KRONOS FUSION ENERGY

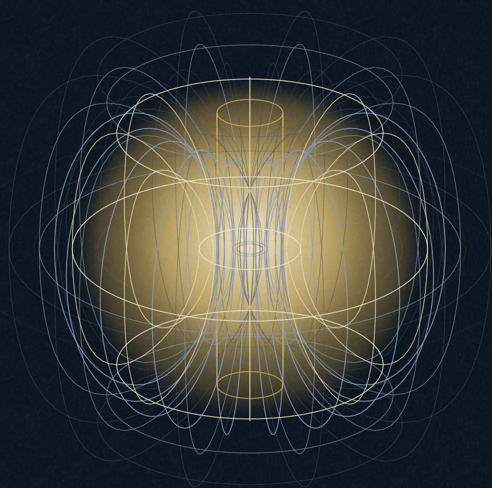

PAPER 2.85 · THE KRONOS 2026 SERIES

Prognostics and Remaining-Useful-Life Estimation for the Consumable Centrepost: A Dual-Channel Fluence-Driven Framework Scheduling Cartridge Replacement in a Spherical-Tokamak Breeder

The inboard centrepost of a compact spherical-tokamak (ST) breeder is not a lifetime component: behind a practical inboard shield the fast-neutron fluence degrades the...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Prognostics and Remaining-Useful-Life Estimation for the Consumable Centrepost: A Dual-Channel Fluence-Driven Framework Scheduling Cartridge Replacement in a Spherical-Tokamak Breeder

The inboard centrepost of a compact spherical-tokamak (ST) breeder is not a lifetime component: behind a practical inboard shield the fast-neutron fluence degrades the no-insulation REBCO conductor and its copper stabiliser to end of life within a fraction of a full-power year, so the machine is designed to swap the centrepost as a scheduled consumable cartridge. *When* to swap it is therefore a control decision, and a single worst-case interval wastes life or courts an unplanned outage. We formalise that decision as a prognostics problem. We derive a state-space damage model in which a physics channel integrates the fast-fluence load into a cumulative-damage variable by a linear (Miner) accumulation with a fracture-mechanics second mode, and cast the remaining-useful-life (RUL) estimate as a Bayesian posterior propagated by a particle filter that fuses the deterministic physics trajectory with a damage-correlated in-situ observable. The load-bearing safety feature is an out-of-distribution (OOD) backstop: when the live signals leave the calibrated envelope — flagged by a Mahalanobis regime score and made rigorous by a distribution-free conformal wrapper — the estimator abstains and returns the conservative physics-fluence bound rather than extrapolating a data model past its support. Grounded in the frozen de-risking gates for the Hyperion breeder ($I_p = 9.66\,MA$, $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, negative triangularity $\delta = -0.30$, $B_0 = 8\,T$, centrepost peak field $16.84\,T$), the framework converts a life estimate that spans more than three orders of magnitude across modelling bases (0.0003–0.765 fpy) into a live, confidence-bounded schedule that tracks the register-resolved design life of **0.58 fpy** behind a 45 cm tungsten-carbide shield, drives a cartridge-swap cadence of order five swaps per full-power year, and feeds the fleet-availability tiers. The one honest open item is the neutron tolerance of the in-winding observable, degraded by up to $39\times$ against the demonstrated diamond-sensor state of the art; the framework is designed so that the protective function survives it by falling back to the physics bound. This is a methods contribution — a prognostics architecture — distinct from the structural-integrity, quench-detection and swap-logistics papers of the series.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645788 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: prognostics and health management · remaining useful life · consumable centrepost · out-of-distribution backstop · conformal prediction · particle filter · cartridge replacement · 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)

prognostics and health management; remaining useful life; consumable centrepost; out-of-distribution backstop; conformal prediction; particle filter; cartridge replacement; spherical-tokamak breeder

§ 1

Introduction

THE CONSUMABLE CENTREPOST AS A DESIGN FACT

A spherical tokamak concentrates a large plasma current in a small volume behind a slender inboard column, and it is that column — the centrepost, carrying the toroidal-field winding at the machine's highest field — that makes the ST an efficient neutron source and breeder platform [1,2]. The same geometry places the centrepost conductor closest to the 14 MeV source with the least room for shielding. In the Hyperion breeder the centrepost holds a peak field of 16.84 T at a major radius of only $R_0 = 1.20\,m$; the winding is no-insulation (NI) REBCO coated conductor [3,4], and its critical current and copper-stabiliser residual-resistance ratio (RRR) are degraded by the accumulated fast-neutron fluence. Behind the practical inboard shield the resulting service life is a fraction of a full-power year (fpy), far shorter than the plant life, so the centrepost is engineered from the outset as a scheduled consumable: a demountable cartridge extracted vertically and replaced, with breeding carried by the outboard blanket [29,30]. This is a deliberate architectural choice that converts magnet lifetime into a recurring operational quantity that must be *managed*.

WHY A FIXED WORST-CASE INTERVAL IS INADEQUATE

The trouble with a consumable is that its life is not a single number. The centrepost life reported across the Kronos de-risking campaign depends on the inboard shield thickness, on the fluence basis, on whether a midplane worst-case disk or a full-height three-dimensional column is modelled, and on the operating current. The campaign's own record makes the point sharply: three contradictory raw figures — 0.22, 0.01 and 0.0003 fpy — collapse to a *resolved* 0.58 fpy only once the fast-fluence transport is run behind a 45 cm tungsten-carbide (WC) shield on a full-height column and the scheduled-cartridge premise is adopted [28,29]. Even the resolved value sits inside a broader family of defensible estimates spanning 0.199–0.765 fpy across shield thickness, current and fluence basis (Table 1) [28]. Scheduling the swap on any single one of these numbers is either wasteful (retiring a cartridge with life left) or unsafe (running one past its support into an unplanned trip during a nuclear campaign). A live estimate that updates on the machine's own measured exposure is worth more than any fixed interval; supplying it, honestly and fail-safe, is the subject of this paper.

PROGNOSTICS AND HEALTH MANAGEMENT

Prognostics and health management (PHM) is the discipline that turns a degrading component into a remaining-useful-life (RUL) estimate and a maintenance decision [5,6,7]. Its two classical model families are directly applicable here. The *physics-of-failure* family accumulates a damage variable from the known load — for fatigue by the linear Miner rule [8], for sub-critical crack growth by the Paris law [9], for irradiation by a fluence- or dose-driven degradation law [10,11] — and reads RUL as the exposure remaining until the variable reaches unity. The *data-driven* family learns a degradation trajectory from a condition-monitoring signal and estimates RUL by Bayesian filtering, with the particle filter the standard tool for the nonlinear, non-Gaussian degradation dynamics that arise in practice [12,13,14]. The benchmark that anchors the field, the NASA C-MAPSS turbofan run-to-failure set [15], is scored by prognostic-specific metrics — notably the α–λ accuracy cone and the prognostic horizon [16] — which we adopt below. What PHM has largely lacked for a *safety-critical, nuclear* asset is a principled statement of what the estimator does when the plant leaves the regime the estimator was built for; the modern answer, calibrated uncertainty with abstention out of distribution, comes from the machine-learning reliability literature [17,18,19,20,21].

CONTRIBUTION

This paper contributes a prognostics architecture for the consumable ST centrepost with four elements, each stated formally and grounded in a frozen de-risking gate. (i) A state-space *physics channel* (§2.1) that integrates the fast-fluence load into a cumulative-damage variable with a fracture-mechanics second mode, calibrated to the resolved 0.58 fpy design life and the conductor break-even fluence. (ii) A *Bayesian dual-channel estimator* (§2.3) that fuses the deterministic physics trajectory with a damage-correlated in-situ observable through a particle filter, returning a RUL posterior with a credible interval rather than a point. (iii) An *out-of-distribution backstop* (§2.6) that gates the data channel by a Mahalanobis regime score and a distribution-free conformal wrapper, and abstains to the conservative physics bound when the live signals leave calibrated support. (iv) A *scheduling and authority* layer (§4.7) that maps the RUL posterior to a cartridge-swap cadence and duty schedule, gated by verification maturity. Section 3 states the GPU-HPC campaign that produced the underlying fluence, damage and current-window gates and the reproducibility discipline; Section 4 reports the estimator on ten figures and four tables; Section 5 benchmarks it against the published PHM literature; Section 6 states the one honest open item. Every quantitative claim traces to the Kronos Physics De-Risking Register [28] by the serial identifiers cited inline (e.g. BR-L1-A13b), so each number is auditable.

§ 2

Governing equations and the formal prognostics problem

THE DAMAGE STATE AND FAST-FLUENCE ACCUMULATION

Let t denote exposure in full-power years (fpy) and let $\Phi(t)$ be the accumulated fast-neutron fluence ($E > 0.1\,MeV$) at the centrepost conductor. The binding degradation is not displacement damage of the structure — at the shielded centrepost the displacement rate is only $\dot{d} \approx 0.0093\text{–}0.0104\ \mathrm{dpa\,fpy}^{-1}$ and is far from any structural limit — but the fluence-driven loss of the REBCO critical current I_c and of the copper-stabiliser RRR (BR-L1-A13b, BR-SX-12) [28,32]. We therefore take the fluence itself as the primary damage driver. The instantaneous fluence rate is set by the neutron source strength, which for a fixed geometry and shield scales with the fusion power, and is attenuated exponentially through the inboard shield of thickness t_{sh},

$$\dot{\Phi}(t) = \dot{\Phi}_0\,\frac{P_{\mathrm{fus}}(t)}{P_{\mathrm{fus}}^{\mathrm{DP}}}\,e^{-\Sigma\,t_{\mathrm{sh}}}, \qquad \Phi(t) = \int_0^t \dot{\Phi}(t')\,\mathrm{d}t',$$

with $\dot{\Phi}_0$ the unshielded source-face rate, Σ the effective removal coefficient of the WC shield, and $P_{\mathrm{fus}}^{\mathrm{DP}} = 85.04\,MW$ the design-point power. The removal law is anchored to the transport result that a bare centrepost sees $\dot{\Phi} \approx 2.9e25\,/m^2 fpy$ (life 0.003 fpy, inadmissible) while a 45 cm WC shield brings this to $\approx 1.7e23\,/m^2 fpy$ (life 0.58 fpy), passing through $1.768e23\,/m^2 fpy$ at 30 cm (life 0.221 fpy, transport CI $0.195\text{–}0.253$) [28,29] (figure 1).

We define the health state as a normalised cumulative-damage variable $D(t) \in [0,1]$ that reaches unity at end of life. Because the dominant mechanism is a monotone fluence-to-I_c degradation with a single end-of-life criterion, a *linear* (Miner) accumulation [8] against the conductor break-even fluence Φ_c is the leading-order model,

$$D(t) = \frac{\Phi(t)}{\Phi_c} = \sum_k \frac{n_k(t)}{N_k}, \qquad \mathrm{RUL}(t) = \inf\{\,\tau \geq 0:\ D(t+\tau) \geq 1\,\},$$

where the right-hand sum writes the same quantity as accumulated exposure n_k against allowable exposure N_k in damage bin k, the standard PHM form for a spectrum of load levels. The break-even fluence Φ_c is not a free parameter: first-principles and irradiation gates fix it at $\Phi_c \approx 3.82e22\,/m^2$ for the artificial-pinning-centre (APC) tape and $5.64e22\,/m^2$ for the un-doped tape whose robustness the campaign recommends (BR-SX-04) [31,32] (figure 6). Under a constant-rate approximation, $\mathrm{RUL}(t) = t\,[1 - D(t)]/D(t)$, the classical linear-damage estimate.

THE SECOND DAMAGE MODE: SUB-CRITICAL CRACK GROWTH

A slender, thermally and electromagnetically cycled column also accumulates mechanical fatigue, and although the structural margins are large — the peak von Mises stress is 181 MPa at the plug-relevant 26.49 T FEA basis (a $6\times$ margin) and the winding-pack operating strain is -0.15% on the reversible I_c plateau with $3.3\times$ margin to the irreversible limit (BR-SX-13, QS-10) [28] — a complete prognostic tracks a crack channel in parallel. Sub-critical growth of an edge flaw of depth a under the per-cycle stress-intensity range ΔK follows the Paris law [9],

$$\frac{\mathrm{d}a}{\mathrm{d}N} = C\,(\Delta K)^m, \qquad \Delta K = Y\,\Delta\sigma\,\sqrt{\pi a},$$

with Y the geometry factor, $\Delta\sigma$ the cyclic stress range, and N the swap/thermal-cycle count. Integrating (3) from an initial flaw a_0 to the critical size $a_c = (K_{\mathrm{Ic}}/Y\sigma_{\max})^2/\pi$ gives a cycle-count RUL that runs in parallel with the fluence RUL of (2); the served RUL is the minimum of the two channels. Because the fluence channel binds first by a wide margin at the design point, the crack channel functions as a guard that would take over only under an off-nominal cyclic-load excursion.

STATE-SPACE FORM

Collect the health variables into a state $x = (D, a, \theta)^\top$, where θ gathers slowly varying model parameters (effective removal Σ, break-even Φ_c, Paris C). The degradation evolves as a stochastic dynamical system driven by the measured operating point $u(t) = (P_{\text{fus}}, I_p, t_{\text{sh}}, \ldots)$,

$$x_{k+1} = f(x_k, u_k) + w_k, \qquad w_k \sim \mathcal{N}(0, Q_k),$$
$$z_k = h(x_k) + v_k, \qquad v_k \sim p_v,$$

with process noise w_k absorbing fluence-rate and material-property uncertainty and measurement noise v_k (not necessarily Gaussian) on the in-situ observable z_k. The prognostic target is the posterior predictive distribution of the first-passage time,

$$p(\text{RUL}(t_k) \mid z_{1:k}) = \int \mathbb{1}[\inf\{\tau : D(t_k + \tau) \geq 1\}]\, p(x_k \mid z_{1:k})\, \mathrm{d}x_k,$$

i.e. the distribution of the exposure at which the damage trajectory forecast from the filtered state $p(x_k \mid z_{1:k})$ first crosses unity. Equation (5) is the formal object the estimator approximates; its *mean* is the scheduling input and its *spread* is what the safety layer reads.

BAYESIAN FILTERING AND THE PARTICLE APPROXIMATION

The filtered state obeys the recursive Bayes update [13]

$$p(x_k \mid z_{1:k}) = \frac{p(z_k \mid x_k) \int p(x_k \mid x_{k-1})\, p(x_{k-1} \mid z_{1:k-1})\, \mathrm{d}x_{k-1}}{p(z_k \mid z_{1:k-1})}.$$

The degradation dynamics (4) are nonlinear (exponential attenuation, power-law crack growth) and the observation noise is generally non-Gaussian, so we approximate (6) with a sequential-importance-resampling particle filter. Carrying N_p weighted samples $\{x_k^{(i)}, w_k^{(i)}\}$, the weight update and normalisation are

$$\tilde{w}_k^{(i)} = w_{k-1}^{(i)} \frac{p(z_k \mid x_k^{(i)})\, p(x_k^{(i)} \mid x_{k-1}^{(i)})}{q(x_k^{(i)} \mid x_{k-1}^{(i)}, z_k)}, \qquad w_k^{(i)} = \frac{\tilde{w}_k^{(i)}}{\sum_j \tilde{w}_k^{(j)}},$$

with proposal q; systematic resampling is triggered when the effective sample size $N_{\text{eff}} = 1/\sum_i (w_k^{(i)})^2$ falls below $N_p/2$. The RUL posterior (5) is then the empirical distribution of per-particle first-passage times, from which the credible interval is read directly. In the linear-Gaussian degenerate limit the recursion (6) reduces to the Kalman filter [22], and the framework retains a dropout-robust Kalman/particle estimator for the sensor-loss case (SH-043) [28].

DUAL-CHANNEL FUSION

The two channels supplying (5) are complementary. The *physics channel* evaluates $D(t) = \Phi(t)/\Phi_c$ from (1)–(2) using the measured operating history; it is unbiased but broad, because Φ_c and Σ carry material and transport uncertainty. The *data channel* maps a damage-correlated in-situ observable (a cartridge-coil condition feature — current-sharing-margin drift, joint resistance, or a strain/magnetization signature, KFE-CF-BR-107/KX-L1-A15) [28] to a sharper but potentially drifting estimate. When both are trustworthy their fusion is the variance-optimal Gaussian product: for physics estimate (μ_p, σ_p^2) and observable estimate (μ_o, σ_o^2),

$$\mu_F = \frac{\sigma_o^2 \mu_p + \sigma_p^2 \mu_o}{\sigma_p^2 + \sigma_o^2}, \qquad \sigma_F^2 = \frac{\sigma_p^2 \sigma_o^2}{\sigma_p^2 + \sigma_o^2} \leq \min(\sigma_p^2, \sigma_o^2),$$

so the fused posterior is never wider than either input (figure 4); the particle filter realises the same fusion without the Gaussian assumption. Equation (8) is the estimator's normal-mode behaviour. Its safety, however, is set entirely by what happens when the observable is *not* trustworthy.

THE OUT-OF-DISTRIBUTION BACKSTOP

A data-driven estimate is reliable only inside the regime for which it was calibrated; extrapolated past its support it can be confidently wrong, and a confidently wrong RUL on a nuclear magnet is a hazard, not an inaccuracy. The architecture therefore carries an explicit backstop that abstains rather than extrapolates [19,20]. Let ϕ_k be the feature vector the data channel consumes, with in-distribution mean μ and covariance Σ_ϕ estimated on the calibration set. The regime score is the Mahalanobis distance

$$d_M(\phi_k) = \sqrt{(\phi_k - \mu)^\top \Sigma_\phi^{-1}(\phi_k - \mu)},$$

and the data channel is admitted only while $d_M \le \tau$ (SH-040, the out-of-distribution regime detector) [28]. The threshold τ is itself calibrated, and the point estimate is wrapped in a *distribution-free conformal* interval that guarantees finite-sample marginal coverage without any distributional assumption [21,20] (SH-073):

$$\mathbb{P}\Big(y \in \hat{\mathcal{C}}_\alpha(\phi_k)\Big) \ge 1 - \alpha, \qquad \hat{\mathcal{C}}_\alpha(\phi_k) = \big[\hat{y} \pm q_{1-\alpha}(\{r_j\})\big],$$

where $q_{1-\alpha}$ is the empirical $(1-\alpha)$-quantile of held-out nonconformity residuals $\{r_j\}$. The served estimate is then

$$\widehat{\mathrm{RUL}}_{\mathrm{served}}(t_k) = \begin{cases} \text{fused posterior\textasciitilde(???)/(???)}, & d_M(\phi_k) \le \tau, \\ \mathrm{RUL}_{\mathrm{phys}}(t_k)\ \text{(conservative fluence bound)}, & d_M(\phi_k) > \tau, \end{cases}$$

so that on any out-of-distribution input the estimator hands authority back to the deterministic physics bound of (2) (figure 5). This is the same abstention discipline that governs the wider Kronos control stack — calibrated uncertainty, model-drift freeze, and graceful degradation toward a certified conservative controller [36,37] — specialised here to the prognostic. It is the difference between a decision aid and a liability.

§ 3

Computational methodology: the NVIDIA GPU HPC campaign

Every quantity the estimator is calibrated to — the fluence–shield–life map, the displacement rates, the conductor break-even fluence, the current-window damage rates, and the structural and quench margins that place the non-binding channels — is a frozen output of the Kronos de-risking campaign, produced on the NVIDIA GPU HPC fleet and independently re-run before being stamped into the register [28,38]. The prognostic itself is a light real-time layer that runs on the in-house edge/HPC node; the heavy transport and materials physics it rests on is what consumes the GPU campaign. We name the codes actually run for the gates this paper uses.

Neutron transport (fluence, dpa, breeding). Fast-fluence and displacement fields at the centrepost were computed with the Monte-Carlo transport code OpenMC [24] on ENDF/B-VIII.0 cross-sections, in both a CSG radial model and a full-height, twenty-axial-bin three-dimensional column, with the STEP-to-DAGMC pathway used for CAD-faithful runs (BR-L1-A13b, BR-SX-12, P3-NEUT). The closing configurations were driven at 5×10^5 particles per batch over twenty batches (seed logged) to converge the deep-penetration inboard tally where the statistical error dominates the life estimate; the batch/seed provenance is carried in the register so the life numbers reproduce arithmetically. Displacement-damage spectra used an NRT-dpa reduction [23] with a SPECTRA-PKA-class primary-knock-on analysis and activation/transmutation by a FISPACT-II-style inventory pass [25].

Conductor degradation. The break-even fluence that fixes Φ_c in (2) was set by a first-principles chain: density-functional-theory-plus-Hubbard-U (DFT+U) chain-oxygen-vacancy formation energies and YBCO displacement-cascade molecular dynamics for the un-doped-versus-APC ordering, GPU-accelerated where the MD and DFT throughput allowed (BR-SX-04) [31,32]. The ordering (un-doped $>$ APC) was robust across 4000 Monte-Carlo samples.

Structure and quench (non-binding channels). The structural margins were computed with finite-element elasticity (sfepy and CalculiX) on the real centrepost CAD with graded per-turn pre-compression; the adiabatic quench/MIITs budget and the fused strain-rate-plus-magnetization precursor were computed in the magnet-twin analysis (BR-SX-13, BR-SX-14, QS-09/10) [28]. These place the mechanical and thermal channels well inside their limits, which is why the fluence channel binds.

Real-time estimator. The particle filter (7), the Mahalanobis/conformal gate (9)–(10) and the scheduler run as a deterministic microsecond-class edge service (bare-metal/FPGA, no garbage collection), consuming validated telemetry and emitting a clamped recommendation [28]. It never actuates outside the certified safe set; it advises the swap cadence and the duty scheduler. Training and calibration of the data channel run offline on the GPU fleet; inference is a milliseconds-class operation dwarfed by the campaign timescale.

Verification and reproducibility. The campaign follows a verification-first discipline: frozen anchors, independent-library re-runs, and a byte-for-byte provenance catalogue [38,39]. The centrepost-life numbers reproduce arithmetically from the deposited OpenMC fluence data; the estimator figures in this paper are regenerated by a single deposited script from those frozen anchors, with no dependence on unpublished or proprietary data. No cost or economic quantity enters the methodology at any point.

§ 4

Results

THE FLUENCE–SHIELD–LIFE MAP THAT CALIBRATES THE PHYSICS CHANNEL

Figure 1 shows the transport-anchored basis of (1): the fast-fluence rate falls by more than two orders of magnitude across the practical shield range, and the centrepost life rises with it from the inadmissible bare value (0.003 fpy) to the resolved design value of 0.58 fpy behind 45 cm of WC. The knee is steep — this is a deep-penetration exponential — so the life is exquisitely sensitive to the as-built shield, which is one reason a live, condition-updated estimate is preferable to a design-book interval.

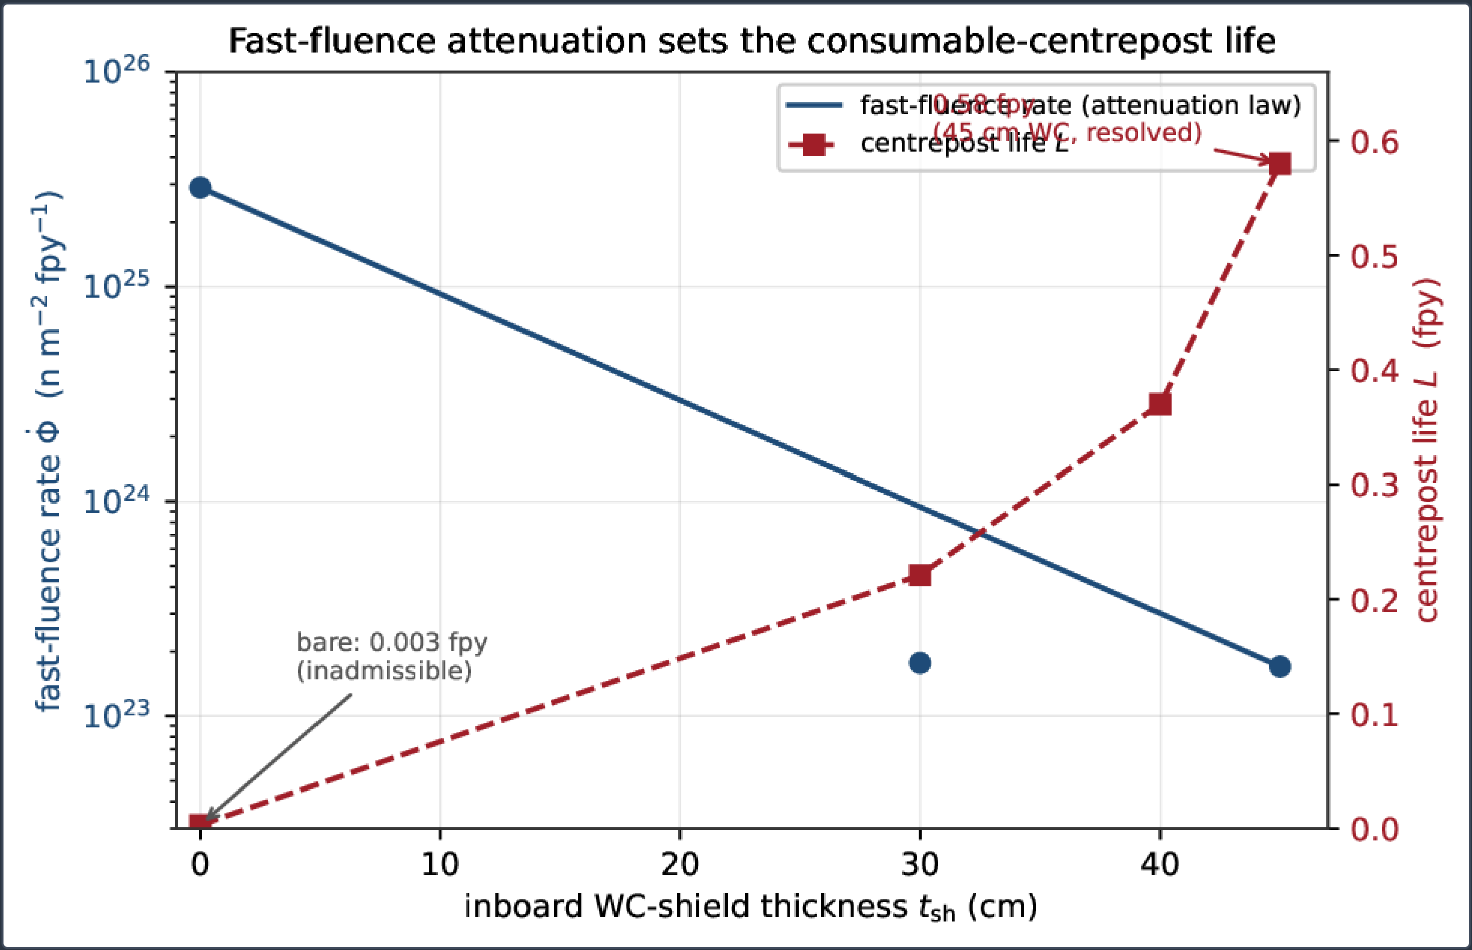

Fast-fluence rate (left axis, log) and consumable-centrepost life (right axis) versus inboard WC-shield thickness. Circles and squares are the frozen transport anchors (BR-L1-A13b, register): bare $2.9e25\ /m^2 fpy$ / 0.003 fpy; 30 cm $1.768e23$ / 0.221 fpy; 45 cm $1.7e23$ / 0.58 fpy. The fluence curve is the exponential removal law of equation (1) fitted to the anchored endpoints.

THE RESOLVED-LIFE SPREAD: THE CASE FOR A LIVE ESTIMATE

Figure 2 collects the family of defensible life estimates on one axis. They span more than three orders of magnitude, from the bare-midplane 0.0003 fpy to the idealised 0.765 fpy, with the register-resolved 0.58 fpy near the top of the credible band and the design-basis synthesis values (0.199–0.21 fpy, rising to 0.46 fpy at reduced current) near the middle. Table 1 lists the bases. No single number is simultaneously safe and efficient across this spread; the estimator's job is to collapse it, per campaign, to a live posterior with a stated interval.

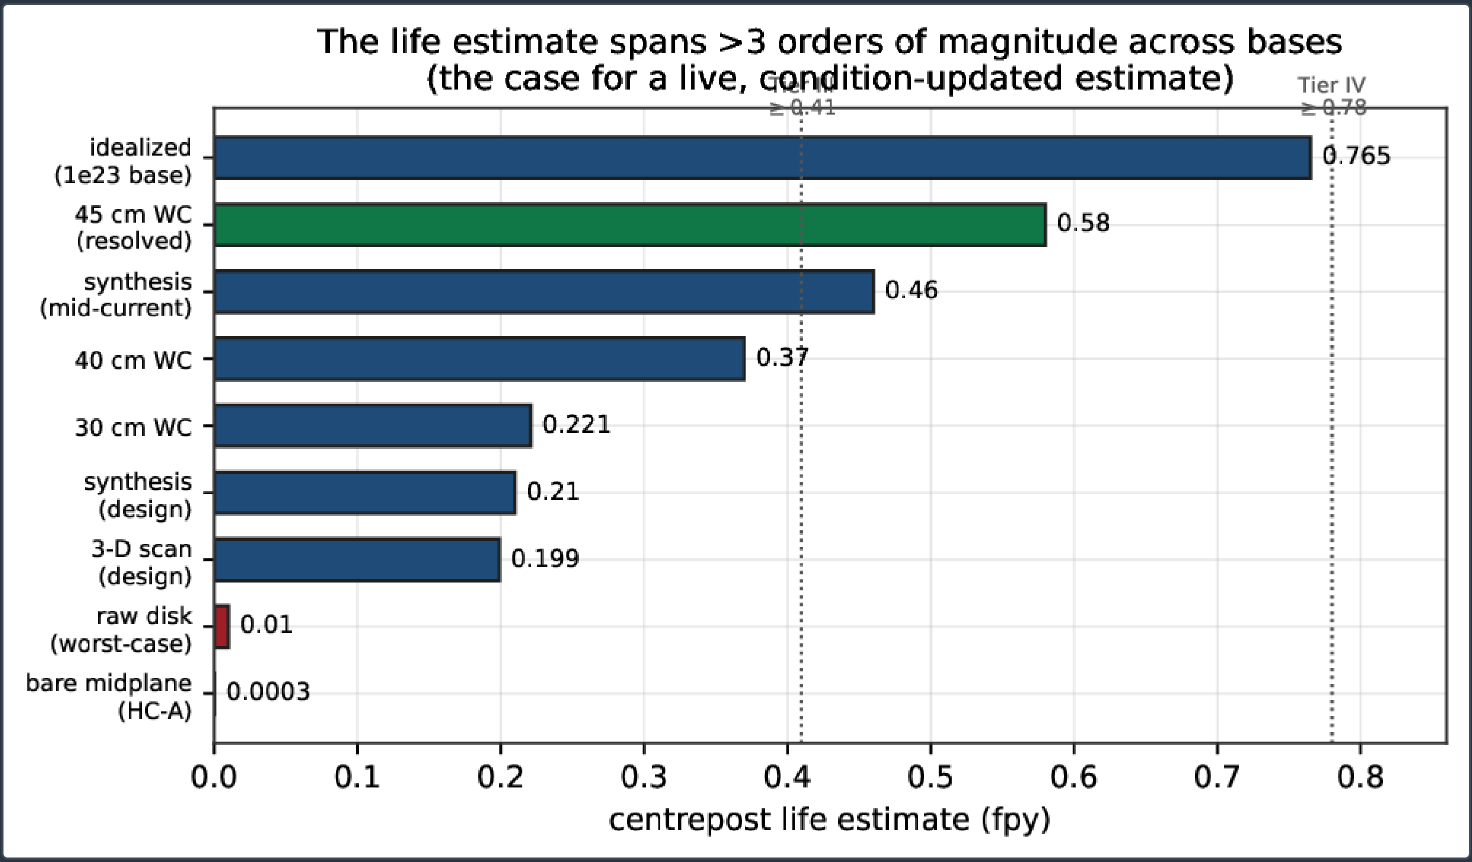

The centrepost-life estimate across modelling bases (Table 1; register). Red bars mark the inadmissible raw figures, green the register-resolved 0.58 fpy, blue the intermediate defensible values. Dotted lines are the N+1 availability tiers (Tier III at life ≥ 0.41 fpy, Tier IV at ≥ 0.78 fpy). The three-order-of-magnitude spread is the argument for a condition-updated estimate over a fixed interval.

BASIS	LIFE (FPY)	SERIAL / NOTE
bare midplane (HC-A)	0.003	inadmissible (BR-L1-A13b)
raw worst-case disk	0.01	pre-resolution (register)
3-D design scan	0.199	BR-SX-12 (dpa 0.0104/fpy)
lifetime synthesis (design)	0.21	BR-SX-16
30 cm WC shield	0.221	CI 0.195–0.253 (register)
40 cm WC shield	0.370	TBR 1.046 (register)
lifetime synthesis (mid-current)	0.46	BR-SX-16 recovery lever
45 cm WC (resolved)	**0.58**	**register MEASURED**
idealised (10^{23} basis)	0.745–0.765	CS-D / P3-NEUT

DAMAGE ACCUMULATION AND THE RUL POSTERIOR

Figure 3 shows the physics-channel trajectory: cumulative damage $D(t)$ rising linearly to unity at the nominal 0.58 fpy end of life, and the RUL posterior $\mathrm{RUL}(t) = 0.58 - t$ with a credible band that widens with forecast horizon and tightens as the cartridge approaches replacement. This is the normal-mode output of equations (2) and (5) before the observable is fused in.

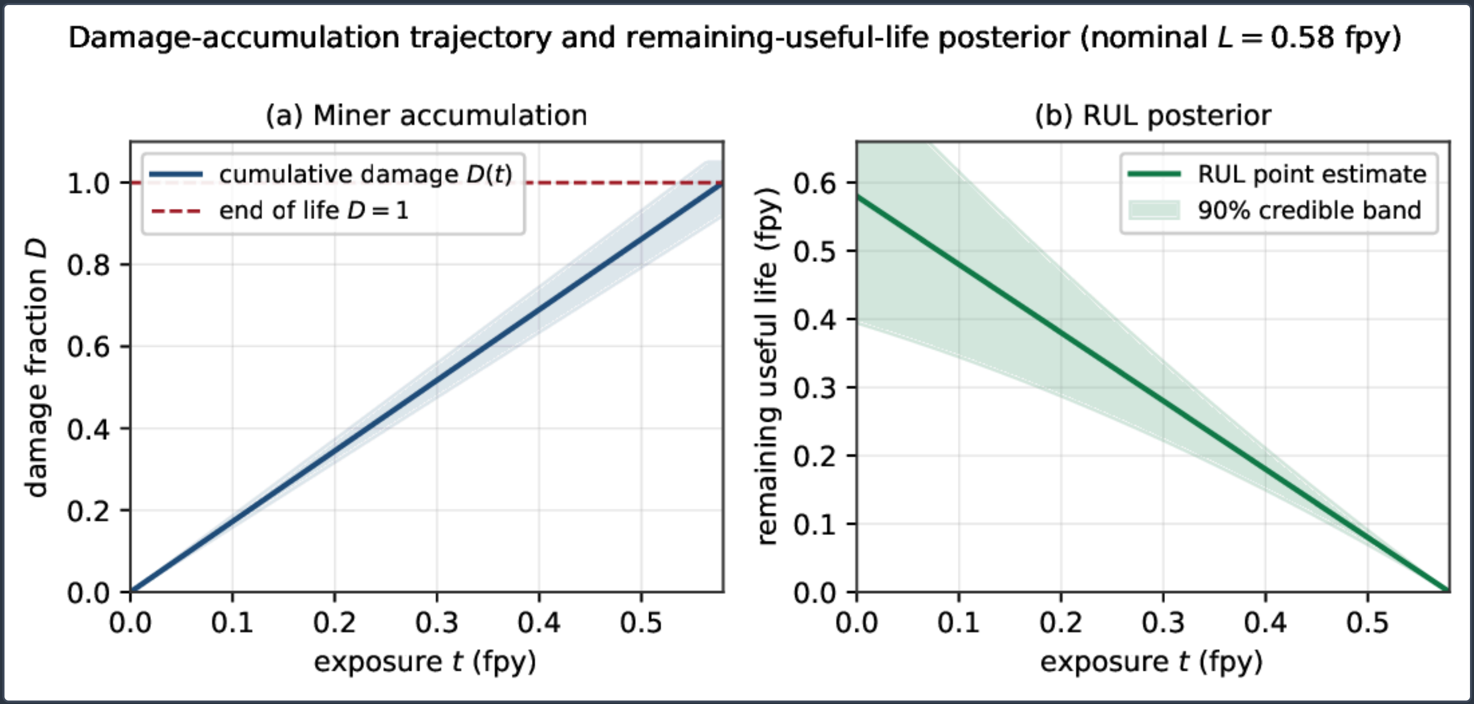

(a) Miner cumulative-damage trajectory $D(t)$ to the end-of-life criterion $D = 1$ at the nominal 0.58 fpy (equation 2); shaded band is the fluence-rate/break-even uncertainty. (b) The corresponding RUL posterior (equation 5) with a 90% credible band that narrows toward end of life. Synthetic trajectory anchored to the frozen design life.

DUAL-CHANNEL FUSION NARROWS THE INTERVAL

Figure 4 shows the fusion of equation (8) at a mid-life epoch. The physics channel is unbiased but broad; the in-situ observable is sharper; their variance-optimal combination is narrower than either and centred near the true RUL. This is the value the observable adds *when it is trustworthy* — and the next subsection shows what protects the estimate when it is not.

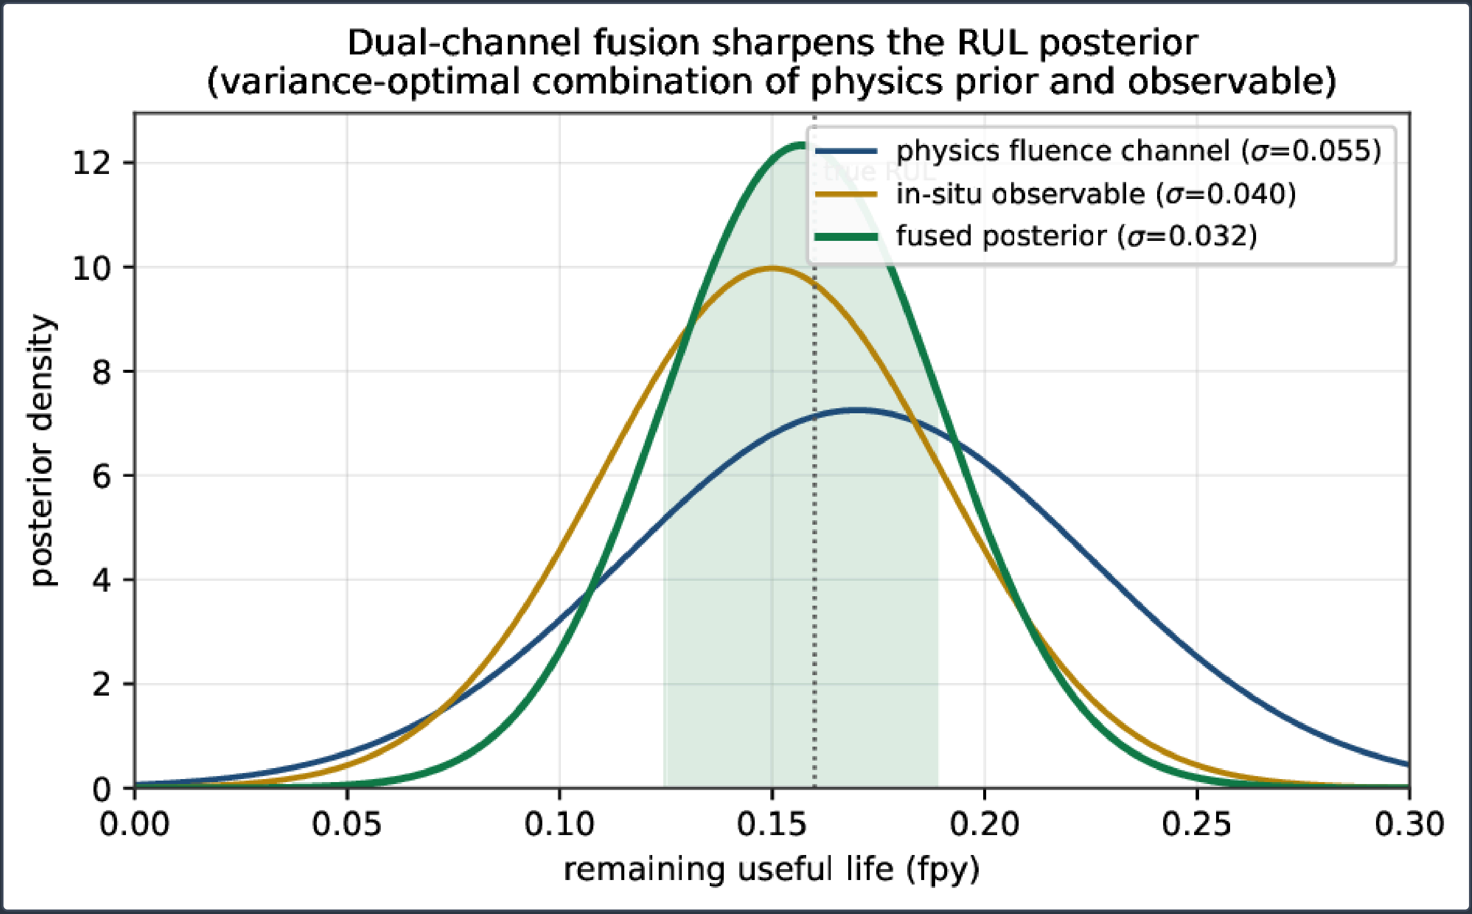

Dual-channel RUL fusion at a mid-life epoch (equation 8). The fused posterior (green, σ_F) is narrower than either the physics fluence channel (blue) or the in-situ observable (gold), and concentrates near the true RUL. The particle filter realises the same fusion without the Gaussian assumption.

THE OUT-OF-DISTRIBUTION BACKSTOP IN ACTION

Figure 5 exercises the backstop of equation (11). A data-model RUL that drifts optimistic inside an out-of-distribution excursion would, unguarded, schedule a swap too late. The Mahalanobis score (9) crosses its threshold in that window; the served estimate abandons the data channel and follows the conservative physics bound, then rejoins the fused estimate once the score falls back inside support. The estimator degrades toward caution, never toward a confident guess.

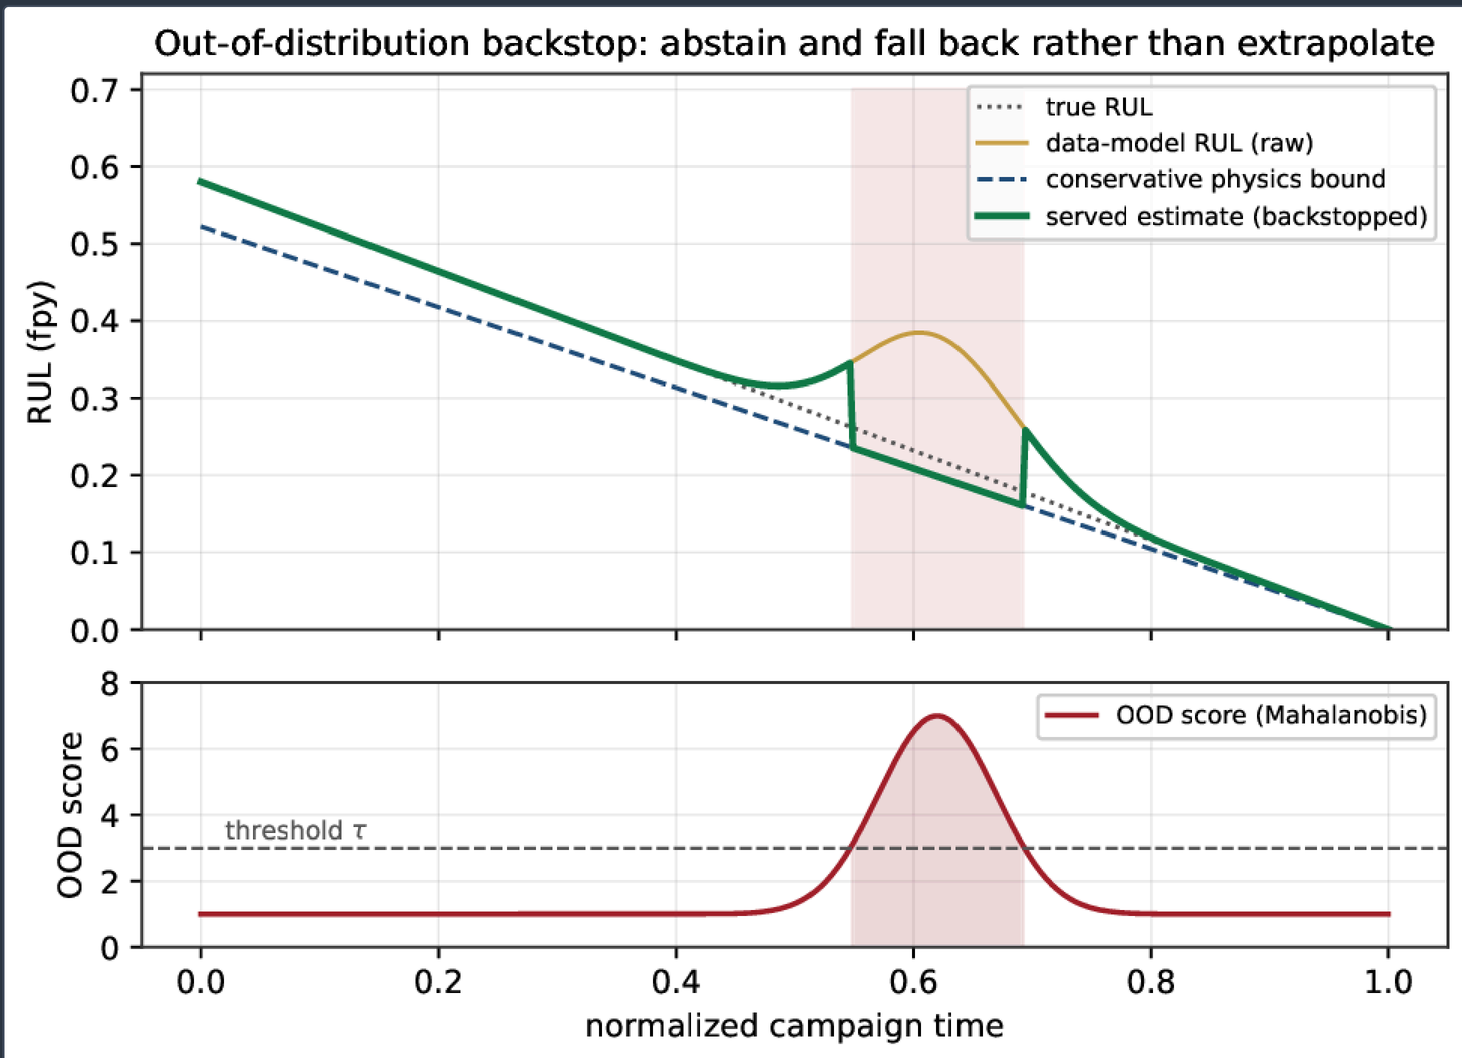

Out-of-distribution backstop (equations 9–11). Top: the served RUL (green) tracks the data model (gold) while in distribution but falls back to the conservative physics bound (blue dashed) inside the OOD window (red band), never following the optimistic excursion. Bottom: the Mahalanobis regime score crossing its calibrated threshold triggers the abstention.

THE END-OF-LIFE CRITERION IS A MEASURED CONDUCTOR PROPERTY

Figure 6 shows the conductor-degradation basis that fixes Φ_c in (2). The un-doped tape reaches its end-of-life criterion at a fast fluence of $5.64e22\ /m^2$, the APC-pinned tape earlier at $3.82e22\ /m^2$ (BR-SX-04); the un-doped-tape rule the campaign recommends therefore buys a $\approx 1.48\times$ fluence budget. The RUL of (2) integrates the fluence of (1) up to whichever break-even applies to the installed conductor.

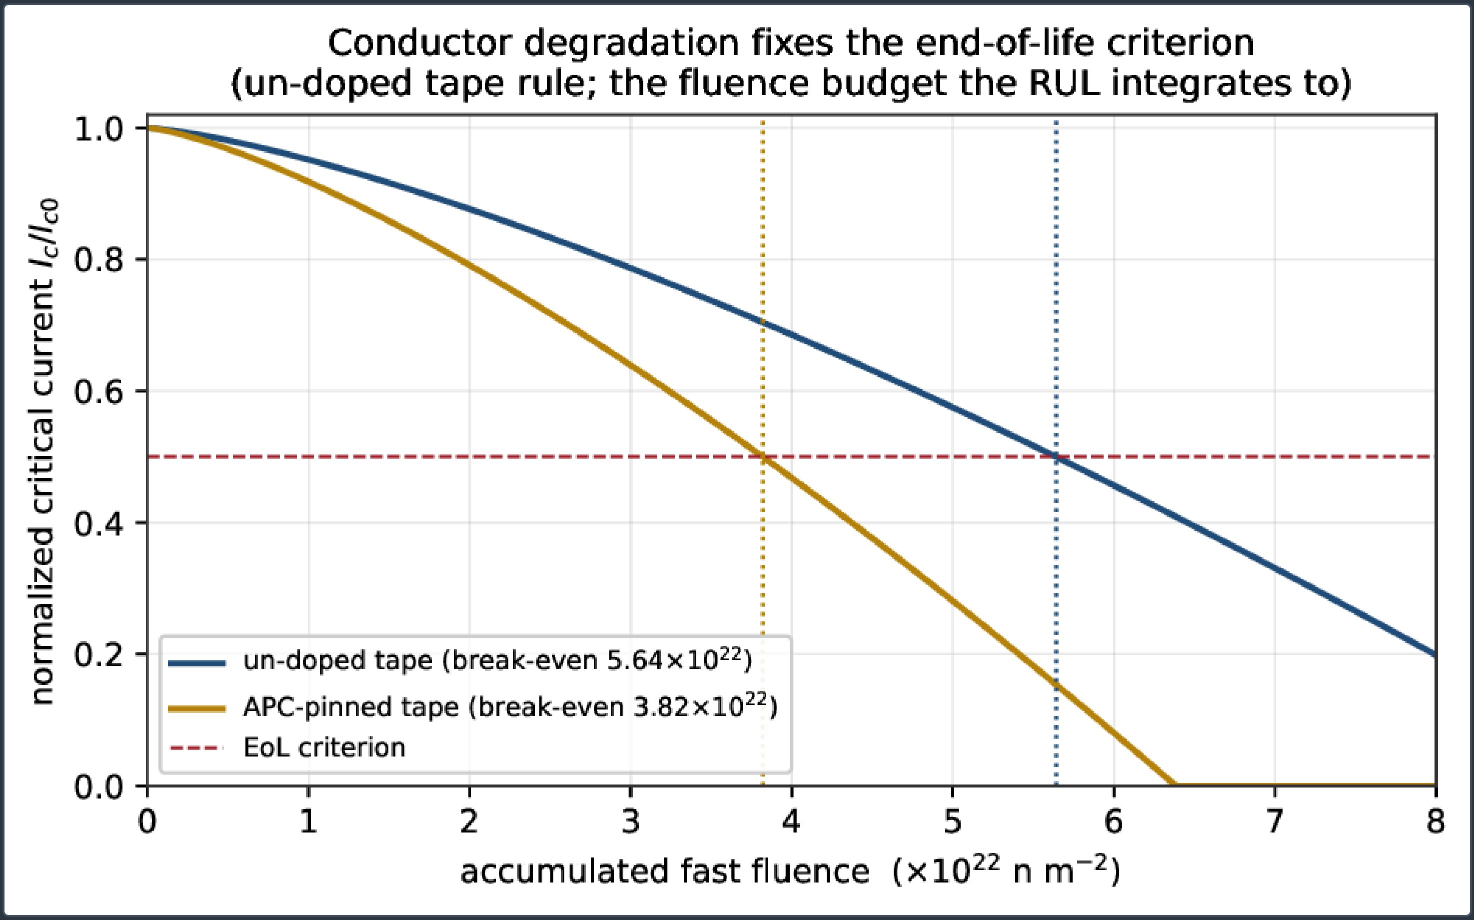

Normalised critical-current degradation versus accumulated fast fluence for un-doped and APC-pinned REBCO tape (BR-SX-04). The end-of-life criterion (dashed) is met at $5.64e22$ and $3.82e22\ /m^2$ respectively; these break-even fluences are the Φ_c the RUL integrates to. Curves are the degradation model anchored to the two measured break-even points.

CURRENT IS THE RECOVERY LEVER THE SCHEDULER PULLS

Because the fluence rate scales with fusion power and thus with plasma current, the operating point is itself a control on life. Figure 7 plots the wall-normalised damage rate and the resulting life across the admissible current window $I_p \in [8.779,\ 9.656]\ MA$ (KX-L1-A21), whose lower bound is MHD-set at $q_{95} = 3.3$ and whose upper bound is the centrepost damage ceiling of $18.85\ \mathrm{dpa\,fpy}^{-1}$ at the design point. Running mid-current roughly doubles the life (0.46 fpy against 0.21 fpy), at the cost of fusion power ($70.28\ MW$ against $85.04\ MW$); the scheduler weighs this trade explicitly through the duty-scheduler gate (KFE-CF-BR-019) [28,33].

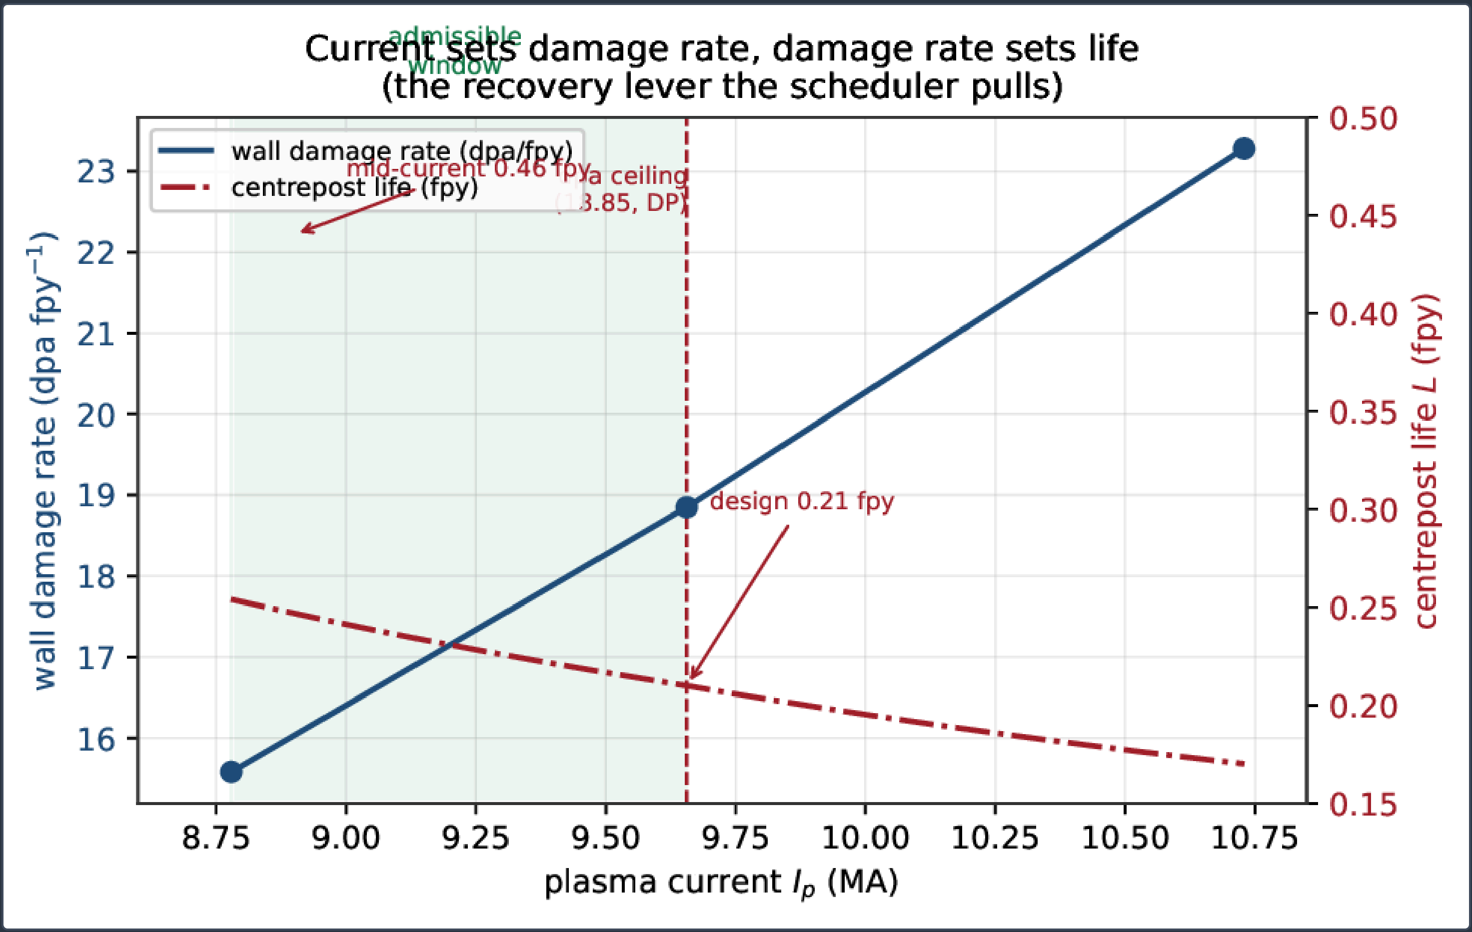

Wall damage rate (left, blue) and centrepost life (right, red) versus plasma current across the admissible window (KX-L1-A21). Anchors: $q_{95} = 3.30 \rightarrow I_p = 8.779$ MA, $15.58\ \mathrm{dpa\,fpy}^{-1}$; design point $I_p = 9.656$ MA, $18.85\ \mathrm{dpa\,fpy}^{-1}$; $q_{95} = 2.70 \rightarrow I_p = 10.729$ MA, $23.28\ \mathrm{dpa\,fpy}^{-1}$. Mid-current operation is the life-recovery lever.

FROM CADENCE TO FLEET AVAILABILITY

The RUL posterior feeds a cartridge-swap cadence — of order five swaps per full-power year at the design point (4.5 at 30 cm, ≈ 4.9 from the 3-D scan) — which, with the 21.5-day critical-path swap outage (40% of it thermal cycling), sets the unit availability. Figure 8 plots availability against coil life with the frozen operating points (design $A \approx 0.582\text{–}0.590$; mid-current $A \approx 0.638\text{–}0.648$, BR-SX-18) and the N+1 tiers. The honest reading is that swap cadence pulls availability below the CF-0.69 baseline at the design point, and that mid-current operation — the same lever as figure 7 — is the recovery [28,34,35]. Extending life toward the Tier IV bound (≥ 0.78 fpy) is the target the whole prognostic serves.

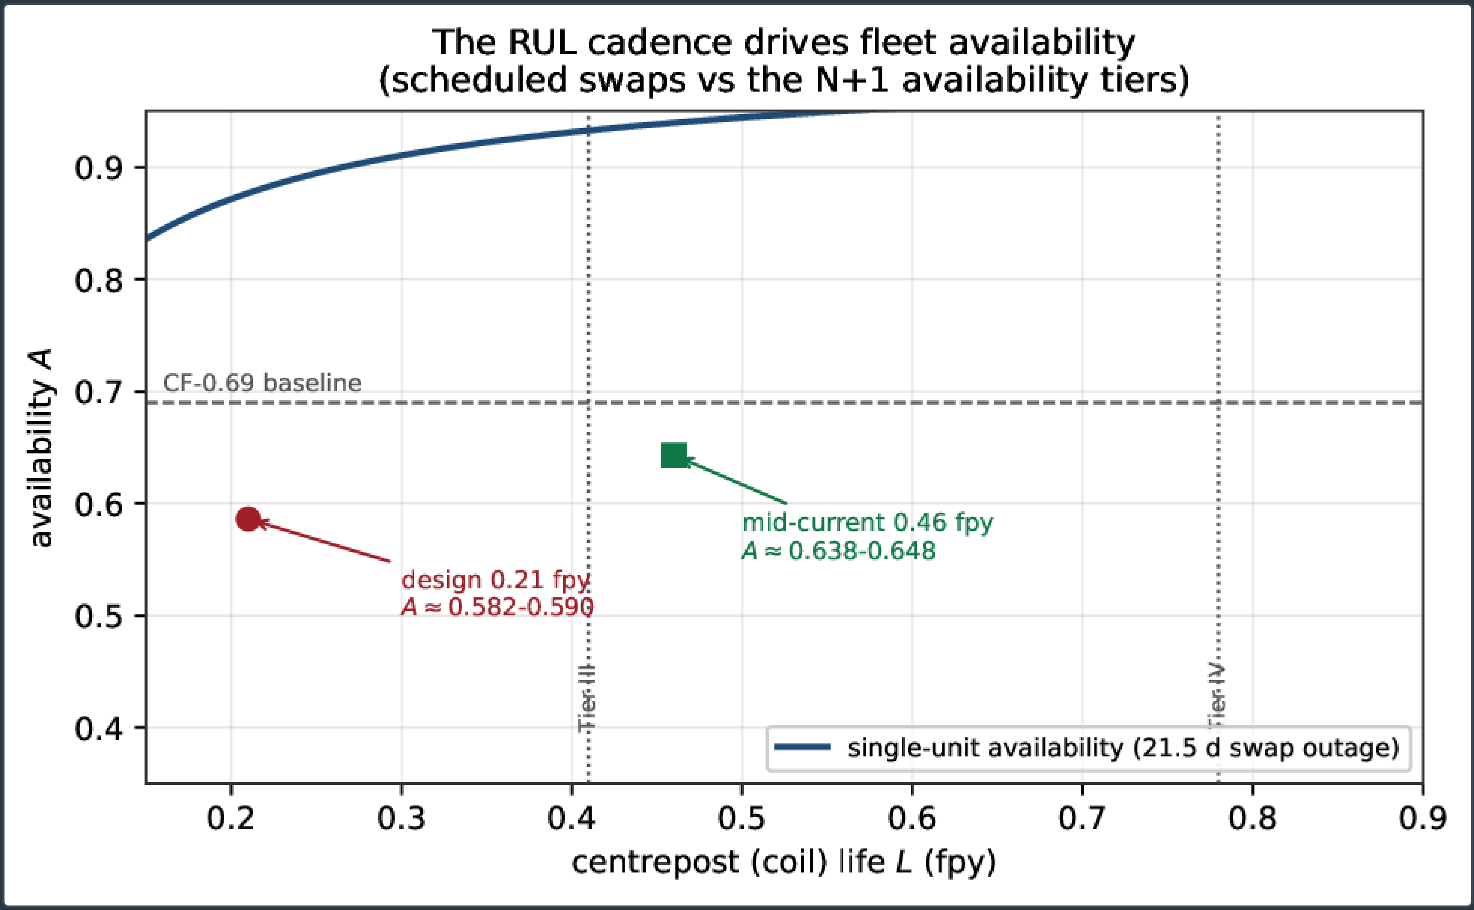

Single-unit availability versus coil life for a 21.5-day swap outage, with the frozen operating points (design and mid-current, BR-SX-18) and the N+1 availability tiers (register). The RUL cadence is the input; the availability tiers are what it must clear. Mid-current is the recovery lever.

PROGNOSTIC ACCURACY AND THE ARCHITECTURE

Figure 9 scores a synthetic estimator run against the α–λ accuracy cone of the standard PHM metrics [16]: the estimate stays within the $\alpha = 0.20$ band and converges to the true RUL as end of life approaches, the desired "tightens-toward-EoL" behaviour. Figure 9 is the architecture schematic: two channels, one backstop, one scheduled decision. Table 2 lists the elements and the failure mode each guards; Table 3 states the acceptance criteria and their status, honestly separating what is a frozen physics anchor from what is a pre-registered target.

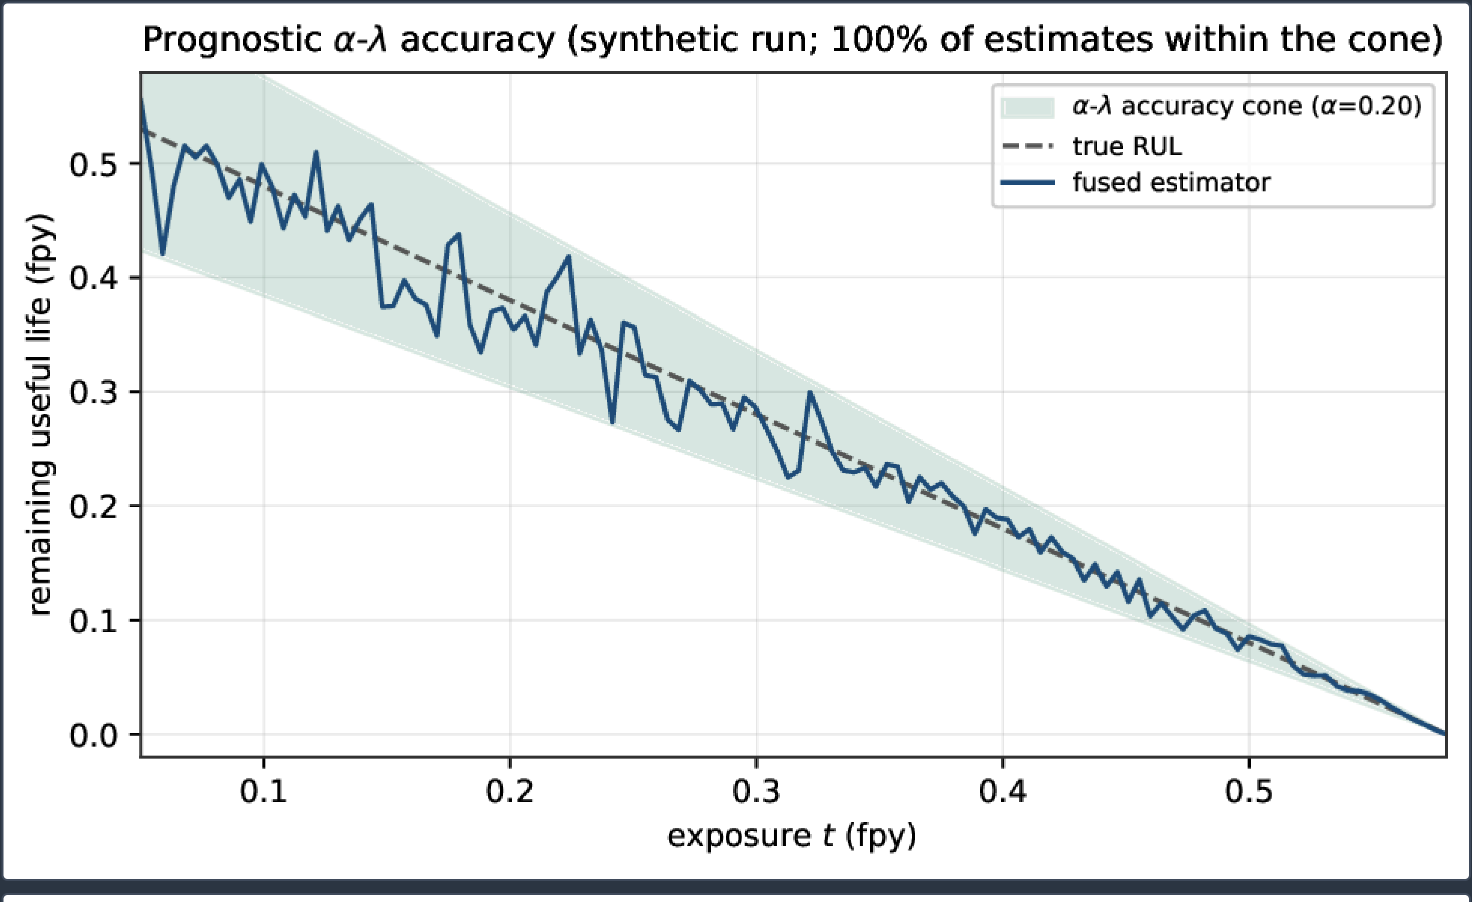

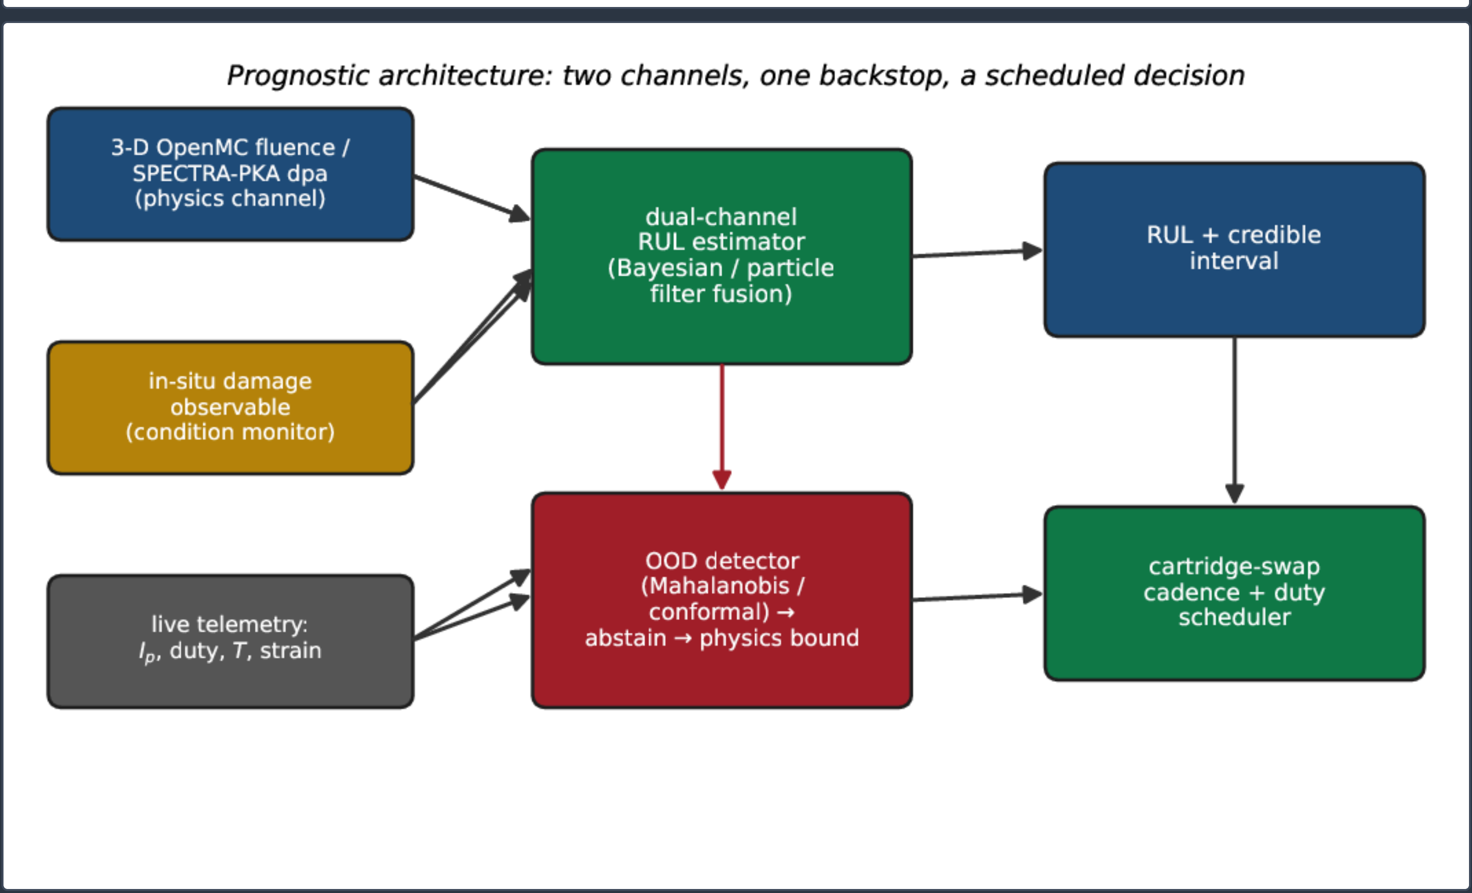

Prognostic α–λ accuracy (equation-agnostic PHM metric [16]) for a synthetic run of the fused estimator. The estimate remains inside the $\alpha = 0.20$ accuracy cone and converges to the true RUL toward end of life. Synthetic trace anchored to the 0.58 fpy design life.

ELEMENT	ROLE	FAILURE MODE IT GUARDS	
physics fluence model (1)–(2)	deterministic damage trajectory	sensor dropout / model absence	
in-situ damage observable	live estimate update	static worst-case interval	
particle-filter fusion (7)–(8)	RUL posterior + credible interval	over-confident point estimate	
Mahalanobis / conformal gate (9)–(10)	regime detection + coverage	silent model drift	
OOD backstop (11)	abstain → conservative bound	confident extrapolation error	
duty scheduler (KFE-CF-BR-019)	maps RUL → swap cadence \	current	wasteful or unsafe cadence

The prognostics architecture and the failure mode each element guards (inventions #35/#69/#71; serials KFE-CF-BR-107/BR-019/BR-021, SH-040/043/073).

CRITERION	TARGET	SERIAL	STATUS
resolved centrepost life	**0.58 fpy** @ 45 cm WC	BR-L1-A13b	**frozen anchor**
conductor break-even fluence	**$3.82/5.64 \times 10^{22}$ m^{-2}**	BR-SX-04	**frozen anchor**
admissible current window	**$[8.779, 9.656]$ MA**	KX-L1-A21	**frozen anchor**
swap cadence	~5 swaps/fpy (design)	BR-SX-12/18	**frozen anchor**
RUL credible-interval coverage	conformal $\geq 1-\alpha$	SH-073	target
OOD abstention	fall back on $d_M > \tau$	SH-040	target
α–λ accuracy	within cone toward EoL	—	target (demonstrated, synthetic)
dropout-robust estimate	graceful under sensor loss	SH-043	target

Acceptance criteria, targets and status. "Frozen anchor" denotes a measured de-risking gate; "target" a pre-registered bar for the real-time layer. Serials refer to the register [28].

LATENCY BUDGET

The prognostic is a slow loop by fusion-control standards — the damage timescale is months of full-power operation — but its guard elements run fast. Table 4 gives the design latency allocation: the OOD gate and the conformal wrapper are microsecond-class edge operations, the particle-filter update is a milliseconds-class operation, and the RUL forecast and scheduling run per shot or per shift. Nothing in the prognostic contends with the microsecond protective loops of the magnet twin.

STAGE	LOOP	LATENCY ALLOCATION	
Mahalanobis / conformal gate (9)–(10)	guard (edge)	microseconds	
particle-filter update (7)	fast	milliseconds	
RUL first-passage forecast (5)	medium	seconds (per shot)	
**cadence **	duty scheduling	slow	per shift / campaign
GPU transport / materials calibration	offline	campaign (register)	

Design latency allocation for the prognostic layer (allocations, not measurements). The estimator is a slow loop; only its OOD guard is edge-fast.

§ 5

Discussion

AGAINST THE PROGNOSTICS LITERATURE

The framework is a faithful instance of the statistical-data-driven RUL paradigm reviewed by Si *et al.* [7] and the condition-based-maintenance paradigm of Jardine *et al.* [5], specialised to a nuclear consumable. Two features distinguish it from the canonical benchmark studies. First, unlike the NASA C-MAPSS turbofan set [15], where the degradation model is hidden and must be learned end-to-end, the centrepost's degradation *driver* is a first-principles transport quantity: the fast fluence of equation (1) is computed, not inferred, so the physics channel is a genuine model rather than a regression, and the data channel is a correction on it rather than the whole estimate. This is closer to the physics-of-failure end of the taxonomy [6] and is exactly why the OOD backstop can be a physics bound rather than a second data model. Second, the estimator is scored on the same prognostic-specific metrics that the field settled on — the α–λ accuracy cone and the prognostic horizon of Saxena *et al.* [16] (figure 9) — rather than a bare error, because in a nuclear maintenance context the shape of the convergence toward end of life matters more than the mean error far from it. The particle filter (7) is the same tool Orchard and Vachtsevanos [12] and the wider literature [14,13] established for nonlinear, non-Gaussian degradation, and the Gaussian-product fusion of (8) is its linear-Gaussian shadow.

THE BACKSTOP AGAINST THE RELIABILITY LITERATURE

The load-bearing element — abstention out of distribution — is drawn directly from the machine-learning reliability literature and is, we argue, the piece most often missing from applied PHM. The Mahalanobis regime score (9) is the detector of Lee *et al.* [19]; the conformal wrapper (10) inherits the finite-sample, distribution-free coverage guarantee of Vovk-style conformal prediction [21,20]; and the ensemble/deep-uncertainty view that motivates treating epistemic uncertainty as a first-class, routable output is that of Gal and Ghahramani and Lakshminarayanan *et al.* [17,18]. The novelty here is not any one of these methods but their composition into a *fail-safe* prognostic: when the data model is out of its depth, the served estimate is not a wider-but-still-data interval, it is the deterministic physics bound. That is a stronger guarantee than calibrated uncertainty alone, and it is what makes the estimator admissible on a nuclear magnet. The discipline is shared with the wider Kronos control stack [36,37] and with the learned-control literature's move toward uncertainty-aware, envelope-bounded action [26,27].

WHAT THE ESTIMATE IS WORTH

The concrete value is the collapse of figure 2. A design book must pick one life; picking the conservative **0.199–0.21 fpy** swaps too often and depresses availability below the CF-0.69 baseline (figure 8), while picking the optimistic **0.58–0.765 fpy** risks an unplanned outage. A live posterior that tracks the machine's measured exposure schedules each swap on that cartridge's own trajectory, with a stated interval the operator can plan against; coupled to the current lever of figure 7, it lets the scheduler trade power for life deliberately — the mechanism by which the fleet reaches the Tier IV availability bound. The prognostic thus sits at the centre of the consumable-centrepost operating case: it consumes the fluence and conductor physics [29,31], the current window of the operating-window synthesis [33], and feeds the swap logistics and fleet RAMI [34,35].

TWO DISTINCT MAGNETS

Throughout, the breeder centrepost is treated as its own device. Its peak field is 16.84 T; the plug coil of the separate D–^{3}He tandem-mirror burner is a distinct, higher-field magnet at 26.49 T (KX-L1-A15) [28]. The prognostic in this paper is the centrepost's; the burner plug carries its own lifetime and integrity case and is never conflated with the centrepost field.

§ 6

Limitations

One honest gate bounds the claims. The in-winding damage observable that sharpens the data channel is only as good as its own survival in the centrepost radiation field, and the leading candidate — an in-winding NV-diamond sensor — carries an open neutron-tolerance item: its required in-service fluence tolerance is about $39\times$ above the demonstrated diamond-sensor state of the art (BR-SX-14, BR-SXD-05) [28]. The framework is deliberately designed so that this does *not* compromise the protective function: on a degraded or failed observable the Mahalanobis gate (9) flags the loss of support and the estimator falls back to the physics bound (11), so the served RUL degrades to the conservative deterministic estimate rather than to a confident wrong one. A second, separate open item on the centrepost — a $\approx 2\,cm$ radial-budget overrun of the winding-plus-shield stack against the inboard envelope (BR-SXD-06) [28] — bounds how much shield, and therefore how much life, is available to the physics channel in the first place; it is a geometry item carried in the swap-logistics paper [34]. Finally, the data channel's fidelity, the OOD threshold τ, and the conformal calibration set are design targets to be validated on a coil-prototype hardware-in-the-loop bench before the estimator earns runtime authority. This is a prognostics architecture demonstrated in model form against frozen physics anchors, not a qualified health-monitoring system.

§ 7

Conclusion

The consumable centrepost of a compact spherical-tokamak breeder retires on a fluence budget, not a calendar, and its life is not one number but a family spanning more than three orders of magnitude across modelling bases. We have formalised the replacement decision as a prognostics problem and given a dual-channel, fluence-driven remaining-useful-life estimator: a physics channel that integrates the transport-anchored fast fluence into a Miner cumulative-damage variable against a measured conductor break-even fluence (equations 1–2); a Bayesian particle-filter fusion with a damage-correlated in-situ observable that returns a RUL posterior with a credible interval (equations 4–8); and an out-of-distribution backstop that abstains to the conservative physics bound whenever the live signals leave calibrated support (equations 9–11). Calibrated to the register-resolved 0.58 fpy design life behind a 45 cm WC shield, the estimator collapses the life spread to a live schedule, drives a cartridge-swap cadence of order five per full-power year, uses the admissible current window as an explicit life-recovery lever, and feeds the fleet-availability tiers. The one open item — the neutron tolerance of the in-winding observable — is contained by the same backstop that makes the estimator fail-safe. The result is a prognostic that turns a worst-case interval into a live, confidence-bounded, honest maintenance decision: a methods contribution distinct from the centrepost's structural-integrity, quench-detection and swap-logistics cases in the Kronos series.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos physics de-risking programme for the frozen fluence, conductor, current-window, structural and quench anchors underlying every quantitative claim, and the NVIDIA GPU HPC campaign on which the underlying transport and materials physics was computed.

Reproducibility. Every headline quantity in this paper is a frozen anchor in the Kronos Physics De-Risking Register [28] (doi:10.5281/zenodo.22133057): the fluence–shield–life map (BR-L1-A13b), the 3-D fluence/dpa scan (BR-SX-12), the conductor break-even fluences (BR-SX-04), the admissible current window (KX-L1-A21), and the swap-cadence/availability points (BR-SX-12/16/18). The centrepost-life numbers reproduce arithmetically from the deposited OpenMC fluence data (ENDF/B-VIII.0), and the ten estimator figures are regenerated by the single deposited script `make_figs_cp_rul_2_85.py` from those anchors. This paper is archived at doi:10.5281/zenodo.22645788 and its official version is at https://www.kronosfusionenergy.com/publications/centrepost-prognostics-and-rul-control.html. No result depends on unpublished or proprietary data, and no economic quantity enters the analysis.

Data availability tocsectionData availability The fluence–shield–life anchors, conductor break-even fluences, current-window damage rates, and swap-cadence/availability points used for figures 1–9 are archived with the deposit at doi:10.5281/zenodo.22645788 and trace to the register [28] 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 (inventions #35/#69/#71) 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 J E Menard *et al*, Fusion nuclear science facilities and pilot plants based on the spherical tokamak, *Nucl. Fusion* **56** 106023 (2016). doi:10.1088/0029-5515/56/10/106023

2 B N Sorbom *et al*, ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* **100** 378 (2015). doi:10.1016/j.fusengdes.2015.07.008

3 S Hahn *et al*, No-insulation multi-width winding technique for high temperature superconducting magnet, *Appl. Phys. Lett.* **103** 173511 (2013). doi:10.1063/1.4826217

4 S Hahn *et al*, 45.5-tesla direct-current magnetic field generated with a high-temperature superconducting magnet, *Nature* **570** 496 (2019). doi:10.1038/s41586-019-1293-1

5 A K S Jardine, D Lin and D Banjevic, A review on machinery diagnostics and prognostics implementing condition-based maintenance, *Mech. Syst. Signal Process.* **20** 1483 (2006). doi:10.1016/j.ymssp.2005.09.012

6 J Lee, F Wu, W Zhao, M Ghaffari, L Liao and D Siegel, Prognostics and health management design for rotary machinery systems, *Mech. Syst. Signal Process.* **42** 314 (2014). doi:10.1016/j.ymssp.2013.06.004

7 X-S Si, W Wang, C-H Hu and D-H Zhou, Remaining useful life estimation — a review on the statistical data driven approaches, *Eur. J. Oper. Res.* **213** 1 (2011). doi:10.1016/j.ejor.2010.11.018

8 M A Miner, Cumulative damage in fatigue, *J. Appl. Mech.* **12** A159 (1945). doi:10.1115/1.4009458

9 P Paris and F Erdogan, A critical analysis of crack propagation laws, *J. Basic Eng.* **85** 528 (1963). doi:10.1115/1.3656900

10 S J Zinkle and G S Was, Materials challenges in nuclear energy, *Acta Mater.* **61** 735 (2013). doi:10.1016/j.actamat.2012.11.004

11 G S Was, *Fundamentals of Radiation Materials Science: Metals and Alloys*, Springer, Berlin (2007). doi:10.1007/978-3-540-49472-0

12 M E Orchard and G J Vachtsevanos, A particle-filtering approach for on-line fault diagnosis and failure prognosis, *Trans. Inst. Meas. Control* **31** 221 (2009). doi:10.1177/0142331208092026

13 M S Arulampalam, S Maskell, N Gordon and T Clapp, A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking, *IEEE Trans. Signal Process.* **50** 174 (2002). doi:10.1109/78.978374

14 N Z Gebraeel, M A Lawley, R Li and J K Ryan, Residual-life distributions from component degradation signals: a Bayesian approach, *IIE Trans.* **37** 543 (2005). doi:10.1080/07408170590929018

15 A Saxena, K Goebel, D Simon and N Eklund, Damage propagation modeling for aircraft engine run-to-failure simulation, *Int. Conf. on Prognostics and Health Management* (2008). doi:10.1109/PHM.2008.4711414

16 A Saxena, J Celaya, E Balaban, K Goebel, B Saha, S Saha and M Schwabacher, Metrics for evaluating performance of prognostic techniques, *Int. Conf. on Prognostics and Health Management* (2008). doi:10.1109/PHM.2008.4711436

17 Y Gal and Z Ghahramani, Dropout as a Bayesian approximation: representing model uncertainty in deep learning, *Proc. 33rd Int. Conf. on Machine Learning (ICML)* 1050 (2016). arXiv:1506.02142

18 B Lakshminarayanan, A Pritzel and C Blundell, Simple and scalable predictive uncertainty estimation using deep ensembles, *Adv. Neural Inf. Process. Syst. (NeurIPS)* **30** (2017). arXiv:1612.01474

19 K Lee, K Lee, H Lee and J Shin, A simple unified framework for detecting out-of-distribution samples and adversarial attacks, *Adv. Neural Inf. Process. Syst. (NeurIPS)* **31** (2018). arXiv:1807.03888

20 A N Angelopoulos and S Bates, Conformal prediction: a gentle introduction, *Found. Trends Mach. Learn.* **16** 494 (2023). doi:10.1561/2200000101

21 J Lei, M G'Sell, A Rinaldo, R J Tibshirani and L Wasserman, Distribution-free predictive inference for regression, *J. Am. Stat. Assoc.* **113** 1094 (2018). doi:10.1080/01621459.2017.1307116

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

23 M J Norgett, M T Robinson and I M Torrens, A proposed method of calculating displacement dose rates, *Nucl. Eng. Des.* **33** 50

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

Prognostics and Remaining-Useful-Life Estimation for the Consumable Centrepost: A Dual-Channel Fluence-Driven Framework Scheduling Cartridge Replacement in a Spherical-Tokamak Breeder

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

Zenodo · doi:10.5281/zenodo.22645788

Read on kronosfusionenergy.com

The Physics De-Risking Register