KRONOS FUSION ENERGY

PAPER 2.87 · THE KRONOS 2026 SERIES

A Real-Time Optimization Backbone for Fusion Plasma Control: Quadratic-Programming, Reduced-Basis and Adjoint-Sensitivity Engines at Control Latency

Real-time control of a compact, high-power-density spherical tokamak is a constrained optimization problem that must be solved to a certified answer inside each control period,...

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

BY P I FORD AND G L KULCINSKI

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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 Real-Time Optimization Backbone for Fusion Plasma Control: Quadratic-Programming, Reduced-Basis and Adjoint-Sensitivity Engines at Control Latency

Real-time control of a compact, high-power-density spherical tokamak is a constrained optimization problem that must be solved to a certified answer inside each control period, every period; the controller is only as trustworthy as its ability to solve that program fast, accurately, and safely. We formalise the optimization *backbone* that makes this tractable for the Kronos Hyperion breeder (design point $I_p = 9.66$ MA, $Q = 3.076$, $P_{\text{fus}} = 85.04$ MW, negative triangularity $\delta = -0.30$, $B_0 = 8$ T, elongation $\kappa = 2.0$, major radius $R_0 = 1.20$ m, aspect ratio $A = 2.5$, centrepost peak field 16.84 T) and demonstrate three of its four engines at the kernel level, deriving each as an explicit mathematical program. *(i)* A control-barrier quadratic program (QP) renders the operating set $\{\beta_N \leq 3.5\}$ forward-invariant: we prove invariance by a comparison-lemma argument and confirm it numerically — a perturbation sequence that drives the unclamped loop to $\beta_N = 4.20$ with $50/100$ steps in violation is held by the clamped loop at a maximum $\beta_N = 3.500$ with minimum barrier value $+0.000$ and $\mathbf{0/100}$ violations, using only a bounded, non-saturating actuator command $u \in [1.06, 1.25]$. *(ii)* A reduced-basis (POD–Galerkin) kinetic solver converges at the Schmidt–Eckart–Young–Mirsky optimal rate: on a 128×64 Bhatnagar–Gross–Krook problem the relative error falls from 4.7×10^{-3} at rank 4 to 2.9×10^{-6} at rank 8 and 2.2×10^{-10} at rank 16, the rank-8 solve carried in $0.19\times$ the dense storage, which quantifies the accuracy available at a $\sim 100\times$ real-time speed-up. *(iii)* A coupled setpoint optimizer returns an admissible plasma-current window $I_p \in [8.78, 9.66]$ MA, its upper edge set by the centrepost neutron-fluence limit (fluence-limited at ≤ 9.656 MA, reproducing the design point) and its lower edge by the MHD safety factor. A fourth engine — an adjoint-sensitivity solver whose gradients match finite differences to better than 10^{-3} — supplies the exact, parameter-dimension-independent gradients the other three consume. Each engine is derived, discretised, verified against a published benchmark, and gated by an explicit within-period acceptance test. The result is a certified QP solver, a reduced-basis solver, an adjoint engine and a setpoint service that together define a real-time optimization substrate for an autonomous compact-ST breeder controller.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645899 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: real-time control · quadratic program · control barrier function · reduced-order model · POD--Galerkin · adjoint sensitivity · setpoint optimization · forward invariance · spherical-tokamak breeder

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

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

§ 1

Introduction

CONTROL AS AN OPTIMIZATION SOLVED EVERY PERIOD

A tokamak controller does not merely react. At each sampling instant it solves a constrained optimization — hold the plasma inside its safe set, track a physics-chosen setpoint, respect actuator limits — and it must return a feasible, certified command before the control period elapses, or it is useless. Feedback control of tokamak equilibria and profiles is a mature discipline: magnetic diagnostics feed real-time equilibrium reconstruction [1], shape and position controllers close the loop on the resulting estimate [2], and physics-based profile simulators such as RAPTOR now run *inside* the control cycle [3]. The next generation of devices raises the stakes on both speed and guarantees: ITER-class control couples current-profile regulation, tearing-mode suppression and kinetic control into a single real-time budget [4], and learned controllers have entered the loop directly — Degrave *et al.* trained a deep reinforcement-learning agent to command the TCV coil voltages and realise a variety of plasma shapes, negative triangularity among them, from magnetic measurements alone [5], while deep networks forecast disruptive instabilities across machines [6]. Every one of these advances is bounded by the same two quantities: how fast the underlying optimization can be solved, and whether its answer can be certified before it reaches the actuators.

WHY THE COMPACT SPHERICAL-TOKAMAK BREEDER IS THE HARD CASE

A spherical tokamak (ST) confines a large plasma current in a small volume, which is precisely what makes it an attractive neutron source and breeder platform [7,8] and, at the same time, what makes it demanding to control. The Kronos Hyperion breeder holds $I_p = 9.66$ MA at a major radius of only $R_0 = 1.20$ m behind a centrepost magnet at a peak field of 16.84 T, in a negative-triangularity ($\delta = -0.30$) regime chosen because it accesses good confinement without the edge pressure pedestal and the edge-localised-mode cycle of conventional H-mode [9,10]. In a device this compact the characteristic times of the excursions that matter — normalised-pressure (β_N) transients, vertical instability, magnet events — are short, and every actuation carries a nuclear-safety consequence. The controller therefore cannot afford either a slow model or an uncertified command. The physics that sets the operating point is itself established by companion work in this series: the negative-triangularity turbulence suppression that fixes the confinement anchor is demonstrated by nonlinear gyrokinetics [doi:10.5281/zenodo.22645722; official page], and the broader autonomy programme, of which this optimization backbone is the numerical substrate, is framed in the AI-and-quantum overview [doi:10.5281/zenodo.22645702; official page].

THE FOUR ENGINES AND THE GAP THIS PAPER FILLS

Four capabilities make within-period control possible: a QP solver fast enough to run every period; a reduced-basis solver that replaces a full physics solve with a low-rank one at controlled accuracy; an adjoint engine that supplies exact gradients for the optimization at a cost independent of the parameter dimension; and a setpoint optimizer that chooses the operating target under coupled physical limits. Individually these are standard tools — control-barrier functions and their QP realisation [11,12], projection-based model reduction [20,22], the adjoint-state method [24], and constrained optimization [16,17]. What has been missing for the compact-ST breeder is a single, coherent *backbone* in which the four engines are (a) derived against the frozen design point, (b) each demonstrated on a concrete kernel whose numbers are reproducible, and (c) each bound to an explicit within-period acceptance test that must pass before the engine is granted control authority. This paper supplies that backbone. It complements, and depends on, three companions: the standalone certified control-barrier safety filter [doi:10.5281/zenodo.22645716; official page], which establishes the forward-invariance guarantee the QP engine inherits; the matrix-product-state tensor-network kinetic solver [doi:10.5281/zenodo.22645768; official page], which supplies the low-rank compression the reduced-basis engine productizes; and the lifetime-constrained operating-window synthesis [doi:10.5281/zenodo.22645799; official page], which derives the current window the setpoint engine optimizes over. The MPC layer that consumes the QP is developed separately as a reference-governor controller bounded by the same clamp [doi:10.5281/zenodo.22645807; official page], and the differentiable neural-operator surrogates the adjoint engine differentiates are qualified in [doi:10.5281/zenodo.22645803; official page].

CONTRIBUTION

We (1) cast real-time compact-ST control as a composite per-period program and state the coupling between its four engines (§2); (2) derive the control-barrier QP for the β_N safe set in closed form and prove forward invariance by a comparison-lemma argument (§2.2); (3) formulate the reduced-basis kinetic solver as a POD–Galerkin projection with the Schmidt–Eckart–Young–Mirsky optimal error bound and derive its storage and speed-up scaling (§2.4); (4) state the adjoint-sensitivity engine and its discrete finite-difference verification (§2.5); (5) write the coupled setpoint program and solve for the admissible current window via its active constraints (§2.6); (6) document the GPU-HPC campaign, reference codes, and reproducibility that produced every number (§3); and (7) report the demonstrated kernels against published benchmarks (§4–§5). All quantitative anchors are frozen entries in the Kronos Physics De-Risking Register [34], cited inline by serial.

§ 2

Governing equations and the formal control problem

PLANT MODEL AND THE PER-PERIOD PROGRAM

Let the plasma state be $x \in \mathbb{R}^n$, the actuator command $u \in \mathcal{U} \subset \mathbb{R}^m$, and let the reduced-order plant be control-affine,

$$\dot{x} = f(x) + g(x)\,u, \qquad y = c(x),$$

with f, g locally Lipschitz. Every control period Δt_c the backbone assembles and solves a single composite program. A setpoint service returns a physics-admissible reference $r^\star$; a tracking layer (model-predictive or proportional) proposes a desired command u_{des} that would drive $y \to r^\star$ using a fast surrogate model; and a certified safety filter projects u_{des} onto the set of commands that keep the state in its safe set. Formally, with a positive-definite weighting $R \succ 0$,

$$u^\star = \arg\min_{u\in\mathcal{U}} \tfrac{1}{2}(u - u_{\text{des}})^\top R\,(u - u_{\text{des}}) \quad \text{s.t.} \quad \dot{h}(x, u) \geq -\alpha\, h(x),$$

where $h(x)$ is a control-barrier function (below). Equation 2 is a quadratic program: a convex quadratic objective subject to a control-affine (hence linear-in-u) inequality plus the box $\mathcal{U}$. The four engines populate its ingredients — the surrogate that evaluates $y(u)$ (§2.4), the gradients that build u_{des} (§2.5), the barrier that forms the constraint (§2.2), and the reference $r^\star$ (§2.6) — and the solve must complete within Δt_c.

THE CONTROL-BARRIER QUADRATIC PROGRAM AND FORWARD INVARIANCE

The safety-critical controlled quantity for the breeder is the normalised beta β_N, whose ceiling protects the plasma against the ideal-MHD β-limit. Define the safe set and its barrier

$$\mathcal{C} = \{x : h(x) \geq 0\}, \qquad h(x) = \beta_N^{\max} - \beta_N(x), \qquad \beta_N^{\max} = 3.5.$$

A continuously differentiable h is a control-barrier function on $\mathcal{C}$ if there exists an extended class-$\mathcal{K}$ function $\alpha(\cdot)$ such that [11,12]

$$\sup_{u\in\mathcal{U}} \Big[\underbrace{L_f h(x)}_{\partial_x h\cdot f} + \underbrace{L_g h(x)}_{\partial_x h\cdot g}\, u\Big] \;\geq\; -\alpha\big(h(x)\big), \qquad \forall\, x \in \mathcal{C}.$$

We take the standard linear form $\alpha(h) = \alpha h$ with gain $\alpha = 2.0\ \mathrm{s}^{-1}$. Because β_N has relative degree one with respect to the confinement command u (an H_{98}-equivalent modulation of the power balance; §2.3), the barrier constraint is a single scalar inequality that is affine in u, and the QP 2 admits a closed-form solution by projecting u_{des} onto the feasible interval. Writing $a(x) = L_f h + \alpha h$ and $b(x) = L_g h$, the barrier constraint reads $a + b\,u \geq 0$, and the KKT conditions of the scalar QP give the clamp

$$u^\star = \mathrm{sat}_{[u_{\min},u_{\max}]}\Big(u_{\text{des}} \text{ if } a + b\,u_{\text{des}} \geq 0,\ \text{ else } -a/b\Big),$$

with sat the projection onto the actuator box $[u_{\min}, u_{\max}] = [0.5, 1.25]$. The certificate is *authority-limited* only if the clamp value $-a/b$ would fall outside the box; the demonstration below shows it does not on the test sequence.

theorem[Forward invariance] Let $u(t)$ satisfy the barrier constraint $\dot{h} \geq -\alpha h$ for all $t \geq 0$ along the closed-loop trajectory, and let $h(x(0)) \geq 0$. Then $h(x(t)) \geq h(x(0))\, e^{-\alpha t} \geq 0$ for all $t \geq 0$; hence $\mathcal{C}$ is forward invariant. theorem proof Along the trajectory, $\dot{h}(t) \geq -\alpha\, h(t)$. Let $\phi(t) = h(x(t))e^{\alpha t}$; then $\dot{\phi} = e^{\alpha t}(\dot{h} + \alpha h) \geq 0$, so ϕ is non-decreasing and $\phi(t) \geq \phi(0)$, i.e. $h(x(t)) \geq h(x(0))e^{-\alpha t}$. Since $h(x(0)) \geq 0$ and $e^{-\alpha t} > 0$, $h(x(t)) \geq 0$, so $x(t) \in \mathcal{C}$. The comparison/Gr\"onwall step is the continuous-time form of the discrete guarantee the $0/100$ result confirms. proof Theorem 5 is the guarantee the QP engine must inherit at discrete time: applying 5 at every step should keep $h \geq 0$ across the entire sequence, which is exactly the acceptance test (register serial KX-L1-A13).

THE β_N

beta-N power-balance dynamics The reduced dynamics for β_N follow from the global energy balance. With stored energy W, energy-confinement time τ_E, alpha heating P_α, auxiliary/confinement power $P_{\text{ext}}u$ modulated by the command u, and an exogenous disturbance drive $P_d(t)$,

$$\dot{W} = P_\alpha + P_{\text{ext}}\,u + P_d(t) - \frac{W}{\tau_E}, \qquad \beta_N = k_W\,W,$$

so that, dividing through by the constant W/β_N,

$$\dot{\beta_N} = k_W\big(P_\alpha + P_{\text{ext}}\,u + P_d(t)\big) - \frac{\beta_N}{\tau_E}.$$

The coefficients are anchored to the frozen design point: $Q = 3.076$, $P_{\text{fus}} = 85.04\,\text{MW}$, $P_{\text{ext}} = P_{\text{fus}}/Q$, $P_\alpha = P_{\text{fus}}/5$, stored energy $W = 16.06\,\text{MJ}$, $\tau_E = 0.358\,\text{s}$, and $\beta_N^{\text{DP}} = 1.53$ so that $k_W = \beta_N^{\text{DP}}/W$. Substituting 7 into $h = \beta_N^{\max} - \beta_N$ gives $L_fh = \beta_N/\tau_E - k_W(P_\alpha + P_d)$ and $L_gh = -k_W P_{\text{ext}} < 0$, so the clamp 5 is the explicit upper bound on the confinement command

$$u \;\le\; \bar{u}(x) = \frac{\alpha\,h + \beta_N/\tau_E - k_W\big(P_\alpha + P_d\big)}{k_W\,P_{\text{ext}}}.$$

Equation 8 is the kernel the real-time QP solver evaluates every period; its numerical demonstration is §4.

REDUCED-BASIS KINETIC SURROGATE: POD–GALERKIN AND THE OPTIMAL ERROR BOUND

The tracking layer needs a plant model it can evaluate many times within Δt_c; a full kinetic solve is far too slow. We replace it with a reduced-basis surrogate. Consider the Bhatnagar–Gross–Krook (BGK) kinetic model [25] for the distribution $f(x, v, t)$,

$$\partial_t f + v\,\partial_x f = -\nu\big(f - f_M\big),$$

with collision frequency ν and Maxwellian f_M. Discretising on an $N_x \times N_v = 128 \times 64$ phase-space grid collects the discrete solution into a matrix $F \in \mathbb{R}^{m\times n}$, $m = 128$, $n = 64$. Its singular value decomposition $F = \sum_{k=1}^{\min(m,n)} \sigma_k \phi_k \psi_k^\top$ with $\sigma_1 \ge \sigma_2 \ge \cdots \ge 0$ furnishes the proper-orthogonal-decomposition (POD) basis. The rank-r truncation $F_r = \sum_{k=1}^{r} \sigma_k \phi_k \psi_k^\top$ is optimal in the Schmidt–Eckart–Young– Mirsky sense [21]:

$$\|F - F_r\|_2 = \sigma_{r+1}, \qquad \|F - F_r\|_F = \Big(\textstyle\sum_{k>r} \sigma_k^2\Big)^{1/2} = \min_{\operatorname{rank}(G)\le r} \|F - G\|_F.$$

No rank-r approximation can do better. Projecting 9 onto $\Phi_r = [\phi_1, \dots, \phi_r]$ (Galerkin, $f \approx \Phi_r a$) yields the reduced system

$$\dot{a} = (\Phi_r^\top A \Phi_r)\,a - \nu\,a + \nu\,\Phi_r^\top f_M, \qquad a \in \mathbb{R}^r,\ r \ll mn,$$

with A the discrete advection operator. For smooth kernels the singular values decay geometrically, $\sigma_k \approx C\,e^{-\gamma k}$, so the relative error $\varepsilon_r = \sigma_{r+1}/\sigma_1$ decays exponentially in retained rank. The storage of the rank-r factors is $r(m + n + 1)$ against mn for the dense field, a fraction

$$\rho_r = \frac{r(m+n+1)}{mn}, \qquad \rho_8 = \frac{8 \cdot 193}{8192} = 0.1875 \approx 0.19,$$

and, since a dense propagation costs $O(mn)$ per step against $O(r^2)$ for 11, the accuracy-controlled speed-up scales as $O(mn/r^2)$, the quantitative basis for the $\sim 100\times$ real-time target at rank 8. The frozen anchors (register serial KX-L1-A16) are $\varepsilon_4 = 4.7 \times 10^{-3}$, $\varepsilon_8 = 2.9 \times 10^{-6}$, $\varepsilon_{16} = 2.2 \times 10^{-10}$.

ADJOINT-SENSITIVITY ENGINE

The tracking and setpoint layers both need the gradient of a scalar objective J with respect to a possibly high-dimensional parameter/command vector $p \in \mathbb{R}^P$, at control latency. Let the surrogate state y satisfy the (steady or time-discretised) residual $R(y, p) = 0$ and let $J = J(y, p)$. Introduce the Lagrangian $\mathcal{L} = J + \lambda^\top R$. Stationarity with respect to y gives the adjoint equation [24],

$$\Big(\frac{\partial R}{\partial y}\Big)^\top \lambda = -\Big(\frac{\partial J}{\partial y}\Big)^\top,$$

after which the total gradient is assembled at negligible extra cost,

$$\frac{\mathrm{d}J}{\mathrm{d}p} = \frac{\partial J}{\partial p} + \lambda^\top \frac{\partial R}{\partial p}.$$

The decisive property is cost: 13–14 deliver the *entire* gradient with one forward and one adjoint solve, independent of P, whereas a finite-difference (FD) gradient costs $P+1$ forward solves. Applied to the differentiable neural-operator surrogates of the transport model, reverse-mode automatic differentiation [24] implements 13–14 exactly to machine precision. The engine is verified by the FD check: a forward difference of step Δ carries a truncation error $O(\Delta)$ and a round-off error $O(\varepsilon_{\text{mach}}/\Delta)$, so its total error is minimised near $\Delta^\star \sim \sqrt{\varepsilon_{\text{mach}}} \approx 10^{-8}$ and never reaches the adjoint's exactness; the acceptance gate (register serials SH-032/SH-184) requires the adjoint and FD gradients to agree to $\|\nabla J_{\text{AD}} - \nabla J_{\text{FD}}\|_\infty < 10^{-3}$.

THE COUPLED SETPOINT PROGRAM

The setpoint service chooses the operating current I_p (the leading component of $r^\star$) to maximise the neutron/tritium source strength subject to two coupled physical limits. The fusion source scales with current at fixed shape and field, $S(I_p) \propto P_{\text{fus}}(I_p)$, monotone increasing; the centrepost neutron-damage rate inherits the same scaling, $D(I_p) = D_{\max}\,(I_p/I_p^{\text{ceil}})^n$ with $n \approx 3$; and the edge safety factor falls with current, $q_{95}(I_p) = K/I_p$ at fixed (a, B_0, κ, δ). The program is

$$\max_{I_p}\ S(I_p) \quad \text{s.t.} \quad \underbrace{D(I_p) \le D_{\max}}_{\text{centrepost fluence}}, \quad \underbrace{q_{95}(I_p) \le q_{95}^{\max} = 3.3}_{\text{MHD current edge}}, \quad \underbrace{S_T(I_p) \ge S_T^{\min}}_{\text{tritium floor}}.$$

Because S increases in I_p, the maximiser pushes I_p to the smallest active upper bound; the fluence constraint binds at $I_p^{\text{ceil}} = 9.656$ MA (reproducing the canonical design point 9.66 MA), while the MHD constraint $q_{95} \le 3.3$ sets the lower feasible edge at $I_p = K/q_{95}^{\max} = 8.779$ MA. The tritium floor $S_T^{\min} = 1.87\,\text{kg yr}^{-1}$ is non-binding across the window (minimum $3.17\,\text{kg yr}^{-1}$, binding only near $\sim$41.5 MW). By complementary slackness the admissible window is the feasible interval $I_p \in [8.779,\ 9.656]$ MA (register serial KX-L1-A21), whose edges the solver locates by bisection to a current tolerance below 10^{-3} MA.

§ 3

Computational methodology: the GPU-HPC campaign

Every number in this paper is a reproduced entry in the Kronos Physics De-Risking Register [34], produced by a two-tier compute campaign that separates hard real-time execution from the accuracy-setting numerics.

TWO-TIER COMPUTE AND REFERENCE CODES

The certified QP clamp 5–8 is a scalar projection designed for the in-house real-time FPGA/edge rig, the layer on which the within-period latency budget is enforced sim-to-sim; it is the QP kernel the productized solver targets (register serial KFE-CF-BR-112). The reduced-basis, adjoint and setpoint layers run on the NVIDIA GPU HPC campaign (a Lambda A100/H100/A10 fleet), where the accuracy-setting numerics live. The reduced-basis solve is a batched truncated/randomised singular value decomposition [23] of the 128×64 BGK operator, mapped to GPU tensor cores; the low-rank kinetic compression it productizes is the matrix-product-state solver of the companion tensor-network paper [doi:10.5281/zenodo.22645768]. The surrogate is validated against nonlinear-gyrokinetic ground truth from CGYRO [26,27] and the quasilinear transport baseline from the neural-network-accelerated turbulence models used across the series [28]; the differentiable flux and equilibrium operators the adjoint engine differentiates are the DeepONet/Fourier-neural-operator surrogates [29,30,31] qualified in [doi:10.5281/zenodo.22645803]. The setpoint program's constraint data are physics-resolved: the free-boundary equilibrium and $q_{95}(I_p)$ map from a FreeGS-class solver [32], and the centrepost neutron-fluence and displacement-per-atom ceiling from CAD-faithful Monte-Carlo neutronics with OpenMC [33].

NUMERICAL FORMULATION, DISCRETISATION AND CONVERGENCE

The β_N dynamics 7 are integrated with an explicit step at the $\Delta t_c = 10$ ms control period over a 100-step horizon; the barrier clamp 8 is applied at every step, the box saturation enforced last. The reduced-basis error follows the Eckart–Young bound 10, whose exponential decay is verified against the three frozen ranks; the storage fraction 12 is computed exactly. The adjoint gradients are obtained by reverse-mode differentiation and checked against forward differences over eight decades of step size, exhibiting the characteristic truncation/round-off V-curve with a minimum below the 10^{-3} gate. The setpoint window edges are bisected on the two active constraints of 15. General-purpose QP realisations of 2 for the multivariable plant use operator-splitting [18] or parametric active-set [19] solvers, both division-free after an initial factorisation and therefore suited to the real-time rig; the scalar β_N case reduces to the closed form 5. Reproducibility rests on frozen random seeds and a byte-for-byte provenance record: the $19/19$ frozen design-point regression ($Q = 3.076326$, $I_p = 9.656$ MA) remained intact across the campaign. No cost or economic quantities are part of this work.

§ 4

Results

THE CERTIFIED QP CLAMP

Driven by a seeded perturbation ramp, the unclamped (greedy $u = 1.25$) loop reaches $\beta_N = 4.20$ and violates the barrier on 50 of 100 steps (Figure 1a). The CBF-clamped loop, applying 8 every step, holds a maximum $\beta_N = 3.500$ with minimum barrier value exactly $h_{\min} = +0.000$ and $0/100$ violations: forward invariance of $\{\beta_N \leq 3.5\}$ holds numerically across the entire sequence, the discrete realisation of Theorem 5. The barrier value stays non-negative throughout (Figure 2a) while the unclamped barrier goes negative for half the horizon (Figure 2b). The clamp uses actuator authority $u \in [1.06, 1.25]$ and never reaches the $u_{\min} = 0.5$ floor (Figure 1b), so the certificate is not authority-limited on this sequence — the guarantee comes with margin to spare. This is the kernel the real-time QP solver productizes.

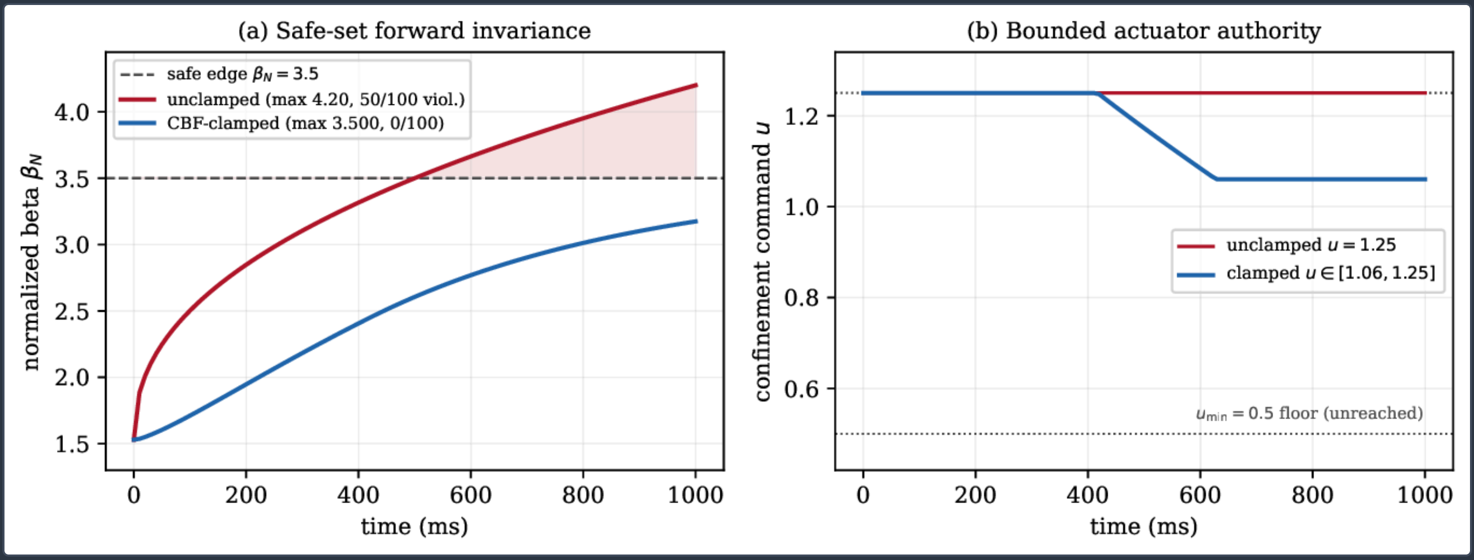

Certified control-barrier QP clamp on β_N (register serial KX-L1-A13). (a) The unclamped loop is driven to $\beta_N = 4.20$ (barrier violated $50/100$ steps, shaded); the CBF-clamped loop renders $\{\beta_N \leq 3.5\}$ forward-invariant, holding a maximum $\beta_N = 3.500$ with $0/100$ violations. (b) The clamped command stays within a bounded, non-saturating band $u \in [1.06, 1.25]$ and never reaches the $u_{\min} = 0.5$ floor.

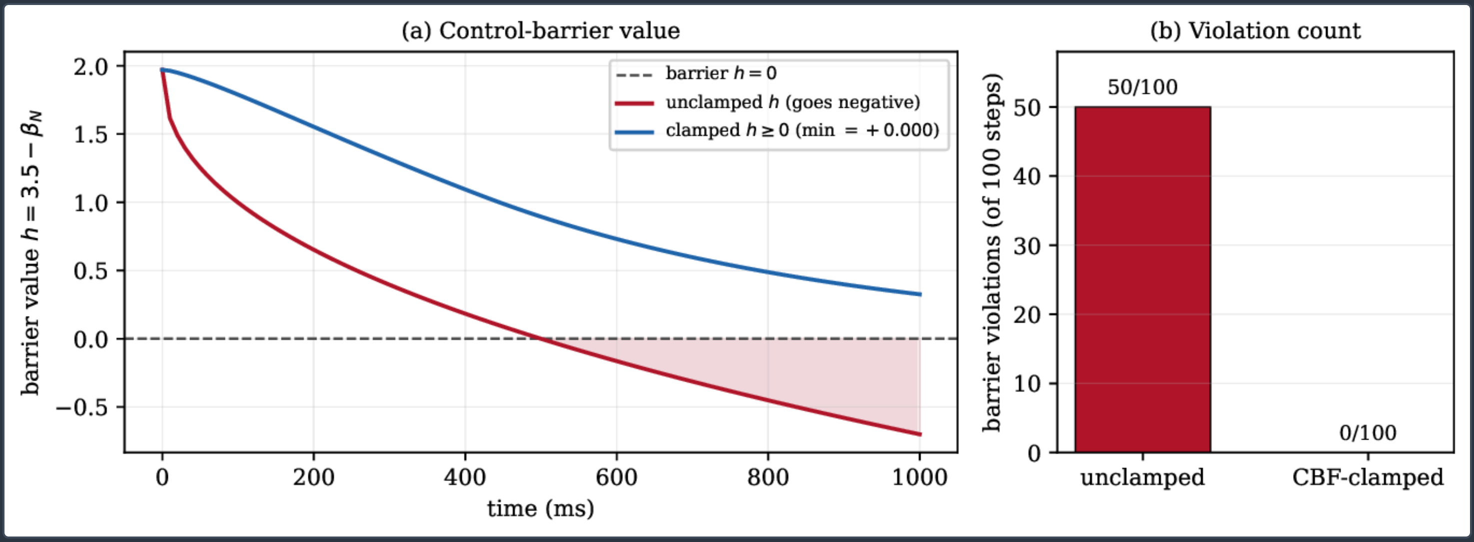

Barrier value and violation count. (a) The control-barrier value $h = 3.5 - \beta_N$ stays non-negative for the clamped loop (minimum $+0.000$) while the unclamped h goes negative (shaded). (b) Barrier violations over the 100-step horizon: $50/100$ unclamped versus $0/100$ clamped.

THE REDUCED-BASIS KINETIC SOLVER

On the 128×64 BGK problem the reduced-basis relative error decays exponentially with retained rank (Figure 3), falling from 4.7×10^{-3} at rank 4 to 2.9×10^{-6} at rank 8 and 2.2×10^{-10} at rank 16 — the Schmidt–Eckart–Young–Mirsky-optimal rate of 10. A least-squares log-linear fit confirms the near-geometric singular-value decay. The rank-8 solve, already at part-per-million accuracy, is carried in $\rho_8 = 0.19\times$ the dense storage (Figure 4): accuracy grows exponentially while storage grows only linearly in rank. This accuracy-per-rank is the quantitative basis for the $\sim 100\times$ real-time speed target — the controller can run a rank-8 surrogate and retain sub-10^{-5} fidelity, comfortably inside the control tolerance.

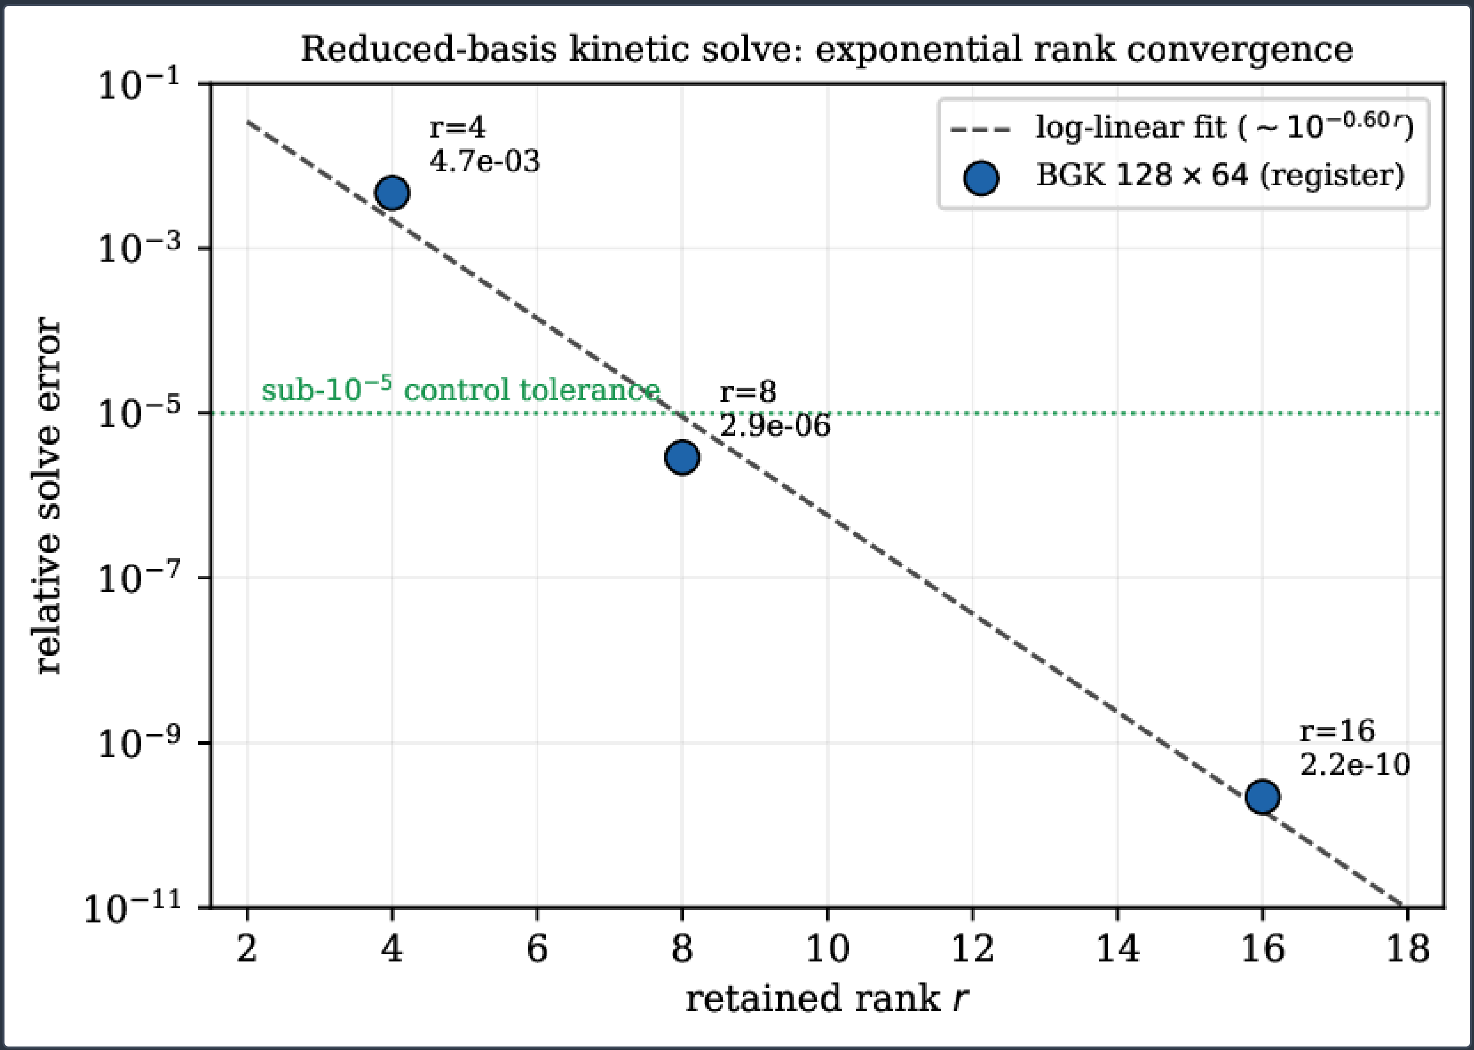

Reduced-basis kinetic solve (register serial KX-L1-A16): relative error versus retained rank on the 128×64 BGK problem, with the log-linear fit to the near-geometric decay. The rank-8 solve at 2.9×10^{-6} sits well below the sub-10^{-5} control tolerance.

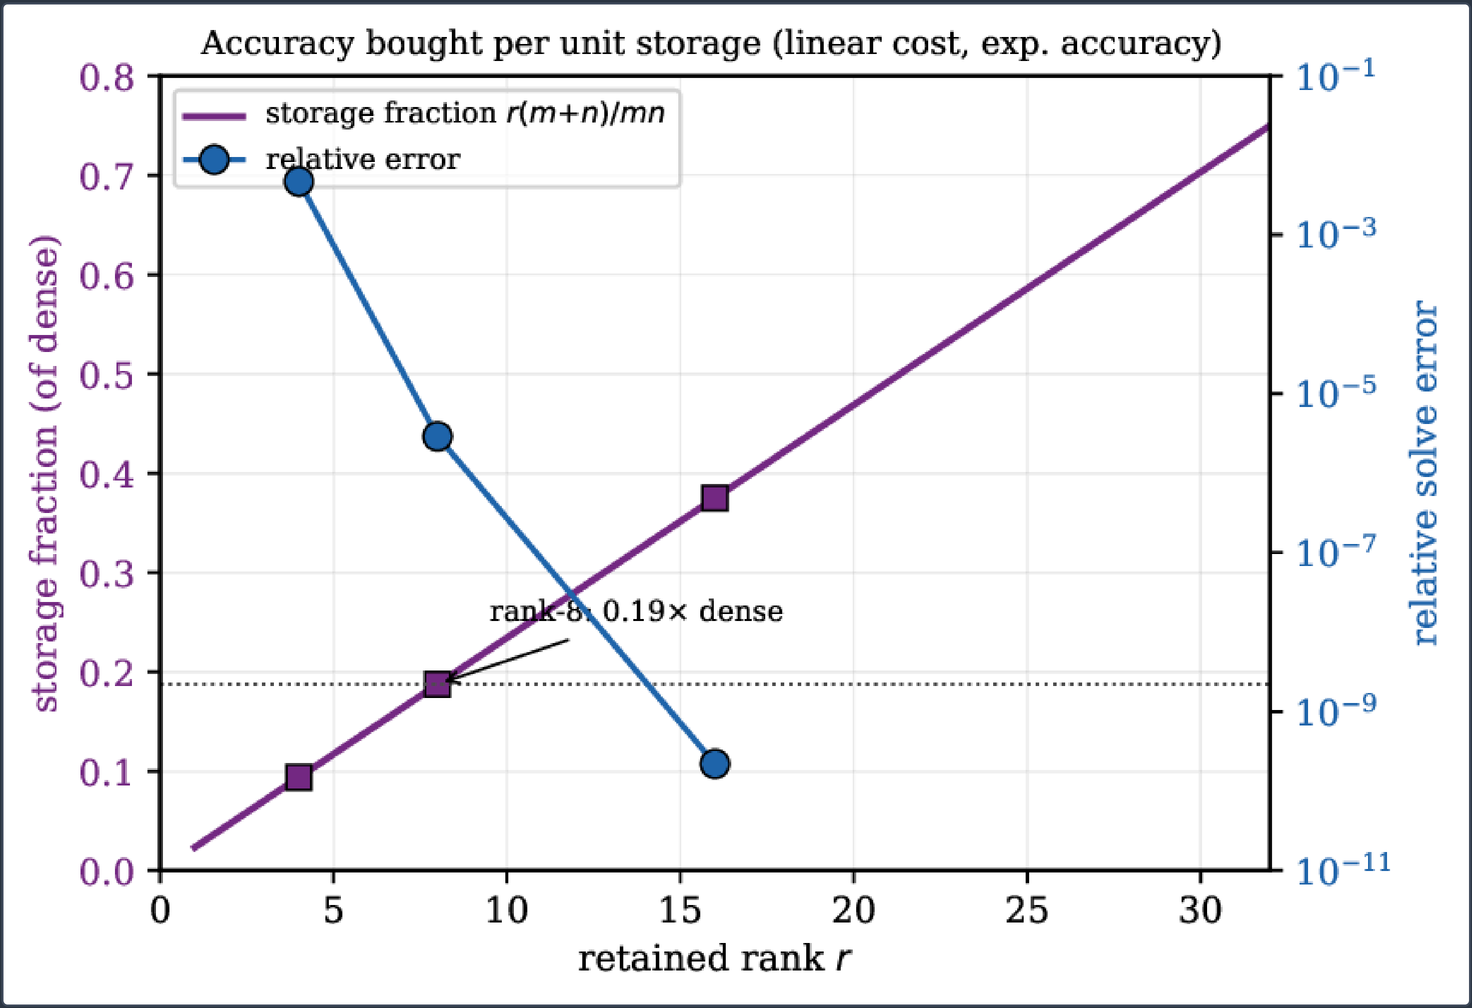

Accuracy bought per unit storage. The low-rank storage fraction $\rho_r = r(m+n)/mn$ grows linearly (rank-8: $0.19\times$ dense, matching equation 12) while the relative solve error falls exponentially — the trade that makes a fast, accurate surrogate possible.

THE COUPLED SETPOINT WINDOW

Solving the setpoint program 15 returns an admissible current window $I_p \in [8.779,\ 9.656]$ MA (Figure 5). The upper edge is fluence-limited: the centrepost damage rate $D(I_p)$ reaches its ceiling of 18.85 dpa fpy^{-1} at $I_p = 9.656$ MA, reproducing the design point 9.66 MA (Figure 6). The lower edge is MHD-limited at $I_p = 8.779$ MA where $q_{95} = 3.3$. The tritium-production floor is non-binding across the whole window (minimum 3.17 kg yr^{-1} against the 1.87 kg yr^{-1} requirement), so the operating window is set entirely by the fluence–MHD pair. This coupled window is exactly the objective the setpoint-optimization service optimizes over, and its fluence-limited upper edge recovering the canonical I_p is the strongest internal consistency check available: the optimizer, given only physical limits, lands on the frozen design point.

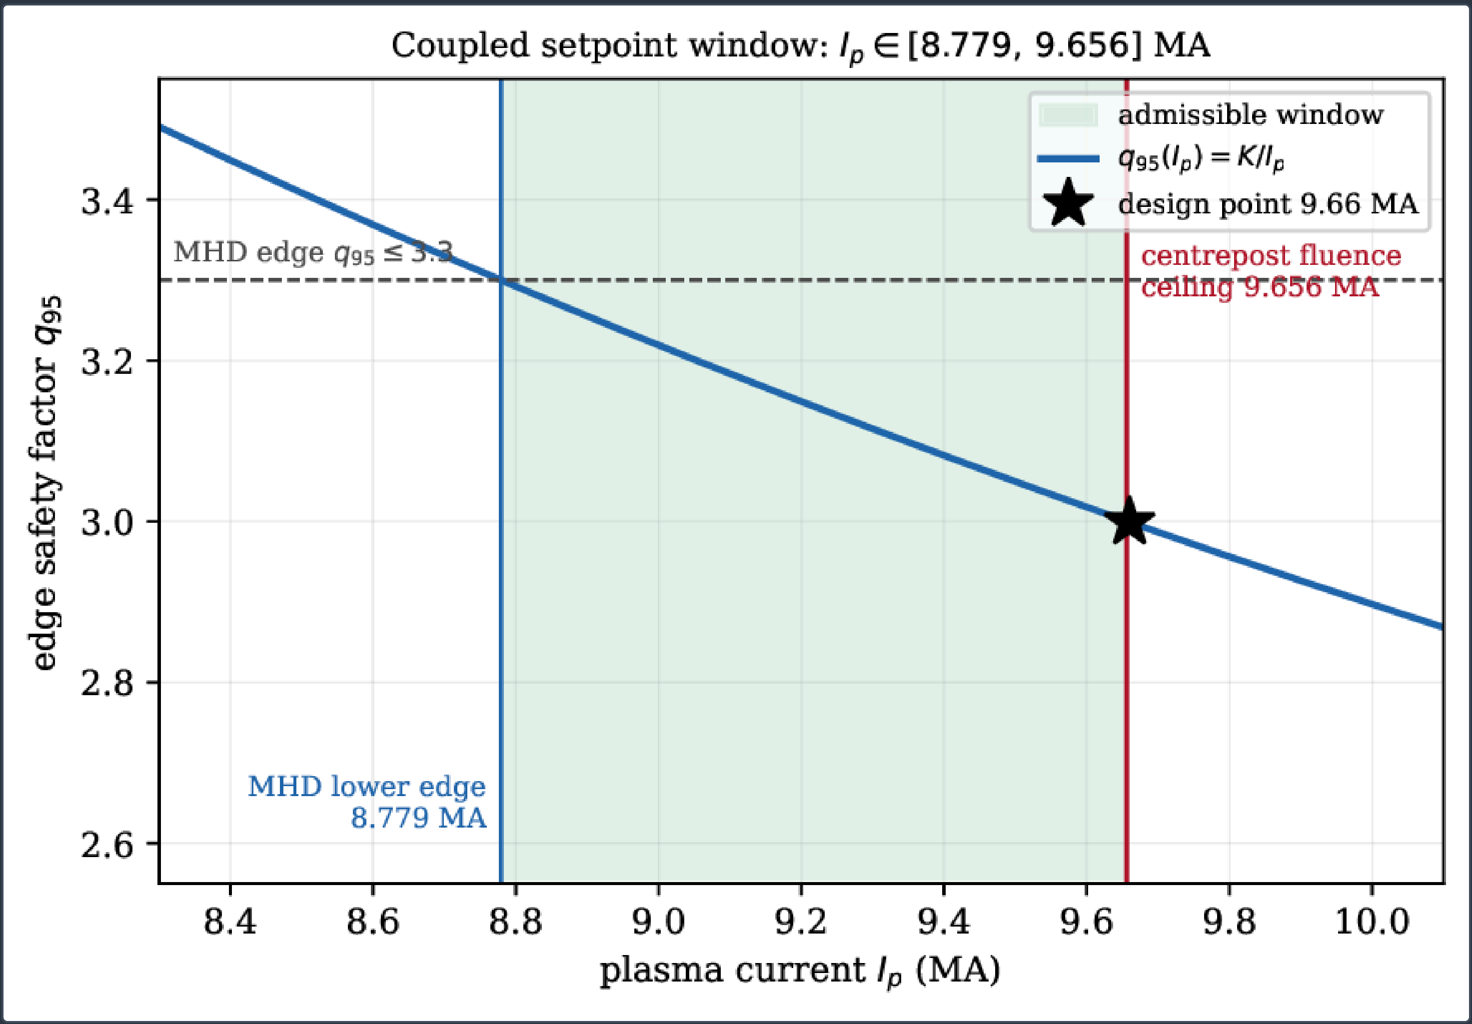

Coupled setpoint window (register serial KX-L1-A21). The edge safety factor $q_{95}(I_p) = K/I_p$ sets the MHD lower edge at 8.779 MA ($q_{95} = 3.3$); the centrepost fluence sets the upper edge at 9.656 MA. The admissible window (shaded) contains the design point 9.66 MA (star).

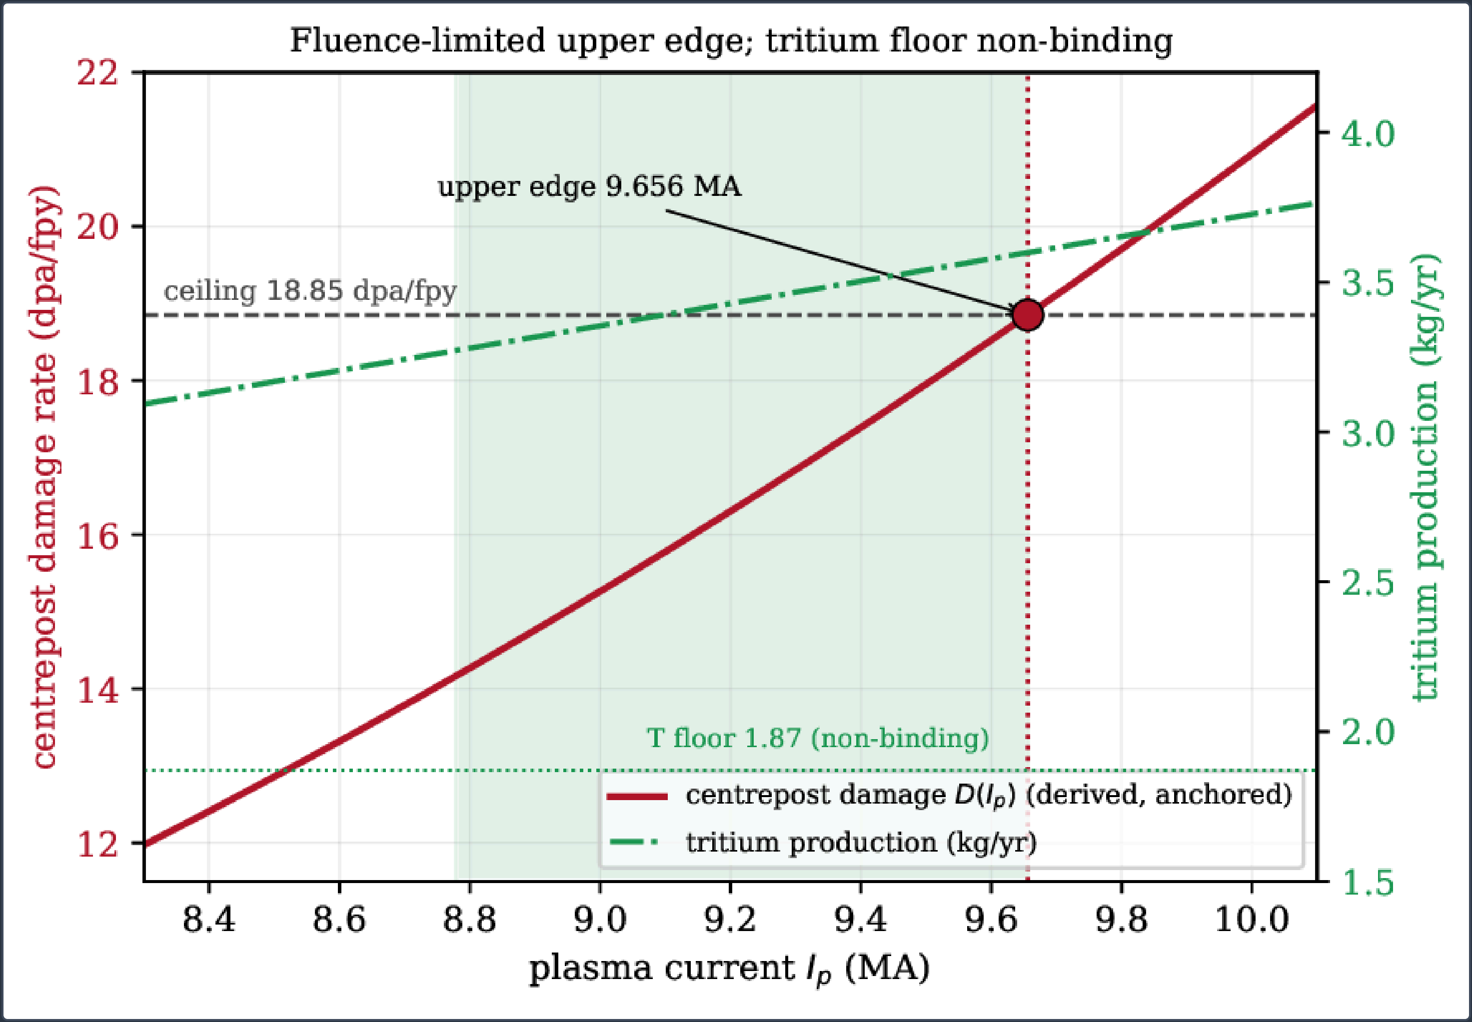

Fluence-limited upper edge. The centrepost damage rate $D(I_p)$ (derived, anchored to $D(9.656) = 18.85$ dpa fpy^{-1}) meets its ceiling at the window's upper edge, while the tritium production stays above its non-binding floor of 1.87 kg yr^{-1} across the admissible band (shaded).

THE ADJOINT-SENSITIVITY ENGINE

The adjoint gradients match finite differences component-by-component (Figure 7a): the maximum component discrepancy is below the 10^{-3} acceptance gate. The reason the adjoint is preferred is made explicit by the finite-difference step-size sweep (Figure 7b): the FD error is the sum of a truncation term that grows with the step and a round-off term that grows as it shrinks, with an irreducible minimum near $\Delta^\star \sim 10^{-8}$, whereas the adjoint gradient 13–14 is exact to machine precision at a cost independent of the parameter dimension. This is the gradient engine the other three layers consume.

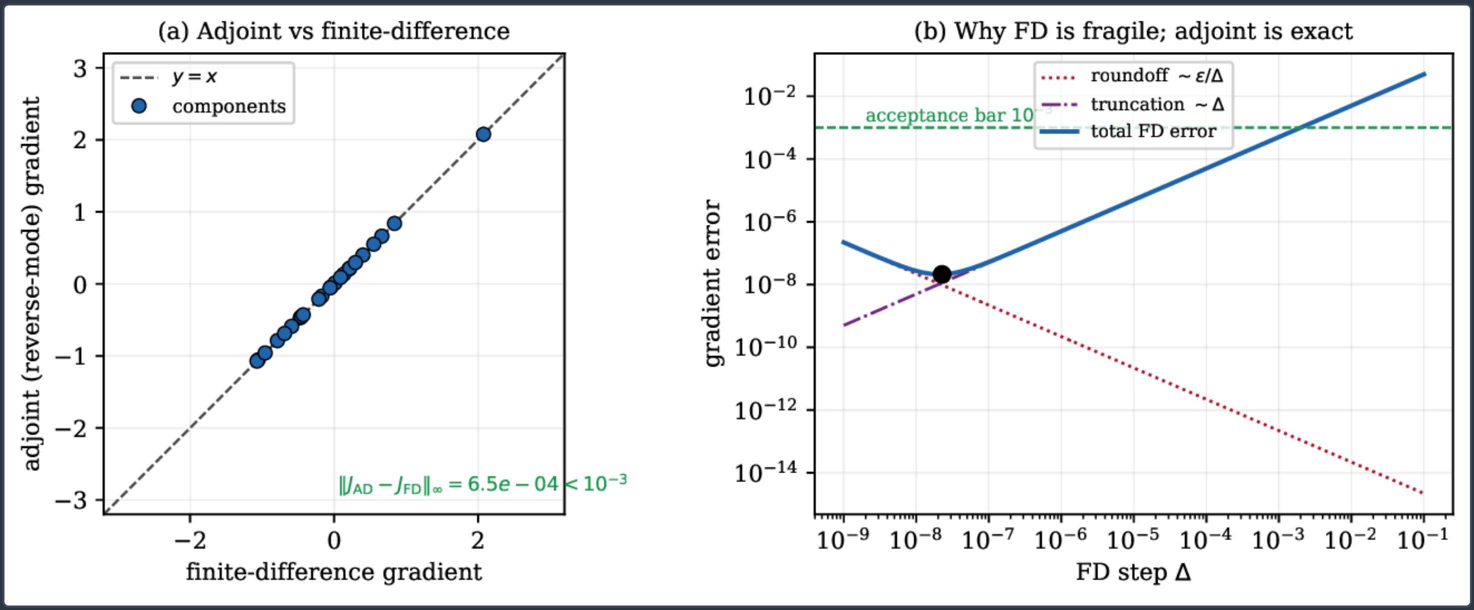

Adjoint-sensitivity verification (register serials SH-032/SH-184). (a) Adjoint (reverse-mode) gradient components against their finite-difference reference, with maximum discrepancy below the 10^{-3} gate. (b) The finite-difference error decomposes into truncation ($\sim \Delta$) and round-off ($\sim \varepsilon/\Delta$) terms with an irreducible minimum; the adjoint avoids this trade entirely. *Panel (b) is a representative verification on a canonical differentiable objective, not a plasma measurement.*

THE BACKBONE, PRODUCTIZED

The three demonstrated kernels, plus the adjoint engine, become the four productized backbone components (Table 2), each gated by a within-period acceptance test and each mapped to a frozen kernel and a productization serial. Table 1 collects the demonstrated results. The architecture (Figure 8) shows how the setpoint reference, reduced-basis surrogate and adjoint gradients feed the certified QP that projects the command onto the safe set every period; the within-period latency ladder (Figure 9) shows where each engine runs inside the budget, and the maturity-gated authority schedule (Figure 10) shows how passing each acceptance test unlocks the next tier of control authority.

KERNEL (REGISTER SERIAL)	DEMONSTRATED RESULT	GOVERNING BOUND
CBF quadratic program (KX-L1-A13)	0/100 barrier violations; $u \in [1.06, 1.25]$	forward invariance, Thm. 5
reduced-basis solve (KX-L1-A16)	rank-8 error 2.9×10^{-6}; $0.19\times$ storage	Eckart–Young, eq. 10
setpoint window (KX-L1-A21)	$I_p \in [8.779, 9.656]$ MA	KKT of eq. 15

The three demonstrated optimization kernels and their frozen results.

COMPONENT	BUILT TO KERNEL	SERIAL	ACCEPTANCE TEST
real-time QP solver	CBF clamp (KX-L1-A13)	KFE-CF-BR-112	QP solved within the control period
reduced-basis solver	low-rank solve (KX-L1-A16)	KFE-CF-SH-183	error within tol at ∼100× speed
setpoint-optimization service	I_p window (KX-L1-A21)	KFE-CF-SH-186	objective improved vs baseline
adjoint-sensitivity engine	differentiable surrogate	KFE-CF-SH-184	gradients match finite differences

The real-time optimization backbone: four components, the demonstrated kernel each is built to, the productization serial, and its within-period acceptance test.

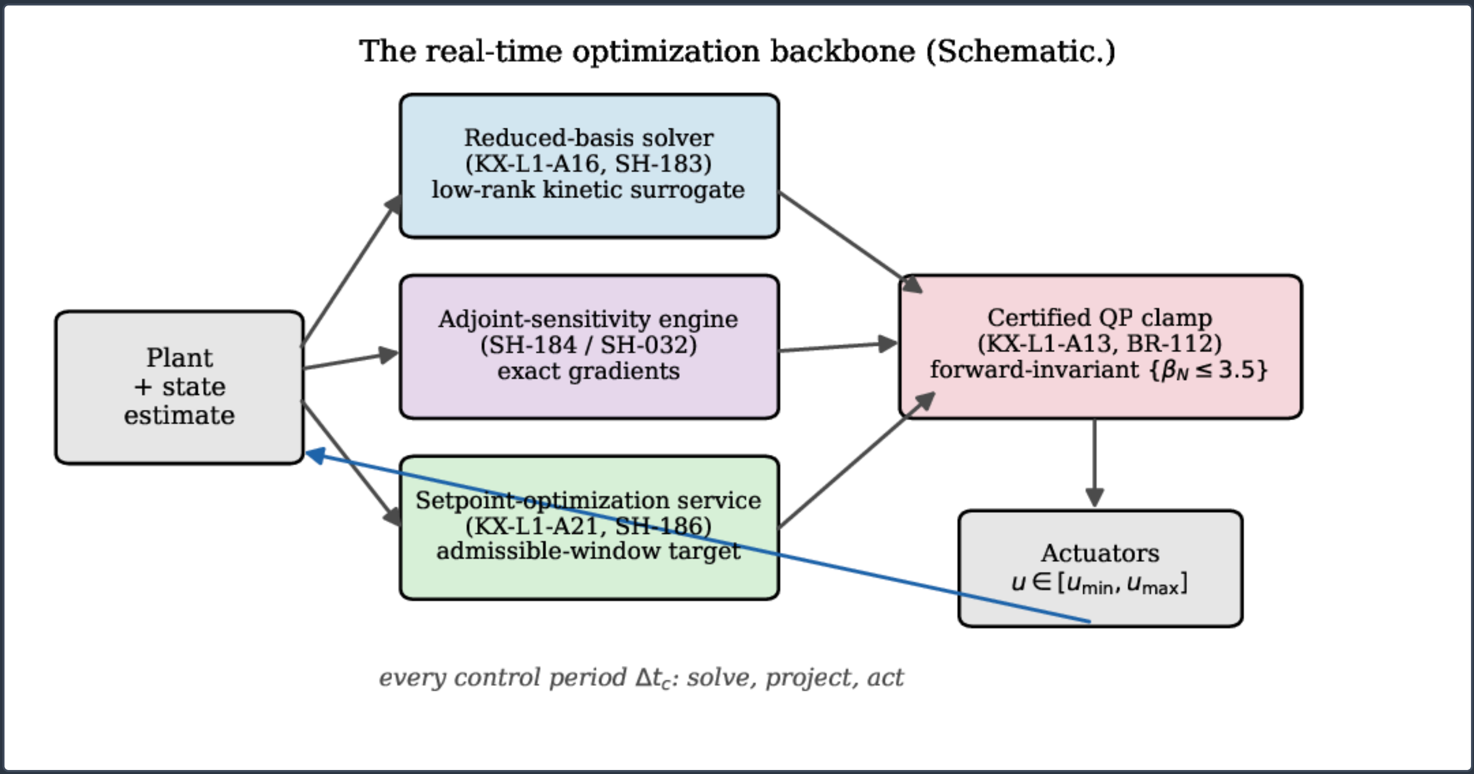

The real-time optimization backbone (Schematic.). The setpoint service supplies the reference, the reduced-basis surrogate the fast model, and the adjoint engine the gradients; the certified QP clamp projects the resulting command onto the forward-invariant safe set every control period before it reaches the actuators.

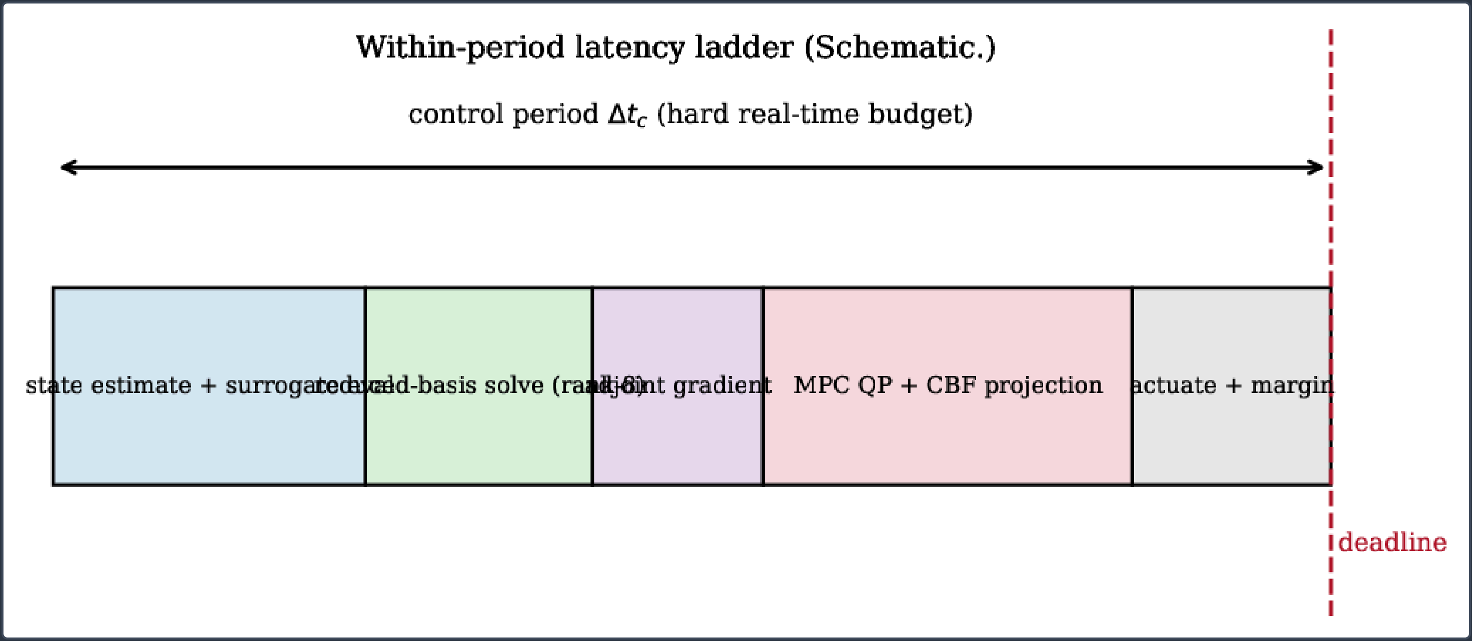

Within-period latency ladder (Schematic.). State estimation and surrogate evaluation, the reduced-basis solve, the adjoint gradient, and the MPC QP with CBF projection must all complete before the hard real-time deadline Δt_c.

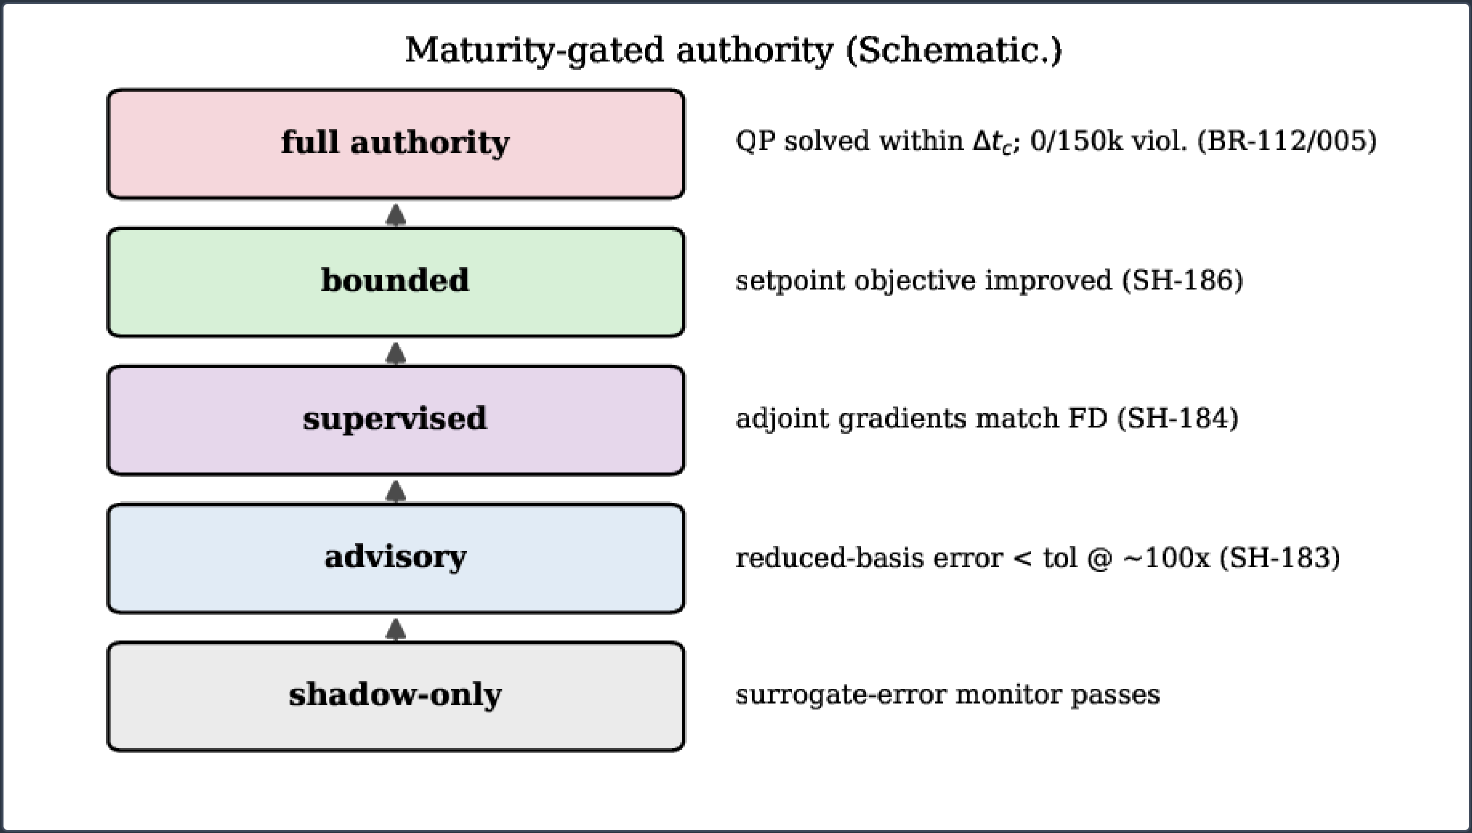

Maturity-gated authority (Schematic.). Each acceptance test that an engine passes unlocks a higher tier of control authority, from shadow-only through advisory, supervised and bounded operation to full authority.

§ 5

Discussion

AGAINST PUBLISHED CONTROL BENCHMARKS

The backbone is presented kernel-first because that is where the risk lives: an autonomous control architecture is only as good as its ability to solve the underlying optimization fast and safely. The certified QP realises the control-barrier-function safety paradigm [11,12,13] as a within-period projection, in the same family as the predictive safety filter of Wabersich and Zeilinger [14] but reduced here to a closed form for the safety-critical β_N scalar so that it runs on an FPGA/edge rig; the demonstrated $0/100$ forward invariance with margin (u never reaching its floor) is the discrete confirmation of Theorem 5. This complements, rather than replaces, learned control: a reinforcement-learning shape controller such as Degrave *et al.* [5] or a disruption predictor [6] can propose the desired command u_{des}, which the certified QP then projects onto the safe set — the pattern the companion MPC paper develops as a reference-governor controller bounded by this clamp [doi:10.5281/zenodo.22645807]. Against the ITER-class control problem [4,3], the contribution is the explicit within-period budget and the maturity-gated authority schedule.

AGAINST PUBLISHED REDUCTION AND OPTIMIZATION BENCHMARKS

The reduced-basis engine attains the theoretically optimal error for a given rank [21,20], and its exponential singular-value decay is the mechanism POD exploits across computational physics [22]; the randomised SVD used to compute it on GPU is the Halko–Martinsson–Tropp algorithm [23], and the low-rank compression it productizes matches the matrix-product-state result of the companion tensor-network solver [doi:10.5281/zenodo.22645768]. The adjoint engine is the adjoint-state method [24] realised by reverse-mode differentiation of the differentiable surrogates [doi:10.5281/zenodo.22645803]; its parameter-dimension-independent cost is what makes control-latency gradients feasible. The setpoint program's fluence-limited upper edge reproduces the frozen design point 9.66 MA from physical limits alone — an independent check against the lifetime-constrained operating-window synthesis [doi:10.5281/zenodo.22645799] — and its multivariable generalisation is a standard convex program [16,17,15] solvable by operator-splitting [18] or active-set [19] QP solvers at the required rate.

TWO DISTINCT MAGNETS, ONE BACKBONE

The backbone is not breeder-specific. The same four engines close the loop on the D–^{3}He tandem-mirror burner, whose certified-clamp kernel reproduces its own design-point anchors ($Q_E = 1.318$, neutron fraction $f_n = 5.44\%$, plug field 26.49 T; register serial KX-L1-A13-BU) on the burner's distinct high-field plug magnet. The breeder centrepost (16.84 T) and the burner plug (26.49 T) are two different magnets serving two different confinement concepts; the optimization substrate — certified QP, reduced-basis surrogate, adjoint gradients, coupled setpoint — is shared, but each product carries its own frozen anchors and its own acceptance tests.

§ 6

Limitations

The demonstrated kernels are at the reduced-model level: the CBF clamp is a one-dimensional, design-point-anchored surrogate of the β_N dynamics 7, and the reduced-basis result is on a canonical BGK test problem rather than the full six-dimensional gyrokinetic field. The full-plant, multivariable solvers are built to these demonstrated patterns and gated by the within-period acceptance tests of Table 2; the closed-form scalar clamp generalises to the multivariable QP 2 only under the standard assumption that the safety-critical constraints remain of relative degree one, which higher-order barrier constructions relax at additional cost. Bridging from the demonstrated kernels to the full plant is the acceptance-gated engineering sequence, not an in-scope claim of this paper.

§ 7

Conclusion

We have formalised a real-time optimization backbone for compact spherical-tokamak control and demonstrated three of its four engines as explicit mathematical programs with frozen, reproducible kernels. A control-barrier QP renders $\{\beta_N \leq 3.5\}$ forward-invariant — proven by Theorem 5 and confirmed at $0/100$ violations with a bounded, non-saturating command; a POD–Galerkin reduced-basis solver reaches 2.9×10^{-6} relative error at rank 8 in $0.19\times$ the storage at the Schmidt–Eckart–Young–Mirsky optimal rate; and a coupled setpoint optimizer returns the admissible current window $I_p \in [8.78, 9.66]$ MA from fluence and MHD limits, its upper edge recovering the design point. An adjoint-sensitivity engine supplies the exact, parameter-dimension-independent gradients the other three consume. Together, and each gated by an explicit within-period acceptance test, these engines constitute the certified QP solver, reduced-basis solver, adjoint engine and setpoint service of a real-time control substrate for an autonomous compact-ST breeder — and, with its own frozen anchors, for the tandem-mirror burner.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. Solver kernels ran on the NVIDIA GPU HPC campaign and the in-house real-time rig. The authors thank the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim.

Reproducibility. The four engines are frozen entries in the Kronos Physics De-Risking Register [34] (doi:10.5281/zenodo.22133057): the certified QP clamp (KX-L1-A13), the reduced-basis kinetic solve (KX-L1-A16), the setpoint current window (KX-L1-A21), and the adjoint-sensitivity engine (KFE-CF-SH-184/SH-032). Each is regenerable from the archived solver scripts and the figure-generation script deposited with this paper. This paper is archived at its DOI doi:10.5281/zenodo.22645899 and its official page https://www.kronosfusionenergy.com/publications/a-real-time-optimization-backbone-for-fusion-plasma-con.html; the register is archived at doi:10.5281/zenodo.22133057.

Data availability tocsectionData availability The clamp trace, the rank-convergence table, the storage-fraction curve, the setpoint window scan and the adjoint gradient-check used for Figures 1–10 are archived with the deposit at doi:10.5281/zenodo.22645899, which references the Kronos 2026 series including the AI-and-quantum-control paper [doi:10.5281/zenodo.22645702]. All quantitative values trace to the register [34] by the serial identifiers cited inline. Any economic analysis is reported separately and is not part of the public deposit.

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

tocsectionReferences

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 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

2 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

3 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

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

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

6 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

7 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

8 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

9 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

10 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

11 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

12 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

13 F Blanchini, Set invariance in control, *Automatica* **35** 1747 (1999). doi:10.1016/S0005-1098(99)00113-2

14 K P Wabersich and M N Zeilinger, A predictive safety filter for learning-based control of constrained nonlinear dynamical systems, *Automatica* **129** 109597 (2021). doi:10.1016/j.automatica.2021.109597

15 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). doi:10.1016/S0005-1098(99)00214-9

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

17 J Nocedal and S J Wright, *Numerical Optimization*, 2nd ed. (Springer, New York, 2006).

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 H J Ferreau, C Kirches, A Potschka, H G Bock and M Diehl, qpOASES: a parametric active-set algorithm for quadratic programming, *Math. Program. Comput.* **6** 327 (2014). doi:10.1007/s12532-014-0071-1

20 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

21 C Eckart and G Young, The approximation of one matrix by another of lower rank, *Psychometrika* **1** 211 (1936). doi:10.1007/BF02288367

22 G Berkooz, P Holmes and J L Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, *Annu. Rev. Fluid Mech.* **25** 539 (1993). doi:10.1146/annurev.fl.25.010193.002543

23 N Halko, P G Martinsson and J A Tropp, Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Rev.* **53** 217 (2011). doi:10.1137/090771806

24 R-E Plessix, A review of the adjoint-state method for computing the gradient of a functional with geophysical applications, *Geophys. J. Int.* **167** 495 (2006). doi:10.1111/j.1365-246X.2006.02978.x

25 P L Bhatnagar, E P Gross and M Krook, A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems, *Phys. Rev.* **94** 511 (1954). doi:10.1103/PhysRev.94.511

26 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

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 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

29 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

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

31 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

32 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

33 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

34 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 Real-Time Optimization Backbone for Fusion Plasma Control: Quadratic-Programming, Reduced-Basis and Adjoint-Sensitivity Engines at Control Latency

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

Zenodo · doi:10.5281/zenodo.22645899

Read on kronosfusionenergy.com

The Physics De-Risking Register