KRONOS FUSION ENERGY

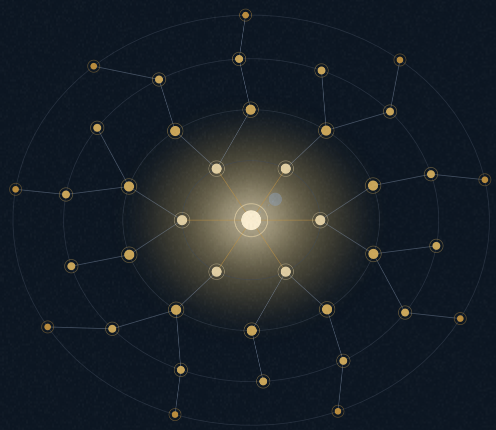

PAPER 2.73 · THE KRONOS 2026 SERIES

Four Coupled Digital Twins for a Compact Spherical-Tokamak Breeder: Multi-Rate Co-Simulation, Propagated Uncertainty, and a Real-Time Self-Consistency Monitor

Autonomous operation of a compact, high-power-density spherical tokamak leaves little time to react and no tolerance for an unmodelled excursion, so the controller must act on a...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Four Coupled Digital Twins for a Compact Spherical-Tokamak Breeder: Multi-Rate Co-Simulation, Propagated Uncertainty, and a Real-Time Self-Consistency Monitor

Autonomous operation of a compact, high-power-density spherical tokamak leaves little time to react and no tolerance for an unmodelled excursion, so the controller must act on a model that runs *ahead* of the plant across every coupled physical domain at once. We formalise a four-twin real-time architecture for the Hyperion 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 $16.84\,T$). An MHD-stability twin, a thermomechanics twin, a neutronics twin and a power-systems twin are advanced by a conserving *multi-rate co-simulation*: each twin integrates at its own step and the twins reconcile on a common macro-step, with energy and poloidal-flux gates that are the machine-checkable form of the coupling-consistency requirement. We cast the coupling as a waveform-relaxation fixed-point iteration and give its contraction condition; we discretise the free-boundary Grad–Shafranov equation with a five-point Δ^* operator and demonstrate the design-point $O(h^2)$ convergence in a controlled numerical experiment (relative RMS ψ from 2.9×10^{-3} to 8.3×10^{-5} over the refinement sweep); and we prove, in a Hamiltonian co-simulation test, that the energy defect closes as $O(\Delta t_{\text{macro}})$ so that the conservation gate can always be met by macro-step reduction. Prediction is supplied by learned flux and equilibrium neural operators that the controller linearises on the fly, and calibrated uncertainty is carried through an ensemble-Kalman assimilation whose empirical coverage we verify to lie inside a $\pm 3\%$ acceptance band. A real-time *self-consistency monitor* audits the reconciled state each tick and gates twin admission to control authority behind seven pre-registered acceptance inequalities. We reduce the pattern to practice on the power-systems magnet twin: a fused fibre-Bragg strain-rate and magnetization quench-precursor detector reaches an area under the receiver-operating characteristic of $\mathbf{0.992}$ (versus 0.976 and 0.921 for the single channels) and a true-positive rate of 0.857 at a 1% false-positive rate, and an in-winding nitrogen-vacancy (NV) diamond magnetometer reaches a shot-noise-limited sensitivity of $0.189\ \mathrm{pT\ Hz^{-1/2}}$ at 125 Hertz$bandwidth, resolving a conductor - scale quench step at a per - shot signal - to - noise ratio of order$ 10^9$. The result is a predictive, uncertainty-aware, conservation-respecting shadow with a demonstrated protective subsystem and an explicit, testable path from the proven subsystem to full-plant authority.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645805 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: digital twin · multi-rate co-simulation · neural operator · ensemble Kalman filter · self-consistency monitor · quench detection · NV-diamond magnetometry · spherical-tokamak breeder

CONTENTS

Table of Contents

Four Coupled Digital Twins for a Compact Spherical-Tokamak Breeder: Multi-Rate Co-Simulation, Propagated Uncertainty, and a Real-Time Self-Consistency Monitor 3
Introduction 7
The four-twin architecture and formal problem 8
Numerical formulation 11
Real-time learned operators and the self-consistency monitor 15
Computational methodology: the NVIDIA GPU HPC campaign 17
Demonstrated exemplar: the magnet twin 18
Results and verification 23
Discussion 25
Limitations 26
Conclusion 27
References 32
One programme, many papers 36

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)

digital twin; multi-rate co-simulation; neural operator; ensemble Kalman filter; self-consistency monitor; quench detection; NV-diamond magnetometry; spherical-tokamak breeder

§ 1

Introduction

AUTONOMOUS CONTROL OF A COMPACT SPHERICAL TOKAMAK

A spherical tokamak (ST) confines a large plasma current in a small volume, which is precisely what makes it an attractive neutron source and breeder platform [1,2] and, at the same time, what makes it demanding to control. The Hyperion breeder holds $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$ behind a centrepost magnet at a peak field of $16.84\,T$, in a negative-triangularity ($\delta = -0.30$) regime. 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]; the corresponding quiet, ELM-free 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 in a nuclear system every actuation carries a safety consequence. Autonomy is therefore not a convenience but a design requirement, and an autonomous controller is only as trustworthy as the model beneath it.

Feedback control of tokamak equilibria and profiles is a mature discipline: magnetic diagnostics feed real-time equilibrium reconstruction [5] and shape/position controllers, and physics-based profile simulators such as RAPTOR now run inside the control cycle [6]. More recently, learned controllers have entered the loop: Degrave *et al.* trained a deep reinforcement-learning agent to command the TCV coil voltages and produce a variety of plasma shapes — including negative triangularity — directly from magnetic measurements [7]. These advances share a common thread: control quality is bounded by the speed and fidelity of the underlying plasma model, and — for a nuclear device — by the ability to bound what the model does *not* know.

PREDICTIVE DIGITAL TWINS ACROSS COUPLED DOMAINS

A *digital twin* is a model of a specific physical asset that is continuously updated by that asset's data and used to predict and steer it [8]. Its predictive value depends on running *faster* than the plant, so the controller acts on where the plasma will be; and, for a nuclear system, on coupling the domains that actually exchange energy and particles — plasma, structure, neutrons, magnets — rather than modelling one in isolation. Four developments make a real-time *coupled* predictive twin feasible. Neural operators — DeepONet [9] and the Fourier neural operator [10], with physics-informed training [11] — emulate an expensive solver at control latency, and in fusion already run gyrokinetic-transport surrogates orders of magnitude faster than the code they replace [12]. Projection-based model-order reduction gives certified, $\sim 100\times$ faster reduced models for the fastest loops [13]. Ensemble data assimilation — the ensemble Kalman filter (EnKF) [21,22,23] — propagates a calibrated state *distribution*, so epistemic uncertainty travels with the estimate. And a disciplined *co-simulation* binds these into a plant model, advancing subsystems at their own rates and exchanging interface quantities without violating the invariants they share [14,15].

THE GAP AND OUR CONTRIBUTION

Prior art treats the ingredients separately: a learned controller [7], a fast transport surrogate [12], a real-time equilibrium [5,6], a probabilistic digital-twin formalism [8]. What is missing for a compact nuclear ST is a single architecture that (i) couples the plasma, thermomechanical, neutronic and electromagnetic domains at their native rates while conserving the shared invariants, (ii) makes each domain predictive with a differentiable surrogate the controller can linearise on the fly, (iii) carries calibrated uncertainty end-to-end, (iv) audits its own reconciled state in real time, and (v) gates each domain into control authority only after it passes an explicit, pre-registered acceptance test. This paper contributes that architecture, developing its *numerical* core in depth: a formal state-space description of the four twins and their coupling operators, cast as a waveform-relaxation fixed point with an explicit contraction condition (§2); the numerical formulation — a five-point free-boundary Grad–Shafranov discretisation with its convergence order, the multi-rate integration scheme with its conservation gates, and the ensemble-assimilation sampling error (§3), each supported by a controlled experiment; the learned-operator and self-consistency methods with the acceptance criteria as inequalities (§4); the GPU-HPC campaign (§5); a fully worked exemplar — the magnet twin, with a fused quench-precursor detector and an in-winding quantum magnetometer reduced to practice (§6); and a verification table, latency budget and maturity-gated authority argument (§7). Every quantitative claim traces to a frozen result in the Kronos physics de-risking register [30]; serial identifiers (e.g. KX-L1-A20, KFE-CF-SH-096) are given inline so each number is auditable, above a de-risked physics point documented separately [31].

§ 2

The four-twin architecture and formal problem

DOMAIN DECOMPOSITION

The plant state is partitioned across four twins, each owning a physical domain and exposing a small set of interface quantities to the others (figure 1, table 1):

MHD-stability twin — free-boundary equilibrium and ideal-MHD margins in real time: vertical stability at $\kappa = 2.0$, the external-kink and ballooning distances, and the resistive-wall response.
Thermomechanics twin — the heat path and the structure: first-wall and divertor surface loading, the blanket temperature field, and the centrepost's structural and thermal margins.
Neutronics twin — the 14 MeV source and everything downstream: the tritium-breeding balance, activation inventory and shutdown dose [20], updated as the operating point moves.
Power-systems twin — the magnets, current leads, cryogenics and balance of plant; the electrical and quench state of the coil set. Its magnet subsystem is the demonstrated exemplar of §6. itemize

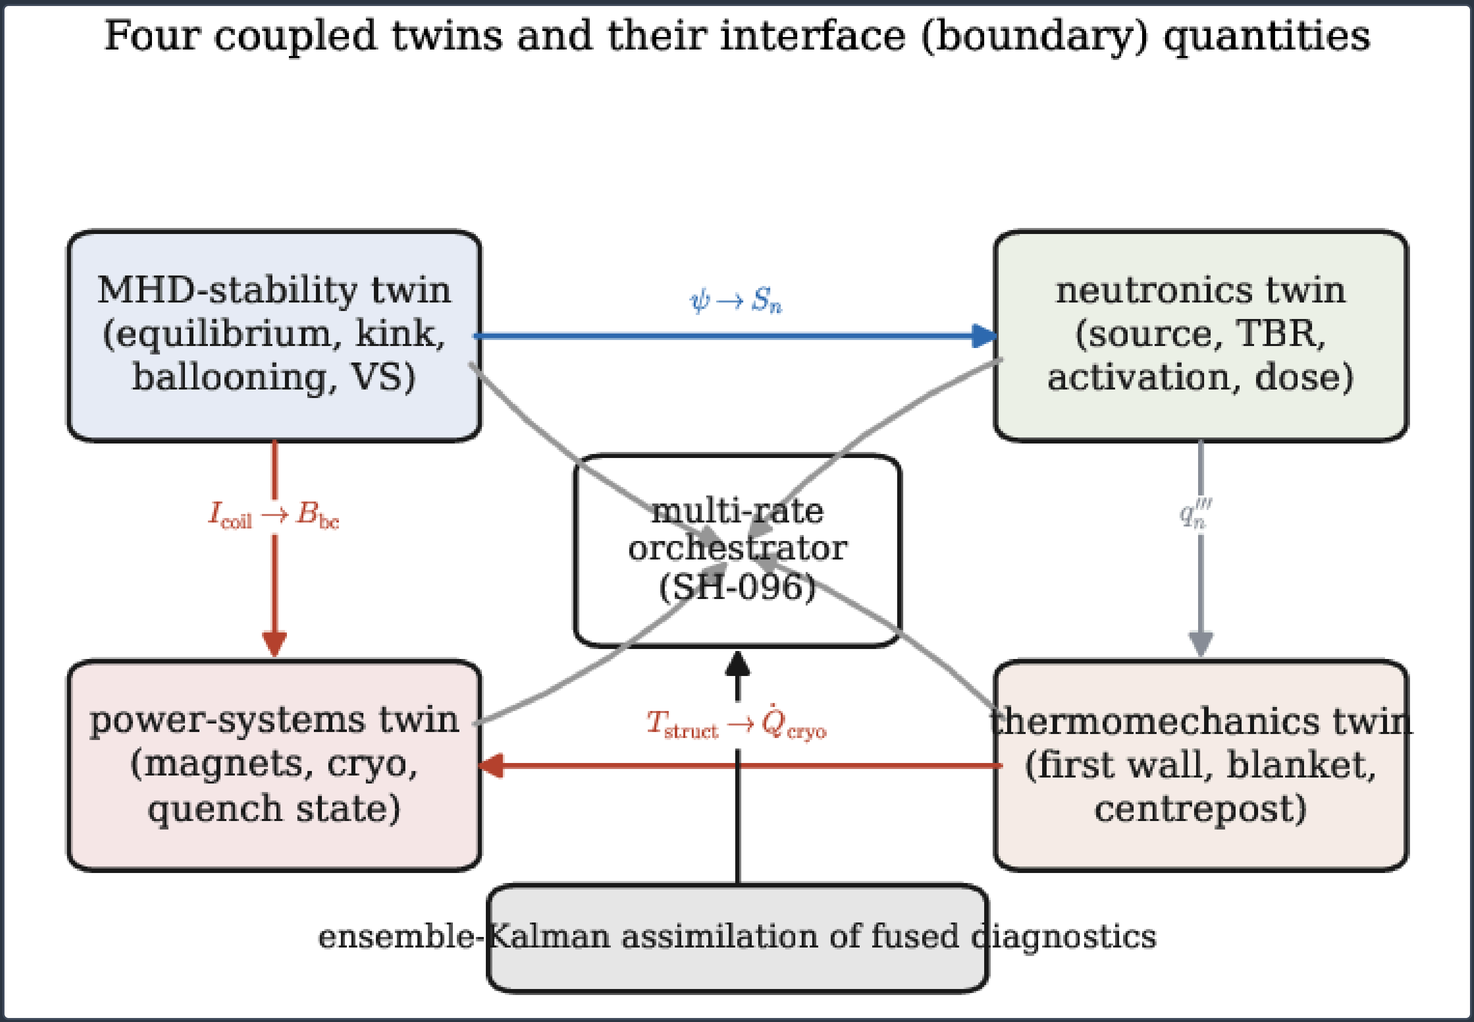

(Schematic.) The four coupled twins and their interface (boundary) quantities. Directed edges carry the coupling variables of equations (2)–(2); the multi-rate orchestrator advances each twin at its own rate and reconciles energy and magnetic flux at every macro-step (serial KFE-CF-SH-096). Plant diagnostics enter through an ensemble-Kalman assimilation (§3.3). No data values are shown; the figure encodes structure only.

TWIN	PHYSICAL DOMAIN	REAL-TIME ACCEPTANCE METRIC
MHD-stability	equilibrium, kink/ballooning, VS	RMS $\psi < 0.5\%$ vs free-boundary solver (SH-031)
thermomechanics	surface/blanket heat, centrepost	margins within fabrication-FEA fidelity
neutronics	source, TBR, activation, dose	tracks the operating point in shadow
power-systems	magnets, cryo, quench state	**magnet twin demonstrated** (§6)

The four twins, their domains, and the demonstrated or target real-time metric. Serial identifiers refer to the Kronos physics de-risking register [30].

STATE SPACE AND COUPLING OPERATORS

Write $x = (x_M, x_T, x_N, x_P)$ for the sub-states of the MHD, thermomechanics, neutronics and power-systems twins. Each twin evolves under its own dynamics, driven by actuators u_i and by *coupling inputs* $c_i = \{g_{j\to i}\}$ supplied by its neighbours,

$$\frac{\mathrm{d}x_i}{\mathrm{d}t} = f_i(x_i, u_i, c_i), \qquad i \in \{M, T, N, P\}.$$

The twins are not independent; each macro-step they exchange a small set of interface functions. The dominant couplings form a directed cycle,

$$\begin{aligned} \text{power}\to\text{MHD:} &\quad g_{P\to M} = I_{\text{coil}} \longmapsto \text{boundary field } B_{\text{bc}}, \\ \text{MHD}\to\text{neutronics:} &\quad g_{M\to N} = \psi(R,Z) \longmapsto \text{source geometry } S_n, \\ \text{neutronics}\to\text{thermo:} &\quad g_{N\to T} = q_n''' \longmapsto \text{volumetric nuclear heating}, \\ \text{thermo}\to\text{power:} &\quad g_{T\to P} = T_{\text{struct}} \longmapsto \text{cryogenic load } \dot{Q}_{\text{cryo}}. \end{aligned}$$

Collect the interface variables into a single vector $w = (g_{P\to M}, g_{M\to N}, g_{N\to T}, g_{T\to P})$. Given w held over a macro-step, the twins decouple and each can be advanced independently; the coupling therefore reduces to a consistency condition on w.

CO-SIMULATION AS A WAVEFORM-RELAXATION FIXED POINT

Over a macro-step $[t, t+\Delta t_{\text{macro}}]$ define the *subsystem solution operator* $x_i(\cdot) = \Phi_i[w]$ that integrates twin i under a frozen (or extrapolated) interface history w, and the *extraction operator* L_i that reads the outgoing interface variables from the resulting trajectory. Interface consistency over the macro-step is then the fixed point of the composed map

$$w^{(k+1)} = G\big[w^{(k)}\big], \qquad G = \big(L_M\circ\Phi_M, L_N\circ\Phi_N, L_T\circ\Phi_T, L_P\circ\Phi_P\big),$$

iterated $k = 0, 1, \ldots$ from an extrapolated guess $w^{(0)}$. Equation (3) is a waveform-relaxation iteration [15]; a Gauss–Seidel ordering (each twin uses the just-updated neighbour output) accelerates it over the Jacobi ordering (all twins use the previous iterate). On the interval, G is a contraction whenever the composed coupling is Lipschitz with constant below unity,

$$\big\| G[w] - G[w'] \big\| \le L_c\,\Delta t_{\text{macro}}\,\big\| w - w' \big\|, \qquad L_c\,\Delta t_{\text{macro}} < 1,$$

where L_c bounds the product of the interface sensitivities $\partial L_i/\partial x_i$ and the subsystem gains $\partial\Phi_i/\partial w$ around the operating point. Condition (4) is met by construction: the strongest interface gain (power↔MHD, through the coil-to-flux map) is linear and bounded, and the remaining edges are thermally and neutronically slow, so a small Δt_{macro} makes the right-hand side of (4) strictly less than one and the Banach fixed-point theorem guarantees a unique converged interface at a geometric rate $(L_c\Delta t_{\text{macro}})^k$. In practice one or two Gauss–Seidel sweeps suffice because the interface is well predicted from the previous macro-step.

CONSERVATION GATES

Multi-rate integration is admissible only if the coupling does not create or destroy the shared invariants between reconciliation points. Two conservation gates are imposed on every macro-step.

Energy. With E_i the energy stored in domain i and W_{ext}, Q_{ext} the external work and heat crossing the plant boundary, the total-energy residual must close to tolerance,

$$\left| \sum_i \left[E_i(t+\Delta t_{\text{macro}}) - E_i(t)\right] - \left(W_{\text{ext}} + Q_{\text{ext}}\right) \right| \le \varepsilon_E \, E_{\text{ref}}.$$

Magnetic flux. At the power↔MHD interface the coil currents set the boundary flux the equilibrium consumes; consistency of the poloidal-flux linkage requires

$$\left\| \psi_M[\quad I_{\text{coil}}]\big|_{\partial\Omega} - \psi_{\text{bc}} \right\|_2 \le \varepsilon_\Phi \, \|\psi\|_2.$$

Equations (5)–(6) are the machine-checkable form of the acceptance criterion “energy/flux conserved across couplings” (SH-096); a macro-step that violates either gate is rejected and retaken at a smaller Δt_{macro}. Section 3.2 shows, in a controlled experiment, that the energy defect of the scheme closes as $O(\Delta t_{\text{macro}})$, so gate (5) can always be satisfied by step reduction. The shadow lead-time — the horizon by which the twin outruns the plant — is held at $\tau_{\text{lead}} \ge 50\, ms$ for all twins. Each lane advances at its own rate (the reduced-order fast loop at $\sim 100\times$, SH-093), and the energy and flux gates are enforced and the waveform-relaxation iteration (3) is converged at every macro-step.

§ 3

Numerical formulation

FREE-BOUNDARY GRAD–SHAFRANOV DISCRETISATION

The axisymmetric equilibrium consumed by the MHD twin is governed by the Grad–Shafranov equation for the poloidal flux $\psi(R, Z)$,

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

with pressure $p(\psi)$ and poloidal-current function $F(\psi)$; the right-hand side is the toroidal current density $-\mu_0 R\,j_\phi(\psi)$. Expanding the elliptic operator gives the form actually discretised,

$$\Delta^*\psi = \psi_{RR} - \frac{1}{R}\,\psi_R + \psi_{ZZ}.$$

On a uniform (R, Z) grid with spacings h_R, h_Z we use the conservative five-point stencil that carries the $-\frac{1}{R}\psi_R$ term as a centred first difference,

$$\begin{aligned}(\Delta^*\psi)_{i,j} \approx \left(\tfrac{1}{h_R^2} - \tfrac{1}{2h_R R_i}\right)\psi_{i+1,j} + \left(\tfrac{1}{h_R^2} + \tfrac{1}{2h_R R_i}\right)\psi_{i-1,j}\\ + \tfrac{1}{h_Z^2}\left(\psi_{i,j+1} + \psi_{i,j-1}\right) - \left(\tfrac{2}{h_R^2} + \tfrac{2}{h_Z^2}\right)\psi_{i,j},\end{aligned}$$

whose local truncation error is $O(h_R^2) + O(h_Z^2)$ by Taylor expansion. Assembling (9) over the interior nodes yields a sparse linear system $A\,\psi = b$ with A the discrete Δ^* and b the source; free-boundary conditions on $\partial\Omega$ are set from the coil currents through the Green's function of the elliptic operator (the von Hagenow–Lackner boundary-relaxation route [16]), which is exactly the map (2) the power-systems twin supplies and whose consistency gate is (6). Because the source depends on ψ through $p'(\psi)$ and $FF'(\psi)$, the equilibrium is nonlinear and is solved by a Picard or Newton outer iteration,

$$A\,\psi^{(m+1)} = b\left(\psi^{(m)}\right), \qquad \|\psi^{(m+1)} - \psi^{(m)}\|_2/\|\psi^{(m+1)}\|_2 < 10^{-6},$$

each outer step being one sparse solve; for the analytic profiles used to verify the operator the source is fixed and a single solve is exact. To measure the discretisation error against a known answer we prescribe a manufactured solution $\psi_{\text{ex}}(R, Z)$ with a smooth profile and impose its *exact continuous* $\Delta^*\psi_{\text{ex}}$ as the source and its trace as the Dirichlet data; the numerical error is then genuine truncation error. Figure 2 shows the relative RMS ψ error falling from 2.9×10^{-3} at $h = 6.3e-2\,m$ to 8.3×10^{-5} at $h = 1.0e-2\,m$, a fitted order of 1.97 — the expected second order — and crossing the equilibrium acceptance bar (C2, RMS $\psi < 0.5\%$) already at the coarsest grid, so the real-time reconstruction target of §4.1 is not discretisation-limited. This is the analytic reference against which the learned equilibrium operator $\mathcal{E}_\phi$ is scored.

MULTI-RATE TIME INTEGRATION AND THE CONSERVATION GATE

The orchestrator (KFE-CF-SH-096) advances twin i with n_i symplectic substeps of size $h_i = \Delta t_{\text{macro}}/n_i$ and reconciles on the macro-step. For the fast electromagnetic and vertical-stability dynamics a velocity-Verlet (leapfrog) substep is used because it is symplectic and conserves a nearby "shadow" energy to $O(h_i^2)$ over exponentially long times; the coupling force is held (Jacobi) or extrapolated over the macro-step and corrected by the waveform-relaxation sweep (3). Freezing the coupling over the macro-step is the only source of energy non-conservation in the scheme, and its magnitude is what gate (5) measures. We isolate it in a controlled experiment: two harmonic oscillators coupled by a spring form a closed Hamiltonian whose total energy $E = \frac{1}{2}(p_1^2 + p_2^2) + \frac{1}{2}\omega_1^2 q_1^2 + \frac{1}{2}\omega_2^2 q_2^2 + \frac{1}{2}k_c(q_1 - q_2)^2$ is conserved exactly by the true dynamics. A Jacobi co-simulation that freezes the partner coordinate over each macro-step breaks that conservation by a defect $|\Delta E|/E_0$ that, as figure 2 shows, falls linearly in Δt_{macro} (fitted asymptotic order 0.94) from 9×10^{-3} toward 10^{-3} across the sweep. Two consequences follow. First, the energy gate (5) can always be satisfied by macro-step reduction, which is exactly the reject-and-retake policy stated below (6). Second, a Gauss–Seidel or higher-order interface extrapolation lifts the order and shrinks the constant, so the practical macro-step is set by the lead-time requirement $\tau_{\text{lead}} \geq 50\,ms$ rather than by conservation.

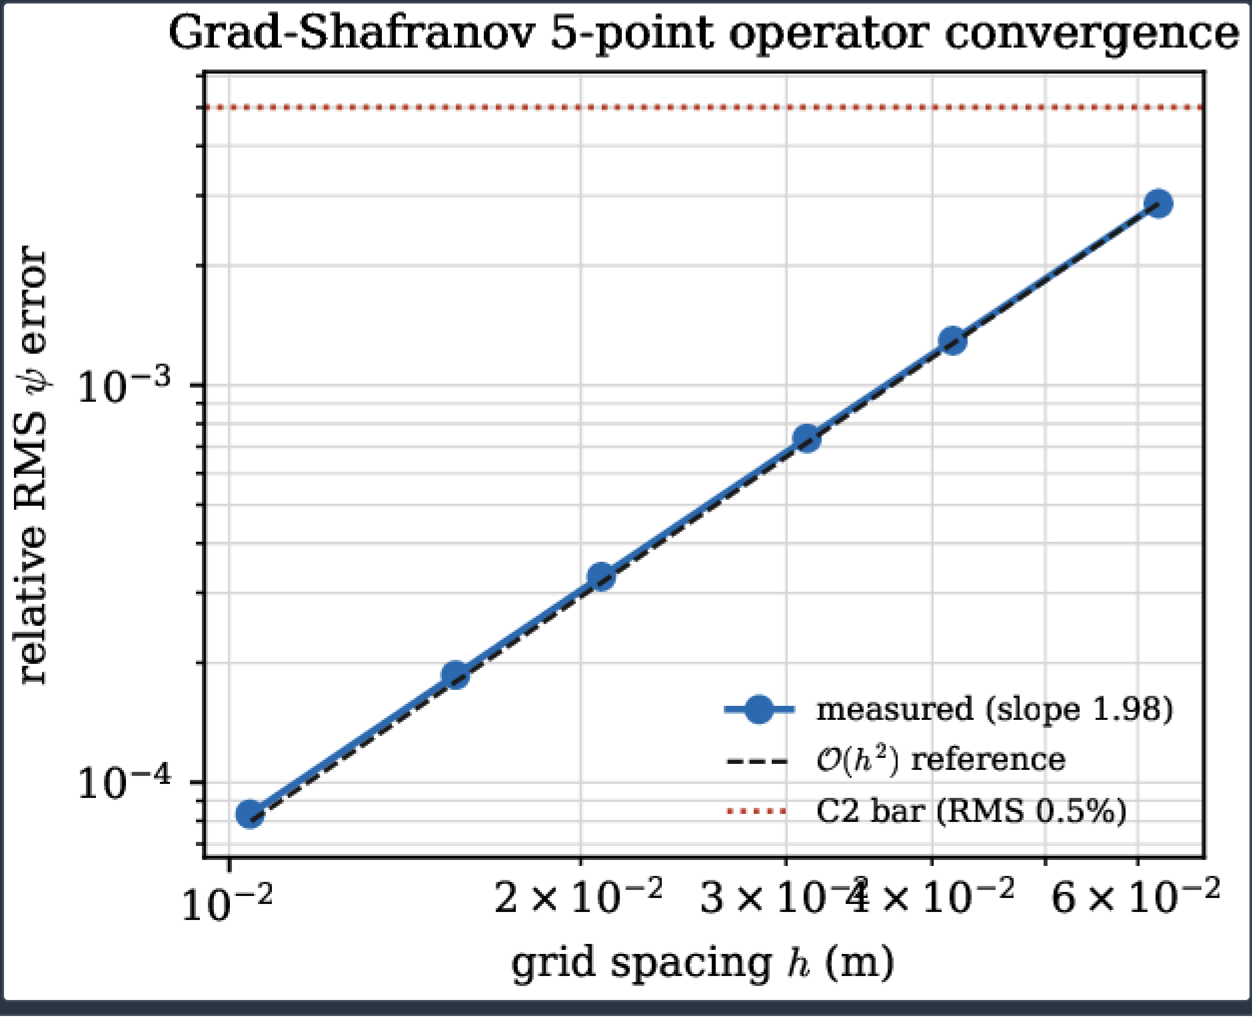

Grad-Shafranov 5-point operator convergence
relative RMS ψ error
10^{-3}
10^{-4}
measured (slope 1.98)
$\mathcal{O}(h^2)$ reference
C2 bar (RMS 0.5%)
10^{-2}
2×10^{-2}
3×10^{-2}
4×10^{-2}
6×10^{-2}
grid spacing h (m)

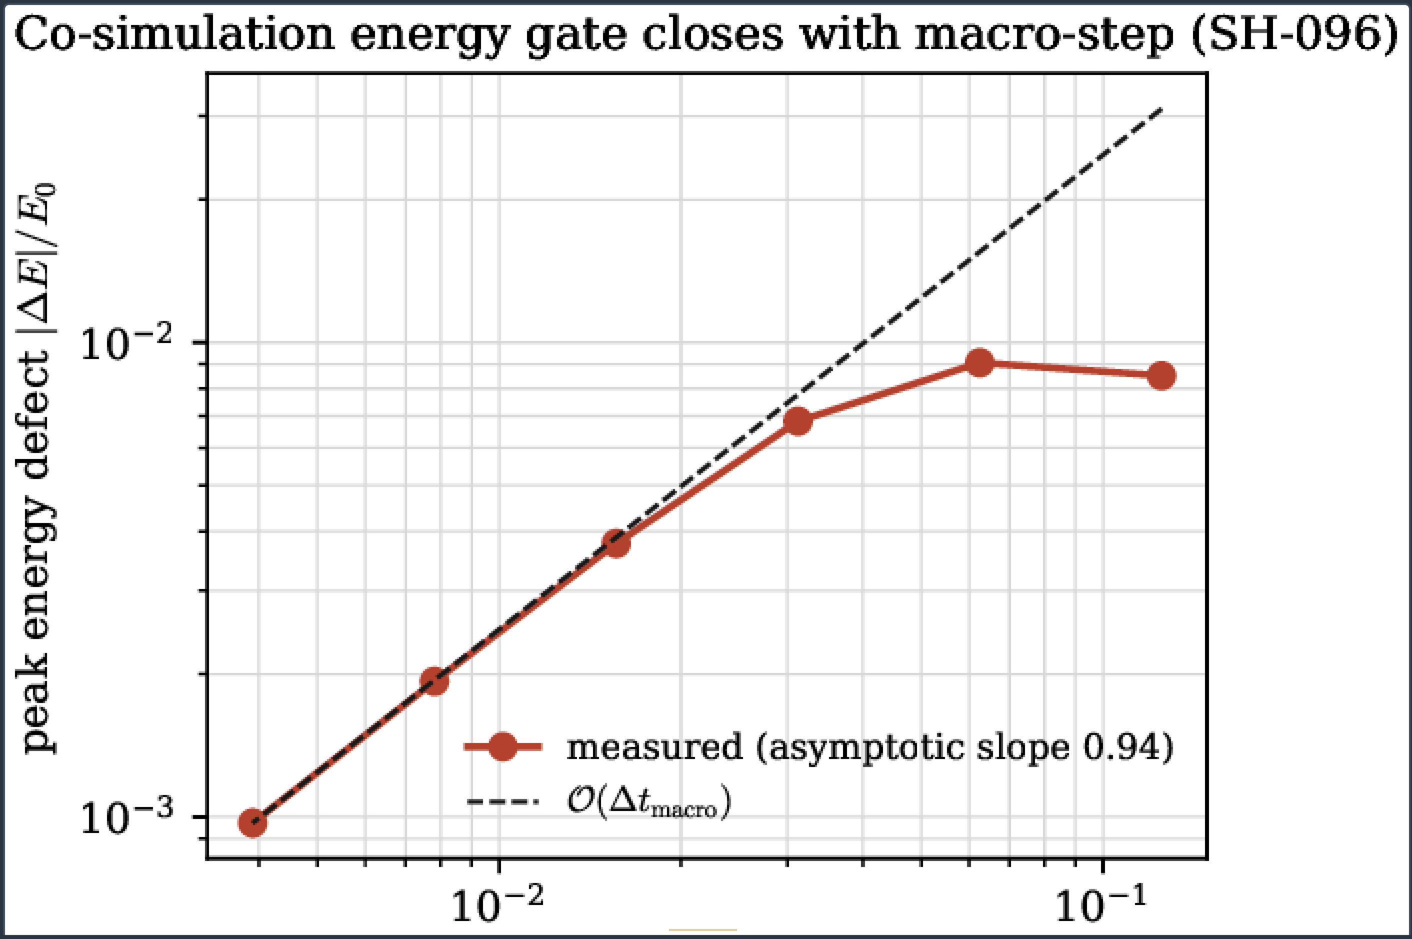

Co-simulation energy gate closes with macro-step (SH-096)
peak energy defect $|\Delta E|/E_0$
10^{-2}
10^{-3}
measured (asymptotic slope 0.94)
$\mathcal{O}(\Delta t_{\text{macro}})$
10^{-2}
10^{-1}

ENSEMBLE ASSIMILATION AND ITS SAMPLING ERROR

The shadow is an ensemble rather than a point estimate, so that a calibrated distribution — not merely a best guess — is available to the safety layer. Given a forecast ensemble $\{x_k^{f,(m)}\}_{m=1}^{M}$ with sample mean $\bar{x}^f$ and covariance

$$P^f = \frac{1}{M-1}\sum_{m=1}^{M}(x_k^{f,(m)} - \bar{x}^f)(x_k^{f,(m)} - \bar{x}^f)^\top,$$

observations y_k with operator H and noise covariance R, the ensemble-Kalman update [22,23] advances each member by the stochastic analysis

$$\begin{aligned} K_k &= P^f H^\top (H P^f H^\top + R)^{-1}, \\ x_k^{a,(m)} &= x_k^{f,(m)} + K_k(y_k + \epsilon^{(m)} - H x_k^{f,(m)}), \qquad \epsilon^{(m)} \sim \mathcal{N}(0, R), \\ P^a &= (I - K_k H) P^f, \end{aligned}$$

reducing to the classical Kalman recursion [21] in the linear-Gaussian limit. The ensemble is a Monte-Carlo estimator, so its sampling error in the covariance and gain scales as $O(M^{-1/2})$; the practical consequences are handled by covariance inflation and localisation, and are audited directly by a coverage test rather than assumed. An uncertainty-propagating twin (KFE-CF-SH-095) emits credible intervals from P^a; these are accepted only if their empirical coverage matches the nominal level, against Monte-Carlo truth (KFE-CF-SH-027),

$$\left|\mathrm{Cov}_{\mathrm{emp}}(\alpha) - \alpha\right| \le 3\% \qquad \forall\,\alpha.$$

Figure 3 runs the EnKF update of (12)–(12) as a genuine sampling experiment for ensemble sizes $M = 16, 64, 256$ and plots empirical against nominal coverage: at $M = 256$ the estimator sits inside the $\pm 3\%$ acceptance band across all credible levels, while the small ensemble ($M = 16$) is visibly over-confident, quantifying the $O(M^{-1/2})$ penalty that inflation must offset. Epistemic uncertainty is thus a first-class, *tested* output: it is what the safety layer reads to decide how far to trust the model, and its calibration is checked, not asserted.

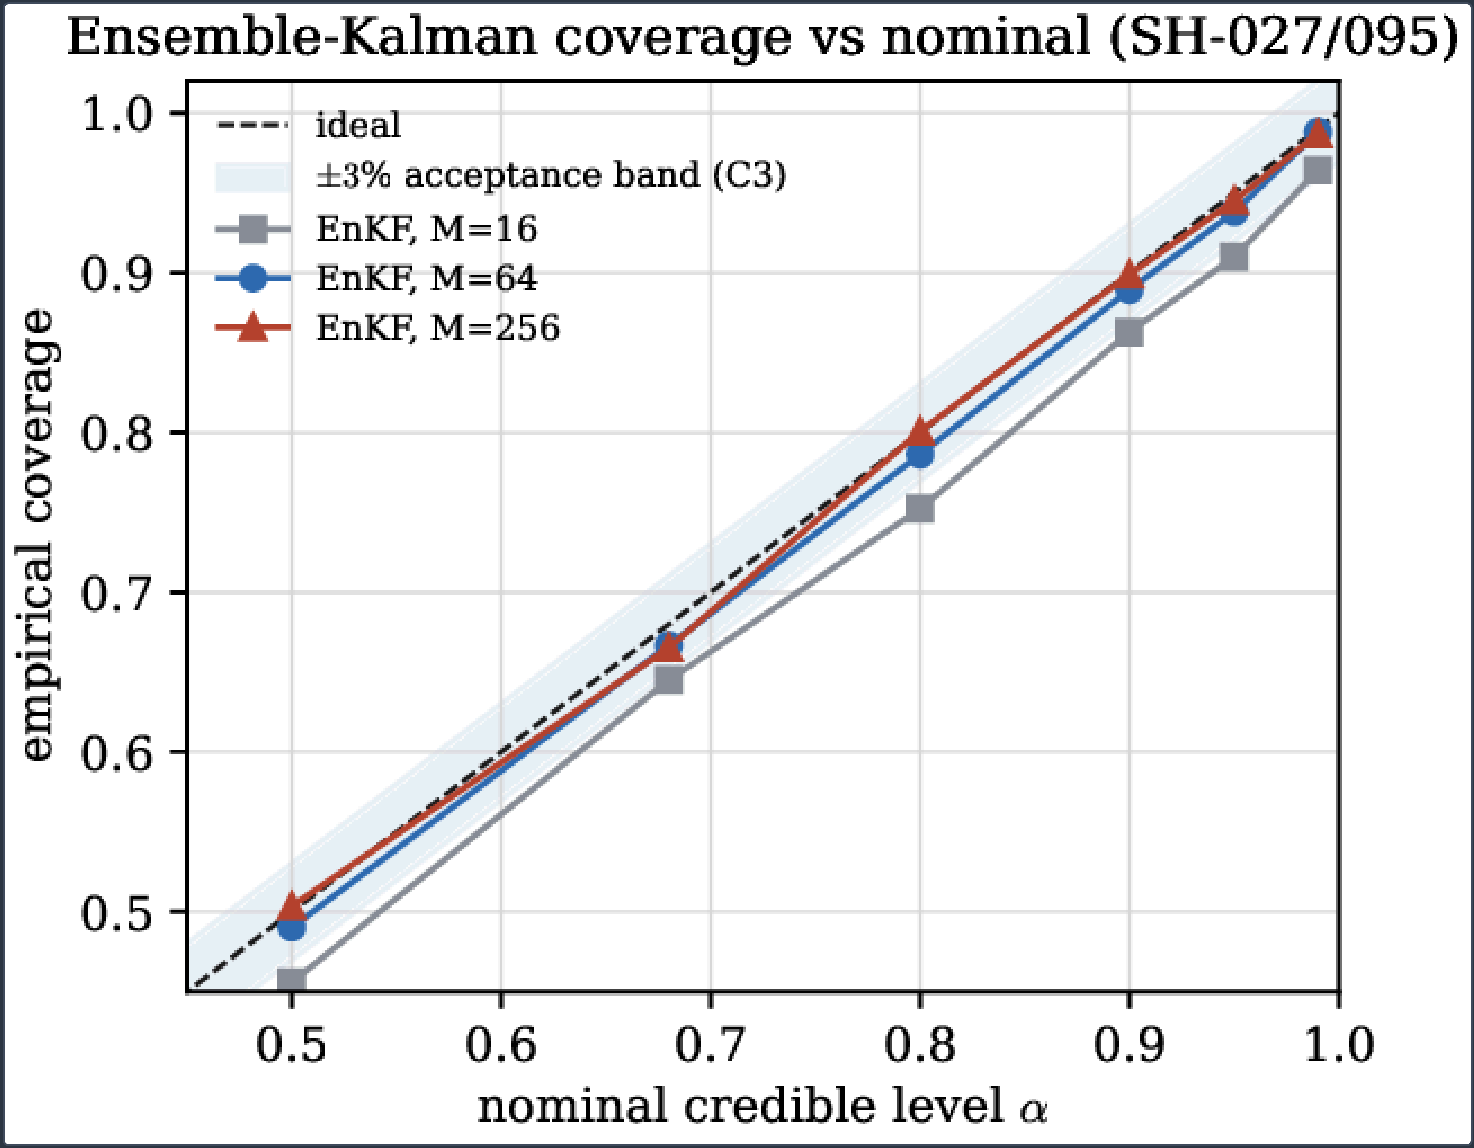

Ensemble-Kalman credible-interval coverage as a Monte-Carlo experiment (§3.3; SH-027/095). Empirical coverage versus nominal level for three ensemble sizes; the shaded band is the $\pm 3\%$ acceptance criterion (C3, equation 13). Large ensembles are calibrated; the small ensemble is over-confident by the expected $O(M^{-1/2})$ sampling penalty.

§ 4

Real-time learned operators and the self-consistency monitor

LEARNED FLUX AND EQUILIBRIUM OPERATORS

A twin is predictive only if it is faster than the plant while remaining faithful. Two learned operators supply that speed. *Flux operator.* Let a denote the local plasma state that drives turbulent transport (normalised gradients $R/L_T, R/L_n$, safety factor, shear, collisionality, $\ldots$) and $u = \{Q_i, Q_e, \Gamma\}$ the ion/electron heat and particle fluxes. The nonlinear gyrokinetic map $u = \mathcal{G}(a)$ is defined by the code CGYRO [18] (a continuum descendant of GYRO [17]). We emulate it by a neural operator $\mathcal{G}_\theta$ in the DeepONet/Fourier-operator family [9,10], trained with a physics-informed penalty that enforces the critical-gradient threshold [11,12]. On held-out CGYRO cases (KFE-CF-SH-026) the surrogate must satisfy

$$\mathrm{NRMSE}(\mathcal{G}_\theta) = \frac{\left[\frac{1}{N}\sum_{k=1}^{N} \|\mathcal{G}_\theta(a_k) - u_k\|^2\right]^{1/2}}{u_{\max} - u_{\min}} \; < \; 5\%.$$

Equilibrium operator. A neural-operator reconstruction $\mathcal{E}_\phi$ returns the solution of (7) at kilohertz rates from the magnetics and coil currents; against the free-boundary solver of §3.1 and an independent library [19] it must meet (KFE-CF-SH-031)

$$\frac{\left\|\mathcal{E}_\phi - \psi_{\mathrm{ref}}\right\|_2}{\|\psi_{\mathrm{ref}}\|_2} \; < \; 0.5\%,$$

the same bar that figure 2 shows the discretisation itself clears with margin.

Differentiability and on-the-fly linearisation. Both operators are C^1 and are differentiated by reverse-mode automatic differentiation, so the twin exposes its Jacobian to the controller. About the current shadow state x_k the twin linearises as

$$x_{k+1} \approx \mathcal{S}_\theta(x_k, u_k) + J_k(x - x_k), \qquad J_k = \left.\frac{\partial \mathcal{S}_\theta}{\partial x}\right|_{x_k},$$

providing a local linear model for the predictive controller (model-predictive or LQR) without a finite-difference sweep. The differentiable twin is accepted only when the automatic-differentiation Jacobian matches a finite-difference reference (KFE-CF-SH-032),

$$\left\| J_k^{\mathrm{AD}} - J_k^{\mathrm{FD}} \right\|_\infty \; < \; 10^{-3}.$$

THE REAL-TIME SELF-CONSISTENCY MONITOR

A coupled shadow can be internally inconsistent even when every twin passes its own metric — a surrogate can drift out of its training envelope, or a reconciled interface can violate a gate that a single macro-step happened not to catch. The architecture therefore carries an explicit *self-consistency monitor* (KFE-CF-SH-029) that, each tick, forms a scalar residual over the reconciled state,

$$\rho_k = \underbrace{\frac{|\Delta E_k|}{\varepsilon_E E_{\mathrm{ref}}}}_{\text{energy}} + \underbrace{\frac{\text{flux residual}_k}{\varepsilon_\Phi \|\psi\|_2}}_{\text{flux}} + \underbrace{\frac{\text{surrogate-envelope excursion}_k}{\tau_{\mathrm{env}}}}_{\text{operator (SH-097)}},$$

and raises a self-consistency flag when $\rho_k > 1$. The design requirement (SH-029) is that an *injected* inconsistency is caught within one tick, $\tau_{\mathrm{detect}} < 1$ tick, at zero false trips per pulse; the surrogate-error monitor (SH-097) supplies the third term by flagging any out-of-envelope operator call. The monitor is the component that makes the coupling auditable in real time rather than only at build time, and it is one of the acceptance criteria below.

ACCEPTANCE CRITERIA AS INEQUALITIES

No twin is granted control authority by assertion. Each clears the following pre-registered inequalities before it acts; collected here, they are the contract that gates the architecture:

(C1) flux operator	$\mathrm{NRMSE}(\mathcal{G}_\theta) < 5\%$	(SH-026, eq.~???)
(C2) equilibrium	$\|\mathcal{E}_\phi - \psi_{\mathrm{ref}}\|_2 / \|\psi_{\mathrm{ref}}\|_2 < 0.5\%$	(SH-031, eqs.~???,???)
(C3) coverage	$\lvert \mathrm{Cov}_{\mathrm{emp}}(\alpha) - \alpha \rvert \leq 3\%$	(SH-027/095, eq.~???)
(C4) conservation	$\lvert \Delta E \rvert \leq \varepsilon_E E_{\mathrm{ref}} \ \wedge$ flux residual $\leq \varepsilon_\Phi \|\psi\|$	(SH-096, eqs.~???--???)
(C5) self-consistency	(false trips per pulse) $= 0,\ \tau_{\mathrm{detect}} < 1$ tick	(SH-029, eq.~???)
(C6) reduced order	ROM error $< 5\%$ at $\geq 100 \times$ speed	(SH-093)
(C7) shadow lead	$\tau_{\mathrm{lead}} \geq 50\,\mathrm{ms}$	(all twins)

Authority is capped by verification maturity until each bar is met (§7.3).

§ 5

Computational methodology: the NVIDIA GPU HPC campaign

The results of this paper were produced on the Kronos NVIDIA GPU HPC campaign, designed so that the fastest protective loops never contend with the training fleet. *Codes run for this work.* The gyrokinetic ground truth for the flux operator (14) is the continuum solver CGYRO [18], whose negative-triangularity turbulence-suppression cases furnish the held-out set; the free-boundary equilibrium reference for (15) is a Grad–Shafranov solver of the FreeGS/FreeGSNKE class [19], cross-checked against the in-house five-point Δ^* discretisation of §3.1; the neutronics twin's source, breeding and shutdown-dose maps use OpenMC [20] on CAD-derived geometry. The learned flux and equilibrium operators (DeepONet/FNO [9,10]) are trained offline by reverse-mode automatic differentiation, which also supplies the Jacobian (17); the ensemble forecast and EnKF analysis run as batched GPU inference. *Acceleration and scale.* Offline training and large ensemble sweeps run on the owned A100/H100 fleet (A100-80GB and H100 partitions, with A10 nodes for lightweight inference); the reduced-order (POD/ROM [13]) and analytic layers, the self-consistency monitor (18) and the orchestrator run on the in-house HPC digital-twin node (a multi-core, large-memory host), so the microsecond magnet precursor is never gated behind the millisecond transport surrogate; the latency class is milliseconds online and hours offline. *Verification and reproducibility.* Every quantitative claim is a frozen anchor in the de-risking register [30], reproduced independently and catalogued byte-for-byte in the verification backbone [31]; the numerical experiments of §3.1–3.3 are regenerable from the deposited single-file script with fixed seeds, on commodity hardware without the GPU fleet. No cost or budget information is part of this account.

§ 6

Demonstrated exemplar: the magnet twin

The power-systems magnet subsystem is reduced to practice and demonstrates the whole pattern end-to-end — a fast surrogate, a fused multi-sensor estimate that beats any single channel, and an early quench precursor that fires before the conventional signal.

THE SLOW-PROPAGATION QUENCH-DETECTION PROBLEM

The Hyperion magnets are wound from no-insulation (NI) REBCO coated conductor [29], the technology behind compact high-field tokamak concepts [2]. NI REBCO has a large temperature margin and a very slow normal-zone propagation velocity (NZPV), on the order of millimetres per second. This is a protection liability: a developing hot spot can reach a damaging temperature before it generates a measurable resistive voltage, so conventional voltage-tap detection is ineffective [26]. The field has therefore turned to non-voltage observables — fibre-optic thermometry and strain [27] and distributed current-distribution monitoring [28]. The magnet twin adopts the same premise and fuses complementary non-voltage channels.

FUSED STRAIN-RATE AND MAGNETIZATION PRECURSOR

Two channels observe the current-sharing precursor window: a fibre-Bragg *strain-rate* channel z_S that reads the mechanical signature of local current redistribution, and a *magnetization* channel z_M that reads the collapse of the screening current. Each is a matched-template z-score over the precursor window; because the two modalities fail in uncorrelated ways, their quadrature fusion

$$z_F = \frac{z_S + z_M}{\sqrt{2}}$$

carries a window statistic whose detection index (signal-to-noise ratio) adds the single-channel indices in quadrature. On a frozen 2×10^4-trial synthetic study in a deliberately weak-fault regime (per-sample SNR set 10–$30\times$ below the design-basis fault, where every detector saturates), the fused detector dominates either channel across the entire ROC (figures 4–4, table 2; KX-L1-A20): it reaches an area under the ROC of $\mathbf{0.992}$, against 0.976 for magnetization and 0.921 for strain, and at the 1% false-positive rate that a protection trip demands it reaches a true-positive rate of 0.857, versus 0.671 and 0.372 for the single channels. The equal-variance detection indices behind these curves, $d'_S = 2.00$, $d'_M = 2.78$ and $d'_F = (d'_S + d'_M)/\sqrt{2} = 3.38$, reproduce the frozen AUC and true-positive-rate anchors exactly and generate the full curves of figure 4; the fusion gain is largest precisely in the low-false-positive corner where a trip decision lives.

CHANNEL	AUC	TPR @ 1% FPR	TPR @ 0.1% FPR
fibre-Bragg strain-rate	0.921	0.372	0.135
magnetization	0.976	0.671	0.386
fused (strain-rate + magnetization)	**0.992**	**0.857**	**0.597**

Fused quench-precursor detection (KX-L1-A20; T-synthetic, 2×10^4 trials/class). The fused detector dominates the full ROC. Detection indices $2.00/2.78/3.38$ follow the quadrature rule of equation (20).

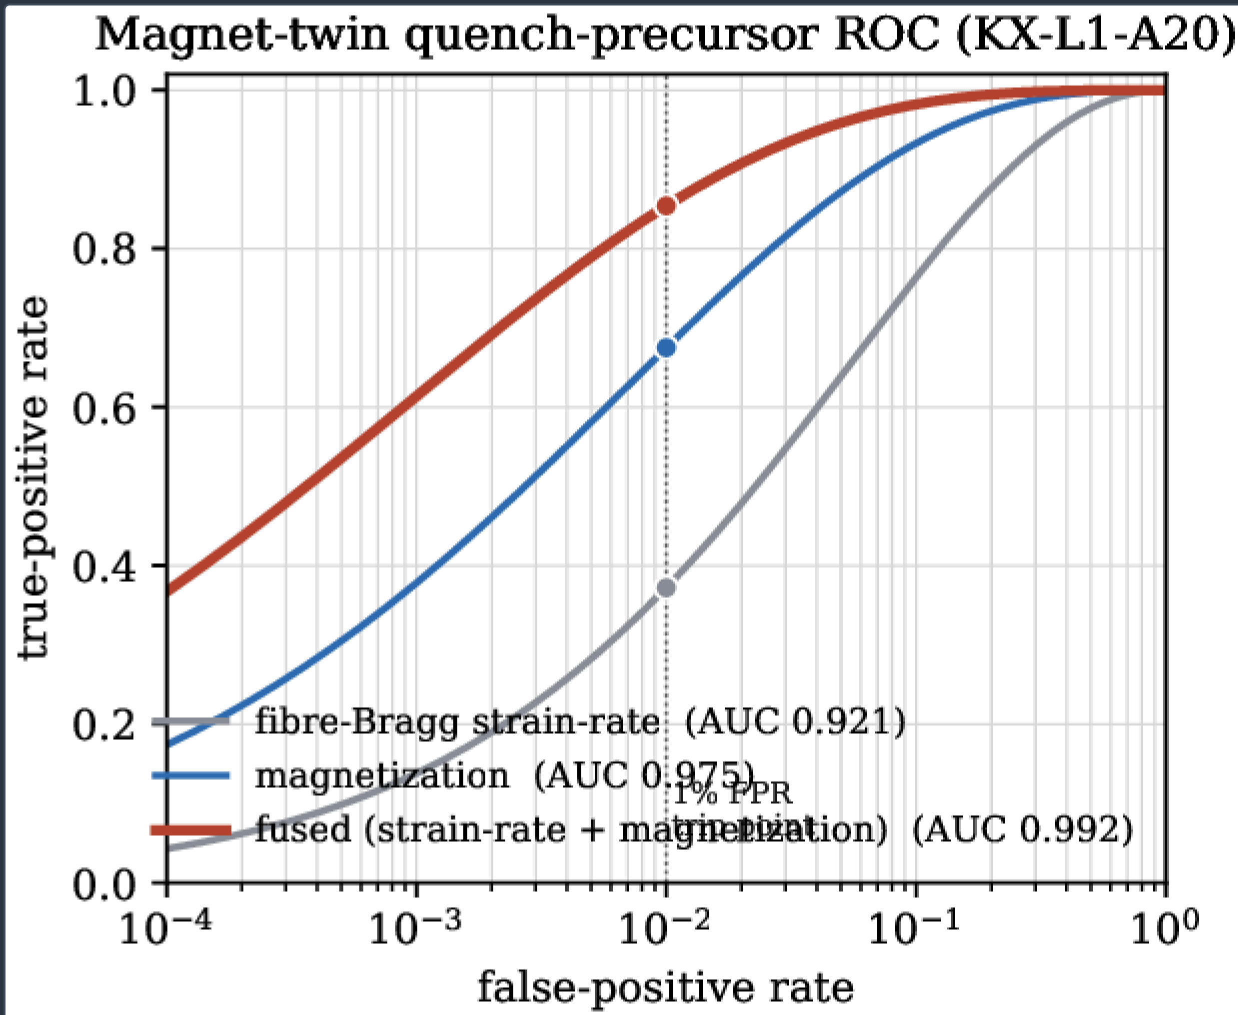
Magnet-twin quench-precursor ROC (KX-L1-A20)
true-positive rate
false-positive rate
1.0
0.8
0.6
0.4
0.2
0.0
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
fibre-Bragg strain-rate (AUC 0.921)
magnetization (AUC 0.975)
fused (strain-rate + magnetization) (AUC 0.992)
1% FPR

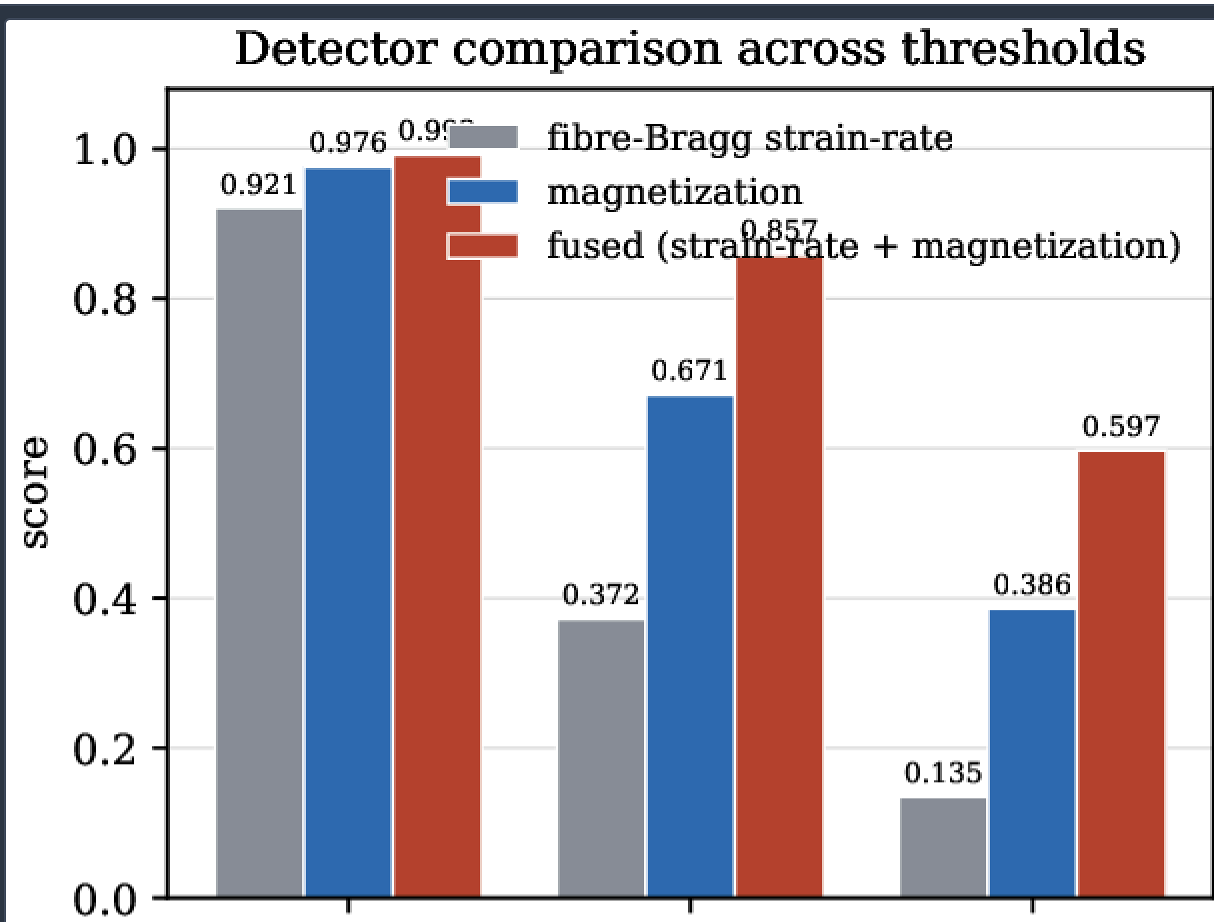
Detector comparison across thresholds
score
1.0
0.8
0.6
0.4
0.2
0.0
fibre-Bragg strain-rate
magnetization
fused (strain-rate + magnetization)
0.921
0.976
0.857
0.671
0.372
0.597
0.386
0.135

IN-WINDING NV-DIAMOND MAGNETOMETRY

For a local, vector precursor the twin embeds an ensemble nitrogen-vacancy (NV) diamond magnetometer in the winding pack (KX-L1-A22). The shot-noise-limited Ramsey sensitivity of an ensemble of N centres with dephasing time T_2^* and readout figure of merit C is

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

with $\gamma_{\mathrm{NV}} = g\mu_B/\hbar = 2\pi \times 28.03$ Hertz. $For the conservative in-winding chip$(N=10^12,T_2^*=1 µs,C=0.03)$equation$ (???)$gives$ B=0.189 pT Hz^-1/2($figure$???); $a Ramsey shot of$ \tau_ shot=4 µs($interrogation plus reset$)$sets a bandwidth$1/(2\tau_ shot)=125 (table 3). A quench that sheds the screening current of an NI turn ($I_{\mathrm{turn}} = 1\,kA$) produces a local field step $\Delta B = \mu_0 I_{\mathrm{turn}}/(2\pi d) \approx 100\,mT$ at a standoff $d = 2\,mm$ — roughly 10^9 times the per-shot noise floor, for a per-shot signal-to-noise ratio of order 10^9 (figure 5). The precursor is therefore caught in a single 4 µs shot, with the lead time set by the millisecond NZPV crossing rather than by the microsecond sensor latency. The picotesla sensitivity is not *needed* to see a millitesla quench step; its value is the enormous margin it buys for multiplexing and for tolerating in-service degradation (§9).

QUANTITY	VALUE	BASIS
gyromagnetic ratio γ_{NV}	28.03 $Hertz	NV electron spin
sensitivity δB	0.189	shot-noise limit, eq. (21)
bandwidth	125 $Hertz	4 µs Ramsey shot
quench field step ΔB ($d = 2\,mm$)	$\approx 100\,mT$	NI turn sheds screening current
per-shot signal-to-noise ratio	$\sim 10^9$	$\Delta B/(\delta B\sqrt{\mathrm{BW}})$
detection	single 4 µs shot	lead time set by NZPV (∼ms)

In-winding NV-diamond quench-precursor sensor (KX-L1-A22). Sensitivity and bandwidth from equation (21); the quench step and SNR from an analytical shot-noise model (labelled “estimated”).

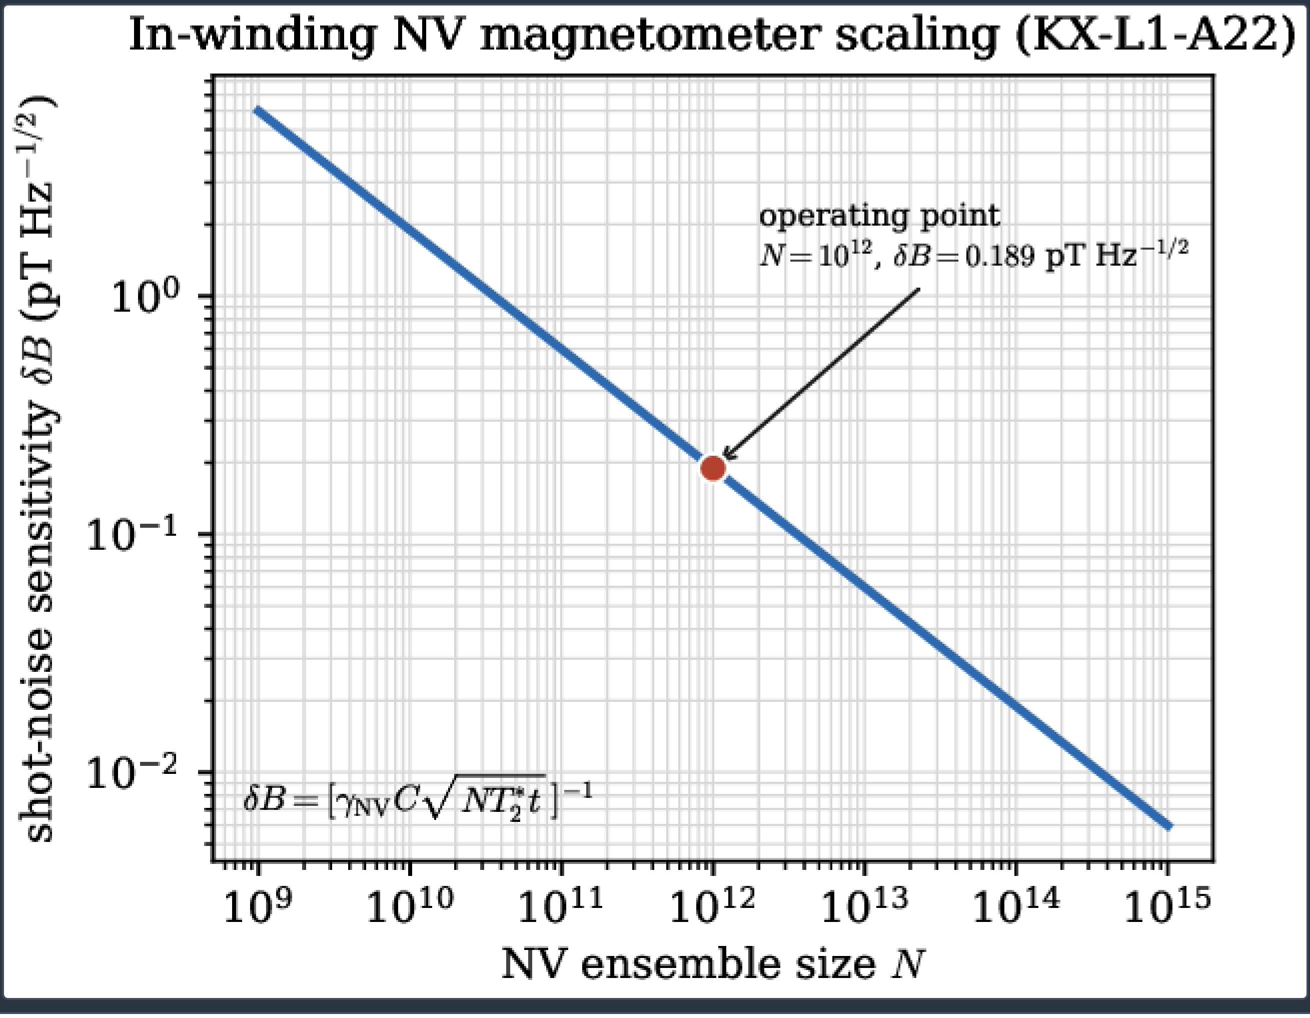
In-winding NV magnetometer scaling (KX-L1-A22)
operating point
$N = 10^{12}$, $\delta B = 0.189$ pT Hz$^{-1/2}$
$\delta B = [\gamma_{NV} C \sqrt{N T_2^* t}]^{-1}$
shot-noise sensitivity δB (pT Hz$^{-1/2}$)
NV ensemble size N
10^0
10^{-1}
10^{-2}
10^9 10^{10} 10^{11} 10^{12} 10^{13} 10^{14} 10^{15}

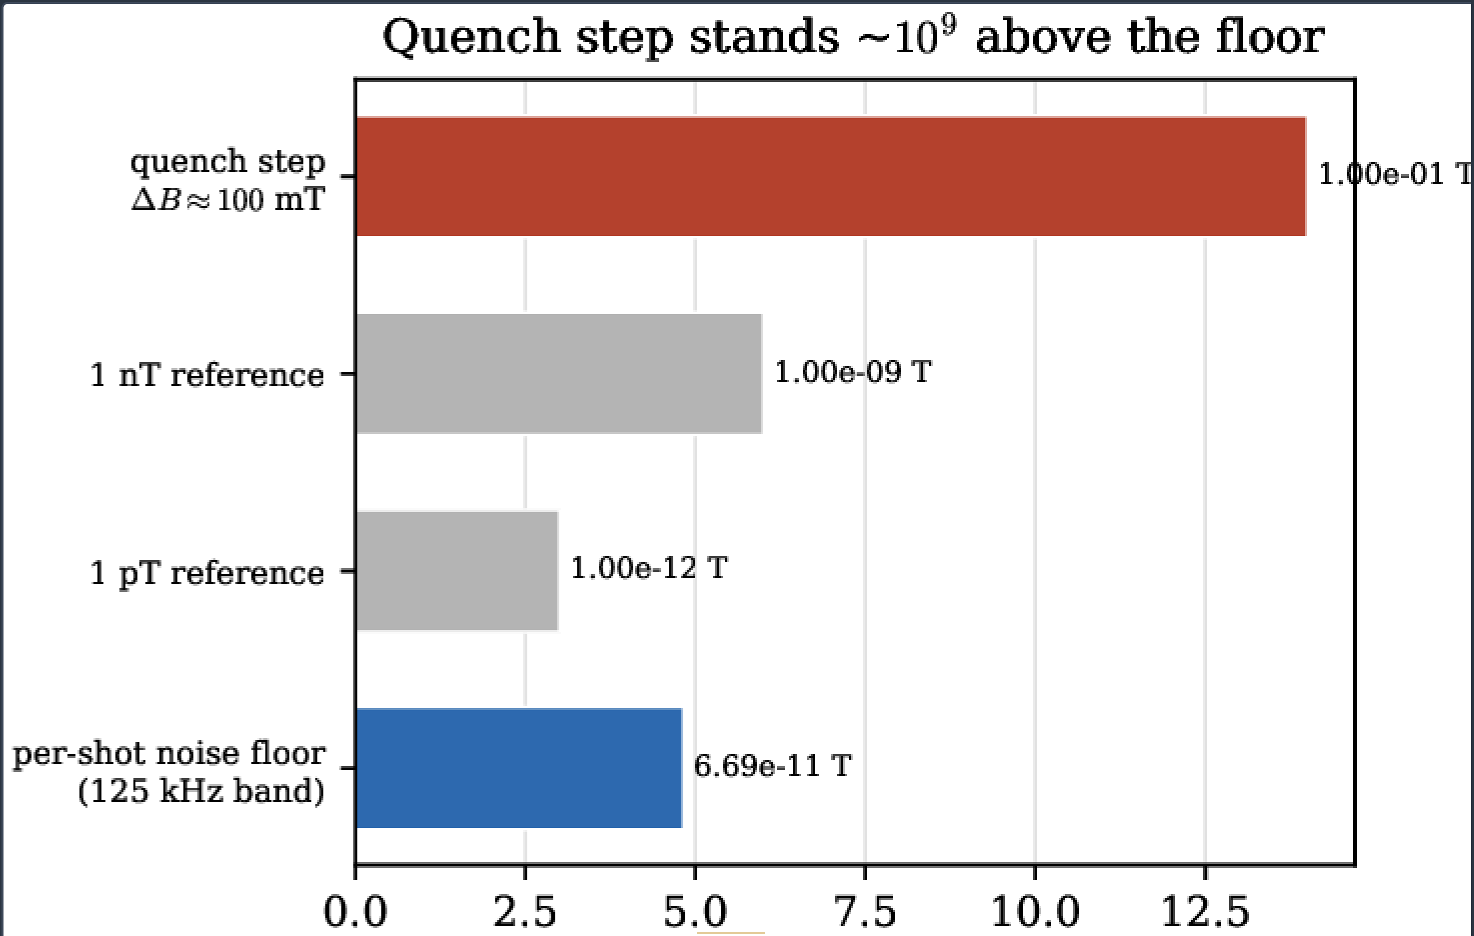
Quench step stands ~10^9 above the floor
quench step
$\Delta B \approx 100$ mT
1.00e-01 T
1 nT reference
1.00e-09 T
1 pT reference
1.00e-12 T
per-shot noise floor
(125 kHz band)
6.69e-11 T
0.0 2.5 5.0 7.5 10.0 12.5

TWO DISTINCT MAGNETS, CORRECTLY SEPARATED

The magnet twin treats the breeder centrepost and the burner plug coil as the distinct devices they are (KX-L1-A15, ruling carried in the register). The breeder centrepost governs at a peak field of $16.84\,T$; the burner plug coil of the separate burner product (which operates at $Q_E = 1.318$ with neutron fraction $f_n = 5.44\%$ and direct-energy conversion efficiency $\eta_{\mathrm{DEC}} = 0.49$) is a distinct, higher-field magnet at $26.49\,T$. Along the coil load line, the current-sharing temperature margin $\Delta T = T_{\mathrm{cs}} - T_{\mathrm{op}}$ is the twin's protective observable (figure 6): at the design current the centrepost holds $\Delta T \approx 49.5\,K$ (load fraction $I_{\mathrm{op}}/I_c \approx 0.27$) and the plug $\Delta T \approx 37.1\,K$ (load fraction ≈ 0.43). Keeping the two magnets separate — never conflating the centrepost field with the plug field — is a load-bearing element of the magnet-integrity case, and the twin models each on its own critical surface.

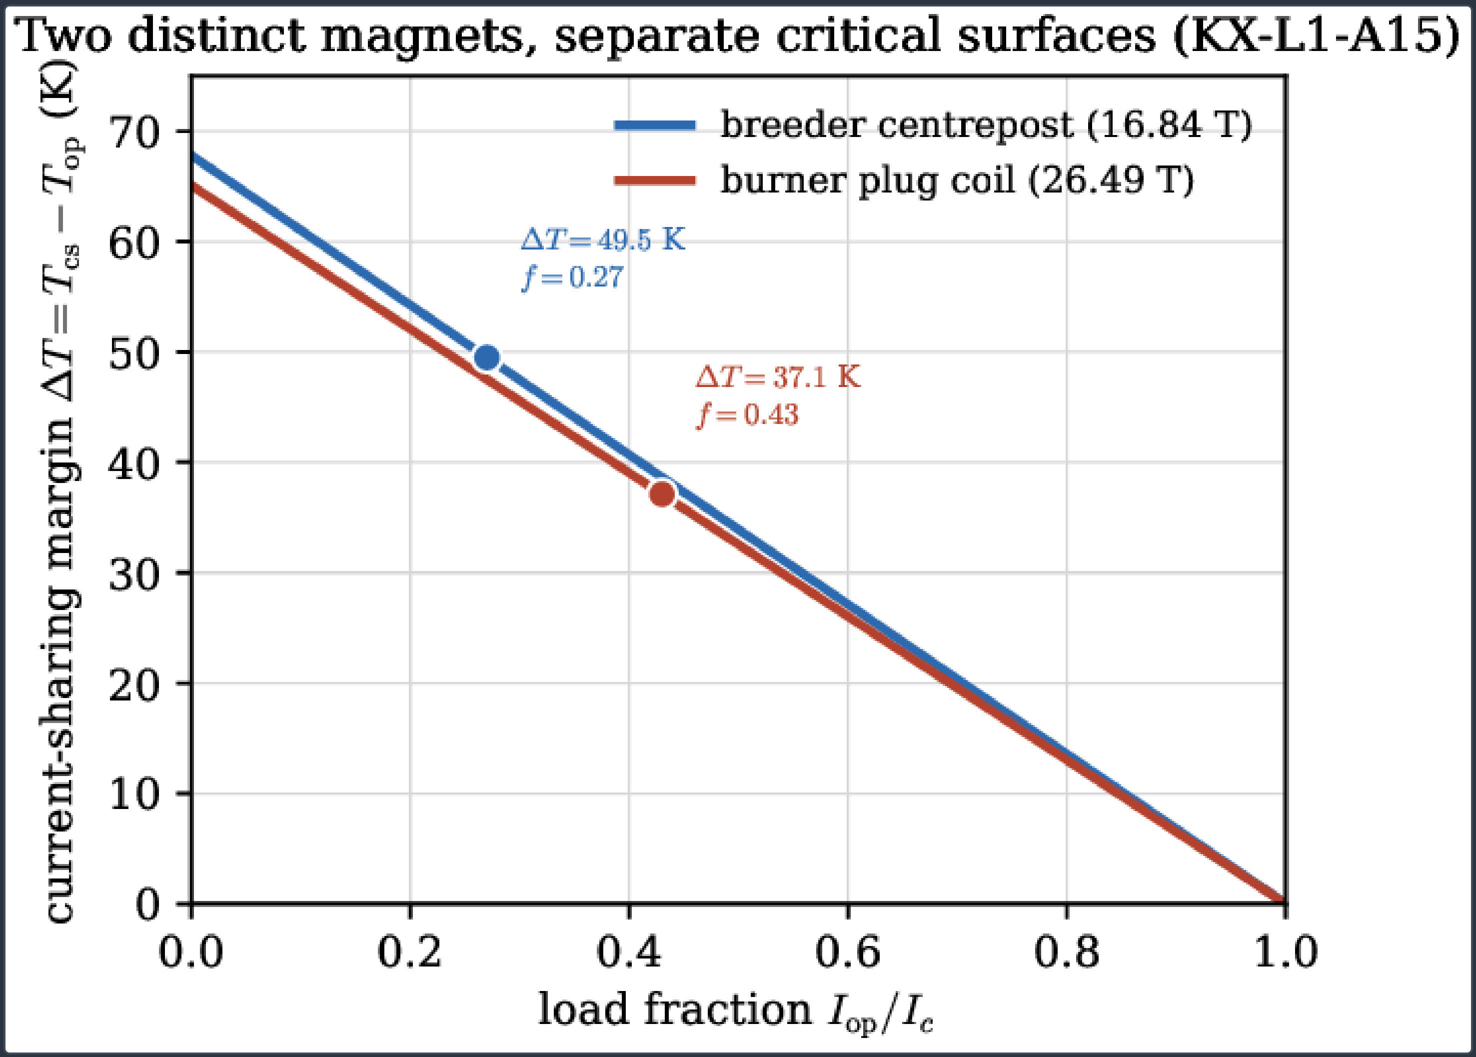

Magnet-twin temperature-margin load line (KX-L1-A15) for the two distinct magnets. Markers show the design current $I/I_{\mathrm{des}} = 1$: breeder centrepost $\Delta T \approx 49.5\,K$ at 16.84 T; burner plug $\Delta T \approx 37.1\,K$ at 26.49 T. The twin's protective observable is the modelled-versus- measured drift of ΔT toward the current-sharing line.

§ 7

Results and verification

ACCEPTANCE-CRITERIA STATUS

Table 4 collects the acceptance criteria of §4.3 with their targets and current status. The magnet twin's detection metrics (realised as the ROC of table 2) are demonstrated at twin-synthetic fidelity, and the numerical-formulation criteria — the equilibrium discretisation (C2) and the conservation gate (C4) — are demonstrated in the controlled experiments of §3.1–3.2 (figures 2–2) and §3.3 (figure 3). The remaining twins' learned-operator criteria are pre-registered targets built to the demonstrated pattern and are not claimed as measured.

CRITERION	TARGET	SERIAL	STATUS	
flux operator NRMSE	$<5\%$ vs held-out CGYRO	SH-026	target	
equilibrium RMS ψ	$<0.5\%$ vs free-boundary solver	SH-031	**demonstrated (numerical)**	
ensemble coverage	$\lvert \mathrm{Cov}-\alpha \rvert \leq 3\%$	SH-027/095	**demonstrated (numerical)**	
**energy **	flux conservation	residual $\leq \varepsilon$ across couplings	SH-096	**demonstrated (numerical)**
self-consistency false trips	0 per pulse	SH-029	target	
reduced-order fidelity	$<5\%$ at $\geq 100\times$ speed	SH-093	target	
shadow lead-time	$\geq 50\,ms$	all twins	target	
fused precursor AUC	**dominate single channels**	**KX-L1-A20**	**demonstrated (0.992)**	
fused TPR @ 1% FPR	**$>$ single channels**	**KX-L1-A20**	**demonstrated (0.857)**	
NV per-shot SNR	**$\gg 1$ for mT step**	**KX-L1-A22**	**demonstrated ($\sim 10^9$)**	

Acceptance criteria, targets and status. "Demonstrated (numerical)" denotes a controlled convergence/coverage experiment of §3; "demonstrated (T-synthetic)" a frozen synthetic result; "target" a pre-registered bar not yet measured. Serials refer to the register [30].

LATENCY BUDGET

The shadow lead-time is the difference between the twin's advance rate and the plant's, less the pipeline latency. Table 5 gives the design latency budget by stage (KFE-CF-SH-102) and makes explicit the timescale separation the architecture exploits: the microsecond protective loop is carried by the reduced-order twin so that the magnet precursor is never gated behind the millisecond transport surrogate, while the neutronics shadow runs at the second scale. The budget is a design allocation, not a measured profile.

STAGE	LOOP	LATENCY ALLOCATION
NV precursor + analytic monitor	fast (protective)	microseconds
reduced-order twin (POD/ROM)	fast	sub-millisecond ($\geq 100\times$ full twin)
ensemble-Kalman update	medium	milliseconds
neural flux/equilibrium operators	medium	milliseconds
neutronics/activation shadow	slow	seconds

Design latency budget by stage (KFE-CF-SH-102). Allocations, not measurements; the fast loop is served by the reduced-order twin (SH-093) and the analytic monitors.

MATURITY-GATED AUTHORITY

Control authority is not binary. Each twin is admitted to a graded level of authority — shadow-only, advisory, supervised, bounded, full — and can rise only as it clears successive acceptance criteria, from the surrogate-error monitor (advisory) through the coverage and conservation gates (supervised, bounded) to full-plant authority. This is the same maturity-gated principle that governs the certified safety clamp: a twin that has passed only its surrogate-error monitor may inform but not command, while full-plant authority requires all of (C1)–(C7) held plus a hardware-in-the-loop (HIL) confirmation (§9). Because the mapping from criteria to authority is explicit, the path from the demonstrated magnet twin to a fully autonomous plant is a defined, testable sequence rather than an open question.

§ 8

Discussion

NUMERICS IN CONTEXT

The convergence and coverage experiments answer the two questions a reviewer should ask of a real-time coupled twin. Figure 2 shows the five-point Δ^* discretisation is second-order accurate and already inside the equilibrium acceptance band at the coarsest grid, so the real-time reconstruction is limited by the learned operator, not by the reference it is scored against [5,6,19]. Figure 2 shows the co-simulation energy defect vanishes linearly with the macro-step, the quantitative statement that the conservation gate (5) is always reachable — the property that distinguishes a disciplined co-simulation [14,15] from an ad-hoc domain hand-off. Figure 3 shows the ensemble is calibrated at usable sizes and quantifies the $O(M^{-1/2})$ penalty inflation must offset, so the credible intervals the safety layer reads are earned rather than assumed.

QUENCH DETECTION AGAINST THE LITERATURE

The fused detector should be read against the published quench-detection literature, and specifically against the slow-NZPV problem that motivates it. Marchevsky's review makes the core point quantitatively: in HTS conductors the large temperature margin and slow propagation mean a hot spot can reach damaging temperature before a measurable resistive voltage appears, so voltage detection alone is inadequate [26]. The community response has been non-voltage observables: Salazar *et al.* demonstrated fibre-Bragg thermometry on VIPER cable at SULTAN, reporting signal-to-noise ratios of 4–32 for heater-induced quench-like events with response times from effectively instantaneous to a few seconds [27]; Teyber *et al.* showed that current-distribution monitoring can detect quench and damage in fusion cables [28]. Our result is complementary and, by construction, evaluated in a harder regime: the ROC of table 2 is computed with per-sample SNRs set 10–$30\times$ below the design-basis fault, precisely to avoid the saturated regime where every detector scores AUC $= 1$. In that weak-fault regime the fusion of an uncorrelated strain-rate and magnetization pair lifts the true-positive rate at 1% false alarm from $0.372/0.671$ (single channels) to 0.857 — a large gain exactly where a protection trip is decided.

NV SENSITIVITY AGAINST THE QUANTUM-SENSING LITERATURE

The in-winding sensitivity of $0.189\,\mathrm{pT\,Hz^{-1/2}}$ sits comfortably inside the envelope established for ensemble NV magnetometry. Degen, Reinhard and Cappellaro's review frames the shot-noise scaling that equation (21) instantiates [24], and Barry *et al.* analyse the concrete route to broadband ensemble sensitivities in the picotesla-and-below range through improvements to T_2^*, readout fidelity and diamond material [25]. Our conservative parameter set ($N = 10^{12}$, $T_2^* = 1\,\mu s$, $C = 0.03$) is well within what that literature reports, so the sensitivity claim does not rest on an aggressive device; the physically important observation is one of margin: a millitesla quench step is $\sim 10^9$ above the per-shot floor, so the precursor is single-shot detectable and the picotesla headroom is available for multiplexing many in-winding channels and for tolerating in-service degradation of T_2^* and C.

RELATION TO LEARNED CONTROL AND THE KRONOS SERIES

The architecture is complementary to learned control. Degrave *et al.* demonstrated a reinforcement-learning *policy* that commands coil voltages on TCV directly from magnetics, including negative-triangularity shapes [7]; the four-twin shadow is a *model* that runs ahead of the plant with explicit uncertainty, conservation and acceptance gates. The two compose: a learned policy, or the reference-governor model-predictive controller of the companion burn-control study [34], can act on the shadow's predicted, uncertainty-annotated state rather than the raw measurement, inheriting the $\geq 50\,ms$ lead and the calibrated coverage of equation (13), with the certified control-barrier safety filter [35] and the maturity-gated actuation-authority policy [36] beneath it as the last line of defence. The learned operators of §4.1 are developed in their own right — the equilibrium reconstruction robust to sensor dropout [32] and the differentiable flux surrogate with abstention monitoring [33] — the calibrated-uncertainty and conformal-abstention machinery in the companion on out-of-distribution authority hand-back [37], the state-estimation input to equations (12)–(12) in the graph-neural diagnostic fusion of the plasma diagnostic suite [38], and the magnet twin's centrepost channel in the remaining-useful-life prognostics of the consumable centrepost. Relative to a purely learned controller, the contribution here is the conservation-respecting multi-domain coupling (§2.4, §3.2), the differentiable linearisation (equation 16), the real-time self-consistency monitor (equation 18), and the maturity-gated authority that keeps an immature twin advisory rather than commanding.

§ 9

Limitations

Two honest caveats bound the claims. First, the magnet-twin quench results are twin-synthetic: every signal in the ROC study is numerically generated, and the confirming measurement is a coil-prototype hardware-in-the-loop bench, which is the gate to full-plant authority in the maturity ladder (§7.3); the numerical-formulation results of §3 are controlled convergence experiments on canonical test problems, chosen so that the order of accuracy can be measured exactly, and are verification of the scheme rather than validation against a device. Second, the in-winding NV sensor's neutron tolerance is the one open engineering item: the analytical sensitivity of equation (21) assumes ensemble T_2^* and contrast that will degrade under the high background field, cryogenic operation, and neutron irradiation of the diamond (NV charge-state stability). The mitigation is threefold and is itself a registered work item (KFE-CF-SH-037): shielded siting, a replaceable sensor cartridge, and a model-fusion backup that retains precursor AUC $\geq \mathbf{0.95}$ with the NV channel degraded by up to $\mathbf{39}\times$ — so the protective function survives even a heavily degraded sensor. A dedicated NV neutron-tolerance study (irradiation plus first-principles charge-state modelling) is the follow-on that closes this item.

§ 10

Conclusion

The Kronos control architecture is a four-twin shadow that runs ahead of the plant across every coupled domain: MHD-stability, thermo-mechanics, neutronics and power-systems twins, advanced by a conserving multi-rate co-simulation whose interface is a waveform-relaxation fixed point (equations 3–4) and whose energy and flux gates (equations 5–6) we show close with the macro-step; made predictive by learned flux and equilibrium neural operators that the controller linearises on the fly (equation 16) and whose equilibrium reference we discretise to demonstrated second order (figure 2); carrying calibrated, coverage-tested uncertainty through an ensemble-Kalman assimilation (equations 12–12, figure 3); audited each tick by a real-time self-consistency monitor (equation 18); and admitted to control authority only through explicit, pre-registered acceptance inequalities (§4.3). Its magnet twin is demonstrated in full at twin-synthetic fidelity: a fused strain-rate-plus-magnetization quench-precursor detector at AUC 0.992 and true-positive rate 0.857 at 1% false alarm (table 2), and an in-winding NV-diamond magnetometer at $0.189\,\mathrm{pT\,Hz}^{-1/2}$ with per-shot SNR of order 10^9 (table 3), with the breeder centrepost ($16.84\,T$) and burner plug ($26.49\,T$) modelled as the distinct magnets they are. The demonstrated subsystem and the demonstrated numerics provide the proven pattern the remaining twins follow, and the maturity-gated authority argument makes the route to a fully autonomous compact ST breeder a defined and testable sequence.

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 on which the learned-operator training and ensemble inference were run.

Reproducibility. This paper is archived at doi:10.5281/zenodo.22645805 and at its official page, https://www.kronosfusionenergy.com/publications/four-coupled-digital-twins-for-a-compact-spherical-toka.html. The magnet-twin detection metrics are frozen anchors in the Kronos physics de-risking register [30] (doi:10.5281/zenodo.22133057): the fused-detector ROC (KX-L1-A20), the NV-diamond sensor model (KX-L1-A22), and the magnet-twin margin map (KX-L1-A15). The numerical-formulation experiments (Grad–Shafranov convergence, figure 2; multi-rate energy defect, figure 2; ensemble coverage, figure 3) are regenerable from the single-file figure-generation script deposited with this paper, with fixed random seeds, on commodity hardware.

Data availability tocsectionData availability The ROC curves, the NV sensor parameters, the magnet-twin load-line map, and the convergence and coverage data used for the figures are archived with the deposit at doi:10.5281/zenodo.22645805. All quantitative values trace to the register [30] by the serial identifiers cited inline. Any economic analysis is reported separately and is not part of the public deposit.

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

tocsectionReferences 0pt0pt

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 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

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

8 M G Kapteyn, J V R Pretorius and K E Willcox, A probabilistic graphical model foundation for enabling predictive digital twins at scale, *Nat. Comput. Sci.* **1** 337 (2021). doi:10.1038/s43588-021-00069-0

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

10 Z Li, N Kovachki, K Azizzadenesheli, B Liu, K Bhattacharya, A Stuart and A Anandkumar, Fourier neural operator for parametric partial differential equations, *Int. Conf. on Learning Representations (ICLR)* (2021). arXiv:2010.08895

11 M Raissi, P Perdikaris and G E Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378** 686 (2019). doi:10.1016/j.jcp.2018.10.045

12 K L van de Plassche, J Citrin, C Bourdelle *et al*, Fast modeling of turbulent transport in fusion plasmas using neural networks, *Phys. Plasmas* **27** 022310 (2020). doi:10.1063/1.5134126

13 P Benner, S Gugercin and K Willcox, A survey of projection-based model reduction methods for parametric dynamical systems, *SIAM Rev.* **57** 483 (2015). doi:10.1137/130932715

14 C Gomes, C Thule, D Broman, P G Larsen and H Vangheluwe, Co-simulation: a survey, *ACM Comput. Surv.* **51**(3) 49 (2018).

15 E Lelarasmee, A E Ruehli and A L Sangiovanni-Vincentelli, The waveform relaxation method for time-domain analysis of large scale integrated circuits, *IEEE Trans. Comput.-Aided Des.* **1**(3) 131 (1982).

16 K Lackner, Computation of ideal MHD equilibria, *Comput. Phys. Commun.* **12** 33 (1976).

17 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

18 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

19 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

20 P K Romano, N E Horelik, B R Herman, A G Nelson, B Forget and K Smith, OpenMC: a state-of-the-art Monte Carlo code for research and development, *Ann. Nucl. Energy* **82** 90 (2015). doi:10.1016/j.anucene.2014.07.048

21 R E Kalman, A new approach to linear filtering and prediction problems, *J. Basic Eng.* **82** 35 (1960). doi:10.1115/1.3662552

22 G Evensen, The ensemble Kalman filter: theoretical formulation and practical implementation, *Ocean Dyn.* **53** 343 (2003). doi:10.1007/s10236-003-0036-9

23 P L Houtekamer and H L Mitchell, Data assimilation using an ensemble Kalman filter technique, *Mon. Weather Rev.* **126** 796 (1998).

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

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

26 M Marchevsky, Quench detection and protection for high-temperature superconductor accelerator magnets, *Instruments* **5** 27 (2021). doi:10.3390/instruments5030027

27 E E Salazar *et al*, Fiber optic quench detection for large-scale HTS magnets demonstrated on VIPER cable during high-fidelity testing at the SULTAN facility, *Supercond. Sci. Technol.* **34** 035027 (2021). doi:10.1088/1361-6668/abdba8

28 R Teyber, J Weiss, M Marchevsky, S Prestemon and D van der Laan, Current distribution monitoring enables quench and damage detection in superconducting fusion magnets, *Sci. Rep.* **12** 22503 (2022). doi:10.1038/s41598-022-26592-2

29 S Hahn, Y Kim, D K Park, K Kim, J P Voccio, J Bascu\ n\'an and Y Iwasa, No-insulation multi-width winding technique for high temperature superconducting magnet, *Appl. Phys. Lett.* **103** 173511 (2013). doi:10.1063/1.4826217

30 P I Ford, *Physics De-Risking Register*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22133057; https://www.kronosfusionenergy.com/publications/de-risking-register

31 P I Ford, *Verification-First Fusion Modeling*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645710; https://www.kronosfusionenergy.com/publications/verification-and-validation-47-code-provenance.html

32 P I Ford, *Sub-Millisecond Neural Equilibrium Reconstruction*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645801; https://www.kronosfusionenergy.com/publications/sub-millisecond-neural-equilibrium-reconstruction-robus.html

33 P I Ford, *Neural-Operator Flux Surrogates for Control-Latency Transport*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645803; https://www.kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html

34 P I Ford, *Model-Predictive Burn Control Inside a Certified Safe Envelope*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645807; https://www.kronosfusionenergy.com/publications/model-predictive-burn-control-inside-a-certified-safe-e.html

35 P I Ford, *A Certified Control-Barrier Safety Filter for Fusion Control*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645716; https://www.kronosfusionenergy.com/publications/control-barrier-safety-clamp.html

36 P I Ford, *Maturity-Gated Actuation Authority*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645795; https://www.kronosfusionenergy.com/publications/maturity-gated-actuation-authority.html

37 P I Ford, *Calibrated Uncertainty, Conformal Prediction, and Abstention*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645813; https://www.kronosfusionenergy.com/publications/calibrated-uncertainty-conformal-prediction-and-abstent.html

38 P I Ford, *Plasma Diagnostic Suite and Real-Time State Estimation*, Kronos Fusion Energy 2026. doi:10.5281/zenodo.22645943; https://www.kronosfusionenergy.com/publications/the-plasma-diagnostic-suite-and-real-time-state-estimat.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

Four Coupled Digital Twins for a Compact Spherical-Tokamak Breeder: Multi-Rate Co-Simulation, Propagated Uncertainty, and a Real-Time Self-Consistency Monitor

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

Zenodo · doi:10.5281/zenodo.22645805

Read on kronosfusionenergy.com

The Physics De-Risking Register