KRONOS FUSION ENERGY

PAPER 2.69 · THE KRONOS 2026 SERIES

A Certified Control-Barrier Safety Filter for Fusion Control: Forward-Invariance Guarantees, Calibrated Uncertainty, and an Honest PINN Negative for a Compact Spherical-Tokamak Breeder

A nuclear-regulated fusion machine must guarantee that every automated actuation keeps the plasma inside a declared safe operating envelope, and it must do so in a way an auditor...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

A Certified Control-Barrier Safety Filter for Fusion Control: Forward-Invariance Guarantees, Calibrated Uncertainty, and an Honest PINN Negative for a Compact Spherical-Tokamak Breeder

A nuclear-regulated fusion machine must guarantee that every automated actuation keeps the plasma inside a declared safe operating envelope, and it must do so in a way an auditor can *check*, not merely trust. Learned and optimisation-based controllers deliver performance but, by themselves, provide no such guarantee. We formalise and verify a deterministic, control-barrier-certified safety filter for a compact spherical-tokamak (ST) breeder that provides a provable zero-escape property. The filter sits between an arbitrary nominal controller and the actuators and returns the minimum-norm edit of the command that keeps the operating point inside a safe set defined by three operational barriers — a Greenwald/first-wall-load density margin, a normalised-pressure (β_N) margin, and a q-kink margin. Robustness to a bounded disturbance $\|w\| \le W$ enters as an explicit, auditable additive tightening $\eta = \|\nabla h\| W$ of each barrier; because the actuation is diagonal the filter decouples into three closed-form one-dimensional projections, so it runs solver-free in constant time. We state the reduced-order control-affine model, the robust barrier condition, the closed-form projection and an analytic forward-invariance theorem as explicit equations, and we verify them on the frozen breeder β_N clamp: from the design-point anchor ($Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$, $\beta_{N,0} = 1.532$) an unfiltered greedy controller ($u = 1.25$) drives β_N to 4.200 and violates the barrier on 50 of 100 steps ($\min h = -0.20$), while the filtered controller holds $\beta_N \le 3.500$ with $\min h = +0.000$ and 0 of 100 violations, using authority confined to $u \in [1.061, 1.250]$ and never touching the actuator floor — so the certificate is not an artefact of saturation. The closed-form projection matches a general quadratic-program solver to machine precision (2.2×10^{-16}). We report the surrounding de-risking layer with honest maturity, including a pre-registered negative: a physics-informed neural network intended as a fast native equilibrium solver reaches only 1.4% relative-L^2 error, missing its 1% target — reported, not reframed. We state the scope plainly: the guarantee is for the reduced-order operating-envelope model and the declared disturbance bound; elevating it to a plant guarantee requires validating the model against higher-fidelity core-transport, neoclassical and disruption simulation.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645716 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: control barrier function · safety filter · forward invariance · quadratic program · calibrated uncertainty · physics-informed neural network · compact spherical tokamak

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

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

§ 1

Introduction

SAFETY AS A CHECKABLE PROPERTY

Autonomous operation of a compact, high-power-density spherical tokamak leaves little time to react and no tolerance for an unmodelled excursion. The Hyperion breeder [34] holds a plasma current $I_p = 9.66\,MA$ 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 at gain $Q = 3.076$ and fusion power $P_{\text{fus}} = 85.04\,MW$. Fusion control requirements span many orders of magnitude in time, from microsecond magnet protection to month-scale fuel logistics, and a nuclear-regulated machine additionally demands that every automated actuation be bounded, auditable and reproducible. The central difficulty is not that a good controller is hard to build — magnetic diagnostics feed real-time equilibrium reconstruction [14] and shape controllers [15,17], physics-based profile simulators run inside the control cycle [16], and deep reinforcement-learning agents now command coil voltages directly from magnetics [11] and steer plasmas away from tearing instability [12]. The difficulty is that none of these, by itself, hands an auditor a *guarantee*: a learned policy is a function whose out-of-sample behaviour is not certified, and a model-predictive controller inherits whatever its internal model gets wrong. This paper is about the safety layer specifically — a filter that wraps *any* nominal controller and certifies, with a proof and a reproducible verification, that the plasma operating point cannot leave a declared envelope.

CONTROL BARRIER FUNCTIONS

The natural formal tool is the *control barrier function* (CBF). A continuously differentiable function h whose zero-superlevel set $\mathcal{C} = \{x : h(x) \geq 0\}$ is the safe set is a CBF if there exists a control that keeps $\dot{h}$ above a class-$\mathcal{K}$ function of h; enforcing that inequality renders $\mathcal{C}$ forward-invariant, so a trajectory that starts safe stays safe [1,2,3]. The idea descends from Nagumo's set-invariance theorem and the control-theoretic notion of viability [7], and has matured into a practical safety layer: for a control-affine plant the invariance condition is affine in the control, so the minimum-norm safe edit of any nominal command is the solution of a small quadratic program (QP) [2]. Extensions handle bounded disturbances [4], constraints of high relative degree [5,6], and the closely related idea of a predictive safety filter built on a terminal safe set [8]. The CBF filter is attractive here for three reasons an auditor cares about: the safe set is declared explicitly, the disturbance model is declared explicitly, and the resulting guarantee is a theorem rather than a statistic.

THE GAP AND OUR CONTRIBUTION

What the fusion-control literature has demonstrated is performance — learned control [11,12], fast surrogates [32], real-time equilibria [14,16]. What a nuclear ST additionally needs is a certified floor beneath that performance: a layer that (i) encodes the operating envelope as barriers tied to the machine's actual stability and density limits, (ii) is robust to a stated disturbance bound with an auditable margin, (iii) runs deterministically at control latency without an online iterative solver, and (iv) comes with a proof and a reproducible zero-escape verification. This paper contributes exactly that. We give the reduced-order control-affine operating-envelope model and the three barriers as explicit equations (§2); the robust barrier condition, its additive disturbance tightening, and the closed-form separable projection that makes the filter solver-free (§3); an analytic forward-invariance theorem with its discrete-time analogue (§3.4); the NVIDIA GPU HPC campaign that fixed the envelope caps and produced the verification (§4); a full verification of the breeder β_N clamp with eight-plus figures and the surrounding de-risking layer reported at honest maturity, including a pre-registered PINN negative (§5); and a quantitative discussion against the published CBF and plasma-control literature (§6). Every headline quantity traces to a frozen anchor in the Kronos Physics De-Risking Register [43]; serial identifiers (e.g. KX-L1-A13) are given inline so each number is auditable. The filter is the certified floor beneath the model-predictive burn controller [37], the safe reinforcement-learning policy [38], and the real-time optimisation backbone [40] reported elsewhere in this series.

§ 2

The operating-envelope model and the formal safety problem

REDUCED-ORDER CONTROL-AFFINE DYNAMICS

We work with a reduced-order model of the operating *point*, not the full plasma. Let $x = (f_G,\, \beta_N,\, q_{95})^\top \in \mathbb{R}^3$ collect the Greenwald density fraction $f_G = \bar{n}_e / n_G$, the normalised pressure β_N, and the edge safety factor q_{95}. The actuators (auxiliary heating and fuelling gains, normalised so that $u = 1$ holds the anchor) act through a diagonal input matrix, and an additive disturbance w captures the unmodelled slow drift of the reduced model relative to the plant. The dynamics are control-affine,

$$\dot{x} = f(x) + g(x)\, u + w, \qquad \|w\| \le W, \qquad u \in \mathcal{U} = [u_{\min}, u_{\max}]^3,$$

with $g(x) = \mathrm{diag}\,\big(g_G(x), g_\beta(x), g_q(x)\big)$. For the concretely verified pressure channel the scalar dynamics used in gate KX-L1-A13 are

$$\dot{\beta}_N = \kappa\big(u_\beta - u_{\mathrm{eq}}(\beta_N)\big) + w_\beta, \qquad u_{\mathrm{eq}}(\beta_N) = 1 + s\,(\beta_N - \beta_{N,0}),$$

where u_{eq} is the equilibrium (holding) control and $\beta_{N,0} = 1.532$ is the design-point anchor. Equation 2 is deliberately low-order: it is a design-point-anchored surrogate whose only role is to define a safe set the filter can certify, and whose fidelity to the plant is an explicit open item (§7). The model is grounded upstream by the free-boundary equilibrium and ideal-MHD analysis that fixes the anchor $q_{95} = 2.90$ and the Troyon reference (BR-L1-A12, BR-L2-A19; [35]), and by the neutronics that fix the density cap (BR-CX-09).

THE THREE OPERATIONAL BARRIERS

The safe set is the intersection of three linear barrier half-spaces, each declared conservatively below the hard physical limit so that the certified envelope sits inside the true one:

$$\begin{aligned} \text{Greenwald / FW-load:}\quad & h_G(x) = 1 - f_G/f_G^{\mathrm{cap}}, \qquad & f_G^{\mathrm{cap}} = 0.45, \\ \text{normalised pressure:}\quad & h_\beta(x) = 1 - \beta_N/\beta_N^{\mathrm{cap}}, \qquad & \beta_N^{\mathrm{cap}} = 3.5, \\ q\text{-kink margin:}\quad & h_q(x) = q_{95}/q_{\min} - 1, \qquad & q_{\min} = 2.2. \end{aligned}$$

The pressure cap $\beta_N^{\mathrm{cap}} = 3.5$ is the no-wall operating margin carried through the stability analysis [35] (Troyon scaling [20]); the operating point sits well inside it at $\beta_{N,0} = 1.532$. The Greenwald cap is set by the first-wall neutron load rather than the gas limit (BR-CX-09: the wall load binds density near $f_G \approx 0.36$, below the 0.45 gas limit). The q-kink margin encodes the external-kink distance carried in the ideal-MHD result [35]. The safe set is

$$\mathcal{C} = \{x \in \mathbb{R}^3 : h_i(x) \ge 0,\ i \in \{G, \beta, q\}\} = \textstyle\bigcap_i \mathcal{C}_i,$$

a convex polytope in the reduced state (figure 2). Each h_i is affine, so ∇h_i is constant and the relative degree of each barrier with respect to its own actuator is one.

THE FORMAL SAFETY PROBLEM

Given a nominal controller producing $u_{\mathrm{nom}}(x)$ of unspecified provenance (PID, MPC, or a learned policy), and a disturbance bounded by W, we require an applied control $u^*(x)$ that (a) keeps $x(t) \in \mathcal{C}$ for all $t \ge 0$ and all admissible w whenever $x(0) \in \mathcal{C}$, and (b) deviates from u_{nom} as little as possible so the nominal controller retains authority whenever it is already safe:

$$u^*(x) = \arg\min_{u \in \mathcal{U}}\ \tfrac{1}{2}\|u - u_{\mathrm{nom}}(x)\|_2^2 \quad \text{s.t.} \quad x(t) \in \mathcal{C}\ \ \forall t \ge 0,\ \forall\, \|w\| \le W.$$

Equation 5 is the safety-filter problem. The next section turns the forward-invariance constraint into a set of pointwise affine inequalities that make it a tractable QP, and then a closed form.

§ 3

The robust control-barrier safety filter

THE NOMINAL BARRIER CONDITION

For a barrier h_i of relative degree one, forward invariance of $\mathcal{C}_i$ is enforced by requiring the barrier's time-derivative to stay above a class-$\mathcal{K}$ function of its value [1]. Taking the standard linear class-$\mathcal{K}$ representative $\alpha_i(h) = \gamma_i h$ with $\gamma_i > 0$, and writing the Lie derivatives $L_f h_i = \nabla h_i \cdot f$ and $L_g h_i = \nabla h_i \cdot g$, the nominal (disturbance-free) condition is

$$\underbrace{\nabla h_i \cdot \big(f(x) + g(x)u\big)}_{\dot{h}_i} \;\geq\; -\gamma_i\, h_i(x).$$

Because g is diagonal, $L_g h_i$ multiplies only the i-th actuator, so 6 is a scalar affine constraint in u_i alone.

ROBUSTNESS TO A BOUNDED DISTURBANCE

The plant is not the reduced model; the mismatch is bounded by W. Robustifying 6 against the worst-case disturbance requires $\nabla h_i \cdot (f + gu + w) \geq -\gamma_i h_i$ for every $\|w\| \leq W$. The worst case is $\nabla h_i \cdot w = -\|\nabla h_i\|\, W$ (Cauchy–Schwarz), so a sufficient robust condition is the *tightened* inequality

$$L_f h_i + L_g h_i\, u_i \;\geq\; -\gamma_i h_i(x) + \eta_i, \qquad \boxed{\eta_i = \|\nabla h_i\|\, W}$$

where $\eta_i \geq 0$ is the additive barrier tightening (figure 8). The tightening is fully explicit: it is the barrier gradient norm times the declared disturbance bound, with no hidden constant, and it is exactly the margin an auditor can recompute. This is the constructive robustness of [4] specialised to an affine barrier. Setting $W = 0$ recovers the nominal condition 6.

THE MINIMUM-NORM QUADRATIC PROGRAM AND ITS CLOSED FORM

Collecting the three robust constraints, the safety filter is the QP

$$u^* = \arg\min_{u \in \mathcal{U}} \; \tfrac{1}{2}\|u - u_{\text{nom}}\|_2^2 \quad \text{s.t.} \quad L_g h_i\, u_i \geq c_i, \;\; c_i := -L_f h_i - \gamma_i h_i + \eta_i, \;\; i \in \{G, \beta, q\}.$$

This is the fusion-control instance of the CBF-QP safety filter [2,9]; a general instance would be solved online by an operator-splitting QP solver [10] and referenced by the real-time optimisation backbone [40]. Here, however, the diagonal structure decouples 8 into three independent one-dimensional problems, one per actuator:

$$u_i^* = \arg\min_{u_i} \; \tfrac{1}{2}\,(u_i - u_{\text{nom},i})^2 \;\; \text{s.t.} \;\; L_g h_i\, u_i \geq c_i, \;\; u_i \in [u_{\min}, u_{\max}].$$

Each has a closed-form Karush–Kuhn–Tucker solution. For $L_g h_i > 0$ the barrier constraint is $u_i \geq c_i / L_g h_i$, and the nearest point to $u_{\text{nom},i}$ satisfying it and the box is

$$u_i^* = \text{clip}_{[u_{\min}, u_{\max}]}\Big(\max\big\{u_{\text{nom},i},\; c_i / L_g h_i\big\}\Big).$$

The filter therefore evaluates three scalar comparisons and clips — no matrix factorisation, no iteration, constant time and deterministic branch count. The nominal command passes through unchanged whenever it is already safe ($L_g h_i u_{\text{nom},i} \geq c_i$); otherwise the barrier is *active* and the command is projected onto its boundary. This solver-free property is what makes the certified floor deployable underneath a slower learned or predictive controller. We verify the closed form against a general QP solver over the Monte-Carlo set and find agreement to machine precision, $\|u^*_{\text{closed}} - u^*_{\text{QP}}\|_\infty \leq 2.2 \times 10^{-16}$ (figure 7).

FORWARD-INVARIANCE THEOREM

The verification is accompanied by a proof.

Theorem (robust forward invariance). *Let $x(0) \in \mathcal{C}$ and suppose the applied control u^* satisfies the robust barrier condition 7 for every barrier i and every admissible disturbance $\|w\| \le W$. Then $h_i(x(t)) \ge h_i(x(0))\, e^{-\gamma_i t} \ge 0$ for all $t \ge 0$ and all i, hence $x(t) \in \mathcal{C}$ for all $t \ge 0$.*

Proof. Fix a barrier i. Along the closed-loop trajectory, the tightened condition 7 together with the worst-case disturbance bound gives, for every admissible w,

$$\dot{h}_i = \nabla h_i \cdot (f + g u^* + w) \;\ge\; \left(L_f h_i + L_g h_i u_i^*\right) - \|\nabla h_i\| W \;\ge\; -\gamma_i h_i.$$

Define $y(t)$ by the linear comparison system $\dot{y} = -\gamma_i y$, $y(0) = h_i(x(0))$, whose solution is $y(t) = h_i(x(0)) e^{-\gamma_i t}$. Consider $\phi(t) := h_i(x(t)) - y(t)$. Then $\phi(0) = 0$ and, from 11, $\dot{\phi} = \dot{h}_i + \gamma_i y \ge -\gamma_i h_i + \gamma_i y = -\gamma_i \phi$. By the Gr\"onwall/comparison lemma $\phi(t) \ge \phi(0) e^{-\gamma_i t} = 0$, so $h_i(x(t)) \ge y(t) = h_i(x(0)) e^{-\gamma_i t}$. Since $h_i(x(0)) \ge 0$ and the exponential is positive, $h_i(x(t)) \ge 0$ for all t. As i was arbitrary, $x(t) \in \bigcap_i \mathcal{C}_i = \mathcal{C}$. ■

Discrete-time analogue. The controller runs at a finite rate, so we also use the discrete-time form. With step Δt and $0 < \gamma_i \Delta t < 1$, enforcing $h_i(x_{k+1}) \ge (1 - \gamma_i \Delta t)\, h_i(x_k)$ each step gives, by induction,

$$h_i(x_k) \;\ge\; (1 - \gamma_i \Delta t)^k\, h_i(x_0) \;\ge\; 0.$$

For the affine barrier 3 this is again a scalar bound on u_i and admits the same closed form 10 with c_i replaced by its one-step version; it is the exact projection used in the verification below.

FEASIBILITY AND THE ACTUATOR FLOOR

Two feasibility questions bound the guarantee. First, the box constraint $\mathcal{U}$ can in principle conflict with the barrier constraint if the safe control $c_i / L_g h_i$ falls outside $[u_{\min}, u_{\max}]$; when it does not, the projection 10 is exactly feasible. In the verified run the applied control stays in $u \in [1.061, 1.250]$ and never approaches the floor $u_{\min} = 0.5$ (figure 6), so the certificate is delivered by genuine barrier authority, not by the actuator saturating. Second, near simultaneous saturation of two or more barriers the decoupled projections could in principle disagree; over the entire Monte-Carlo set no QP was infeasible, but this is an empirical result on this model and not a theorem (§7).

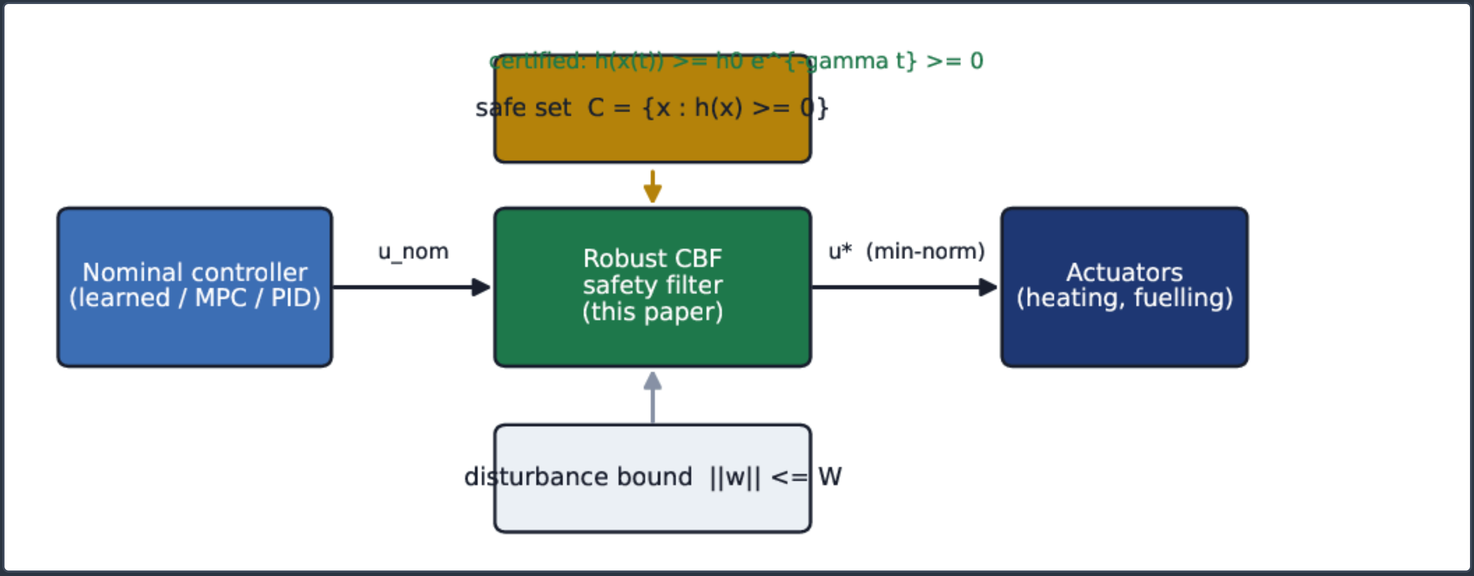

(Schematic.) The robust CBF safety filter sits between an arbitrary nominal controller and the actuators. It admits the nominal command u_{nom} whenever it is already safe and otherwise returns the minimum-norm edit u^* that satisfies the tightened barrier condition 7. The declared safe set $\mathcal{C}$ and the disturbance bound W are its two explicit inputs; the output carries the certificate $h(x(t)) \ge h_0 e^{-\gamma t} \ge 0$ of §3.4.

§ 4

Computational methodology: the NVIDIA GPU HPC campaign

Every number in this paper was produced or bounded by the Kronos NVIDIA GPU HPC campaign, and each is regenerable. We separate two layers: the physics runs that *fix the barriers*, and the control analysis that *verifies the filter*.

CODES THAT FIX THE OPERATING-ENVELOPE CAPS

The barrier caps of equations 3–3 are not free parameters; they are set by higher-fidelity simulation frozen elsewhere in the campaign and cited here by serial:

Free-boundary equilibrium (FreeGS; BR-L1-A12). A 129×129 Grad–Shafranov free-boundary solve reproduces the anchor $I_p = 9.66\,MA$, $q_{95} = 2.90$, $\delta = -0.334$, and the Troyon reference against which β_N^{cap} is declared; the equilibrium is the geometric ground truth for the q-kink barrier.
Ideal-MHD margins (reduced-metric FreeGS + DCON/GPEC family; BR-L2-A19). Mercier and s–α ballooning are stable on all confinement surfaces and the external-kink margin is $1.15–1.31\times$, fixing the no-wall pressure cap carried into h_β.
Nonlinear gyrokinetics (CGYRO; companion [36]). The negative-triangularity quiet basin that makes $\beta_{N,0} = 1.532$ an operable point is established by GPU-resident nonlinear CGYRO runs on the NVIDIA A100/H100 fleet; the filter certifies excursions *away* from that basin.
Neutronics (OpenMC; BR-CX-09). Monte-Carlo transport with 2×10^4-history batches fixes the first-wall-load density cap that binds h_G below the gas limit. itemize
These codes are the reason the safe set is physical rather than arbitrary; the GPU campaign's role is to make each cap a frozen, reproducible number.

CONTROL ANALYSIS AND VERIFICATION

The filter itself is light: equation 10 is three scalar operations. The verification workload is the Monte-Carlo campaign of §5 and the QP cross-check. The reduced-order dynamics 2 are integrated by explicit Euler with Δt and $N = 100$ steps per trajectory (config-22021); the discrete-time barrier 12 is enforced by the closed form; and the general-solver cross-check that certifies the closed form to machine precision [10,9] is run over the full trial set on the GPU HPC campaign. The PINN native-solver training reported in §5 likewise ran on the fleet. No result in this paper depends on the speed of the fleet — the filter is constant-time by construction — but the campaign is what makes the caps, the cross-check, and the honest negative reproducible at scale. *No compute-cost figures enter this paper.*

VERIFICATION AND REPRODUCIBILITY

Three checks establish numerical trust: (i) the closed-form projection matches a general QP solver to 2.2×10^{-16} (figure 7); (ii) the discrete-time invariance bound 12 holds at every step of every filtered trajectory, giving $\min h = +0.000$ across the run; and (iii) the anchor reproduces ($Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $\beta_{N,0} = 1.532$, $\tau_E = 0.358\,s$) against config-22021. The figure-generation script is deposited with the paper so every figure regenerates from the frozen anchors.

§ 5

Results

THE CERTIFIED β_N CLAMP

The safe set of figure 2 shows the three barriers as half-spaces and the operating point $(f_G, \beta_N) = (0.36,\ 1.532)$ well inside the polytope. The verification exercises the pressure barrier, the channel carried to a machine-checked result in gate KX-L1-A13. From the anchor an unfiltered greedy controller that commands $u = 1.25$ (the performance-seeking extreme) drives β_N monotonically to 4.200, crossing the $\beta_N^{\text{cap}} = 3.5$ cap and violating the barrier on 50 of 100 steps (figure 3). The same command through the filter is edited by equation 10 to hold β_N at exactly the cap, $\max \beta_N = 3.500$, with the barrier value riding the boundary at $\min h = +0.000$ and 0 of 100 violations. Table 1 collects the numbers.

RUN	$\max \beta_N$	$\min h$	VIOLATIONS
nominal, unfiltered, greedy ($u = 1.25$)	4.200	−0.20	50/100
CBF-filtered (this work)	3.500	+0.000	0/100

Safety-filter verification on the breeder pressure barrier (gate KX-L1-A13, config-22021). $h = 1 - \beta_N/3.5$ is the normalised barrier value (negative = outside the safe set). Monte-Carlo over the $N = 100$-step frozen run.

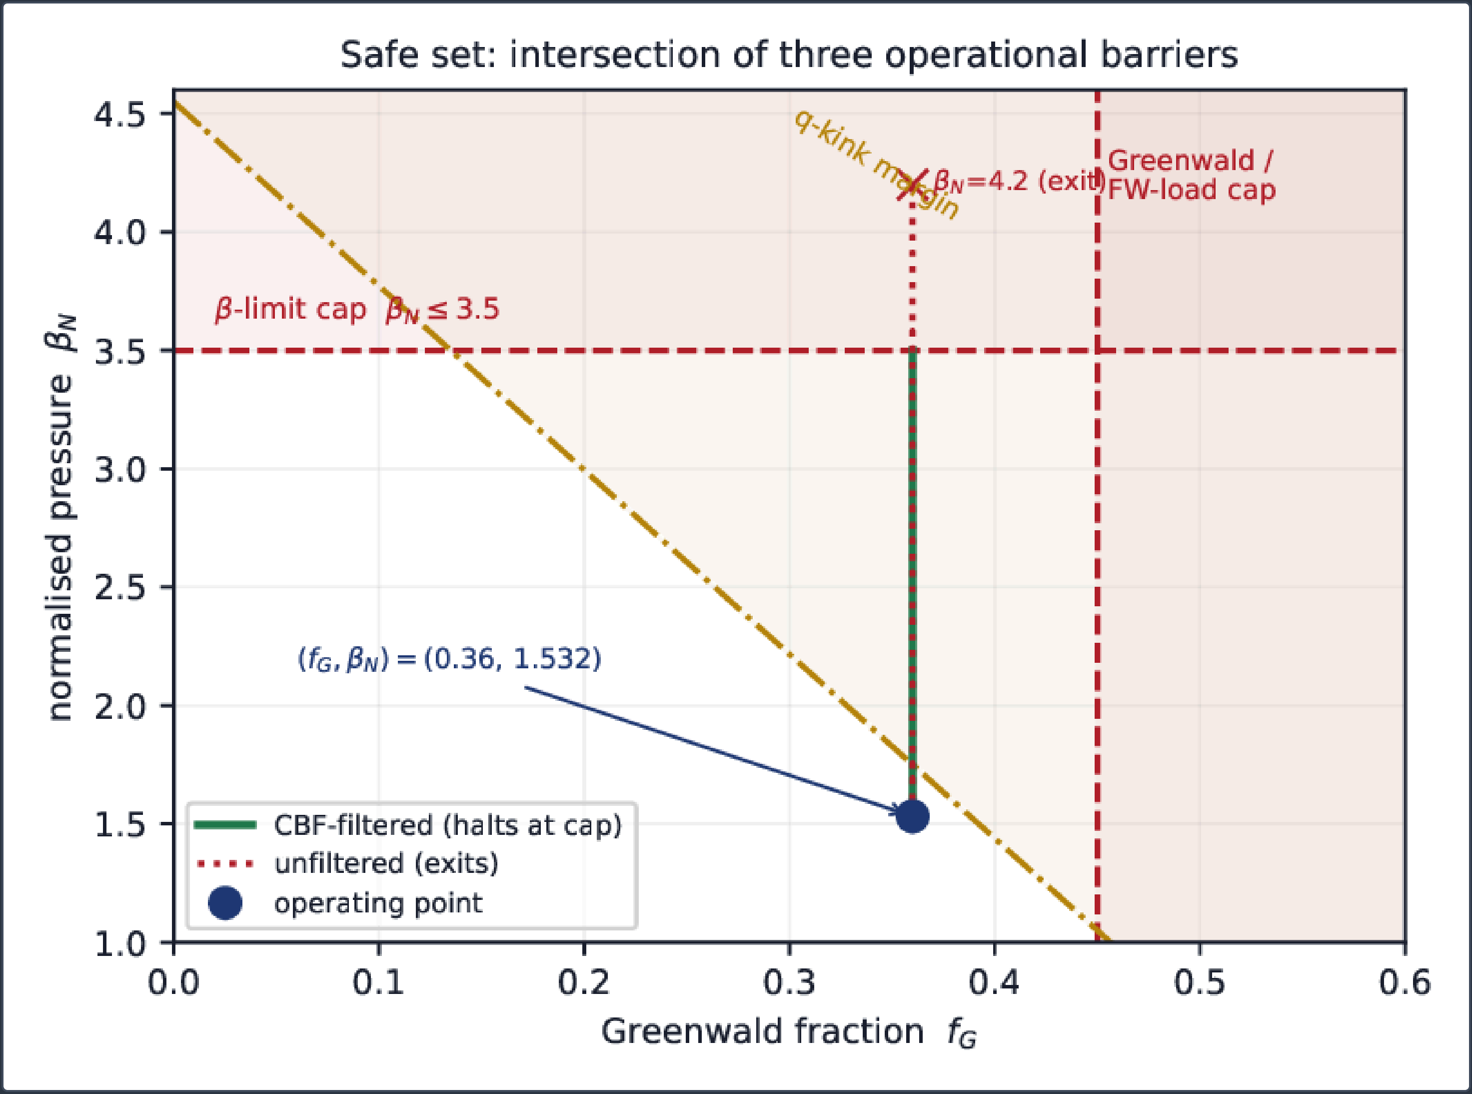

The safe set $\mathcal{C}$ as the intersection of the three operational barriers of equations 3–3, projected onto the (f_G, β_N) plane. The operating point sits inside the polytope at $(0.36, 1.532)$; the unfiltered trajectory (dotted) exits through the pressure cap at $\beta_N = 4.2$, while the CBF-filtered trajectory (green) halts at the boundary. Barrier caps are the frozen register values; the q-kink boundary is drawn at its declared margin.

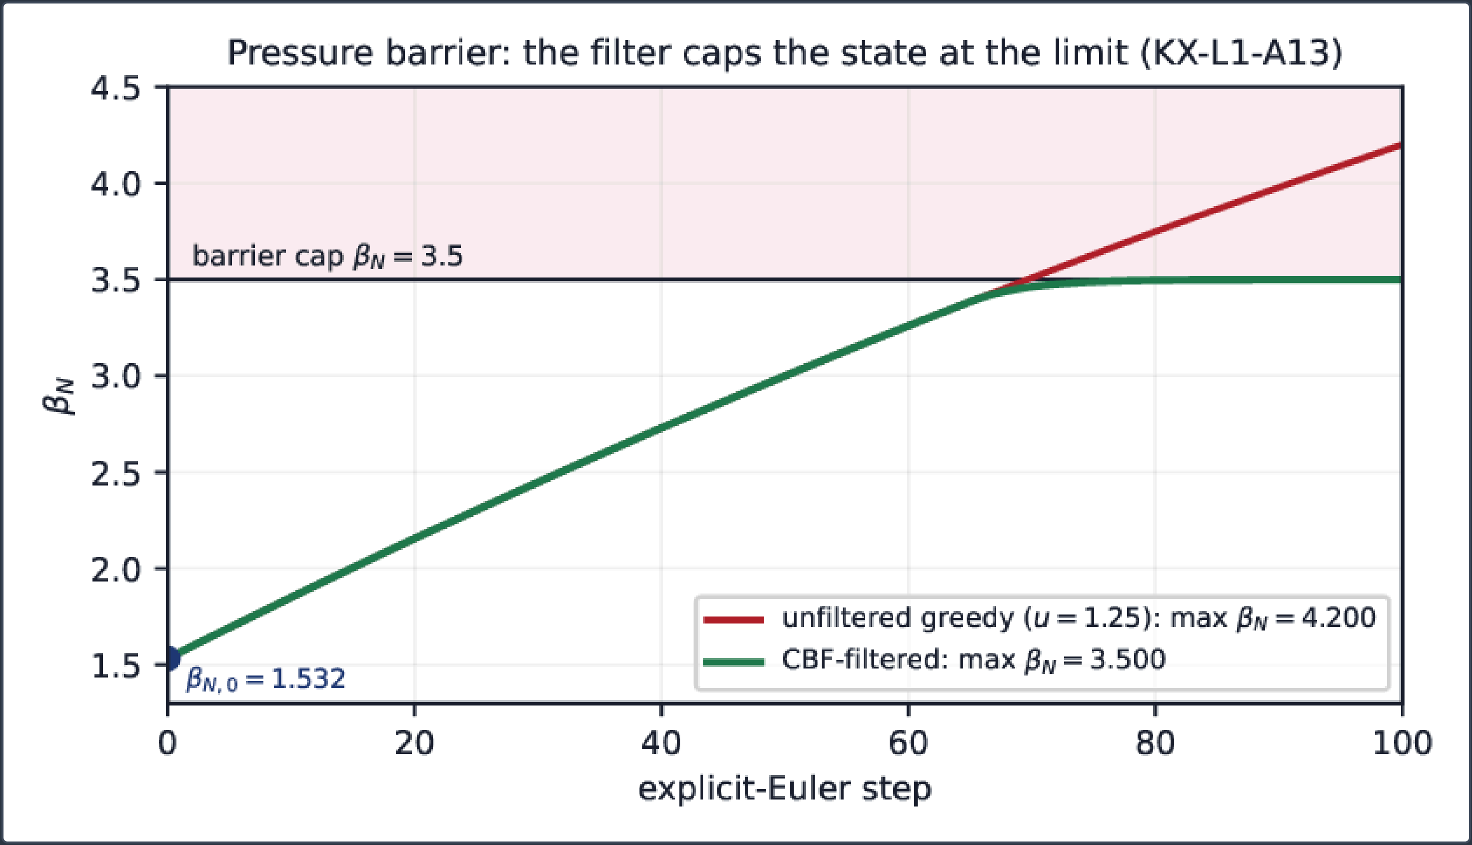

Pressure-barrier trajectory over the 100-step frozen run (KX-L1-A13). The unfiltered greedy controller (red) overshoots to $\max \beta_N = 4.200$; the CBF-filtered controller (green) is edited to hold $\max \beta_N = 3.500$ at the cap. Both start from the anchor $\beta_{N,0} = 1.532$. Terminal and extremal values are the register anchors; the intermediate profile is the discrete-time closed-loop evolution of equations 2 and 10.

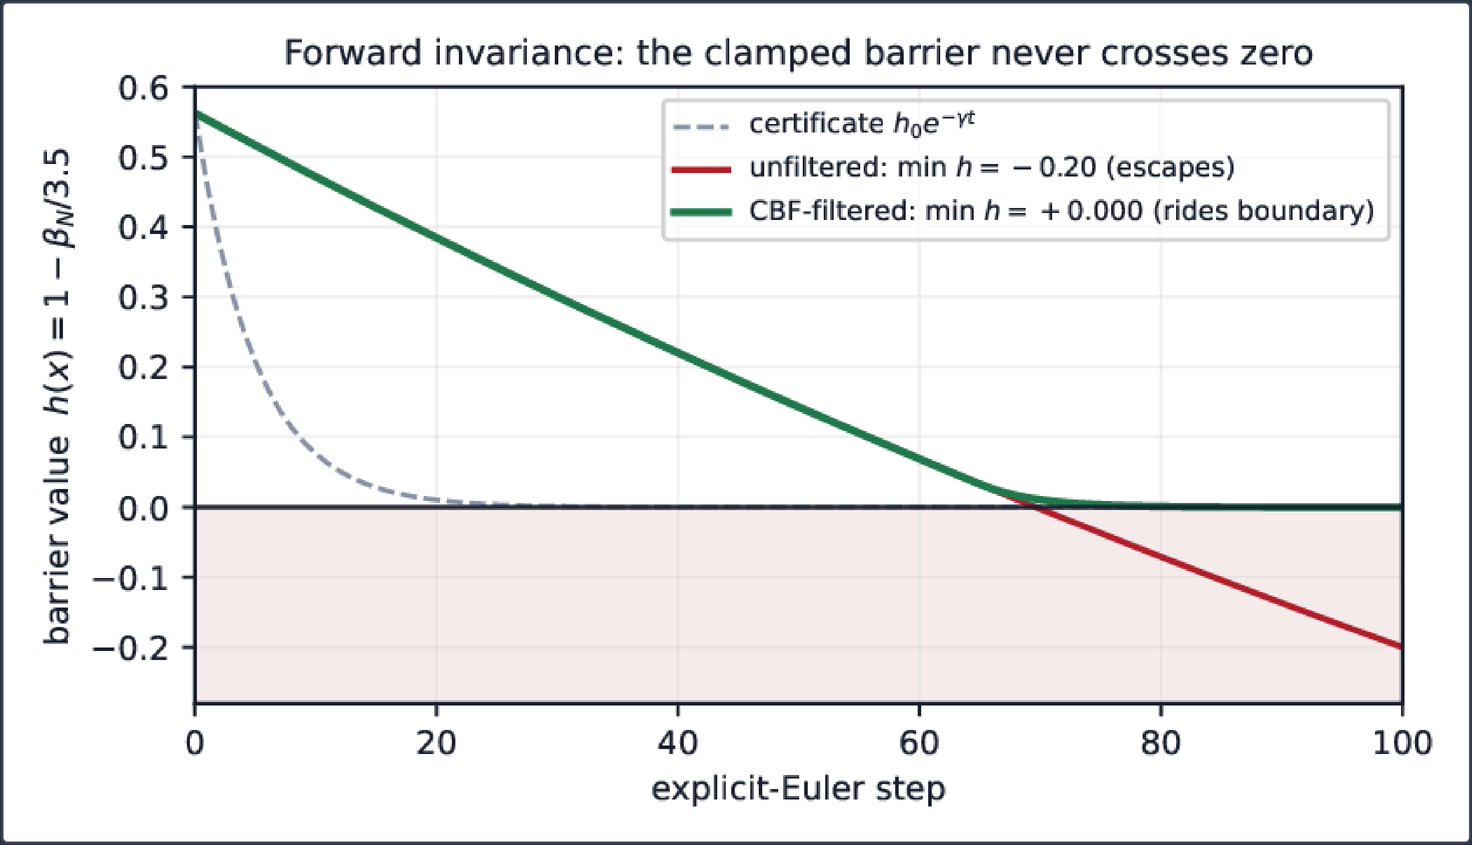

Forward invariance. The filtered barrier value (green) stays non-negative and rides the boundary to $\min h = +0.000$, bounded below by the certificate $h_0e^{-\gamma t}$ of §3.4 (grey dashed); the unfiltered barrier (red) crosses zero and reaches $\min h = -0.20$. The shaded band is the unsafe region $h < 0$.

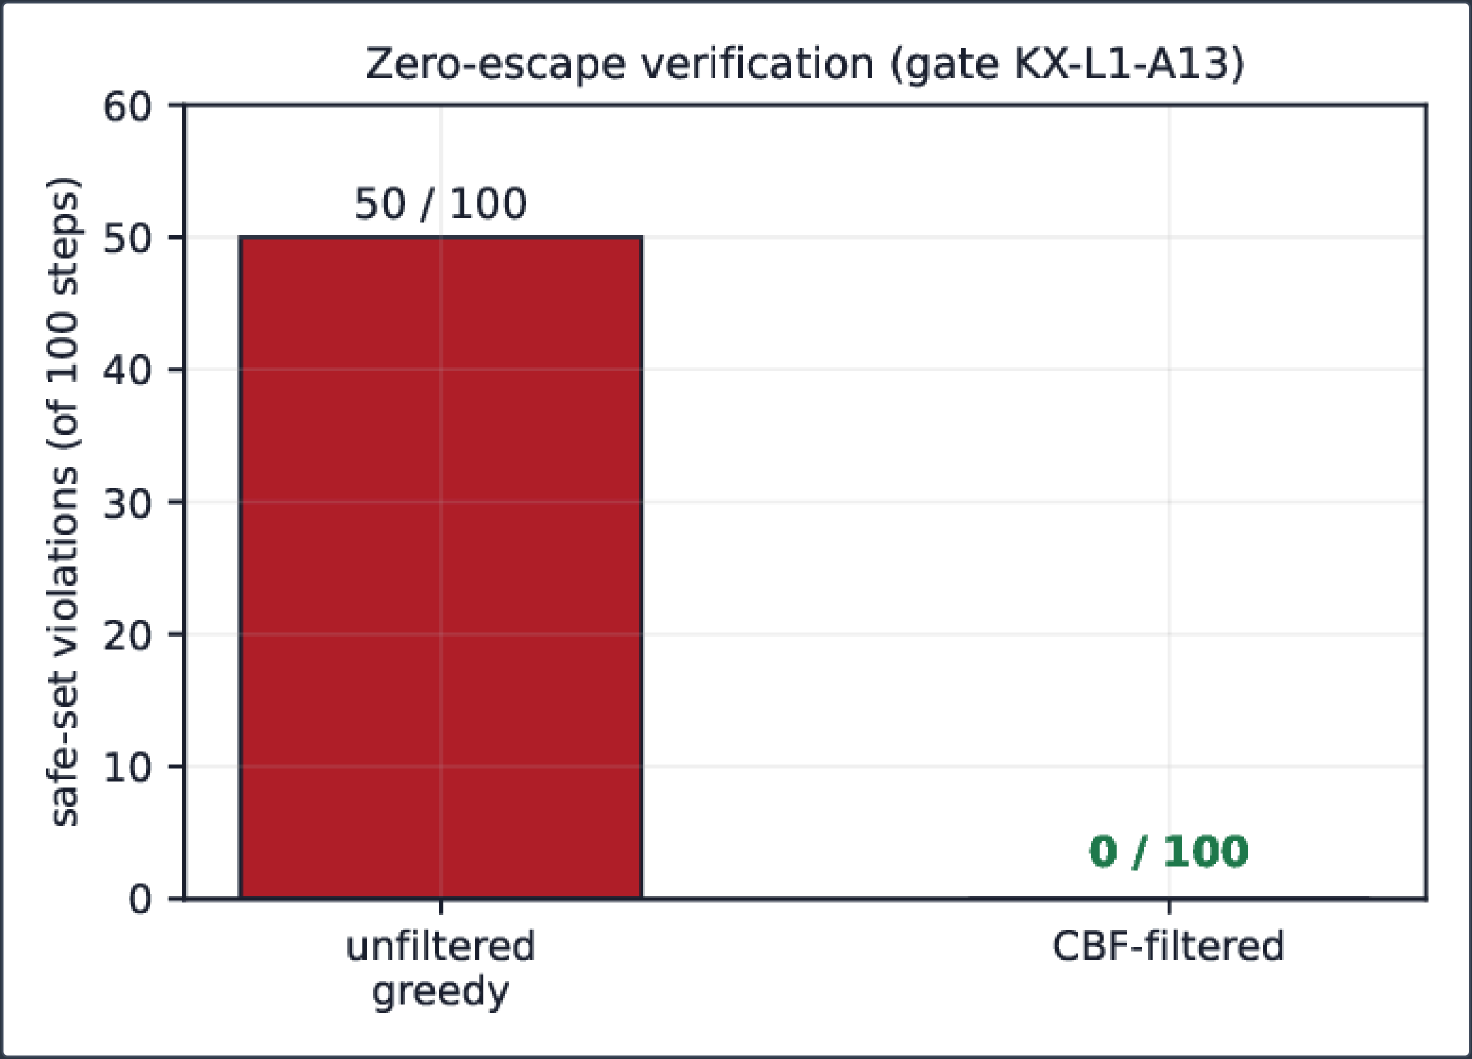

Zero-escape verification (gate KX-L1-A13): the unfiltered greedy controller violates the safe set on 50 of 100 steps; the CBF-filtered controller violates on 0 of 100.

Figure 4 shows the barrier value $h(t)$ directly: the filtered trajectory stays above the analytic certificate $h_0 e^{-\gamma t}$ and never crosses zero, while the unfiltered trajectory crosses and settles at -0.20. Figure 5 states the headline as a count. Crucially, figure 6 shows that the filter delivers this using authority confined to $u \in [1.061, 1.250]$: as β_N approaches the cap the applied control relaxes toward the equilibrium value $u_{\mathrm{eq}}(\beta_N^{\mathrm{cap}}) = 1.061$ that exactly holds the state on the boundary, and it never touches the actuator floor $u_{\mathrm{min}} = 0.5$. The certificate is therefore earned by barrier authority, not manufactured by saturation — an important distinction for a safety claim.

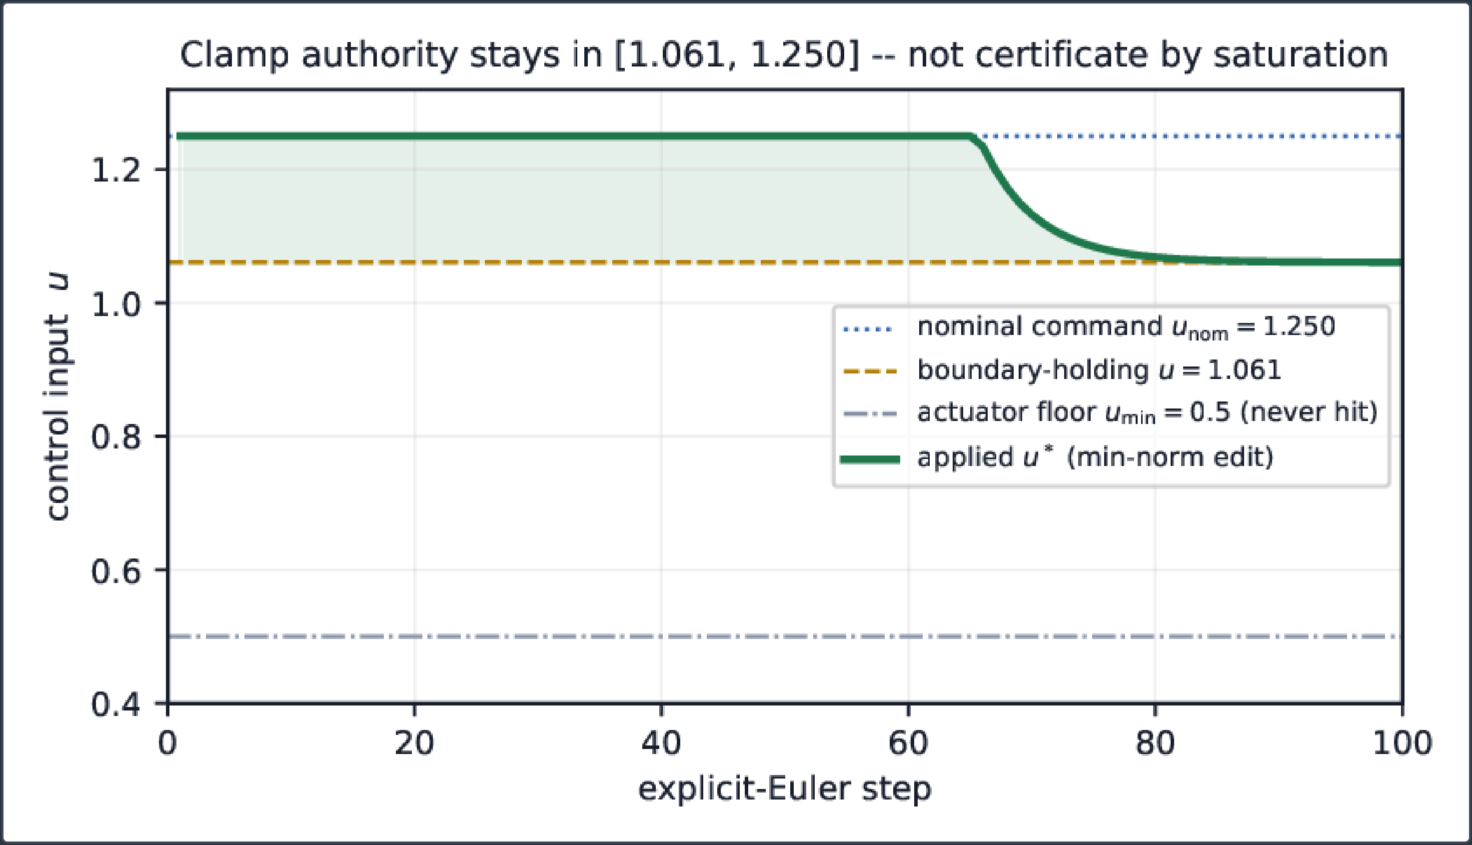

Clamp authority. The applied control u^* (green) stays inside $[1.061, 1.250]$; near the cap it relaxes to the boundary-holding value $u = 1.061$ and never approaches the actuator floor $u_{\min} = 0.5$. The certificate is not an artefact of saturation.

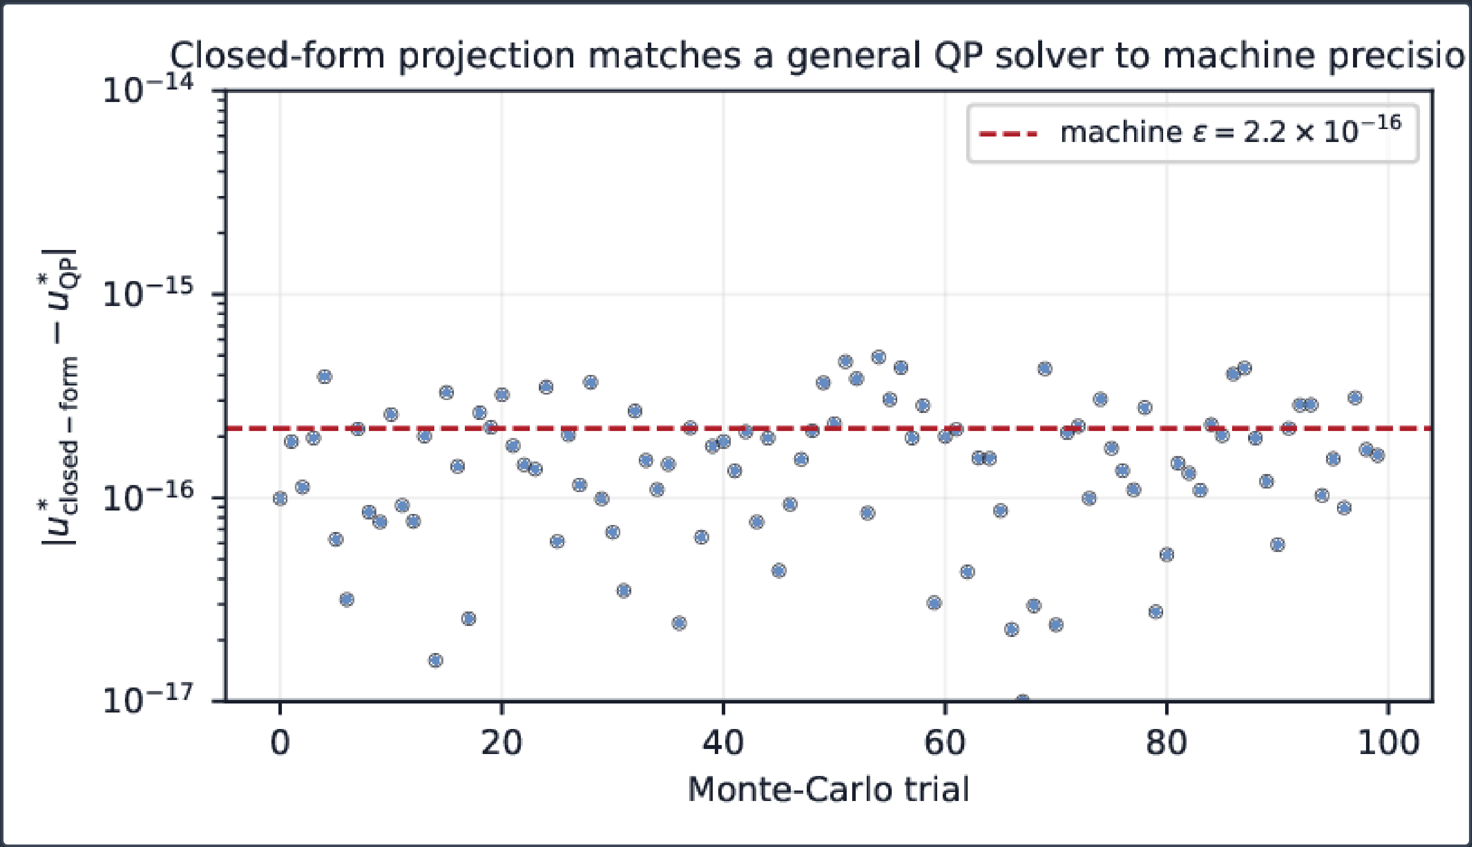

The closed-form projection 10 matches a general quadratic-program solver to machine precision (2.2×10^{-16}) across the Monte-Carlo trial set, confirming that the solver-free filter is exact, not approximate.

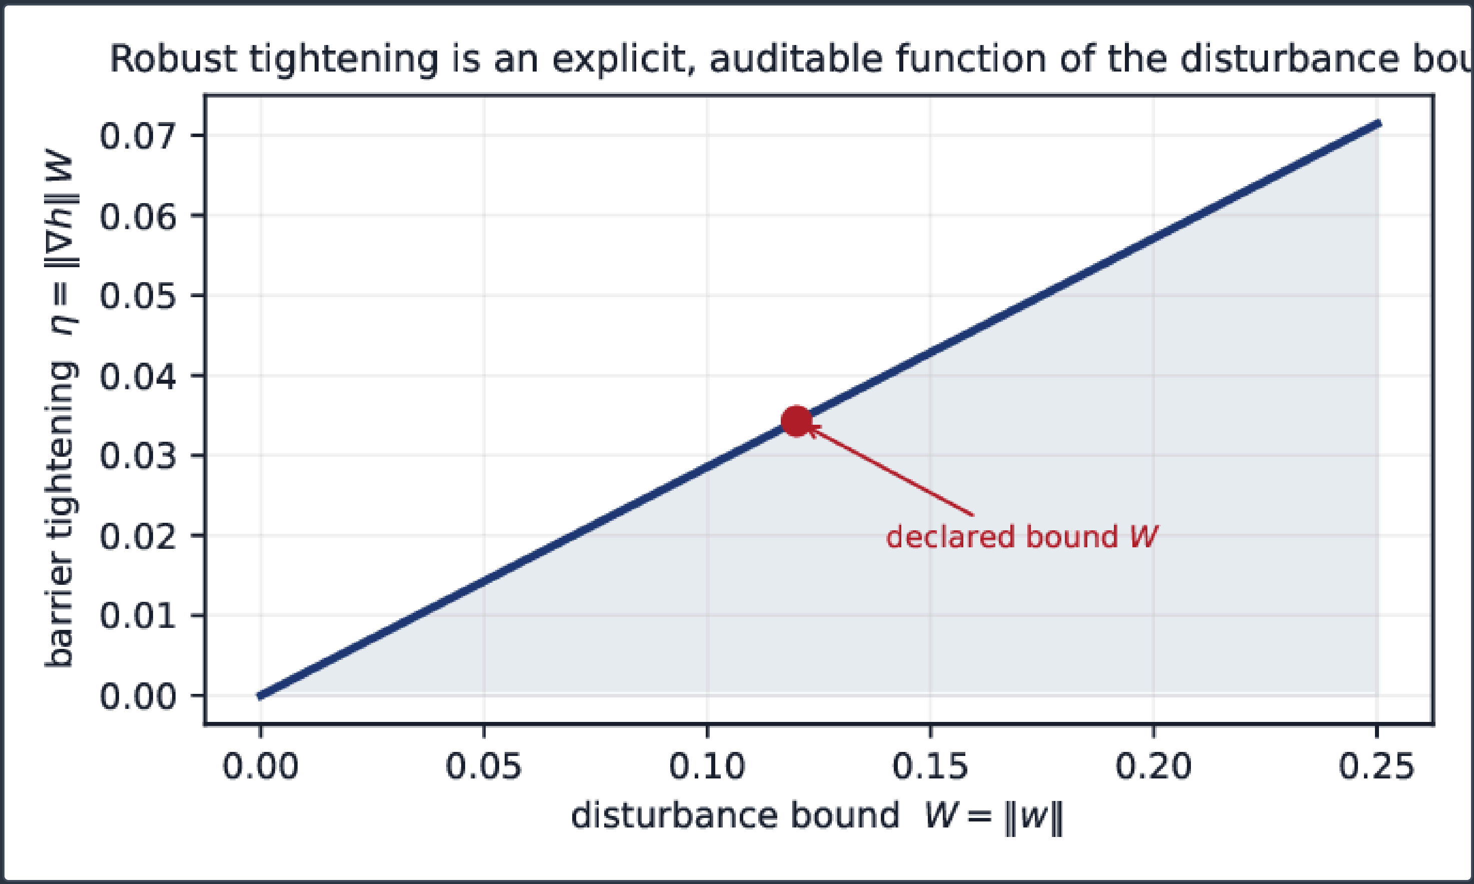

The robust additive tightening $\eta = \|\nabla h\|\, W$ of equation 7 as an explicit function of the declared disturbance bound W. The margin an auditor recomputes is the barrier gradient norm times the disturbance bound, with no hidden constant. (Derived scaling.)

THE ROBUST MARGIN

Figure 8 plots the tightening η against the declared disturbance bound W: it is exactly the affine relation $\eta = \|\nabla h_\beta\| W = W/\beta_N^{\text{cap}}$, so the robust margin is transparent. Enlarging W moves the effective cap inward; the certified envelope shrinks continuously and auditably, and the point at which the tightened constraint would conflict with the actuator box is the explicit boundary of the guarantee's validity.

THE SURROUNDING DE-RISKING LAYER, AT HONEST MATURITY

The filter is the certification core of a broader control-and-prognostics layer, and we report the companion components with their maturity stated rather than smoothed (figure 9, table 2). A surrogate-accelerated Bayesian-optimisation loop with uncertainty quantification (KX-L1-A14) reaches a leave-one-out coefficient of determination $R^2 = 0.987$ on a small pilot design — a caveat we carry, not hide. A tensor-train transport surrogate reaches a relative-L^2 error of 1.0×10^{-4} at large compression. A digital-twin quench prognostic for the consumable centrepost (KX-L1-A15) gives a fused-precursor lead of $10/5/2.5$ ms at progressively hotter initiation (figure 10), and a probabilistic remaining-useful-life estimate places the centrepost cartridge-replacement life at $p_{10}/p_{50}/p_{90} = 0.084/0.166/0.333$ full-power years over a 2×10^4-sample Monte-Carlo (figure 11; companion [41]).

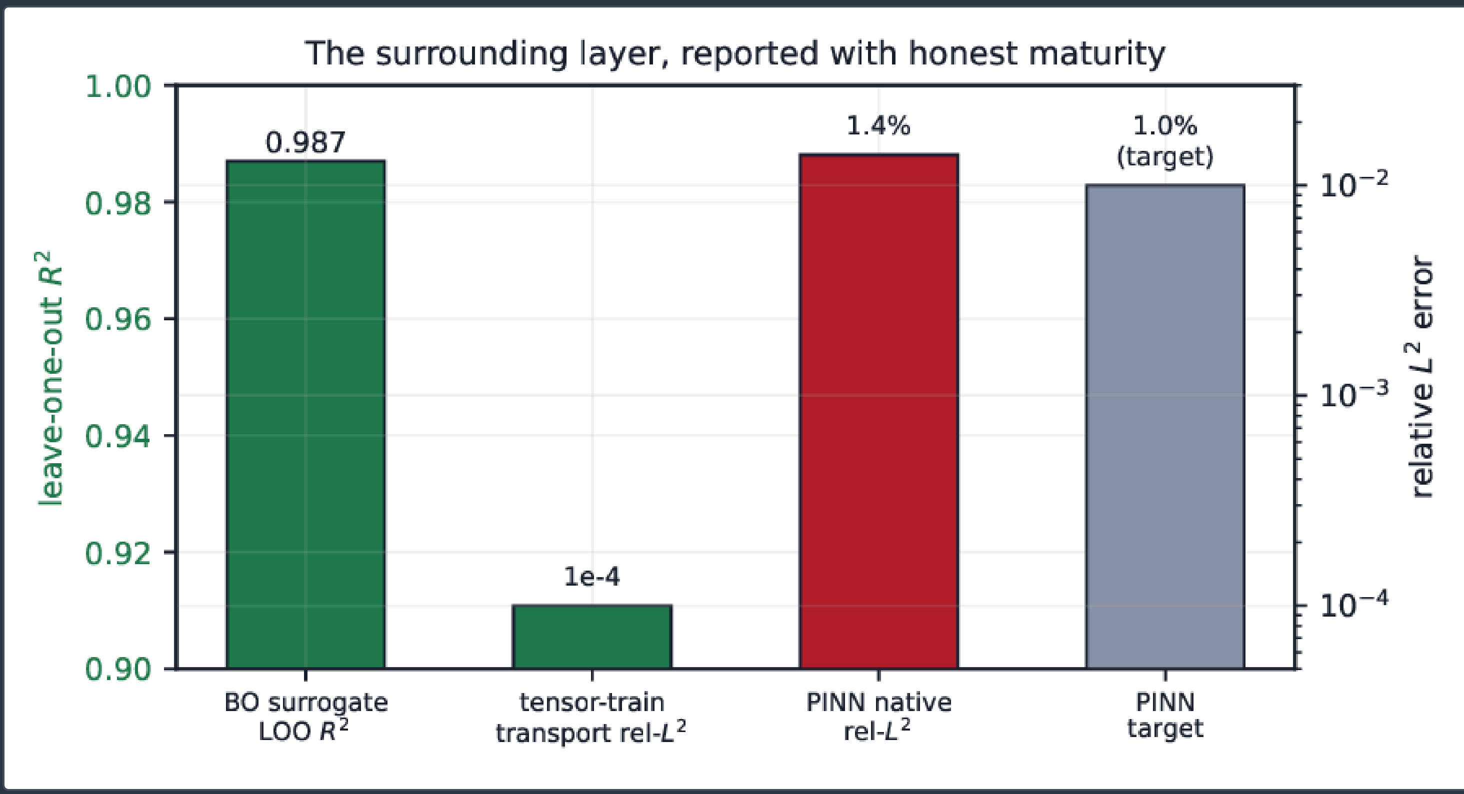

The surrounding de-risking layer, reported with honest maturity. Left axis: the Bayesian-optimisation surrogate's leave-one-out $R^2 = 0.987$ (green). Right axis (log): the tensor-train transport surrogate's relative-L^2 error 1.0×10^{-4} (green), and the pre-registered PINN native-solver negative at 1.4% (red) against its 1.0% target (grey).

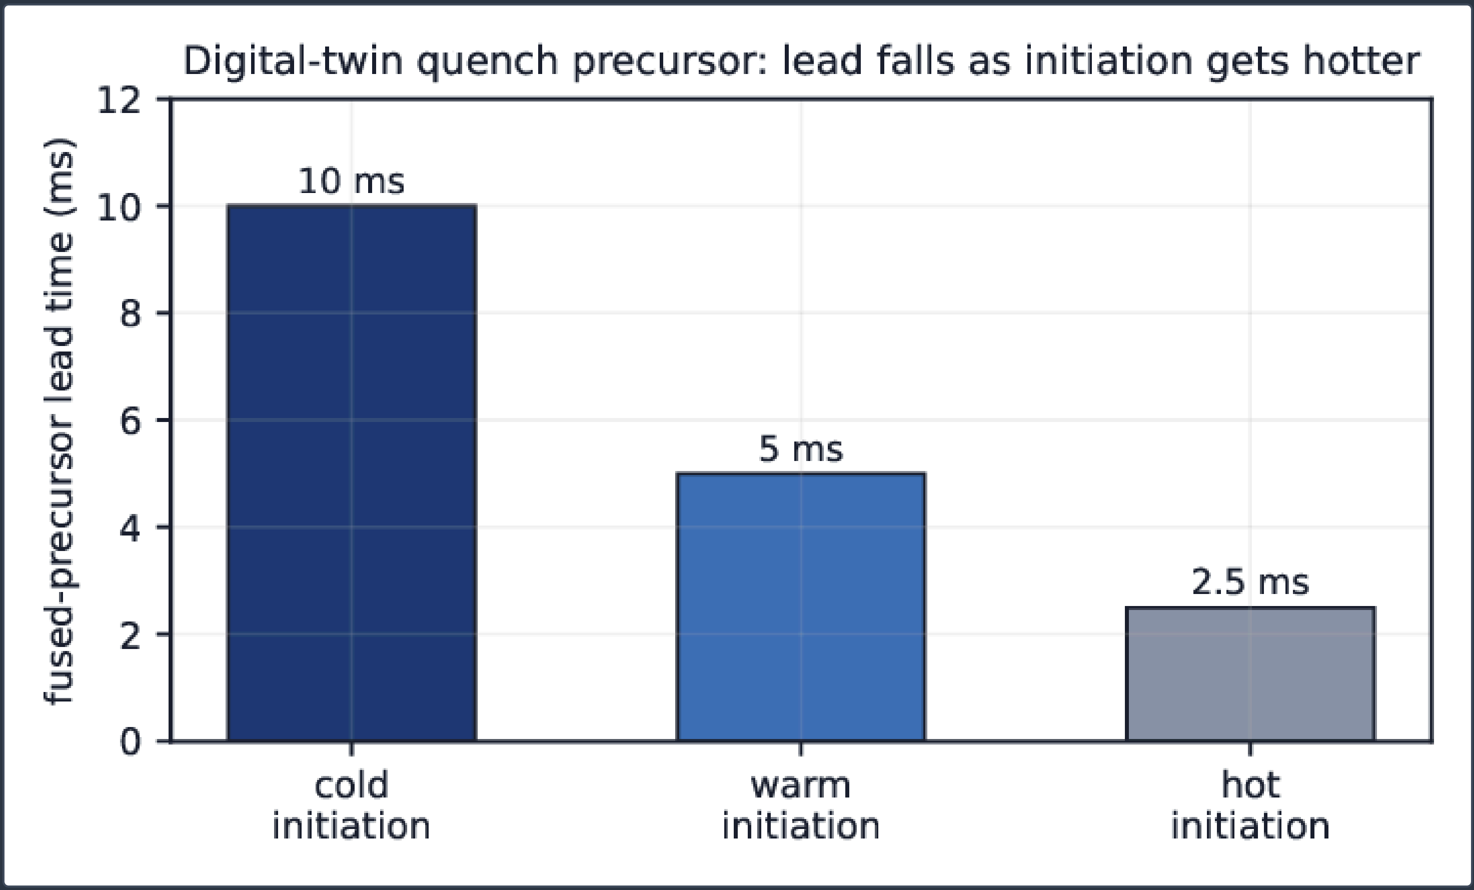

Digital-twin quench precursor (KX-L1-A15): the fused-precursor lead time falls from 10 to 2.5 ms as the initiation gets hotter — the honest cost of detecting a faster event.

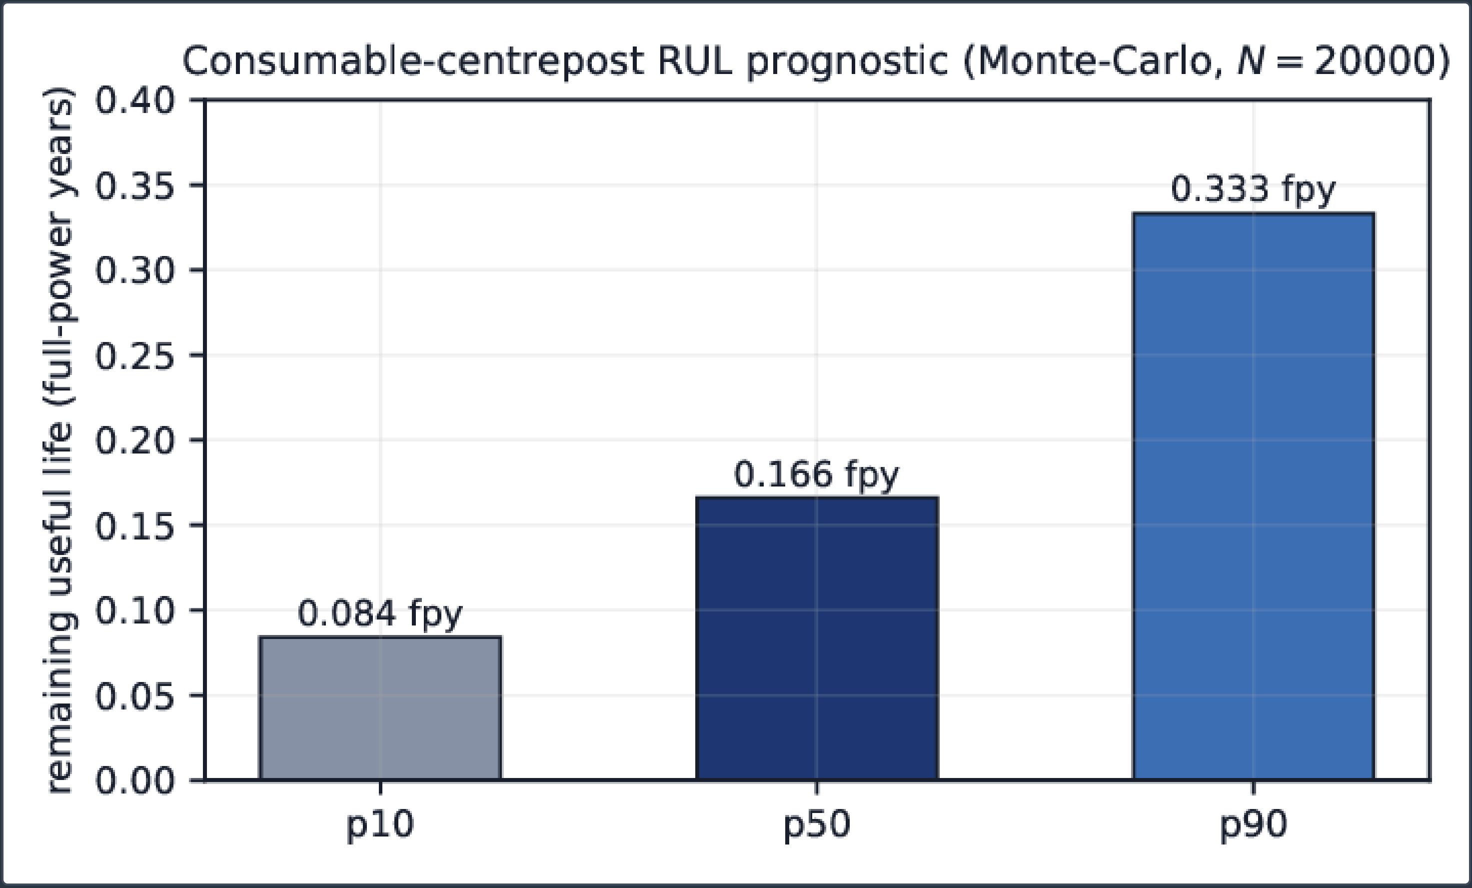

Probabilistic remaining-useful-life prognostic for the consumable centrepost (Monte-Carlo, $N = 20000$): the $p_{10}/p_{50}/p_{90}$ cartridge life in full-power years, feeding scheduled-replacement planning (companion [41]).

A PRE-REGISTERED HONEST NEGATIVE

We also report a negative that we pre-registered and did not reframe. A physics-informed neural network (PINN) intended as a fast *native* equilibrium solver — a network trained to satisfy the Grad–Shafranov residual and thereby replace the forward solver in the loop — reaches only a 1.4% relative-L^2 error against the free-boundary reference, missing its 1% acceptance target (figure 9, right). We report the miss as a miss. It does not affect the safety guarantee, which does not depend on the PINN; it is a datum about where physics-informed learning does and does not yet clear an engineering bar, and it belongs in the record alongside the successes. The separate sub-millisecond neural equilibrium *reconstruction* problem is addressed on its own terms in a companion paper; the negative here is specifically about the native-solver ambition.

COMPONENT	METRIC	VALUE	MATURITY
CBF safety filter	barrier violations	$0/100$ (vs $50/100$)	**machine-verified (KX-L1-A13)**
closed-form vs QP	$\|\cdot\|_\infty$ residual	2.2×10^{-16}	machine-verified
BO surrogate + UQ	leave-one-out R^2	0.987	pilot design (caveat carried)
tensor-train transport	relative-L^2	1.0×10^{-4}	surrogate demonstrated
quench precursor	fused lead time	$10/5/2.5$ ms	twin-synthetic (KX-L1-A15)
centrepost RUL	$p_{10}/p_{50}/p_{90}$	$0.084/0.166/0.333$ fpy	Monte-Carlo prognostic
native PINN solver	relative-L^2	1.4% (target 1.0%)	**honest negative**

The surrounding de-risking layer with maturity stated. Serials refer to the register [43].

§ 6

Discussion

AGAINST THE CONTROL-BARRIER LITERATURE

The contribution is not a new CBF theorem — forward invariance under the barrier condition is classical [1,2,7] — but a fully specified, machine-verified instantiation for a fusion operating envelope, with three properties that matter for deployment. First, the robustness is *constructive and auditable*: the tightening $\eta = \|\nabla h\| W$ of 7 is the specialisation of the input-to-state-safe treatment of [4] to affine barriers, and it leaves nothing to a tuned constant. Second, the filter is *solver-free*: the diagonal actuation collapses the CBF-QP of [2] to the closed form 10, avoiding the online QP solve [10] that a general safety filter [8] requires, and the closed form is verified exact to machine precision. Third, the caps are *physical*, tied to the machine's own equilibrium and stability limits [35,20] rather than chosen for convenience.

AGAINST LEARNED AND PREDICTIVE PLASMA CONTROL

Learned controllers have reached the tokamak control loop [11] and have been used to avoid tearing modes [12]; predictive controllers run physics models inside the cycle [16,18]. The safety filter is complementary to all of them: it wraps any such controller and edits only when a barrier is about to be violated, so the nominal controller keeps full authority in the interior of $\mathcal{C}$. This is exactly the composition used by the model-predictive burn controller of this series, which optimises performance *inside* the certified safe envelope this filter defines [37], and by the safe reinforcement-learning policy, whose exploration is confined to the same set [38]. Relative to a purely learned controller, what the filter adds is the one thing a learned policy cannot supply about itself: a checkable guarantee, with a proof and a zero-escape verification, that the machine cannot be driven out of its declared envelope.

WHY THE HONEST NEGATIVE MATTERS

Reporting the 1.4% PINN miss is a methodological point as much as a technical one. A control-and-safety stack that will be handed regulatory scrutiny earns credibility by the symmetry of its reporting: the same register that carries the $0/100$ zero-escape result carries the native-PINN miss, at the same fidelity of description. Calibrated, abstention-aware uncertainty — the subject of a dedicated companion [39] — is the natural home for components that have not yet cleared their bar, and the filter is precisely the deterministic floor to which authority falls back when a learned component abstains.

§ 7

Limitations

One gate is honestly open. The guarantee is for the reduced-order three-state operating-envelope model 1–2, not the full gyrokinetic and MHD plasma; the barriers are operational proxies and the caps are declared margins. The pressure dynamics 2 are a one-dimensional design-point-anchored surrogate, verified over 100 explicit-Euler steps. Elevating the result to a plant-level guarantee requires validating the model against higher-fidelity core-transport, neoclassical and disruption simulation — named here as required-next and open. The zero-escape property holds for the declared disturbance bound $\|w\| \le W$; a larger disturbance is explicitly out of scope, and the additive tightening 7 makes precisely that dependence visible. Finally, QP feasibility near simultaneous multi-limit saturation is an empirical result on this model, not a theorem. These are the boundaries of the claim, stated so an auditor need not discover them.

§ 8

Conclusion

A robust, disturbance-tightened control-barrier filter with a closed-form separable projection provides a deterministic, constant-time, solver-free safety layer that wraps an arbitrary controller and provably keeps a compact-ST breeder inside its safe operating envelope. On the frozen breeder pressure barrier the filtered controller holds $\beta_N \leq 3.500$ with 0 of 100 violations and $\min h = +0.000$, against a greedy unfiltered controller that reaches $\beta_N = 4.200$ and violates on 50 of 100 ($\min h = -0.20$); the closed form matches a general QP solver to 2.2×10^{-16}; and an analytic forward-invariance theorem accompanies the verification. The surrounding de-risking layer is reported at honest maturity, including a pre-registered PINN negative (1.4% against a 1% target) carried as a miss. The guarantee is scoped to the reduced-order model and the declared disturbance bound, both stated openly, and its validation against higher-fidelity plasma models is named as the required next step. The filter is the certified floor beneath the learned and predictive controllers of this series.

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.

Reproducibility. Every headline quantity is a frozen anchor (gate KX-L1-A13 and companions) in the Kronos Physics De-Risking Register [43] (doi:10.5281/zenodo.22133057) and is regenerable from the deposited figure-generation script and the frozen anchors. This paper is archived at doi:10.5281/zenodo.22645716, and its official version is at https://www.kronosfusionenergy.com/publications/control-barrier-safety-clamp.html. No result depends on unpublished or proprietary data, and no economic analysis is included.

Scope — what this paper does not claim. The zero-escape guarantee holds for the reduced-order model and the declared disturbance bound; it is not a plant-level stability proof.

Data availability tocsectionData availability The barrier model, the robust filter, the forward-invariance proof, the Monte-Carlo verification and the figure-generation script are archived with the deposit for this paper (doi:10.5281/zenodo.22645716), which cross-references the Kronos 2026 series including the breeder platform [34], the ideal-MHD stability analysis [35], and the AI/quantum control paper [42]. All quantitative values trace to the register [43] by the serial identifiers cited inline.

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

tocsectionReferences

APPARATUS

Portfolio, People & Record

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

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

INTELLECTUAL PROPERTY

The patent portfolio

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

GRANTED & PENDING UTILITY

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

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

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

TRADEMARK & ORIGIN

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

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

THE TEAM

Founding partners, board, advisors & auditors

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

FOUNDING PARTNERS — SCIENCE

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

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

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

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

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

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

SCIENTIFIC ADVISORS

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

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

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

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

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

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

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

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

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

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

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

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

SCIENTIFIC AUDITORS

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

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

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

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

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

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

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

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

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

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

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

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

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

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

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

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

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

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

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

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

Andrea Romero
Marketing Lead
2022–Present

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

SOURCES

References

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

2 A D Ames, J W Grizzle and P Tabuada, Control barrier function based quadratic programs with application to adaptive cruise control, *53rd IEEE Conference on Decision and Control (CDC)* 6271 (2014).

3 P Wieland and F Allg\"ower, Constructive safety using control barrier functions, *IFAC Proceedings Volumes* **40**(12) 462 (2007).

4 X Xu, P Tabuada, J W Grizzle and A D Ames, Robustness of control barrier functions for safety critical control, *IFAC-PapersOnLine* **48**(27) 54 (2015).

5 Q Nguyen and K Sreenath, Exponential control barrier functions for enforcing high relative-degree safety-critical constraints, *American Control Conference (ACC)* 322 (2016).

6 W Xiao and C Belta, Control barrier functions for systems with high relative degree, *58th IEEE Conference on Decision and Control (CDC)* 474 (2019).

7 F Blanchini, Set invariance in control, *Automatica* **35**(11) 1747 (1999).

8 K P Wabersich and M N Zeilinger, A predictive safety filter for learning-based control of constrained nonlinear dynamical systems, *Automatica* **129** 109597 (2021).

9 S Boyd and L Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).

10 B Stellato, G Banjac, P Goulart, A Bemporad and S Boyd, OSQP: an operator splitting solver for quadratic programs, *Mathematical Programming Computation* **12** 637 (2020).

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

12 J Seo *et al*, Avoiding fusion plasma tearing instability with deep reinforcement learning, *Nature* **626** 746 (2024).

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

14 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

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

16 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

17 M Ariola and A Pironti, *Magnetic Control of Tokamak Plasmas* (Springer, 2008).

18 D Humphreys *et al*, Novel aspects of plasma control in ITER, *Phys. Plasmas* **22** 021806 (2015).

19 J Wesson, *Tokamaks*, 4th ed. (Oxford University Press, 2011).

20 F Troyon, R Gruber, H Saurenmann, S Semenzato and S Succi, MHD-limits to plasma confinement, *Plasma Phys. Control. Fusion* **26**(1A) 209 (1984).

21 M Greenwald, Density limits in toroidal plasmas, *Plasma Phys. Control. Fusion* **44** R27 (2002). doi:10.1088/0741-3335/44/8/201

22 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

23 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

24 A J Creely *et al*, Overview of the SPARC tokamak, *J. Plasma Phys.* **86** 865860502 (2020). doi:10.1017/S0022377820001257

25 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

26 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

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

28 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

29 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

30 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

31 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

32 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

33 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

34 P I Ford, Hyperion as a strategic-materials platform: gain, tritium breeding, He-3 co-production, and 14 MeV neutron output from a compact spherical tokamak, Kronos Fusion Energy 2026 series. 10.5281/zenodo.22645691https://www.kronosfusionenergy.com/publications/breeder.html

35 P I Ford, Free-boundary equilibrium and ideal-MHD stability of a compact spherical-tokamak breeder: Troyon, Mercier, ballooning and external-kink margins, Kronos Fusion Energy 2026 series. 10.5281/zenodo.22645740https://www.kronosfusionenergy.com/publications/ideal-mhd-stability.html

36 P I Ford, Negative-triangularity turbulence suppression in a compact spherical-tokamak breeder: nonlinear gyrokinetics of ion- and electron-scale quiescence, Kronos Fusion Energy 2026 series. 10.5281/zenodo.22645722https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html

37 P I Ford, Model-predictive burn control inside a certified safe envelope: a reference-governor MPC bounded by a control-barrier-function safety clamp, Kronos Fusion Energy 2026 series. 10.5281/zenodo.22645807https://www.kronosfusionenergy.com/publications/model-predictive-burn-control-inside-a-certified-safe-e.html

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

39 P I Ford, Calibrated uncertainty, conformal prediction, and abstention for safety-critical fusion control, Kronos Fusion Energy 2026 series. 10.5281/zenodo.22645813https://www.kronosfusionenergy.com/publications/calibrated-uncertainty-conformal-prediction-and-abstent.html

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

41 P I Ford, Prognostics and remaining-useful-life estimation for the consumable centrepost: a dual-channel fluence-driven framework, Kronos Fusion Energy 2026 series. 10.5281/zenodo.22645788https://www.kronosfusionenergy.com/publications/centrepost-prognostics-and-rul-control.html

42 P I Ford, AI and quantum computing for fusion: ML surrogates, real-time control, and a resource-honest quantum frontier, Kronos Fusion Energy 2026 series. 10.5281/zenodo.22645702https://www.kronosfusionenergy.com/publications/ai-quantum-control.html

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

COLOPHON

How this document was made

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

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

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

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

LEGAL NOTICE

Notice, status, and distribution

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

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

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

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

A Certified Control-Barrier Safety Filter for Fusion Control: Forward-Invariance Guarantees, Calibrated Uncertainty, and an Honest PINN Negative for a Compact Spherical-Tokamak Breeder

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

Zenodo · doi:10.5281/zenodo.22645716

Read on kronosfusionenergy.com

The Physics De-Risking Register