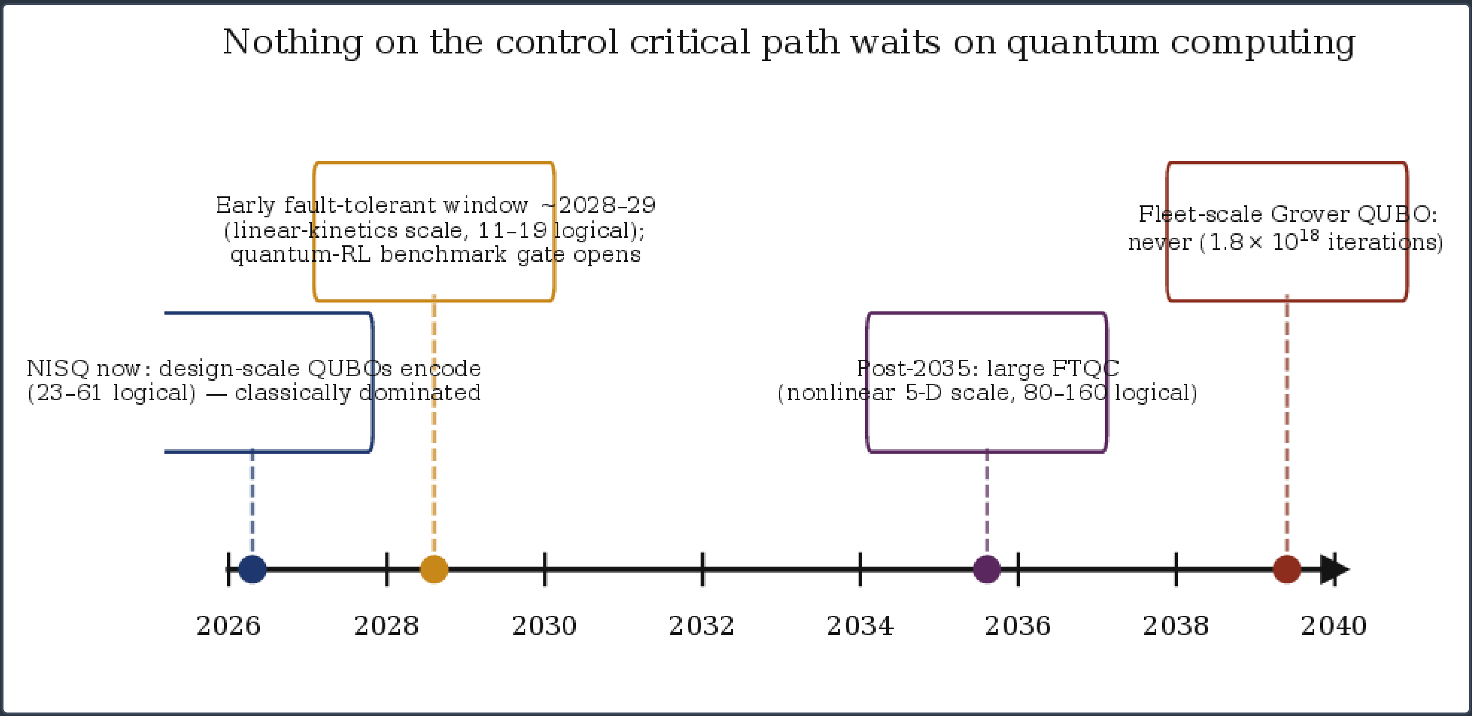

Dated fault-tolerant roadmap for quantum computing in the control stack (KX-L3-A7): design-scale QUBOs available now but classically dominated; an early fault-tolerant window ~**2028–29** at which the quantum-RL benchmark gate opens; large fault-tolerant computation post-**2035**; and the fleet-scale Grover QUBO never practical. Milestone years and logical/T counts are frozen register values.

§ 6

Discussion

QUANTITATIVE COMPARISON TO THE PUBLISHED LITERATURE

The finding here is consistent with, and quantitatively sharper than, the trajectory of the near-term quantum-optimization literature. QAOA was introduced as a shallow variational route to combinatorial optima [11] and remains an active research target [12], but its practical advantage over strong classical heuristics on generic QUBOs has not been demonstrated, and the field has grown cautious about advantage claims on NISQ hardware [16]. Two obstructions in the literature match our measurements precisely. First, variational objectives suffer barren plateaus and cost concentration, with gradients that vanish exponentially in qubit number for expressive ans\"atze [17,18]; our approximation-ratio plateau at 0.90–0.95 and the loss of ground-truth verification past $n \approx 24$ are the operational face of the same difficulty. Second, fault-tolerant accounting is unforgiving: even the most-optimized resource estimates for celebrated speedups reach tens of millions of physical qubits and enormous non-Clifford budgets [15,19], and our fleet-scale Grover figure — 1.2×10^{23} T-gates and 1.8×10^{18} sequential iterations from equations (15)–(16) — sits well beyond any practical horizon. Against the classical baseline the comparison is decisive: simulated annealing [21], a 1983-vintage heuristic, matches or beats QAOA on every breeder QUBO to $n = 120$, and below $n = 16$ the exact optimum is available in seconds. Where the literature offers a provable but only quadratic speedup [13,14], the quadratic factor is overwhelmed by the exponential base 2^n at fleet scale.

WHAT THE ENCODINGS ARE FOR

The negative on advantage does not make the encodings idle. They are the interface through which any future advantage would arrive, and freezing them — with footprints known to the logical qubit and the T-gate — means the program can re-run the advantage benchmark against improved hardware without re-deriving the problem. This is the same logic that motivates the companion quantum-algorithms study for the linear-algebra primitives of the control loop [32] and the sister variational-chemistry and kinetics work [31], and it is the discipline the umbrella AI-and-quantum paper applies across the Kronos frontier [30]. The control-relevant QUBO of §3.4 is, in the meantime, solved classically at control latency by the real-time optimization backbone [38] and, for the equilibrium currents, by the PF-coil optimizer [37], both of which act on the same shape-response Jacobian (6).

WHERE QUANTUM REINFORCEMENT LEARNING FITS

The QRL element is complementary to, not a competitor of, classical learned control. A deep RL policy already commands coil voltages on a real tokamak from magnetics, including negative-triangularity shapes [8], and the Kronos classical control suite is the operating baseline [34]. A variational quantum policy (19) is a candidate that must clear that baseline on the same task, at equal wall-clock, with the CBF safe-set filter (22) inviolate — the same forward-invariant clamp and safe-RL discipline that govern the rest of the stack [35,36,26]. Framing QRL as a gated candidate rather than a claim keeps the quantum path available and small without asserting an advantage the current evidence does not support; the parameter-shift training (21) and softmax read-out (19) are standard and hardware-ready [24,25,22,23], so when early fault-tolerant hardware arrives the benchmark can be run without further design.

§ 7

Limitations

The quantum-advantage benchmark is a pre-registered negative and stands as such. The single genuinely open item is that variational quantum RL for shape and vertical-position control has an encoding (19)–(22) and an acceptance test (23) but not yet a head-to-head result against the classical controller; that confirming comparison is gated on early fault-tolerant hardware availability (∼**2028–29**, Figure 7) and on the trainability of the policy circuit surviving the barren-plateau regime at the relevant qubit count [17,18]. Until that comparison is run, the classical controller remains the sole authority on the control critical path, and the safety filter (22) is inviolate regardless of which policy proposes the action.

§ 8

Conclusion

Real-time control of the Kronos breeder reduces, at its combinatorial core, to three QUBO problems — pulse cadence, blanket layout, and coil-current/actuator scheduling — and we derive and encode all three from first principles, culminating in the control-relevant QUBO built on the free-boundary equilibrium shape-response Jacobian (6)–(9). Their fault-tolerant footprints are compact and verified: 23–61 logical qubits and 5×10^3–5.5×10^7 T-gates at design scale, every row reproduced from the D\"urr–Hyer and Fowler formulas across 50 resource rows with zero discrepancy. On the question of quantum advantage we report an explicit negative: classical simulated annealing matches or beats QAOA to $n = 120$; the exact-to-heuristic crossover is $n = 12$–16 (0.03–292 s to certified optimality); the QAOA approximation ratio plateaus at 0.90–0.95 and optimality is unverifiable past $n \approx 24$; and fleet-scale Grover (1.8×10^{18} iterations, 1.2×10^{23} T-gates) is a resource boundary. Variational quantum reinforcement learning for shape and position control is placed on a dated, benchmark-gated schedule against the classical baseline inside a control-barrier-certified safe set. The deliverable is the honest, resource-grounded map of where quantum computing does — and does not — pay for real-time fusion control: available and small, not needed and not overclaimed.

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 acknowledges the NVIDIA GPU HPC campaign for the compute on which the quantum-circuit simulations, classical baselines, and resource recomputation were carried out.

Reproducibility. The QUBO resource footprints and the quantum-advantage benchmark are frozen entries (BR-SX-06, BR-QX-02, BR-QX-04, KX-L3-A6, KX-L3-A7) in the Kronos Physics De-Risking Register; the 50-row resource recomputation reproduced every value from the stated D\"urr–Hyer and Fowler formulas with zero failures. This paper is archived at its DOI doi:10.5281/zenodo.22645903 and at its official page https://www.kronosfusionenergy.com/publications/quantum-optimization-and-reinforcement-learning-for-rea.html; the de-risking register is archived at doi:10.5281/zenodo.22133057. The QUBO encoders, the QAOA-versus-classical benchmark, the figure-generation script, and the resource-recomputation harness are deposited with the paper.

Data availability tocsectionData availability The QUBO encodings, the resource-recomputation table, and the QAOA-versus-classical benchmark are archived with the deposit at doi:10.5281/zenodo.22645903, which references the Kronos 2026 series including the umbrella AI-and-quantum study [30] and the sister quantum-computing paper [31]. Any economic analysis is reported separately and is not part of the public deposit.

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

tocsectionReferences

APPARATUS

Portfolio, People & Record

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

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

INTELLECTUAL PROPERTY

The patent portfolio

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

GRANTED & PENDING UTILITY

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

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

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

TRADEMARK & ORIGIN

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

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

THE TEAM

Founding partners, board, advisors & auditors

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

FOUNDING PARTNERS — SCIENCE

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

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

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

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

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

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

SCIENTIFIC ADVISORS

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

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

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

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

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

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

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

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

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

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

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

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

SCIENTIFIC AUDITORS

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

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

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

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

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

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

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

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

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

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

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

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

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

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

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

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

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

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

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

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

Andrea Romero
Marketing Lead
2022–Present

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

SOURCES

References

1 J 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 M E Austin *et al*, Achievement of reactor-relevant performance in negative triangularity shape in the DIII-D tokamak, *Phys. Rev. Lett.* **122** 115001 (2019). doi:10.1103/PhysRevLett.122.115001

4 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

5 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

6 A Pironti and M Walker, Fusion, tokamaks, and plasma control: an introduction and tutorial, *IEEE Control Syst. Mag.* **25**(5) 30 (2005). doi:10.1109/MCS.2005.1512794

7 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

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

9 A Lucas, Ising formulations of many NP problems, *Front. Phys.* **2** 5 (2014). doi:10.3389/fphy.2014.00005

10 T Kadowaki and H Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58** 5355 (1998). doi:10.1103/PhysRevE.58.5355

11 E Farhi, J Goldstone and S Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028 (2014). doi:10.48550/arXiv.1411.4028

12 K Blekos, D Brand, A Ceschini *et al*, A review on Quantum Approximate Optimization Algorithm and its variants, *Phys. Rep.* **1068** 1 (2024). doi:10.1016/j.physrep.2024.03.002

13 L K Grover, A fast quantum mechanical algorithm for database search, *Proc. 28th Annual ACM Symp. Theory of Computing (STOC)* 212 (1996). doi:10.1145/237814.237866

14 C D\"urr and P Hyer, A quantum algorithm for finding the minimum, arXiv:quant-ph/9607014 (1996). doi:10.48550/arXiv.quant-ph/9607014

15 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

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

17 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

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

19 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

20 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

21 S Kirkpatrick, C D Gelatt and M P Vecchi, Optimization by simulated annealing, *Science* **220** 671 (1983). doi:10.1126/science.220.4598.671

22 K Mitarai, M Negoro, M Kitagawa and K Fujii, Quantum circuit learning, *Phys. Rev. A* **98** 032309 (2018). doi:10.1103/PhysRevA.98.032309

23 M Schuld, V Bergholm, C Gogolin, J Izaac and N Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99** 032331 (2019). doi:10.1103/PhysRevA.99.032331

24 S Jerbi, C Gyurik, S C Marshall, H J Briegel and V Dunjko, Parametrized quantum policies for reinforcement learning, *Advances in Neural Information Processing Systems (NeurIPS)* **34** (2021). arXiv:2103.05577. doi:10.48550/arXiv.2103.05577

25 S Y-C Chen, C-H H Yang, J Qi, P-Y Chen, X Ma and H-S Goan, Variational quantum circuits for deep reinforcement learning, *IEEE Access* **8** 141007 (2020). doi:10.1109/ACCESS.2020.3010470

26 A D Ames, S Coogan, M Egerstedt, G Notomista, K Sreenath and P Tabuada, Control barrier functions: theory and applications, *Proc. 18th European Control Conf. (ECC)* 3420 (2019). doi:10.23919/ECC.2019.8796030

27 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

28 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

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

30 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. Official page: https://www.kronosfusionenergy.com/publications/ai-quantum-control.html

31 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. Official page: https://www.kronosfusionenergy.com/publications/quantum-computing-for-fusion.html

32 P I Ford, *Quantum algorithms for the real-time control loop: HHL equilibrium reconstruction, disruption state-estimation, and edge-turbulence simulation with resource-honest advantage accounting*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645925. Official page: https://www.kronosfusionenergy.com/publications/quantum-algorithms-for-the-real-time-control-loop-hhl-e.html

33 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. Official page: https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html

34 P I Ford, *A real-time control suite for a compact negative-triangularity spherical-tokamak breeder: profile, shape, vertical-stability, q-profile, and solenoid-free ramp controllers*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645809. Official page: https://www.kronosfusionenergy.com/publications/a-real-time-control-suite-for-a-compact-negative-triang.html

35 P I Ford, *A certified control-barrier safety filter for fusion control: forward-invariance guarantees, calibrated uncertainty, and an honest PINN negative for a compact spherical-tokamak breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645716. Official page: https://www.kronosfusionenergy.com/publications/control-barrier-safety-clamp.html

36 P I Ford, *Safe reinforcement learning from operator feedback for fusion control: imitation and policy exploration confined to a control-barrier-certified safe set*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645823. Official page: https://www.kronosfusionenergy.com/publications/safe-reinforcement-learning-from-operator-feedback-for-.html

37 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. Official page: https://www.kronosfusionenergy.com/publications/poloidal-field-coil-set-and-equilibrium-current-optimiz.html

38 P I Ford, *A real-time optimization backbone for fusion plasma control: quadratic-programming, reduced-basis and adjoint-sensitivity engines at control latency*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645899. Official page: https://www.kronosfusionenergy.com/publications/a-real-time-optimization-backbone-for-fusion-plasma-con.html

39 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. Official page: https://www.kronosfusionenergy.com/publications/verification-and-validation-47-code-provenance.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 Optimization and Reinforcement Learning for Real-Time Plasma Control: QUBO/QAOA Actuator Scheduling and Quantum RL for Shape Control, without Claimed Advantage

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

Zenodo · doi:10.5281/zenodo.22645903

Read on kronosfusionenergy.com

The Physics De-Risking Register