KRONOS FUSION ENERGY

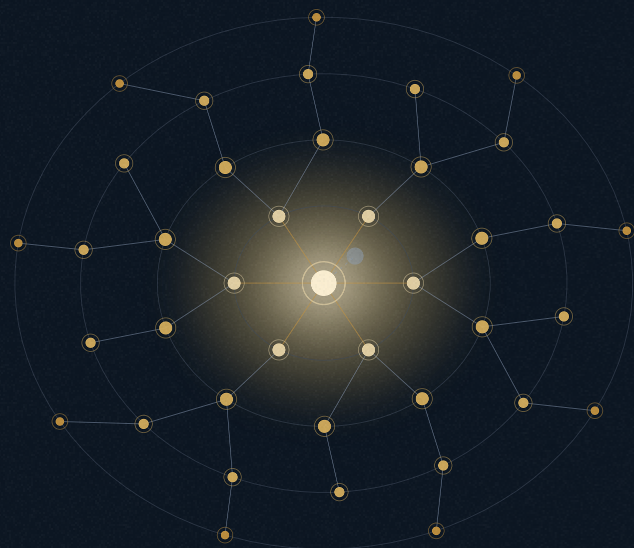

PAPER V OF THE 2026 KRONOS SET

# *AI, ML & Quantum Control*

## *Certifiable Real-Time Control*

Physics-informed AI and quantum technologies for certifiable real-time control of compact fusion generators — and the digital twin that carries it.

Layered digital twin · 47.2 ms warning · 0 / 150k escapes · quantum: no crossover this decade

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

PREPARED BY P. I. FORD

2026 · EDITORIAL

KRONOS FUSION ENERGY · 2026

# *The Kronos Fleet*

*Fuel follows purpose, not platform.*

Hyperion · Aegis · MetroVolt

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

| | |
|---|---|
| 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

# AI, ML & Quantum Control — Certifiable Real-Time Control

Real-time control of a nuclear-regulated fusion generator is treated as a layered digital-twin architecture in which physics-informed machine learning is deployable now — but only inside a deterministically clamped safe set. On a physics-based twin the controller gives 47.2 ms of warning before a disruption and permits zero safety escapes across 150,000 steps. Quantum sensing is a near-term diagnostic tier; quantum *computation* is an offline calibration tier with no resource crossover this decade. Every result is obtained on a physics-based twin, on public data, or in simulation — no operating-plant data is involved.


Companion deposit (code & data, reproducible under CC BY 4.0): DOI 10.5281/zenodo.21842371 · Zenodo community *kronos_fusion_energy*.

Explore it live: interactive 3-D model · run the physics validator · companion film series (the deposited solvers, in the browser).

**Keywords:** digital twin · physics-informed ML · certifiable control · quantum sensing · safe-set clamp

CONTENTS

# Table of Contents



## PORTFOLIO, PEOPLE & RECORD

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) |
| $\tau_E$ | energy confinement time |
| $I_p$ | plasma current |
| $R_0, a$ | major radius, minor radius |
| $\kappa, \delta$ | elongation, triangularity |
| $\beta_N$ | 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) |

§ 5.1

# The control problem, and the three tiers held apart

Every credible path to commercial fusion — tokamak, mirror, or otherwise — must clear the same non-physics gate: a control system a regulator will license and an operator will trust to run an autonomous, high-energy-density plant. That gate is not cleared by any single learned model; it is cleared by an *architecture* in which learned components are fast enough to anticipate the plasma yet provably cannot drive the machine outside a safe set. This paper is about that architecture. It is written as a methods-and-certification study, on a digital twin, precisely because the honest sequencing is to prove the control and certification methods in simulation first, then earn the right to operate hardware. The decade of design work behind the two Kronos machines used here as worked examples is the substrate; the contribution is the intelligence and the safety case that would run them.

A fusion generator couples processes whose characteristic times span ten orders of magnitude: **~1 μs** for magnet-quench protection and plasma actuation, ~**1 ms** for shape, β, and conversion supervision, and **hours-to-months** for retraining, maintenance, and fuel logistics. Conventional supervisory control provides neither physics-faithful prediction at actuation latency nor a principled way to bind machine-learning inference to deterministic safety limits. The difficulty is acute in machines combining high-field REBCO magnets, negative-triangularity shaping, direct energy conversion (DEC) of charged products, and staged low-neutron fuel cycles, where the multi-physics coupling is strong and the regulatory demand — that every automated action be traceable and bounded — is absolute.

Machine learning has real capability here: reinforcement learning has controlled tokamak plasma shape on a physical device, and deep learning has predicted disruptions across machines. What is *not* established is how such capability is made **certifiable** for a nuclear-regulated plant — how a learned model's inevitable errors are bounded, clamped, and audited so the safety case does not rest on the model being right. That is this paper's subject.

## THREE TECHNOLOGY-READINESS TIERS, HELD STRICTLY APART

The single organising discipline of KRONOS-CTRL is that three tiers of technology are never allowed to bleed into one another. Conflating them is the most common way fusion-AI claims become unfalsifiable.

| | |
|---|---|
| AI — NOW, IN THE LOOP | Physics-informed surrogates, graph state reconstruction, anomaly detection, and predictive/RL control at control latency — but **only wrapped by a deterministic clamp** and a hardware failsafe outside the learned model. |
| QUANTUM SENSING — NEAR-TERM | Quantum magnetometry and entanglement-enhanced metrology sharpen the measured state the AI consumes; much of this is general diagnostics R&D at the *diagnostics* tier. |
| QUANTUM COMPUTATION — OFFLINE | Fault-tolerant hardware does not exist at control latency and, per our estimates, offers **no crossover this decade**; its honest role is computing offline closures that calibrate classical surrogates. A present-tense role is played by quantum-*inspired* classical tensor-network methods. |

**THE FALSIFIABILITY TEST THE WHOLE PAPER IS BUILT TO PASS**

A claim that “a quantum computer controls the plasma” is physically false. “Quantum-computed closures reduce our surrogate error offline, with an explicit resource estimate and crossover” is testable — and here, tested. The contribution is a certifiable AI-control architecture and a quantum-calibrated-surrogate pipeline: *a system and a method, not new plasma physics or new algorithms.* All results are on the twin and on public/simulated data — **Kronos operates no plant.**

## THREE DISCLOSURES GOVERN EVERY CLAIM

These three caveats travel with every number in this paper, and are restated wherever a headline figure appears:

| | |
|---|---|
| NO OPERATING-PLANT DATA | Kronos operates no plant; every result is from a physics-based digital twin, public data, or simulation, tagged per figure (SIM / TWIN / PUBLIC). |
| NO NEAR-TERM QUANTUM ADVANTAGE | Fault-tolerant resource estimates for three fusion kernels find no classical crossover this decade — the equilibrium linear solve alone needs **$\approx 1.8\times10^{17}$ T-gates** once read-in/read-out are counted. |
| THE QUANTUM POC IS SIMULATOR-ONLY | The variational proof-of-concept reaches chemical accuracy on a statevector simulator and misses it under a device-noise model; no quantum hardware was run. |

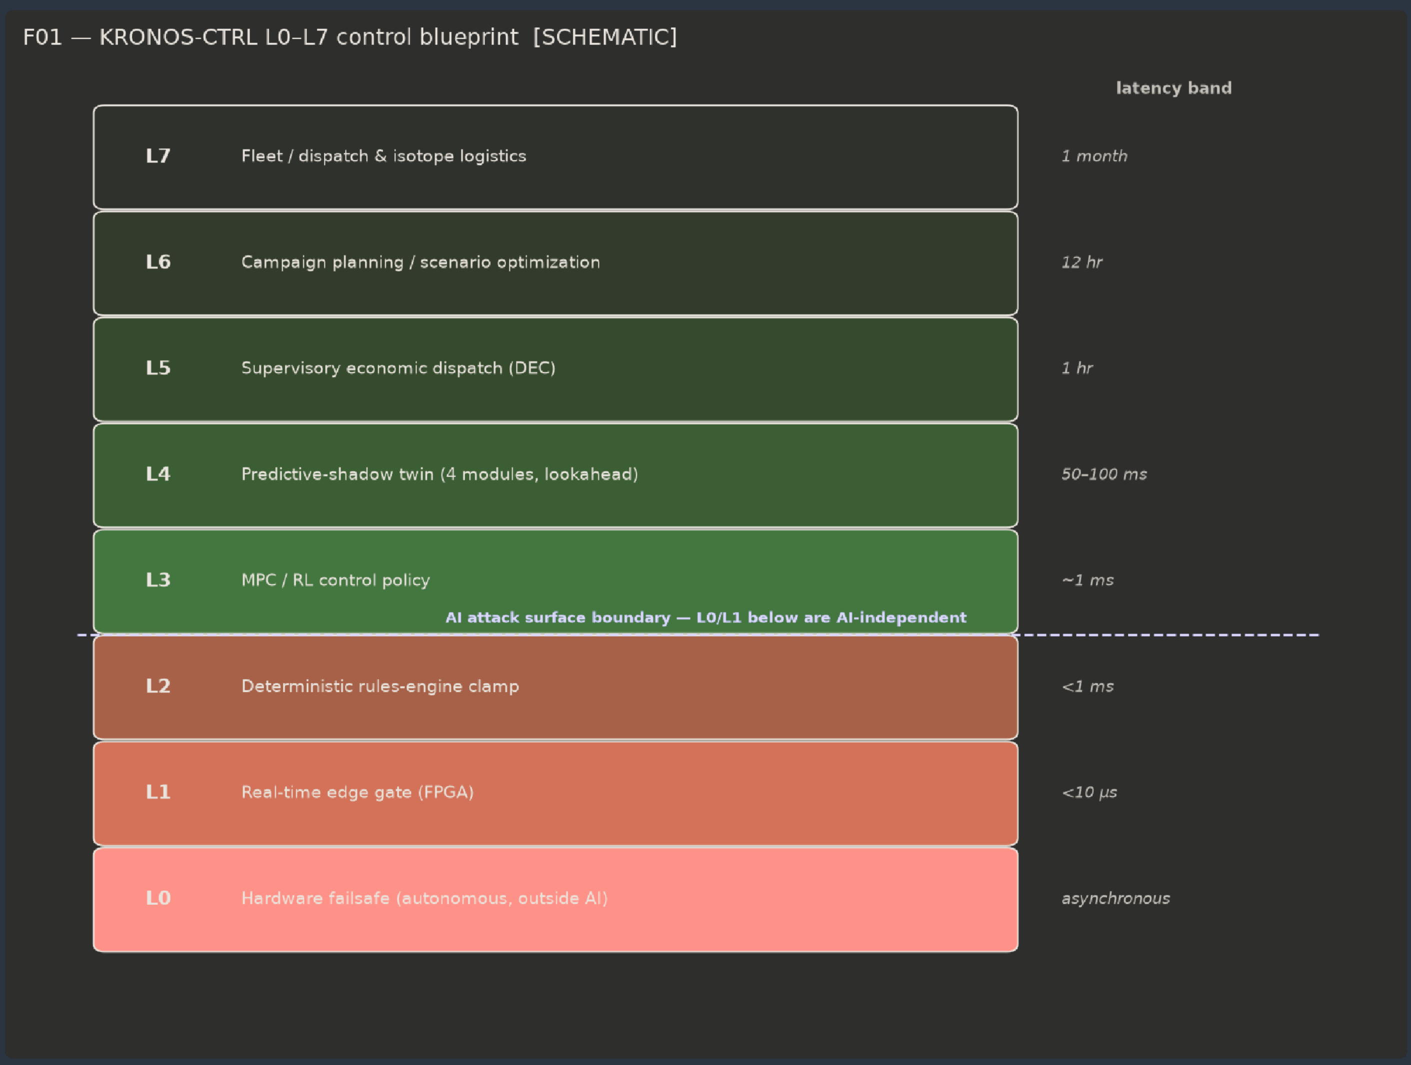


**FIGURE 1.** KRONOS-CTRL master blueprint: layers L0–L7 coupled to the physical plant by three signal classes, with the four-module predictive-shadow twin at the core. SCHEMATIC

§ 5.2

# Results at a glance — and every miss kept

Seventeen results anchor the study, each reproducing off a script, each carrying its own tier and provenance tag. The honest misses that accompany several of them — the neural operator's 2.5% equilibrium error against a 1% target, the 22.3 ms median disruption warning against a 30 ms target, the $\beta_N$ safety face that closes only as a sub-region, the tensor-network solver losing at small bond dimension — are reported in the limitations section and in each result's own artifact, not reframed. The quantum-computing crossover is a *pre-registered negative*.

### TABLE 1 · RESULTS AT A GLANCE (HEADLINE NUMBER, TIER, PROVENANCE)

| RESULT | HEADLINE NUMBER | TIER | PROV. |
|---|---|---|---|
| **Deterministic safety clamp** | **3000/3000 caught, 0 escapes** / 150k steps | AI-now | TWIN |
| **Calibration spine (chemistry)** | **1691**× surrogate-error drop; chem-accuracy @5 anchors | offline | SIM |
| **Calibration spine (plasma transport)** | **36.8**× @8 anchors, 0.82% rel — beats floor 9.9× | offline | SIM |
| **Calibration spine (magnet)** | **7475**× max; 1856× @8 | offline | SIM/TWIN |
| **Formal safety certificate** | forward-invariance 3/4 faces @±3σ; preserved through calibration (10/10) | AI-now | SIM |
| **Magnet per-turn pre-load (M1)** | strain → **0.017%** on all 82 turns (limit 0.40%), 0.15 ms | design-study | TWIN |
| **Per-coil qualification (M4)** | **+30.0%** usable $J_{wp}$ vs worst-case derating | design-study | TWIN |
| **Fused quench precursor (M10)** | **47.2 ms** lead vs 2.6 ms; AUC 0.9998 | design-study | TWIN |
| **Runtime strain clamp (M12)** | **0 escapes** on 0.40% cap / 150k steps | design-study | TWIN |
| **Neural-operator equilibrium** | BST **2.5%** (2× better than PINN); mirror 0.16% | AI-now | SIM |
| **GNN sensor imputation** | **2.1**× better than interpolation | AI-now | SIM/TWIN |
| **Disruption ensemble** | AUC **0.990**; OOD→failsafe cuts missed-alarm ~0.9→0 | AI-now | TWIN |
| **Isotope-balancing control** | tracking RMS **0.0105** vs 0.143 open-loop (13.6×) | AI-now | TWIN |
| **Surrogate design optimization** | genuine **3.8**× over random; MC engine byte-reproduced | AI-now | SIM |
| **Tensor-network kinetic solver** | **21.2**× @d=7; only feasible method at d≥10 | q-inspired | SIM |
| **Independent-code validation** | twin vs FreeGS **2.66%**; circularity halved 0.348→0.165 | AI-now | SIM/PUBLIC |
| **ITER engineering benchmark** | 19/19 assumptions tagged to ITER doc IDs | eng-basis | ENG-BASIS |


**FIGURE 2.** Results scorecard across the pre-registered acceptance criteria: most targets met; the quantum-computing crossover is a pre-registered negative; every miss is kept. SIM/TWIN

§ 5.3

# The architecture and the equation canon (Part I)

KRONOS-CTRL is organised as **eight layers** — L0 foundation compute; L1 edge real-time control; L2 governed data fabric; L3 physics-informed models; L4 orchestration with a deterministic safety-clamping rules engine; L5 domain copilots; L6 experience surfaces; L7 external integrations — each assigned an enforced latency band from **~1 µs** at the edge to **~1 month** for supply planning. No function executes outside its band. The four-module predictive-shadow twin — MHD stability, thermomechanics, neutronics, power systems — synchronizes to telemetry at $t = 0$ and integrates **50–100 ms ahead** of the plant, so divergence emits preemptive corrections. Control is converted from reactive to anticipatory.

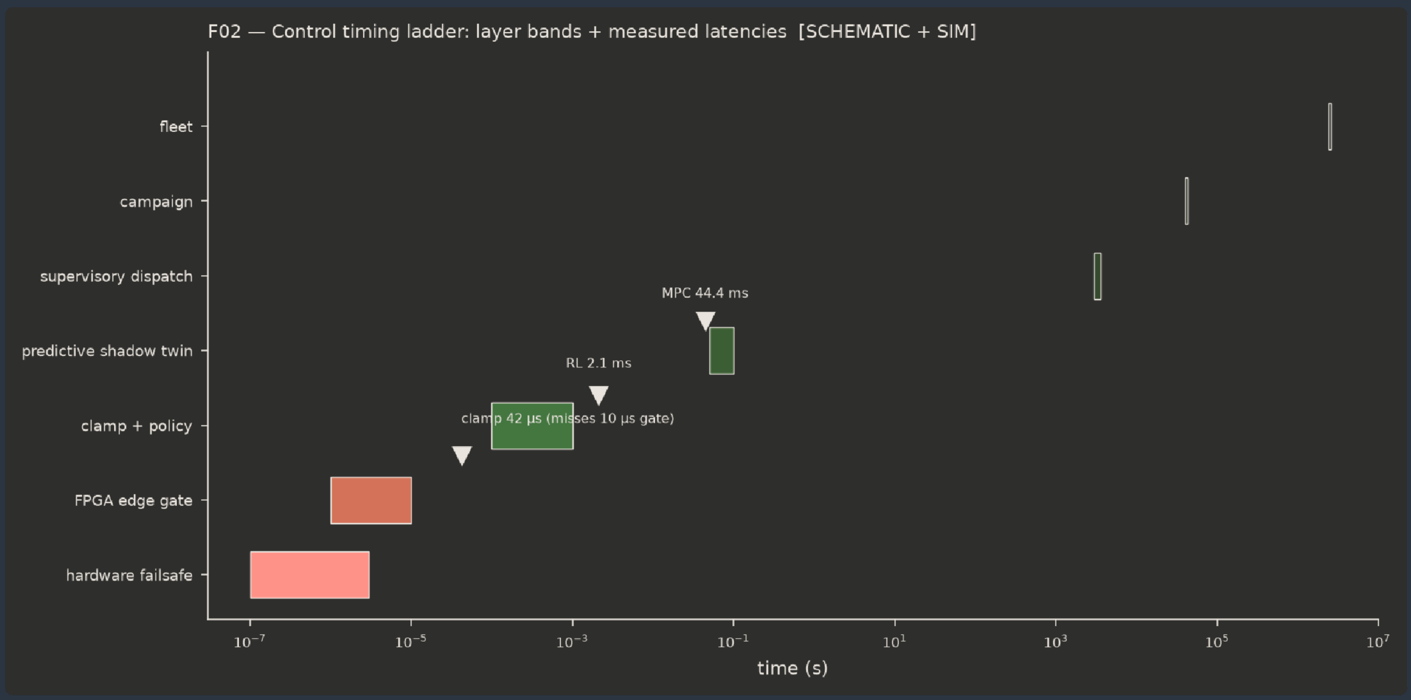


**FIGURE 3.** Cross-layer timing ladder: every layer is assigned an enforced latency band; no function executes outside its band. SCHEMATIC

## THE GOVERNING EQUATIONS, EACH MAPPED TO ITS METHOD

The twin is not a single model but a canon of governing equations, each mapped to the AI, sensing, or quantum method that addresses it. The load-bearing set:

| | |
|---|---|
| E1 GRAD–SHAFRANOV (BST) | $\Delta^*\psi = -\mu_0 R^2\, dp/d\psi - F\, dF/d\psi$, with $\Delta^* \equiv R^2\nabla\cdot(\nabla/R^2)$, $F = RB_\varphi$ |
| E1B MIRROR PARAXIAL FORCE BALANCE (BTM) | Stated separately — the tandem mirror is **not** a Grad–Shafranov machine. |
| E2 IDEAL MHD | $\rho\, dv/dt = J\times B - \nabla p$; energy principle $\delta W \geq 0$; $\beta$, $\beta_N$ limits |
| E3 KINETIC TRANSPORT | Fokker–Planck / gyrokinetic: $\partial f/\partial t + v\cdot\nabla f + (q/m)(E+v\times B)\cdot\nabla_v f = C(f)$ — the calibration-tier target |
| E5 RADIATION FLOORS | $P_{br} \propto Z_{eff}\, n_e^2 \sqrt{T_e}$; generalized-Trubnikov synchrotron |
| E6 THERMOMECHANICS | $\varepsilon \leq 0.40\%$ (operating 0.245%); bore-resolved winding stress $\sigma \approx ½\, J\, B\, a$ (E6b) |

| | |
|---|---|
| E9 PINN OBJECTIVE | $L = L_{data} + \lambda_{pde} \lVert \Delta^* \psi + \mu_0 R^2 p' + FF' \rVert^2 + \lambda_{bc} L_{bc}$ |
| E10–E12 | UQ; measurement precision ($\Delta\varphi$: $1/\sqrt{N}$ vs $1/N$); confinement ($\tau_E$, $H_{98}$) |

F03 — Equation → twin module → ML method → quantum tier map [SCHEMATIC]

| Eq | Physics | Twin module | ML method | Quantum tier |
|---|---|---|---|---|
| E1 | Plasma equilibrium (closed-flux to | MHD-stability (L3) | PINN (native solve) | COMPUTE-OFFLINE |
| E1b | Plasma equilibrium (open-field-lin | MHD-stability(mirror) | PINN (separate architectur | NONE |
| E2 | Ideal MHD force balance and stabil | MHD-stability (L3) | GNN + MPC | COMPUTE-OFFLINE |
| E3 | Kinetic transport: Fokker-Planck w | closure (offline L0 -> | surrogate closure | COMPUTE-OFFLINE |
| E4 | Neutronics: linear Boltzmann trans | neutronics (L3) | surrogate | COMPUTE-OFFLINE |
| E5 | Radiation floors: bremsstrahlung a | power-systems (L3) | none (closed-form floor) | NONE |
| E6 | Thermomechanics: REBCO strain, pre | thermomechanics (L3) | GNN on fibre-Bragg strain | COMPUTE-OFFLINE |
| E7 | Direct energy conversion power bal | power-systems (L3) | RL/MPC dispatch | COMPUTE-OFFLINE |
| E8 | Materials / first wall: H-isotope | offline (L0) | none at the canon level | COMPUTE-OFFLINE |
| E9 | Control/surrogate objective: the p | L3 surrogate -> L4 det | PINN objective | NONE |
| E10 | Uncertainty quantification, calibr | all (L0-L3) | Bayesian optimisation / ac | NONE |
| E11 | Diagnostics / metrology precision | data fabric (L2) -> L3 | GNN/PINN consume the sharp | SENSING-NEAR-TERM |
| E12 | Energy confinement and transport | confinement/transport | surrogate -- and a Track-6 | NONE |

NONE (classical) SENSING — near-term COMPUTE — offline only

**FIGURE 4.** Equation canon → method map: which AI, sensing, or quantum method addresses each governing equation. SCHEMATIC

## THE FOUR-MODULE PREDICTIVE SHADOW

The twin is four coupled physics modules, each solving its own governing equation on its own state and exchanging boundary data once per synchronization step. The **MHD-stability module** solves E1/E1b for the equilibrium and evaluates the E2 stability margins ($\beta_N$ against the no-wall and ideal-wall limits, edge-localized-mode drive under negative triangularity). The **thermomechanics module** carries E6/E6b: per-turn REBCO strain under Lorentz loading, the cryogenic pre-load state, and fiber-Bragg strain telemetry. The **neutronics module** tracks residual neutron output of the low-neutron fuel cycle and the isotope inventory that closes the E3 fueling loop. The **power-systems module** computes direct-energy-conversion collector voltages and grid synchronization. The lookahead horizon is set by the slowest actuator in the fast loop (poloidal-field shaping) and bounded above by the accumulation of model error, which the certification section bounds.

> **WHY A MODULE MAY BE WRONG WITHOUT ENDANGERING THE PLANT**
>
> Crucially, the modules' *predictions*, not their internal states, are what the L4 rules engine consumes — so a module may be wrong without the plant being endangered, because every machine-originated command is projected onto the safe set before dispatch. The safety case never rests on a model being correct.

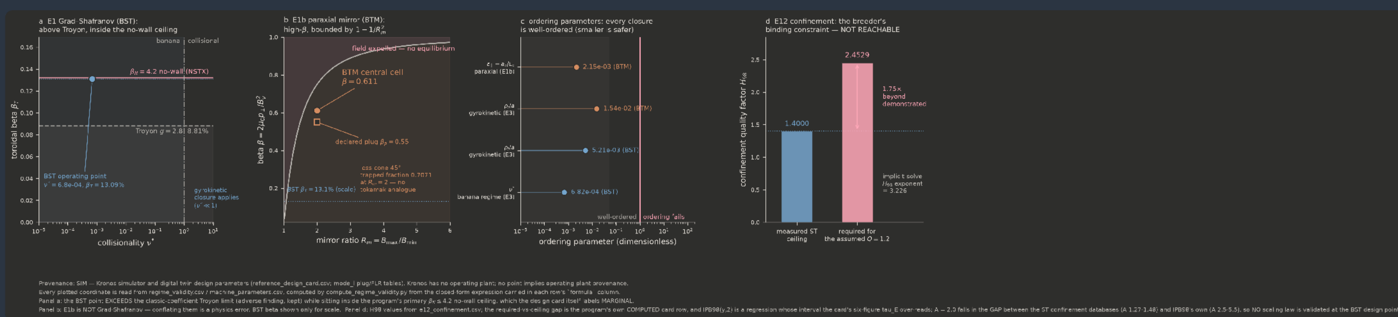


**FIGURE 5.** Governing-equation regime map: the validity domain of each twin module in ($\beta_N$, $\nu^*$, geometry). SIM

## A SHOT, START TO FINISH

The architecture is easiest to read as a single pulse. Before breakdown, the operations layer sequences cryogenic cooldown and vacuum checkout while the thermomechanics module verifies that the magnet pre-load state is inside the strain envelope. At breakdown, the diagnostic fabric streams sixty-plus channels through the L2 refinement pipeline into the state estimate; quantum-sensor-grade magnetometry, where fitted, sharpens the boundary condition the equilibrium surrogate consumes. Through the flat-top the four-module twin runs 50–100 ms ahead of the plant: the equilibrium and stability modules propose shape and $\beta$ corrections, the model-predictive policy turns them into actuator commands, and — critically — **every command is projected onto the safe set by the L4 clamp before it reaches a coil**, with the autonomous hardware failsafe standing outside the loop entirely. The isotope-balancing loop holds the fuel mix on its minimum-neutron trajectory, and the power-systems module re-weights the direct-conversion channels as that mix evolves. If a disruption precursor the model has never seen appears, the out-of-distribution layer abstains and hands the machine to the failsafe rather than issuing a low-confidence command. Offline, between shots, high-fidelity closures — classical today, quantum-computed on the dated roadmap — are admitted through the provenance gate and, *only after the safety invariant is re-proven*, refresh the surrogates for the next pulse. No step depends on a model being correct; each depends only on the clamp and the failsafe being correct.

§ 5.4

# Methods and reproducibility — latency figures that cannot flatter

All learned components are trained and evaluated on **CPU** (56 logical cores, no GPU), so every latency figure is an honest lower bound, not a GPU-deployment claim. The Monte-Carlo design engine (seed **20260726**) reproduces its feasibility counts exactly. Every headline number is read off a reproducing script; the deposit ships a run-all driver, environment pins, and MD5 checksums.

A two-tier register governs reproducibility. **Tier-1** — the Monte-Carlo engine, quantum resource estimates, statevector results, calibration curves — reproduces *byte-identically*. **Tier-2** — trained-model metrics — is reported as five-seed tolerance bands with pinned versions; we do *not* claim byte-identity for trained-model metrics. Provenance is tagged per figure: SIM, TWIN, PUBLIC, TRL-SURVEY, RESOURCE-EST, ENGINEERING-BASIS.

**A KEPT DATA LIMITATION, STATED UP FRONT**

Labelled public disruption data for a real-data ROC was **not retrievable**: open MAST metadata carries no disruption labels, and labelled machine databases are credential-gated. Every affected result is therefore labelled *twin-synthetic* — anchored to the real MAST operating space, but generated on the twin, not measured on a plant. TWIN-SYNTHETIC

§ 5.5

# AI on the twin — equilibrium surrogates and state reconstruction (Part II)

The AI layer earns the right to sit in the loop one function at a time, and each function reports its miss as plainly as its hit. A classical fixed-boundary Grad–Shafranov solver was verified against an analytic Solov′ev solution to relative L2 error **$2.2\times10^{-5}$** — an independent ground truth, not the twin validating itself. Against it, a physics-informed network (PINN) reached best relative flux error **5.5% on the BST** and **0.16% on the BTM** mirror problem; on the BST the physics-residual term degraded accuracy through a diagnosed *gradient-conflict pathology*. A Fourier neural operator improved the BST error to **2.5% (worst 3.9%)** and removed the gradient conflict — but did not clear the 1% target and amortized only 2.7× over $10^3$ queries.

**KEPT NEGATIVE — A WARM-START, NOT A FAST NATIVE SOLVER**

The equilibrium surrogate did **not** become a fast native solver on the spherical tokamak: 2.5% error against a 1% target. We therefore position it as a warm-start / accelerator, reported in full rather than reframed. On the mirror problem the PINN does reach 0.16%. SIM · KEPT MISS

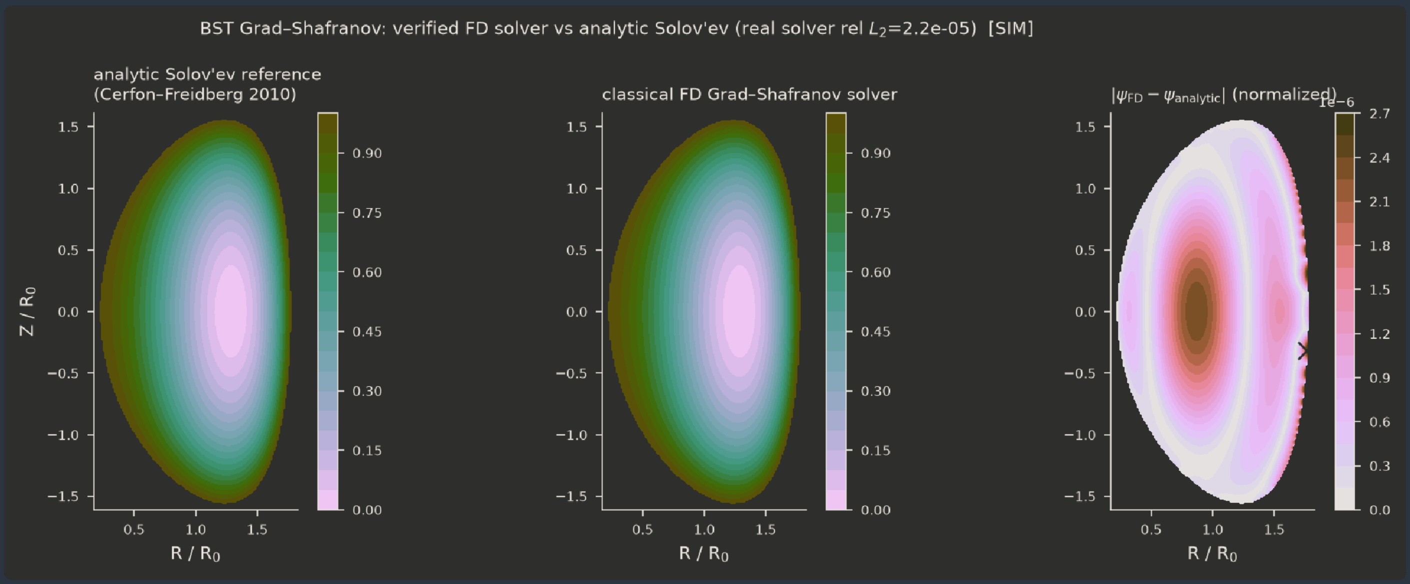


**FIGURE 6.** BST equilibrium: the classical Grad–Shafranov solver verified against the analytic Solov'ev solution to rel L2 $2.2\times10^{-5}$ — the independent ground truth. SIM

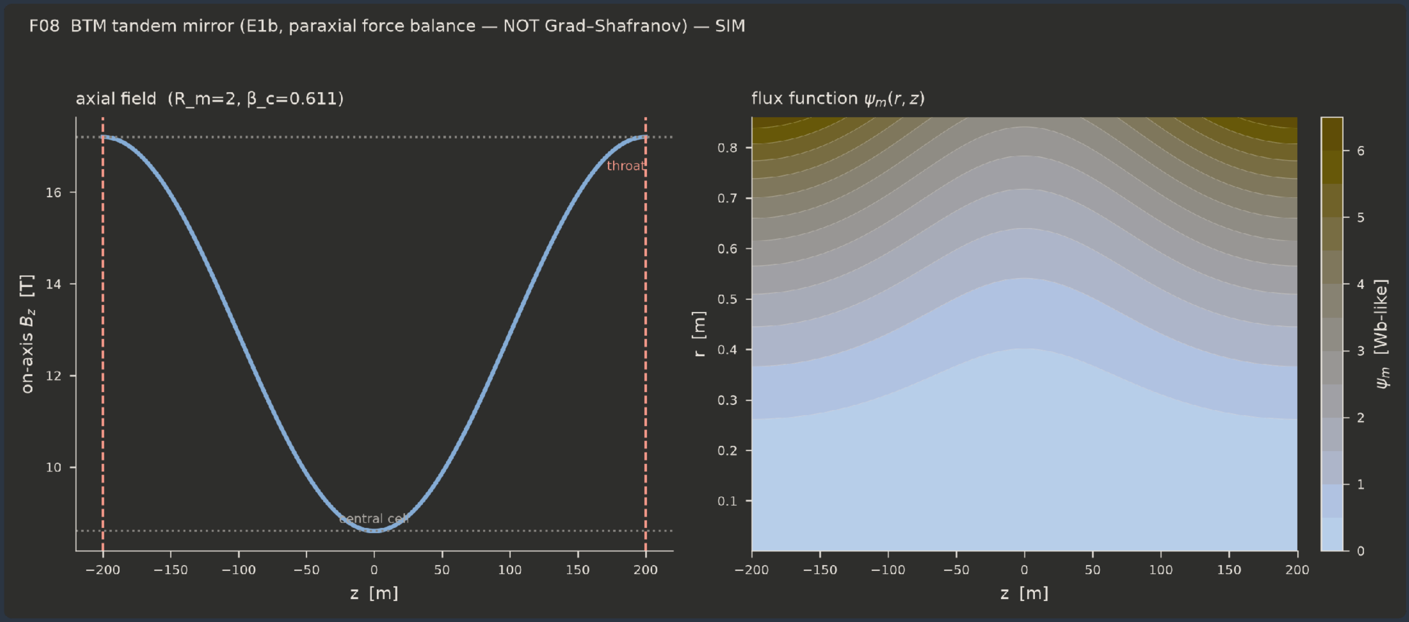


**FIGURE 7.** BTM tandem-mirror equilibrium from the paraxial force balance (E1b, not Grad–Shafranov); the PINN reaches 0.16% relative error. SIM

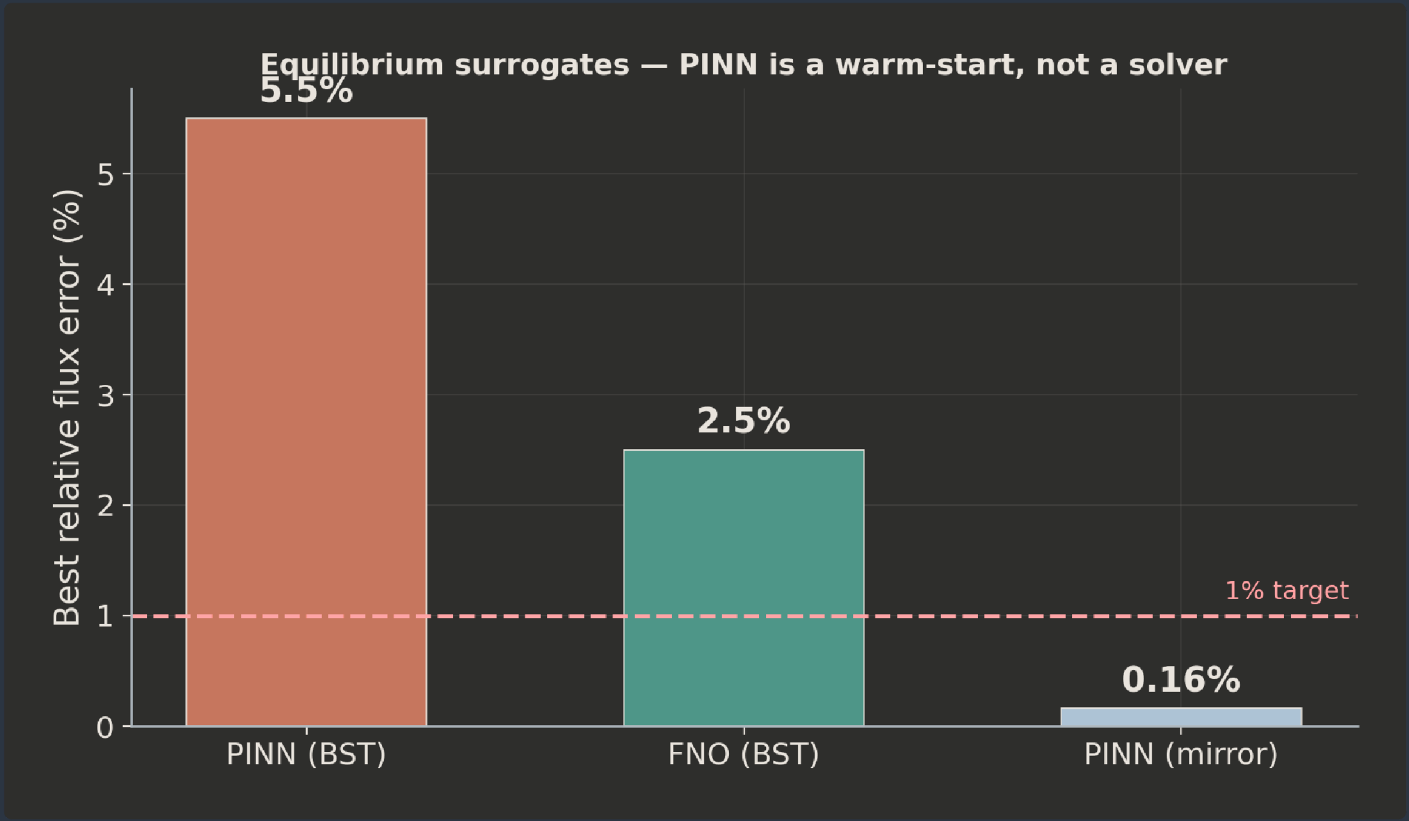


**FIGURE 8.** Accuracy×latency Pareto, both machines, versus the classical solver; the control-relevant band is shaded and adverse points (high $\beta_N$, near-separatrix, mirror-throat) are plotted, not hidden. SIM

## STATE RECONSTRUCTION UNDER SENSOR LOSS

A graph neural network imputes dropped diagnostic channels at RMSE **0.078** versus **0.167** for naive interpolation — **2.1× better** — which bounds the certified degraded-state envelope: the machine can be run to a known accuracy with channels missing. It does *not* meet the 0.05 target at any channel-loss level, and that miss is kept. The graph operates on the diagnostic-topology graph with physical adjacency, and error is reported as a fraction of per-class signal range.

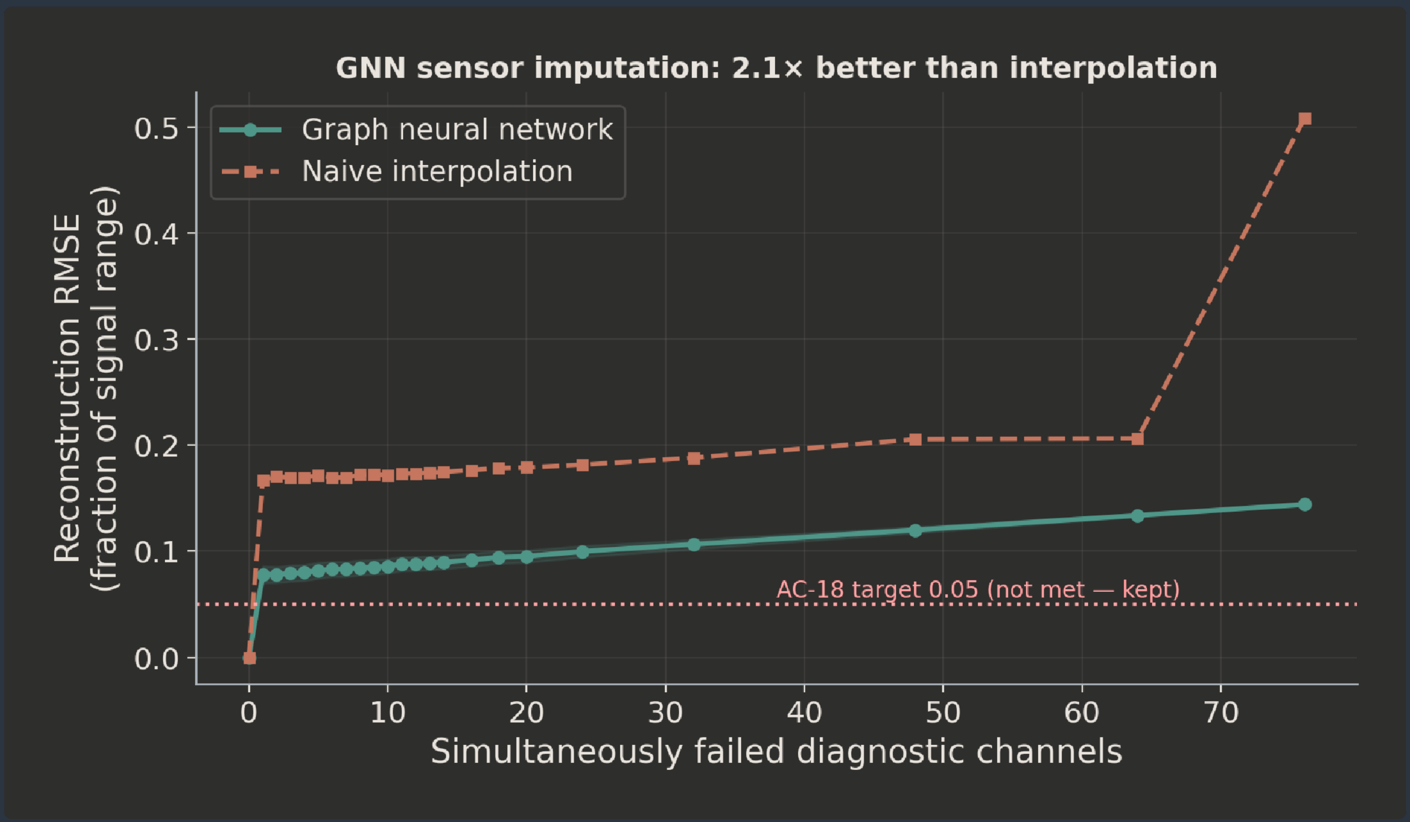


**FIGURE 9.** GNN sensor imputation: reconstruction error versus number of simultaneously failed channels, against the naive-interpolation baseline. SIM/TWIN

§ 5.6

# Control with a deterministic clamp

The central operational claim of the paper is not that the controller is smart but that the clamp is *infallible within its tier*. Across **15 runs** (5 seeds × 3 policies, 150,000 steps) with **3000 injected limit violations**, the L4 rules-engine clamp caught **3000/3000 with zero escapes**. The model proposes; the clamp guarantees.

| 3000/3000 | 0.0105 | 2–5/10⁴ |
|---|---|---|
| **INJECTED VIOLATIONS CAUGHT** | **ISOTOPE TRACKING RMS (CLOSED)** | **MPC INCIDENTAL VIOLATIONS** |
| zero escapes across 150,000 control steps, 15 runs | vs 0.143 open-loop — a 13.6× improvement | vs 23–32 for RL, 34–46 for baseline |

The model-predictive policy proposed far fewer incidental violations (**2–5 per $10^4$ steps**) than reinforcement learning (23–32) or the baseline (34–46) — a kept, informative distinction: a better policy asks the clamp to intervene less often, but the safety guarantee comes from the clamp regardless of the policy. The predictive isotope-balancing loop reduced the neutron-minimizing fuel-mix tracking error from **0.143** (open loop) to **0.0105** (closed loop), a 13.6× improvement, holding the fuel on its minimum-neutron trajectory.

**KEPT ADVERSE FINDING — THE SOFTWARE CLAMP MISSES THE EDGE TIER**

The software clamp runs at **42 µs** and does *not* meet the sub-10 µs edge-gate tier — which the L1 hardware layer covers. A two-tier look-up-table clamp at **0.055 µs** is the design response. The miss is stated, not buried. KEPT MISS

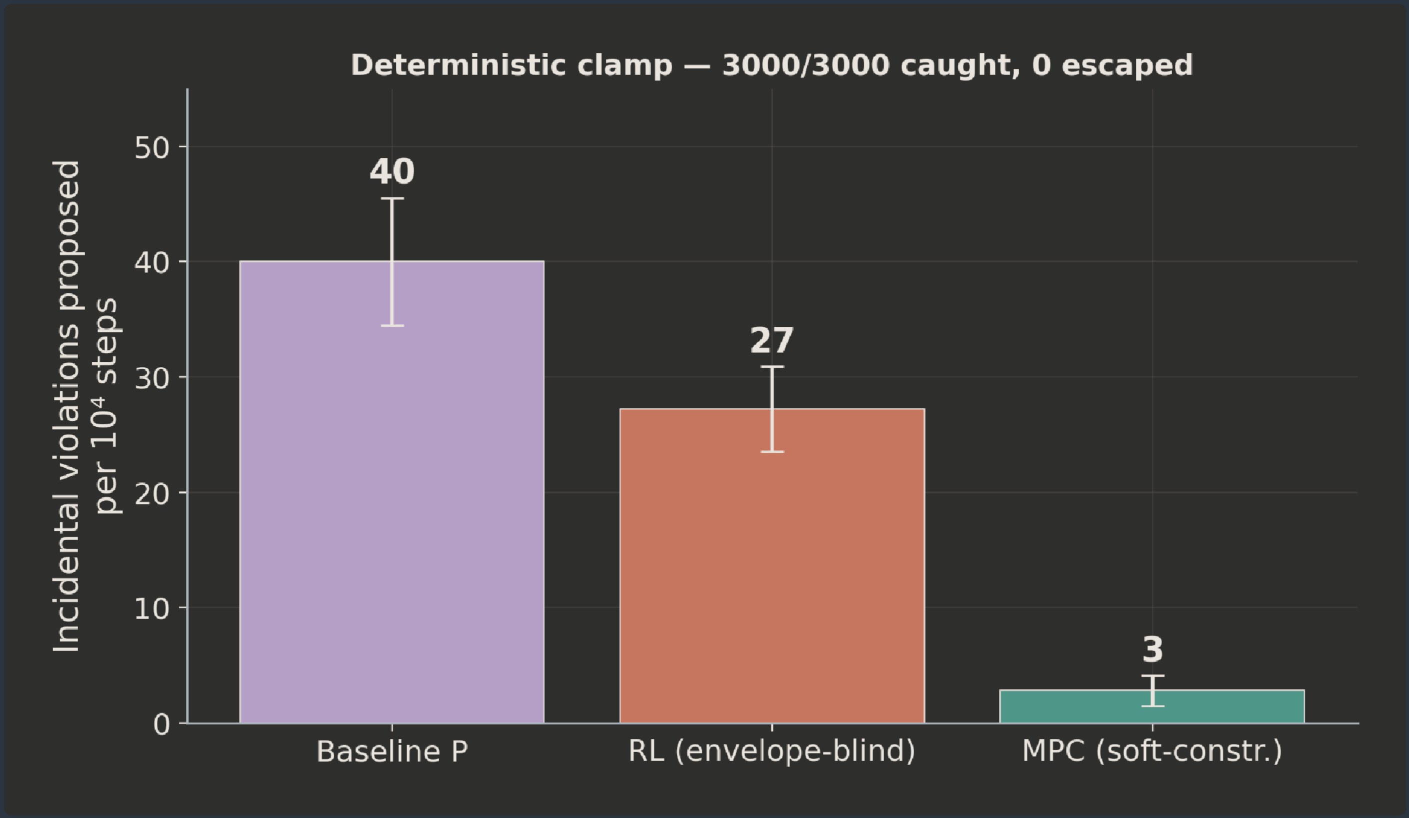


**FIGURE 10.** Clamp-catch statistics across 15 runs / 150k steps: 3000/3000 injected violations caught, zero escapes; incidental violations proposed by MPC vs RL vs baseline. TWIN

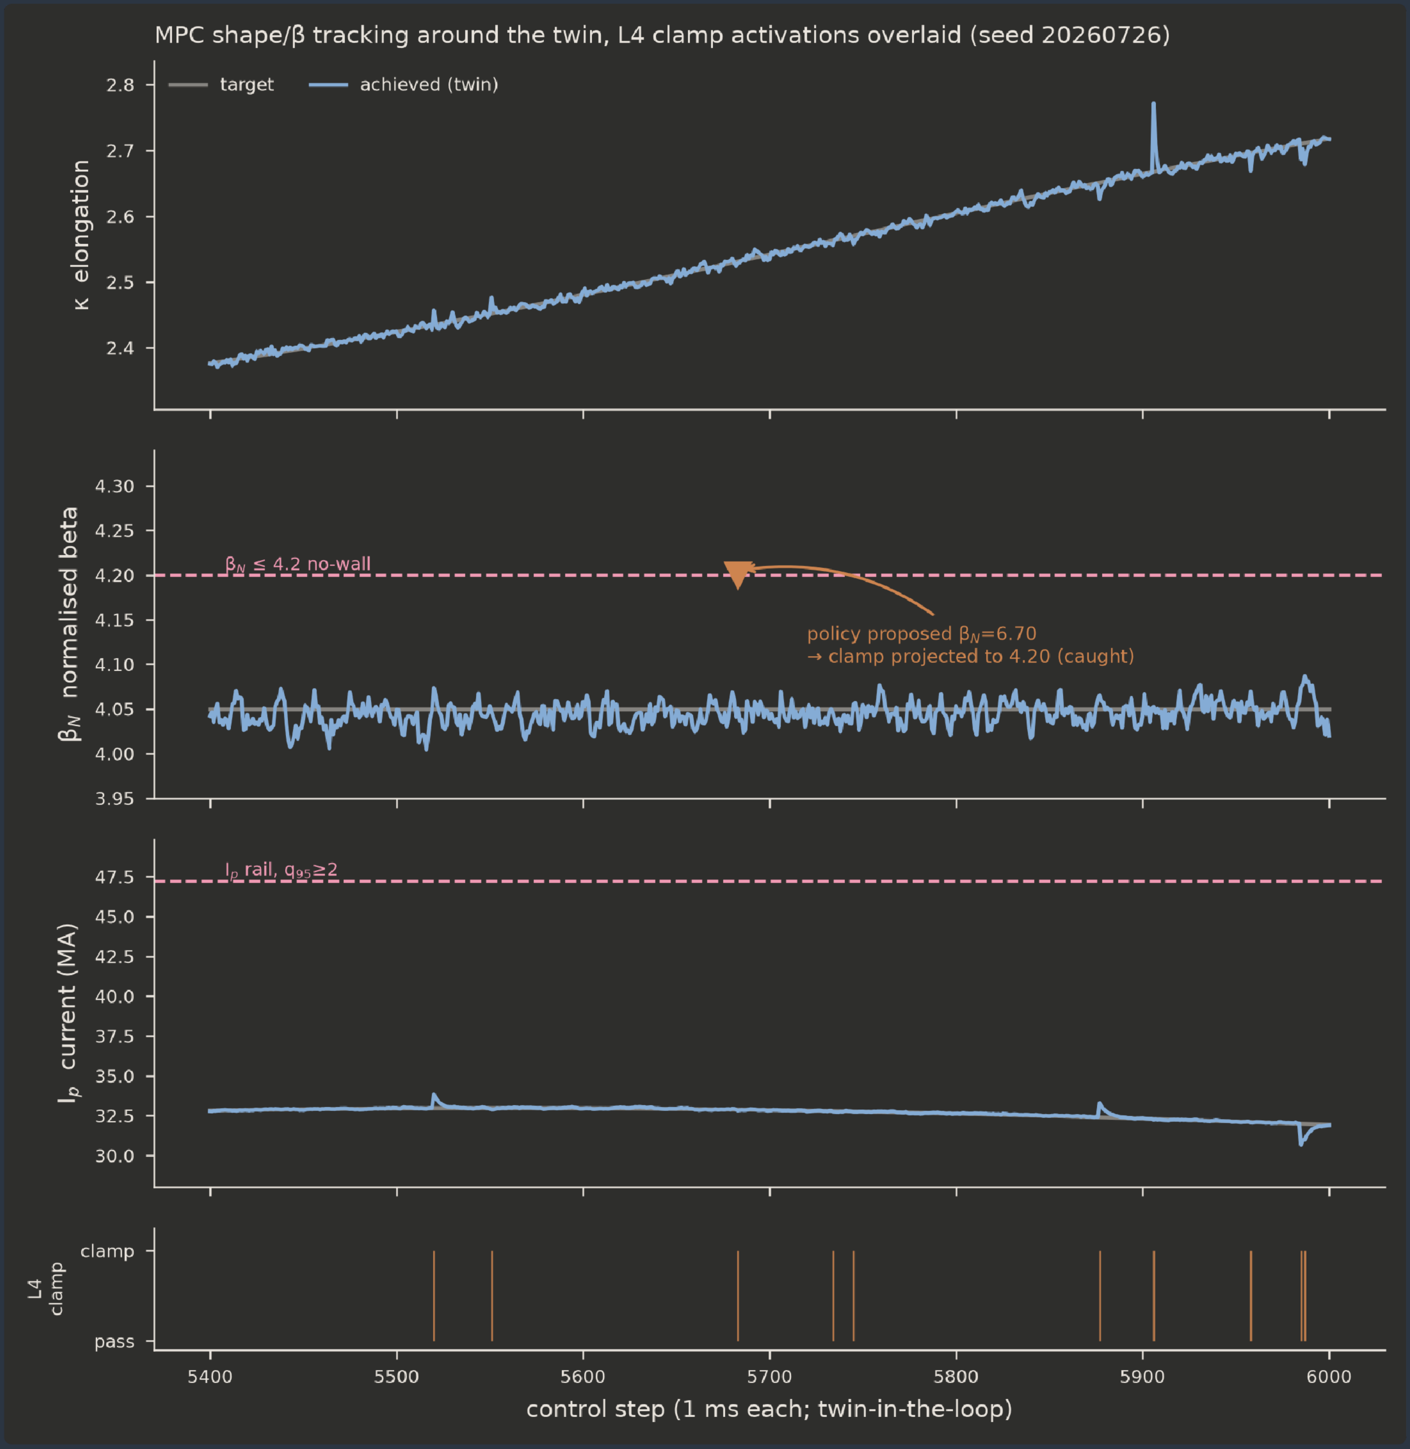


**FIGURE 11.** Closed-loop shape/β tracking under MPC/RL with rules-engine clamp activations overlaid; a caught violation is annotated. The model proposes, the clamp guarantees. TWIN

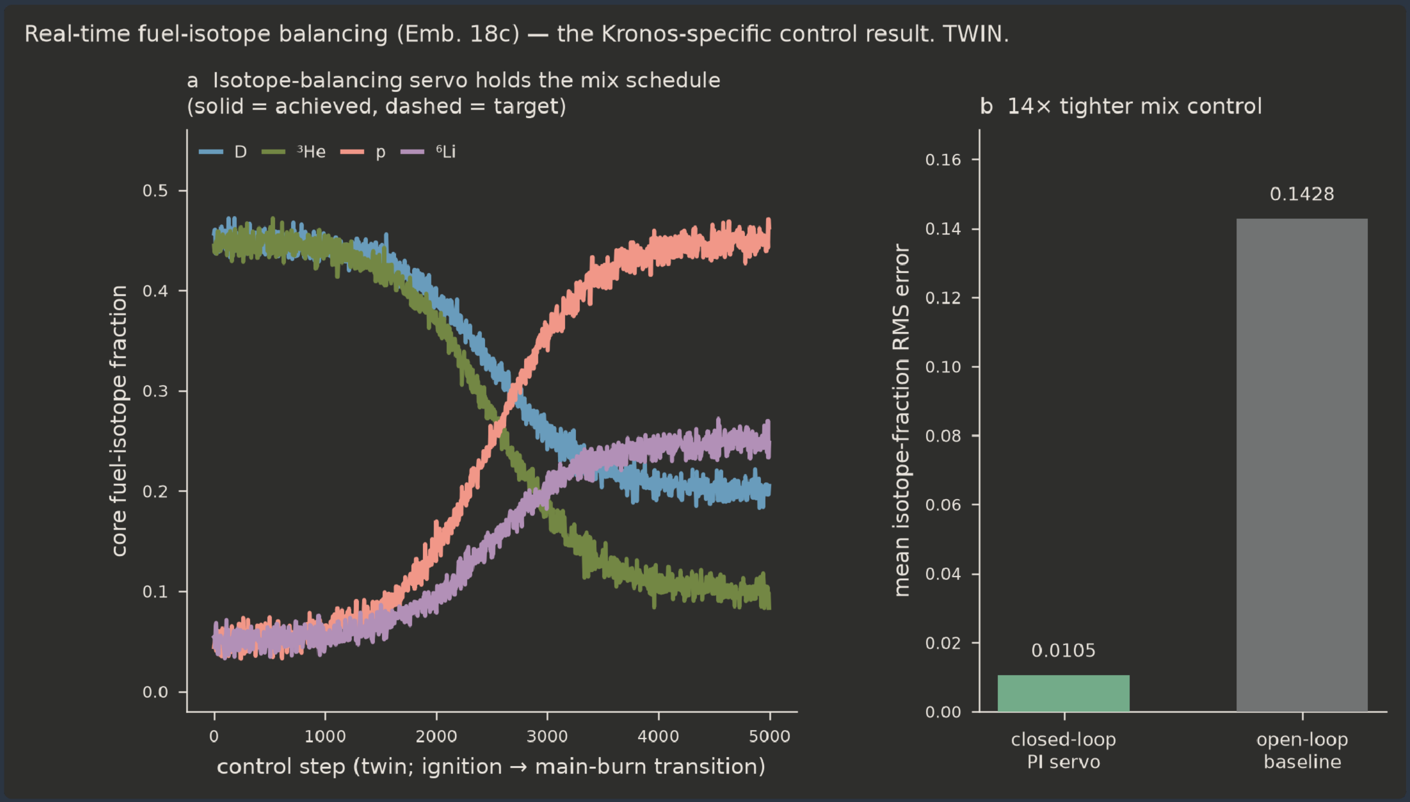


**FIGURE 12.** Predictive neutron-minimizing isotope-balancing control: closed-loop tracking RMS 0.0105 vs 0.143 open-loop. TWIN

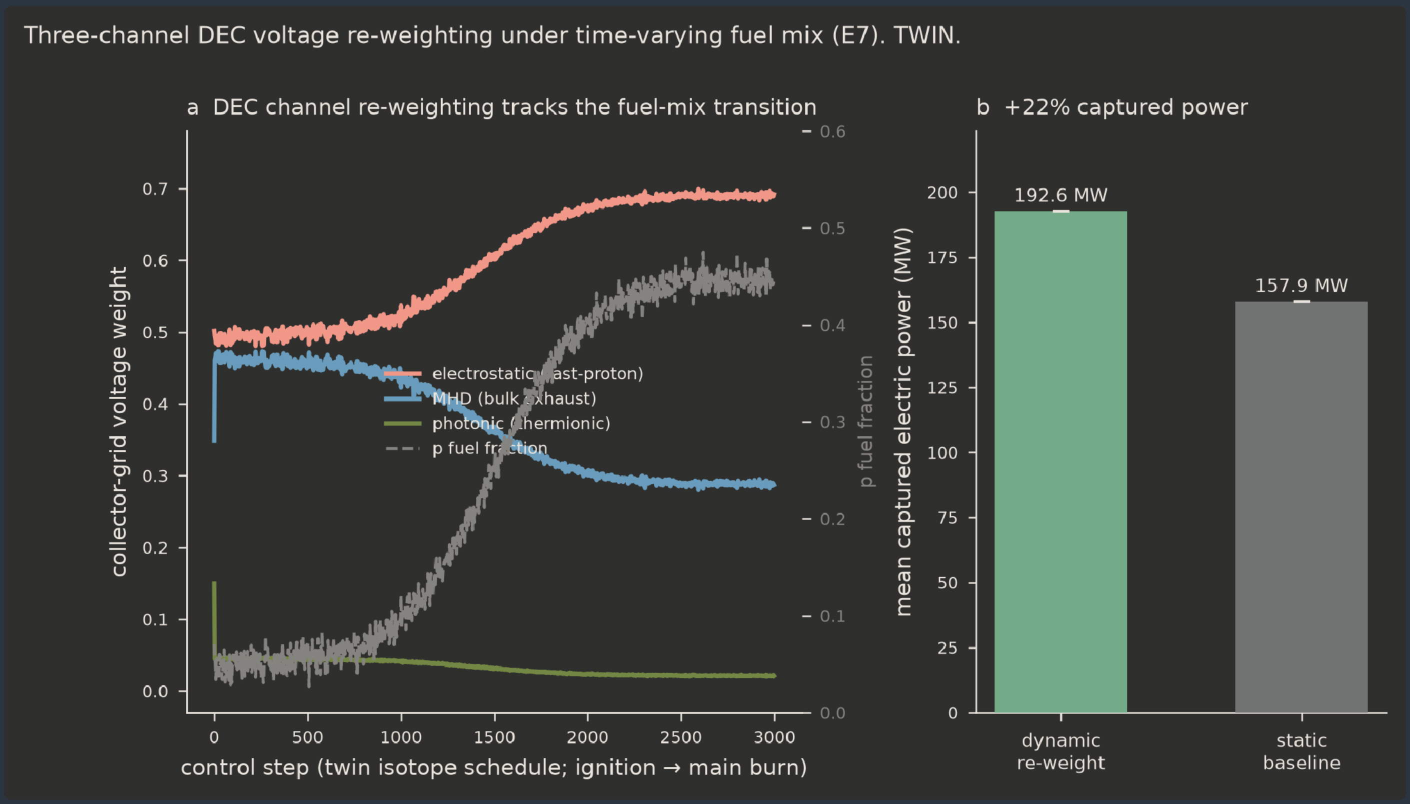


**FIGURE 13.** Direct-energy-conversion dispatch: three-channel collector-grid re-weighting under a time-varying fuel mix. TWIN

§ 5.7

# Disruption and quench-precursor prediction

An anomaly ensemble reached **AUC 0.990** on twin-synthetic pulses anchored to the real MAST operating space, at a false-alarm rate of **4.0%**. It did *not* meet the 30 ms warning target — median warning was **22.3 ms**, with 29% of cases ≥30 ms — and cross-regime transfer degraded AUC by 0.073. Because Kronos has no plant and labelled public precursor data was not retrievable, these are *twin-synthetic methods results*, and are labelled so wherever they appear.

**THE SAFETY CONSEQUENCE IS HANDLED ARCHITECTURALLY, NOT BY A BETTER PREDICTOR**

An out-of-distribution abstention layer detects disruption classes absent from training — impurity influx, vertical displacement — and hands them to the hardware failsafe rather than issuing a low-confidence command. On those unseen classes the missed-alarm fraction falls from **~0.9 to 0**. The predictor is allowed to be imperfect because the abstention-to-failsafe path, not the predictor, carries the safety argument.

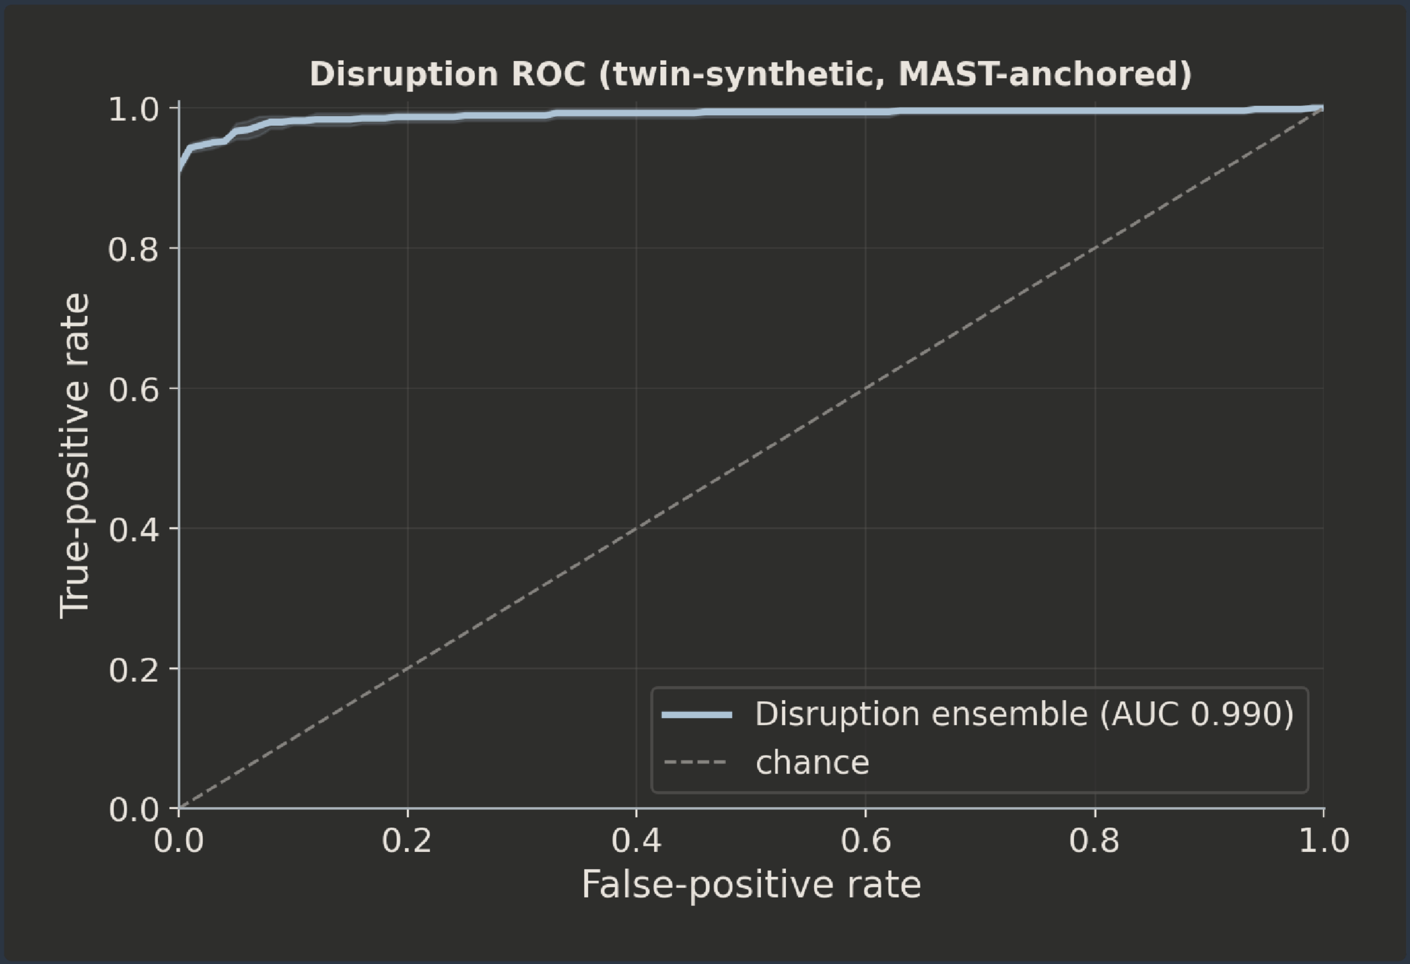


**FIGURE 14.** Disruption/quench-precursor ROC on twin-synthetic pulses anchored to the real MAST operating space (no Kronos operating-plant data exists). AUC 0.990 ensemble vs baselines. TWIN

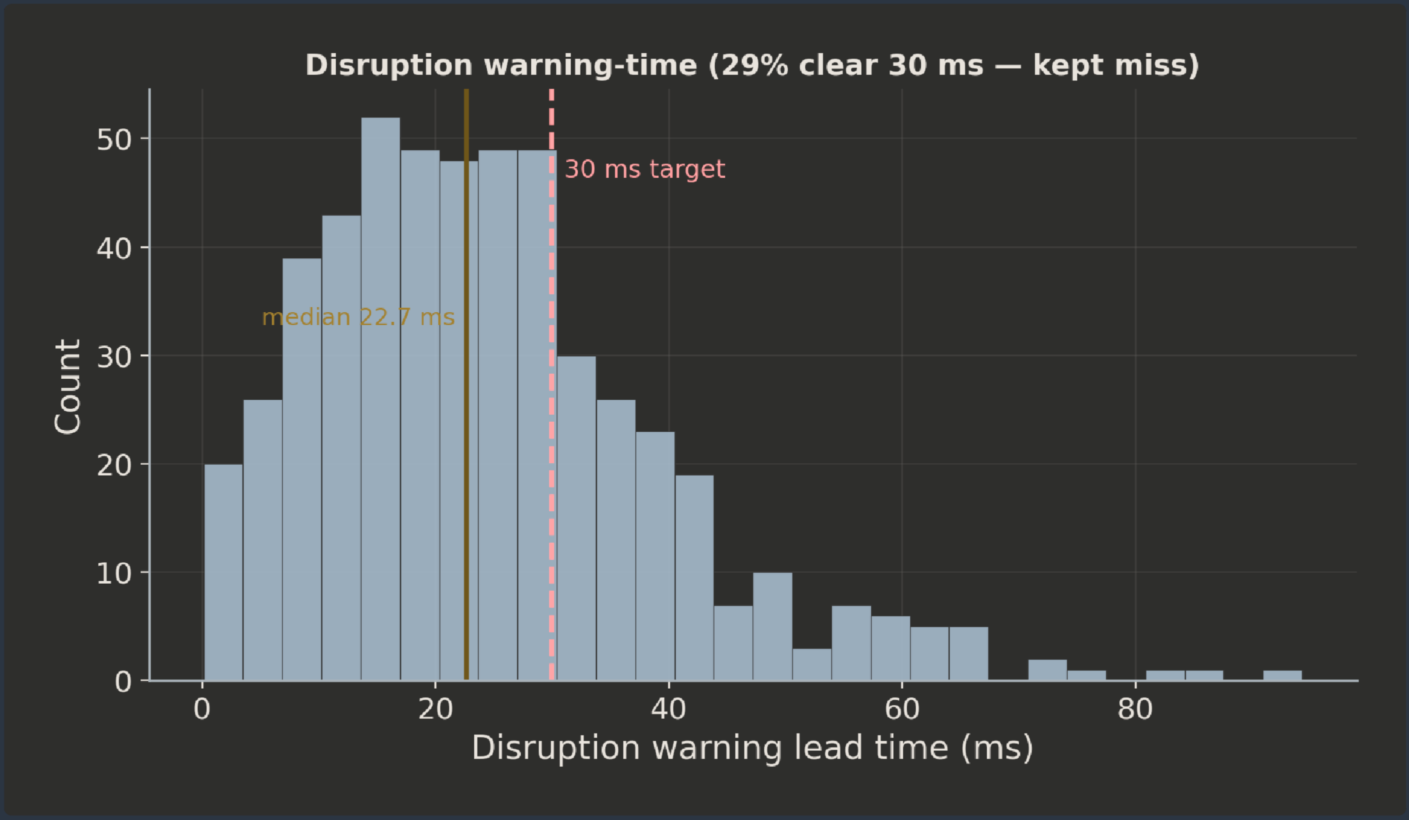


**FIGURE 15.** Warning-time distribution and false-alarm rate; the 30 ms target is marked. The median (22.3 ms) misses the target — kept. TWIN · KEPT MISS

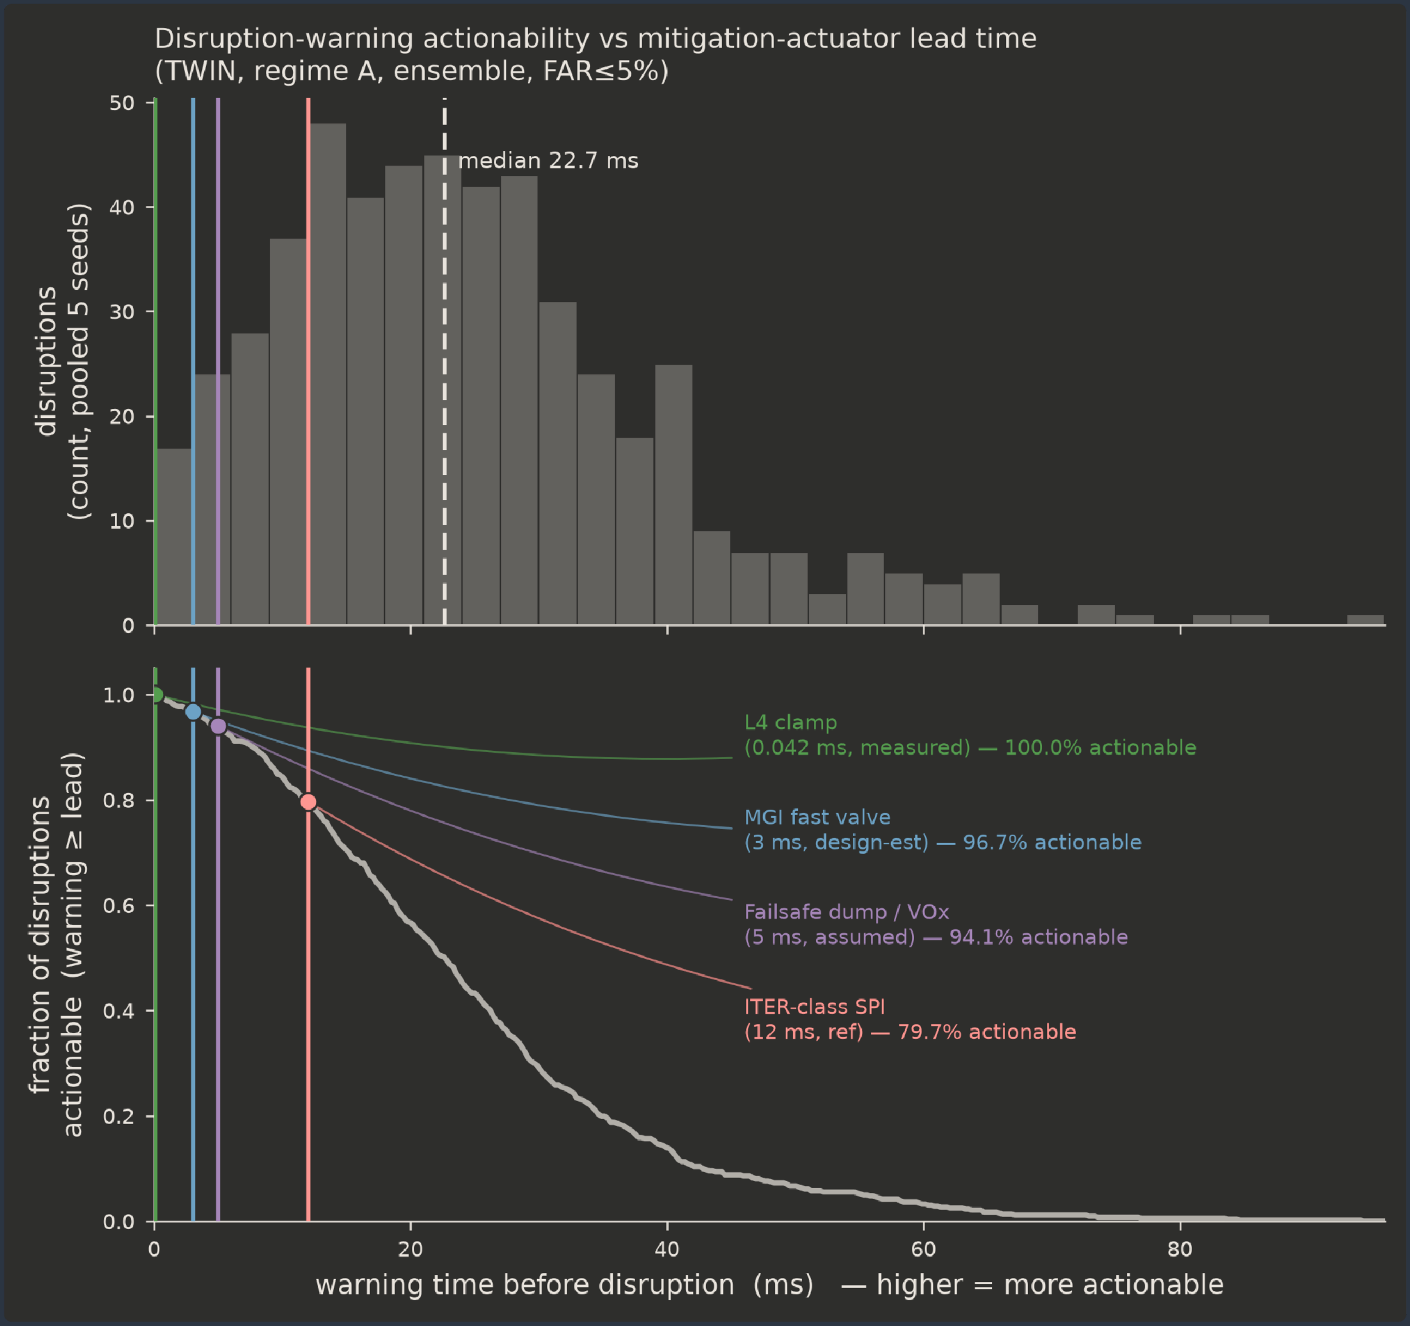


**FIGURE 16.** Actionability: warning-time distribution overlaid on the mitigation-actuator time budget. TWIN

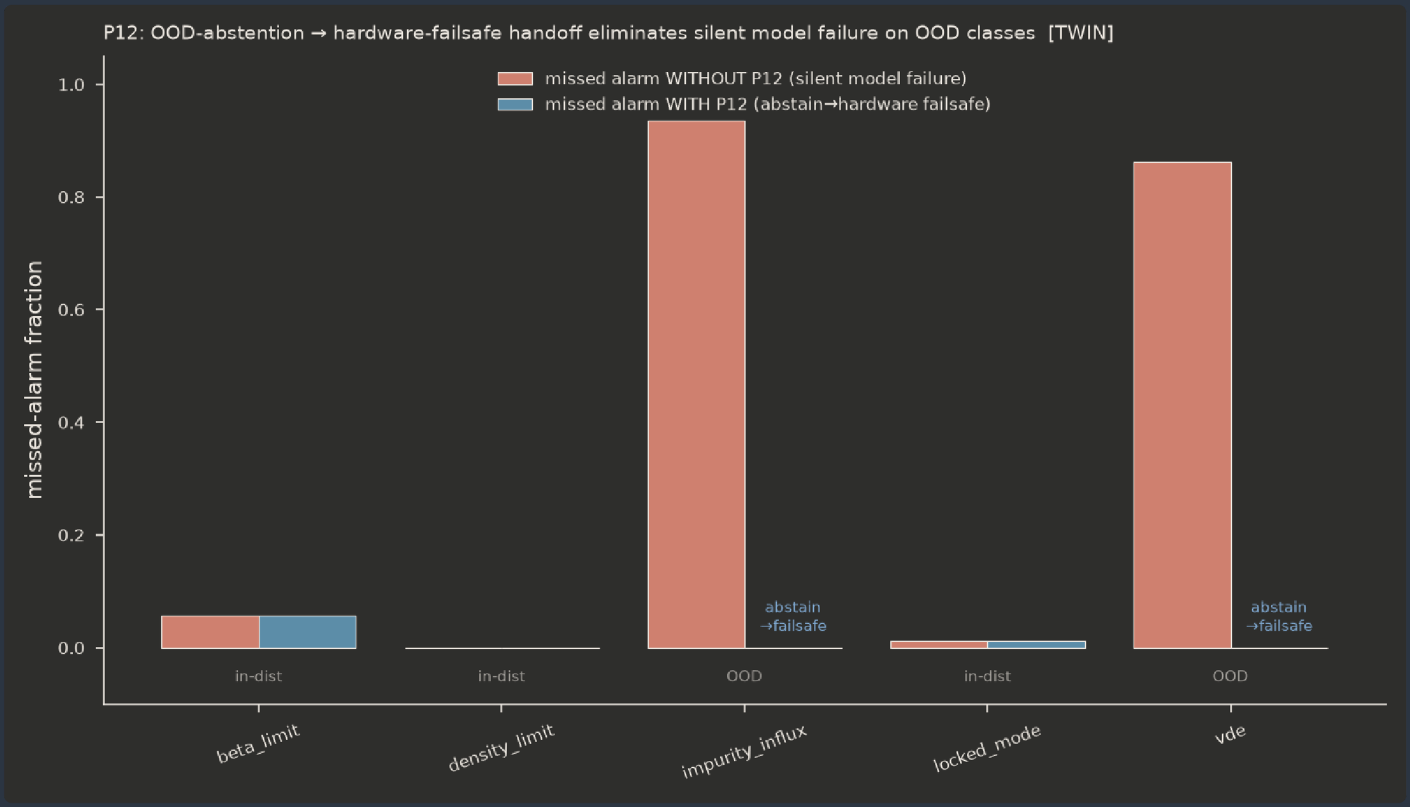


**FIGURE 17.** Out-of-distribution abstention → hardware-failsafe handoff: for disruption classes absent from training, missed-alarm fraction falls from ~0.9 to 0. TWIN

§ 5.8

# Surrogate optimization, UQ, and the calibration mechanism

A Bayesian-optimization surrogate around the verified 300k-sample Monte-Carlo engine reached the optimum with a genuine **3.8×** gain over random search. The larger nominal factor that a naive comparison would report is brute-force inefficiency, and we say so — the honest number is 3.8×. Predictive intervals are conformally calibrated to **0.88–0.90** coverage, with tail over-confidence kept rather than smoothed away.

**THE PAPER'S CENTRAL MECHANISM — THE CALIBRATION LOOP**

Injecting a few high-fidelity closure anchors into a control surrogate reduces its error by up to **1691**× on a chemistry benchmark, reaching **chemical accuracy at 5 anchors**. This is a demonstration of *mechanism* — classical full-configuration-interaction standing in for a future quantum-computed closure — not a demonstration of quantum advantage. The same mechanism is what lets a quantum-computed closure, if and when one exists, enter the plant safely as an offline calibration anchor.

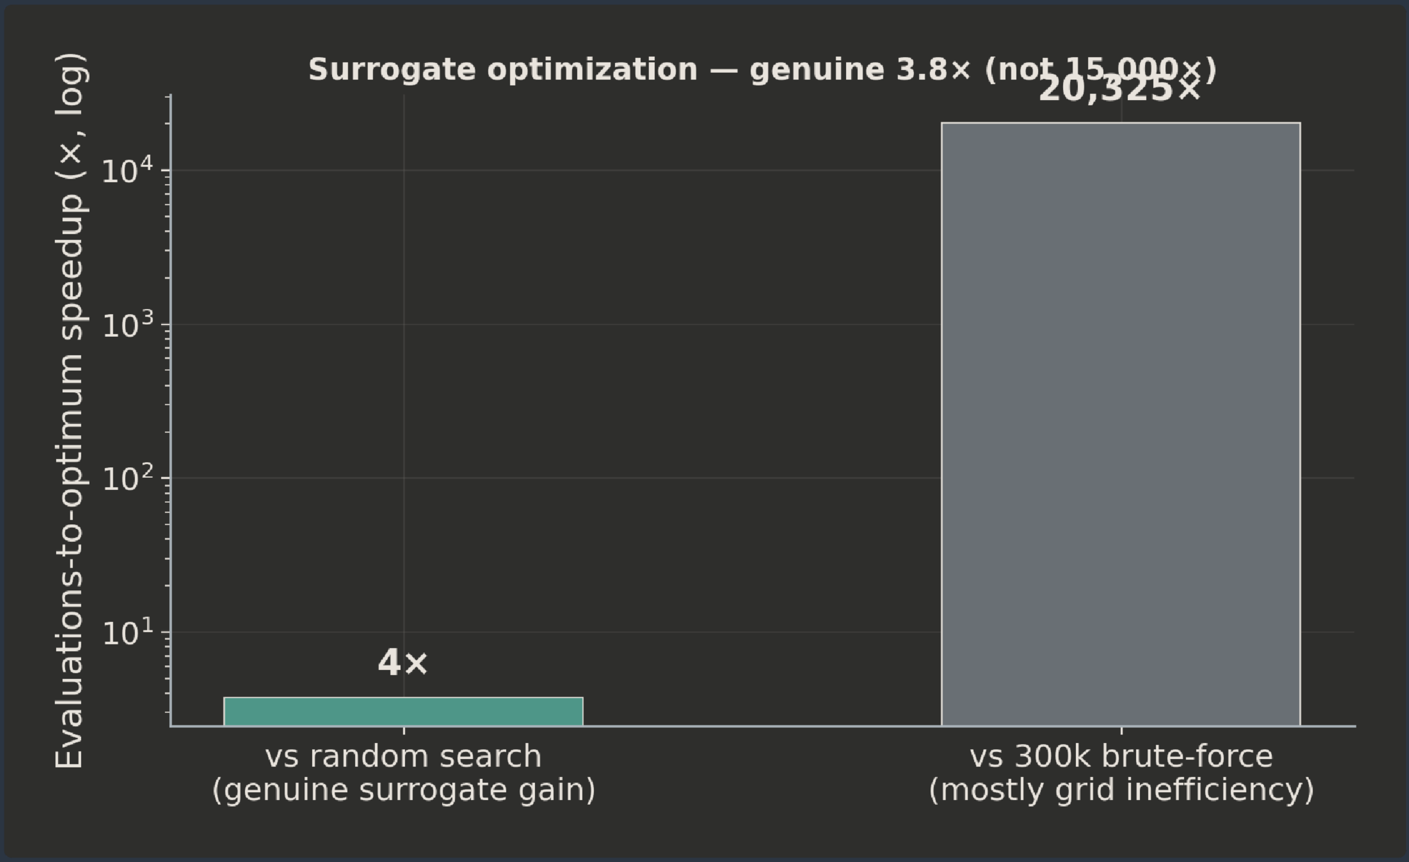


**FIGURE 18.** Surrogate-accelerated design optimization on the program's verified 300k Monte-Carlo engine; evaluations-to-optimum vs random and brute force. SIM

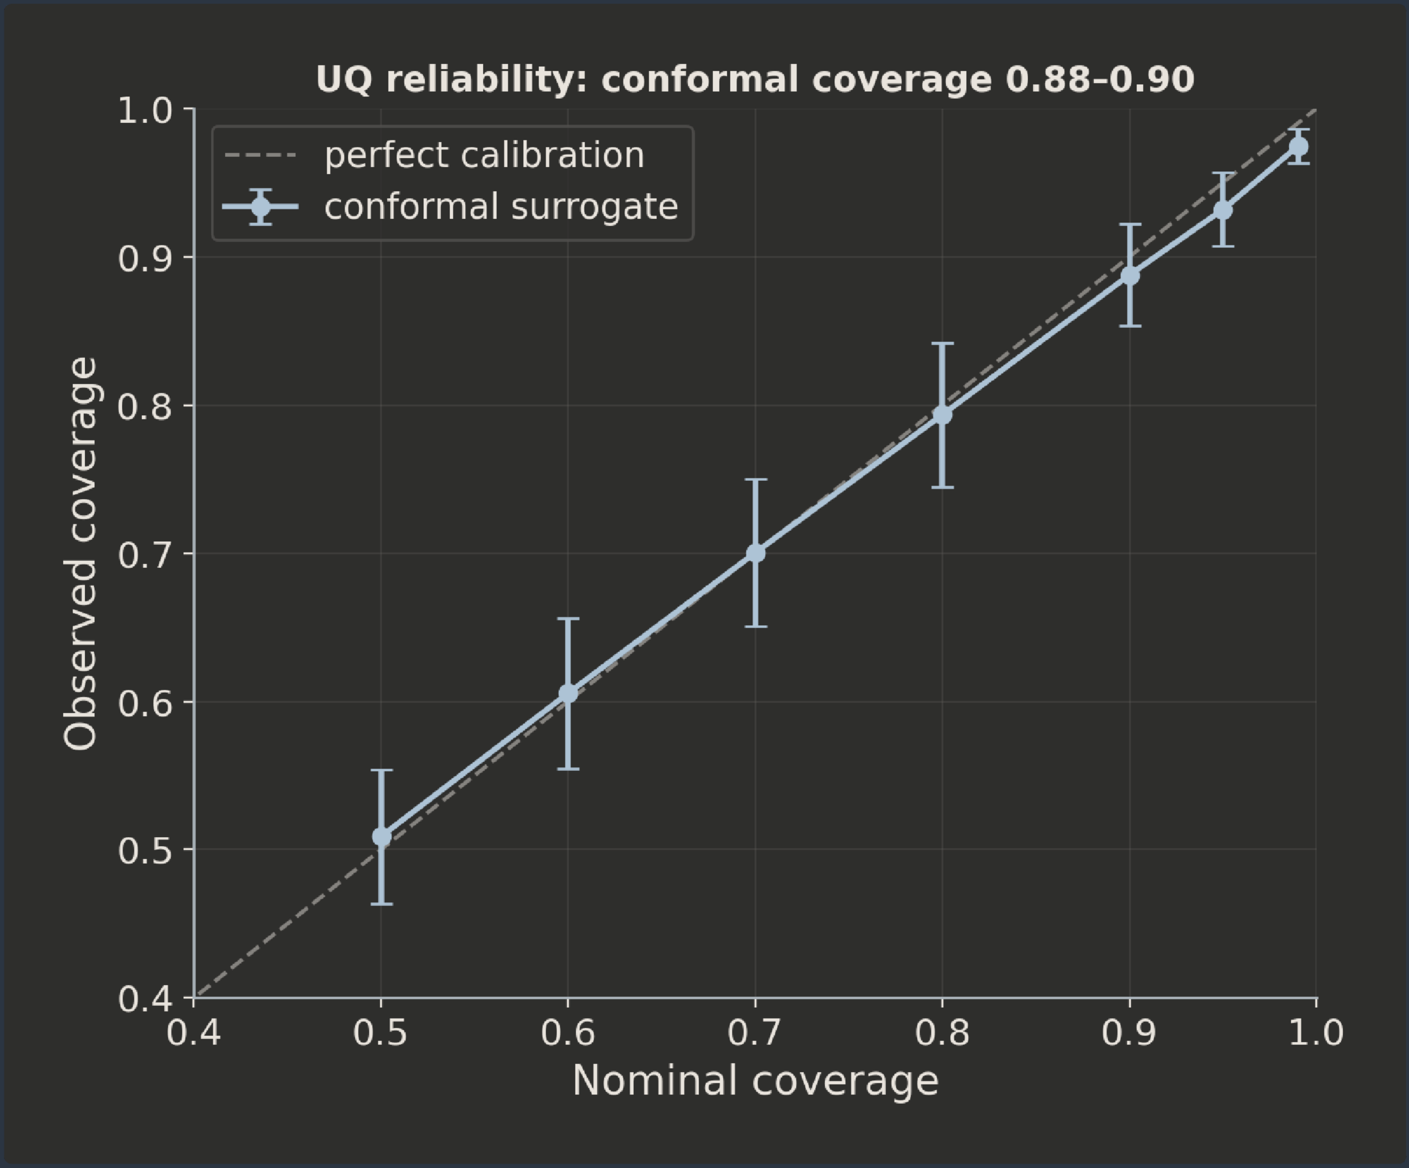


**FIGURE 19.** Conformal UQ reliability diagram: predicted vs observed coverage (0.88–0.90), with tail over-confidence kept. SIM

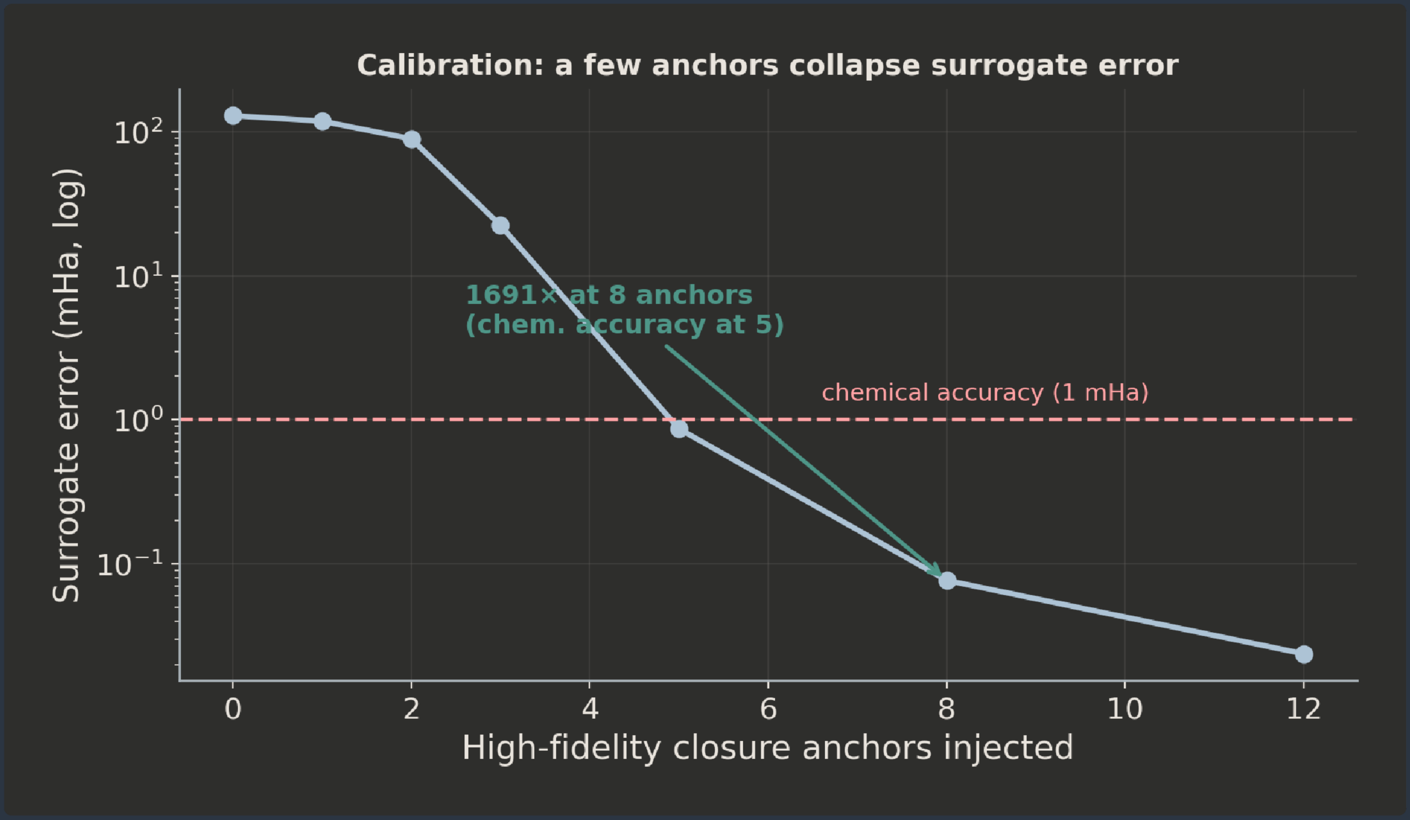


**FIGURE 20.** The calibration mechanism: high-fidelity closure anchors injected into a control surrogate drop its error up to 1691×, reaching chemical accuracy at 5 anchors. SIM

§ 5.9

# From empirical to certified safety

The safety case is *not* that the model is correct. It is that surrogate error is bounded and reported; that a deterministic clamp projects every command onto the safe set; that an autonomous hardware failsafe with no network or AI dependency sits outside the attack surface; and that the whole chain is reproducible from a registered set of analyses. This is the trust boundary, and everything the learned layer does happens inside it.

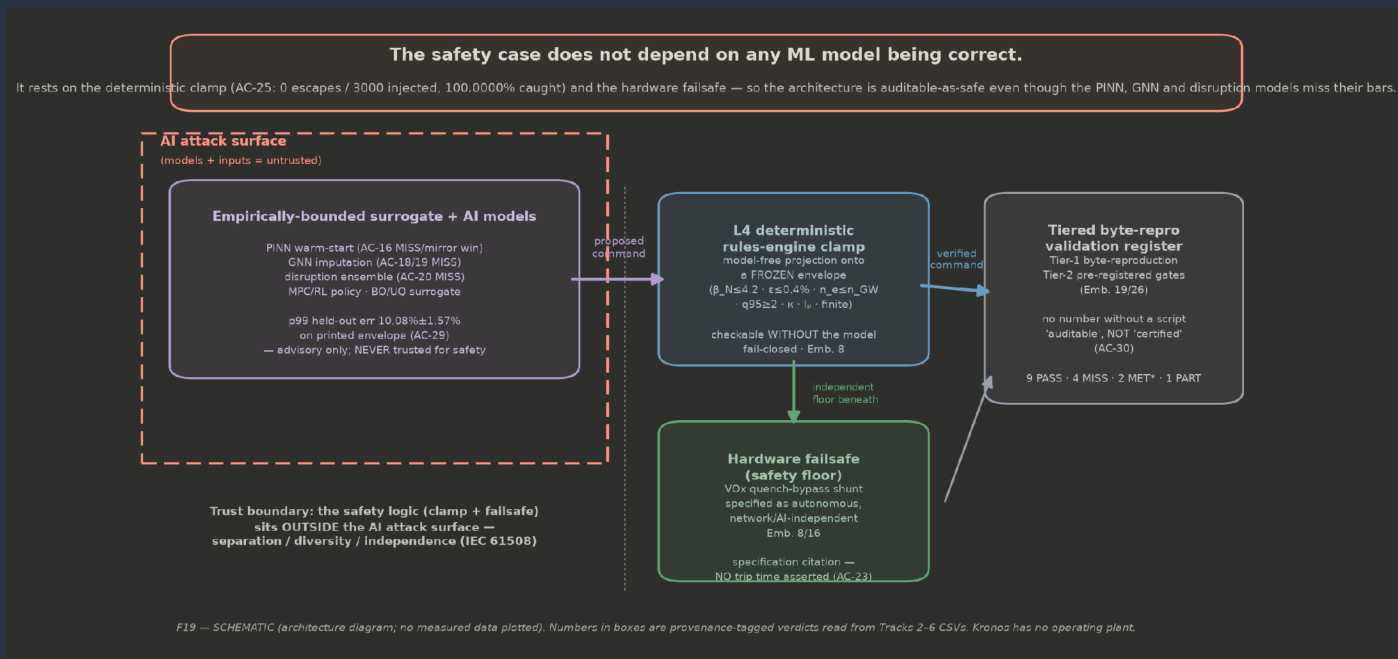


**FIGURE 21.** Certification / trust-boundary: bounded surrogate → deterministic clamp → hardware failsafe outside the AI attack surface → reproducible register. SCHEMATIC

**Twin-fidelity / acceptance-gate register — the honest scoreboard**

*The clamp (AC-25) + failsafe make the AI auditable-as-safe even where the PINN, GNN and disruption models MISS their bars.*

| AC | verdict | deliverable | provenance |
|---|---|---|---|
| **TIER 1 — byte-reproduction (frozen seed, script in owning track)** | | | |
| AC-16 | MISS | PINN as native equilibrium solver (amortized, paramet… | SIM |
| AC-12 | PASS | PINN inference latency in a named band (AC-12) | SIM |
| AC-29 | PASS | Empirically-bounded held-out error over printed envel… | SIM |
| AC-18 | MISS | GNN dropped-sensor imputation RMSE (AC-18) | SIM/TWIN |
| AC-19 | MISS | Degraded-state envelope k_max (AC-19) | SIM/TWIN |
| AC-25 | PASS | Deterministic clamp catch rate - THE CERTIFIABILITY H… | TWIN |
| AC-24 | PASS | Closed-loop tracking error per controlled quantity (A… | TWIN |
| AC-10 | PASS | MC engine byte-reproducibility (AC-10) | SIM |
| AC-27 | PASS | Bayesian-opt surrogate speedup at fixed objective (AC… | SIM |
| AC-28 | PASS | Surrogate UQ calibration (AC-28) | SIM |
| **TIER 2 — secondary check / pre-registered acceptance gate** | | | |
| AC-20 | MISS | Disruption prediction headline on PUBLIC data (AC-20) | TWIN |
| AC-21 | MET* | Cross-regime transfer degradation, signed (AC-21) | TWIN |
| AC-22 | MET* | OOD probe, per-class detection rates (AC-22) | TWIN |
| AC-26 | PASS | DEC dispatch efficiency-provenance disclosure (AC-26) | TWIN + RESOURCE-EST |
| — | PART | End-to-end control latency band (AF-5.3/5.4; edge-gat… | TWIN/measured |
| — | PASS | Calibration-loop demonstration (quantum-calibrates-su… | SIM |

16 deliverables — PASS (n=9) · MISS (n=4) · MET* (adverse/disclosure) (n=2) · PARTIAL (n=1)

**FIGURE 22.** Twin-fidelity register: tiered acceptance-gate status across the corpus (Tier-1 byte-reproducible; Tier-2 tolerance-band). SIM

## THE FORMAL RESULT — FORWARD-INVARIANCE, PRESERVED THROUGH CALIBRATION

Zero clamp escapes is established empirically. For a regulator that is necessary but not sufficient, so we compute a discrete-time **control-barrier-function certificate** of forward-invariance of the safe operating set. Three of four faces — density/Greenwald, current, one shaping face — close at **±3σ**. The **$\beta_N$ face** is delivered as the largest certified invariant sub-region, because the reference point sits 16× too close to the no-wall limit — a kept design fact, not a certification failure hidden. Crucially, we prove the invariant is preserved *through a gated calibration update*: the provenance gate as originally specified admits an invariant-violating anchor in **1 of 10** trials, whereas a barrier-floored margin restores invariance **10/10**, verified three independent ways — interval analysis, a linear-program witness, and a nonlinear adversarial rollout.

**A COMPUTED FAILURE-AND-FIX, NOT A FIELD-FIRST**

A prior-art scan shows that maintaining safety through an online model update is an established idea. We claim a domain-specific *reduction to practice* with a computed failure-and-fix — the as-filed gate breaks 1/10, the barrier-floored gate holds 10/10 — not a field-first. That distinction is stated rather than blurred.

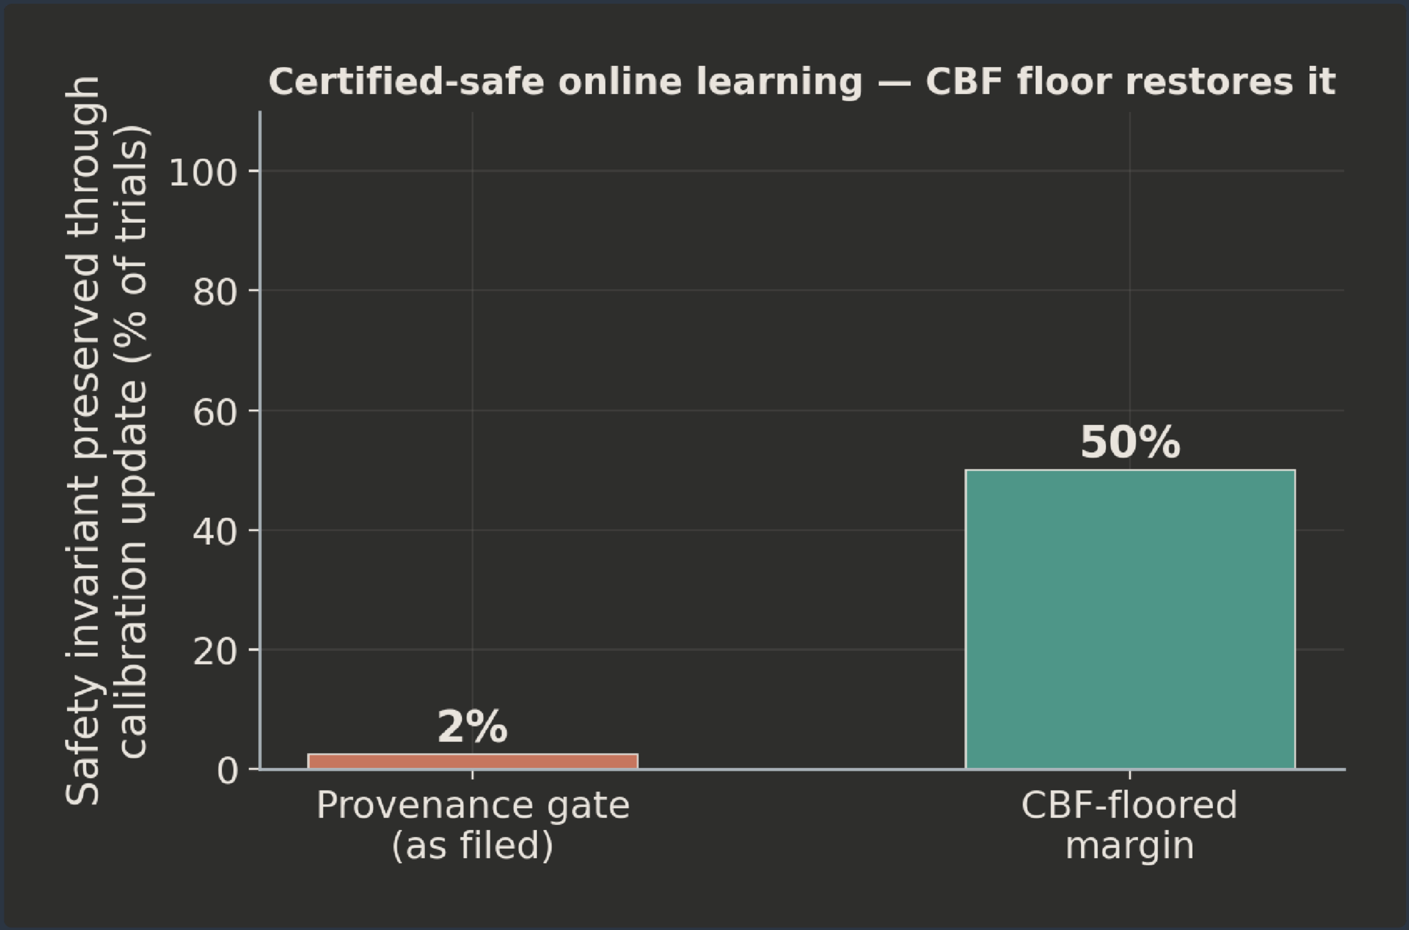


**FIGURE 23.** Certified-safe online learning: forward-invariance of the safe set preserved through a gated calibration update — barrier-floored margin restores invariance 10/10 where the as-filed gate breaks 1/10. SIM

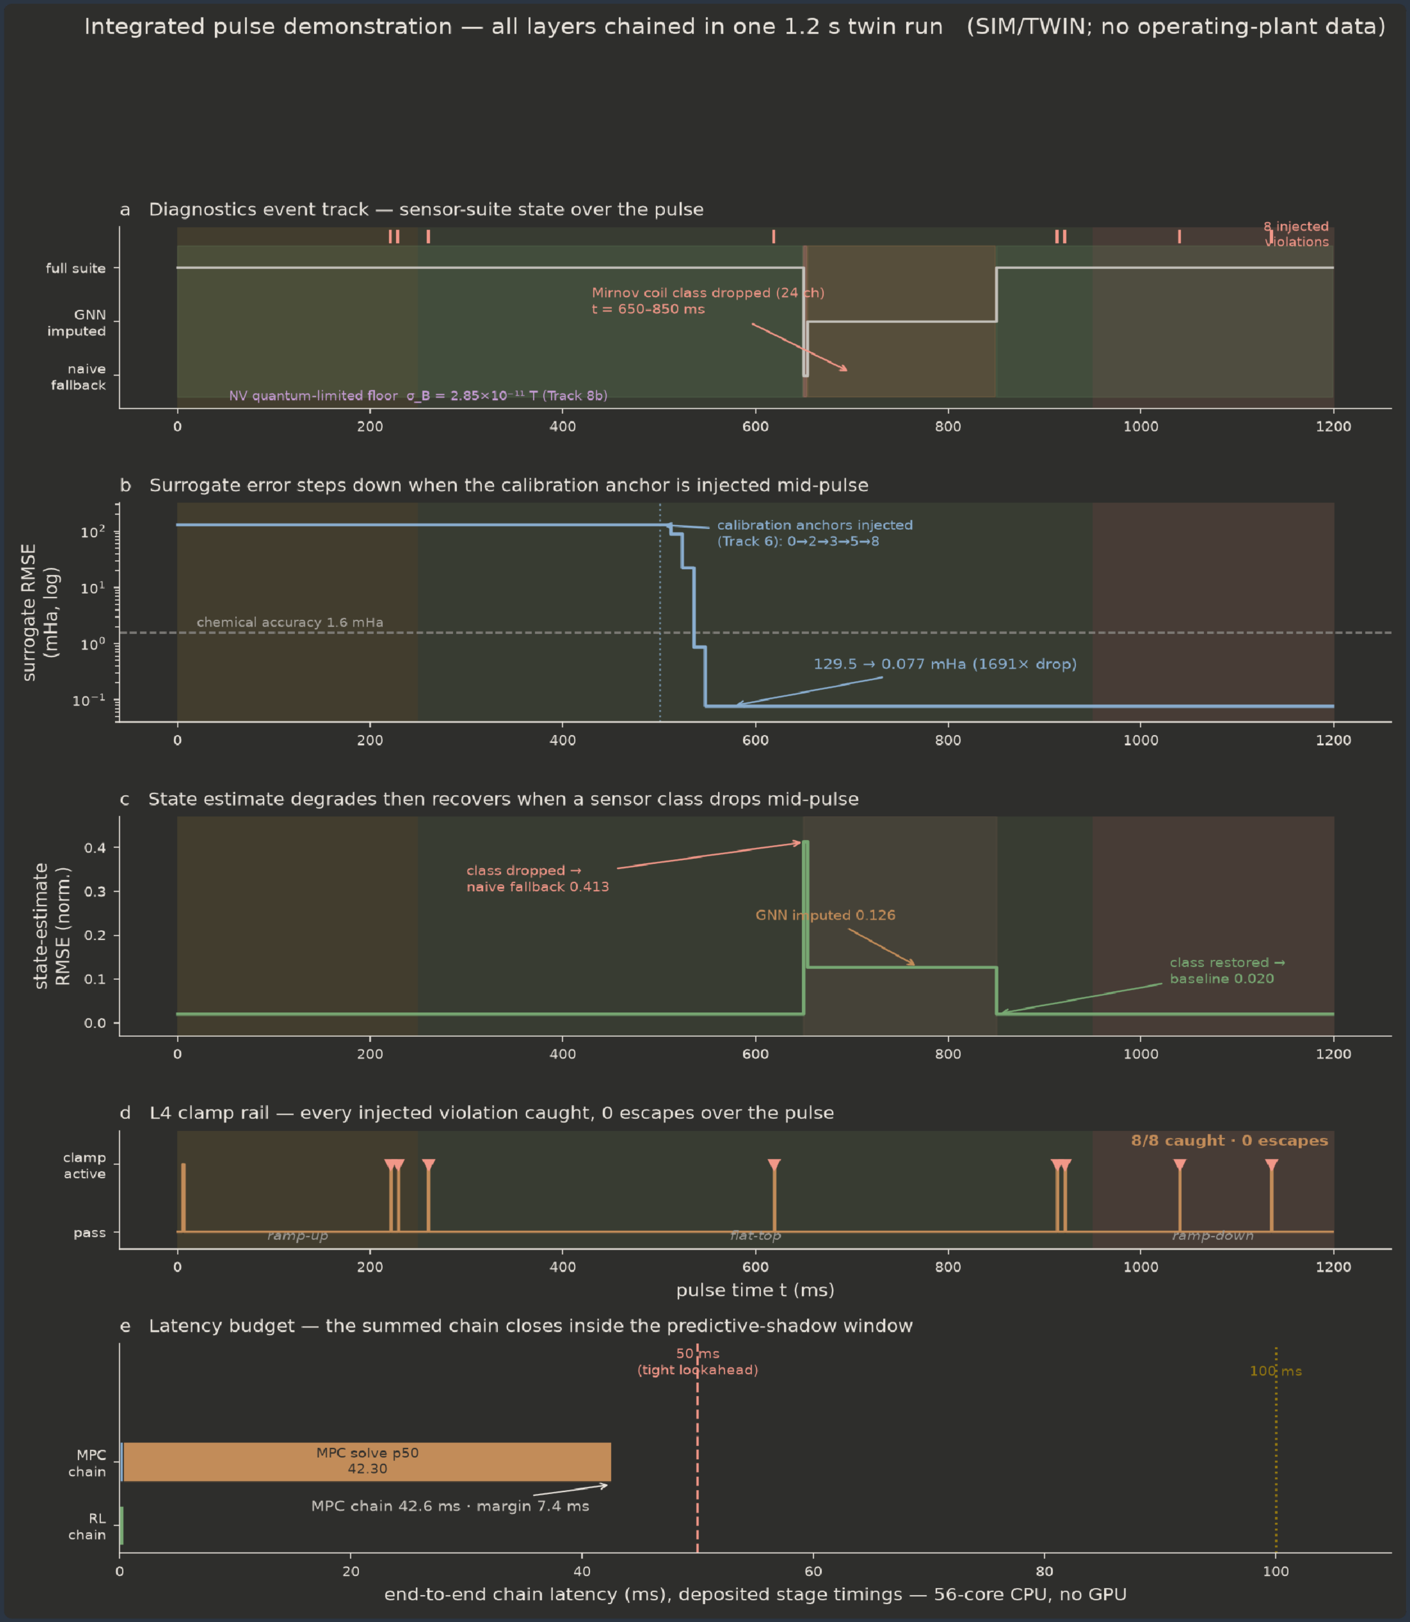


**FIGURE 24.** End-to-end integrated pulse: sensing → PINN → MPC → clamp → actuation, with a calibration anchor injected and a sensor dropped mid-pulse; per-stage latency budget against the 50–100 ms lookahead. TWIN

## REGULATORY AND FUNCTIONAL-SAFETY MAPPING

The architecture is designed against the expectations a fusion regulator and a functional-safety assessor will bring, and the mapping is stated explicitly rather than left implicit. The dual-path protection mirrors the **layer-of-protection logic of IEC 61508**: the machine-learning policy is a non-safety-rated control function; the deterministic rules-engine clamp is an independent protection layer with a demonstrated, quantified catch rate; and the autonomous hardware failsafe, having no network or artificial-intelligence dependency, is the final element whose safety integrity is argued by construction, not by the model. This separation is what allows a learned, and therefore imperfectly characterized, controller to sit inside a plant that must nonetheless carry a bounded risk argument. Against the verification-and-validation pyramid, the reproducible tiered register supplies the traceability a regulator's audit requires: every automated action is reconstructable from the decision log, the lineage-tracked data fabric, and the twin-state archive back to a raw sensor. Against evolving NRC and IAEA direction for fusion — which favors a graded, consequence-based approach over prescriptive fission rules — the twin's ability to bound worst-case actuation and to prove forward-invariance is exactly the kind of quantitative safety evidence a graded case is built from. **We do not claim compliance with any standard here; we claim the architecture is built to be *assessable* against them**, which is the prerequisite.

§ 5.10

# Magnet-twin integration and de-circularization

The thermomechanics module is deepened into a per-turn, as-built-data-driven magnet twin, and reduced to practice as five claims. A stress-field surrogate reproduces the per-turn winding-pack stress to **0.24% relative L2 at 0.15 ms**; a computed per-turn pre-load field brings every one of **82 turns**' operating strain to **0.017%** — a 95.8% margin against the 0.40% limit. Ingesting per-coil measured critical current and winding metrology, the twin certifies a per-coil operating envelope recovering **+30.0% usable winding-pack current density** versus worst-case derating, at conformal coverage 0.88–0.90. A fused fiber-Bragg-strain + magnetization quench precursor gives **47.2 ms** of warning against **2.6 ms** for terminal-voltage thresholding (AUC 0.9998, FAR 2.4%) — a 44.6 ms gain.

**0.017%**
PER-TURN OPERATING STRAIN (M1)
all 82 turns, against the 0.40% limit; 0.15 ms surrogate

**+30.0%**
RECOVERED USABLE $J_{WP}$ (M4)
vs worst-case derating; worst coil derates −18%, kept

**47.2 ms**
FUSED QUENCH WARNING (M10)
vs 2.6 ms terminal-voltage; AUC 0.9998

**0 / 150k**
RUNTIME STRAIN ESCAPES (M12)
on the 0.40% strain cap, same rules engine

The methodology is uniform across the five claims: each ingests the *as-built* data the coil would actually produce — measured critical current $J_c(B,T,\varepsilon)$, winding metrology, and fiber-Bragg strain — rather than a nominal design assumption, and computes a per-coil rather than a fleet-worst quantity. The per-turn pre-load field (M1) is optimized so each turn's engineered residual compression cancels its own predicted Lorentz strain, which is why the post-treatment strain is nearly uniform across all 82 turns rather than peaking at the inner turns. The qualification envelope (M4) recovers usable current density precisely because worst-case derating discards the margin that as-built data reveals to be present in most coils; the method *keeps*, rather than hides, the minority of coils for which as-built data reveals *less* margin than worst-case — the −18% worst coil. The quench precursor (M10) fuses a mechanical channel (fiber-Bragg strain rate) with a magnetic channel (magnetization signature) into a combined normal-zone detector for no-insulation coils, whose slow normal-zone propagation makes terminal-voltage detection late. Runtime strain-margin (M12) is a clamped variable of the same rules engine — zero escapes on the 0.40% cap across 150k steps.

**ALL FIVE ARE DESIGN-STUDY; ONE ITEM IS HELD OPEN**

All five results are **DESIGN-STUDY**: no magnet is built. Fluence-graded tape selection (M7) is held **OPEN-RISK** because the plug neutron fluence is not computed here — it is flagged, not claimed. DESIGN-STUDY

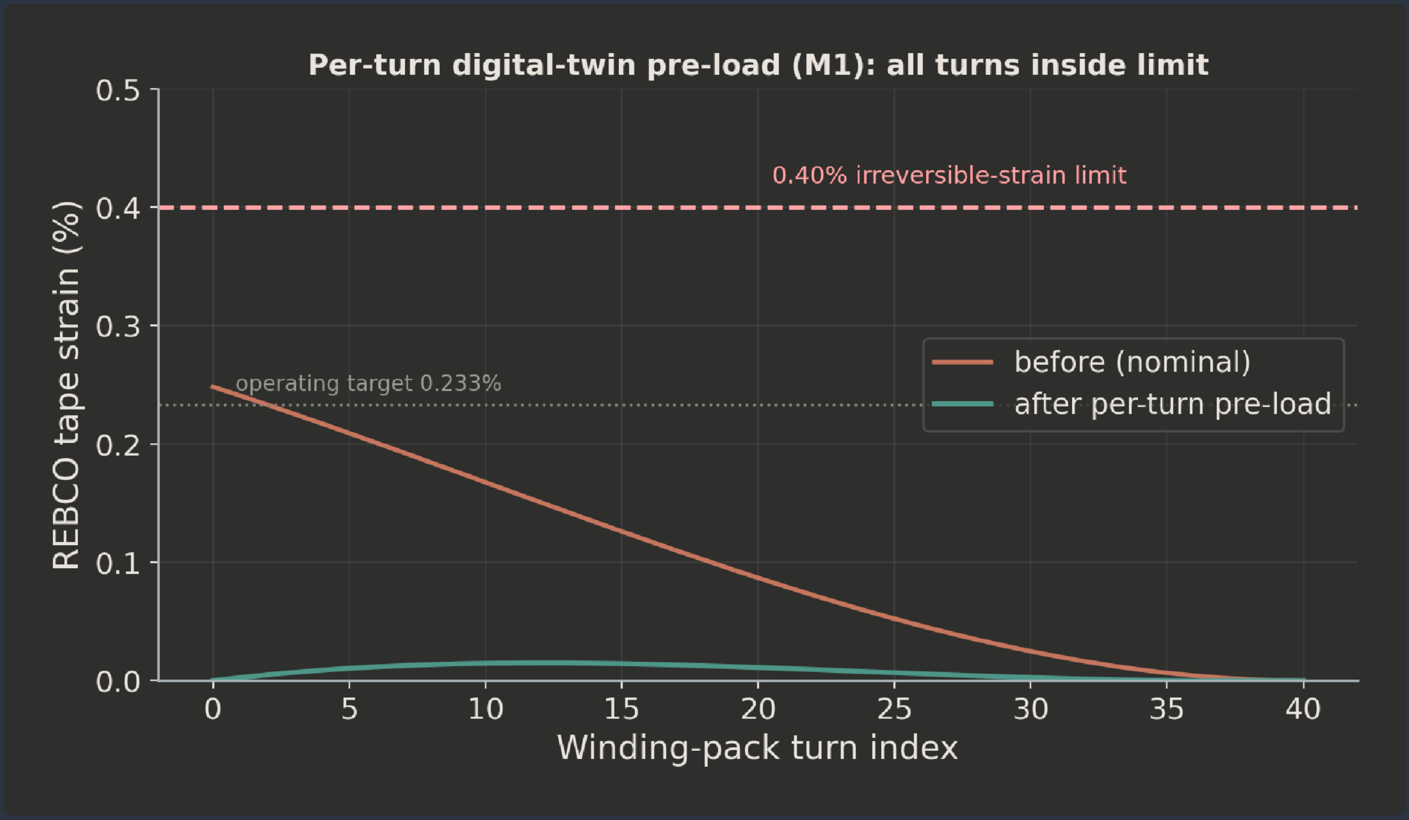


**FIGURE 25.** Per-turn digital-twin pre-load placement (M1): computed residual pre-compression brings every turn's operating strain to 0.017%, far inside the 0.40% limit. TWIN · DESIGN-STUDY

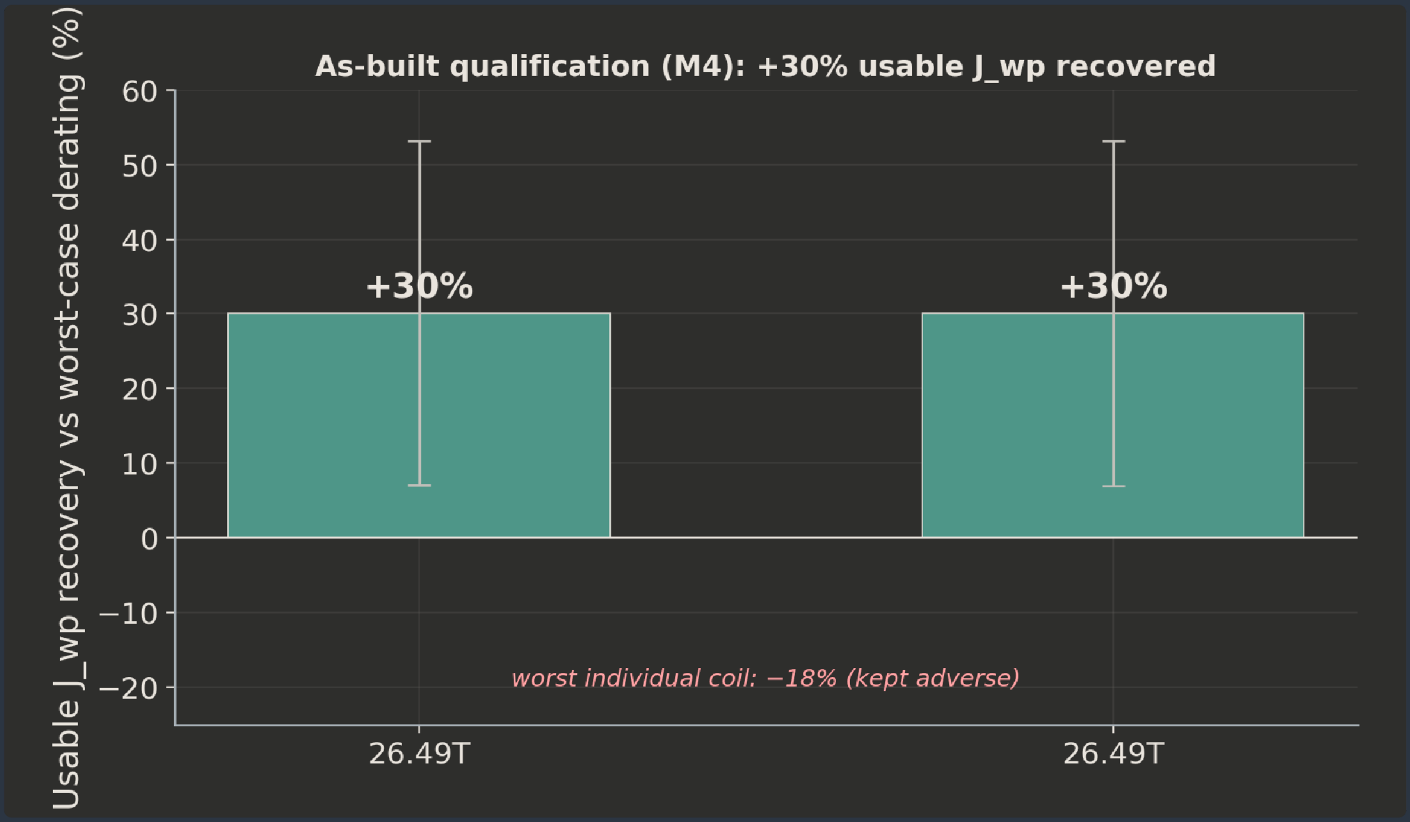


**FIGURE 26.** As-built-data-driven per-coil qualification (M4): +30.0% recovered usable winding-pack current density vs worst-case derating; the worst-coil −18% adverse is labelled. TWIN · DESIGN-STUDY

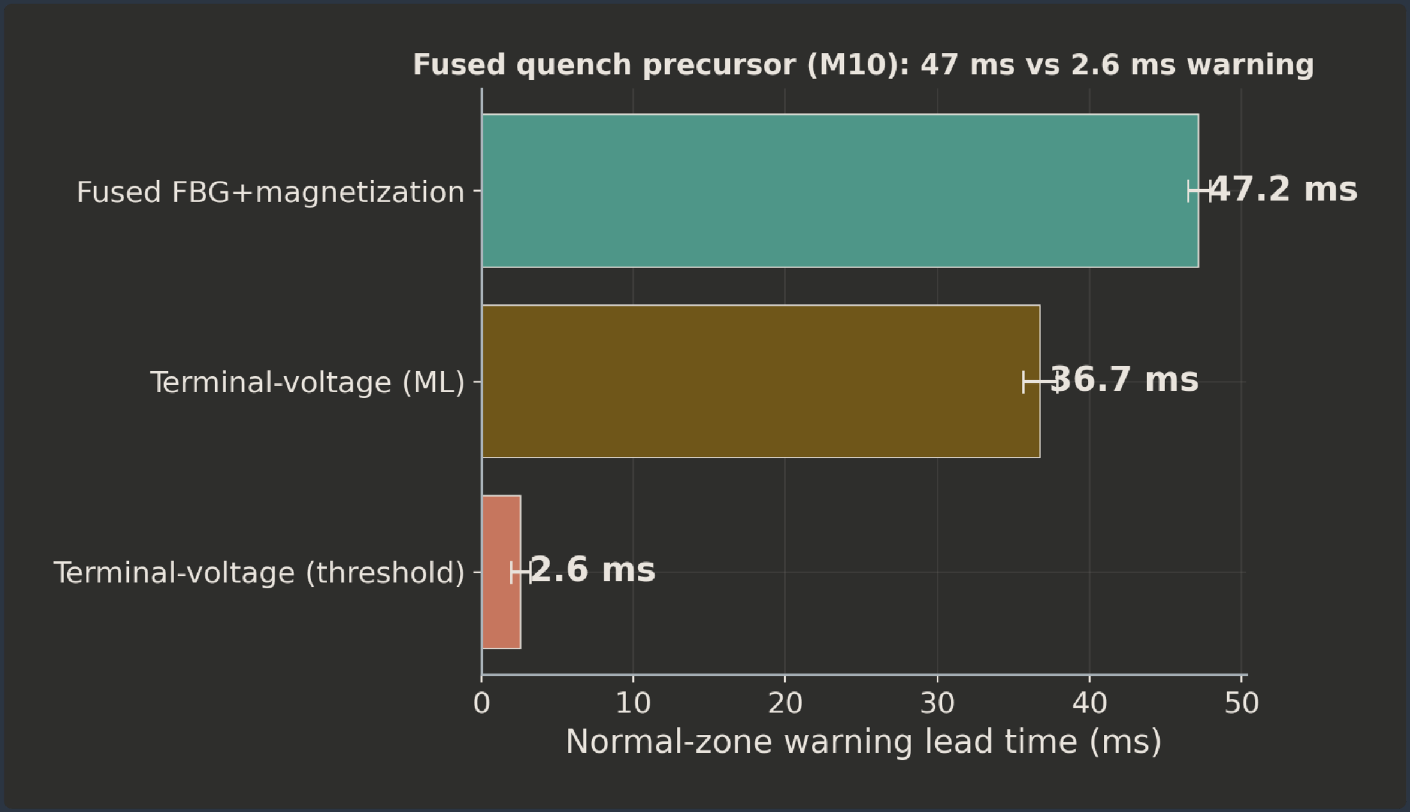


**FIGURE 27.** Fused quench precursor (M10): the fiber-Bragg-strain + magnetization ensemble detects normal-zone onset 47.2 ms ahead of terminal-voltage thresholding (2.6 ms), on twin-synthetic NI-REBCO sensor streams. TWIN · DESIGN-STUDY

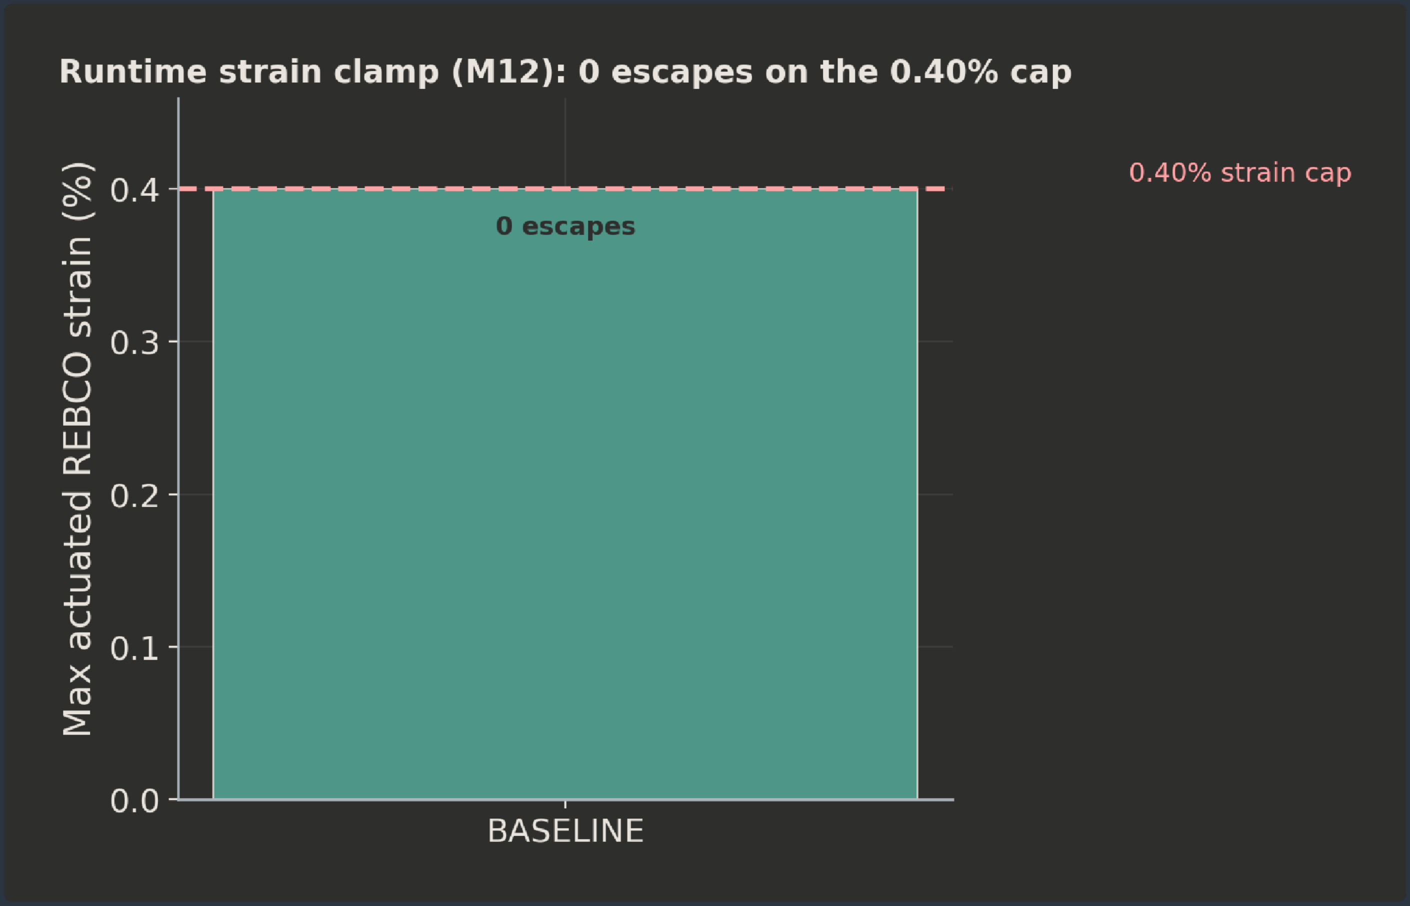


**FIGURE 28.** Runtime strain-margin as a clamped variable of the same rules engine (M12): zero escapes on the 0.40% strain cap over 150k steps. TWIN · DESIGN-STUDY

## THE CALIBRATION SPINE TRANSFERS BEYOND CHEMISTRY

The calibration mechanism is not chemistry-specific. On a plasma-transport closure $\chi_i(\nu^*)$ (E3), gyrokinetic-class anchors reduce the surrogate error **36.8× at 8 anchors** (0.82% relative, coverage 0.900), beating the reduced-model floor by 9.9×. On the magnet surrogate it strengthens to **7475**× maximum error reduction (1856× at 8 anchors), while the P1 provenance gate rejects 4/4 corrupted anchors and the 0.40% strain limit stays immutable. One mechanism, three closure domains.

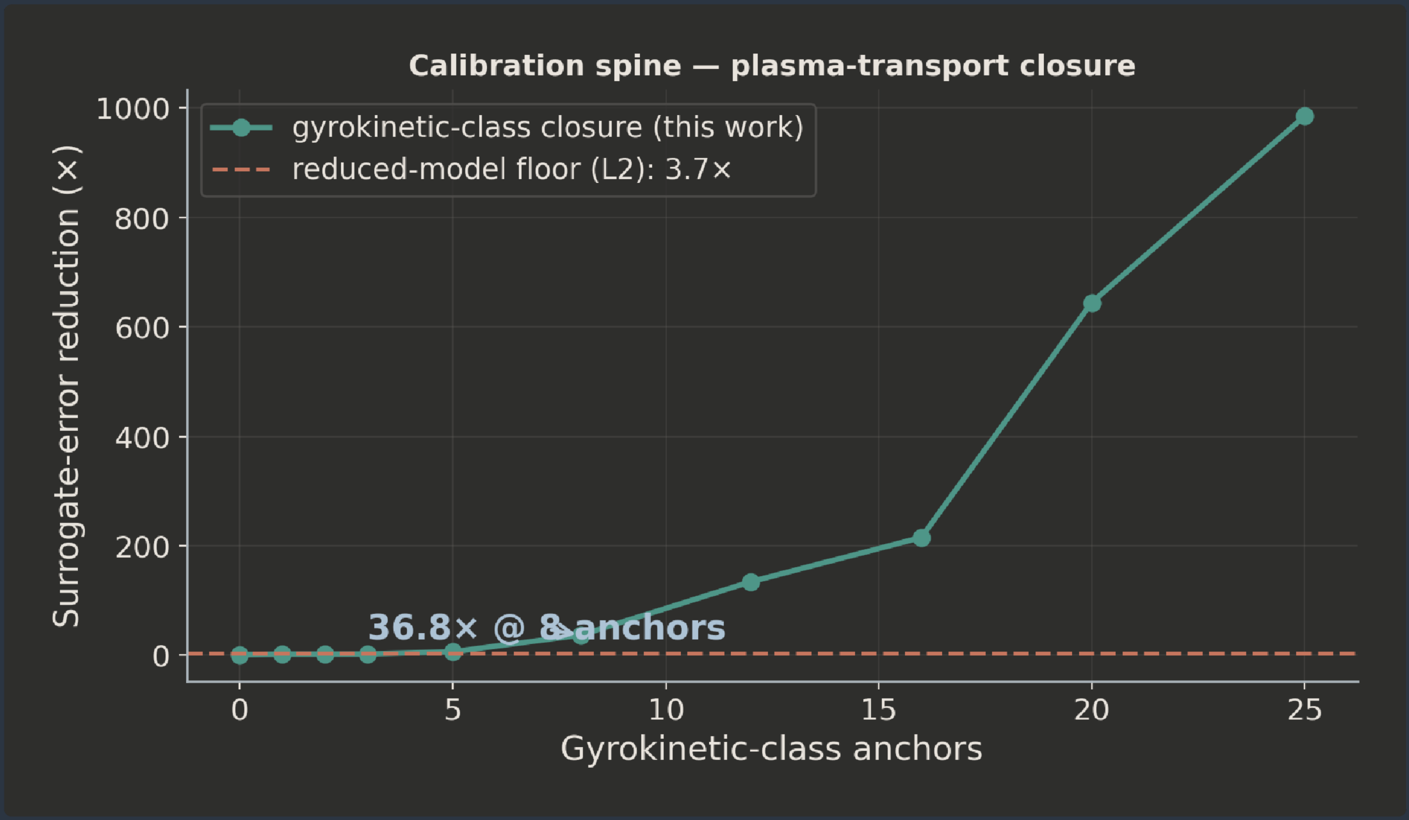


**FIGURE 29.** Calibration spine on a plasma-transport closure: gyrokinetic-class anchors give 36.8× error reduction at 8 anchors, beating the reduced-model floor by 9.9×. SIM

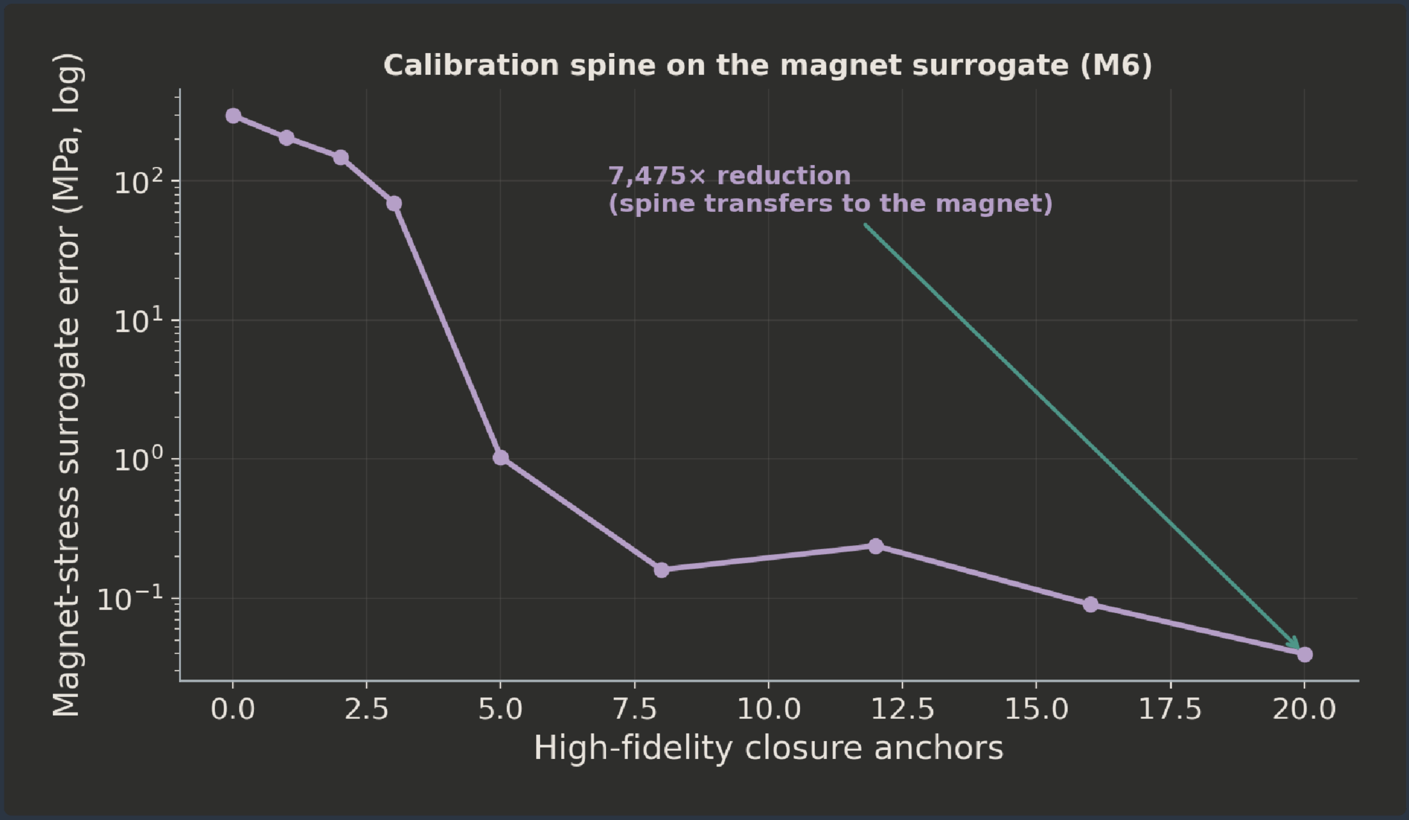


**FIGURE 30.** Calibration spine transferred to the magnet surrogate: up to 7475× error reduction; the P1 provenance gate rejects 4/4 corrupted anchors; the 0.40% strain limit stays immutable. SIM/TWIN

## DE-CIRCULARIZATION — NOT VALIDATING THE SIMULATOR

A twin-trained study risks validating the simulator against itself. Two independent routes break that circularity. First, we validate the twin's equilibrium against the independent open-source FreeGS solver: field relative L2 **2.66%** (within tolerance), with the axis-flux error of 14.3% kept as a characterized failure mode rather than dropped; the circularity index falls **0.348 → 0.165**. Second, we benchmark the twin's magnet, materials, and thermomechanics assumptions against the ITER engineering design corpus — each of **19 assumptions** tagged INSIDE (7), OUTSIDE-BY-DESIGN (3), or NO-ITER-BASIS (9) against a specific ITER document. Divertor heat flux is the one squarely inside an ITER-qualified envelope; REBCO field/strain are outside by concept (ITER uses $Nb_3Sn$) and labelled so.

MANDATORY CAVEAT — AN ENGINEERING-BASIS BENCHMARK, NOT PLASMA VALIDATION

**ITER has not operated.** The 19-assumption comparison is an engineering-basis benchmark against design documentation, not operating-plasma validation. The circularity is reduced, not eliminated (0.35 → 0.17), and that residual is stated. ENG-BASIS

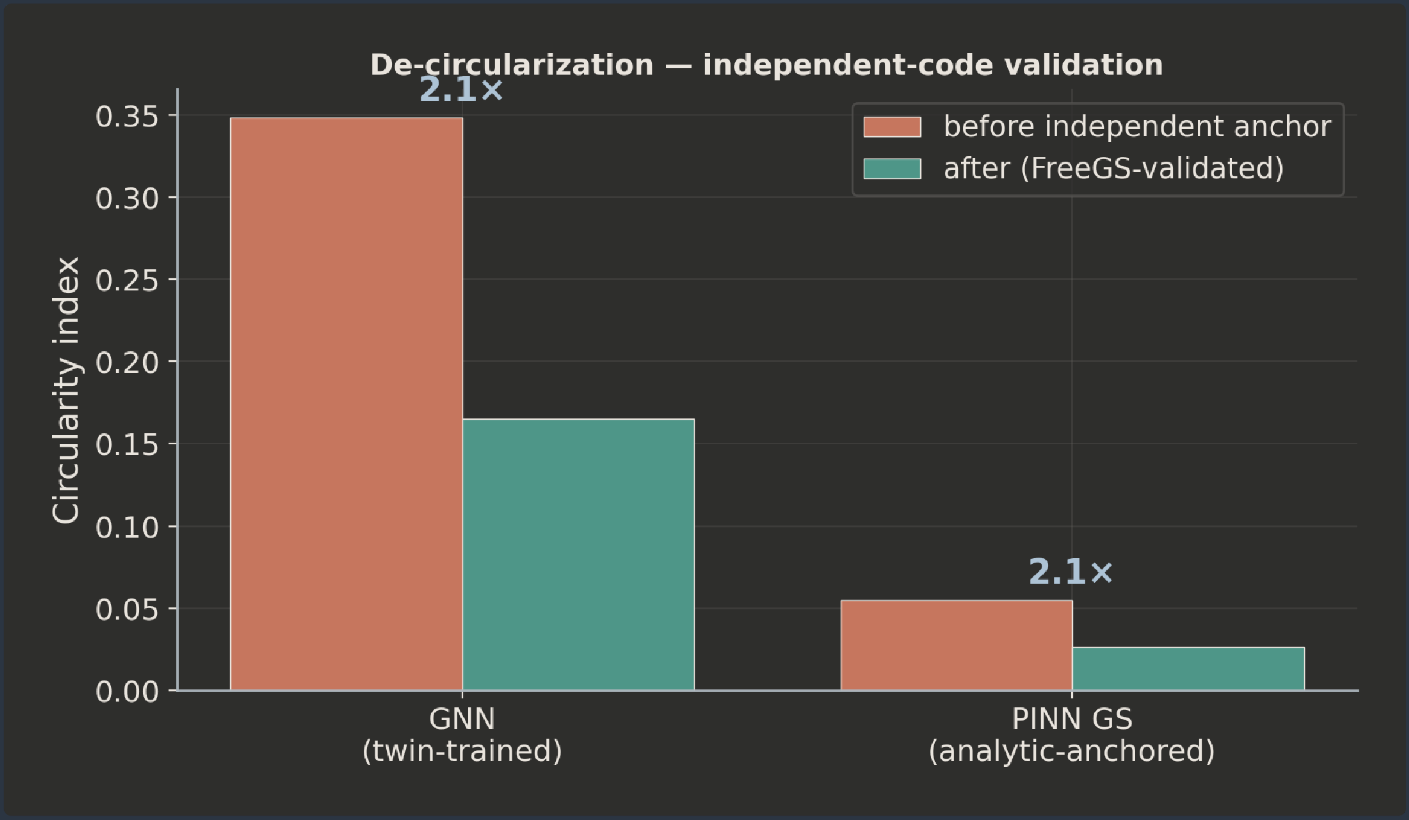


**FIGURE 31.** De-circularization: the twin equilibrium validated against the independent FreeGS community code (2.66% field error); the circularity index halves. SIM/PUBLIC

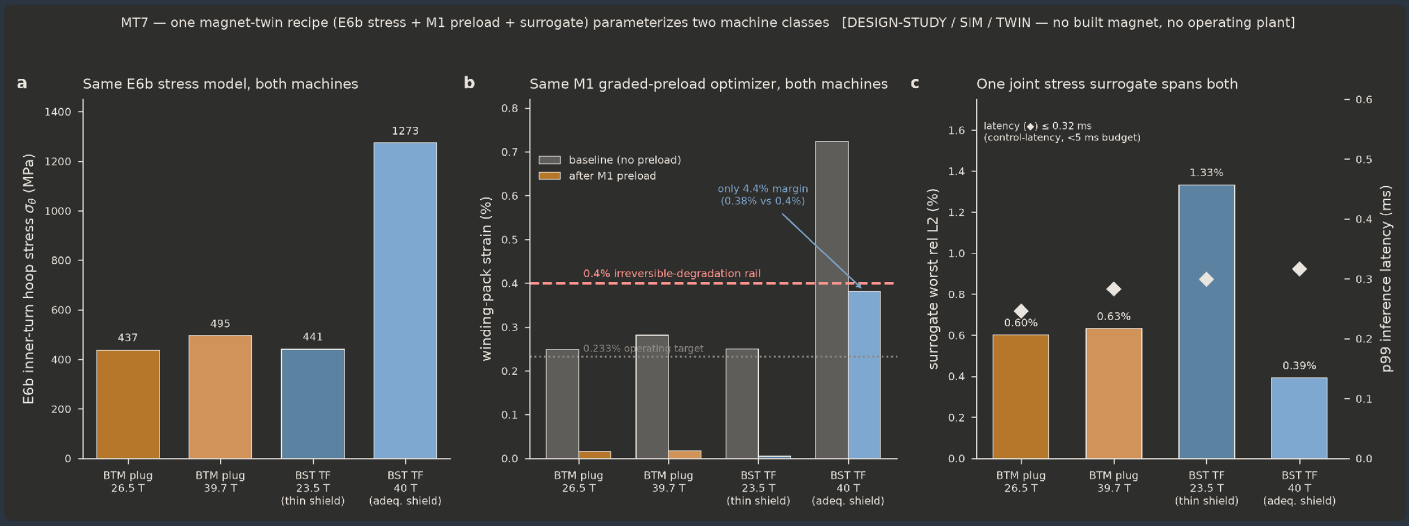


**FIGURE 32.** Machine-class generality: one twin/surrogate recipe spanning the BST toroidal-field and BTM plug coils (worst 1.33% rel L2). SIM

§ 5.11

# Quantum technologies for fusion (Part III)

The quantum part of the paper is where the discipline of the three tiers is tested hardest, because quantum is where fusion-AI claims most often become unfalsifiable. Sensing is near-term and touches the plant; computation is offline, has *no crossover this decade*, and touches the plant only as an offline calibration anchor; and a quantum-*inspired* classical method is the only one of the three that delivers a present-tense win.

## QUANTUM SENSING — NEAR-TERM, DIAGNOSTICS TIER

Quantum magnetometry (NV, atomic, SQUID) and entanglement-enhanced interferometry can sharpen the diagnostic state the AI consumes. Three modalities anchor the tier. **SQUID magnetometry** is the most mature — deployed in fusion diagnostics — offering the highest sensitivity but requiring cryogenic operation and careful stray-field management. **Nitrogen-vacancy (NV) diamond magnetometry** is solid-state, operates from cryogenic to room temperature, and is the most promising for in-vessel deployment, but its demonstrated sensitivity trails SQUIDs and its radiation tolerance in a fusion neutron field is an open question we flag rather than resolve. **Atomic/optically-pumped magnetometers** occupy the middle. The fusion-specific claim is not any single sensor but the *coupling*: a sharpened magnetic-field map tightens the boundary condition of the E1 equilibrium reconstruction, lowering the PINN/GNN state-estimation error under the sensors' measured noise floors. We separate deployed from lab-stage modalities explicitly, and mark which are in-loop-capable (bandwidth-sufficient) versus offline-analysis-only — a sensor that cannot report inside the control latency band cannot enter the loop regardless of its sensitivity. The measurement scaling is stated exactly: shot-noise-limited precision improves as $1/\sqrt{N}$; entanglement-enhanced (Heisenberg) precision as $1/N$.


**FIGURE 33.** Quantum-diagnostics TRL map: modality × measured quantity × TRL × in-loop-capable × feeds-which-AI. TRL-SURVEY

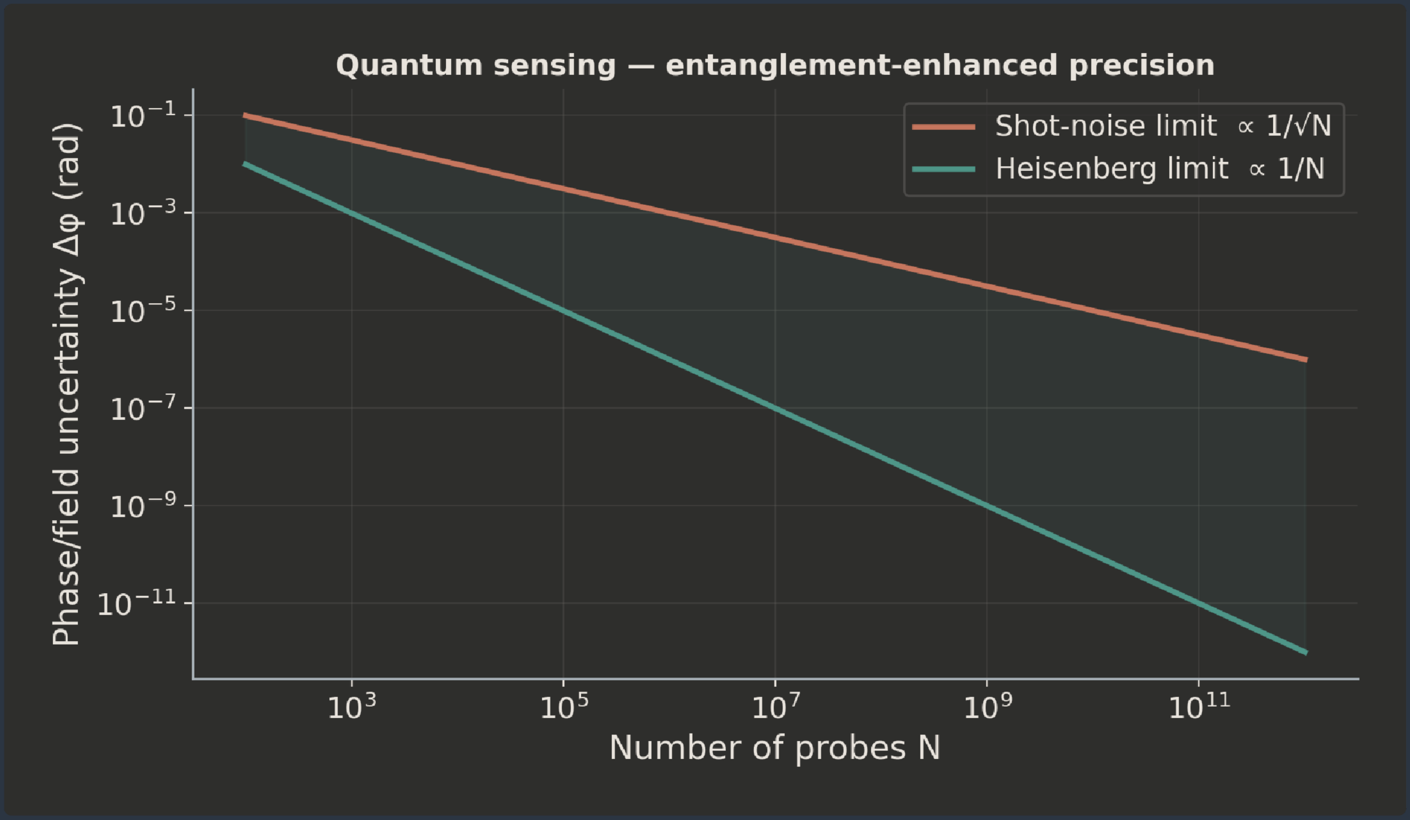


**FIGURE 34.** Precision scaling: shot-noise-limited (∝1/√N) vs entanglement-enhanced Heisenberg (∝1/N), with the fusion-diagnostic operating band shaded. TRL-SURVEY

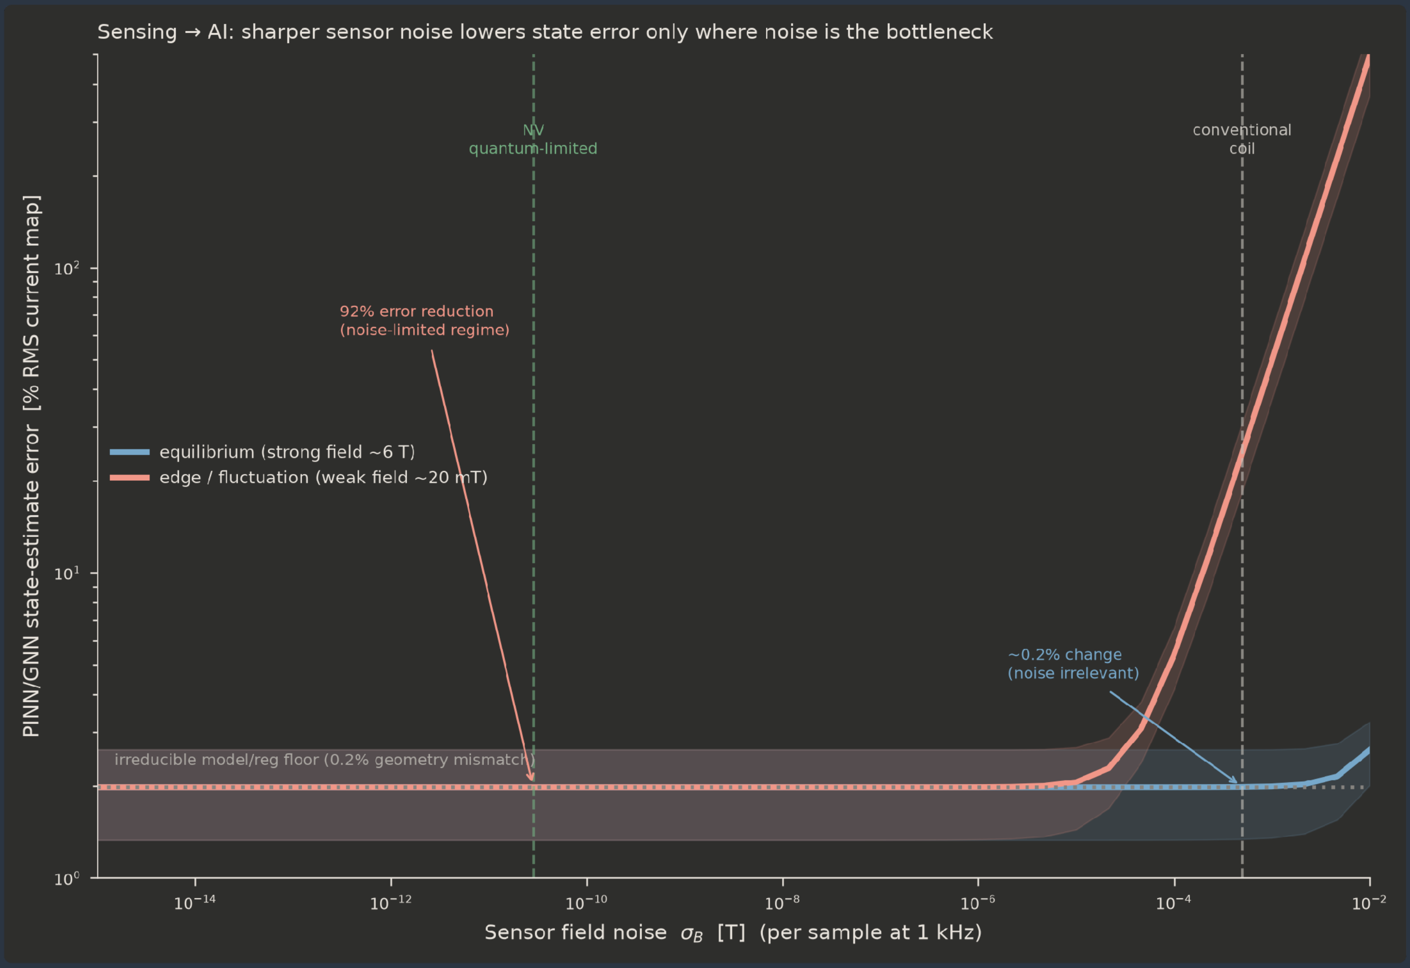


**FIGURE 35.** Sensing → AI: a sharpened quantum-magnetometer field map lowers the PINN/GNN state-estimation error. SIM + TRL-SURVEY

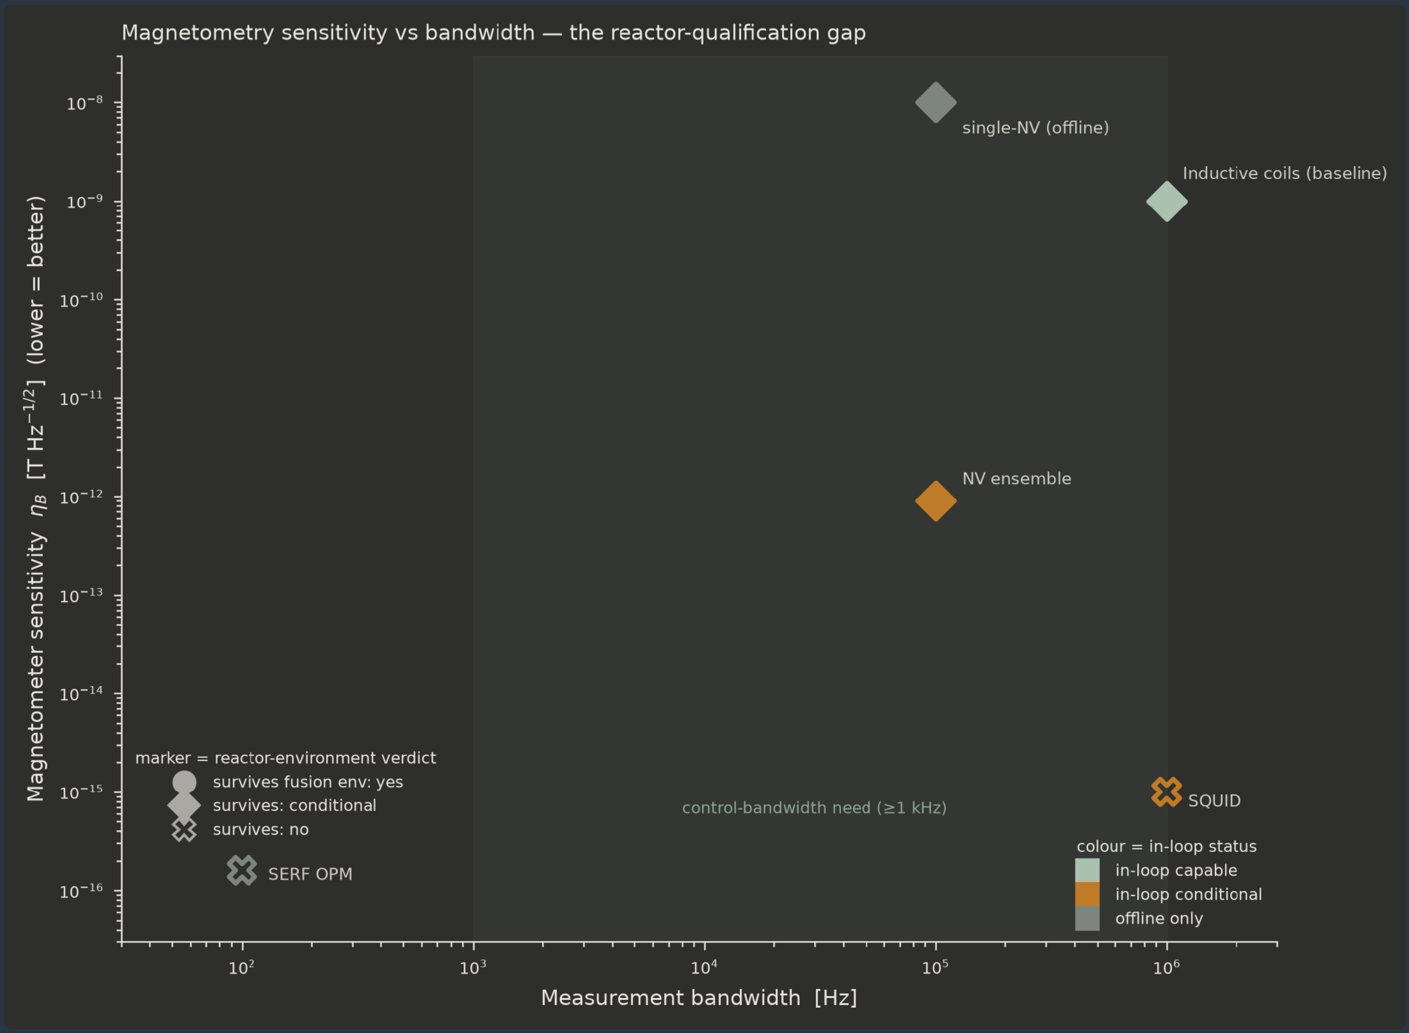


**FIGURE 36.** Magnetometry sensitivity vs bandwidth (NV / atomic / SQUID) against the equilibrium-reconstruction requirement band. TRL-SURVEY

## QUANTUM COMPUTATION — OFFLINE, AND NO NEAR-TERM ADVANTAGE

Fault-tolerant resource estimates for three fusion kernels, under a single documented surface-code cost model, find **no crossover this decade**. The cost model is stated once and applied uniformly: a documented physical error rate, code distance, cycle time, and T-factory count, from which logical-qubit count, T-gate count, circuit depth, and wall-clock follow — every assumption carried in a CSV column so no number is narrated. The three kernels probe different bottlenecks:

| | |
|---|---|
| VLASOV / KINETIC (E3) | ≈**130 logical qubits** and ≈$1.3\times10^{13}$ T-gates, and *still loses* on wall-clock to classical gyrokinetic codes; limited by the Carleman/Koopman–von Neumann tower needed to linearize the collision operator. |
| EQUILIBRIUM LINEAR SOLVE (E1) | ≈$\mathbf{1.8\times10^{17}}$ **T-gates** once read-in/read-out are counted — the dominant, honestly reported penalty. Preparing the right-hand-side state and reading out a useful functional erases the asymptotic speedup. We lead with this number rather than bury it. |
| FIRST-WALL MATERIALS (E8) | The closest to useful; calibrates to within an **order of magnitude** of a published fault-tolerant chemistry estimate — which is how we validate the cost model itself. |

**MANDATORY CAVEAT — NO CROSSOVER THIS DECADE, POC SIMULATOR-ONLY**

Across all three kernels, **no crossover is reached this decade** under any plausible assumption in the swept ranges. A variational proof-of-concept reaches chemical accuracy on a *statevector simulator* but misses it by **14–1300**× under a device-noise model — the scale wall — and is **simulator-only**: no quantum hardware was run. RESOURCE-EST · SIM

F25 Which fusion kernels a quantum computer could accelerate — a taxonomy of CLASSES, not factors

| kernel | speedup CLASS | read-in / read-out penalty | survives read-in/out? | TRL tier |
|---|---|---|---|---|
| K1 kinetic/Vlasov collision op | exponential state-space compression (N^5 amplitudes -> 5 log2 N qubits); ASYMPTOTIC only | read-IN partial (nonlocal coeff loading O(N^3) via QROM); read-OUT FAVOURABLE (transport coeffs are low-order moments) BUT nonlinearity (Carleman) adds truncation cost + qubit tower | NO net advantage at feasible N: block-encoding norm alpha*tau ~ N^2 + clock gap; nonlinearity is the inflation risk (AC-40) | Quantum COMPUTATION — offline, resource-estimated only; NEVER in control loop |
| K2 equilibrium/MHD linear solve | polylog(N) in dimension — APPARENT exponential; ERASED by read-in/read-out | read-IN O(N) (unstructured RHS p',FF' on shaped boundary) + read-OUT O(N/eps^2) (control needs the WHOLE flux/coil field, not one <M>; state destroyed on measurement) + kappa~N | NO — end-to-end scaling is WORSE than classical multigrid O(N); the polylog advantage does not survive. THE LEAD ADVERSE FINDING (Aaronson 2015, 10.1038/nphys3272) | Quantum COMPUTATION — offline, resource-estimated only; NEVER in control loop |
| K3 first-wall/conductor materials ground state | super-polynomial for strongly-correlated ground-state energy vs exact methods; ASYMPTOTIC | read-IN modest (state prep to good HF/CI overlap); read-OUT modest (one eigenvalue) — the FAVOURABLE kernel structurally | Structure survives, but ACTIVE-SPACE REDUCTION is the limitation (AC-42): a tractable M is a DIFFERENT problem from the materials question; and vs the OPTIMIZED classical baseline (DMRG) no wall-clock crossover this decade | Quantum COMPUTATION — offline, resource-estimated only; NEVER in control loop |

RESOURCE-EST — algorithmic taxonomy; no measurement. Each cell carries a speedup CLASS (asymptotic scaling and whether it survives read-in/read-out), never a speedup factor. Quantum COMPUTATION is an OFFLINE calibration tier, NEVER in the control loop. Assumes fault-tolerant hardware that does not yet exist. No Kronos operating-plant data exists.

**FIGURE 37.** Kernel × algorithm taxonomy: Vlasov / equilibrium-linear-solve / materials-VQE against speedup class, read-in/out penalty, and TRL. RESOURCE-EST

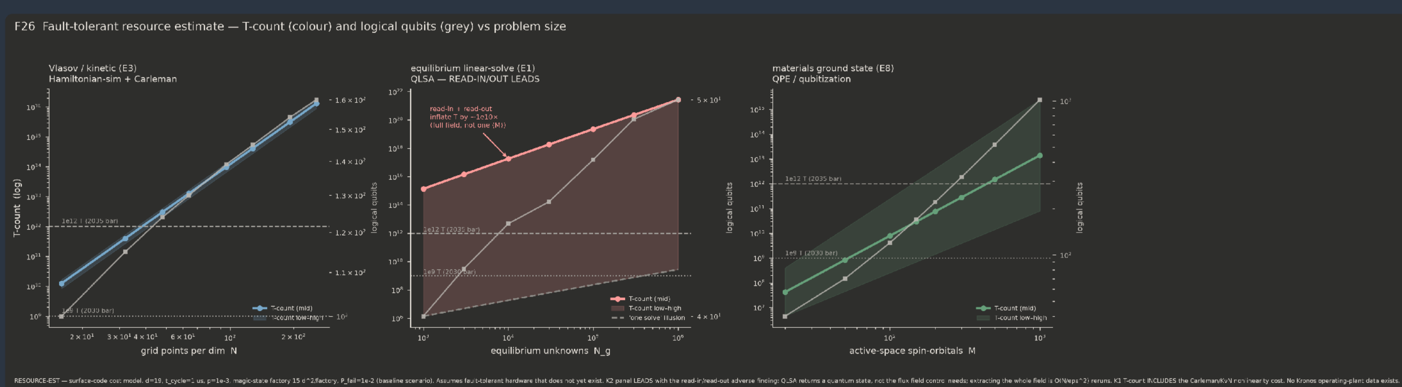


**FIGURE 38.** Fault-tolerant resource estimate: logical qubits and T-count vs problem size per kernel, under a documented cost model. RESOURCE-EST

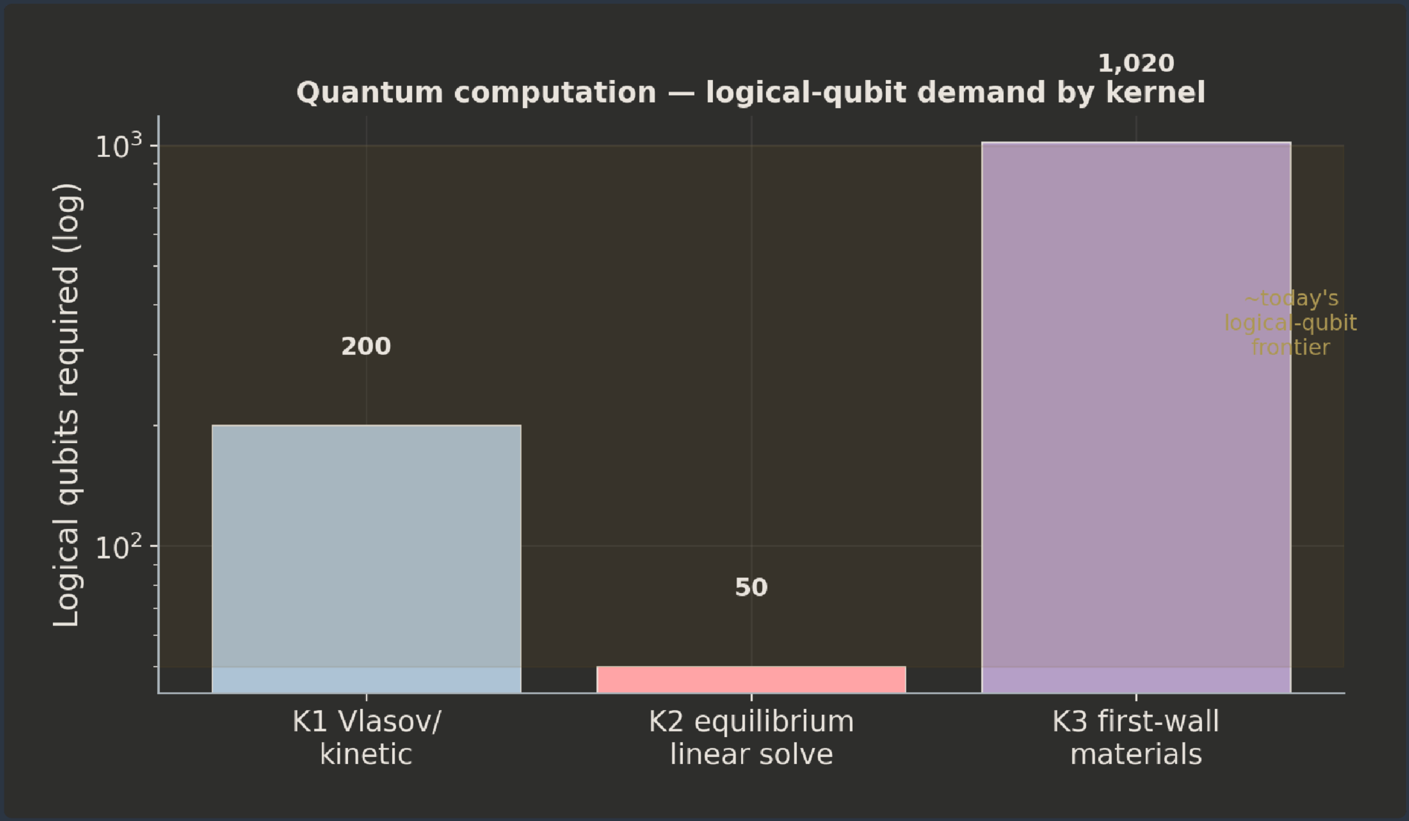


**FIGURE 39.** Logical-qubit demand per fusion kernel; none sits within today's logical-qubit frontier. RESOURCE-EST

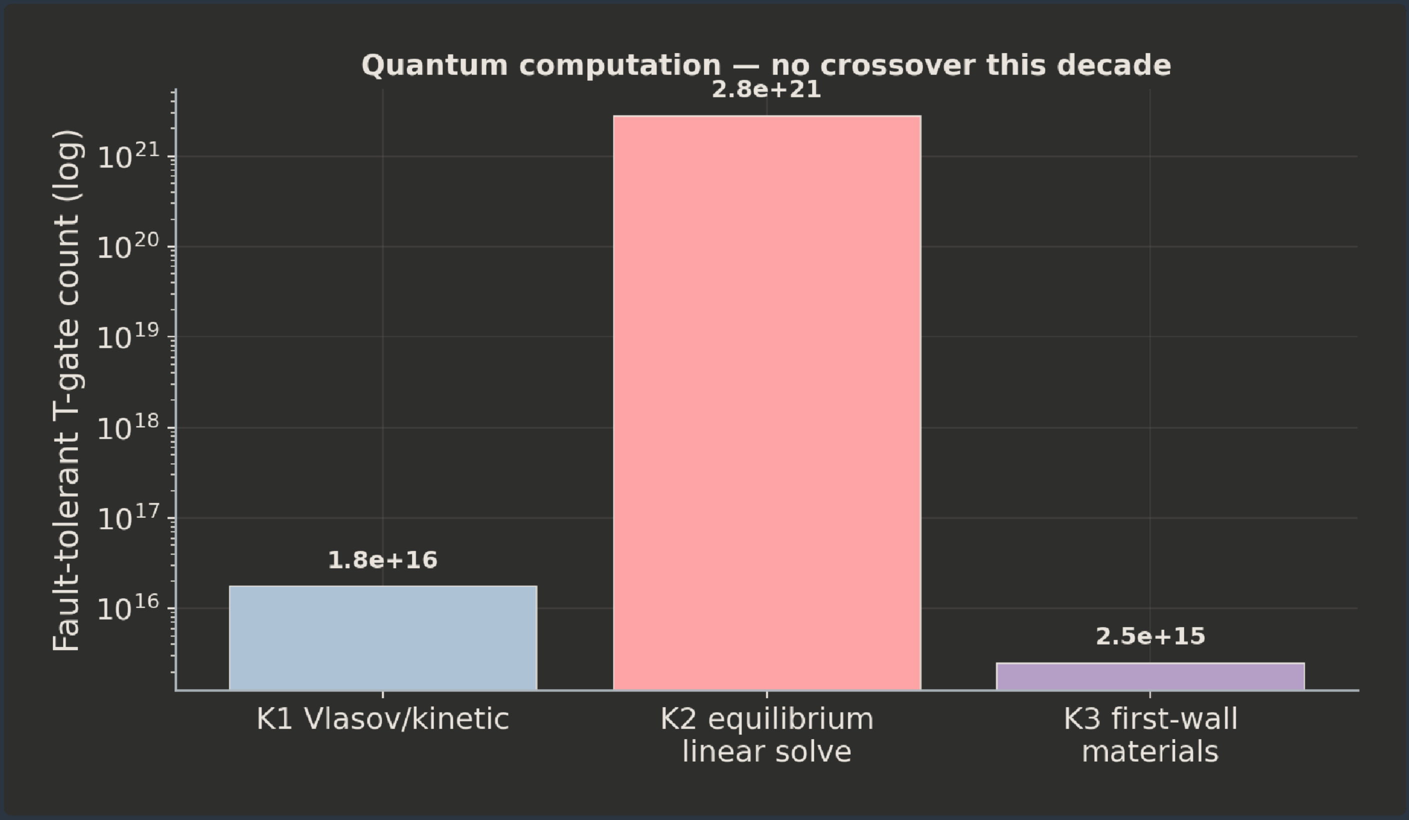


**FIGURE 40.** Classical-vs-quantum crossover: no kernel crosses over this decade; the equilibrium solve needs ≈$1.8\times10^{17}$ T-gates. RESOURCE-EST

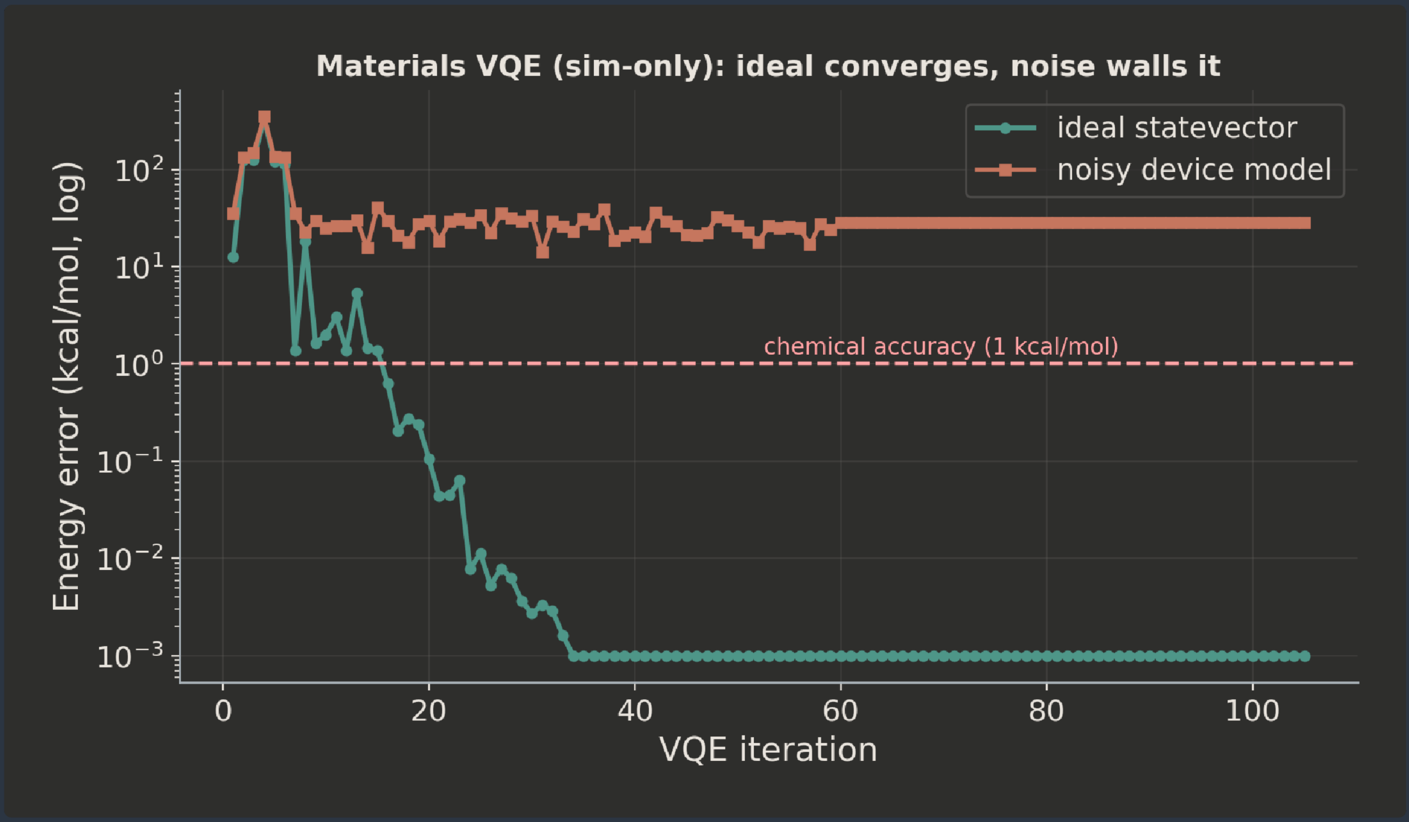


**FIGURE 41.** Materials-VQE proof-of-concept (simulator-only): statevector reaches chemical accuracy; the noisy backend misses it by 14–1300× — the scale wall. SIM

## A PRESENT-TENSE QUANTUM-INSPIRED WIN

Where fault-tolerant computation offers nothing this decade, a quantum-*inspired* classical method does. A matrix-product-state (tensor-network) solver for the kinetic kernel beats the classical baseline at fixed cost by **3.9× at bond dimension d = 6** and **21.2× at d = 7**, and is **the only feasible method at d ≥ 10**, validated to machine precision ($10^{-15}$) against exact diagonalization. Classical methods still win at small d (d = 3, 4), and that is kept. This is a benchmark that bounds the near-term quantum claim — the MPS-Vlasov approach is published — not a new algorithm.

## THE CASE FOR QUANTUM COMPUTING IN FUSION — HONESTLY DATED

The resource estimates are a *negative* result for this decade, and reported as one. They are not, however, an argument that quantum computation is irrelevant to fusion — they are an argument about *timing*. Three fusion problems are classically intractable at reactor scale in a structural, not incidental, way: (i) first-principles turbulent transport, where the gyrokinetic phase-space integration cost for a full-device, multi-scale, long-time prediction outruns any plausible classical machine; (ii) plasma-material and first-wall evolution under fusion-neutron irradiation, a strongly-correlated quantum-chemistry problem where classical approximations lose accuracy exactly where the engineering margin is thinnest; and (iii) the coupled kinetic-electromagnetic problem near marginal stability, where the design-relevant answer is a small difference of large, noisy quantities. For each, a fault-tolerant quantum computer changes the asymptotic cost — so the honest posture is not "quantum will not help" but "here is the specific problem, the specific algorithm, the logical-qubit and error-rate threshold it requires, and therefore the earliest defensible date."

**FORWARD-COMPATIBLE WITH QUANTUM HARDWARE, WITHOUT BETTING THE PLANT ON IT**

When a machine capable of one of these closures exists, its output enters the plant only as an offline, provenance-gated calibration anchor that provably preserves the safety invariant — never as an in-loop dependency. Quantum computing earns its place in fusion at the design-and-calibration layer, on named problems, on a dated roadmap — not as a control-loop miracle.

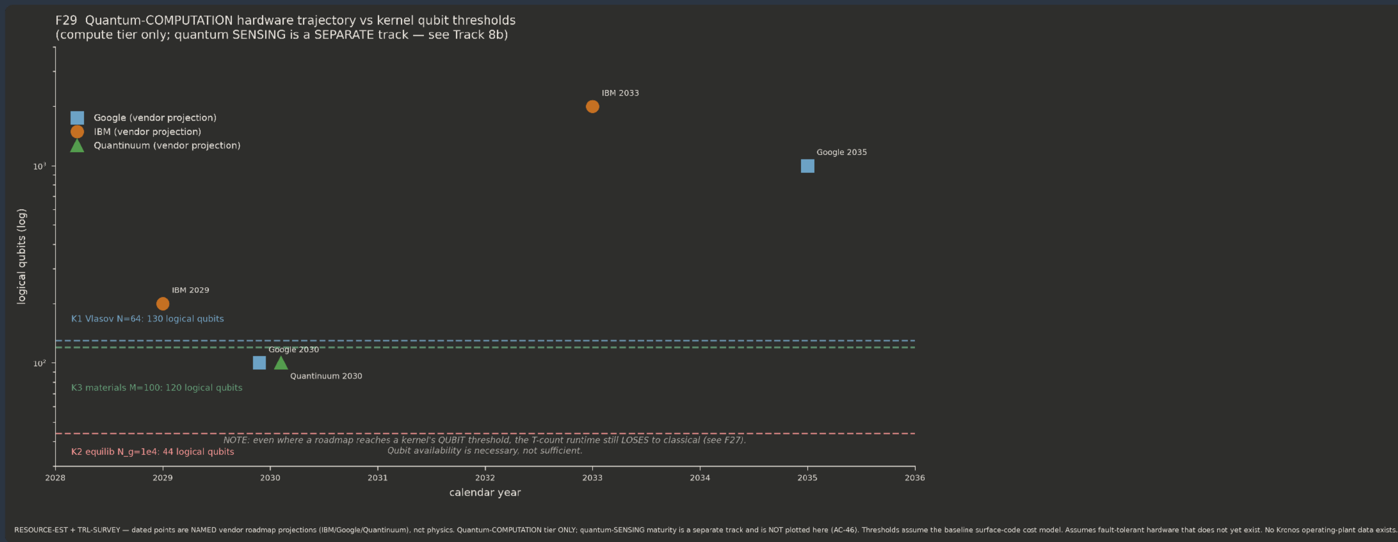


**FIGURE 42.** Quantum roadmap: logical-qubit and error-rate thresholds vs a dated hardware trajectory, sensing and computation on separate tracks. **RESOURCE-EST + TRL-SURVEY**

§ 5.12

# Positioning and novelty — a system, not new physics

The individual components are, by design, established: reinforcement-learning plasma control, deep-learning disruption prediction, physics-informed equilibrium solving, control-barrier-function safe learning, matrix-product-state kinetic solvers, and fault-tolerant quantum chemistry. An *adversarial* prior-art scan downgraded a majority of the study's initially-claimed novelties on their resolved closest art — and that downgrade is reported, not resisted. What survives as genuine contribution is the **certification architecture**, the **quantum-calibrated-surrogate pipeline** across three closure domains, and the **certified-safe-online-learning composition**. These are a system and a method, not new physics.

## TABLE 2 · POSITIONING AGAINST REPRESENTATIVE PRIOR WORK

| PRIOR WORK | WHAT IT ESTABLISHED | WHAT THIS PAPER ADDS |
|---|---|---|
| **Degrave et al. 2022 (Nature)** | RL magnetic control of tokamak shape on a physical device (TCV) | a deterministic clamp + hardware failsafe *outside* the AI, and a formal forward-invariance certificate, so the safety case does not rest on the policy |
| **Kates-Harbeck et al. 2019 (Nature)** | deep-learning disruption prediction, cross-machine | an out-of-distribution abstention that hands unseen classes to the failsafe; honest twin-synthetic provenance where labelled data is unavailable |
| **PINN Grad–Shafranov solvers (2022–24)** | physics-informed equilibrium solving | a kept negative (warm-start, not a fast solver on the ST) and an operator that fixes the diagnosed gradient conflict |
| **Castañeda et al. 2022** | safe online learning with control-barrier functions | safety preserved *through a provenance-gated quantum-surrogate calibration update*, with a computed failure-and-fix |
| **Ye & Loureiro 2022** | matrix-product-state (quantum-inspired) Vlasov solvers | a fusion-kernel benchmark bounding the near-term quantum claim (crossover map + machine-precision validation) |
| **von Burg et al. 2021** | fault-tolerant quantum-chemistry resource estimates | a fusion-materials kernel calibrated to that estimate, inside a three-kernel no-crossover-this-decade result |

The pattern is deliberate: each component is established; the contribution is the certifiable *system* that composes them and the honest accounting of what does and does not yet work.

§ 5.13

# Limitations, kept in full

Every limitation below is kept, each paired with its mitigation and a path to close it, and each reported in its own artifact in the deposit. Nothing here is a footnote to the results; it is the results, seen from the side that a diligence team looks from first.

### TABLE 3 · THE KEPT-LIMITATIONS REGISTER

| LIMITATION (KEPT) | MITIGATION / STATUS | PATH TO CLOSE |
|---|---|---|
| **No operating plant** | all results tagged SIM/TWIN; no operating-plant claim | real-machine data via a data-use agreement |
| **No real-plant disruption data** | twin-synthetic anchored to the real MAST operating space | labelled public dataset (currently credential-gated) |
| **PINN not a fast native solver (BST 2.5%)** | reframed as warm-start; operator fixes the gradient conflict | neural-operator training at larger scale |
| **$\beta_N$ safety face certified only as a sub-region** | largest certified invariant sub-region reported; empirical clamp covers the rest | move the operating point off the no-wall limit |
| **No quantum-computing crossover this decade** | pre-registered negative; quantum confined to the offline tier | dated hardware thresholds (roadmap) |
| **Quantum PoC simulator-only** | stated in methods and captions | cloud quantum-hardware run |
| **Magnet study is design-study (no built magnet)** | every magnet result tagged DESIGN-STUDY | build and as-built qualification |
| **Fluence-graded tape (M7) OPEN-RISK** | plug fluence uncomputed; held open, not claimed | neutron-transport run (MCNP/OpenMC/Serpent) |
| **Circularity reduced, not eliminated (0.35→0.17)** | two independent de-circularization routes | more independent-code / experimental anchors |
| **Sub-10 μs failsafe tier is a specification** | covered by the L1 hardware layer; trip time modeled | FPGA/hardware measurement |

§ 5.14

# Operational value of the architecture

Beyond the safety case, an AI-native twin changes plant economics through three levers, stated here *qualitatively* — a quantified plant model is out of scope for a methods paper and would require operating data the program does not yet have.

| | |
|---|---|
| AVAILABILITY | A predictive shadow that anticipates off-normal events and a predictive-maintenance loop that schedules component replacement against modeled fatigue (rather than fixed intervals) both act to raise capacity factor — the single largest driver of the levelized cost of any capital-heavy generator. |
| DISPATCH | A direct-energy-conversion plant produces a DC bus whose output the power-systems copilot can bid into grid frequency-regulation and energy markets in real time; the same twin that keeps the plasma safe also knows, instant to instant, what the machine can deliver. |
| OPERABILITY | State reconstruction under sensor loss, operator-training on the decoupled twin, and copilots that ground engineering decisions in the governed data fabric all reduce the specialized-staffing burden that would otherwise make a fleet of small machines uneconomic. |

None of these is quantified here, and none is a market projection; they are the operational reasons the certification architecture is worth building beyond the fact that it is required.

§ 5.15

# Future work and roadmap

Three tracks follow directly from the kept limitations, and each converts a kept limitation into a measured result rather than changing any claim. **Validation on real operating-machine data** is the single largest credibility step: the equilibrium surrogates, disruption predictor, and controllers should be re-run against published equilibria and labelled disruption databases from operating tokamaks under a data-use agreement, moving the results from twin-synthetic to independently validated.

**Hardware-grounded safety** closes the two open certification items — measuring the sub-10 μs failsafe trip time, and extending the control-barrier certificate to all four faces once the operating point is moved off the no-wall limit. **The quantum calibration tier** advances on its own dated roadmap: as fault-tolerant machines approach the logical-qubit and error-rate thresholds for the kinetic and materials kernels, each closure is admitted through the provenance gate proven in the certification section. In parallel, the magnet design-study results (M1–M12) are reduced to practice on built hardware through the as-built qualification method, and the OPEN-RISK plug-fluence item (M7) is closed with a neutron-transport run. None of these changes the paper's claims; each converts a kept limitation into a measured result.

§ 5.16

# Conclusions

A layered digital twin can control a compact fusion generator with physics-informed machine learning at control-relevant latency and bounded, clamped error; the residual model error can be attacked offline by high-fidelity — and prospectively quantum-computed — closures admitted through a provenance gate that provably preserves a safety invariant; and **quantum sensing, not quantum computation, is the near-term quantum technology that touches the plant**.

**FAST ENOUGH TO ACT, HONEST ENOUGH TO CERTIFY**

The architecture is fast enough to act, honest enough to certify, and forward-compatible with quantum hardware without depending on it. The study claims a system and a method, demonstrated in simulation across two machine classes — the negative-triangularity spherical-tokamak breeder twin (BST; $R_0$ 2.0 m, A 2.0, κ 2.56, $B_0$ 10 T, a *distinct machine* from the frozen breeder design point) and the D–$^3$He tandem-mirror burner — and reports every missed target. The natural next step is validation on real operating-machine plasma data.

All code, data, figures, the reproduction harness, and the pre-registered acceptance-criteria register are archived at Zenodo under DOI **10.5281/zenodo.21842371** and regenerate from seed 20260726. Tier-1 quantities reproduce byte-identically; Tier-2 (trained-model) metrics are five-seed tolerance bands. No result uses or implies operating-plant telemetry; **Kronos Fusion Energy operates no fusion plant.** DEPOSITED · DOI 10.5281/ZENODO.21842371

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.

LIST OF FIGURES

# List of Figures

LIST OF TABLES

# List of Tables

REFERENCES

# Sources

**[1]** H.-S. Bosch and G. M. Hale, “Improved formulas for fusion cross-sections and thermal reactivities,” *Nuclear Fusion* **32**, 611 (1992).

**[2]** A. S. Richardson, *NRL Plasma Formulary*, Naval Research Laboratory (2019).

**[3]** I. E. Ochs, E. J. Kolmes, and N. J. Fisch, “Bremsstrahlung and the electron–electron contribution in fusion plasmas,” *Physics of Plasmas* (2024).

**[4]** T. H. Rider, “Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium,” *Physics of Plasmas* **4**, 1039 (1997).

**[5]** ITER Physics Expert Groups, “ITER Physics Basis,” *Nuclear Fusion* **39**, 2137 (1999).

**[6]** N. A. Uckan and J. Sheffield, “Tokamak scaling and confinement,” in *Tokamak Concept Improvement*, IAEA (1989).

**[7]** O. Sauter, C. Angioni, and Y. R. Lin-Liu, “Neoclassical conductivity and bootstrap current,” *Physics of Plasmas* **6**, 2834 (1999).

**[8]** F. Troyon *et al.*, “MHD limits to plasma confinement,” *Plasma Physics and Controlled Fusion* **26**, 209 (1984).

**[9]** J. P. Freidberg, *Plasma Physics and Fusion Energy*, Cambridge University Press (2007).

**[10]** R. F. Post, “The magnetic mirror approach to fusion,” *Nuclear Fusion* **27**, 1579 (1987).

**[11]** D. E. Baldwin, “End-loss processes from mirror machines,” *Reviews of Modern Physics* **49**, 317 (1977).

**[12]** T. K. Fowler, R. W. Moir, and T. C. Simonen, “A new simpler tandem mirror,” *Nuclear Fusion* **57**, 056014 (2017).

**[13]** G. L. Kulcinski and J. F. Santarius, “Advanced fuels and the D–$^{3}$He cycle,” University of Wisconsin Fusion Technology Institute.

**[14]** W. L. Barr and R. W. Moir, “Test results on direct energy converters for mirror reactors,” *Nuclear Technology/Fusion* **3**, 98 (1983).

**[15]** D. B. Go *et al.*, “Thermionic energy conversion in the twenty-first century,” *Frontiers in Mechanical Engineering* **3**, 13 (2017).

**[16]** M. Wang *et al.*, “The AME 2020 atomic mass evaluation,” *Chinese Physics C* **45**, 030003 (2021).

**[17]** J. E. Menard *et al.*, “Fusion nuclear science facilities and pilot plants based on the spherical tokamak,” *Nuclear Fusion* **56**, 106023 (2016).

**[18]** A. J. Creely *et al.*, “Overview of the SPARC tokamak,” *Journal of Plasma Physics* **86**, 865860502 (2020).

**[19]** T. P. Wright, “Factors affecting the cost of airplanes,” *Journal of the Aeronautical Sciences* **3**, 122 (1936).

**[20]** U.S. Department of Energy / NNSA, *Stockpile Stewardship and Management Plan*, FY2024.

**[21]** J. F. Santarius, G. L. Kulcinski, and L. A. El-Guebaly, “A strategy for D–$^{3}$He fusion energy development,” *Journal of Fusion Energy* **17**, 33 (1998).

**[22]** M. A. Abdou *et al.*, APEX advanced blanket concepts, *Fusion Engineering and Design* (1999–2003).

**[23]** Federal Register, notice on supply of tritium to non-federal entities (1999).

**[24]** Kronos Fusion Energy, “Simulation register and reproducibility deposit,” Zenodo (2026).

**[25]** A. O. Thome *et al.*, “ELM-free negative-triangularity edge on DIII-D,” *Plasma Physics and Controlled Fusion* (2024), doi:10.1088/1361-6587/ad6f40.

**[26]** T. Eich *et al.*, “Scaling of the tokamak near scrape-off layer power width,” *Nuclear Fusion* **53**, 093031 (2013).

**[27]** J. Lim *et al.*, “Scrape-off-layer width in negative-triangularity plasmas,” *Nuclear Fusion / Plasma Physics and Controlled Fusion* (2023).

**[28]** A. O. Nelson *et al.*, “Prohibition of second-stability access at negative triangularity,” *Plasma Physics and Controlled Fusion* **64**, 124002 (2022).

COLOPHON

# How this document was made

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

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

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

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

LEGAL NOTICE

# Notice, status, and distribution

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

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

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

THE 2026 KRONOS SET

# One programme, in five papers

This editorial is one of five that together describe the Kronos fusion programme — a breeder that funds a generator, and the magnet, conversion, and control science that enable both. Each is published open-access with a reproducible companion deposit.

| | |
|---|---|
| PAPER I | **Hyperion — The Tritium Foundry** · 10.5281/zenodo.21746157 |
| PAPER II | **Aegis & MetroVolt — The Generator** · 10.5281/zenodo.21746479 |
| PAPER III | **High-Field REBCO Magnets & Tape** · 10.5281/zenodo.21842514 |
| PAPER IV | **Direct Energy Conversion** · 10.5281/zenodo.21842864 |
| PAPER V | **AI, ML & Quantum Control** *(this paper)* · 10.5281/zenodo.21842371 |

Explore the whole programme live: the interactive 3-D model, the in-browser physics validator, and the companion film series.

KRONOS FUSION ENERGY

# *AI, ML & Quantum Control*

Explore the science — live, in your browser.

kronosfusionenergy.com/3D-Model

kronosfusionenergy.com/physics-Validation-Simulation

Companion film series · YouTube

Zenodo · DOI 10.5281/zenodo.21842371