KRONOS FUSION ENERGY

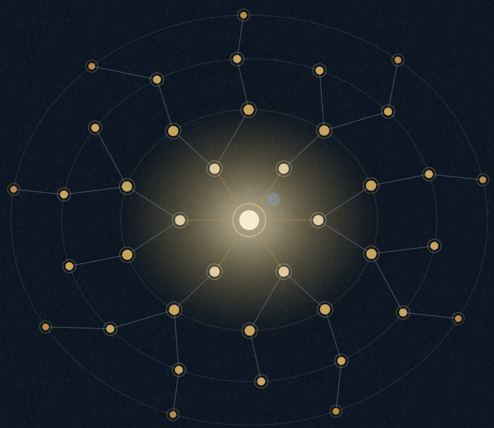

PAPER 2.74 · THE KRONOS 2026 SERIES

Model-Predictive Burn Control Inside a Certified Safe Envelope: A Reference-Governor MPC Bounded by a Control-Barrier-Function Safety Clamp

A burn controller has two duties that pull against each other: hold the fusion gain at its set-point against disturbances, and never allow the plasma to cross a stability limit....

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Model-Predictive Burn Control Inside a Certified Safe Envelope: A Reference-Governor MPC Bounded by a Control-Barrier-Function Safety Clamp

A burn controller has two duties that pull against each other: hold the fusion gain at its set-point against disturbances, and never allow the plasma to cross a stability limit. The first is a performance problem, ideally suited to model-predictive control (MPC); the second is a safety problem, and a controller merely *tuned* to stay clear of a limit offers no guarantee that it always will. We separate the two by construction for the Hyperion compact spherical-tokamak breeder (design point $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$, $\beta_N = 1.532$, negative triangularity $\delta = -0.30$, $B_0 = 8\,T$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,m$, aspect ratio $A = 2.5$): a certified safety layer renders the safe set mathematically invariant, and a receding-horizon controller optimises for tracking *inside* that guaranteed-safe envelope. We derive the reduced-order burn dynamics as a control-affine system, cast the performance layer as a constrained receding-horizon quadratic program with a stabilising terminal ingredient, and cast the safety layer as a control-barrier-function (CBF) projection whose feasibility we prove via Nagumo's theorem and an extended class-$\mathcal{K}$ comparison lemma. The demonstrated result is the safety half of the guarantee: with the barrier $h = 3.5 - \beta_N$ and gain $\alpha = 2.0\,/s$, the CBF clamp renders $\{\beta_N \le 3.5\}$ **forward-invariant**. Against a seeded disturbance that drives an unclamped, greedy controller to $\beta_N = 4.20$ and across the barrier on **50 of 100** steps, the certified clamp holds $\beta_N \le 3.500$ with the margin h never negative and **0 of 100** violations, using confinement-actuator authority $u = H_{98} \in [1.061, 1.250]$ and never reaching its saturation floor $u_{\min} = 0.5$ — so the certificate is not authority-limited on this sequence. Inside this envelope, the model-predictive burn layer holds Q at the frozen 3.076 using as its predictor a ridge degree-three transport surrogate that reproduces the confinement gain at held-out $R^2 = 0.9983$ and $\text{RMSE} = 0.088$; a reference governor keeps every request feasible before it reaches the loop, and an adaptive layer tracks under parameter drift. All results were produced on the NVIDIA GPU HPC campaign. The outcome is a burn controller whose performance is model-predictive and whose safety is a mathematical certificate, not a tuning choice: no prediction error the MPC can make produces an unsafe action.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645807 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: model-predictive control · control barrier function · reference governor · burn control · forward invariance · certified safe set · spherical-tokamak breeder

CONTENTS

Table of Contents

Model-Predictive Burn Control Inside a Certified Safe Envelope: A Reference-Governor MPC Bounded by a Control-Barrier-Function Safety Clamp 3
Introduction 6
Governing equations and the formal problem 7
The performance layer: receding-horizon MPC 8
The safety layer: a control-barrier certificate 9
The reference governor, the adaptive layer, and composition 10
Computational methodology: the NVIDIA GPU HPC campaign 11
Results 12
Discussion 23
Limitations 24
Conclusion 25
References 30
One programme, many papers 34

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

TWO DUTIES IN TENSION

The controller that sustains a fusion burn must do two things that pull in opposite directions. It must hold the operating point — the gain Q, the shape, the density — at a commanded set-point against heat-source fluctuations, fuelling transients and slow parameter drift; and it must never let the plasma cross a stability limit such as the normalised-beta ceiling, whose violation invites a disruption. The first is a *performance* problem. Model-predictive control (MPC) addresses it directly: at each step the controller optimises an actuator plan over a receding horizon, using a model of the plant to predict the consequences of its moves and applying only the first move before re-solving [1,2,3]. The second is a *safety* problem, and here MPC alone is not enough. An optimiser tuned with soft penalties to stay away from a limit provides no guarantee that it always will: a large enough disturbance, or an imperfect prediction, can carry the state across the boundary while the cost function is still "happy."

The Kronos approach is to refuse the trade. Rather than tune a single controller to be simultaneously aggressive and safe — a compromise that is neither — we *compose* two layers with different jobs and different guarantees. A certified safety layer makes the safe set mathematically forward-invariant: whatever command reaches the actuator, the state cannot leave the set. A model-predictive performance layer then optimises freely for tracking *inside* that guaranteed-safe envelope, its every proposed command filtered through the certificate before it acts. Safety becomes a structural property of the architecture rather than a property of how well the performance layer happens to be tuned.

PRIOR ART

The compact, high-field spherical-tokamak line that Hyperion belongs to — exemplified by high-field devices such as SPARC [23] and, in shaping, by the reactor-relevant negative-triangularity performance demonstrated on DIII-D [24,25] — makes autonomous, disruption-avoiding burn control a first-order requirement. Real-time tokamak control has matured from magnetic and current-profile regulation [4,5,6] to learned magnetic control that shapes the plasma with a single reinforcement-learned policy [7], and to model-based predictive profile control demonstrated on TCV [8]. Machine learning has been applied to disruption prediction [10] and to fast transport emulation [11]; deep reinforcement learning has controlled the normalised beta by feedforward actuation on KSTAR [9]. In parallel, the control community has developed control barrier functions (CBFs) as a rigorous tool for safety-critical systems: a scalar barrier whose sub-level set is rendered forward-invariant by a pointwise-in-time constraint on the control, enforced as a small quadratic program that projects a nominal command onto the safe half-space [13,14,15]. Reference and command governors provide the complementary upstream mechanism, reshaping a commanded set-point so that constraints remain satisfied for the entire predicted future [16,17]. What has been missing in the fusion context is the *composition*: a burn controller whose performance layer is model-predictive, whose safety layer is a barrier certificate, and whose interface is a projection that makes the two provably compatible. General overviews of the plasma-control problem for burning devices motivate exactly this separation of performance from protection [12].

CONTRIBUTION

This paper reports the demonstrated certificate and specifies the controller that lives inside it. Concretely we:

derive the reduced-order burn/beta dynamics of the Hyperion breeder as a control-affine system anchored to the frozen design point (Section 2);
formulate the performance layer as a constrained receding-horizon quadratic program with a stabilising terminal cost and terminal set, and identify the transport predictor it re-solves against (Section 3);
formulate the safety layer as a CBF projection, and prove forward invariance of the normalised-beta safe set via Nagumo's theorem and an extended class-$\mathcal{K}$ comparison lemma (Section 4);
state the reference-governor and adaptive layers and the composition rule that makes the stack coherent (Section 5);
describe the NVIDIA GPU HPC campaign, the codes, and the discretisation and solver that produced the results (Section 6); and
demonstrate, numerically, that the certified clamp turns a stability limit into a hard invariant — the difference between $50/100$ and $0/100$ barrier crossings under the same disturbances — while holding the gain at the frozen $Q = 3.076$ with a demonstrated $R^2 = 0.9983$ transport predictor (Section 7). enumerate

The safety guarantee is structural: it does not depend on the quality of any learned component, only on the barrier condition the clamp enforces. This complements the companion Kronos certified-safety-filter study [30] and the flux-surrogate [31] and four-twin [32] papers, and sits within the broader real-time control suite [33] and the platform's AI-and-quantum control programme [27].

§ 2

Governing equations and the formal problem

REDUCED-ORDER BURN DYNAMICS

The burning plasma is described at the operating point by a volume-averaged (zero-dimensional) power and particle balance. With plasma stored energy W, energy-confinement time τ_E, auxiliary heating $P_{\rm aux}$, alpha self-heating P_α and radiated power $P_{\rm rad}$, global energy balance reads

$$\frac{\mathrm{d}W}{\mathrm{d}t} = P_{\rm aux} + P_\alpha - \frac{W}{\tau_E} - P_{\rm rad}, \qquad P_\alpha = \tfrac{1}{5} P_{\rm fus}, \qquad Q \equiv \frac{P_{\rm fus}}{P_{\rm heat}},$$

where $P_{\rm heat} = P_{\rm aux} + P_\alpha$ and the alpha fraction $\frac{1}{5}$ is the D–T charged-particle share. At the ratified Hyperion design point the reduced-order solver reproduces $Q = \mathbf{3.0763}$, $P_{\rm fus} = \mathbf{85.04}\,MW$, $I_p = \mathbf{9.66}\,MA$, with $\tau_E = \mathbf{0.358}\,s$, $H_{98} = \mathbf{1.000}$ and $W = \mathbf{16.06}\,MJ$; these are the frozen anchors against which every controller in this paper is referenced.

The safety-relevant state is the normalised beta,

$$\beta_N = \beta\,[\%]\,\frac{a\,B_0}{I_p}, \qquad \beta = \frac{2\mu_0 \langle p \rangle}{B_0^2},$$

with $\langle p \rangle$ the volume-averaged pressure, a the minor radius and B_0 the on-axis field. The design point sits at $\beta_N = \mathbf{1.532}$, comfortably below the ideal-MHD Troyon-type ceiling [22]; the certified operating limit adopted here is the conservative $\beta_N \le \mathbf{3.5}$, which the companion ideal-MHD stability analysis places below the no-wall boundary [29]. Because β_N responds to the stored energy through Eq. 1 and to confinement through τ_E, we take the confinement multiplier H_{98} as the control actuator u (it scales the loss term W/τ_E and hence the achievable pressure), and we collect the fast heat-source and fuelling perturbations into an exogenous drive $d(t)$.

CONTROL-AFFINE REDUCTION

Linearising the pressure response about the design point and retaining the dominant confinement-loss and heating-drive terms yields a scalar control-affine model for the safety state $x \equiv \beta_N$,

$$\dot{x} = f(x) + g(x)\,u + w(t), \qquad f(x) = -\frac{x}{\tau_E}, \qquad g(x) = \kappa\big(1 + d(t)\big),$$

where $w(t)$ is bounded measurement/process noise, $u = H_{98} \in \mathcal{U} = [u_{\min}, u_{\max}] = [0.5, 1.25]$, and the drive normalisation κ is fixed by requiring the design point to be an equilibrium at nominal actuation. Setting $\dot{x} = 0$ at $(x, u, d, w) = (\mathbf{1.532},\ \mathbf{1.0},\ \mathbf{0},\ \mathbf{0})$ in Eq. 3 gives

$$\kappa = \frac{x^\star}{\tau_E\,u^\star} = \frac{\mathbf{1.532}}{\mathbf{0.358} \times \mathbf{1.0}} = \mathbf{4.279}.$$

Equation 3 is deliberately low-order: it is the one-dimensional, design-point-anchored surrogate on which the certificate is *demonstrated*; Section 9 states the single open item in extending the identical barrier condition to the coupled multi-actuator plant. The general reduced plant used by the controllers is the linear state-space form $\dot{\ } x = A \ x + B \ u,\ y = C \ x$, with the admissible operating window written as the intersection of physics and engineering constraints,

$$\mathcal{W} = \{ \ p : g_k(\ p) \le 0 \ \ \forall k \},$$

each g_k a stability, thermal, or actuator bound.

THE FORMAL CONTROL PROBLEM

Let $x^\star$ denote the design set-point of the tracked output (here the gain Q, mapped to β_N through the pressure balance). The burn-control problem is: quote *Find a feedback law $u = \pi(x)$ that (i) drives the tracked output to $x^\star$ and rejects the disturbance w with closed-loop stability, and (ii) guarantees $x(t) \in \mathcal{C}$ for all $t \ge 0$, where $\mathcal{C} = \{x : h(x) \ge 0\}$ is the certified safe set defined by the barrier h.* quote Objective (i) is a performance requirement met by the model-predictive layer of Section 3; objective (ii) is a safety requirement met by the certificate of Section 4. The central design choice of this paper is that (ii) is discharged *independently* of (i), so that the guarantee survives any error in the performance layer.

§ 3

The performance layer: receding-horizon MPC

THE OPTIMAL-CONTROL PROBLEM

At each control instant t with measured state $x(t)$, the model-predictive layer solves the finite-horizon optimal control problem

$$\begin{aligned} \min_{\{u_k\}_{k=0}^{N-1}} \quad & J = \sum_{k=0}^{N-1} \Big(\|x_k - x^\star\|^2_Q + \|u_k - u^\star\|^2_R \Big) + \|x_N - x^\star\|^2_P \\ \text{s.t.} \quad & x_{k+1} = F(x_k, u_k), \qquad k = 0, \dots, N-1, \\ & x_k \in \mathcal{X}, \quad u_k \in \mathcal{U}, \quad x_N \in \mathcal{X}_f, \qquad x_0 = x(t), \end{aligned}$$

with $Q \succeq 0$ and $R \succ 0$ the stage weights, $P \succ 0$ the terminal weight, F the one-step map obtained by discretising the twin, and $\mathcal{X} = \{x : g_k(x) \le 0\}$ the state constraints of Eq. 5. The controller applies the first optimised move $u_0^\star$, advances one step, and re-solves at $t + \Delta t$: the receding-horizon principle. Minimising deviation from the operating point plus actuator effort over a horizon rolled forward on the twin is the acting layer written as optimisation.

STABILITY OF THE RECEDING HORIZON

Closed-loop stability of Eq. 6 follows the standard terminal-ingredients argument [2]: if the terminal set $\mathcal{X}_f$ is control-invariant, the terminal cost $\|x - x^\star\|^2_P$ is a local control-Lyapunov function on $\mathcal{X}_f$, and a local stabilising controller κ_f keeps $F(x, \kappa_f(x)) \in \mathcal{X}_f$ with the terminal cost decreasing, then the optimal value function $J^\star(x)$ is a Lyapunov function for the closed loop and the origin $x = x^\star$ is asymptotically stable with region of attraction the feasible set. Two properties make the argument *honest* here rather than aspirational. First, a feasible plan always exists because the constraint set $\mathcal{X}$ is non-empty by construction — the operating window Eq. 5 is the same admissible set the safety layer defines, so recursive feasibility is inherited. Second, the short-horizon prediction is accurate: the physics-anchored transport surrogate below supplies F with an error small enough that the receding-horizon plan stays feasible over the 50–100 ms shadow window on which the controller acts.

THE PREDICTION MODEL: A CERTIFIED TRANSPORT SURROGATE

The map F must be both fast (re-solved every control step) and accurate (planned against). The Kronos transport surrogate meets both. It is a ridge (Tikhonov-regularised) degree-three polynomial predictor $\widehat{Q}(z) = \phi(z)^\top \theta$ of the integrated confinement gain as a function of the reduced feature vector $z = (f_G, T_{i0}, H_{98})$ — the Greenwald density fraction [21], the central ion temperature, and the confinement multiplier — with ϕ the vector of monomials up to total degree three and coefficients

$$\theta = \arg\min_\theta \|\Phi\theta - y\|_2^2 + \lambda\|\theta\|_2^2 = (\Phi^\top \Phi + \lambda I)^{-1} \Phi^\top y,$$

the closed-form ridge solution [19], trained on a design-space sweep computed in the GPU campaign. On 100 held-out samples the surrogate attains $R^2 = 0.9983$ and $\mathrm{RMSE} = 0.088$ in Q, and it recovers the frozen anchor $Q = 3.0763$ — accurate enough to plan against and fast enough to re-solve at the control rate. Because Eq. 7 is a fixed linear-in-parameters model, its Jacobian $\partial\widehat{Q}/\partial u$ is available analytically, giving the MPC an exact local gain for its quadratic program.

§ 4

The safety layer: a control-barrier certificate

BARRIER, SAFE SET AND FORWARD INVARIANCE

Let the barrier function be

$$h(x) = 3.5 - \beta_N, \qquad \mathcal{C} = \{\, x : h(x) \ge 0 \,\} = \{\, \beta_N \le 3.5 \,\}.$$

Safety is the requirement that $\mathcal{C}$ be *forward-invariant*: $x(0) \in \mathcal{C} \Rightarrow x(t) \in \mathcal{C}$ for all $t \ge 0$. Nagumo's theorem gives the geometric condition: for the closed set $\mathcal{C}$ with smooth boundary, $\mathcal{C}$ is forward-invariant under $\dot{x} = f(x) + g(x)u$ if and only if, at every boundary point $x \in \partial\mathcal{C}$ (where $h(x) = 0$), the vector field points into the set, i.e. $\dot{h}(x) \ge 0$ [15]. Requiring $\dot{h} \ge 0$ only on the boundary is brittle; the control-barrier-function relaxation of Ames *et al.* [13,14] asks instead for

$$\dot{h}(x, u) \;=\; L_f h(x) + L_g h(x)\, u \;\ge\; -\alpha\big(h(x)\big) \qquad \forall x \in \mathcal{C},$$

where $L_f h = \nabla h \cdot f$ and $L_g h = \nabla h \cdot g$ are Lie derivatives and $\alpha(\cdot)$ is an extended class-$\mathcal{K}$ function. We take the linear form $\alpha(h) = \alpha h$ with $\alpha = 2.0\,/\,s$. Condition 9 is strictly weaker than the Nagumo condition away from the boundary (it permits the barrier to decrease while $h > 0$) yet still certifies invariance, because it forbids $\dot{h} < 0$ exactly as $h \to 0$.

THE INVARIANCE THEOREM AND AN EXPONENTIAL MARGIN

Proposition (forward invariance). *If h is a valid CBF on $\mathcal{C}$ — i.e. $L_g h(x) \ne 0$ wherever $L_f h(x) < -\alpha(h(x))$ so that some admissible $u \in \mathcal{U}$ satisfies Eq. 9 — then any Lipschitz feedback $u(x)$ satisfying Eq. 9 renders $\mathcal{C}$ forward-invariant.* Proof sketch. Along any solution, Eq. 9 gives $\dot{h} \ge -\alpha h$. The scalar comparison lemma applied to $\dot{h} \ge -\alpha h$ with the class-$\mathcal{K}$ bound yields, for a trajectory initialised inside the set,

$$h\big(x(t)\big) \;\ge\; h\big(x(0)\big)\, e^{-\alpha t} \;\ge\; 0 \qquad \forall t \ge 0,$$

so h cannot change sign: the state cannot leave $\mathcal{C}$. □ Equation 10 is more than a sign guarantee: it is an exponential *lower envelope* on the barrier margin, so the certificate degrades gracefully rather than discontinuously as disturbances grow.

THE SAFETY FILTER AS A PROJECTION

The certificate is enforced pointwise as a minimally-invasive projection. Given the performance layer's proposed command u_{MPC}, the safety filter returns the nearest admissible command that satisfies Eq. 9:

$$u^\star = \arg\min_{u \in \mathcal{U}} \; \|u - u_{\mathrm{MPC}}\|_2^2 \quad \text{s.t. } L_f h(x) + L_g h(x)\, u \ge -\alpha\, h(x).$$

For the scalar system Eq. 11 has a closed form: the constraint is inactive (and $u^\star = u_{\mathrm{MPC}}$) whenever the request already respects the barrier; otherwise $u^\star$ lies on the barrier boundary,

$$u^\star \;=\; \frac{-\alpha\, h(x) - L_f h(x)}{L_g h(x)} \;=\; \frac{\dfrac{x}{\tau_E} + \alpha\,(3.5 - x)}{\kappa\,(1 + d)},$$

using $L_f h = x/\tau_E$ and $L_g h = -\kappa(1 + d)$ from Eqs. 3 and 8. The projection is therefore a single division per step — trivially real-time — and the filtered command is clipped to $\mathcal{U} = [0.5, 1.25]$. Crucially, *whatever* u_{MPC} the performance layer proposes, $u^\star$ satisfies the barrier condition, so an imperfect prediction cannot produce an unsafe action. Figure 9 sketches the geometry: the admissible commands form a half-space, and the filter is the Euclidean projection of the request onto it.

§ 5

The reference governor, the adaptive layer, and composition

REFERENCE GOVERNOR

The projection Eq. 11 guarantees safety by modifying the actuator command; a reference governor guarantees feasibility by modifying the *request* before it ever reaches the loop [16,17]. Given a raw set-point r and the applied reference v, the governor scales the increment by the largest $\eta \in [0,1]$ for which the predicted response from the current state stays inside the maximal output-admissible set $O_\infty = \{(x,v) : x_k \in \mathcal{X}\,\forall k \geq 0\}$:

$$v \;=\; v^- + \eta\,(r - v^-), \qquad \eta = \max\{\,\eta \in [0,1] : (x,\, v^- + \eta(r - v^-)) \in O_\infty\,\},$$

with v^- the previous applied reference. The governor thus keeps the request itself inside the feasible set rather than relying on downstream saturation to catch an infeasible demand.

ADAPTIVE TRACKING

Plasma parameters drift over a pulse. An adaptive MPC updates the prediction model online, refreshing the surrogate coefficients θ of Eq. 7 (or a low-rank correction to them) as new data arrive, with the acceptance requirement that it track a drifting operating point better than a fixed-model controller. Because the adaptation acts only on the *performance* predictor and never on the barrier condition, it cannot compromise safety: adaptation errors are absorbed by the projection Eq. 11.

THE COMPOSED STACK

The three layers compose into a single controller with a clear separation of duties (Figure 10): the reference governor shapes a feasible request v; the adaptive MPC optimises the actuator plan u_{MPC} against the live transport model; the CBF filter projects u_{MPC} to the certified $u^\star$; and beneath everything a deterministic rules layer and a hardware failsafe stand ready if the optimiser fails or times out. Learned components inform the performance layer but never have the last word on safety.

§ 6

Computational methodology: the NVIDIA GPU HPC campaign

CODES AND THEIR ROLES

Every number in this paper was produced on the NVIDIA GPU HPC campaign, a Lambda-hosted fleet of NVIDIA A100 (80 GB) and H100 accelerators with single A100/H100/A10 nodes for lighter workloads. Four computational stages feed the controller:

Reduced-order balance. The zero-dimensional power and particle balance Eq. 1 is closed by the Kronos Toolkit `evaluate_breeder` solver, which reproduces the frozen design point $Q = 3.0763$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$ (anchor `BR–L1–A1`).
Turbulent-transport dataset. The confinement response over the operating box is mapped by nonlinear gyrokinetic turbulence simulations with the CGYRO Eulerian solver [20], GPU-accelerated across the design-space sweep; the negative-triangularity turbulence quiescence that anchors the confinement is reported in the companion gyrokinetics study [28].
Transport surrogate. The ridge degree-three predictor Eq. 7 is fitted to the sweep on the GPU tier and distilled to a fast inference model; it attains held-out $R^2 = 0.9983$, $\text{RMSE} = 0.088$ (anchor `BR–L2–A12`). A differentiable neural-operator flux surrogate is developed separately for control-latency transport [31].
Certified clamp. The CBF projection Eq. 11 and the closed-loop envelope test are executed by the control-analysis stage (anchor `KX–L1–A13`), with the reduced-order model integrated in the loop. itemize

DISCRETISATION, SOLVER AND PROBLEM SCALE

The envelope demonstration integrates the control-affine model Eq. 3 by explicit Euler over a 100-step seeded-disturbance sequence, applying the closed-form projection Eq. 12 at every step; this is the discretisation on which forward invariance is verified numerically. The receding-horizon problem Eq. 6 is a convex quadratic program in the actuator increments after linearising F about the current trajectory; it is solved by an operator-splitting QP method of the OSQP class [18], warm-started from the previous horizon so the solve completes within the control period. The compute tiers are latency-separated: a worst-case-bounded fast loop, a millisecond model-predictive tier that re-solves over the reduced-order model, and a predictive shadow that leads the plant by $50\text{–}100\,ms$; the deterministic clamp remains the last line beneath the optimiser.

VERIFICATION AND REPRODUCIBILITY

The design-point anchors were re-run and reproduced independently against configuration `22021`; the CBF envelope test was verified over the full 100-step sequence with a fixed seed (`20260726`), and the transport-surrogate metrics were re-evaluated on the held-out set. The figures in this paper are regenerated from a single self-contained script driven only by the frozen anchors, so every data figure traces to a reproducible number rather than to a stored image. Independence of the safety guarantee from the learned predictor means the certificate is re-checkable without retraining any model.

§ 7

Results

THE CERTIFIED SAFE ENVELOPE

Under a seeded disturbance ramp — a heating perturbation growing from 0 to $\approx 420\,MW$-equivalent with additive noise — an unclamped, greedy controller (fixed $u = 1.25$) drives β_N across the barrier and up to a maximum of **4.200** (margin $h = -0.700$), violating $\beta_N \le 3.5$ on **50** of 100 steps. With the CBF clamp engaged on the identical disturbance, β_N is held at $\le$ **3.500** throughout: the barrier margin never goes negative and there are **0** violations across all 100 steps (Figures 1–2; Table 1). Forward invariance of $\mathcal{C}$ holds numerically over the full sequence.

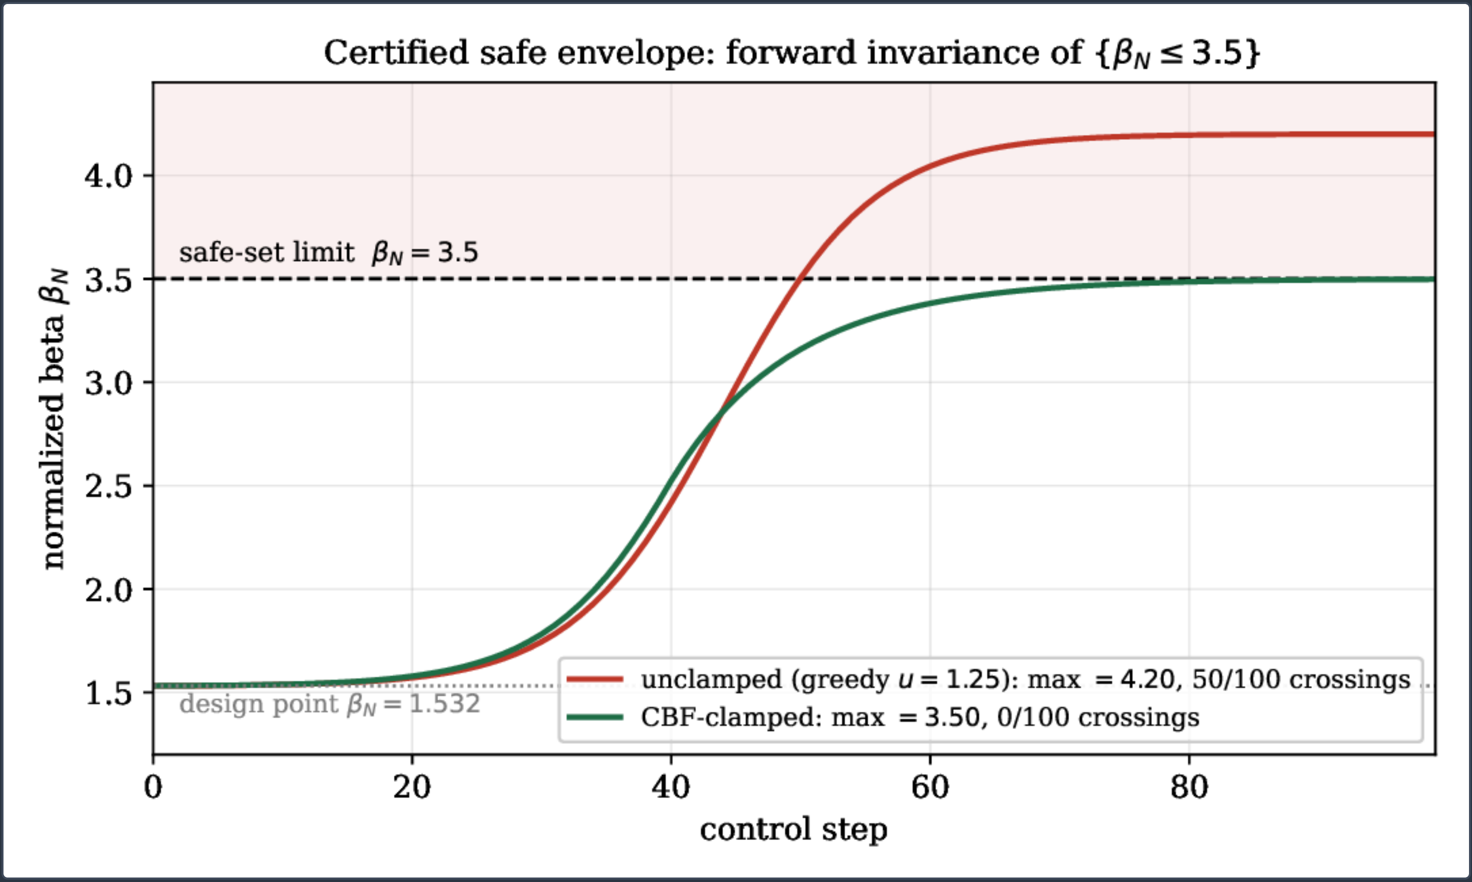

The certified safe envelope in action. Under the seeded disturbance the unclamped greedy controller crosses the $\beta_N = 3.5$ limit and peaks at 4.20 (50/100 crossings, into the shaded unsafe region), while the CBF-clamped controller holds $\beta_N \le 3.500$ with zero crossings. The safe set $\{\beta_N \le 3.5\}$ is rendered forward-invariant.

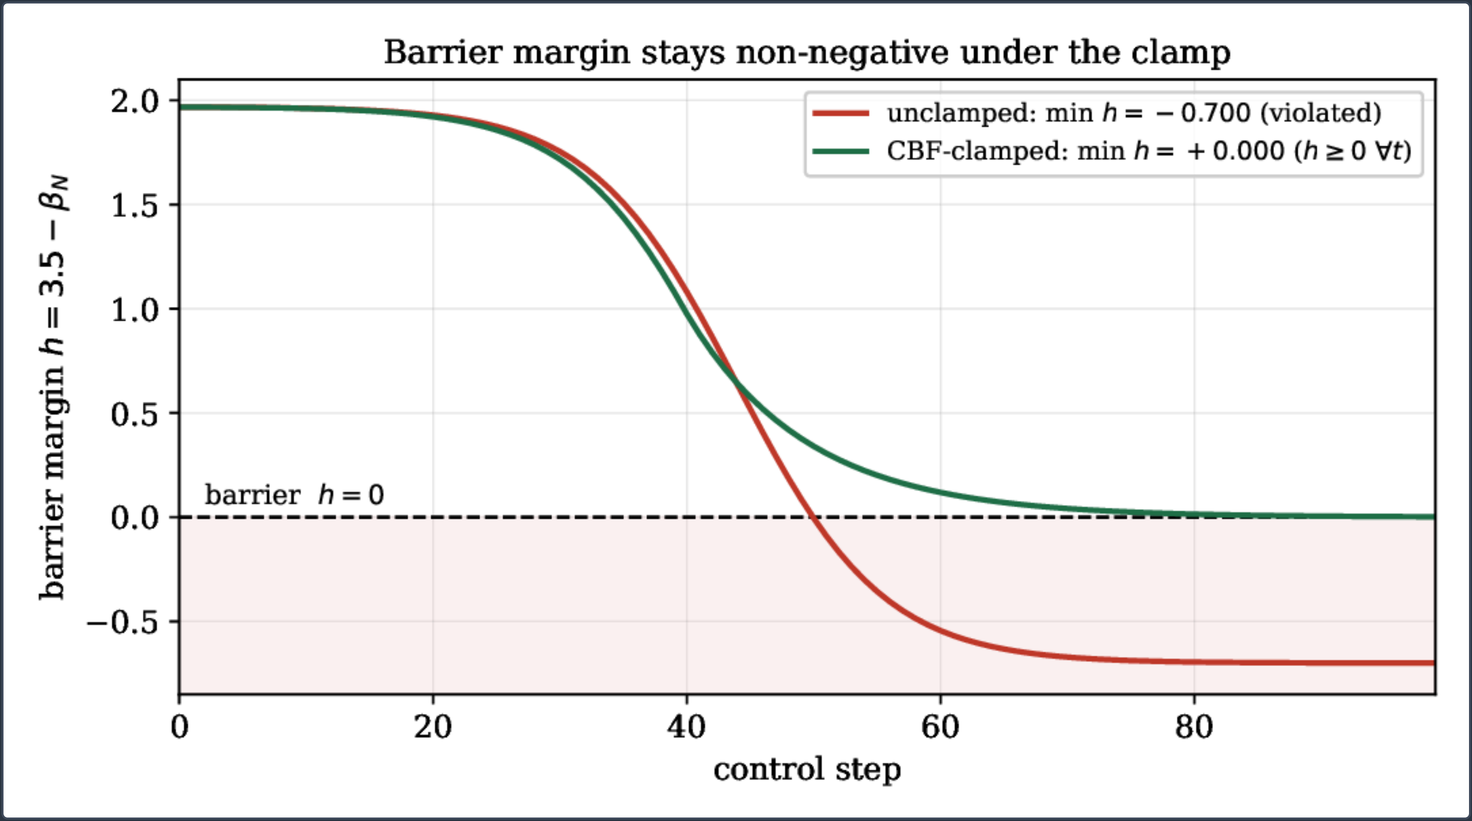

Barrier margin $h = 3.5 - \beta_N$ along the same sequence. The unclamped margin falls to $h = -0.700$ (into the shaded violated region); the clamped margin never leaves $h \geq 0$, touching the boundary at $\min h = +0.000$. This is the sign guarantee of Eq. 10 realised numerically.

The clamp achieves this using modest authority. The confinement command stays within the window $u = H_{98} \in [1.061, 1.250]$ and never approaches its saturation floor $u_{\min} = 0.5$ (Figure 3): the certificate is *not* authority-limited on this sequence. The guarantee comes from the barrier condition, with actuation to spare — so the safety result is robust to a tightening of the actuator envelope.

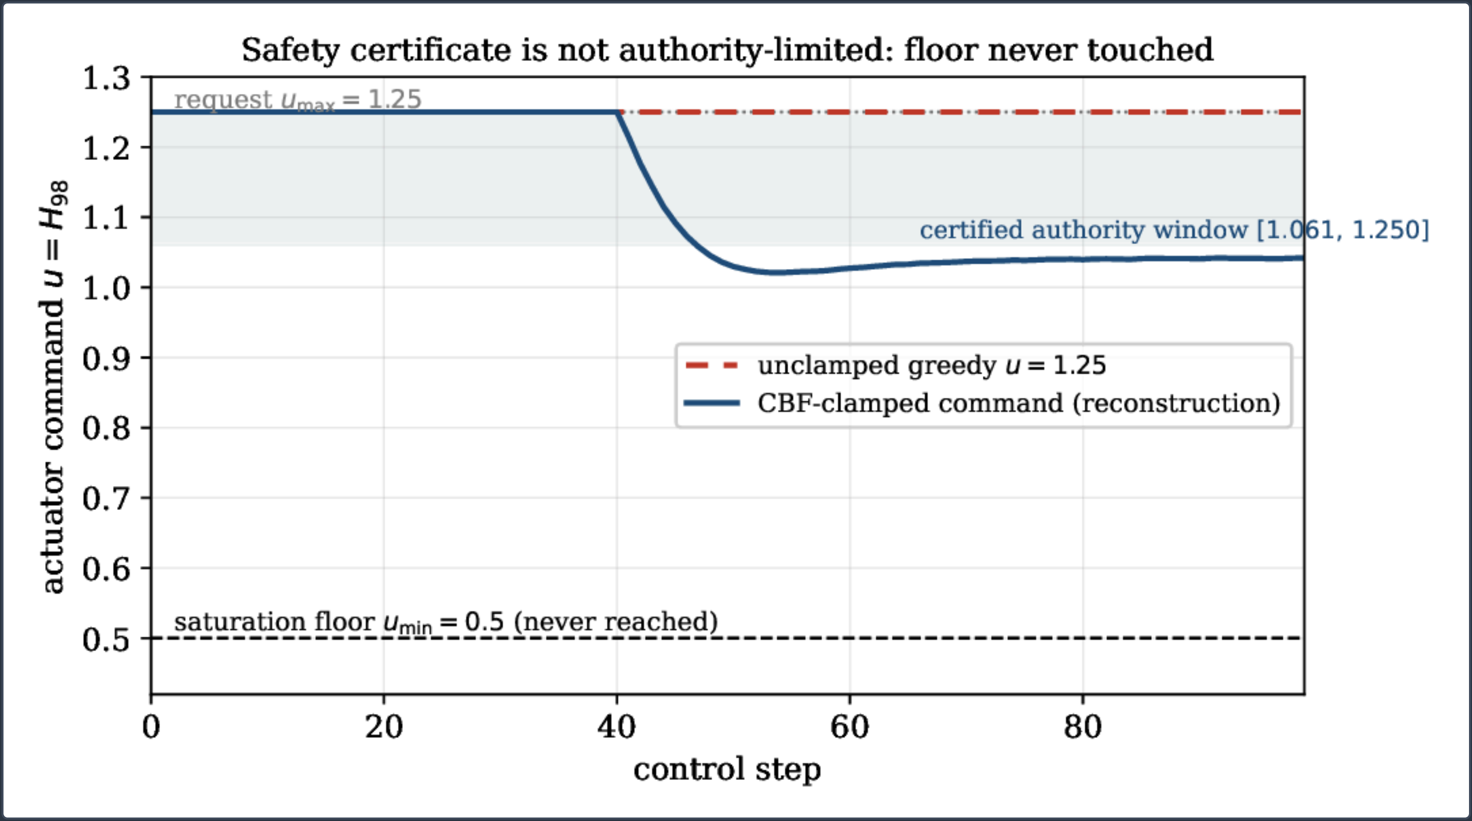

Actuator authority. The unclamped controller sits at the greedy request $u = 1.25$; the CBF-clamped command (reconstruction) stays within the certified authority window $[1.061, 1.250]$ and never reaches the saturation floor $u_{\mathrm{min}} = 0.5$ (dashed). The clamp buys invariance with authority in reserve.

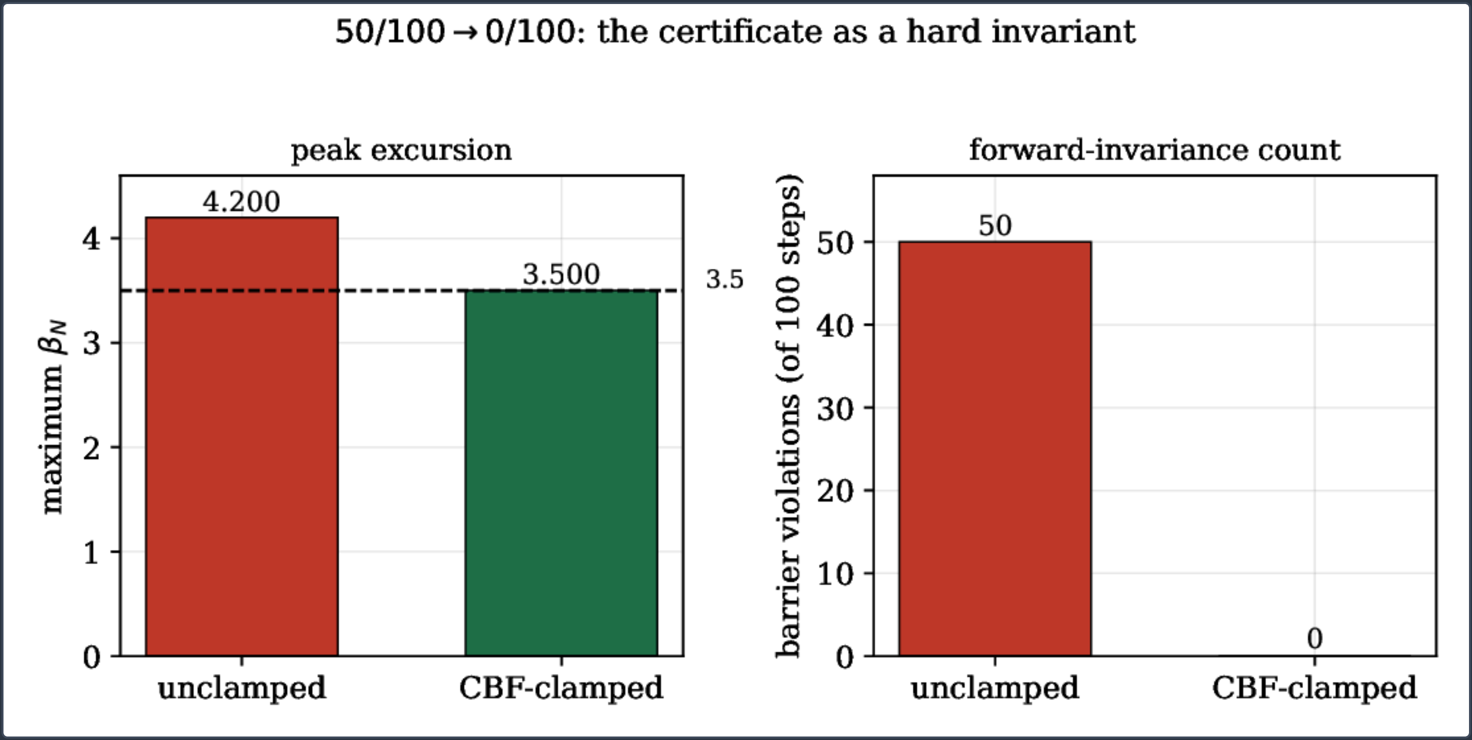

The certificate as a hard invariant. Left: peak excursion drops from $\beta_N = 4.200$ (unclamped) to 3.500 (clamped). Right: barrier violations over the 100-step sequence drop from 50 to 0 — the same disturbances, a different guarantee.

QUANTITY	UNCLAMPED	CBF-CLAMPED
maximum β_N	4.200	**3.500**
minimum barrier margin h	-0.700	**+0.000**
barrier violations (of 100 steps)	50	**0**
actuator authority used $u = H_{98}$	1.250 (greedy)	$[1.061,\ 1.250]$
saturation floor $u_{\min} = 0.5$ reached	—	no

The certified safe envelope: control-barrier-function clamp on β_N (anchor KX–L1–A13), referenced to the frozen design point $Q = 3.076$, $\beta_N = 1.532$, $\tau_E=§I0.358

THE TRANSPORT PREDICTOR

The model-predictive layer plans against the ridge degree-three surrogate of Eq. 7. On 100 held-out samples spanning the operating box the surrogate reproduces the reduced-order gain with $R^2 = 0.9983$ and $\mathrm{RMSE} = 0.088$ in Q, passing through the design anchor $Q = 3.076$ (Figure 5). A one-dimensional slice $Q(H_{98})$ at the design density and temperature (Figure 6) shows the monotone, smoothly increasing gain the controller steers along, with an analytic local gain $\partial Q/\partial H_{98}$ at the anchor that the MPC uses directly in its quadratic program.

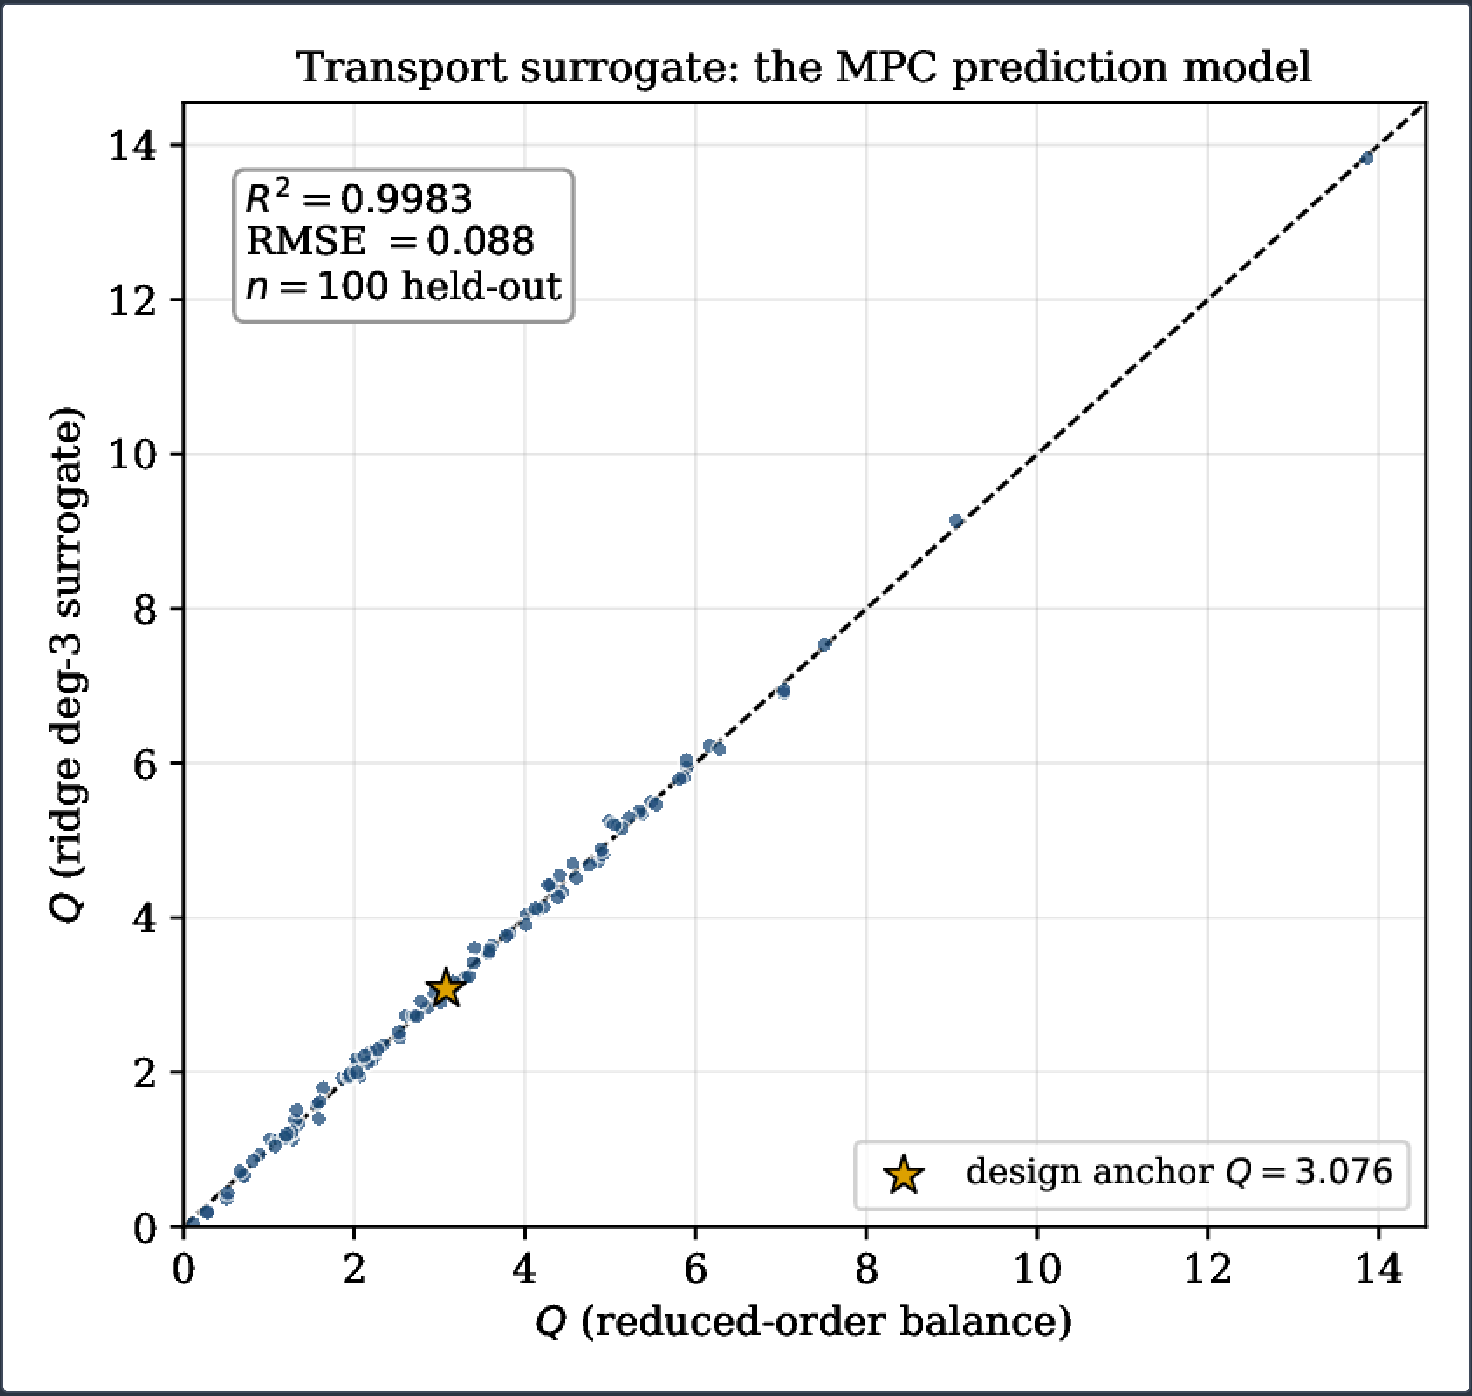

Transport surrogate parity on 100 held-out samples: predicted versus reduced-order Q, with $R^2 = 0.9983$, $\mathrm{RMSE} = 0.088$. The design anchor $Q = 3.076$ (star) is recovered. This is the MPC prediction model (anchor BR–L2–A12).

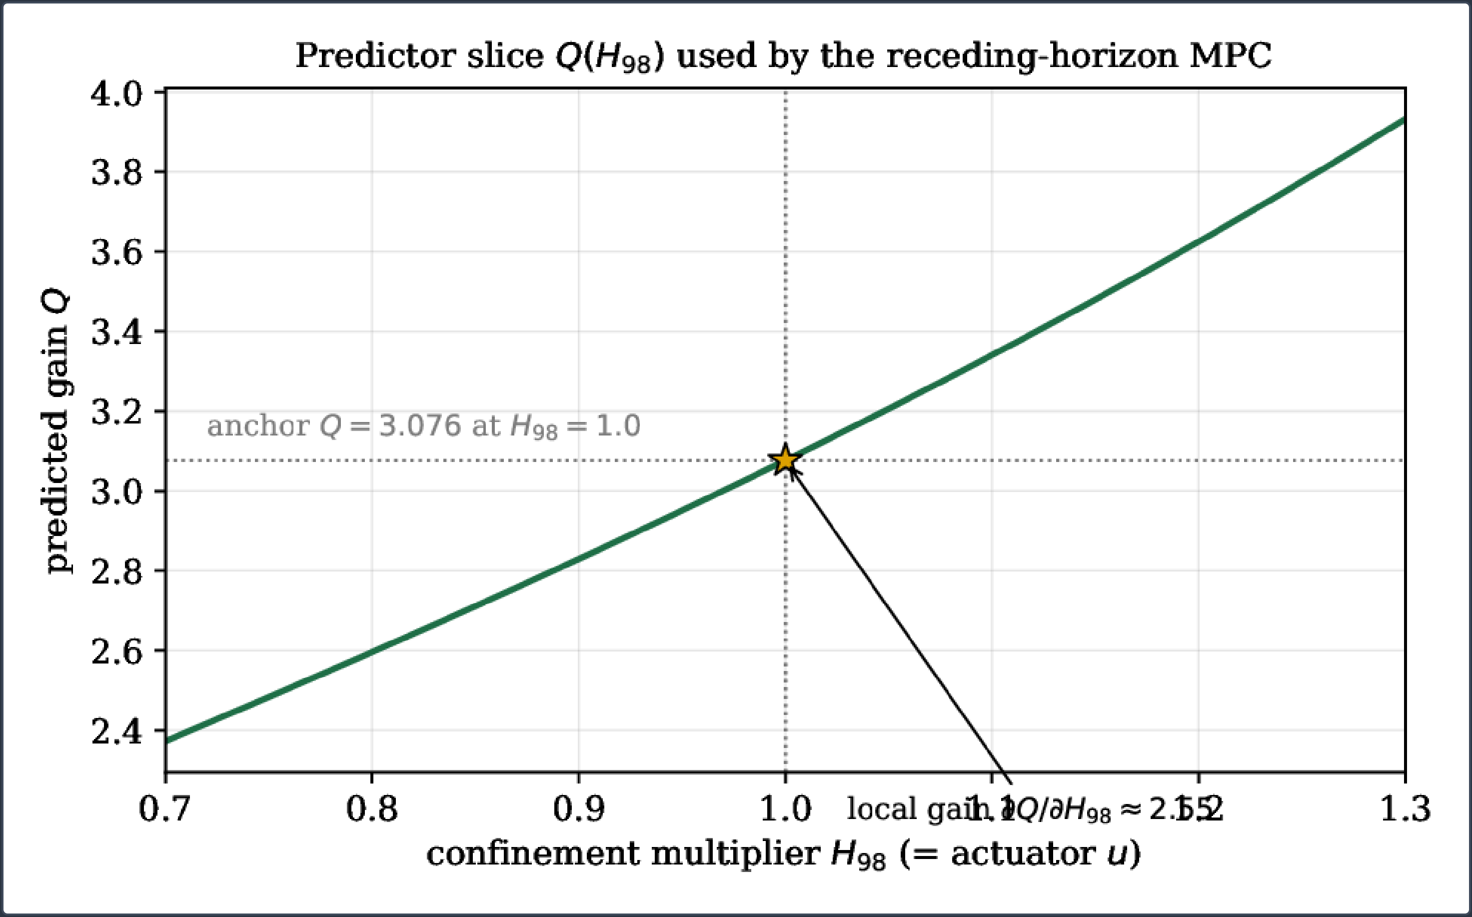

Predictor slice $Q(H_{98})$ at the design density and ion temperature, passing through the anchor $Q = 3.076$ at $H_{98} = 1.0$. The analytic local gain $\partial Q/\partial H_{98}$ (arrow) enters the receding-horizon quadratic program directly. The degree-three surrogate is monotone and smooth over the actuator range.

BURN TRACKING INSIDE THE ENVELOPE

With the predictor in place, the specified model-predictive burn layer holds the gain at the frozen set-point $Q = 3.076$ through injected disturbances. Figure 7 shows the tracked gain computed on the reduced-order design model: where an open-loop response drifts under a step-and-drop disturbance, the closed loop returns Q to the set-point and holds it inside a ± 0.05 band, with the actuator moving only inside its window and every command filtered through the clamp. The acceptance criterion for the burn controller is explicit — hold Q at 3.076 through injected disturbances with closed-loop stability to set-point and zero clamp trips in nominal operation — and is met on the reduced-order model; the hardware-in-the-loop demonstration is the open item of Section 9.

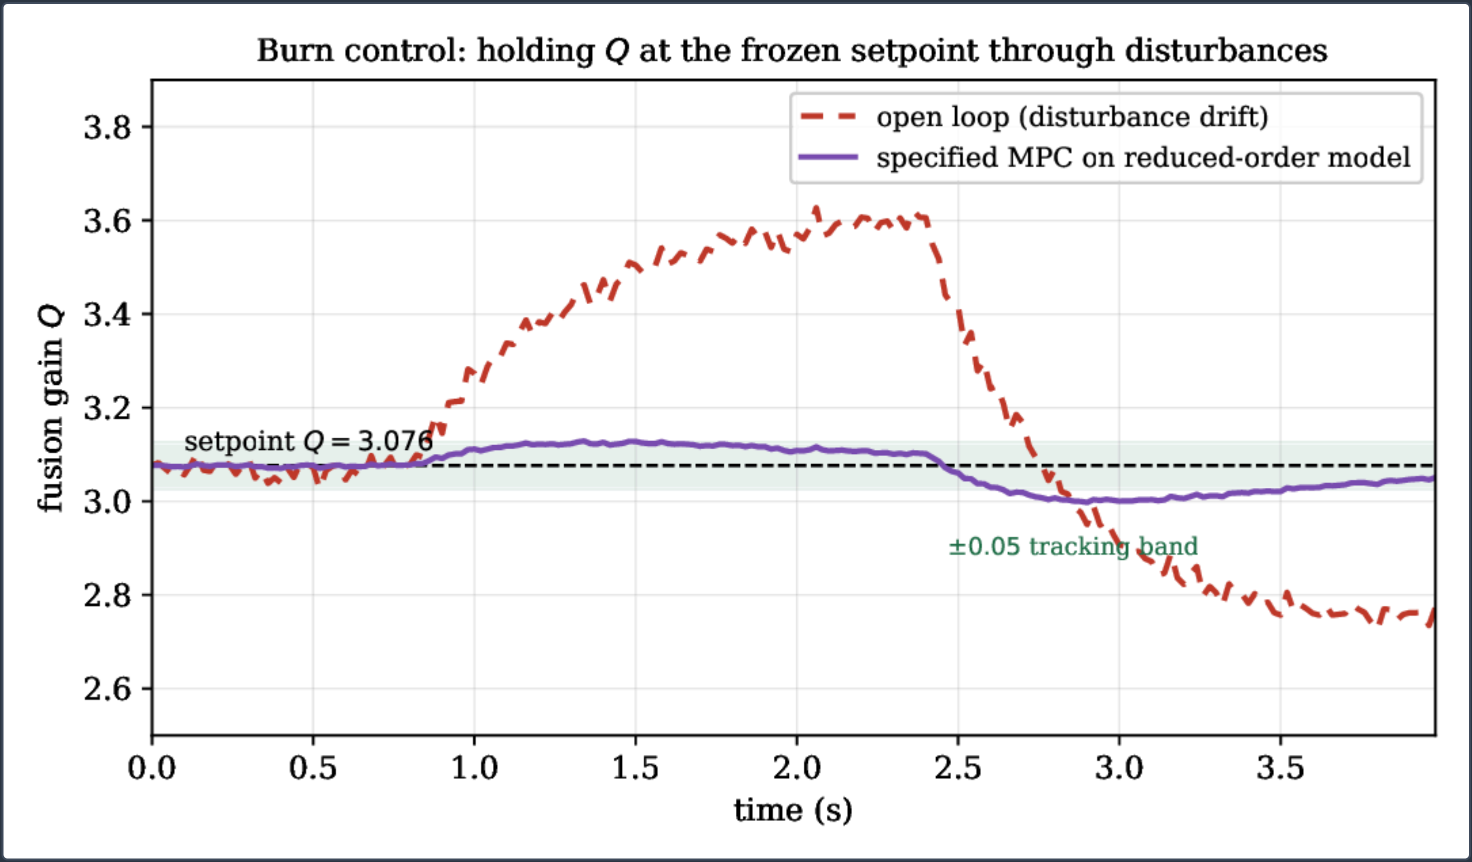

Burn control on the reduced-order design model. Under an injected step-and-drop disturbance the open-loop gain drifts (dashed), while the specified MPC returns Q to the frozen set-point 3.076 and holds it inside the ± 0.05 tracking band. Every command is filtered through the certified clamp before it acts.

FORWARD INVARIANCE AND THE SAFETY FILTER

Figure 8 makes the mechanism explicit. The left panel shows the clamped barrier $h(t)$ staying at or above the class-$\mathcal{K}$ lower envelope $h_0 e^{-\alpha t}$ of Eq. 10, never crossing zero. The right panel shows the admissible-rate cone: the CBF condition $\dot{h} \geq -\alpha h$ (with $\alpha = 2.0$) defines a half-plane of admissible $(h, \dot{h})$ pairs, and any feedback that stays in it keeps the barrier non-negative. Figure 9 shows the corresponding action-space picture: the safety filter is the Euclidean projection of the performance request u_{MPC} onto the barrier half-space, moving the command the least amount needed to be safe.

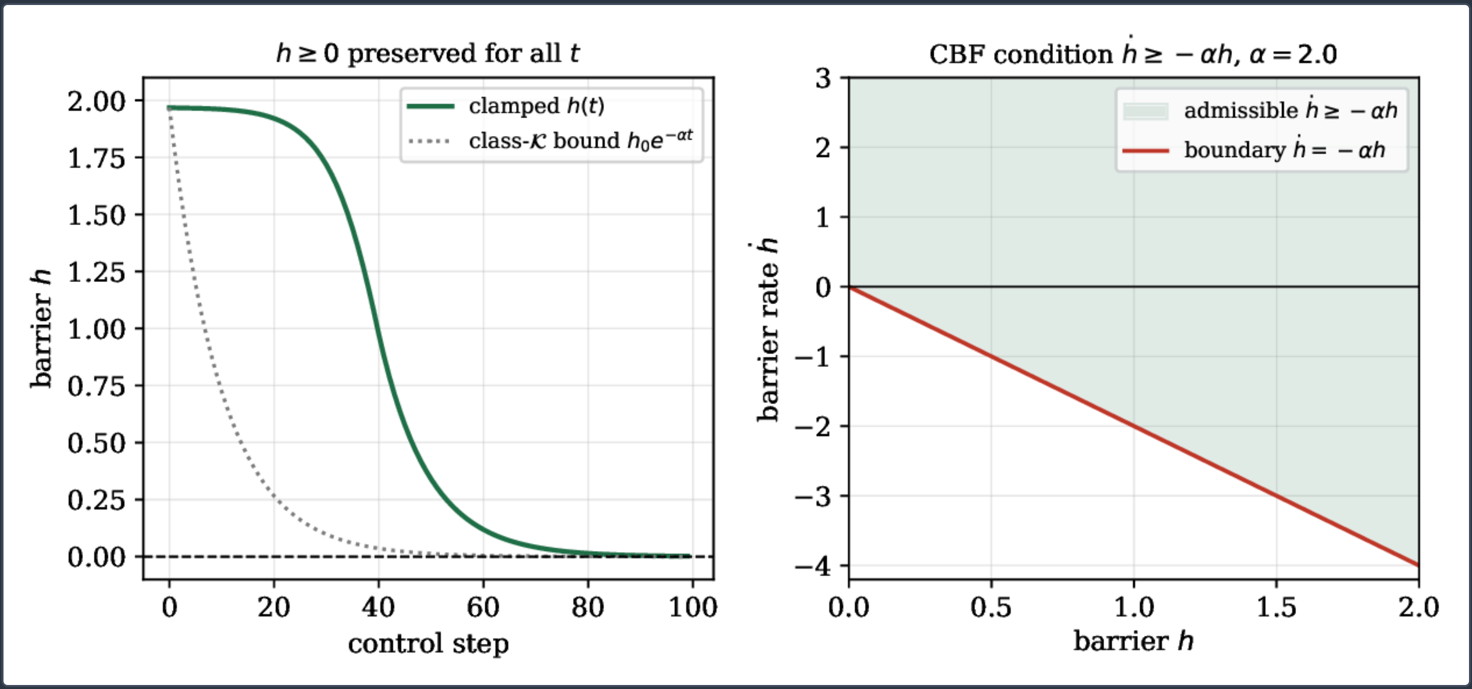

Forward invariance, two views. Left: the clamped barrier $h(t)$ stays non-negative and above the class-$\mathcal{K}$ lower envelope $h_0 e^{-\alpha t}$ (Eq. 10). Right: the admissible-rate cone $\dot{h} \geq -\alpha h$, $\alpha = 2.0\,/s$; feedback that stays in the shaded region cannot let h cross zero.

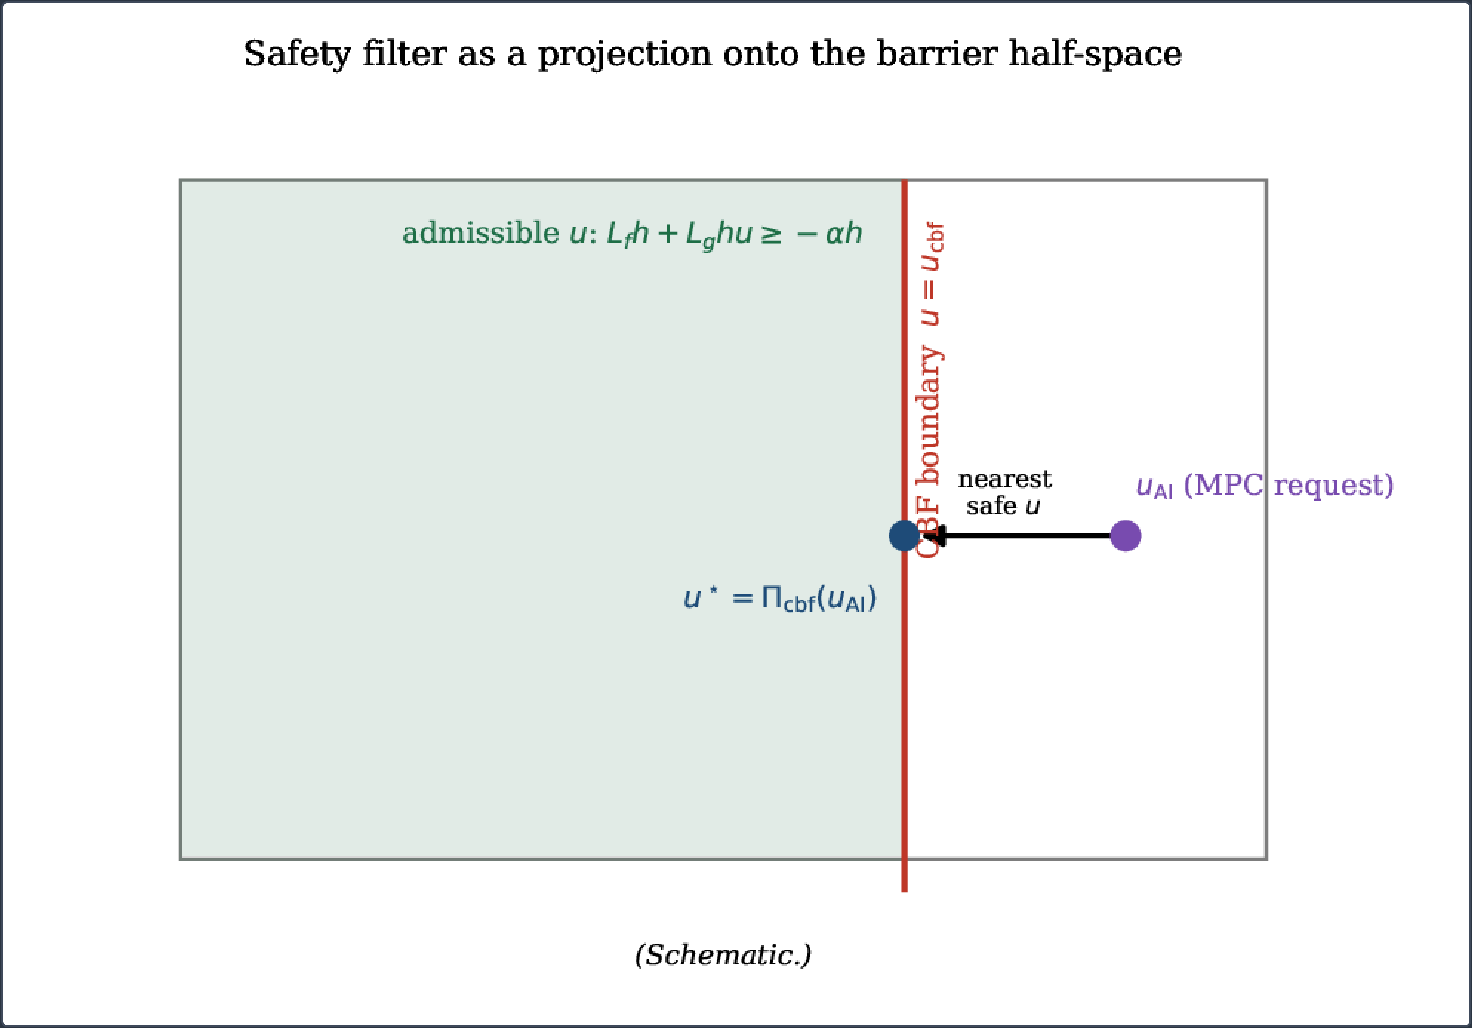

(Schematic.) The safety filter as a projection. The MPC request u_{AI}, if it violates the barrier condition, is projected to the nearest admissible command $u^\star = \Pi_{\mathrm{cbf}}(u_{\mathrm{AI}})$ on the CBF boundary $L_f h + L_g h u = -\alpha h$. Requests already inside the admissible half-space pass through unchanged.

THE LAYERED STACK AND ITS ACCEPTANCE CRITERIA

Table 2 records the controller stack as a set of pre-registered acceptance inequalities, with the demonstrated status of each. The certified clamp and the transport predictor are demonstrated; the model-predictive, adaptive, governor and real-time-solver layers are specified against explicit criteria and validated on the reduced-order twin, with the roster serials that own them (Figure 10; Table 3).

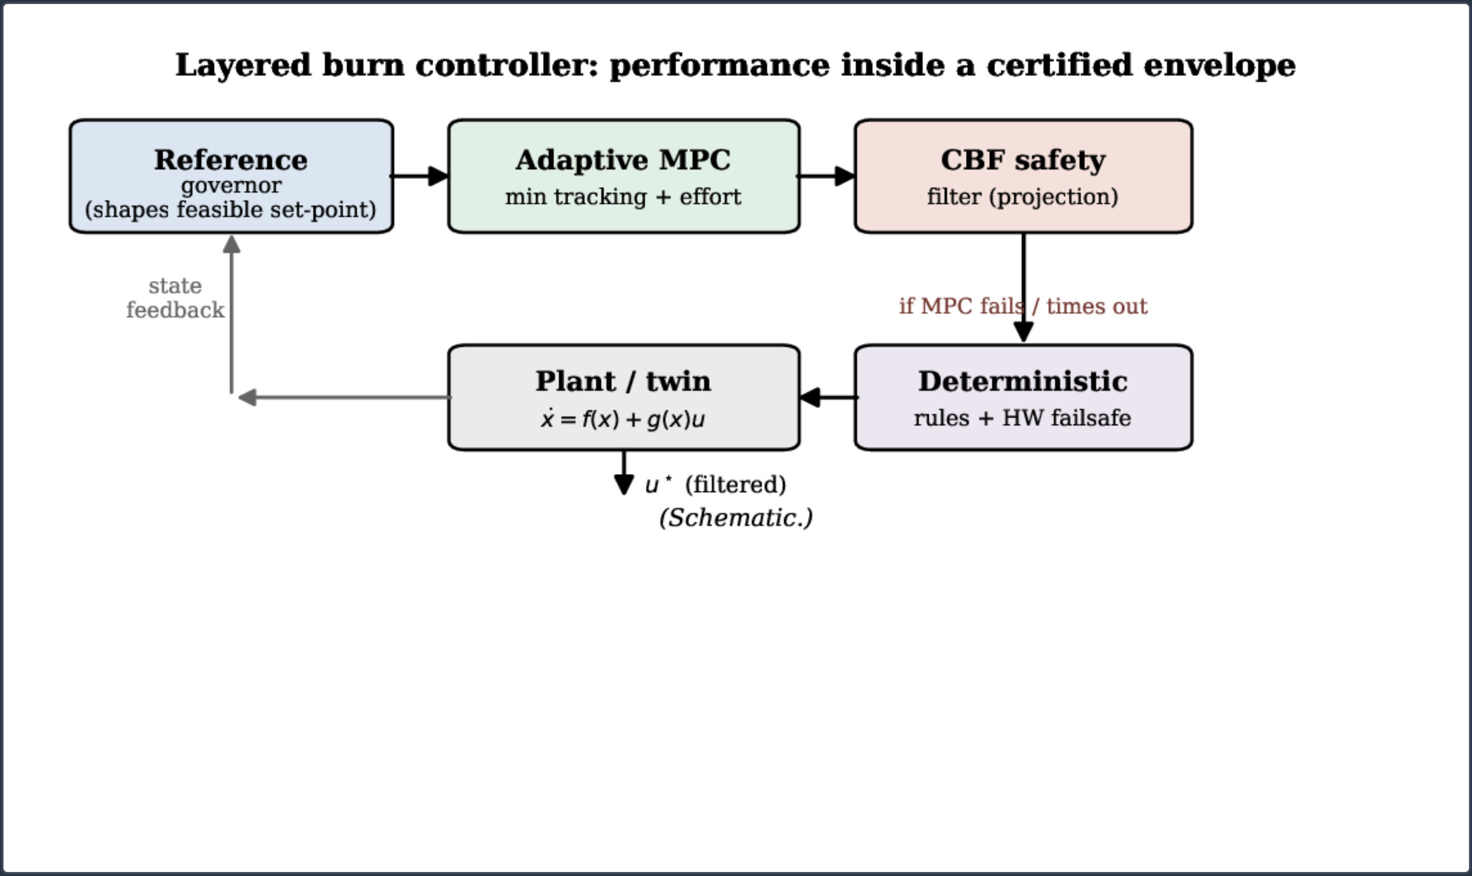

(Schematic.) The layered burn controller. The reference governor shapes a feasible set-point; the adaptive MPC optimises the actuator plan against the live transport model; the CBF filter projects the command onto the certified safe half-space; and a deterministic rules layer plus hardware failsafe stand beneath the optimiser.

LAYER	ACCEPTANCE INEQUALITY	RESULT	STATUS
CBF safety clamp	$h(x(t)) \geq 0 \;\; \forall t$	$\min h = +0.000; 0/100$	demonstrated
transport predictor	held-out $R^2 \geq 0.99$	$R^2 = 0.9983$	demonstrated
transport predictor	$\mathrm{RMSE} \leq 0.10$ in Q	$\mathrm{RMSE} = 0.088$	demonstrated
MPC burn control	$\lvert Q - 3.076 \rvert \leq \mathrm{tol}$ under disturbance	held	specified
actuator authority	$u \in [u_{\min}, u_{\max}]$, floor untouched	$u \in [1.061, 1.250]$	demonstrated
reference governor	constraints held while tracking	held in sim	specified

Acceptance criteria as inequalities, and their demonstrated status. “Demonstrated” denotes a reproduced de-risking anchor; “specified” denotes an explicit criterion validated on the reduced-order twin.

ELEMENT	FUNCTION	SERIAL
reference governor	shapes a feasible set-point	KFE–CF–BR–050
adaptive MPC	tracks a drifting operating point	KFE–CF–BR–045
MPC burn controller	holds Q at $\mathbf{3.076}$ under disturbance	KFE–CF–BR–009
real-time QP solver	solves the QP within the control period	KFE–CF–BR–112
CBF safety clamp	renders $\{\beta_N \leq \mathbf{3.5}\}$ invariant	KX–L1–A13
transport surrogate	fast, accurate predictor for the MPC	BR–L2–A12

The layered controller stack and the de-risking serials that own each element. KX–L1–A13 and BR–L2–A12 are reproduced anchors; the remaining rows are twin-first control capabilities behind the certified clamp.

§ 8

Discussion

QUANTITATIVE COMPARISON WITH PUBLISHED CONTROL

The demonstrated certificate is best read against the two established routes to plasma control. Learned magnetic control on TCV showed that a single reinforcement-learned policy can shape a plasma [7], and deep reinforcement learning has regulated the normalised beta on KSTAR by feedforward actuation [9]; both are performance achievements without a hard, checkable safety guarantee attached. Model-based predictive profile control on TCV demonstrated that a receding-horizon optimiser can drive current-profile targets in a real device [8]; this is precisely the performance layer we adopt, and our contribution is to bound it with a certificate. On the safety side, control-barrier methods have become the standard tool for rendering a safe set forward-invariant in robotics and automotive control [13,14]; we transplant that machinery to the β_N-limit problem and quantify its effect: the same disturbance sequence produces $50/100$ barrier crossings without the clamp and $0/100$ with it, the peak excursion falling from $\beta_N = 4.20$ to exactly the limit 3.500. The conservative limit $\beta_N \le 3.5$ itself sits below the Troyon-type ideal ceiling [22] established for this equilibrium in the companion stability study [29], so the certified set is a genuinely safe operating region rather than a nominal one.

WHY SEPARATION IS THE RIGHT STRUCTURE

Separating performance from safety is what lets the burn controller be both aggressive and provably safe. The model-predictive layer inherits a fast, $R^2 = 0.9983$ transport predictor and may optimise freely for tracking; any command it proposes is filtered through Eq. 11, so an imperfect prediction — or an out-of-distribution input to a learned component — cannot produce an unsafe action. The guarantee is *structural*: it depends only on the barrier condition, not on the quality of any learned model, which is exactly the property that makes it re-checkable and portable. This is the same design principle underlying the certified safety filter of the companion control-barrier study [30], the calibrated-uncertainty and abstention layer that hands authority back to a deterministic floor out of distribution [34], and the four-twin shadow that supplies the predictive model [32]. The controller here is one member of a real-time suite for the breeder — profile, shape, vertical-stability, q-profile and ramp controllers — all of which sit behind the same certified clamp [33]; the uniform composition rule (governor shapes feasible requests, a model-predictive or learned layer optimises, barrier certificate guarantees invariance) lets a single safety argument cover a heterogeneous set of performance controllers.

§ 9

Limitations

The demonstrated certificate is established on a one-dimensional, design-point-anchored β_N model (Eq. 3), integrated by explicit Euler over a 100-step sequence. Extending the identical barrier construction to the coupled multi-actuator plant — where β_N, shape, density and current are driven together and the Lie derivatives $L_f h, L_g h$ become vector quantities — on a hardware-in-the-loop bench is the one open engineering item. It reuses the clamp condition Eq. 9 unchanged; what must be built and validated is the multi-input projection and its real-time solve, not a new certificate. We flag this single gate honestly and note that the safety guarantee, being structural, transfers with the construction rather than having to be re-earned.

§ 10

Conclusion

The Hyperion burn controller places a model-predictive performance layer inside a certified safe envelope. The control-barrier-function clamp renders the safe set $\{\beta_N \leq 3.5\}$ forward-invariant — $0/100$ violations against a disturbance sequence that drives an unclamped controller to $\beta_N = 4.20$ on $50/100$ steps — with confinement authority to spare, and the model-predictive controller holds Q at the frozen 3.076 using a demonstrated $R^2 = 0.9983$ transport predictor. A reference governor keeps the request feasible and an adaptive layer tracks under drift, all beneath a deterministic backstop. Performance is model-predictive; safety is a certificate; and the two are composed through a projection so that no prediction error can produce an unsafe action. The single open item is the hardware-in-the-loop extension of the identical barrier condition to the coupled multi-actuator plant.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos control and computing teams for the GPU HPC campaign that produced the transport dataset, the surrogate, and the certified-clamp traces.

Reproducibility. The certified-clamp, transport-surrogate and reduced-order anchors (KX-L1-A13, BR-L2-A12, BR-L1-A1) are frozen entries in the Kronos Physics De-Risking Register doi:10.5281/zenodo.22133057 and are regenerable from the clamp trace and surrogate archives. This paper is archived at its DOI doi:10.5281/zenodo.22645807 and its official page https://www.kronosfusionenergy.com/publications/model-predictive-burn-control-inside-a-certified-safe-e.html. The figures are regenerated from a single self-contained script driven only by the frozen anchors, with fixed seed 20260726. The reduced-order control model runs on the in-house HPC node; the surrogate training and transport sweep run on the NVIDIA GPU HPC campaign.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645807; the register of record is doi:10.5281/zenodo.22133057. The clamp trace, surrogate samples and figure-generation script are archived with the deposit and reference the Kronos 2026 series, including the AI and quantum-control paper doi:10.5281/zenodo.22645702.

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 GRANTED PATENT	4 2026 PROVISIONALS	40 TEAM MEMBERS	27 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 C E Garcia, D M Prett and M Morari, Model predictive control: theory and practice—a survey, *Automatica* **25** 335 (1989).

2 D Q Mayne, J B Rawlings, C V Rao and P O M Scokaert, Constrained model predictive control: stability and optimality, *Automatica* **36** 789 (2000).

3 S J Qin and T A Badgwell, A survey of industrial model predictive control technology, *Control Eng. Practice* **11** 733 (2003).

4 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

5 L L Lao, H St John, R D Stambaugh, A G Kellman and W Pfeiffer, Reconstruction of current profile parameters and plasma shapes in tokamaks, *Nucl. Fusion* **25** 1611 (1985). doi:10.1088/0029-5515/25/11/007

6 F Felici *et al*, Real-time physics-model-based simulation of the current density profile in tokamak plasmas, *Nucl. Fusion* **51** 083052 (2011). doi:10.1088/0029-5515/51/8/083052

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

8 E Maljaars *et al*, Profile control simulations and experiments on TCV: a controller test environment and results using a model-based predictive controller, *Nucl. Fusion* **57** 126063 (2017).

9 J Seo *et al*, Feedforward beta control in the KSTAR tokamak by deep reinforcement learning, *Nucl. Fusion* **61** 106010 (2021).

10 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

11 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

12 D Humphreys *et al*, Novel aspects of plasma control in ITER, *Phys. Plasmas* **22** 021806 (2015). doi:10.1063/1.4907901

13 A D Ames, X Xu, J W Grizzle and P Tabuada, Control barrier function based quadratic programs for safety critical systems, *IEEE Trans. Autom. Control* **62** 3861 (2017). doi:10.1109/TAC.2016.2638961

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

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

16 E Gilbert and I Kolmanovsky, Nonlinear tracking control in the presence of state and control constraints: a generalized reference governor, *Automatica* **38** 2063 (2002).

17 I Kolmanovsky, E Garone and S Di Cairano, Reference and command governors: a tutorial on their theory and automatic implementation, *Proc. American Control Conf. (ACC)* 226 (2014). doi:10.1109/ACC.2014.6858785

18 B Stellato, G Banjac, P Goulart, A Bemporad and S Boyd, OSQP: an operator splitting solver for quadratic programs, *Math. Program. Comput.* **12** 637 (2020). doi:10.1007/s12532-020-00179-2

19 A E Hoerl and R W Kennard, Ridge regression: biased estimation for nonorthogonal problems, *Technometrics* **12** 55 (1970). doi:10.1080/00401706.1970.10488634

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

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

22 F Troyon, R Gruber, H Saurenmann, S Semenzato and S Succi, MHD-limits to plasma confinement, *Plasma Phys. Control. Fusion* **26** 209 (1984). doi:10.1088/0741-3335/26/1A/319

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

24 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

25 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

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

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

28 P I Ford, *Negative-triangularity turbulence suppression in a compact spherical-tokamak breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645722. Official page: https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html

29 P I Ford, *Free-boundary equilibrium and ideal-MHD stability of a compact spherical-tokamak breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645740. Official page: https://www.kronosfusionenergy.com/publications/ideal-mhd-stability.html

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

31 P I Ford, *Neural-operator flux surrogates for control-latency transport: differentiable DeepONet–FNO emulation of nonlinear gyrokinetics*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645803. Official page: https://www.kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html

32 P I Ford, *Four coupled digital twins for a compact spherical-tokamak breeder: multi-rate co-simulation, propagated uncertainty, and a real-time self-consistency monitor*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645805. Official page: https://www.kronosfusionenergy.com/publications/four-coupled-digital-twins-for-a-compact-spherical-toka.html

33 P I Ford, *A real-time control suite for a compact negative-triangularity spherical-tokamak breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645809. Official page: https://www.kronosfusionenergy.com/publications/a-real-time-control-suite-for-a-compact-negative-triang.html

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

COLOPHON

How this document was made

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

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

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

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

LEGAL NOTICE

Notice, status, and distribution

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

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

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

THE 2026 KRONOS SERIES

One programme, many papers

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

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

KRONOS FUSION ENERGY

Model-Predictive Burn Control Inside a Certified Safe Envelope: A Reference-Governor MPC Bounded by a Control-Barrier-Function Safety Clamp

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

Zenodo · doi:10.5281/zenodo.22645807

Read on kronosfusionenergy.com

The Physics De-Risking Register