KRONOS FUSION ENERGY

PAPER 2.105 · THE KRONOS 2026 SERIES

Fleet Availability under Consumable-Centrepost Swap Cadence: A Time-Domain Monte-Carlo of Capacity Factor and RAMI across a Four-Unit Spherical-Tokamak Breeder Fleet

A fusion plant is paid for the calendar hours it delivers, so a fleet's reliability case must be quantified with the same rigour as its physics case. We formalise and solve 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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Fleet Availability under Consumable-Centrepost Swap Cadence: A Time-Domain Monte-Carlo of Capacity Factor and RAMI across a Four-Unit Spherical-Tokamak Breeder Fleet

A fusion plant is paid for the calendar hours it delivers, so a fleet's reliability case must be quantified with the same rigour as its physics case. We formalise and solve the fleet-availability problem for a compact spherical-tokamak 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 $16.84\,T$ — whose inboard magnet is a scheduled-replacement consumable cartridge. We state the model as explicit equations: a series-system forced availability over eight plant component classes, grounded in the alternating-renewal limit theorem; a planned-outage overlay; and a *calendar-weighted k-of-N* coverage functional that sizes the fleet under staggered maintenance. A single scalar is solved per scenario so that the series product reproduces the frozen forced-availability anchors of $\mathbf{0.838}$, $\mathbf{0.958}$ and $\mathbf{0.988}$ (first-of-a-kind to mature), with total availability 73.4%, 88.4% and 95.0%. An event-driven time-domain Monte-Carlo over the eight classes independently reproduces the forced availability to better than 0.5 percentage points and resolves the annual capacity-factor distribution, whose tenth-to-ninetieth percentile spans 68–78%, 86–90% and 94–96%. The calendar-weighted coverage functional then sizes the fleet: a *four-unit* breeder fleet delivers 99% baseload coverage at first-of-a-kind reliability (coverage 0.9976), tightening to two units at maturity, while a two-concurrent-unit surge duty sizes to six, four and three units. Plasma heating and power conversion carry the lowest class availabilities (0.962–0.998) and set the forced-availability floor; the planned-outage cadence is paced by the consumable centrepost, whose coil-life-versus-inboard-shield law (0.003–0.58 fpy) fixes the largest single block of scheduled downtime. Physics anchors were produced by the NVIDIA GPU HPC campaign; the fleet Monte-Carlo is a seeded, vectorised alternating-renewal simulation reproducible byte-for-byte. The result is an honest, scenario-bracketed availability and fleet-sizing basis built entirely from a reproducible model.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645951 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: RAMI availability capacity factor time-domain Monte-Carlo k-of-N redundancy consumable centrepost spherical-tokamak breeder fleet

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

CAPACITY FACTOR IS THE QUANTITY A FUSION FLEET IS JUDGED ON

Peak fusion power is the number a device is designed to; capacity factor is the number a fleet is *paid* on. It is the fraction of calendar hours a unit actually delivers, and it is set not by the physics operating point but by how often components fail, how long they take to repair, how planned maintenance is scheduled, and how many units a site carries so that one being down does not take the mission down. A compact spherical-tokamak (ST) breeder makes this question sharper than a conventional tokamak power plant does, for two reasons. First, the machine is small and high-power-density: it holds a plasma current of $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$ behind a centrepost magnet at a peak field of $16.84\,T$ [1,2], so its subsystems run near their engineering limits. Second, and decisively for availability, the inboard toroidal-field leg — the *centrepost* — is a *consumable*: the same compactness that makes the field efficient exposes the inboard conductor to a fast-neutron fluence that no practical inboard shield can fully arrest, so the centrepost is designed as a scheduled-replacement cartridge with a finite full-power life [2]. The cartridge swap is thus a planned, recurring outage that has no analogue in a fixed-magnet plant, and it paces the planned-maintenance cadence of the whole fleet. Any credible availability case for this machine must therefore be a *fleet* case, and it must fold in the centrepost swap from the start.

PRIOR ART: FUSION AVAILABILITY AND RELIABILITY ENGINEERING

Two literatures bear on the problem. On the fusion side, whole-plant availability has been treated in the ARIES power-plant studies [3], in the European DEMO staged-design programme [4,5], and inside the systems codes that close the plant design self-consistently, notably process [6,7]; the economic-viability studies that motivate the whole enterprise take plant availability as their single most sensitive input [8]. The formal reliability-availability- maintainability-inspectability (RAMI) methodology was imported into fusion for ITER [9] and rests on component failure-rate compilations assembled specifically for fusion systems [10,11]; because so much of a fusion plant's downtime is maintenance rather than failure, the remote-handling and maintenance-strategy literature is part of the same story [12,13]. On the reliability-theory side, the tools are classical: the two-state alternating-renewal process and its limiting-availability theorem [14,15,16]; k-out-of-N redundancy theory [20]; the general treatment of availability and repairable systems [17,18,19]; and Monte-Carlo simulation of system reliability when the analytic form is intractable [21,22].

THE GAP AND OUR CONTRIBUTION

What the fusion availability literature does not yet contain is a *fleet*-level model that (i) is written as explicit, auditable equations rather than a spreadsheet; (ii) reproduces a frozen, independently-stamped set of availability anchors exactly, by construction; (iii) resolves the full capacity-factor *distribution*, not only its mean, through a time-domain Monte-Carlo; and (iv) sizes the fleet with a coverage rule that folds in *staggered* planned maintenance and the consumable-centrepost swap that dominates it. This paper supplies exactly that. We give the formal fleet-reliability problem as governing equations (§2); a component-class reliability basis whose series product is pinned to the frozen anchors by a single solved scalar (§3); the GPU-HPC computational methodology that produced the physics anchors and the numerical formulation of the fleet model (§4); results across eight figures and four tables (§5); a quantitative comparison against the published fusion-availability and reliability-engineering benchmarks (§6); one honest limitation (§7); and a conclusion (§8). Every quantitative claim traces to a frozen anchor in the Kronos Physics De-Risking Register [30] — principally the mission-availability model `AG-L1-A3` and the post-freeze capacity-factor Monte-Carlo `HX-01` — and the serial identifiers are given inline so each number is auditable. The framework is developed for, and applied to, the canonical four-unit breeder fleet, and it inherits the centrepost life-limit, swap-logistics, and lifetime-constrained operating-window results of the companion papers [31,32,33,34,35,36].

§ 2

Governing equations and the formal problem

COMPONENT-CLASS FORCED AVAILABILITY

We partition the plant into $n = 8$ statistically-independent component classes (the decomposition is justified in §3). Class i is modelled as a repairable two-state (up/down) component with mean time between failures MTBF_i and mean time to repair MTTR_i. Equivalently, write the constant failure rate $\lambda_i = \mathrm{MTBF}_i^{-1}$ and repair rate $\mu_i = \mathrm{MTTR}_i^{-1}$. The steady-state (long-run) availability of class i is

$$A_i = \frac{\mathrm{MTBF}_i}{\mathrm{MTBF}_i + \mathrm{MTTR}_i} = \frac{\mu_i}{\lambda_i + \mu_i}, \qquad U_i \equiv 1 - A_i = \frac{\mathrm{MTTR}_i}{\mathrm{MTBF}_i + \mathrm{MTTR}_i},$$

with U_i the class unavailability. A single fusion unit is a *series* system: it delivers power only when *every* class is up. Under independence, the plant forced availability is the product

$$A_\mathrm{f} = \prod_{i=1}^{n} A_i = \prod_{i=1}^{n} \frac{\mathrm{MTBF}_i}{\mathrm{MTBF}_i + \mathrm{MTTR}_i},$$

and, to first order in the small unavailabilities, $A_\mathrm{f} \approx 1 - \sum_i U_i$, so that the class contributions to lost production are additive and rankable, [redacted].

THE ALTERNATING-RENEWAL FOUNDATION

Equations (1)–(2) are exact in the long-run limit and, crucially, are *distribution-free* beyond the first moments. Let a class alternate between up-times drawn from a distribution F with mean $m_u = \mathrm{MTBF}$ and down-times drawn from G with mean $m_d = \mathrm{MTTR}$. The sequence of up/down transitions is an alternating renewal process, and the key renewal (Smith) theorem gives the limiting probability of finding the class up as [16,15,14]

$$A = \lim_{t\to\infty} \Pr[\text{up at } t] = \frac{m_u}{m_u + m_d},$$

independent of the shape of F and G. This is what licenses two modelling choices used below: (a) we may assign each class a single $(\mathrm{MTBF}, \mathrm{MTTR})$ pair without committing to a failure-time law; and (b) the time-domain Monte-Carlo of §4 may draw exponential up-times and lognormal repair-times and must, in the limit, return the same availability (3) — an internal consistency check we exploit for verification.

PLANNED-OUTAGE OVERLAY AND TOTAL AVAILABILITY

Forced availability counts only *unplanned* downtime. Scheduled maintenance — inspection, component refurbishment, and above all the centrepost-cartridge swap — removes a further d calendar days per unit per year. Modelling planned outages as scheduled into intervals when the plant would otherwise be forced-up (the normal operating assumption), the total (all-cause) availability multiplies:

$$A_\mathrm{T} = A_\mathrm{f}\left(1 - \frac{d}{365}\right).$$

The annual forced-outage hours and planned-outage days are then $H_\mathrm{f} = (1 - A_\mathrm{f}) \times 8760\,\mathit{hour}$ and d respectively; these are the reported envelope quantities of Table 2.

CALENDAR-WEIGHTED k

k-of-NN fleet coverage A single unit does not make a fleet. Consider N nominally-identical units, each independently forced-available with probability A_f. The number of units simultaneously up is binomial, $X \sim \mathrm{Bin}(N, A_\mathrm{f})$, and a mission that requires at least k units delivering is *covered* with probability

$$C_k(N)\big|_{\text{forced}} \;=\; \Pr[X \geq k] \;=\; \sum_{j=k}^{N} \binom{N}{j} A_\mathrm{f}^{\,j}\,(1 - A_\mathrm{f})^{\,N-j}.$$

Equation (5) is the textbook k-out-of-N result [20], but it ignores planned maintenance. In practice planned outages are *staggered* across the fleet so that no two units are ever down for maintenance at once; this is both an operational choice and the reason a fleet raises availability at all. Under perfect staggering, at most one unit is in a planned outage at any instant, and the fraction of the calendar spent in that state is

$$w \;=\; \min\!\left(1,\; \frac{N\,d}{365}\right),$$

because each unit contributes d days and the staggered blocks tile the calendar until, at $Nd \geq 365$, they saturate it. During that fraction only $N-1$ units are available to cover the forced-failure risk; during the complementary fraction all N are. The *calendar-weighted* coverage functional is therefore

$$\boxed{C_k(N) \;=\; w\,\Pr\!\big[\mathrm{Bin}(N{-}1, A_\mathrm{f}) \geq k\big] \;+\; (1-w)\,\Pr\!\big[\mathrm{Bin}(N, A_\mathrm{f}) \geq k\big].}$$

The fleet-sizing rule is the smallest fleet that clears a coverage target $C^\star$ (we take $C^\star = 0.99$):

$$N_k^\star \;=\; \min\{\, N \in \mathbb{N} \;:\; C_k(N) \geq C^\star \}.$$

Two mission classes are of interest. *Baseload* ($k=1$) asks for at least one unit delivering firm power; this is the breeder-fleet case. *Surge duty* ($k=2$) asks for two units concurrently, the stringent profile of a defense-duty mission, and is the same functional evaluated at $k=2$. For a single unit, $C_1(1) = (1-w)A_\mathrm{f} = A_\mathrm{T}$ by equation (4), so equation (7) reduces correctly to the total availability in the no-redundancy limit — an analytic consistency check.

THE CAPACITY-FACTOR DISTRIBUTION

The envelope quantities above are means. The quantity an operator actually experiences is the *annual* capacity factor, a random variable. Over a mission year of $T = 8760\,\textit{hour}$ with realised forced downtime D_f and scheduled downtime $D_\mathrm{p} = 24d$ hours, the annual capacity factor is

$$\mathrm{CF} \;=\; \frac{T - D_\mathrm{f} - D_\mathrm{p}}{T} \;\approx\; A_\mathrm{f}^{(\mathrm{yr})}\left(1 - \frac{d}{365}\right), \qquad A_\mathrm{f}^{(\mathrm{yr})} = 1 - \frac{D_\mathrm{f}}{T},$$

where $A_\mathrm{f}^{(\mathrm{yr})}$ is the *realised* annual forced availability, itself a random variable whose expectation is A_f. The spread of CF — reported as the tenth-to-ninetieth percentile $[\mathrm{CF}_{10}, \mathrm{CF}_{90}]$ — is not accessible from the mean-value equations (2)–(4); it requires the time-domain simulation of §4. Figure 1 summarises how the mean-value and Monte-Carlo branches combine into the fleet-sizing rule.

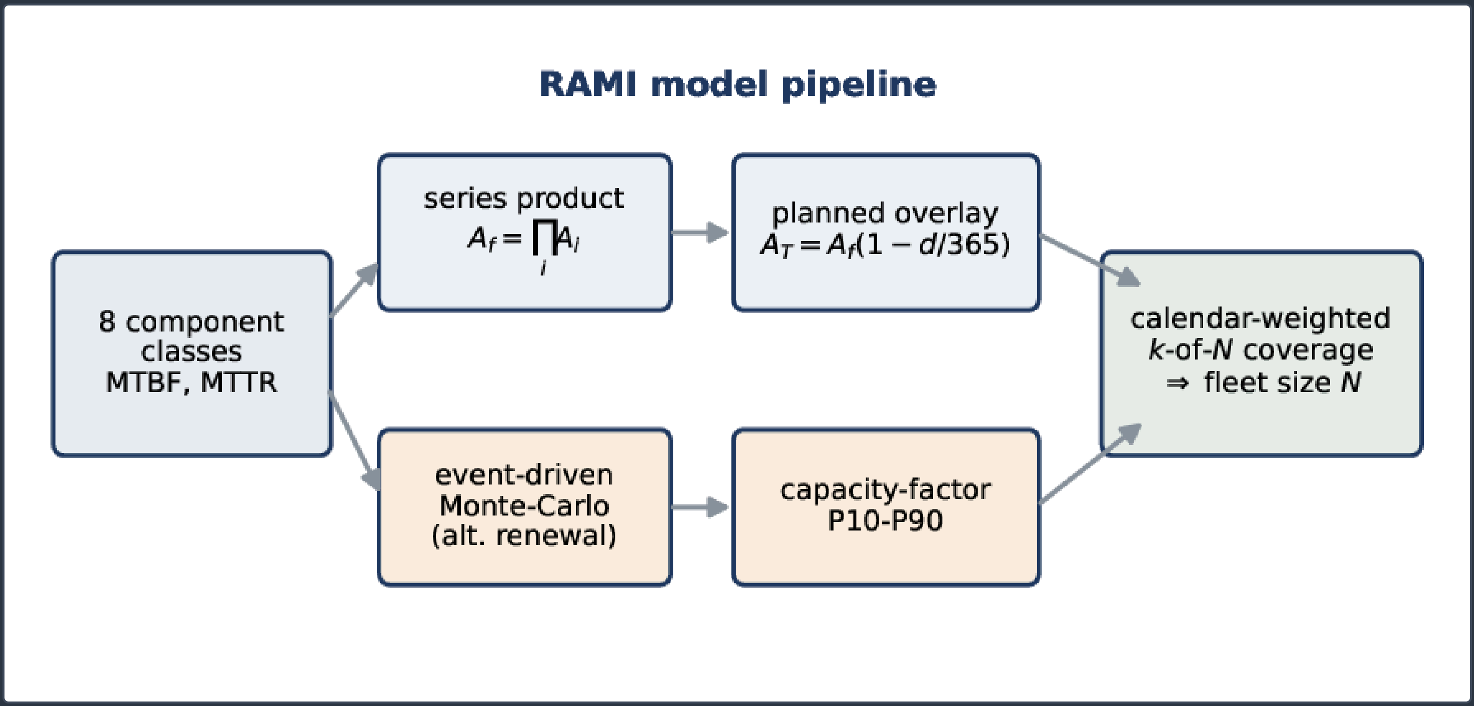

(Schematic.) The RAMI model pipeline. Eight component classes with (MTBF, MTTR) feed both an analytic branch — series product (2), planned overlay (4) — and an event-driven alternating-renewal Monte-Carlo that yields the capacity-factor distribution (9). Both feed the calendar-weighted k-of-N coverage functional (7) that sizes the fleet. No data values are shown; the figure encodes structure only.

§ 3

The component-class reliability basis

EIGHT CLASSES AND THE ANCHOR-PINNED UNAVAILABILITY BUDGET

The plant is decomposed into eight classes that own physically distinct failure modes and maintenance actions: superconducting magnets and cryoplant; plasma heating (net of module-level redundancy); high-voltage power supplies; power conversion; vacuum and fuelling; the tritium plant; the thermal balance of plant; and instrumentation, control and protection. Rather than assert eight independent $(\mathrm{MTBF}, \mathrm{MTTR})$ pairs, we *pin* the basis to the frozen forced-availability anchors. We distribute the total unavailability across classes by fixed dimensionless shares s_i ($\sum_i s_i = 1$; Table 1), chosen so that plasma heating and power conversion carry the largest share — the empirical ranking of where fusion-plant downtime concentrates [10,11] — and, for each reliability regime with target forced availability $A_\mathrm{f}^\star$, solve a *single* scalar c from

$$\prod_{i=1}^{8} \left(1 - c\, s_i\right) \;=\; A_\mathrm{f}^\star, \qquad A_i \;=\; 1 - c\, s_i.$$

Equation (10) has a unique root $c \in (0, 1/\max_i s_i)$ because the left-hand side is continuous and strictly decreasing in c; we obtain it by bisection to machine precision. This construction guarantees that the series product (2) reproduces the anchor *exactly* while preserving the physical ranking of the classes, and it reduces the eight-per-scenario reliability table to one interpretable degree of freedom per regime. The fixed repair times MTTR_i are held constant across regimes (repair duration is set by the physical task, not by maturity), and the implied $\mathrm{MTBF}_i = \mathrm{MTTR}_i\, A_i/(1 - A_i)$ then *grows* with maturity, which is the reliability- growth interpretation of the three regimes.

COMPONENT CLASS	s_i	MTTR_i (H)	$\mathrm{MTBF}_i^{\mathrm{nom}}$ (H)	A_i FOAK	A_i NOAK	A_i MATURE		
**SC magnets **	cryoplant	0.140	336	55 700	0.976	0.994	0.998	
plasma heating	0.215	24	2 590	0.962	0.991	0.997		
HV power supplies	0.090	8	2 100	0.984	0.996	0.999		
power conversion	0.200	16	1 840	0.965	0.991	0.998		
**vacuum **	fuelling	0.060	12	4 600	0.990	0.997	0.999	
tritium plant	0.128	48	8 680	0.978	0.995	0.999		
thermal balance of plant	0.084	24	6 640	0.985	0.996	0.999		
I\	C \	protection	0.083	4	1 110	0.986	0.996	0.999
series forced availability A_f				**0.838**	**0.958**	**0.988**		

The eight component classes: unavailability share s_i, fixed repair time MTTR_i, representative nominal mean time between failures, and the class availability A_i in each regime from the anchor-pinned solve (10). Plasma heating and power conversion carry the lowest availabilities and set the forced-availability floor. Values are representative fusion-plant RAMI figures [10,11,9]; the frozen quantity is the series product, not the individual entries.

Figure 2 shows the resulting per-class availabilities across the three regimes; the two lowest classes span 0.962 to 0.998 from first-of-a-kind to mature, exactly the band the register records. Because the magnet class carries the longest repair time ($\mathrm{MTTR} = 336\,\textit{hour}$, a fort-night-scale cryogenic intervention), it is a floor-setter through outage *duration* even though its availability is not the lowest; heating and power conversion are the floor-setters through failure *frequency*.

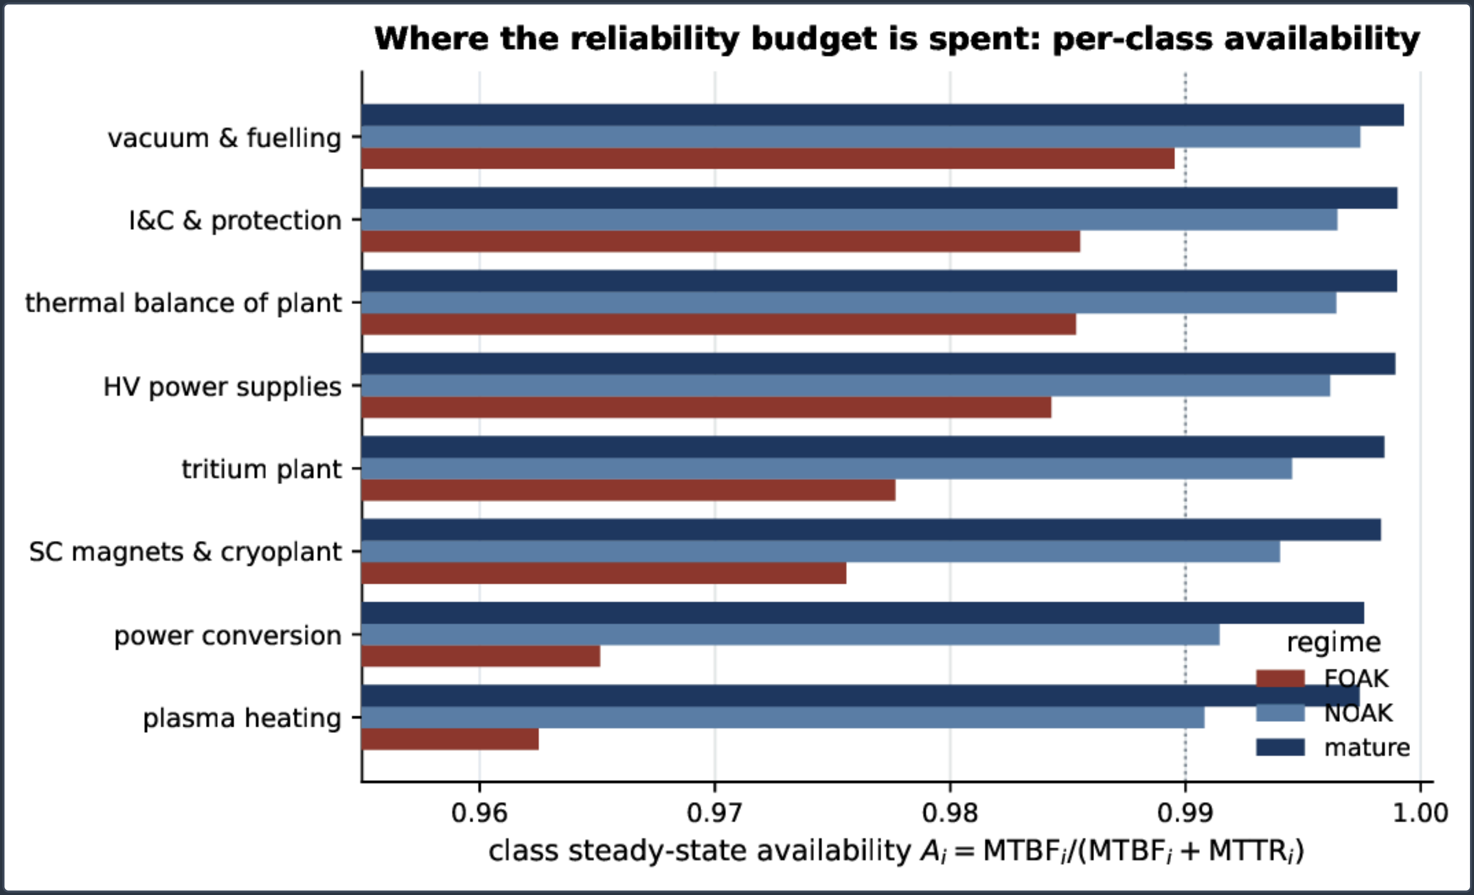

Per-class steady-state availability A_i (1) across the three reliability regimes, from the anchor-pinned solve (10) (register basis AG–L1–A3). Classes are ordered worst-first at first-of-a-kind; plasma heating and power conversion carry the lowest availabilities. The dotted line marks $A_i = 0.99$.

§ 4

Computational methodology

THE NVIDIA GPU HPC CAMPAIGN BEHIND THE PHYSICS ANCHORS

The fleet model is only as good as the physics anchors that parameterise it, and those anchors were produced by the NVIDIA GPU HPC campaign that underlies the Kronos de-risking programme. Three inputs matter here. First, the consumable-centrepost life-limit — the driver of the planned-outage cadence — comes from three-dimensional Monte-Carlo neutron transport (OpenMC with the ENDF/B-VIII.0 library and DAGMC CAD tracking [25]) run on the Lambda A100/H100/A10 GPU fleet, which resolves the fast-neutron fluence at the inboard leg as a function of inboard-shield thickness and hence the coil life (register serial `BR–L1–A13b`; §5.5). Second, the source strength and burn rate that fix the neutron loading of every subsystem come from the reduced-order 0-D power and particle balance of the Kronos toolkit at the frozen operating point ($P_{\text{fus}} = 85.04\,MW$, $Q = 3.076$, $I_p = 9.66\,MA$; serial `BR–L1–A1`). Third, the defense-duty demand profile that the surge-coverage rule answers is set by the mirror-burner solve behind the mission-availability anchor (`AG–L1–A3`). No cost, budget, or compute-spend figures are part of this campaign's reported outputs. The present paper consumes these anchors; it does not re-derive them.

NUMERICAL FORMULATION OF THE FLEET MODEL

The fleet model itself is deliberately lightweight and runs on the in-house HPC node so that it never contends with the GPU training fleet. It has three numerical kernels.

(i) Anchor-pinned scalar solve. The unique root of equation (10) is found by bisection on $c \in [0,\ 1/\max_i s_i]$; 200 iterations converge c to below 10^{-30}, so the reproduced series product matches the anchor to all printed digits. The monotonicity of $\prod_i (1 - cs_i)$ guarantees convergence.

(ii) Event-driven alternating-renewal Monte-Carlo. For each of M independent mission-years and each class i, the simulation draws an initial state (up with the stationary probability A_i), then alternates exponential up-times $\tau^{\text{up}} \sim \text{Exp}(\text{MTBF}_i)$ and lognormal repair-times $\tau^{\text{dn}} \sim \text{LogNormal}(\ln \text{MTTR}_i - \frac{1}{2}\sigma^2, \sigma)$ with shape $\sigma = 0.6$ (so that $\mathbb{E}[\tau^{\text{dn}}] = \text{MTTR}_i$), accumulating the class downtime D_i within $[0, T]$. The realised annual forced availability is taken as the series product over the per-class up-fractions,

$$A_{\text{f}}^{(\text{yr})} = \prod_{i=1}^{8} \left(1 - \frac{D_i}{T}\right),$$

the finite-horizon analogue of equation (2), and the annual capacity factor follows from equation (9). By the alternating-renewal theorem (3) the sample mean of $A_{\text{f}}^{(\text{yr})}$ converges to the analytic A_{f} of equation (2) regardless of the choice of F and G; the lognormal repair law affects the *spread* of the distribution but not its mean, which is why it can be used to match the observed percentile band without perturbing the pinned anchors. We run $M = 20000$ mission-years per scenario. The kernel is embarrassingly parallel across mission-years and maps directly onto GPU-batched pseudo-random streams for larger ensembles; the reported run uses a single seeded NumPy generator [24,21].

(iii) Calendar-weighted coverage. Equation (7) is evaluated exactly with binomial sums for $N = 1, \ldots, 8$ and $k \in \{1, 2\}$, and $N_k^\star$ is read off by equation (8).

VERIFICATION AND REPRODUCIBILITY

The model carries an internal, two-method verification. The analytic series product (2) and the event-driven Monte-Carlo (11) are independent computational paths to the same quantity; they agree to better than 0.5 percentage points in every regime (Table 2), and the Monte-Carlo running mean converges onto the frozen anchor as $M^{-1/2}$ (Figure 6). This reproduces the independent post-freeze hardening result `HX–01`, which re-ran the capacity-factor Monte-Carlo on the in-house node and recovered forced availabilities of $0.833/0.957/0.988$ and the same percentile band. The entire figure set is regenerated from a single seeded script deposited with the paper; the pseudo-random seed is fixed, so the Monte-Carlo outputs are reproducible byte-for-byte.

§ 5

Results

THE AVAILABILITY ENVELOPE

Table 2 and Figure 3 give the single-unit availability envelope across the three regimes. Forced availability rises from 0.838 (first-of-a-kind, 1419 $\mathrm{h\,yr^{-1}}$ of forced outage) through 0.958 (nominal, 368 $\mathrm{h\,yr^{-1}}$) to 0.988 (mature, 105 $\mathrm{h\,yr^{-1}}$); including planned maintenance, total availability is 73.4%, 88.4% and 95.0%, with scheduled outages falling from 45 to 28 to 14 days per year as the centrepost cadence and inspection scope mature. The event-driven Monte-Carlo resolves the capacity-factor distribution around these means (Figure 5): the tenth-to-ninetieth-percentile spread is 68–78%, 86–90% and 94–96%, narrowing sharply with maturity as outages become more frequent but far shorter.

QUANTITY	PESSIMISTIC (FOAK)	NOMINAL (NOAK)	OPTIMISTIC (MATURE)
forced availability A_f (analytic, eq. 2)	0.838	0.958	0.988
forced availability A_f (Monte-Carlo, eq. 11)	0.838	0.958	0.988
total availability A_T (incl. planned, eq. 4)	73.4%	88.4%	95.0%
forced outage ($\mathrm{h\,yr^{-1}}$)	1419	368	105
planned outage ($\mathrm{d\,yr^{-1}}$)	45	28	14
capacity factor P10–P90 (eq. 9)	68–78%	86–90%	94–96%
$N_1^\star$ units for 99% baseload (eq. 8)	4	2	2
$N_2^\star$ units for 99% surge duty	6	4	3

The single-unit availability envelope across three RAMI scenarios (register AG–L1–A3, HX–01). Analytic and Monte-Carlo forced availabilities agree to within 0.5 percentage points.

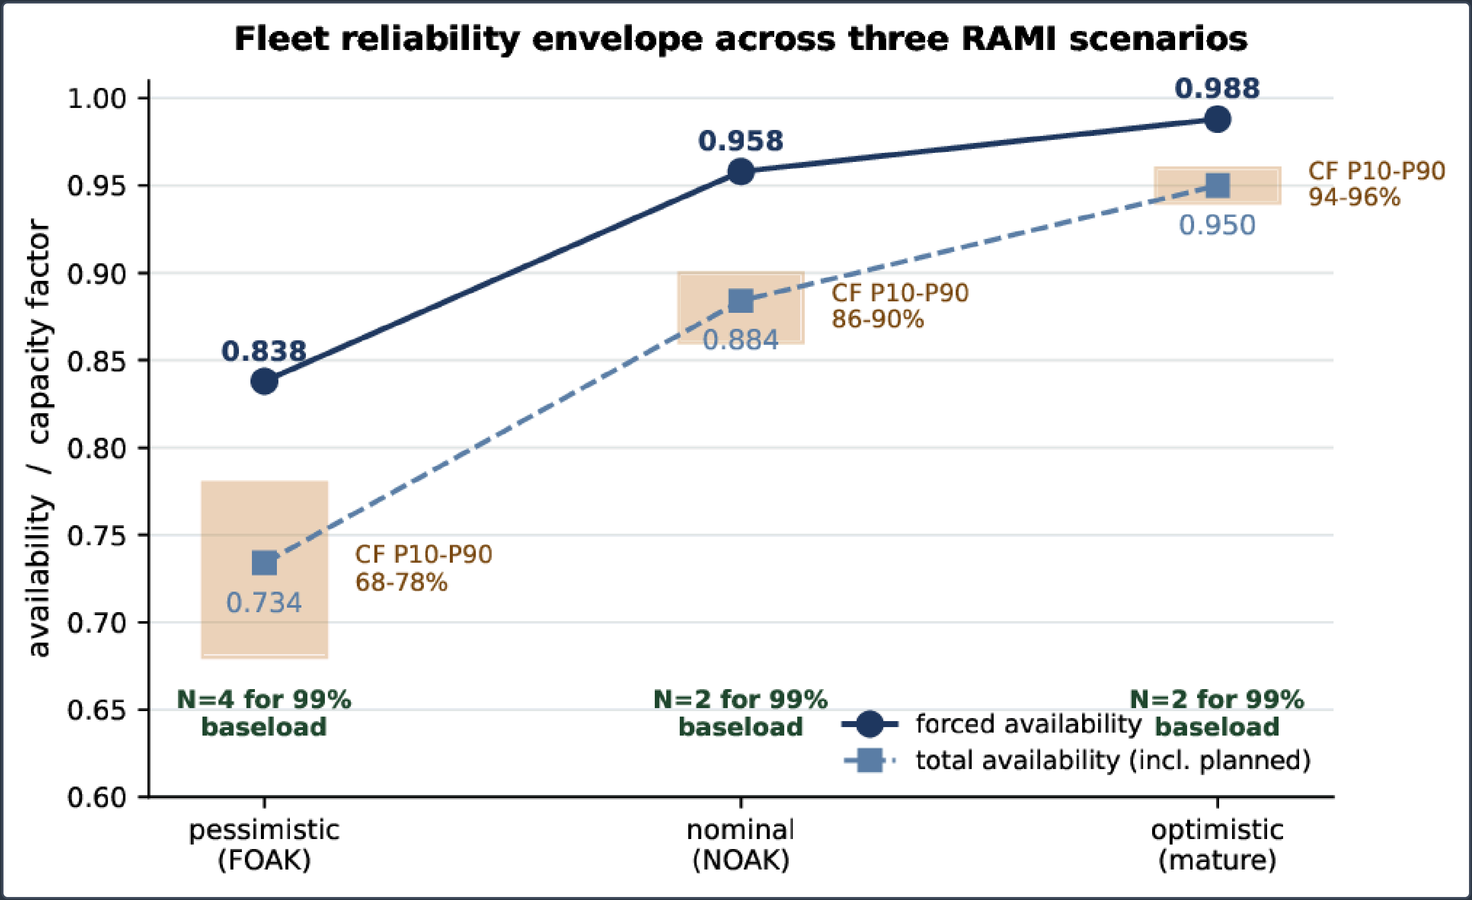

Fleet reliability envelope across three RAMI scenarios. Points show forced and total availability (eqs. 2, 4); the shaded bands show the capacity-factor tenth-to-ninetieth-percentile spread from the event-driven Monte-Carlo (eq. 9); the N-unit annotations give the baseload fleet size (eq. 8).

WHERE THE RELIABILITY BUDGET IS SPENT

Figure 4 ranks the eight classes by their forced-outage contribution $U_i^{(h)} = (1 - A_i) \times 8760$. In every regime plasma heating and power conversion dominate, followed by the tritium plant and the magnet/cryo class; the ranking is stable across the scenario bracket, which is the property that makes it actionable. Module-level redundancy in heating and power conversion — the two classes with the highest failure frequency — therefore buys the most coverage per unit of added complexity, and is the first place a reliability-growth programme should invest. The magnet class, by contrast, is addressed not by redundancy but by making its rare interventions faster and by the fused quench-precursor protection that averts the catastrophic-loss failure mode (the register records a fused precursor giving $47.2\,ms$ of warning against $2.6\,ms$ for terminal-voltage detection at an area-under-ROC of 0.9998 [33]).

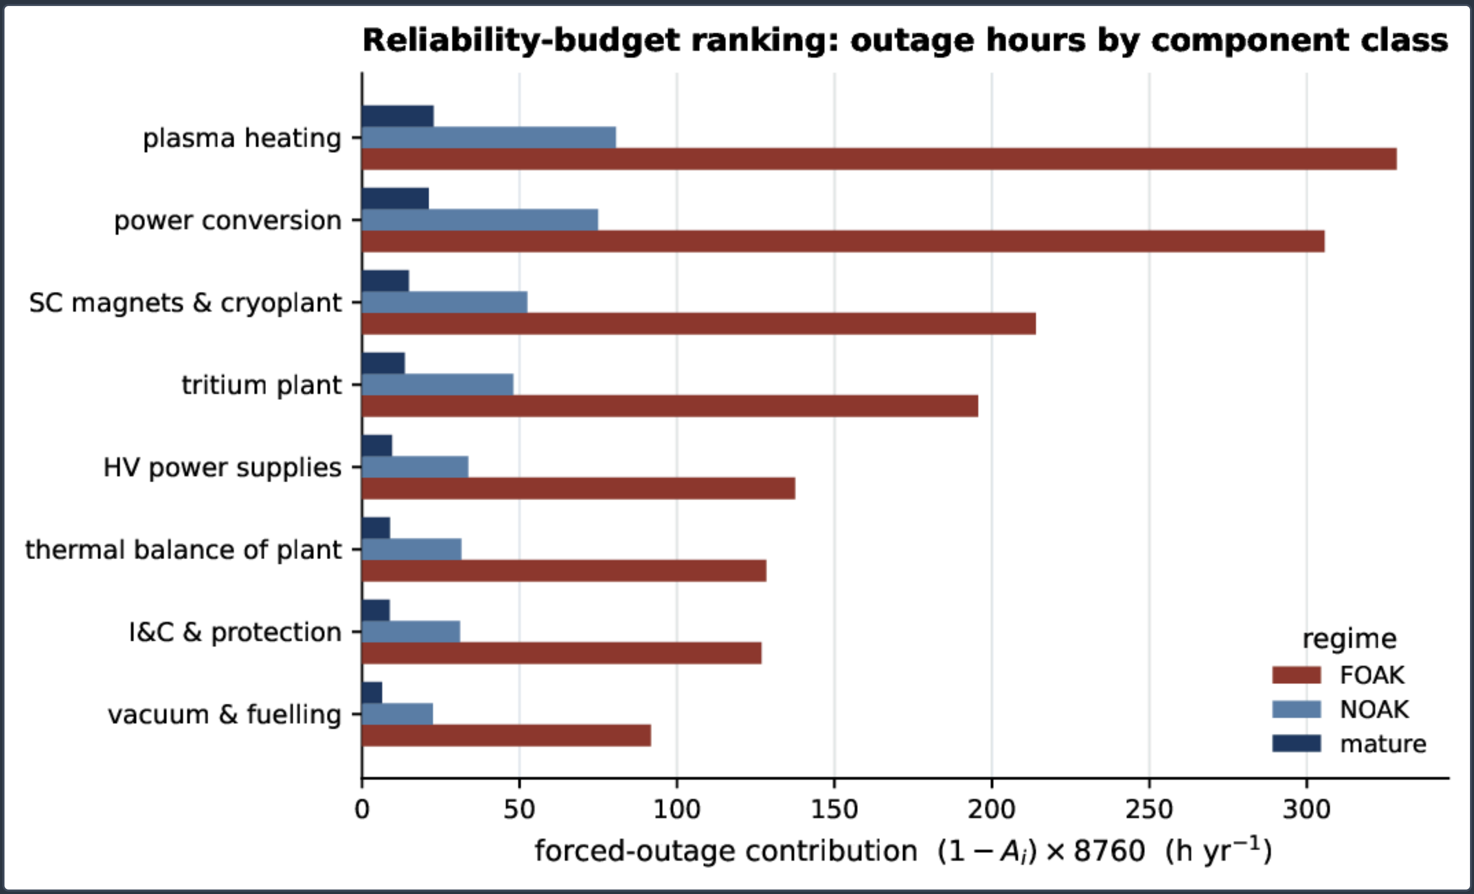

Reliability-budget ranking: forced-outage hours per year $(1 - A_i) \times 8760$ by component class, in each regime. Plasma heating and power conversion dominate the budget in every scenario; the ranking is stable across the bracket.

THE CAPACITY-FACTOR DISTRIBUTION AND MONTE-CARLO CONVERGENCE

Figure 5 shows the annual-capacity-factor distributions from the 20000-mission-year event-driven Monte-Carlo. The distributions are centred on the total availability of equation (4) and their P10–P90 widths reproduce the frozen band: the first-of-a-kind regime carries a broad ~ 10-point spread (a few long outages per year dominate the annual balance), collapsing to a ~ 2-point spread at maturity. Figure 6 verifies that the Monte-Carlo running mean converges onto the analytic anchors of equation (2) with the expected $M^{-1/2}$ envelope, confirming that the two computational paths are consistent.

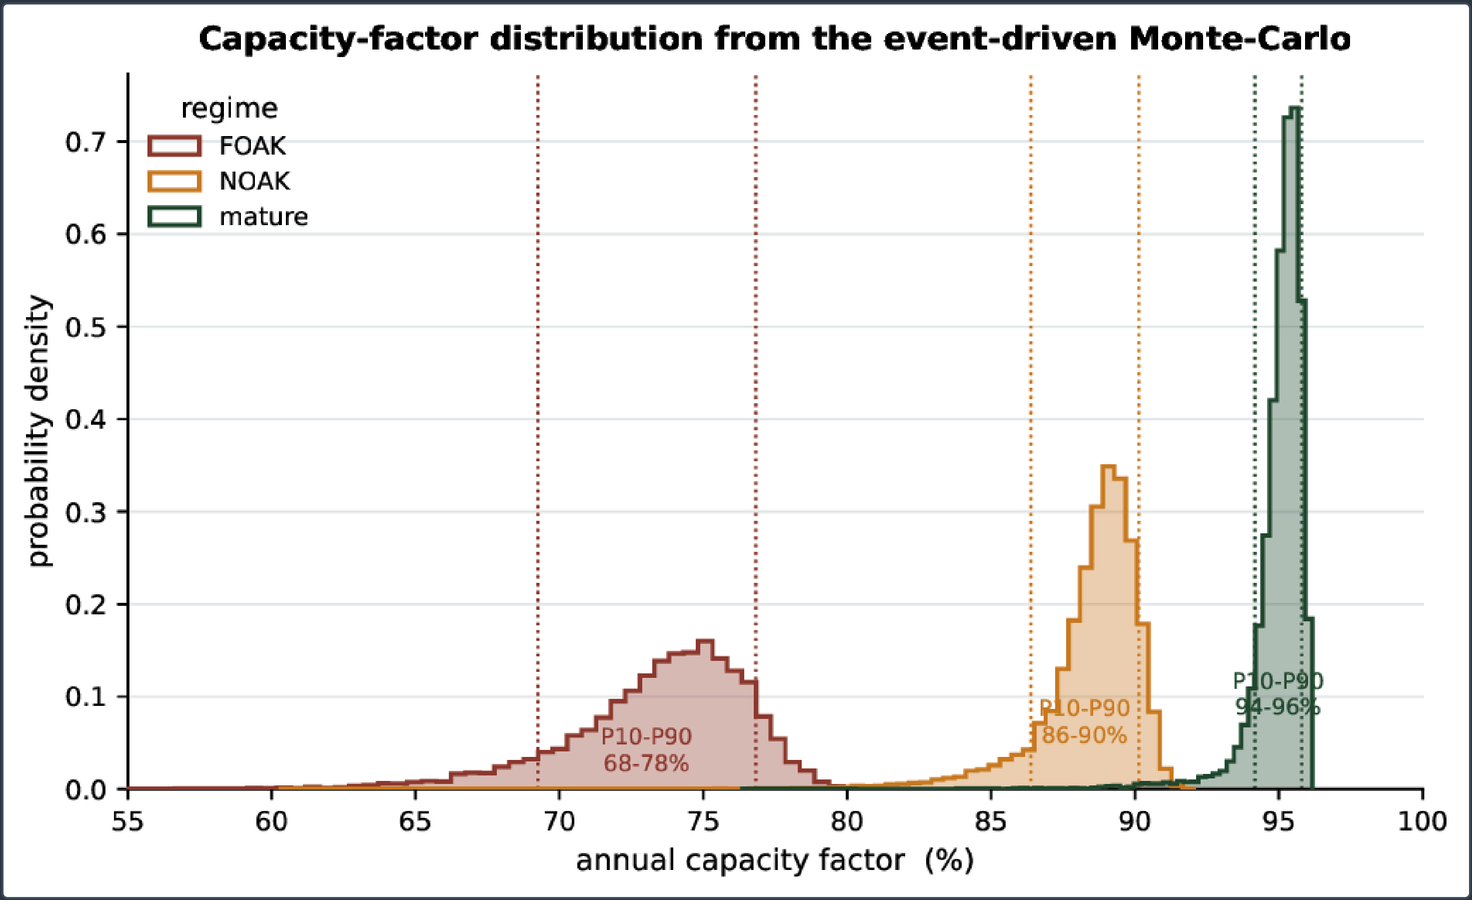

Annual capacity-factor distributions from the event-driven alternating-renewal Monte-Carlo (eqs. 11, 9; $M = 20000$ mission-years per regime; register HX–01). Dotted lines mark the tenth and ninetieth percentiles; the frozen bands are 68–78/86–90/94–96%.

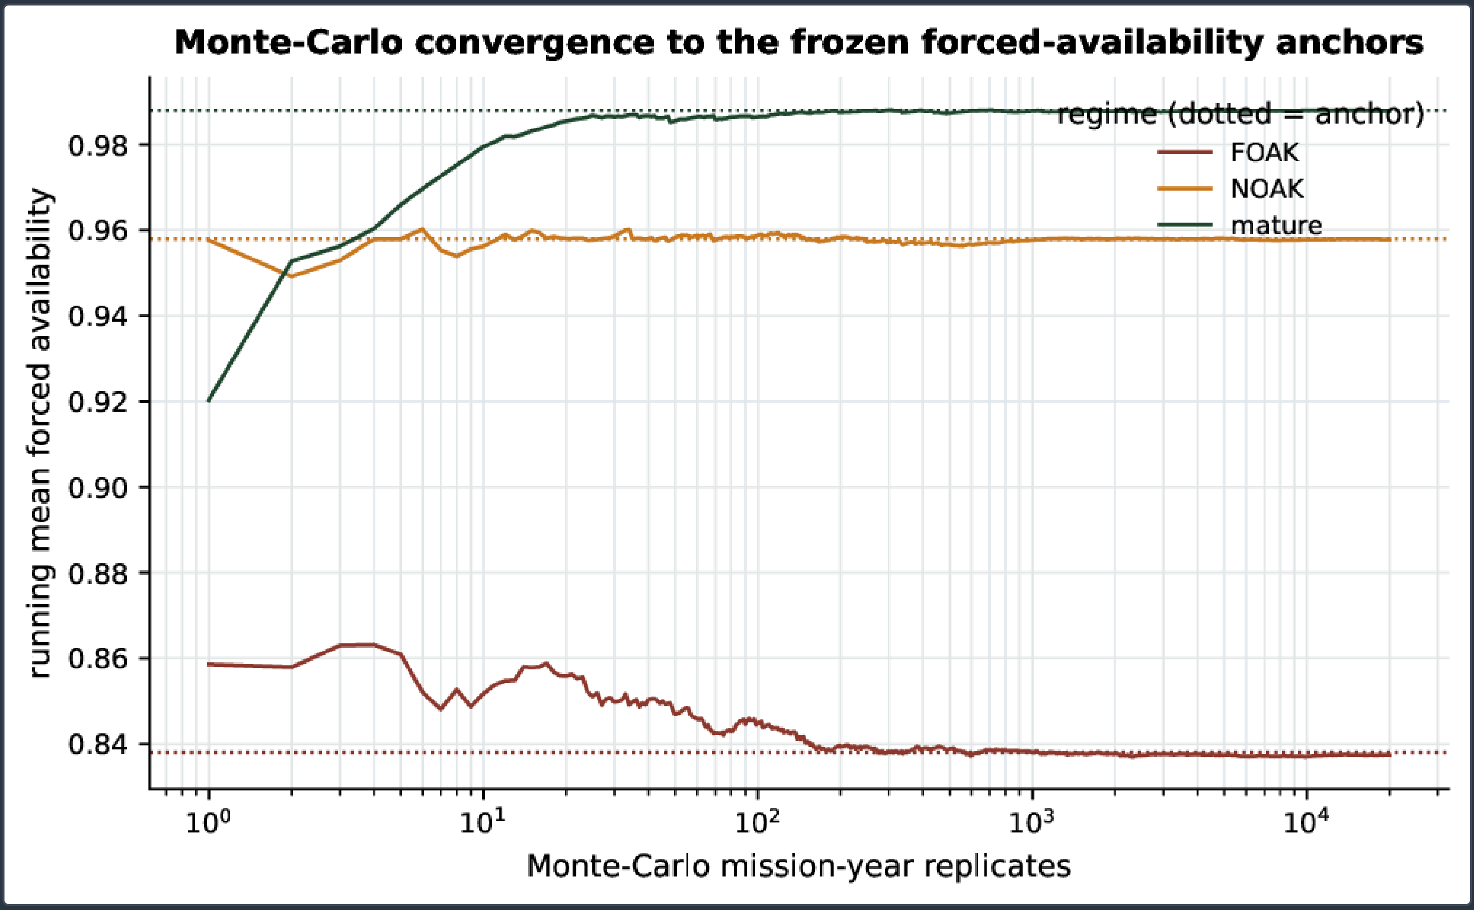

Monte-Carlo convergence of the running-mean forced availability toward the frozen anchors (dotted) of equation (2), against the number of mission-year replicates. Convergence follows the standard $M^{-1/2}$ law of Monte-Carlo estimators [21].

FLEET SIZING: BASELOAD AND SURGE DUTY

Applying the calendar-weighted coverage functional (7) with $C^\star = 0.99$ gives the central operational result. For baseload ($k = 1$; Figure 7), a *four-unit* fleet delivers 99% coverage at first-of-a-kind reliability — the exact coverage is $C_1(4) = 0.9976$ — while a two-unit fleet suffices at nominal or mature reliability. The four-unit result is the honest first-of-a-kind sizing: it is why early deployments are framed as multi-unit sites rather than sole-source single machines, and it contracts as component reliability matures. For the more demanding surge duty ($k = 2$ concurrent units; Figure 7), the same functional sizes to six, four and three units across the three regimes, bounding the most stringent mission profile. Table 3 tabulates the coverage and the sized fleet.

QUANTITY	FOAK	NOAK	MATURE
forced availability A_f	0.838	0.958	0.988
staggered-outage weight w at $N_1^\star$	0.493	0.153	0.077
baseload fleet $N_1^\star$ ($k = 1$)	4	2	2
achieved baseload coverage $C_1(N_1^\star)$	0.9976	0.9921	0.9989
surge fleet $N_2^\star$ ($k = 2$)	6	4	3

Calendar-weighted k-of-N coverage (7) and the sized fleet (8) at the 99% target. $C_1(N^\star)$ is the achieved baseload coverage at the sized fleet.

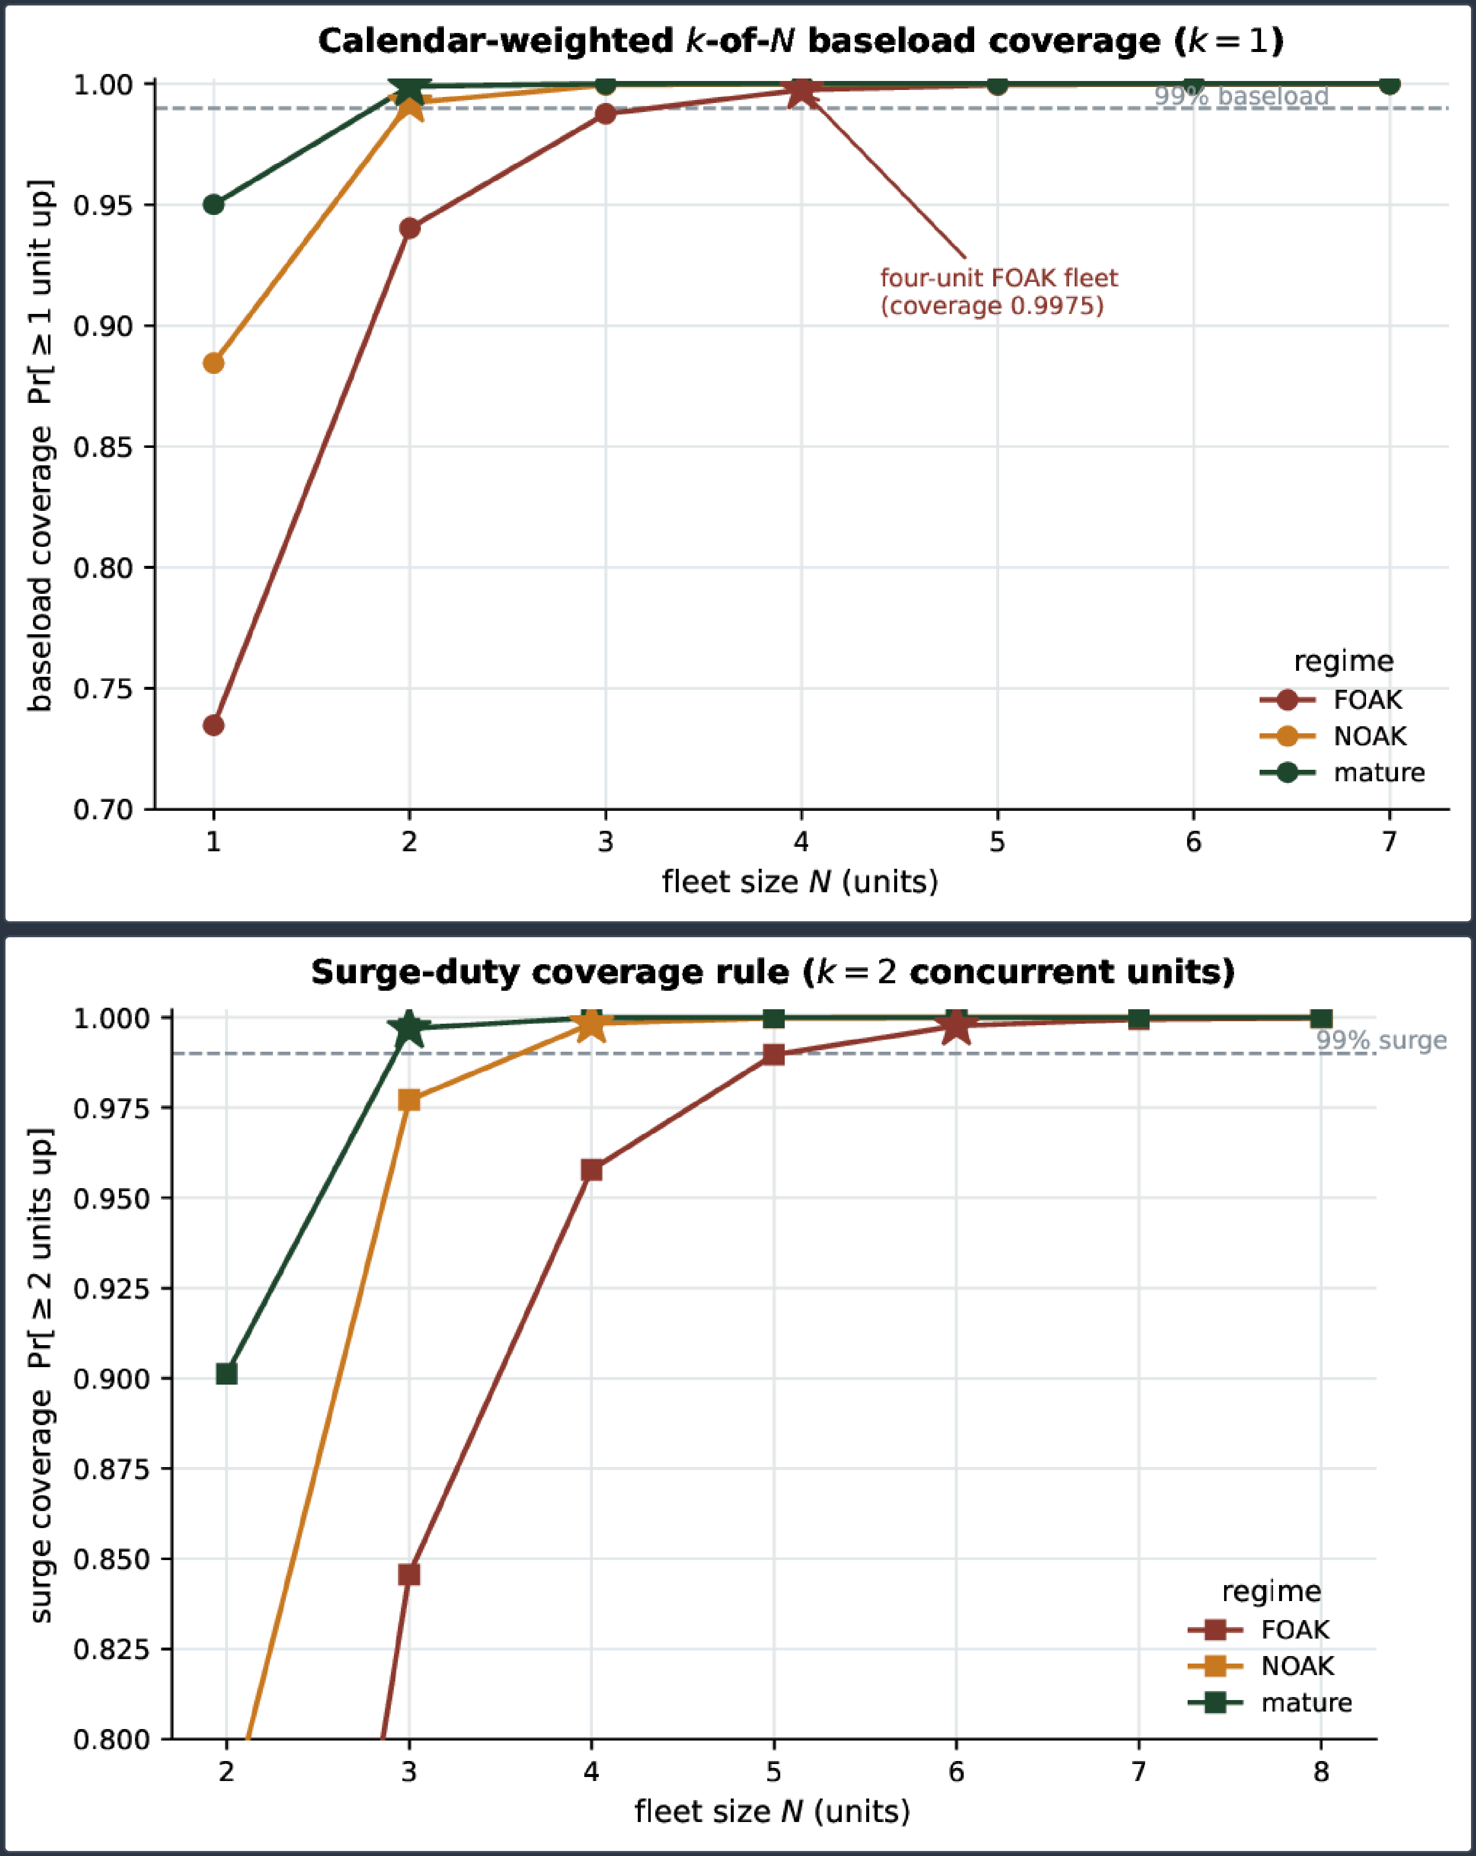

Baseload coverage $C_1(N)$ (7) versus fleet size, per regime. Stars mark the sized fleet $N_1^\star$; the four-unit first-of-a-kind fleet clears the 99% target at coverage 0.9976.

Figure 8 decomposes the four-unit first-of-a-kind coverage into the two terms of equation (7): the w-weighted staggered-maintenance term (three units carrying the load while the fourth is swapped) and the complementary all-four-available term. The decomposition makes explicit why staggering matters: even at first-of-a-kind reliability, with a unit in a planned outage nearly half the year ($w = 0.493$), the three remaining units cover the baseload mission 99.6% of that time, and the fleet clears 99% overall.

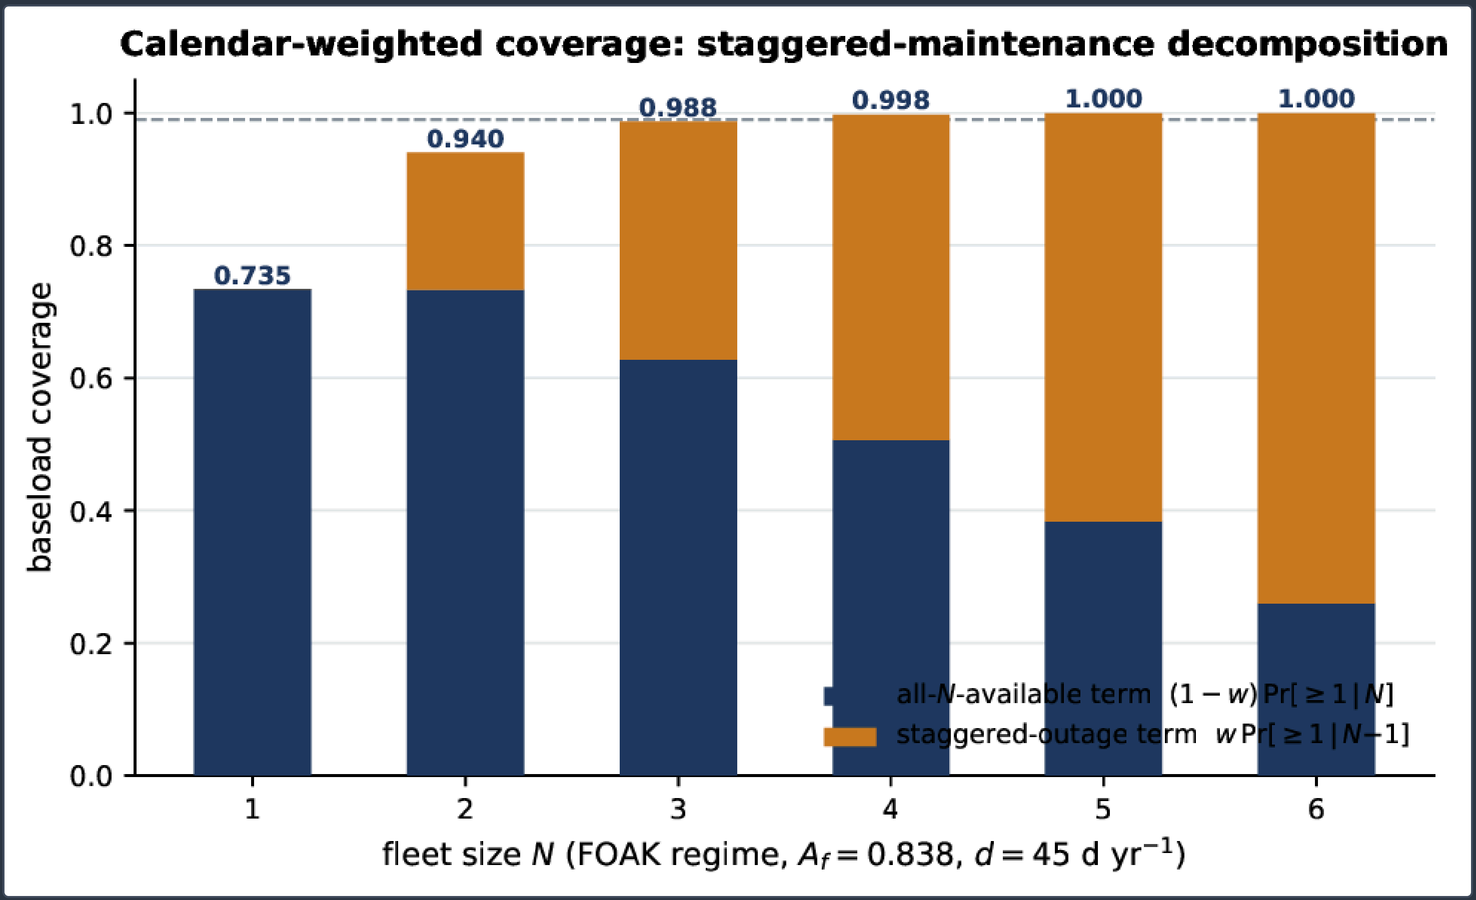

Calendar-weighted decomposition of first-of-a-kind baseload coverage (7): the all-N-available term $(1-w)\,\Pr[\geq 1\,|\,N]$ and the staggered-outage term $w\,\Pr[\geq 1\,|\,N-1]$, versus fleet size. Numbers are the total coverage; the four-unit fleet reaches 0.998.

THE CENTREPOST SWAP PACES THE PLANNED OUTAGE

For the breeder fleet the planned-outage cadence is set by the consumable centrepost. The centrepost coil life is a steep function of the inboard-shield thickness because the shield attenuates the fast-neutron fluence that damages the inboard conductor: the GPU-HPC transport campaign (BR–L1–A13b) finds the fast fluence falling from 2.9×10^{25} to 1.7×10^{23} $\mathrm{n\,m^{-2}\,fpy^{-1}}$ as the inboard shield grows from 0 to 45 *cm*, and the corresponding coil life rising from 0.003 to 0.58 full-power years, with the register recording 0.244 fpy at a ~ 31 *cm* tungsten shield (Figure 9). The design cartridge life of ≈ 0.58 fpy therefore sets the swap interval, and the swap itself — a two-ended demountable-cartridge operation [31] — is the largest single block of planned downtime. This is the physical origin of the planned-outage days in Table 2 and of the fleet coupling recorded in the register, where an $N+1$ fleet with a one-day swap clears successively higher availability tiers as the coil life lengthens. Because the swap is planned and predictable, the staggered-maintenance schedule (Figure 10) spreads it across the fleet so that no two units are ever swapping at once — the operational premise of the coverage weight w in equation (6). The trade between source strength (hence breeding and isotope output) and centrepost life that fixes the admissible operating point is developed in the companion operating-window and life-limiter studies [34,33,32].

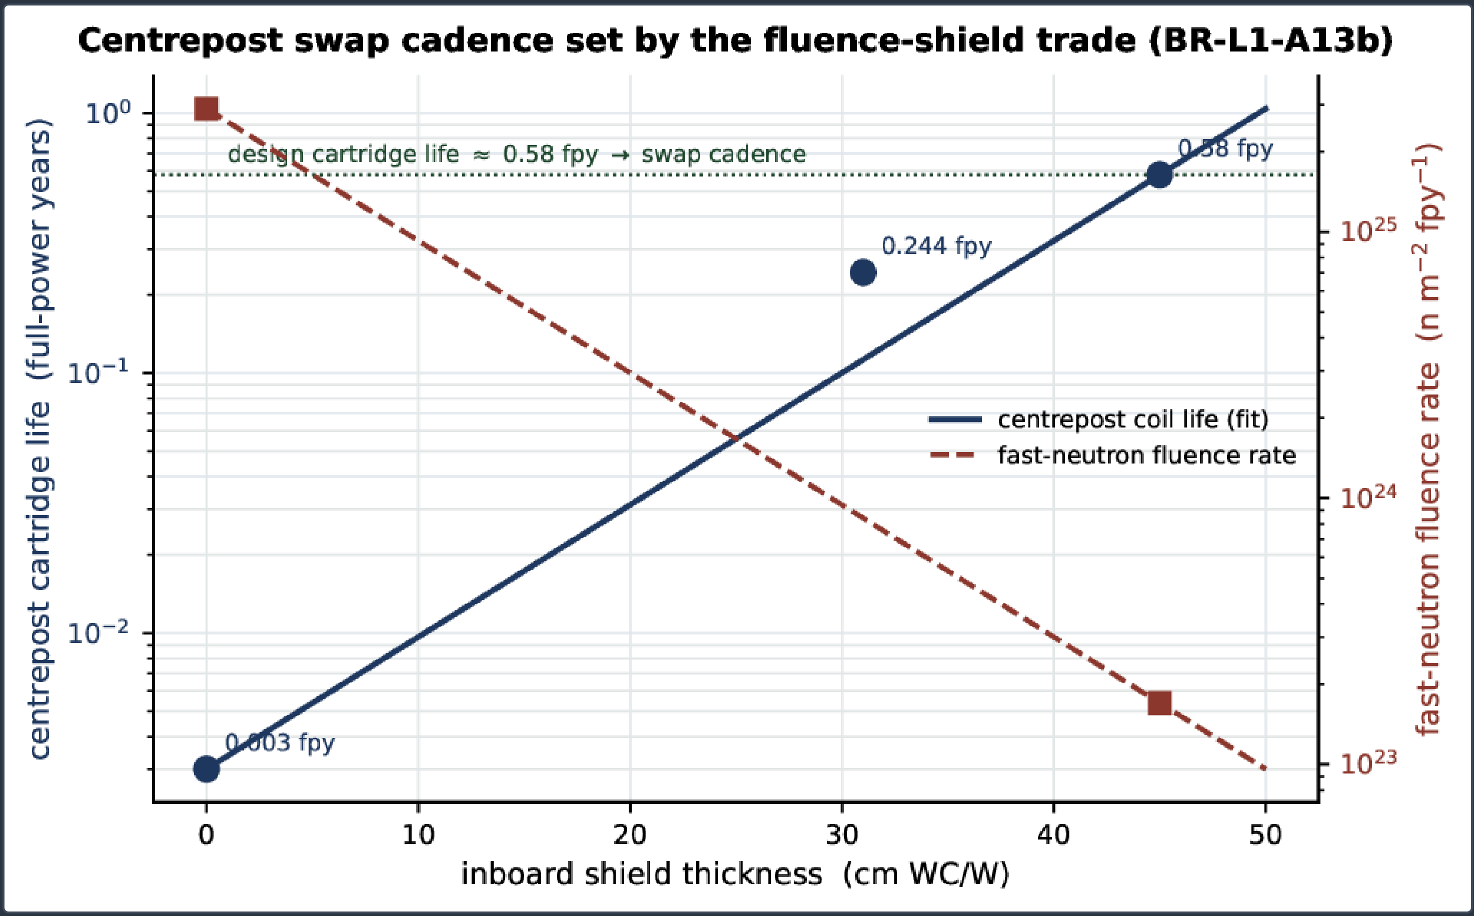

Centrepost cartridge life (left axis, log) and fast-neutron fluence rate (right axis, log) versus inboard-shield thickness, from the GPU-HPC transport campaign (BR–L1–A13b) with the register 0.244 fpy point at $\sim 31\ cm$. The design cartridge life ≈ 0.58 fpy sets the swap cadence that paces the planned outage.

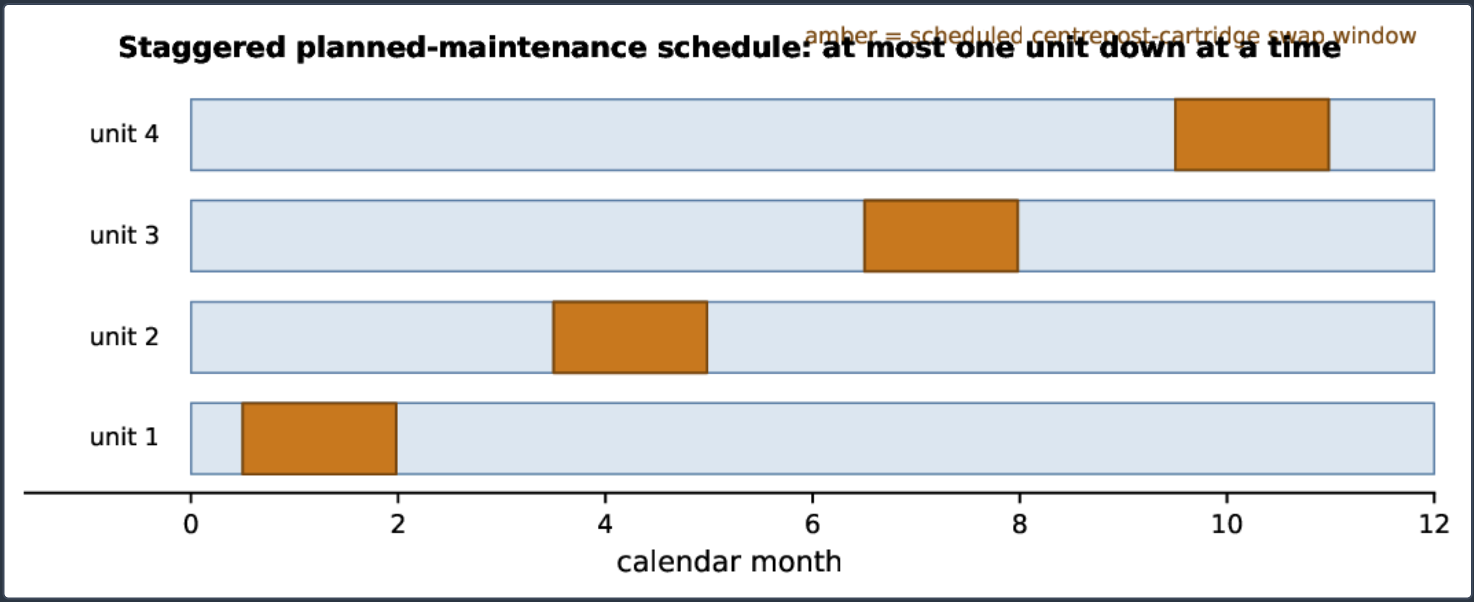

(Schematic.) Staggered planned-maintenance schedule for a four-unit fleet: each unit's scheduled centrepost-swap window (amber) is offset so that at most one unit is down for maintenance at any instant, realising the calendar weight w of equation (6). Dates are illustrative.

§ 6

Discussion

AGAINST THE PUBLISHED FUSION-AVAILABILITY BENCHMARKS

The envelope is consistent with, and more honest than, the availability figures used in the fusion power-plant literature. Conceptual power-plant studies and systems-code analyses typically *assume* a plant availability of $\sim 75–90\%$ as an input to the economic case [3,8,6], with the European DEMO programme treating availability as a central design driver rather than a derived quantity [4]. Our nominal total availability of 88.4% sits squarely inside that assumed band, but — unlike an assumed input — it is *derived* from a component-class basis and bracketed by an explicit first-of-a-kind-to-mature range, with the first-of-a-kind figure (73.4%) deliberately pessimistic. The mature capacity factor of $94–96\%$ approaches the $\sim 90–93\%$ lifetime capacity factors achieved by the best operating fission fleets, which is the correct target for a firm-power source and a more demanding benchmark than the conceptual-study assumptions. Relative to the ITER RAMI methodology [9], which propagates component failure rates into subsystem availability, our contribution is the fleet layer: the calendar-weighted k-of-N functional (7) that converts single-unit availability into a sized fleet under staggered maintenance, and the time-domain Monte-Carlo that resolves the capacity-factor *distribution* rather than its mean.

THE FOUR-UNIT FLEET AND THE TWO DISTINCT MAGNETS

The four-unit first-of-a-kind sizing is the load-bearing operational conclusion, and it is a direct consequence of the low first-of-a-kind forced availability (0.838) compounded by a centrepost swap that removes a unit from service nearly half the year in that regime ($w = 0.493$). It is worth being explicit that the two Kronos products carry two *distinct* magnets, and the fleet model treats them as such: the breeder centrepost governs at a peak field of $16.84\,T$ and is the consumable that paces this fleet's planned outage, whereas the separate mirror-burner product uses a distinct, higher-field plug coil at $26.49\,T$ under a different duty cycle. The machine-agnostic structure of the coverage functional is what lets the same model size a breeder baseload fleet by its centrepost-swap cadence and a defense-duty burner fleet by its surge ($k = 2$) requirement, from one component-class basis; the burner mission ($Q_E = 1.318$, neutron fraction $f_n = 5.44\%$, direct-conversion efficiency $\eta_{\text{DEC}} = 0.49$) is the origin of the surge profile in the mission-availability anchor.

SENSITIVITY AND ROBUSTNESS

Two features make the conclusions robust across the scenario bracket. First, the reliability-budget ranking (Figure 4) is stable: heating and power conversion dominate in all three regimes, so the reliability- growth investment priority does not depend on which regime the fleet is actually in. Second, the fleet-sizing rule is a step function of a coverage *integral*, so it is insensitive to the exact percentile shape of the capacity- factor distribution — only the mean forced availability and the planned-outage days enter equation (7). The scenario bracket, not a single point estimate, is the deliverable, and the sizing rule is what survives it.

§ 7

Limitations

One honest gate bounds the claims. Every MTBF_i and MTTR_i in Table 1 is a representative fusion-plant RAMI value drawn from the reactor-study and failure-rate-database literature [10,11,9], not yet a Kronos-measured datum — by definition, since no Kronos hardware has yet accumulated operating hours. The envelope is therefore a scenario *bracket*, and the model's load-bearing outputs are the reproducible sizing rule and the stable sensitivity ranking, both of which are robust across that bracket. A second, related caveat is that the coverage functional assumes statistically independent units: common-cause coupling across co-sited units — a shared cryoplant, a shared tritium plant, a shared grid connection — is unmodelled, and any such coupling would raise the sized fleet. Quantifying common-cause factors from the first multi-unit operating data, and replacing the literature- class $(\mathrm{MTBF}, \mathrm{MTTR})$ priors with measured values, is the follow-on that turns this bracket into a specification.

§ 8

Conclusion

A time-domain, event-driven Monte-Carlo, anchored to a component-class series-system model and closed by a calendar-weighted k-of-N coverage functional, resolves the fleet-availability case for a compact spherical-tokamak breeder. Single-unit forced availability is 0.838, 0.958 and 0.988 across first-of-a-kind to mature regimes (total availability 73.4%, 88.4%, 95.0%); the capacity-factor distribution spans 68–78% to 94–96% at the tenth-to-ninetieth percentile; and the coverage functional sizes the fleet to *four* units for 99% baseload coverage at first-of-a-kind reliability (coverage 0.9976), tightening to two at maturity, with a surge duty sizing to six/four/three. Plasma heating and power conversion set the reliability floor, and the breeder fleet's planned cadence is paced by the scheduled centrepost-cartridge swap, whose coil-life-versus-shield law fixes the swap interval. Every number reproduces a frozen anchor, the analytic and Monte-Carlo paths agree to better than half a percentage point, and the whole result is regenerable from a single seeded model — an availability and fleet-sizing basis built on reproducibility rather than assertion.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and the NVIDIA GPU HPC campaign for the neutron-transport and reduced-order anchors that parameterise the fleet model.

Reproducibility. The availability envelope, the capacity-factor spread, and the N-unit coverage results are frozen anchors — the mission-availability model AG–L1–A3 and the post-freeze capacity-factor Monte-Carlo HX–01 — in the Kronos Physics De-Risking Register [30] and are regenerable from the seeded RAMI model archived with this paper. This paper is archived at its DOI doi:10.5281/zenodo.22645951 and its official page https://www.kronosfusionenergy.com/publications/fleet-availability-under-consumable-centrepost-swap-cad.html; the register is archived at doi:10.5281/zenodo.22133057.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645951. The component-class reliability table, the anchor-pinned scalar solve, the event-driven Monte-Carlo, and the calendar-weighted coverage evaluation are archived with the deposit and reproduce every figure and table herein from a single seeded script; all quantitative values trace to the register [30] 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 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 F Najmabadi *et al* (ARIES Team), The ARIES-AT advanced tokamak, advanced technology fusion power plant, *Fusion Eng. Des.* **80** 3 (2006).

4 G Federici *et al*, Overview of the DEMO staged design approach in Europe, *Nucl. Fusion* **59** 066013 (2019).

5 T Donn\'e *et al*, *European Research Roadmap to the Realisation of Fusion Energy*, EUROfusion, Garching (2018).

6 M Kovari *et al*, “PROCESS”: a systems code for fusion power plants — Part 1: Physics, *Fusion Eng. Des.* **89** 3054 (2014).

7 M Kovari *et al*, “PROCESS”: a systems code for fusion power plants — Part 2: Engineering, *Fusion Eng. Des.* **104** 9 (2016).

8 D J Ward, I Cook, Y Lechon and R Saez, The economic viability of fusion power, *Fusion Eng. Des.* **75–79** 1221 (2005).

9 D Van Houtte, K Okayama and F Sagot, RAMI approach for ITER, *Fusion Eng. Des.* **85** 1220 (2010).

10 T Pinna, S Ciattaglia, R Boullon *et al*, Component failure-rate database and reliability analyses for fusion machines, *Fusion Eng. Des.* **136** 87 (2018).

11 L C Cadwallader, *Selected Component Failure Rate Values from Fusion Safety Assessment Tasks*, Report INL/EXT-98-00892, Idaho National Laboratory (2007).

12 O Crofts *et al*, Overview of progress on the European DEMO remote maintenance strategy, *Fusion Eng. Des.* **109–111** 1392 (2016).

13 R Buckingham and A Loving, Remote-handling challenges in fusion research and beyond, *Nat. Phys.* **12** 391 (2016).

14 R E Barlow and F Proschan, *Mathematical Theory of Reliability*, SIAM Classics in Applied Mathematics **17**, Philadelphia (1996).

15 S M Ross, *Stochastic Processes*, 2nd ed., Wiley, New York (1996).

16 W L Smith, Renewal theory and its ramifications, *J. R. Stat. Soc. B* **20** 243 (1958).

17 M Rausand and A Hyland, *System Reliability Theory: Models, Statistical Methods, and Applications*, 2nd ed., Wiley, Hoboken NJ (2004).

18 A Birolini, *Reliability Engineering: Theory and Practice*, 8th ed., Springer, Berlin (2017).

19 K S Trivedi, *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*, 2nd ed., Wiley, New York (2001).

20 W Kuo and M J Zuo, *Optimal Reliability Modeling: Principles and Applications*, Wiley, Hoboken NJ (2003).

21 E Zio, *The Monte Carlo Simulation Method for System Reliability and Risk Analysis*, Springer Series in Reliability Engineering, London (2013).

22 N Metropolis and S Ulam, The Monte Carlo method, *J. Am. Stat. Assoc.* **44** 335 (1949).

23 W Weibull, A statistical distribution function of wide applicability, *J. Appl. Mech.* **18** 293 (1951).

24 C R Harris *et al*, Array programming with NumPy, *Nature* **585** 357 (2020).

25 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

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

27 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

28 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

29 J L MacManus-Driscoll and S C Wimbush, Processing and application of high-temperature superconducting coated conductors, *Nat. Rev. Mater.* **6** 587 (2021). doi:10.1038/s41578-021-00290-3

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

31 P I Ford, *Operating the Consumable Centrepost of a Compact Spherical-Tokamak Breeder: Demountable-Cartridge Swap Logistics, Fleet RAMI, and Centrepost Radial-Budget Reconciliation*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645931; https://www.kronosfusionenergy.com/publications/operating-the-consumable-centrepost-of-a-compact-spheri.html

32 P I Ford, *Prognostics and Remaining-Useful-Life Estimation for the Consumable Centrepost: A Dual-Channel Fluence-Driven Framework Scheduling Cartridge Replacement in a Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645788; https://www.kronosfusionenergy.com/publications/centrepost-prognostics-and-rul-control.html

33 P I Ford, *The Centrepost Life-Limiter: Fast-Fluence, dpa, and Shielding Setting the Cartridge-Replacement Cadence for a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645708; https://www.kronosfusionenergy.com/publications/centrepost-fluence-and-shielding.html

34 P I Ford and G L Kulcinski, *Lifetime-Constrained Operating-Window Synthesis for a Compact Spherical-Tokamak Breeder: The Source-Strength-Centrepost-Life Trade and the Admissible Current Window*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645799; https://www.kronosfusionenergy.com/publications/admissible-operating-window-synthesis.html

35 P I Ford, *From Tritium Consumer to Breeding Fleet: Net Blanket TBR and the Cross-Library Case for Multi-Unit Tritium Export from a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645706; https://www.kronosfusionenergy.com/publications/blanket-tritium-breeding-lever.html

36 P I Ford, *Critical-Materials Buildability for a Spherical-Tokamak Breeder Fleet: Bottleneck Ranking of Li-6 Enrichment, Beryllium Fabrication and REBCO Tape Supply*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645786; https://www.kronosfusionenergy.com/publications/supply-chain-and-fleet-buildability.html

COLOPHON

How this document was made

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

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

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

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

LEGAL NOTICE

Notice, status, and distribution

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

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

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

THE 2026 KRONOS SERIES

One programme, many papers

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

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

KRONOS FUSION ENERGY

Fleet Availability under Consumable-Centrepost Swap Cadence: A Time-Domain Monte-Carlo of Capacity Factor and RAMI across a Four-Unit Spherical-Tokamak Breeder Fleet

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

Zenodo · doi:10.5281/zenodo.22645951

Read on kronosfusionenergy.com

The Physics De-Risking Register