KRONOS FUSION ENERGY

PAPER 2.9 · THE KRONOS 2026 SERIES

Disruptions, Runaway-Electron Avalanche and Electromagnetic Loads in a Compact Spherical-Tokamak Breeder: The Design Case for Shattered-Pellet Mitigation

The disruption response is the credibility gap a compact, high-current spherical-tokamak (ST) programme is expected to answer, because such a machine stores a large magnetic and...

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

BY P I FORD AND G L KULCINSKI

2026 · EDITORIAL EDITION

THE OFFICIAL RECORD

What this document is, and how to read its numbers

This is a combined editorial. It gathers, in one volume, the two design studies Kronos Fusion Energy released in 2026 — a compact spherical-tokamak tritium/helium-3 breeder and a deuterium–helium-3 tandem-mirror generator — together with the intellectual-property estate, the people, and, in the internal edition, the full commercial case. The scientific text is the same text that appears in the arXiv preprints and the journal submissions; the editorial adds the apparatus a reader needs to hold the whole program at once, and nothing in the physics is altered to fit it.

THE HONESTY STANCE IS THE ASSET

Every headline quantity carries an **evidence class** — DERIVED RECOMPUTED MEASURED REQUIREMENT — and, where a number is a modelled extrapolation rather than a demonstrated value, it is labelled as such and its downside is shown next to it. Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly. This is deliberate: a diligence team will find the load-bearing assumptions, so the document surfaces them first.

TWO EDITIONS

This volume is issued in two editions. The **public edition** (light) carries the physics and the design only — no dollar, price, or per-kilogram figure appears anywhere in it. The **internal edition** (dark) adds the economics, the funding architecture, and the defense-commercial analysis, and is marked *Confidential* accordingly; it is not for public distribution. Where the commercial case rests on a modelled assumption rather than a demonstrated one, it is labelled as such and its downside is shown alongside it, on the same terms as the physics.

TITLE	The Kronos Fleet — Hyperion · Aegis · MetroVolt: A Combined Design & Commercialization Study
AUTHOR	Priyanca Ford, Founder & CEO, on behalf of Kronos Fusion Energy
EDITION	2026 Combined Editorial · first issue
COMPANION WORKS	The five-paper arXiv drop — (1) Breeder, (2) Burner, (3) REBCO magnet & tape, (4) Direct Energy Conversion, (5) AI/ML/Quantum control — with matching numbered Zenodo deposits, plus the IOP journal and IAEA-FEC submissions. Patents were filed before the drop.
NAMING	Hyperion = the breeder; Aegis (defense) and MetroVolt (data-center) = one generator in two housings. Internal mode letters (D1, L) do not appear on public surfaces.
STATUS	Conceptual design & simulation study. Nothing herein is a construction commitment.

HOW TO READ THIS VOLUME

The fleet at a glance, and the way its numbers are marked

This volume opens with *Orientation*, presents the five research papers as a bound sequence, and closes with the *Apparatus*. *Foundations* sets out the three general results both machines rest on. *Paper One* is the Hyperion breeder; *Paper Two* the Aegis / MetroVolt generator; *Papers Three, Four and Five* are the enabling science — high-field REBCO magnets and tape, direct energy conversion, and the AI / ML / quantum control stack. A *Digital Twin & 3-D Model* part follows, then the *Environmental* profile. The *Economics* and levelized cost and the *Business Case* and diligence record appear in the internal edition only; the *Simulations* proof record and the Apparatus — patent portfolio, team, limitations, and sources — close both editions.

THE FLEET AT A GLANCE

	HYPERION	AEGIS	METROVOLT
Role	Strategic-isotope foundry	Defense installation power	Campus power
Machine	**Spherical tokamak**	**Tandem mirror**	**Tandem mirror**
Fuel	Deuterium–tritium	Deuterium–helium-3	Deuterium–helium-3
Delivers	Tritium · ^{3}He · 14 MeV n	Resilient installation power	Firm campus power
Physics bar	**Fusion gain only**	Closure (gates named)	Closure + lunar ^{3}He
In the fleet	Breeds the fuel	Proves the generator	Commercial destination

HOW THE NUMBERS ARE MARKED

Every headline quantity carries an evidence class — DERIVED from first principles, RECOMPUTED from raw data, MEASURED in hardware, or a REQUIREMENT yet to be met. Financial figures carry a case label and are confined to the internal (confidential) edition. Values that moved after first publication are shown as withdrawn or restated rather than quietly changed. A capability you can evaluate is one whose limits are on the page.

ABSTRACT

Disruptions, Runaway-Electron Avalanche and Electromagnetic Loads in a Compact Spherical-Tokamak Breeder: The Design Case for Shattered-Pellet Mitigation

The disruption response is the credibility gap a compact, high-current spherical-tokamak (ST) programme is expected to answer, because such a machine stores a large magnetic and thermal energy in a small volume and releases it fast. We formalise the disruption, runaway-electron (RE), and electromagnetic-load problem for the frozen Hyperion breeder design point ($I_p = 9.66$ MA, $R_0 = 1.20$ m, aspect ratio $A = 2.5$, elongation $\kappa = 2.0$, triangularity $\delta = -0.30$, on-axis field $B_0 = 8$ T, centrepost peak field 16.84 T) and solve it with a hierarchy of reduced-order models anchored on a NIMROD-verified free-boundary equilibrium. We derive the inductive current-quench (CQ) circuit model, the Connor–Hastie and Aleynikov–Breizman runaway threshold fields, the Rosenbluth–Putvinski close-collision avalanche, and the halo-current force balance as explicit equations, then evaluate them across the plasma-resistivity and quench-rate parameter space on the NVIDIA GPU HPC campaign that also ran the reference extended-MHD (NIMROD) and fluid-kinetic runaway (DREAM) codes. The result is honest and uncomfortable. The CQ time is $\tau_{CQ} = 7\text{–}19$ ms, set by the plasma inductance-over-resistance ratio ($L_p \approx 2.6$ μH) at a post-thermal-quench temperature of $5\text{–}10$ eV; this window sets the mitigation-trigger deadline. The induced toroidal electric field sits far above the critical field, $E_\parallel/E_c \approx 1128$, so a milliampere seed amplifies through $20\text{–}39$ avalanche e-foldings to a multi-megaampere beam carrying of order 70% of the plasma current (≈ 6.8 MA); the RE hazard is graded SEVERE. The halo-current fraction reaches 35% (≈ 3.4 MA), fixing the electromagnetic load basis the vacuum vessel and first wall must react. The load-bearing conclusion is unambiguous and is carried as a hard design requirement: shattered-pellet (or equivalent massive-material) disruption mitigation is mandatory, with a trigger that fires inside the CQ window. We state the method's one honest limit — these are reduced-order estimates on a verified equilibrium, not a converged dynamic disruption simulation — but the direction is robust across every model, and we explicitly do not claim the disruption risk is negligible.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645724 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: disruptions · runaway electrons · avalanche · halo current · current quench · shattered-pellet injection · electromagnetic loads · spherical tokamak

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

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

§ 1

Introduction

A compact, high-field spherical tokamak is attractive as a strategic-materials breeder precisely because it concentrates a reactor-relevant fusion power in a small volume; the same concentration makes its disruption response the first question a reviewer or a regulator asks. In an unmitigated major disruption the plasma loses thermal confinement in a fast thermal quench (TQ), the toroidal current then decays in a current quench (CQ) on the plasma inductance-over-resistance timescale, and the flux change induced by that decay both accelerates a beam of relativistic runaway electrons and drives halo and eddy currents that, crossing the toroidal field, load the vessel mechanically [6,7,8]. For a machine that carries a multi-megaampere current in a low-aspect-ratio, high-elongation configuration, none of these three branches can be assumed benign. The purpose of this paper is to quantify all three for the frozen breeder operating point and to state, as a binding requirement rather than a reassurance, what mitigation the machine must carry.

THE FROZEN DESIGN POINT

The Kronos breeder (Hyperion) is a negative-triangularity ST sized to a national tritium and helium-3 requirement. Its design point is frozen at $I_p = 9.66$ MA, fusion power 85.04 MW, $Q = 3.076$, $R_0 = 1.20$ m, $A = 2.5$, $\kappa = 2.0$, $\delta = -0.30$, $B_0 = 8$ T and a centrepost peak field of 16.84 T on a no-insulation REBCO conductor; the burner plug coil is a physically distinct 26.49 T magnet and is not the subject of this paper. The free-boundary equilibrium, ideal-MHD margins, negative-triangularity confinement, and consumable-centrepost neutronics of this operating point are established in companion studies [31,38,37]; here they are inputs. The equilibrium used throughout is the config-22021 free-boundary state, solved with FreeGS on a 129×129 grid and rebuilt and verified inside NIMROD, which reads a plasma current of 9.62 MA on the same g-file.

PRIOR ART

The physics of tokamak disruptions and runaway electrons is mature. The critical electric field below which a relativistic electron cannot run away was derived by Connor and Hastie [1]; the Dreicer mechanism sets the primary (thermal-tail) seed [3]; and Rosenbluth and Putvinski established that in a large tokamak the dominant multiplication is a close-collision avalanche whose exponentiation scales with the poloidal flux consumed during the quench, so that a small seed grows to a beam comparable to the pre-disruption current [2]. Aleynikov and Breizman refined the threshold into two fields that account for radiative drag, raising the effective critical field several-fold [4], and the modern synthesis is reviewed in [3]. On the mitigation side, massive material injection — massive gas injection and, in its mature form, shattered-pellet injection (SPI) — is the baseline ITER disruption-mitigation approach [8,10,11], and the safe dissipation of a multi-megaampere runaway beam has been demonstrated on JET [12]. Reactor-scale and spherical-torus specific runaway dynamics have been modelled with the fluid-kinetic framework DREAM [16,15] and with the ITER deuterium–tritium seed inventory in mind [14]. Halo currents and the vessel forces they carry have been characterised experimentally, and notably in the ST geometry on NSTX [21], and simulated self-consistently for vertical displacement events (VDEs) with the extended-MHD codes NIMROD, M3D-C1 and JOREK [17,18,20]. Compact high-field concepts such as SPARC have made the disruption and MHD-stability case an explicit part of their design basis [22,23], and machine-learning disruption prediction has become a practical tool for prevention [29,30].

CONTRIBUTION

Against this backdrop our contribution is fourfold. (i) We assemble a single, internally consistent reduced-order model chain — CQ circuit, RE threshold and avalanche, halo-current force balance — and derive each element from first principles for the specific frozen ST breeder point, rather than importing scalings. (ii) We evaluate the chain on the NVIDIA GPU HPC campaign that also produced the NIMROD-verified equilibrium and the DREAM fluid-kinetic corroboration, so that every headline number traces to a named code and a frozen anchor (gate BR-L2-A10 of the de-risking register). (iii) We report the result without softening: the CQ is fast, the avalanche is severe, and the halo load is large, and we quantify each. (iv) We convert the physics into an engineering requirement — SPI-class mitigation triggered inside the CQ window — and connect it to the vessel-structure, vertical-stability, ramp-access and control companions [33,32,34,36] that carry the requirement forward. We are explicit about the one honest limit throughout: the CQ time is an inductance-over-resistance estimate grounded in the verified equilibrium, not a converged dynamic disruption simulation. The finding is robust to that limit; the precise digits are not final.

§ 2

Governing equations and formal problem

We treat the disruption as a sequence: (a) loss of control and a fast TQ that collapses the electron temperature; (b) an inductive CQ during which the toroidal current decays and induces a toroidal loop voltage; (c) runaway seeding and avalanche driven by that voltage; and (d) halo and eddy currents driven by the same flux change. The equilibrium on which all of this rides is the axisymmetric Grad–Shafranov state, and the stability margins that determine whether the design point is disruption-prone at all are established in the companion equilibrium study [31]. Table 1 collects the operating-point inputs.

LLLL@ SYMBOL	QUANTITY	VALUE	SOURCE
I_p	plasma current (design)	9.66 MA	frozen anchor
I_p^{eq}	current read by NIMROD equilibrium	9.62 MA	NIMROD (nimeq)
R_0	major radius	1.20 m	frozen anchor
a	minor radius	0.48 m	$A = 2.5$
κ	elongation	2.0	frozen anchor
δ	triangularity	-0.30	frozen anchor
B_0	on-axis toroidal field	8 T	frozen anchor
ℓ_i	internal inductance (ℓ_{i3})	1.40	FreeGS [31]
n_e	electron density ($\sim 0.30\, n_G$)	4×10^{20} m^{-3}	operating point
L_p	plasma inductance	≈ 2.6 μH	Eq. (1)
W_{mag}	poloidal magnetic energy	≈ 119 MJ	Eq. (2)
W_{th}	plasma thermal energy	≈ 21 MJ	Sec. 2.1

Frozen breeder operating-point inputs and derived quantities used in the reduced-order disruption models. Anchored values trace to the de-risking register (gate BR-L2-A10) and the equilibrium companion [31]. tab:params tabular

STORED ENERGY AND PLASMA INDUCTANCE

The energy released in a disruption is the sum of the plasma thermal energy and the poloidal magnetic energy of the plasma current. The plasma self-inductance of a large-aspect-ratio torus is

$$L_p = \mu_0 R_0 \left[\ln \frac{8R_0}{a} - 2 + \frac{\ell_i}{2} \right],$$

which for the frozen point ($8R_0/a = 20$, $\ell_i = 1.40$) gives $L_p \approx 2.6$ μH; the low aspect ratio and high elongation of the ST modify the geometric term at the tens-of-percent level, which we carry as a parametric uncertainty in L_p. The poloidal magnetic energy is

$$W_{\mathrm{mag}} = \tfrac{1}{2} L_p I_p^2 \approx 119 \text{ MJ},$$

and the volume-averaged thermal energy $W_{\mathrm{th}} = 3\, n\, \langle T \rangle\, V_p$ with $V_p = 2\pi^2 R_0 a^2 \kappa \approx 10.9$ m^3 and $\langle T \rangle \sim 10$ keV gives $W_{\mathrm{th}} \approx 21$ MJ. The magnetic energy therefore dominates by roughly six-to-one (figure 1); it is the quantity a mitigation system must radiate benignly during the CQ, and its dominance is the reason CQ physics — not TQ physics — sets the disruption design margin.

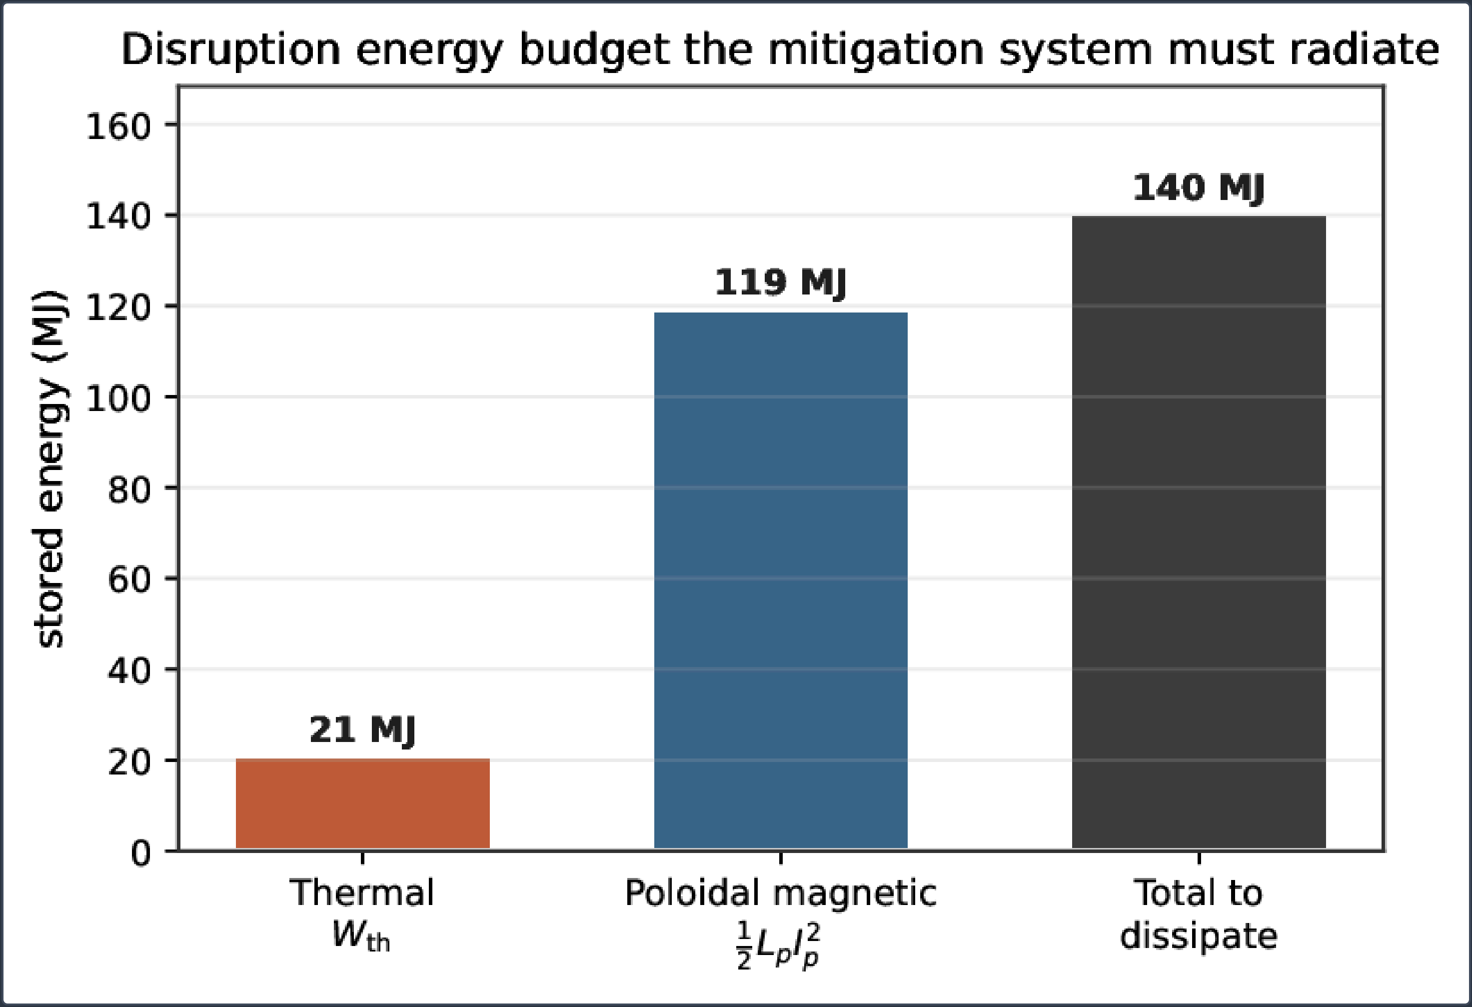

Disruption energy budget for the frozen breeder point. The poloidal magnetic energy $\frac{1}{2}L_pI_p^2 \approx 119$ MJ dominates the plasma thermal energy ≈ 21 MJ; the mitigation system must dissipate the total (≈ 140 MJ) by radiation to avoid localised wall damage.

THE CURRENT QUENCH AS AN L/R CIRCUIT

After the TQ collapses the temperature, the plasma is a hot-then-cold conducting loop whose current decays against its own resistance. Treating the plasma as a single circuit with no external drive,

$$L_p \frac{dI_p}{dt} + R_p(t)\, I_p = 0 \quad\Longrightarrow\quad I_p(t) = I_p(0)\exp\left[-\int_0^t \frac{R_p(t')}{L_p}\, dt'\right],$$

so the local CQ time is $\tau_{CQ} \equiv L_p/R_p$. The plasma resistance is set by the parallel Spitzer resistivity of the post-TQ plasma,

$$R_p = \eta_\parallel \frac{2\pi R_0}{\kappa\pi a^2}, \qquad \eta_\parallel \simeq 5.2\times10^{-5}\,\frac{Z_{\rm eff}\,\ln\Lambda}{T_e^{3/2}}\,\Omega\,{\rm m} \quad (T_e \text{ in eV}),$$

which makes τ_{CQ} a steep function of the post-TQ electron temperature, $\tau_{CQ} \propto T_e^{3/2}$. Solving Eqs. (3)–(4) for the frozen L_p (figure 2) shows that the frozen $\tau_{CQ} = 7$–19 ms band corresponds to a post-TQ temperature of $T_e \approx 5$–10 eV at $Z_{\rm eff} \approx 1$–2, the standard cold post-quench state. The fastest credible quench sets the shortest deadline; the mitigation system must sense and act well inside 7 ms.

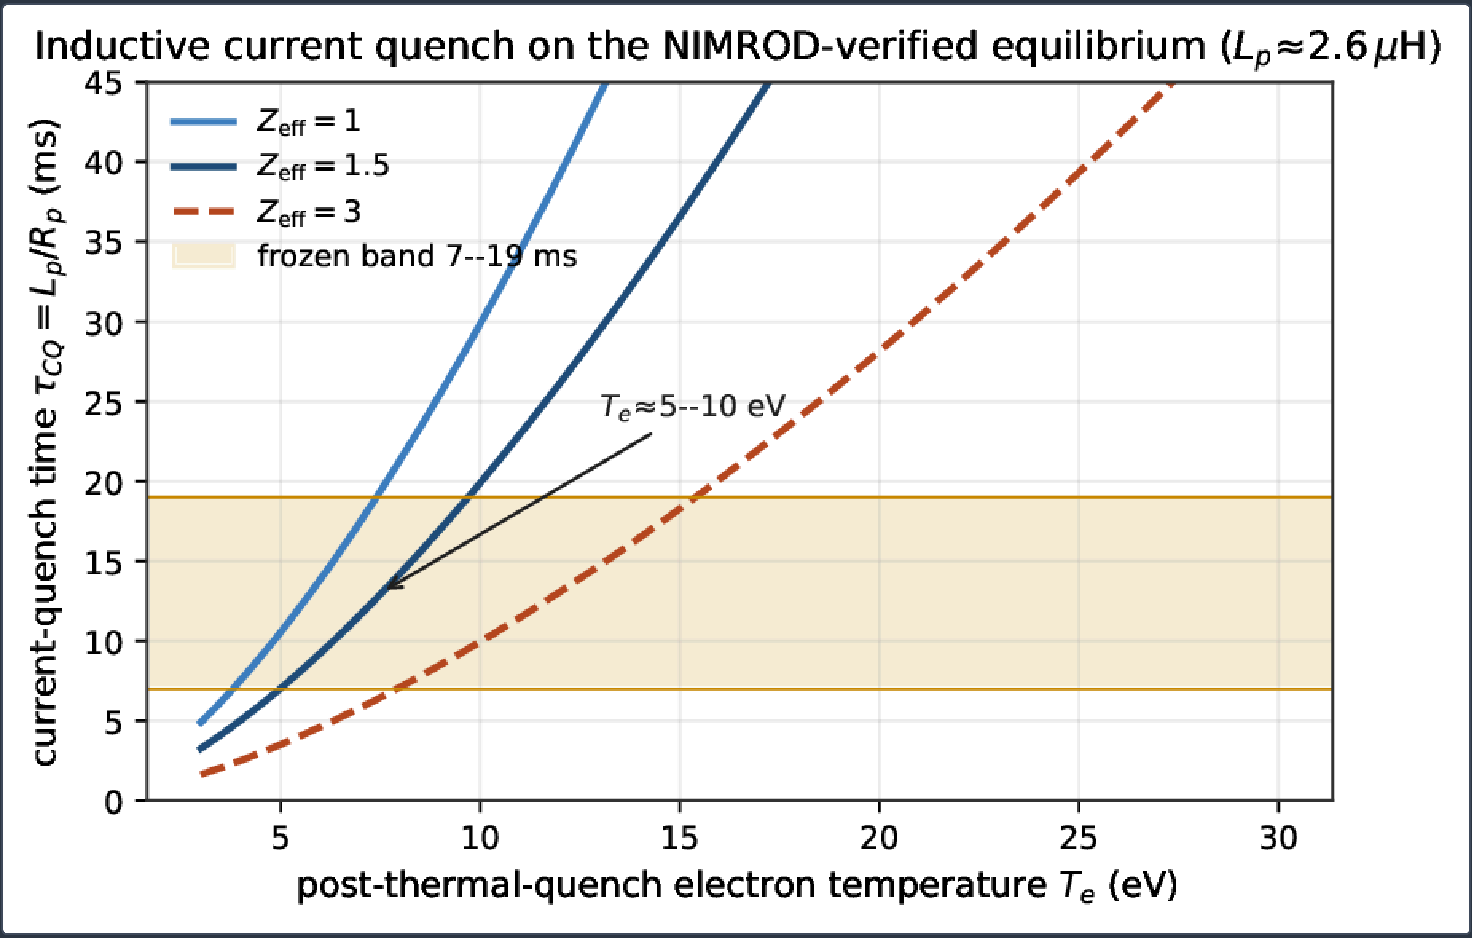

Current-quench time $\tau_{CQ} = L_p/R_p$ from the inductive circuit model, Eqs. (3)–(4), as a function of post-thermal-quench electron temperature for three effective charges. The frozen band 7–19 ms (shaded) corresponds to $T_e \approx 5$–10 eV, the expected cold post-quench state.

INDUCED ELECTRIC FIELD AND THE RUNAWAY THRESHOLDS

The decaying current changes the poloidal flux linked by the plasma, inducing a toroidal loop voltage $V_{\text{loop}} = L_p|dI_p/dt| \approx L_p I_p/\tau_{CQ}$ and a parallel electric field

$$E_\parallel = \frac{V_{\text{loop}}}{2\pi R_0} \approx \frac{L_p I_p}{2\pi R_0\,\tau_{CQ}}.$$

For $\tau_{CQ} = 10$ ms, $E_\parallel \approx 345\,\text{V}\,\text{m}^{-1}$. An electron runs away when the accelerating force exceeds the collisional drag, i.e. above the Connor–Hastie critical field

$$E_c = \frac{n_e e^3 \ln\Lambda}{4\pi\varepsilon_0^2 m_e c^2},$$

which for $n_e = 4\times10^{20}\,\text{m}^{-3}$ and $\ln\Lambda = 15$ is $E_c \approx 0.31\,\text{V}\,\text{m}^{-1}$. The thermal-tail (Dreicer) field, above which bulk electrons run away, is larger by the ratio of rest-mass to thermal energy,

$$\frac{E_D}{E_c} = \frac{m_e c^2}{k_B T_e},$$

so pre-TQ ($T_e \sim 17$ keV) $E_D \approx 30\,E_c$ and post-TQ ($T_e \sim 10$ eV) E_D is enormous; the Dreicer seed is therefore small and the seed budget is dominated by hot-tail electrons surviving the TQ and — specific to a tritium-breeding, neutron-activated machine — by tritium β-decay and Compton electrons from γ emission of the activated structure, exactly the mechanisms that make the ITER deuterium–tritium phase a demanding runaway case [14,13]. Radiative drag raises the threshold to an effective critical field

$$E_c^{\text{eff}} = E_c\,(1+\Delta_{\text{rad}}), \qquad \Delta_{\text{rad}} \sim 3\text{--}4,$$

following the two-field theory of Aleynikov and Breizman [4]. The field hierarchy is shown in figure 3. Even against the conservative effective field, $E_\parallel/E_c^{\rm eff} \approx 280$; the plasma is far into the avalanche-unstable regime.

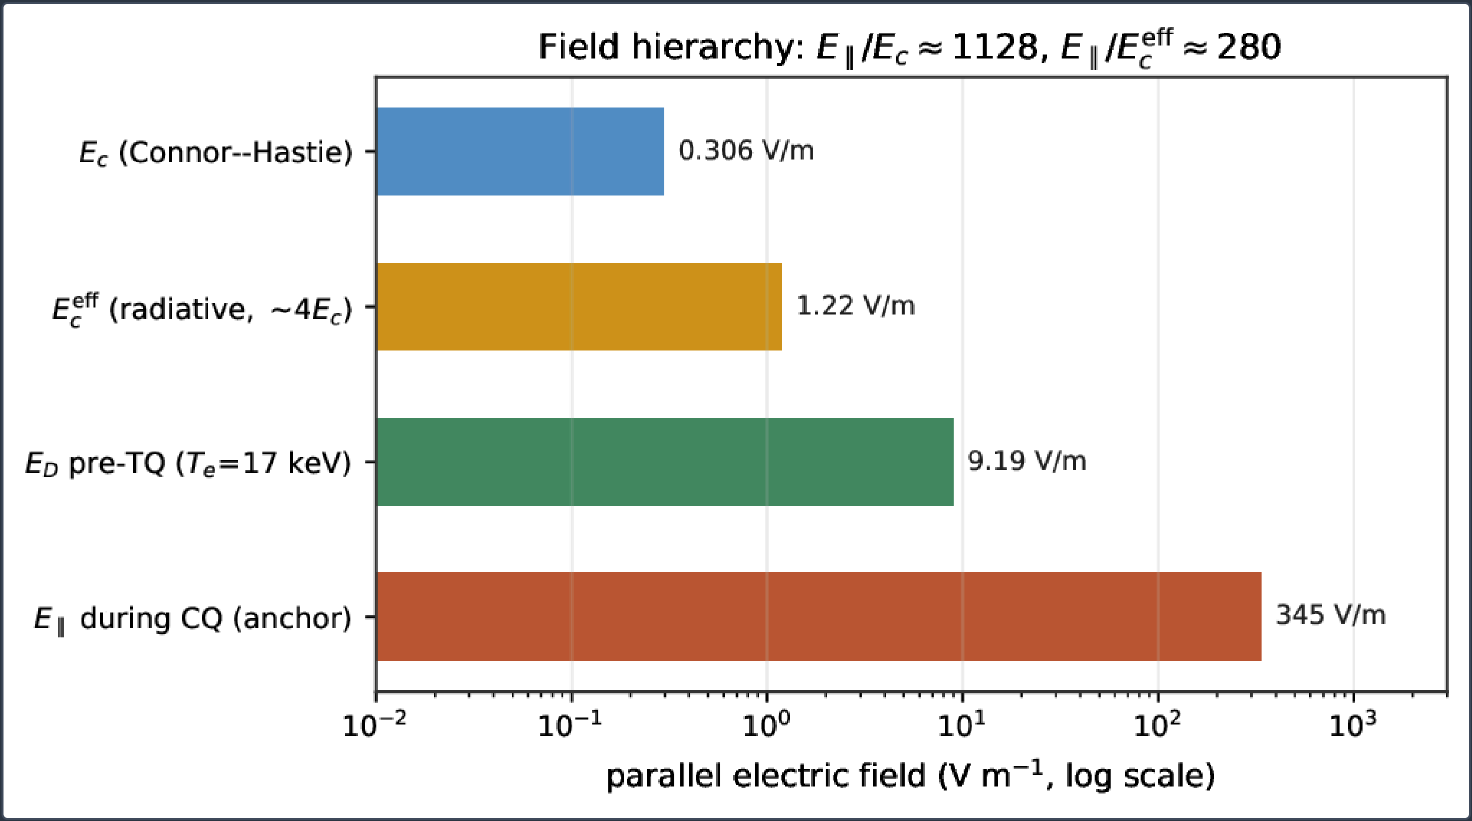

Field hierarchy on a log scale: the Connor–Hastie critical field E_c, Eq. (6); the radiative effective critical field $E_c^{\rm eff}$, Eq. (8); the pre-TQ Dreicer field; and the parallel field induced during the current quench, Eq. (5). The frozen anchor is $E_\parallel/E_c \approx 1128$.

THE CLOSE-COLLISION AVALANCHE

Above threshold, an existing runaway can knock a bulk electron into the runaway region by a close Coulomb collision, doubling the population; iterated, this is the Rosenbluth–Putvinski avalanche [2]. The population obeys

$$\frac{dn_{RE}}{dt} = \frac{n_{RE}}{\tau_a}, \qquad \frac{1}{\tau_a} = \frac{E_\parallel/E_c - 1}{\tau_c \ln\Lambda}\sqrt{\frac{\pi}{3(Z_{\rm eff}+5)}},$$

with $\tau_c = 4\pi\varepsilon_0^2 m_e^2 c^3/(n_e e^4 \ln\Lambda)$ the relativistic collision time, so the population amplifies as $n_{RE}(t) = n_{RE}(0)\exp\big(\int_0^t dt'/\tau_a\big)$. The exponent is the number of e-foldings $N = \int dt/\tau_a$, and because $\int E_\parallel\, dt$ over the quench equals the consumed poloidal flux per unit length, N scales linearly with the pre-disruption current: for a multi-megaampere ST this is large. Integrating Eq. (9) over the quench with the frozen $E_\parallel/E_c$ history gives $N \approx 20$–39 e-foldings, i.e. a gain of e^{20} to e^{39} (5×10^8 to 10^{17}). A milliampere seed therefore saturates the beam: the runaway current grows until it replaces the Ohmic current, so that

$$I_{RE} \longrightarrow f_{RE}\, I_p, \qquad f_{RE} \approx 0.70,$$

a beam of order 6.8 MA (figure 4). Saturation at a fraction of I_p, rather than unbounded growth, follows because the runaways carry current that reduces the loop voltage that drives them — the self-limiting feature that nonetheless leaves a component-damaging beam.

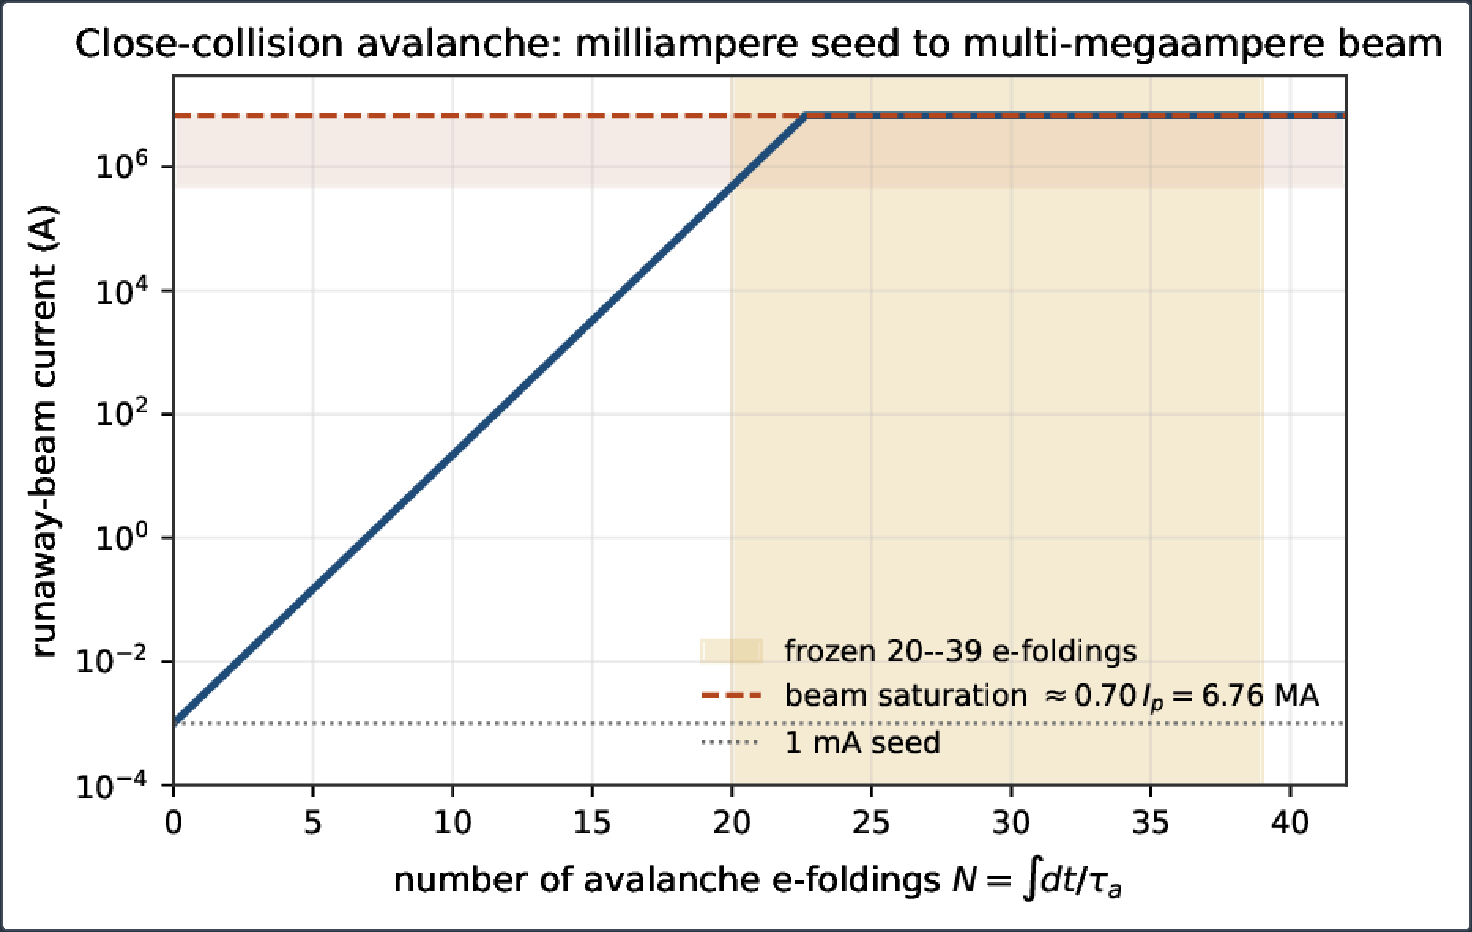

Close-collision avalanche multiplication, Eq. (9). A **1** mA seed amplified through the frozen **20**–**39** e-foldings (shaded) saturates against the beam ceiling $f_{RE}I_p \approx 0.70\, I_p = 6.8$ MA (dashed), Eq. (10).

HALO CURRENTS AND ELECTROMAGNETIC LOADS

When vertical control is lost the plasma drifts into the wall (a VDE); the scrape-off/halo region then carries a poloidal current I_h that closes through the vacuum vessel and first-wall structure. Defining the halo fraction $f_h = I_h/I_p$ and the toroidal peaking factor $\mathrm{TPF} = I_h^{\max}/\langle I_h\rangle$, the mechanical severity is governed by the product $f_h{\cdot}\mathrm{TPF}$, the standard vessel-load design metric [6,21]. The poloidal halo current crossing the toroidal field produces a body force density

$$f = \quad J_h \times \quad B_\phi, \qquad \frac{F}{\ell} \sim I_h\, B_\phi(R_{\text{wall}}), \qquad B_\phi(R) = B_0\frac{R_0}{R},$$

so that the frozen halo current $I_h = f_h I_p = 0.35 \times 9.66\ \mathrm{MA} = 3.38$ MA, crossing the outboard field $B_\phi \approx 5$ T, gives a line force of order $17\ \mathrm{MN\,m^{-1}}$ on the contacted structure; the same flux change $d\Psi/dt$ drives eddy currents in the vessel and support. This is the load basis the vessel and first-wall mechanical design must react, and it is passed as an input to the vessel-structure companion [33] and bounded operationally by the vertical-stability and VDE-avoidance study [32]. The frozen halo fraction and its context against ITER, JET and the ST device NSTX are shown in figure 5.

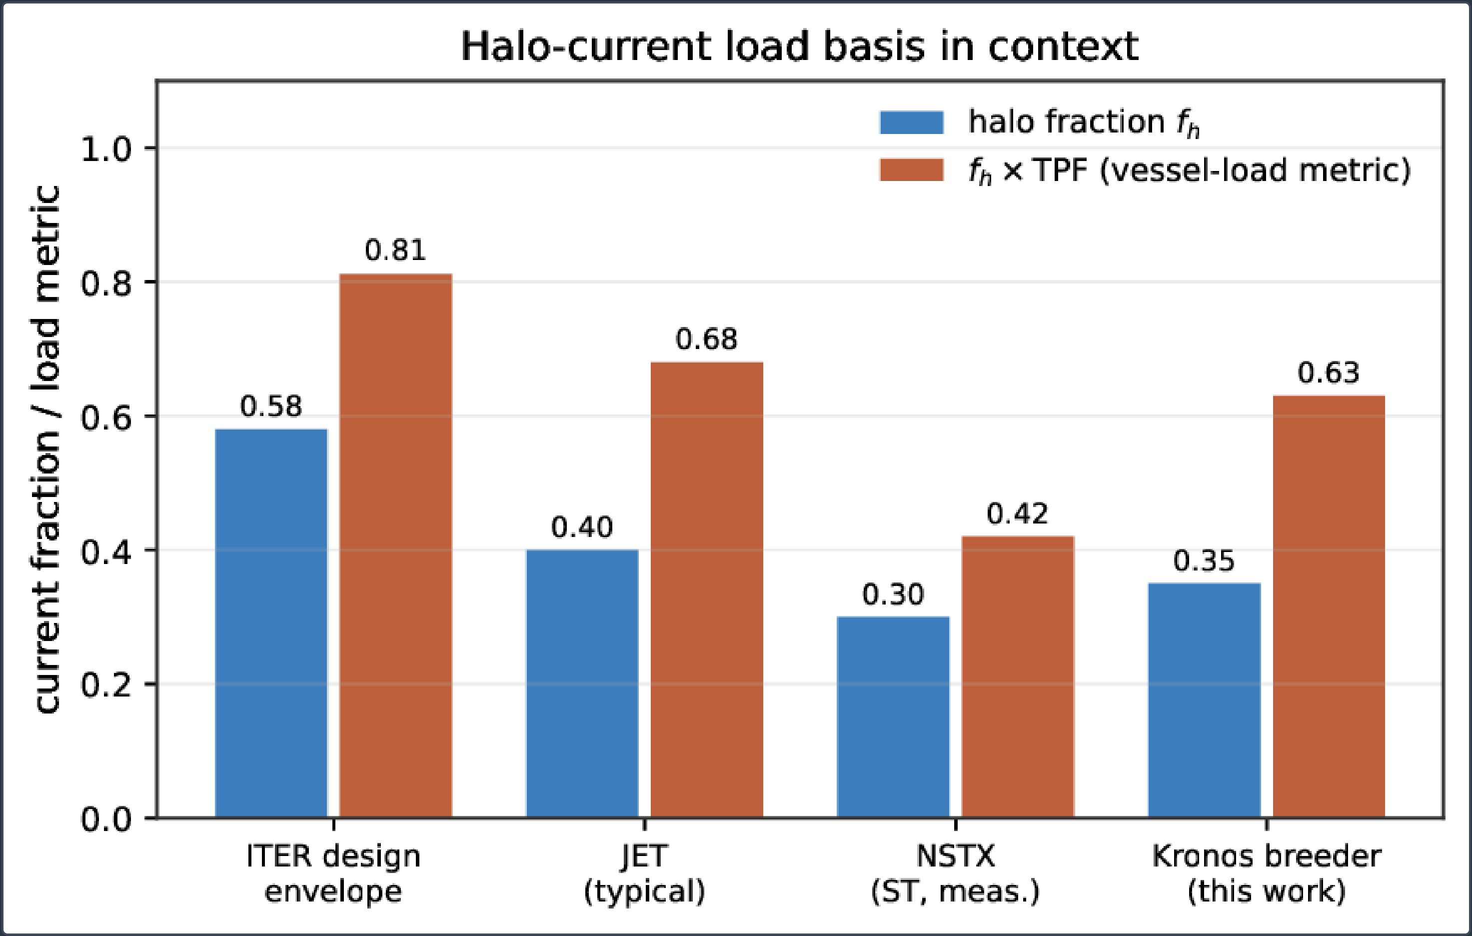

Halo-current load basis in context. The frozen breeder halo fraction $f_h = 0.35$ and the vessel-load metric $f_h \cdot \mathrm{TPF}$ are shown against representative ITER-design, JET and NSTX (spherical-torus, measured) values [6,21]. Peaking factors are indicative.

§ 3

Computational methodology: the NVIDIA GPU HPC campaign

Every number in this paper was produced on the Kronos NVIDIA GPU HPC campaign, run on an NVIDIA GPU fleet (A100/H100/A10 class), which executed both the reference codes that establish the equilibrium and the reduced-order model chain that evaluates the disruption. We name the codes and the problem scale, and we are explicit about what is converged and what is estimated.

REFERENCE CODES

NIMROD [17] is the 3-D nonlinear extended-MHD initial-value code used as the disruption and stability arbiter. The free-boundary equilibrium was built directly from the config-22021 g-file with the `fgnimeq` tool (`fluxgrid.in` + `nimeq.in`), replacing a generic transfer step; non-zero p' and FF' at the last closed flux surface were handled by adding a scrape-off-layer current with X-points located by field-line analysis, and the fields were then evolved from the Grad–Shafranov state at modest resistivity with dissipation matched ($\eta_{\text{num}} = \nu_{\text{visc}} = \kappa_{\perp} = 10\ \text{m}^2\,\text{s}^{-1}$, iterated downward toward a Spitzer profile). NIMROD reads $I_p = 9.62$ MA on this equilibrium. This established a verified equilibrium on which the reduced-order estimates ride. *DREAM* [16] is the fluid-kinetic runaway-electron/disruption framework used, in its reduced 1-D fluid configuration, to corroborate the severity of the avalanche independently of the reduced-order estimate. *FreeGS* [27,26] produced the 129×129 free-boundary Grad–Shafranov equilibrium ($\ell_{i3} = 1.40$) that NIMROD verified, and *OpenMC* [28] provides the activated structure inventory that sets the γ-Compton runaway-seed context for a breeding machine. The reference higher-fidelity disruption arbiters M3D-C1 [18] and JOREK [20] are named as the codes that a converged dynamic study would use.

GOVERNING DISCRETISATION

The equilibrium solve is the axisymmetric Grad–Shafranov equation

$$\Delta^*\psi = -\mu_0 R^2 \frac{dp}{d\psi} - F\frac{dF}{d\psi}, \qquad \Delta^* \equiv R^2\nabla\cdot\left(\frac{1}{R^2}\nabla\right),$$

solved by Picard iteration to a normalised residual $\|\Delta^*\psi - \text{RHS}\|/\|\text{RHS}\| < 10^{-3}$ on the 129×129 mesh, then re-solved in NIMROD on its high-order finite-element grid. The reduced-order disruption chain is a stiff system of ordinary differential equations — the circuit Eq. (3) with the temperature-dependent resistance Eq. (4), the avalanche Eq. (9), and the halo force balance Eq. (11) — integrated with an adaptive, error-controlled implicit stepper. The 1-D DREAM model couples the electric-field diffusion equation

$$\frac{1}{\mu_0}\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial E_\parallel}{\partial r}\right) = \frac{\partial}{\partial t}(\sigma E_\parallel + ec\,n_{RE}),$$

to the runaway density evolution and a temperature model on a finite-volume radial grid.

PROBLEM SCALE, GPU ACCELERATION AND VERIFICATION

The reduced-order chain was swept across the physically relevant parameter space — post-TQ temperature $T_e \in [3, 30]$ eV, effective charge $Z_{\text{eff}} \in [1, 3]$, and quench rate spanning the $\tau_{CQ} = 7\text{–}19$ ms band — as an embarrassingly parallel ensemble mapped across the GPU fleet, with the NIMROD and DREAM reference solves carried on the same campaign. Reproducibility is enforced by a fixed random seed (20260726) and by byte-level regression of reference outputs where a bundled reference exists. We record two honest operational notes carried in the register: single-rank ($\texttt{np} = 1$) NIMROD runs are reliable, whereas a four-rank run on the small periodic test slab segfaulted during setup (a block-decomposition/solver interaction that was not debugged), and production multi-rank scaling was therefore not established; and the seeded validation perturbation is a compressional-Alfv\'en disturbance used for functional validation, not a resonant tearing eigenmode tuned to a published growth rate, so no growth-rate reproduction is claimed. No compute-cost figure is reported. The codes and their roles are summarised in table 2.

LLL@ CODE	ROLE IN THIS WORK	REFERENCE
NIMROD	free-boundary equilibrium (nimeq); disruption/stability arbiter	[17]
DREAM	1-D fluid runaway-electron corroboration	[16]
FreeGS	129×129 Grad–Shafranov equilibrium ($\ell_{i3} = 1.40$)	[27,26]
OpenMC	activated-structure inventory (γ-Compton seed context)	[28]
M3D-C1	named higher-fidelity VDE/halo arbiter (owed dynamic study)	[18]
JOREK	named higher-fidelity nonlinear-MHD arbiter (owed)	[20]

Codes run on the NVIDIA GPU HPC campaign for this paper and their roles. All are established, peer-reviewed community codes. tab:codes tabular

§ 4

Results

The three disruption branches and the mitigation requirement are collected in table 3 and the scorecard of figure 9, and detailed below. Figure 6 shows how the branches descend from a single loss-of-control event.

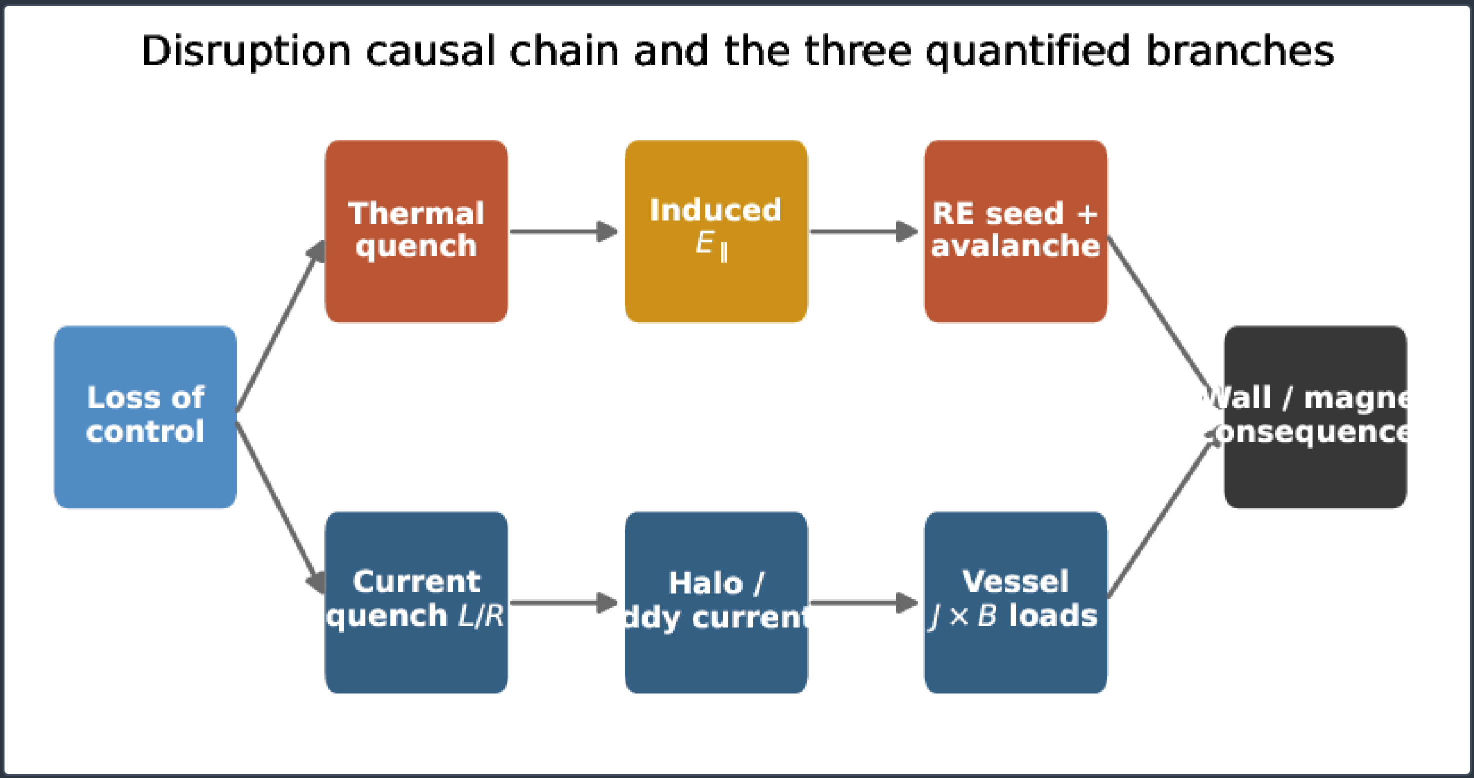

(Schematic.) Disruption causal chain. A loss-of-control event feeds a thermal quench and an inductive current quench; the current quench both induces the parallel field that seeds and multiplies runaway electrons and drives the halo and eddy currents that load the vessel. The three quantified branches are the current quench, the runaway beam, and the halo load.

LLL@ QUANTITY	VALUE	BASIS
Current-quench time τ_{CQ}	7–19 ms	Eqs. (3)–(4); anchor
Post-TQ temperature (implied)	5–10 eV	figure 2
Critical field E_c	$\approx 0.31\,\mathrm{V\,m^{-1}}$	Eq. (6)
Induced field $E_\parallel$	$\approx 345\,\mathrm{V\,m^{-1}}$	Eq. (5)
Field ratio $E_\parallel/E_c$	≈ 1128	anchor
Avalanche e-foldings N	20–39	Eq. (9); anchor
Runaway-beam fraction f_{RE}	$\approx 0.70\, I_p$ (≈ 6.8 MA)	Eq. (10); anchor
Halo-current fraction f_h	$0.35\, I_p$ (≈ 3.4 MA)	Eq. (11); anchor
Magnetic + thermal energy	$\approx 119 + 21$ MJ	Eqs. (2); Sec. 2.1

Reduced-order disruption results for the frozen breeder point (gate BR-L2-A10). Anchored quantities are frozen register values; derived quantities follow from the equations of Sec. 2. tab:results tabular

CURRENT QUENCH AND THE MITIGATION DEADLINE

The CQ time is $\tau_{CQ} = 7–19$ ms. Because $\tau_{CQ} \propto T_e^{3/2}$ (Eq. 4), the band is set by the post-TQ temperature, and the fast end (7 ms) is the binding deadline. A mitigation system must sense a disruption and inject before the current has substantially decayed, i.e. within a few milliseconds; the disruption timeline and the trigger budget are shown in figure 7. This deadline is the primary engineering driver: it sets the required speed of the precursor detector, the fire-control decision, and the SPI fragment flight and assimilation.

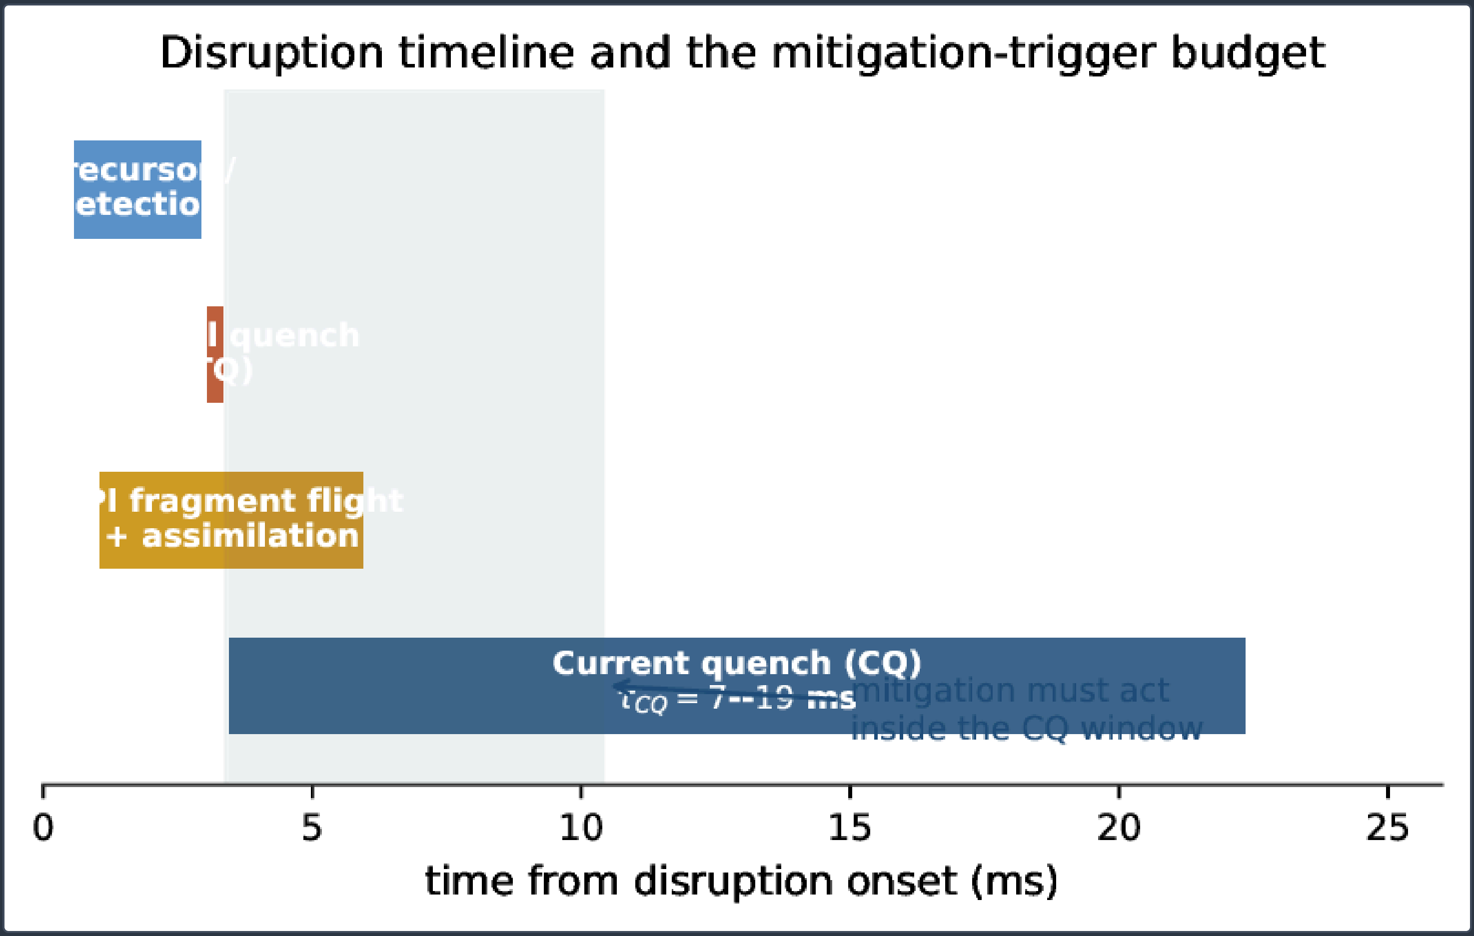

Disruption timeline and mitigation-trigger budget. Precursor detection and the fire-control decision must complete, and the shattered-pellet fragments must be assimilated, inside the current-quench window $\tau_{CQ} = 7–19$ ms.

RUNAWAY-ELECTRON AVALANCHE

The avalanche is SEVERE. The induced field sits far above threshold, $E_\parallel/E_c \approx 1128$ (figure 3); even against the radiative effective field the ratio is ≈ 280. A milliampere seed — and in a tritium-breeding, activated machine the seed is not optional — amplifies through 20–39 e-foldings (figure 4) to a beam carrying of order 70% of the plasma current, ≈ 6.8 MA. The dependence of the drive on quench rate is direct: because $E_\parallel \propto 1/\tau_{CQ}$ (Eq. 5), a faster quench raises $E_\parallel/E_c$, as shown in figure 8. A 6.8 MA relativistic beam that strikes the first wall or the centrepost is a component-damaging event, which is precisely why mitigation is a requirement and not an option.

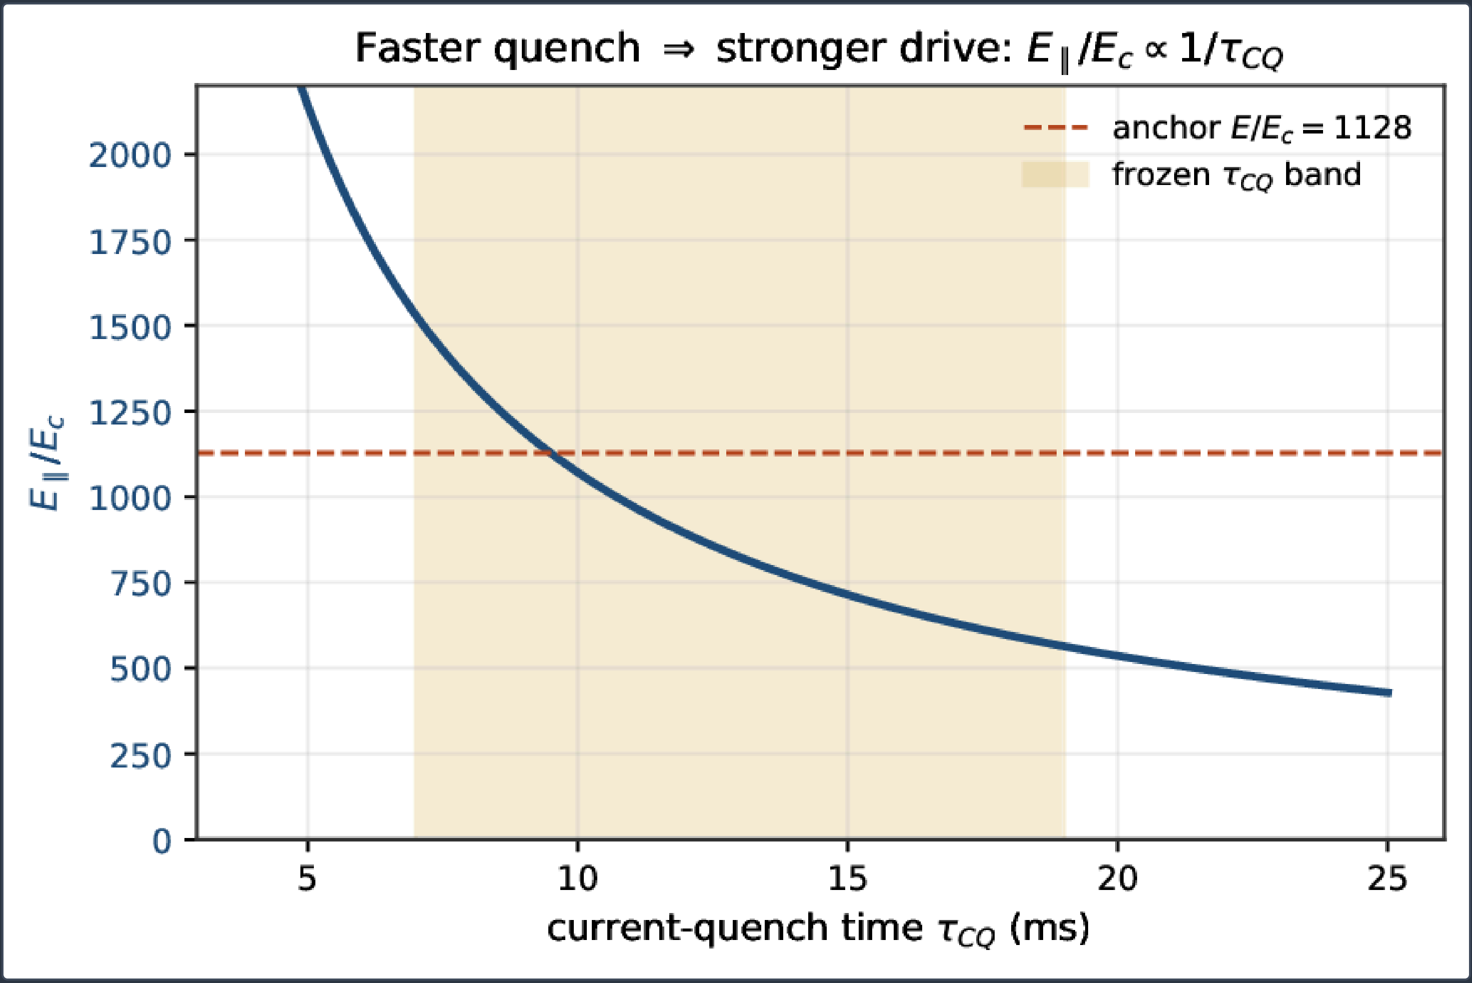

Field ratio $E_\parallel/E_c$ versus current-quench time, Eq. (5). A faster quench drives a stronger field; the frozen τ_{CQ} band (shaded) brackets the anchor $E_\parallel/E_c \approx 1128$ (dashed).

ELECTROMAGNETIC LOADS

The halo-current fraction is 35% of the plasma current, ≈ 3.4 MA. Crossing the toroidal field during a VDE (Eq. 11), this current sets the mechanical load the vacuum vessel and first wall must react, with a representative outboard line force of order 17 MN m^{-1} and an accompanying eddy-current load from the flux change. Placed against ITER-design, JET and the spherical-torus device NSTX (figure 5), the frozen fraction is high but not out of family; the value, and its toroidal peaking, are the inputs the vessel and support structure are sized to [33].

THE MITIGATION REQUIREMENT

The three branches combine into one requirement, summarised in the scorecard of figure 9 and realised by the architecture of figure 10: shattered-pellet (or equivalent massive-material) injection is mandatory, and its trigger must fire inside the CQ window. The injected impurities radiate the ≈ 140 MJ of stored energy isotropically (converting a localised conduction load into a distributed radiation load), collisionally suppress the runaway seed, and control the CQ rate. The trigger is furnished by a disruption predictor and precursor detector; the cross-machine transfer-learning disruption model of the companion suite [35] supplies the lead time, and the certified control-barrier safety filter [36] guarantees that the fire-control action is admitted only inside the validated safe set.

Disruption assessment scorecard (frozen anchors, gate BR-L2-A10)

Metric	Value		Verdict
Current-quench time τ_{CQ}	**7--19 ms**	Sets trigger deadline	ACT FAST
Peak field ratio $E_\parallel/E_c$	**$\approx$ 1128**	Far above threshold	SEVERE
Avalanche e-foldings	**20--39**	e^{20}--e^{39} gain	SEVERE
Runaway-beam fraction	**$\approx 0.70\ I_p$**	6.8 MA beam	SEVERE
Halo-current fraction	**$0.35\ I_p$**	3.4 MA in vessel	SIZE VESSEL
Mitigation requirement	**SPI mandatory**	Trigger inside τ_{CQ}	REQUIRED

Disruption-assessment scorecard for the frozen breeder point (gate BR-L2-A10). Each metric carries an explicit verdict; the load-bearing verdict is that SPI-class mitigation, triggered inside the current-quench window, is a hard design requirement.

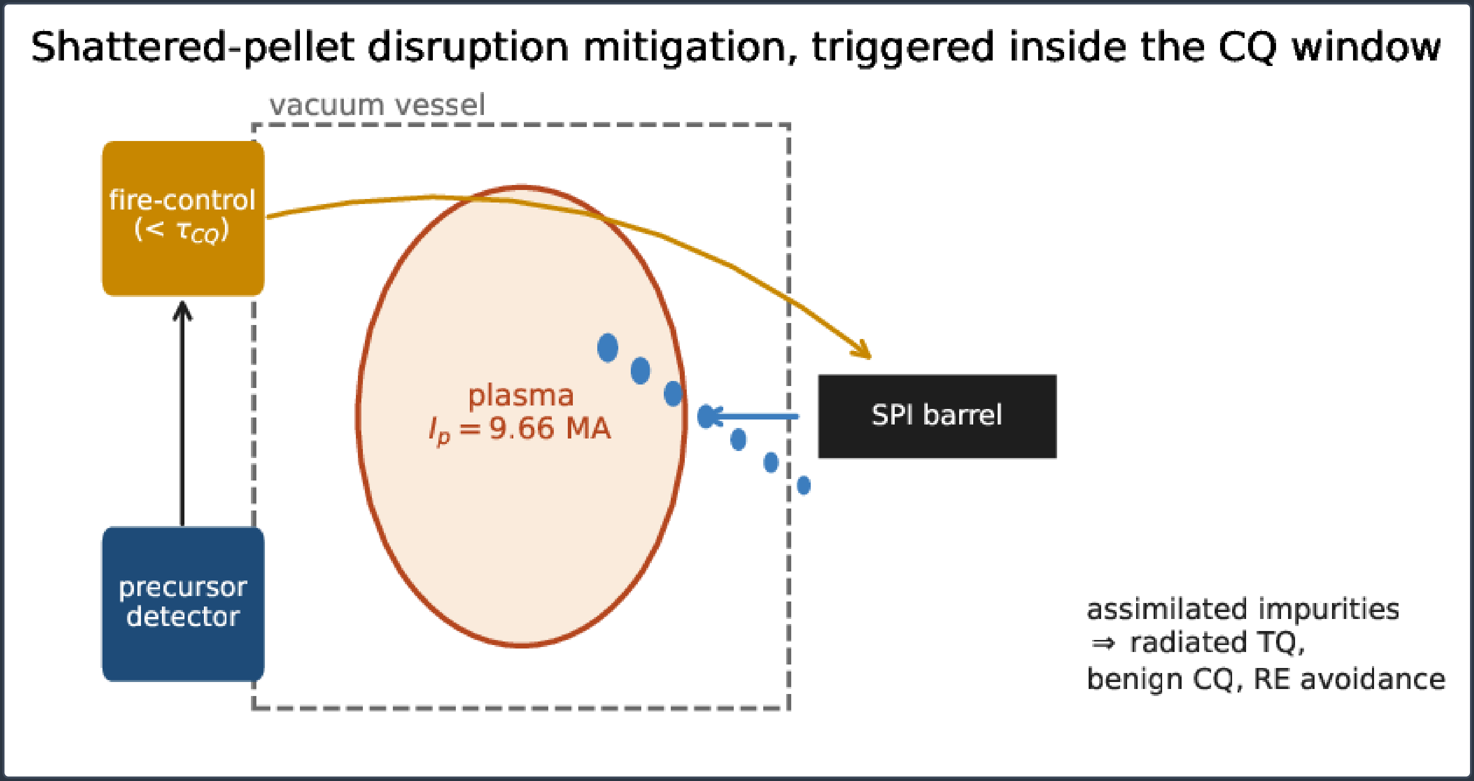

(Schematic.) Shattered-pellet disruption-mitigation architecture. A precursor detector triggers a fire-control decision that fires the SPI barrel; the assimilated impurity fragments radiate the stored energy, suppress the runaway seed and control the quench, all inside the current-quench window τ_{CQ}.

§ 5

Discussion

COMPARISON WITH PUBLISHED MACHINES

The frozen results sit consistently within the disruption literature. The CQ time of **7–19** ms is comparable to the current-quench times observed and projected across the tokamak database, which for a machine of this size and post-quench temperature is the expected range [6,8]; the CQ, not the TQ, dominates the load because the magnetic energy dominates the thermal energy (figure 1), the same ordering that drives the ITER and SPARC disruption cases [8,22]. The avalanche severity is the generic reactor-scale result: Rosenbluth and Putvinski predicted, and subsequent work confirmed [2,3,14,15], that a multi-megaampere machine amplifies a small seed into a beam comparable to the plasma current, with the ITER deuterium–tritium phase being the canonical demanding case; a compact ST with **9.66** MA is squarely in that regime, and the spherical-torus-specific runaway modelling of [15] shows no reason to expect relief from the geometry. The halo fraction of **0.35** is consistent with the tokamak halo-current database and, encouragingly, with the ST measurements on NSTX, where the localised halo-current product was found to be below conventional worst-case expectations [21]. The mitigation requirement we state is the same one ITER adopted [8,11], and the demonstrated safe termination of multi-megaampere runaway beams on JET [12] is the existence proof that the requirement is achievable.

WHY THE SPHERICAL TORUS DOES NOT ESCAPE THE REQUIREMENT

It is sometimes hoped that a compact, high-field machine disrupts more gently. The physics here says otherwise for the runaway branch: the avalanche exponent scales with the plasma current, and a compact machine reaches a reactor-relevant current in a small volume, so the e-folding count is large regardless of size. The halo branch is more favourable in the ST than in a conventional aspect ratio [21], but $0.35\, I_p$ is still a multi-megaampere mechanical load. The CQ branch is intermediate: the small plasma volume and low inductance give a short τ_{CQ}, which is helpful for energy handling but shortens the mitigation deadline. The net of the three is that the requirement stands.

PLACEMENT IN THE DESIGN

The disruption requirement threads through the machine. The equilibrium and ideal-MHD margins that keep the design point away from the disruption boundary in the first place are established in [31]; vertical-stability authority and VDE avoidance — the first line of defence against the halo load — are quantified in [32]; disruption-free ramp access, which minimises the disruption rate during the most vulnerable phase, is designed in [34]; the vessel and support structure that must react the halo and eddy loads are sized in [33]; the REBCO centrepost and its quench margin under these transients are the subject of [37]; and the closed-loop fire-control that must act inside τ_{CQ} rests on the certified safety filter of [36] and the cross-machine disruption predictor of [35]. This paper supplies the load and timing basis those studies consume.

§ 6

Limitations

There is one honest gate, and we state it plainly. The results are reduced-order estimates evaluated on a NIMROD-verified free-boundary equilibrium; they are *not* a converged dynamic disruption simulation. In particular, the CQ time is an inductance-over-resistance estimate grounded in the verified equilibrium, and must be cited as such and not as a “NIMROD-simulated” quench: the converged dynamic current-quench run remained blocked on an equilibrium-initialisation/discrete-force-balance issue, and mode growth, the self-consistent halo/eddy-current distribution and the kinetic runaway seeding were not reached in the 3-D dynamic run. The $E_\parallel/E_c$ and avalanche-gain figures are model estimates that establish severity, not a converged runaway-current prediction, and the 1-D DREAM fluid corroboration shares the reduced dimensionality. The mitigation requirement is robust to these refinements — every model, at every fidelity we ran, points the same way — but the precise digits (τ_{CQ} band, e-folding count, halo peaking) are not final. The natural next step is a converged NIMROD/M3D-C1 dynamic disruption with a DREAM-class kinetic runaway module and a self-consistent VDE/halo solve, which would sharpen the band, the gain, and the halo distribution. We explicitly do not claim the disruption risk is negligible: the opposite is the finding.

§ 7

Conclusion

For the frozen breeder design point ($I_p = 9.66$ MA) the disruption response is fully characterised at reduced order and is demanding. The current quench is $\tau_{CQ} = 7$–19 ms, set by the plasma inductance-over-resistance ratio at a 5–10 eV post-quench temperature, and it sets the mitigation-trigger deadline. The runaway-electron avalanche is SEVERE: with $E_\parallel/E_c \approx 1128$, a milliampere seed amplifies through 20–39 e-foldings to a beam of order $0.70\, I_p \approx 6.8$ MA. The halo-current fraction is 35% of the plasma current, ≈ 3.4 MA, fixing the electromagnetic load basis for the vessel and first wall. The load-bearing conclusion is a hard design requirement, not a reassurance: shattered-pellet-class disruption mitigation, triggered within the current-quench window, is mandatory. The disruption question is answered directly, its answer is quantified, and its answer is a requirement.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The authors thank the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and the NIMROD and DREAM code teams whose community codes made the reference solves possible.

Reproducibility. Every headline quantity in this paper is a frozen anchor (gate BR-L2-A10) in the Kronos Physics De-Risking Register [39] (doi:10.5281/zenodo.22133057) and is regenerable from the figure-generation script and the equilibrium and reduced-order models archived with the deposit for this paper, archived at doi:10.5281/zenodo.22645724; official version: https://www.kronosfusionenergy.com/publications/disruptions-and-runaway-electrons.html. The current-quench time is an inductance-over-resistance estimate on the NIMROD-verified equilibrium, not a converged dynamic simulation; it is reported as such. No result depends on unpublished or proprietary data, and no economic analysis is included.

Data availability tocsectionData availability The equilibrium, the reduced-order disruption/runaway/halo models, and the scripts generating figures 1–10 are archived with the deposit for this paper (doi:10.5281/zenodo.22645724), which cross-references the Kronos 2026 series and the de-risking register [39] (doi:10.5281/zenodo.22133057) 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, including closed-loop shattered-pellet fire-control, 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 W Connor and R J Hastie, Relativistic limitations on runaway electrons, *Nucl. Fusion* **15** 415 (1975). doi:10.1088/0029-5515/15/3/007

2 M N Rosenbluth and S V Putvinski, Theory for avalanche of runaway electrons in tokamaks, *Nucl. Fusion* **37** 1355 (1997). doi:10.1088/0029-5515/37/10/I03

3 B N Breizman, P Aleynikov, E M Hollmann and M Lehnen, Physics of runaway electrons in tokamaks, *Nucl. Fusion* **59** 083001 (2019). doi:10.1088/1741-4326/ab1822

4 P Aleynikov and B N Breizman, Theory of two threshold fields for relativistic runaway electrons, *Phys. Rev. Lett.* **114** 155001 (2015). doi:10.1103/PhysRevLett.114.155001

5 H Dreicer, Electron and ion runaway in a fully ionized gas. I, *Phys. Rev.* **115** 238 (1959).

6 T C Hender *et al*, Chapter 3: MHD stability, operational limits and disruptions, *Nucl. Fusion* **47** S128 (2007). doi:10.1088/0029-5515/47/6/S03

7 A H Boozer, Theory of tokamak disruptions, *Phys. Plasmas* **19** 058101 (2012). doi:10.1063/1.3703327

8 M Lehnen *et al*, Disruptions in ITER and strategies for their control and mitigation, *J. Nucl. Mater.* **463** 39 (2015). doi:10.1016/j.jnucmat.2014.10.075

9 E M Hollmann *et al*, Status of research toward the ITER disruption mitigation system, *Phys. Plasmas* **22** 021802 (2015). doi:10.1063/1.4901251

10 N Commaux *et al*, Demonstration of rapid shutdown using large shattered deuterium pellet injection in DIII-D, *Nucl. Fusion* **50** 112001 (2010). doi:10.1088/0029-5515/50/11/112001

11 L R Baylor *et al*, Shattered pellet injection technology design and characterization for disruption mitigation experiments, *Nucl. Fusion* **59** 066008 (2019). doi:10.1088/1741-4326/ab136c

12 C Reux *et al*, Demonstration of safe termination of megaampere relativistic electron beams in tokamaks, *Phys. Rev. Lett.* **126** 175001 (2021). doi:10.1103/PhysRevLett.126.175001

13 J R Mart\'in-Sol\'is, A Loarte and M Lehnen, Formation and termination of runaway beams in ITER disruptions, *Nucl. Fusion* **57** 066025 (2017). doi:10.1088/1741-4326/aa6939

14 O Vallhagen, O Embr\'eus, I Pusztai, L Hesslow and T F\"ul\"op, Runaway dynamics in the DT phase of ITER operations in the presence of massive material injection, *J. Plasma Phys.* **86** 475860401 (2020). doi:10.1017/S0022377820000859

15 E Berger *et al*, Runaway dynamics in reactor-scale spherical tokamak disruptions, *J. Plasma Phys.* **88** 905880611 (2022). doi:10.1017/S0022377822001209

16 M Hoppe, O Embr\'eus and T F\"ul\"op, DREAM: a fluid-kinetic framework for tokamak disruption runaway-electron simulation, *Comput. Phys. Commun.* **268** 108098 (2021). doi:10.1016/j.cpc.2021.108098

17 C R Sovinec *et al*, Nonlinear magnetohydrodynamics simulation using high-order finite elements, *J. Comput. Phys.* **195** 355 (2004). doi:10.1016/j.jcp.2003.10.004

18 N M Ferraro, S C Jardin, L L Lao, M S Shephard and F Zhang, Multi-region approach to free-boundary three-dimensional tokamak equilibria and resistive wall instabilities, *Phys. Plasmas* **23** 056114 (2016). doi:10.1063/1.4948722

19 N M Ferraro and S C Jardin, Calculations of two-fluid magnetohydrodynamic axisymmetric steady-states, *J. Comput. Phys.* **228** 7742 (2009). doi:10.1016/j.jcp.2009.07.015

20 M Hoelzl *et al*, The JOREK non-linear extended MHD code and applications to large-scale instabilities and their control in magnetically confined fusion plasmas, *Nucl. Fusion* **61** 065001 (2021). doi:10.1088/1741-4326/abf99f

21 S P Gerhardt *et al*, Disruptions, disruptivity and safer operating windows in the high-β spherical torus NSTX, *Nucl. Fusion* **53** 043020 (2013). doi:10.1088/0029-5515/53/4/043020

22 R Sweeney *et al*, MHD stability and disruptions in the SPARC tokamak, *J. Plasma Phys.* **86** 865860507 (2020).

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

Disruptions, Runaway-Electron Avalanche and Electromagnetic Loads in a Compact Spherical-Tokamak Breeder: The Design Case for Shattered-Pellet Mitigation

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

Zenodo · doi:10.5281/zenodo.22645724

Read on kronosfusionenergy.com

The Physics De-Risking Register