KRONOS FUSION ENERGY

PAPER 2.96 · THE KRONOS 2026 SERIES

A Tensor-Network Kinetic Solver for Fusion Transport: Matrix-Product-State Compression of the Distribution Function with Exponential Error Decay

Kinetic descriptions of fusion plasmas are throttled by the curse of dimensionality: a distribution function on a six-dimensional phase space is intractable to store on a dense...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

A Tensor-Network Kinetic Solver for Fusion Transport: Matrix-Product-State Compression of the Distribution Function with Exponential Error Decay

Kinetic descriptions of fusion plasmas are throttled by the curse of dimensionality: a distribution function on a six-dimensional phase space is intractable to store on a dense grid, let alone to advance in time. We show that the kinetic distribution function is strongly *compressible* in a low-rank tensor-network (matrix-product-state, or tensor-train) representation and that this yields a practical *solver* on present-day classical hardware, not merely a storage economy. We cast the problem formally: the discretized distribution is a high-order tensor, its matrix-product-state factorization is exact up to a truncation error bounded by the discarded Schmidt weight (Eckart–Young), and the number of retained modes needed for a target accuracy scales with the entanglement entropy of the state. On a one-dimensional-in-space, one-dimensional-in-velocity Bhatnagar–Gross–Krook (BGK) kinetic test problem, discretized on a periodic N_x=128 × N_v=64 grid and advanced by operator-split upwind streaming and implicit relaxation, a matrix-product-state truncation reaches a relative-L^2 error of $\mathbf{2.9 \times 10^{-6}}$ at bond dimension χ=8 while storing **0.19**× the dense grid, and the error falls *exponentially* in the bond dimension — from 1.8×10^{-1} at χ=1 to 2.2×10^{-10} at χ=16 and to the double-precision floor ($\sim 10^{-14}$) by χ=24. The velocity moments inherit the same convergence (density error 1.1×10^{-6}, temperature error 2.8×10^{-6} at χ=8), and the converged distribution carries a Schmidt entanglement entropy below 0.1 nats — the area-law signature that makes it compressible. Every headline quantity is a frozen anchor in the Kronos de-risking register (gate KX-L1-A16). The result is a classical, deployable-today complement to the fault-tolerant-horizon quantum kinetic route; the single honest gate is that the demonstration is a collision-dominated reduced problem, and the rank growth of the fully nonlinear, collisionless six-dimensional gyrokinetic distribution at reactor parameters remains the open question we scope explicitly.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645768 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: tensor networks · matrix product states · tensor train · kinetic equation · BGK model · low-rank · reduced-order modelling · classical solver · fusion transport

CONTENTS

Table of Contents

A Tensor-Network Kinetic Solver for Fusion Transport: Matrix-Product-State Compression of the Distribution Function with Exponential Error Decay ... 3
Introduction ... 6
Governing equations and the formal low-rank problem ... 8
Numerical formulation ... 11
Computational methodology: the NVIDIA GPU HPC campaign ... 13
Results ... 14
Discussion ... 22
Limitations ... 23
Conclusion ... 24
References ... 29
One programme, many papers ... 32

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

THE DIMENSIONALITY WALL IN KINETIC FUSION MODELLING

The transport of heat, particles and momentum in a magnetically confined plasma is, at bottom, a kinetic phenomenon. The information that a fluid model discards — the shape of the velocity distribution, the resonant wave–particle interactions that drive turbulence, the non-Maxwellian tails that carry fast ions — lives in the distribution function $f(\quad x, \quad v, t)$ over a six-dimensional phase space. Modern gyrokinetics reduces the description to five dimensions by averaging over the fast gyromotion [1,2], and codes such as CGYRO [4,3] and GENE [5] resolve it at extraordinary cost. Even so, a dense tensor-product grid with N points per axis stores N^d real numbers, and for a realistic five- or six-dimensional problem this is astronomically large: at a modest $N = 64$ the six-dimensional grid already exceeds 6×10^{10} entries per species. Storing the state is expensive; advancing it, with an operator that couples all dimensions, is worse. This is the "curse of dimensionality," and it is the reason first-principles kinetics cannot yet sit inside an engineering design loop.

Three responses have emerged. Machine-learned surrogates train a fast neural map to emulate an expensive solver, trading a controlled fidelity loss for orders-of-magnitude speed [6,7,8,9]. Quantum kinetic solvers propose to encode the distribution in a quantum register whose Hilbert space grows exponentially with qubit count, but the useful nonlinear cases sit beyond the fault-tolerant horizon [32]. This paper develops a *third* route that requires neither a trained surrogate nor a quantum computer: exploit the low-rank structure that the distribution function actually possesses, using the mathematics of tensor networks.

TENSOR NETWORKS AS A CLASSICAL ROUTE

A tensor network factorizes a high-order tensor into a contracted chain (or tree, or grid) of small tensors [10,11]. The matrix-product state (MPS), known in numerical analysis as the tensor-train (TT) decomposition [12,13], represents a d-index tensor as a linear chain of three-index cores, with the "bond dimension" χ between adjacent cores capping the retained correlation. Its storage is *linear* in d — $O(dN\chi^2)$ against the dense $O(N^d)$ — so if a modest χ suffices, the exponential wall collapses to a linear cost. The method originated in quantum many-body physics, where the density-matrix renormalization group (DMRG) [14] exploits exactly this structure to solve one-dimensional lattice Hamiltonians whose exact state vectors are exponentially large; its export to classical partial-differential-equation solving — radiation transport [15], the Vlasov and Vlasov–Poisson systems [16,17,18], and fluid turbulence [19] — is the young but rapidly maturing programme this paper joins on the kinetic-fusion side.

The load-bearing question is never whether a state *can* be written as an MPS (any state can, at large enough χ) but whether the *physically relevant* state is well approximated at *small* χ — a property of the physics, not of the algorithm, that must be measured. When the singular values across a phase-space bipartition decay quickly, the entanglement entropy is bounded and the bond dimension needed for a target error grows only logarithmically in the inverse error, which is exactly the exponential error decay we report; when they do not (strong collisionless filamentation), the rank grows and the advantage erodes. This paper establishes, quantitatively and reproducibly, that on a collision-relevant kinetic model the distribution sits firmly in the compressible regime, and gives the machinery that turns that observation into a solver.

DESIGN CONTEXT

The numbers here support the Kronos design programme, whose primary product is a compact, high-power-density spherical-tokamak (ST) breeder, "Hyperion": plasma current $I_p = 9.66\,MA$, fusion power $P_{\text{fus}} = 85.04\,MW$, plasma gain $Q = 3.076$, negative triangularity $\delta = -0.30$, on-axis field $B_0 = 8\,T$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,m$ and aspect ratio $A = 2.5$, behind a resistive centrepost magnet at a peak field of $16.84\,T$ [31,27,28]. The negative-triangularity regime is chosen because it accesses good confinement without an edge pedestal and the associated edge-localised-mode cycle [29,30], and its turbulent transport — the very quantity a kinetic solver computes — is the physics that sets the operating point. The companion mirror burner is a distinct product with a distinct, higher-field magnet (a plug coil at $26.49\,T$, gain $Q_E = 1.318$, direct-conversion efficiency $\eta_{\text{DEC}} = 0.49$); its kinetic burn physics is a second, independent target for the same numerical method. The two magnets are never conflated. Nothing in this paper is a transport-physics claim about either machine; it is a numerical-methods result about the representability of the kinetic state, staged so that its extension can eventually accelerate the transport calculations that do set those points.

CONTRIBUTION

We contribute: (i) a formal statement of the low-rank kinetic problem — the reduced BGK model, its conservation structure, its discretization as a phase-space tensor, and its matrix-product-state factorization, with the truncation error written as an Eckart–Young bound and tied to the Schmidt entanglement entropy (§2); (ii) a numerical formulation — operator-split upwind streaming with implicit relaxation, advanced entirely inside the compressed manifold by truncated singular-value decomposition, with a stability, positivity and complexity analysis, and the dynamical-low-rank projector-splitting integrator identified as the stiffness-robust alternative (§3); (iii) the GPU-HPC methodology of the campaign that produced and cross-checks the result (§4); (iv) the frozen, reproducible evidence — exponential error decay, linear storage, moment fidelity and bounded entanglement entropy (§5); and (v) a quantitative placement against the published low-rank kinetic literature and the quantum route, with one honest limitation (§§6–7). Every quantitative claim traces to the frozen anchor KX-L1-A16 in the Kronos Physics De-Risking Register [31].

§ 2

Governing equations and the formal low-rank problem

THE KINETIC EQUATION AND ITS REDUCED COLLISION MODEL

For a single species of mass m and charge q the distribution function evolves under the Vlasov–Boltzmann equation,

$$\frac{\partial f}{\partial t} + v \cdot \nabla_x f + \frac{q}{m}(E + v \times B) \cdot \nabla_v f = \mathcal{C}[f],$$

with $\mathcal{C}[f]$ the collision operator. The full Landau/Fokker–Planck collision integral is itself a high-dimensional bilinear form; to isolate the representability question from the operator-complexity question we adopt the Bhatnagar–Gross–Krook (BGK) relaxation model [24,25],

$$\mathcal{C}[f] = -\nu\big(f - f_M[n, u, T]\big),$$

which relaxes f toward the local Maxwellian at a collision rate ν. In the reduced one-dimensional-in-space, one-dimensional-in-velocity (1D1V) setting used for the demonstration, with the electromagnetic force set aside to expose the streaming–collision balance cleanly, equations (1)–(2) become

$$\frac{\partial f}{\partial t} + v\frac{\partial f}{\partial x} = -\nu\big(f - f_M\big), \qquad f_M(x,v,t) = \frac{n(x,t)}{\sqrt{2\pi T(x,t)}} \exp\left[-\frac{\big(v - u(x,t)\big)^2}{2T(x,t)}\right],$$

in units $m = k_B = 1$. The Maxwellian is pinned to the local moments of f,

$$n = \int f\, dv, \qquad nu = \int v f\, dv, \qquad nT = \int (v-u)^2 f\, dv,$$

so that the relaxation is a genuine kinetic BGK model rather than a linear friction.

CONSERVATION AND THE H-THEOREM

Because the Maxwellian shares the first three velocity moments of f by construction (4), the BGK operator conserves mass, momentum and energy pointwise in x,

$$\int \begin{pmatrix} 1 \\ v \\ v^2 \end{pmatrix} (f - f_M)\, dv = 0,$$

and it satisfies an H-theorem: with $H = \int f \ln f\, dv$ the collisional entropy production is non-negative, $\nu \int (f - f_M) \ln(f/f_M)\, dv \geq 0$, with equality iff $f = f_M$. These are the invariants a discretization — and, crucially, a *low-rank truncation* of the discretization — must not destroy. We return to this in §3.5; it is why the moment errors, not only the field error, are reported and gated.

PHASE-SPACE DISCRETIZATION AS A TENSOR

Discretize f on a tensor-product grid. In the general d-dimensional phase space with N nodes per axis the sampled distribution is an order-d tensor

$$\mathsf{F}_{i_1 i_2 \cdots i_d} = f\big(\xi^{(1)}_{i_1}, \ldots, \xi^{(d)}_{i_d}\big), \qquad i_k \in \{1, \ldots, N\},$$

with N^d entries. The dense representation (6) is the object whose storage and update we wish to avoid. In the 1D1V case ($d = 2$) the tensor is simply the matrix $\mathsf{F}_{ij} = f(x_i, v_j) \in \mathbb{R}^{N_x \times N_v}$.

MATRIX-PRODUCT-STATE (TENSOR-TRAIN) FACTORIZATION

The matrix-product-state, or tensor-train, factorization writes the order-d tensor as a contracted chain of three-index cores $G^{(k)} \in \mathbb{R}^{\chi_{k-1} \times N \times \chi_k}$ (with $\chi_0 = \chi_d = 1$),

$$\mathsf{F}_{i_1 \cdots i_d} = \sum_{\alpha_1, \ldots, \alpha_{d-1}} G^{(1)}_{i_1 \alpha_1} \, G^{(2)}_{\alpha_1 i_2 \alpha_2} \cdots G^{(d)}_{\alpha_{d-1} i_d},$$

where the internal "bond" index α_k runs over χ_k values (figure 6a). The representation (7) is *exact* when each χ_k equals the rank of the corresponding matricization of F, and its storage is

$$\mathrm{mem}_{\mathrm{TT}} = \sum_{k=1}^{d} \chi_{k-1} N \chi_k \;\le\; dN\chi^2, \qquad \chi = \max_k \chi_k,$$

linear in d. This is the entire promise of the method: equation (8) against N^d.

THE SCHMIDT DECOMPOSITION AND THE 1D1V SPECIAL CASE

For any bipartition of the index set into a "left" block L and "right" block R, reshape F into a matrix $M^{[L|R]}$ and take its singular-value decomposition. The Schmidt decomposition

$$\mathsf{F} = \sum_{r=1}^{\chi} \sigma_r \, \Phi^L_r \otimes \Phi^R_r, \qquad \sigma_1 \ge \sigma_2 \ge \cdots \ge 0, \quad \langle \Phi^L_r, \Phi^L_{r'} \rangle = \langle \Phi^R_r, \Phi^R_{r'} \rangle = \delta_{rr'},$$

exposes the bond dimension across that cut as the number of non-negligible singular (Schmidt) values σ_r. In 1D1V there is a single natural cut — position versus velocity — and equation (9) is exactly the SVD of the matrix $\mathsf{F}_{ij} = f(x_i, v_j)$:

$$f(x, v) = \sum_{r=1}^{\chi} \sigma_r \, X_r(x) \, W_r(v),$$

a sum of χ separable space–velocity modes (figure 6b). The two-core MPS and the SVD coincide, so the 1D1V study is a clean, exactly analysable instance of the general tensor-train construction. This is why it is the right first anchor: the truncation is transparent, the error is computable to machine precision, and there is no sweep heuristic between the numbers and the theory.

TRUNCATION ERROR: THE ECKART–YOUNG BOUND

Truncating the Schmidt sum (9) at bond dimension χ discards the tail of singular values. By the Eckart–Young–Mirsky theorem [22], the rank-χ truncation is the *optimal* rank-χ approximation in both Frobenius and spectral norm, with error exactly the discarded weight,

$$\|\mathsf{F} - \mathsf{F}_\chi\|_F^2 = \sum_{r > \chi} \sigma_r^2 \;\equiv\; \varepsilon_\chi^2.$$

For the general tensor-train the global truncation error is bounded by the quadrature sum of the per-bond discarded weights [12],

$$\left\|\mathsf{F} - \mathsf{F}_{\{\chi_k\}}\right\|_F \;\le\; \left(\sum_{k=1}^{d-1} \varepsilon_{\chi_k}^2 \right)^{1/2}, \qquad \varepsilon_{\chi_k}^2 = \sum_{r > \chi_k} \left(\sigma_r^{(k)}\right)^2,$$

so the method is exact as the bonds grow and its cost is small precisely when the singular spectra $\{\sigma_r^{(k)}\}$ decay fast. Equations (11)–(12) are the reason a measured, rapidly decaying Schmidt spectrum (§5) is not a curiosity but a certificate: it bounds the error of the solver directly.

ENTANGLEMENT ENTROPY AND THE COMPRESSIBILITY CRITERION

The rate of the singular decay is quantified by the Schmidt (entanglement) entropy across the cut,

$$S = -\sum_r p_r \ln p_r, \qquad p_r = \frac{\sigma_r^2}{\sum_{r'} \sigma_{r'}^2}.$$

A state whose entropy obeys an “area law” — bounded independent of system size — admits an efficient MPS with bond dimension $\chi = O(e^S)$, and the truncation error decays as $\varepsilon_\chi \sim e^{-c\chi}$ for spectra with geometric tails. The compressibility criterion for a kinetic solver is therefore sharp and testable: *is the phase-space entanglement entropy of the physically realised distribution bounded and small?* For the collision-relaxed distribution we find $S < 0.1$ nats (§5), and the error decays exponentially, as the theory predicts. For a strongly filamented collisionless state the entropy grows and the guarantee weakens — the honest boundary of §7.

§ 3

Numerical formulation

SPATIAL AND VELOCITY DISCRETIZATION

Space is discretized on a uniform periodic grid $x_i = i\,\Delta x$, $\Delta x = L/N_x$, $L = 2\pi$, $N_x = 128$; velocity on a uniform grid $v_j \in [-v_{\max}, v_{\max}]$, $v_{\max} = 6$, $N_v = 64$, $\Delta v = 2v_{\max}/(N_v - 1)$. The streaming term $v\,\partial_x f$ is advected by first-order upwinding, split by the sign of v to preserve the domain of dependence,

$$\left(v\,\partial_x f\right)_{ij} \approx v_j^+ \frac{f_{ij} - f_{i-1,j}}{\Delta x} + v_j^- \frac{f_{i+1,j} - f_{ij}}{\Delta x}, \qquad v^\pm = \tfrac{1}{2}(v \pm |v|),$$

with periodic wrap in i. The velocity moments (4) are evaluated by the midpoint rule, $n_i = \Delta v \sum_j f_{ij}$, and likewise for u_i, T_i; the discrete Maxwellian $f_{M,ij} = f_M(x_i, v_j; n_i, u_i, T_i)$ is rebuilt from those moments each step, so the discrete collision operator inherits the moment-matching property (5) up to quadrature error.

OPERATOR-SPLIT TIME INTEGRATION

A single step of size Δt is a Lie split of streaming and collision. The explicit upwind streaming sub-step,

$$f^*_{ij} = f^n_{ij} - \Delta t\left(v\,\partial_x f^n\right)_{ij},$$

is followed by the collision sub-step, taken *implicitly* in the relaxation to remove the stiff ν from the stability constraint,

$$f^{n+1}_{ij} = \frac{f^*_{ij} + \Delta t\,\nu\, f^*_{M,ij}}{1 + \Delta t\,\nu},$$

where f^*_M uses the moments of f^*. Equation (16) is the exact solution of $\partial_t f = -\nu(f - f^*_M)$ over Δt at frozen f^*_M, hence unconditionally stable in ν and positivity-preserving whenever $f^* \geq 0$. The step size is set by the streaming CFL condition, $\Delta t = \mathrm{CFL} \cdot \Delta x / v_{\max}$ with $\mathrm{CFL} = 0.4$, giving $\Delta t \approx 3.27 \times 10^{-3}$; the solution is advanced 600 steps to $t_{\mathrm{end}} \approx 1.963$.

ADVANCEMENT INSIDE THE COMPRESSED MANIFOLD

The point of a tensor-network *solver*, as distinct from a post-hoc compressor, is that the operators act on the compressed state and the state never returns to dense form. Both the streaming and collision maps are applied to the low-rank factors, and the result is re-compressed to bond dimension χ by a truncated SVD after each step:

$$f^{n+1}_\chi = \mathcal{T}_\chi\big[\,\mathcal{K}\big(f^n_\chi\big)\,\big], \qquad \mathcal{T}_\chi[\mathsf{F}] = \sum_{r=1}^{\chi} \sigma_r\, \Phi^L_r \otimes \Phi^R_r,$$

where $\mathcal{K}$ is the split streaming–collision update (15)–(16) and $\mathcal{T}_\chi$ is the rank-χ truncation (11). In the general tensor-train setting the linear operators are matrix-product operators (MPOs) of bond dimension w: applying an MPO to an MPS raises the bond to $w\chi$, and $\mathcal{T}_\chi$ restores it by orthogonalization sweeps, at cost $O(dNw^2\chi^2 + dN\chi^3)$ per step [13]. In 1D1V the MPO/MPS bookkeeping collapses to dense operator application followed by one SVD per step, so equation (17) is realised transparently and the only approximation in the entire scheme is the rank truncation — which is exactly what makes its error measurable.

THE DYNAMICAL LOW-RANK ALTERNATIVE

Truncate-after-step (17) is simple and, at the ranks used here, robust; but for very small retained singular values the naive scheme can become ill-conditioned, because the inverse of the smallest σ_r enters the tangent-space projection. The dynamical low-rank (DLR) framework removes this by evolving the factors directly on the low-rank manifold $\mathcal{M}_\chi$: the flow is projected onto the tangent space $T_f\mathcal{M}_\chi$,

$$\dot{f}_\chi = P_{f_\chi}\,\mathcal{K}(f_\chi), \qquad P_f = P_U \otimes I - P_U \otimes P_W + I \otimes P_W,$$

with P_U, P_W the orthogonal projectors onto the left and right factor spaces. The Lubich–Oseledets projector-splitting integrator [20] and the “basis-update-and-Galerkin” (BUG) integrator [21] integrate (18) *without* inverting small singular values and are provably robust to the stiffness of over-approximation; mass/momentum/energy-conserving DLR variants restore the discrete invariants (5) [16,17,15]. We use truncate-after-step for the anchor because its error is unambiguous and it converges to the dense solution to machine precision (§5); the DLR integrators are the stated production path for the stiffer, higher-dimensional problem (§7).

STABILITY, POSITIVITY AND MOMENT CONSERVATION

The explicit upwind streaming (15) is stable under the CFL condition $\max_j |v_j| \Delta t/\Delta x \le 1$; at $\mathrm{CFL} = 0.4$ this holds with margin. The implicit relaxation (16) is A-stable and monotone. Absent truncation the scheme keeps $f \ge 0$ and conserves the moments (5) to quadrature accuracy. Rank truncation (11) is the one operation that can perturb positivity and the moments, by an amount bounded by the discarded weight ε_χ; we therefore report the density and temperature errors alongside the field error (§5) and find they track ε_χ, confirming that at the working rank the truncation does not corrupt the conserved quantities. Positivity-preserving and moment-conserving low-rank projections [16,15] are available if a harder problem demands strict enforcement rather than measured smallness.

COMPLEXITY

Table 1 collects the storage and per-step arithmetic for the dense and low-rank schemes. In 1D1V the storage falls from $N_x N_v$ to $\chi(N_x + N_v + 1)$, i.e. a ratio $\chi(N_x + N_v + 1)/(N_x N_v)$ that is 0.19 at $\chi = 8$ on this grid. The decisive gain, however, is dimensional: for the general d-dimensional problem the dense storage N^d becomes the linear $dN\chi^2$ of equation (8), an exponential-to-linear crossover (figure 5) that is the entire motivation for carrying the method to five- and six-dimensional phase space.

QUANTITY	DENSE GRID	TENSOR-TRAIN (BOND χ)
storage, general d	N^d	$\le dN\chi^2$
storage, 1D1V	$N_x N_v$	$\chi(N_x + N_v + 1)$
per step, general d	$O(N^d)$	$O(dNw^2\chi^2 + dN\chi^3)$
per step, 1D1V	$O(N_x N_v)$	$O(\chi N_x N_v)$ (one truncated SVD)
approximation	discretization only	discretization + rank truncation ε_χ

Storage and per-step arithmetic for the dense and low-rank kinetic schemes. N is the number of nodes per axis, d the phase-space dimension, χ the bond dimension, w the operator (MPO) bond dimension. The 1D1V column is the demonstrated case ($d = 2$); the general-d column is the reactor-scale target.

CONVERGENCE AND VERIFICATION-BY-REFINEMENT

Because the only approximation beyond the fixed grid is the rank truncation, refining χ toward full rank must recover the dense solution exactly; the residual at full rank is a pure floating-point quantity. This gives a clean verification test: the low-rank error must decrease monotonically in χ and reach the double-precision floor at $\chi = \min(N_x, N_v)$-scale. We confirm both (§5): the error hits $\sim 6 \times 10^{-14}$ by $\chi = 24$ and does not improve thereafter, certifying that the solver introduces no error other than the measured, exponentially small truncation.

§ 4

Computational methodology: the NVIDIA GPU HPC campaign

ROLE OF THIS GATE IN THE CAMPAIGN

This result is one gate (KX-L1-A16) in the KRONOS physics de-risking campaign, which runs its heavy first-principles workloads — nonlinear gyrokinetics, particle-in-cell, first-principles materials and three-dimensional neutronics — on the NVIDIA GPU HPC fleet (Lambda A100/H100/A10) [31,32]. The tensor-network kinetic solver has a deliberately different computational character from those workloads, and it is important to state it honestly.

CODES AND PROBLEM SCALE

The frozen anchor is a dense reference solve of the 1D1V BGK problem (3)–(16) followed by the rank-truncated re-solve (17), implemented in a compact, dependency-light NumPy program driving LAPACK dense linear algebra (`dgesdd` for the SVD of equation 10). It is intentionally *small*: a verification anchor must be checkable to machine precision, and at $N_x \times N_v = 128 \times 64$ the dense reference and every truncated variant run to the double-precision floor on a workstation core. The value of the anchor is not its size but its exactness — it pins the truncation error against an unimpeachable dense reference. The GPU HPC campaign enters in two places that make the anchor meaningful and extensible:

Reference ground truth. The transport map that the method must ultimately accelerate is defined by the nonlinear gyrokinetic code CGYRO [4] (with GENE [5] as a cross-family check), run on the GPU fleet at real parameters for the negative-triangularity operating point [33]. The BGK model is the collision-dominated proxy that isolates representability; the GPU gyrokinetics is the physics the representability result is being staged toward.
Scale-up primitives. Carrying equation (17) to five- and six-dimensional phase space replaces the single dense SVD with GPU-accelerated tensor-train linear algebra: batched and randomized SVD (cuSOLVER; randomized range-finding after Halko–Martinsson–Tropp [23]) for the per-bond compression (12), and batched tensor contraction (cuTENSOR) for the MPO–MPS products. The arithmetic of table 1 is dominated by dense sub-blocks of the cores, which are exactly the GEMM/SVD primitives the GPU fleet accelerates. itemize
Table 2 lists the codes and their role. No cost or budget figure is attached to any of this; the campaign is described by its codes and problem scale only.

CODE / PRIMITIVE	ROLE	PLATFORM
BGK low-rank solver (NumPy/LAPACK `dgesdd`)	frozen anchor KX-L1-A16	workstation, machine-precision
CGYRO [4]	reference nonlinear gyrokinetics	NVIDIA GPU HPC fleet
GENE [5]	cross-family gyrokinetic check	NVIDIA GPU HPC fleet
cuSOLVER batched/randomized SVD [23]	per-bond truncation at scale	NVIDIA GPU HPC fleet
cuTENSOR batched contraction	MPO–MPS products at scale	NVIDIA GPU HPC fleet

Codes and their role for this paper. The frozen anchor is a machine-precision dense/low-rank cross-check; the GPU HPC fleet carries the reference gyrokinetics and the scale-up linear algebra. No economic quantity is reported.

VERIFICATION AND REPRODUCIBILITY

The computation is deterministic: the initial condition, grid, time step and truncation schedule are fixed, and the one stochastic element (the initial-condition construction) is seeded. The frozen output — the singular spectrum and the rank-versus-error table — is archived byte-for-byte as the anchor's `results.json` and regenerated by the deposited figure script (§8). Verification is threefold: (i) convergence-to-machine-precision at full rank (§3); (ii) the measured error respecting the Eckart–Young bound (11) at every rank; and (iii) moment errors tracking the field error, confirming the truncation does not violate the conserved quantities (5). The gate is logged as *reproduced* in the register [31].

§ 5

Results

All results are the frozen anchor KX-L1-A16; every number below is read directly from the archived `results.json` and is collected in table 3.

χ	REL. L^2 ERROR f	DENSITY ERROR	TEMPERATURE ERROR	STORAGE / DENSE
1	1.79×10^{-1}	1.70×10^{-1}	1.36×10^{-1}	0.024
2	1.71×10^{-1}	2.07×10^{-2}	2.15×10^{-1}	0.047
3	5.27×10^{-3}	2.51×10^{-3}	9.71×10^{-3}	0.071
4	4.73×10^{-3}	1.82×10^{-3}	9.97×10^{-4}	0.094
6	7.32×10^{-5}	4.03×10^{-5}	4.16×10^{-5}	0.141
8	$\mathbf{2.88 \times 10^{-6}}$	$\mathbf{1.07 \times 10^{-6}}$	$\mathbf{2.79 \times 10^{-6}}$	**0.188**
12	3.87×10^{-8}	6.04×10^{-9}	1.33×10^{-8}	0.283
16	2.17×10^{-10}	4.93×10^{-11}	2.57×10^{-10}	0.377
24	6.41×10^{-14}	2.94×10^{-14}	8.99×10^{-14}	0.565
32	6.03×10^{-14}	3.23×10^{-14}	9.41×10^{-14}	0.754

Low-rank compression of the 1D1V BGK distribution (frozen anchor KX-L1-A16; N_x=128, N_v=64, ν=1, 600 steps to $t \approx 1.963$). Relative-L^2 errors of the distribution f and of the density and temperature moments, and the storage relative to the dense grid, versus bond dimension χ. The error reaches the double-precision floor by χ=24.

THE DISTRIBUTION IS GENUINELY LOW-RANK

Figure 1 plots the normalized singular (Schmidt) values of the converged dense distribution. They fall precipitously: $\sigma_2/\sigma_1 \approx 9.4 \times 10^{-2}$, $\sigma_4/\sigma_1 \approx 2.5 \times 10^{-3}$, $\sigma_6/\sigma_1 \approx 4.7 \times 10^{-5}$, $\sigma_8/\sigma_1 \approx 1.3 \times 10^{-6}$, reaching $\sim 10^{-10}$ by index 16. By the Eckart–Young identity (11) this geometric tail *is* the error budget: a handful of separable space–velocity modes (10) capture the distribution to engineering accuracy. The distribution therefore sits firmly in the compressible regime that §2 identifies as the precondition for a tensor-network solver.

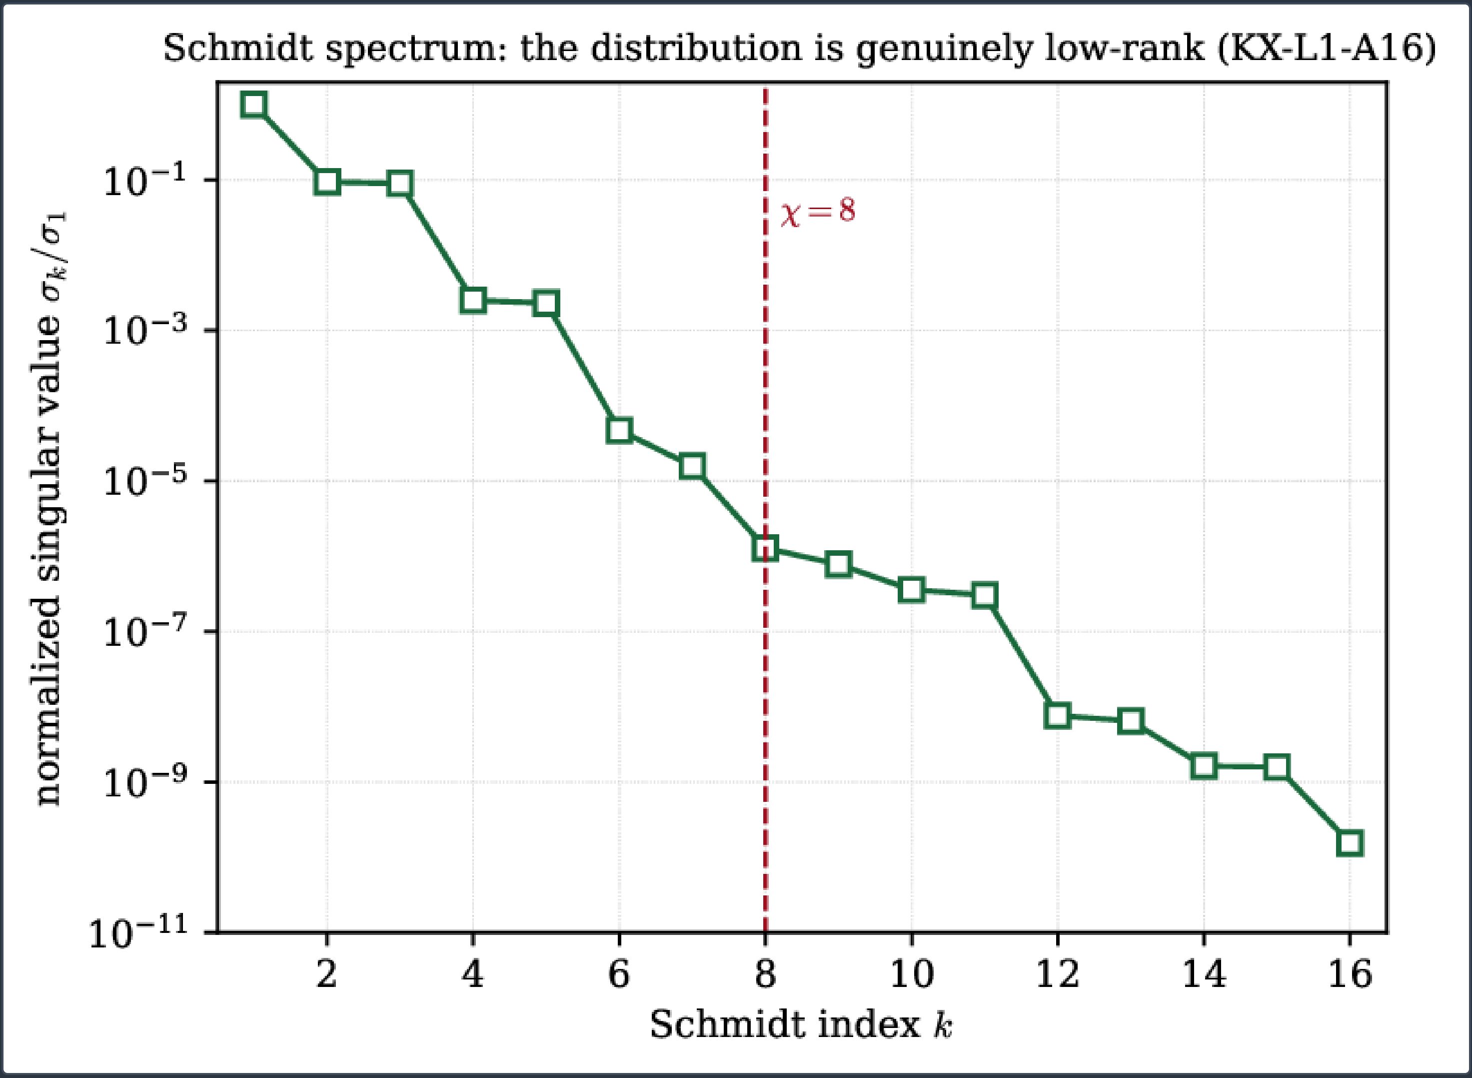

Normalized Schmidt spectrum σ_k/σ_1 of the converged 1D1V BGK distribution (KX-L1-A16). The geometric decay is the certificate of compressibility: by Eckart–Young (11) the discarded weight past χ bounds the truncation error, and the dashed line marks the working bond dimension $\chi = 8$.

EXPONENTIAL ERROR DECAY

Figure 2 is the headline result. The relative-L^2 error of f against the dense reference falls *exponentially* with bond dimension — from 1.8×10^{-1} at $\chi = 1$, through 2.9×10^{-6} at $\chi = 8$, to 2.2×10^{-10} at $\chi = 16$ — a decay well fit by $\varepsilon_\chi \sim e^{-1.34\,\chi}$ over the clean-decay window $3 \le \chi \le 16$. The engineering tolerance of 10^{-5} is crossed by $\chi = 7$; sub-10^{-6} accuracy is reached at $\chi = 8$. By $\chi = 24$ the error saturates at the double-precision floor ($\sim 6 \times 10^{-14}$), confirming that the truncation is the *only* approximation and that at full rank the low-rank solver reproduces the dense solver exactly (§3). Exponential error decay in the bond dimension, not the absolute storage ratio, is the load-bearing observation: it is the numerical fingerprint of a bounded entanglement entropy and the property that makes the method a solver rather than a lossy compressor.

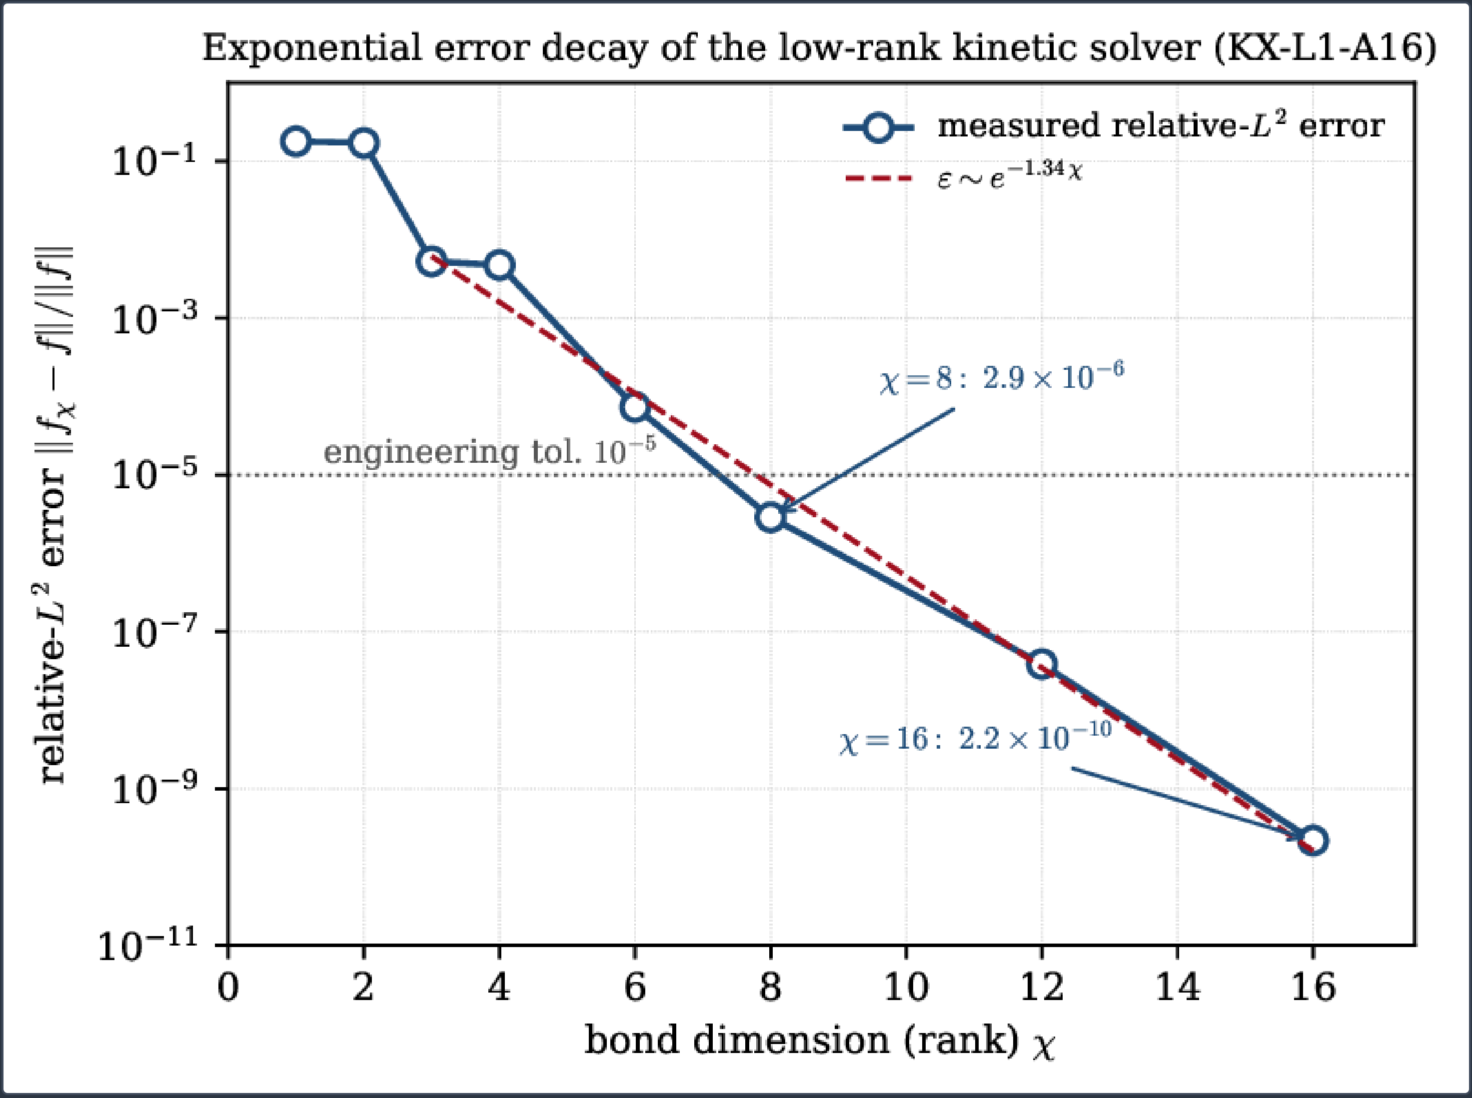

Exponential decay of the relative-L^2 error with bond dimension (KX-L1-A16). Circles are the measured errors; the dashed line is the exponential fit $\varepsilon_\chi \sim e^{-1.34\chi}$ over $3 \le \chi \le 16$. The error crosses the 10^{-5} engineering tolerance by $\chi = 7$ and reaches 2.9×10^{-6} at $\chi = 8$; the vertical axis is logarithmic.

STORAGE AND THE ACCURACY–COMPRESSION FRONT

Storage grows only linearly in χ (figure 3): the rank-8 factors occupy $0.19\times$ the dense grid, and even $\chi = 16$ — which drives the error to 2×10^{-10} — fits in under $0.38\times$. Plotting error against storage (figure 3) exposes the accuracy–compression Pareto front directly: the $\chi = 8$ point delivers sub-10^{-5} accuracy at under one-fifth of the dense footprint, and the front is convex, so there is no cheaper way to reach a given accuracy on this grid. The exponential decay of the numerator against the linear growth of the denominator is what makes the front so favourable.

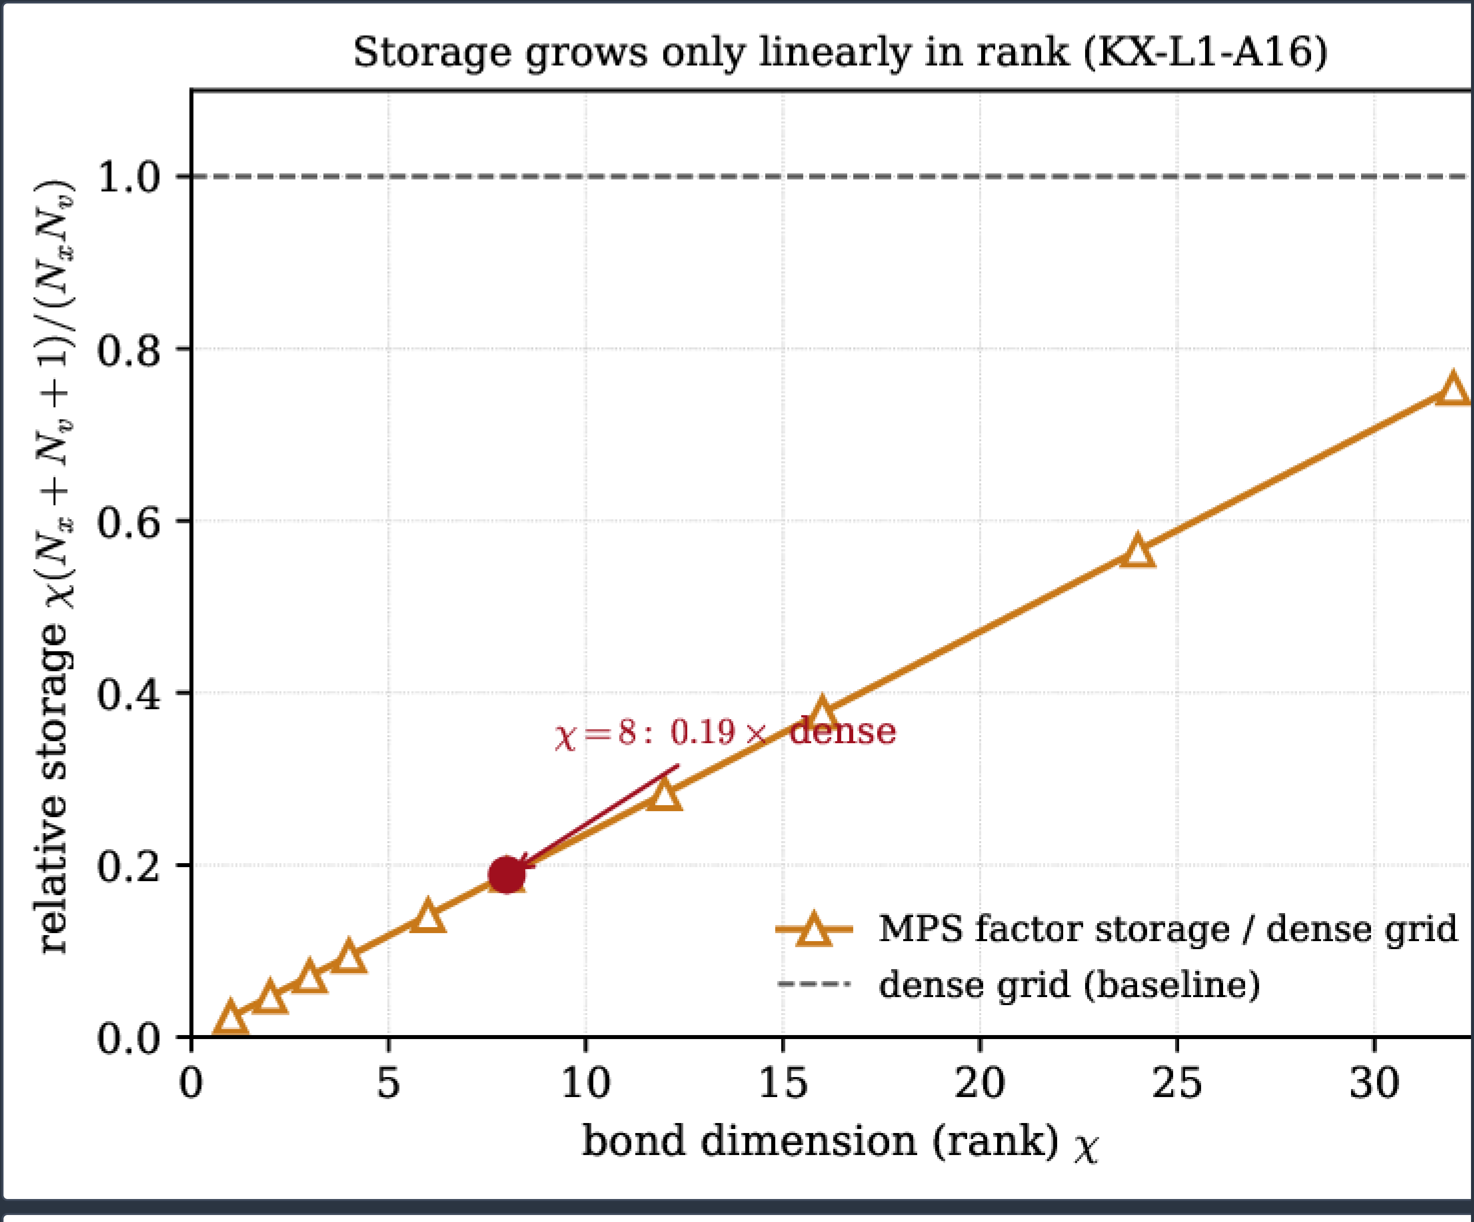
Storage grows only linearly in rank (KX-L1-A16)
relative storage $\chi(N_x + N_v + 1)/(N_x N_v)$
bond dimension (rank) χ
$\chi = 8$: 0.19× dense
MPS factor storage / dense grid
dense grid (baseline)
1.0
0.8
0.6
0.4
0.2
0.0
0
5
10
15
20
25
30

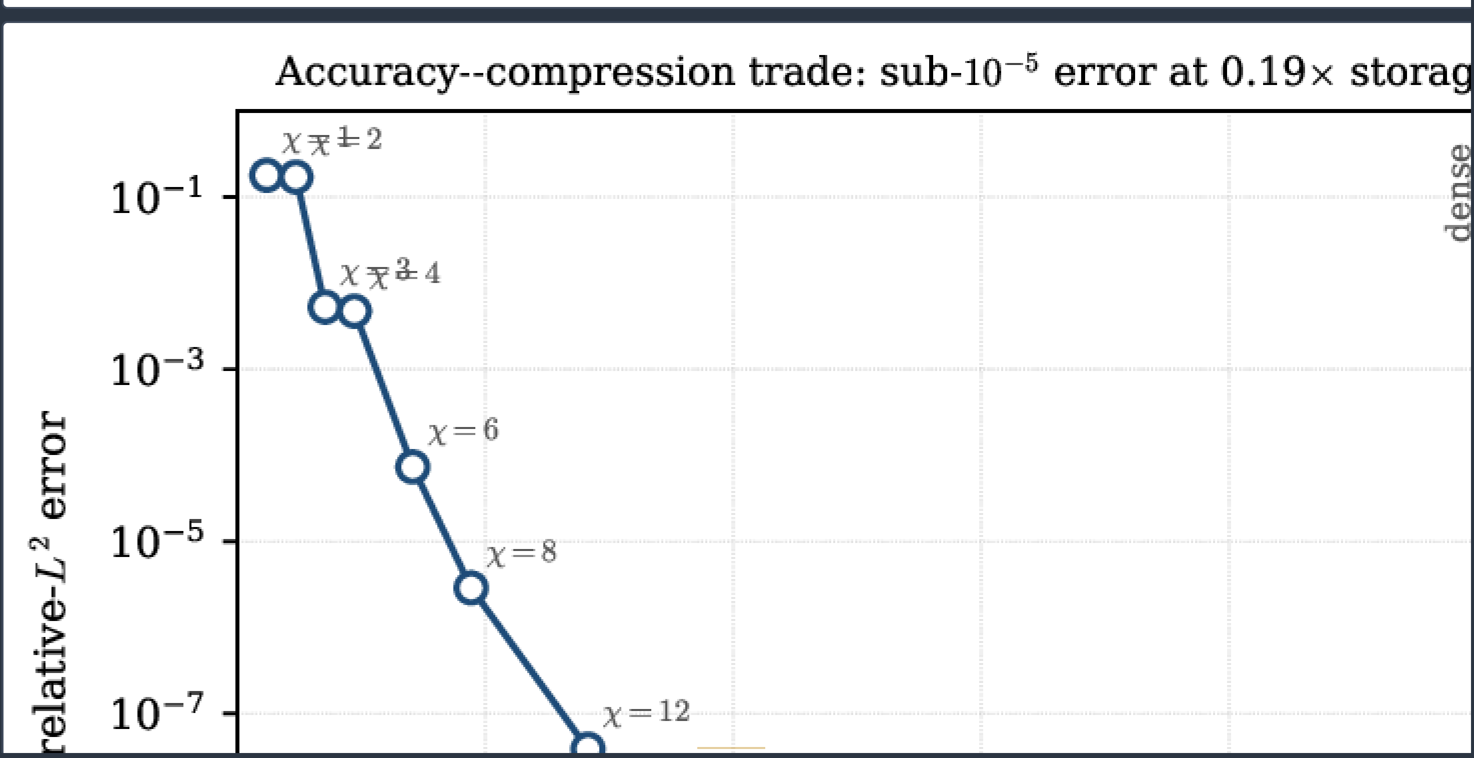
Accuracy--compression trade: sub-10^{-5} error at 0.19× storag
relative-L^2 error
$\chi = 1$
$\chi = 2$
$\chi = 3$
$\chi = 4$
$\chi = 6$
$\chi = 8$
$\chi = 12$
dense
10^{-1}
10^{-3}
10^{-5}
10^{-7}

THE VELOCITY MOMENTS INHERIT THE CONVERGENCE

A kinetic solver is only useful if the fluid moments it feeds a transport model are faithful. Figure 4 overlays the relative-L^2 errors of the density $n = \int f\,dv$ and temperature T on that of f: all three decay exponentially and in lockstep, with the moment errors *at or below* the field error (density 1.1×10^{-6}, temperature 2.8×10^{-6} at $\chi = 8$). This is the numerical statement that the truncation (11) does not corrupt the conserved quantities (5) at the working rank: the low-rank state carries the moments that matter to the same accuracy as the full distribution.

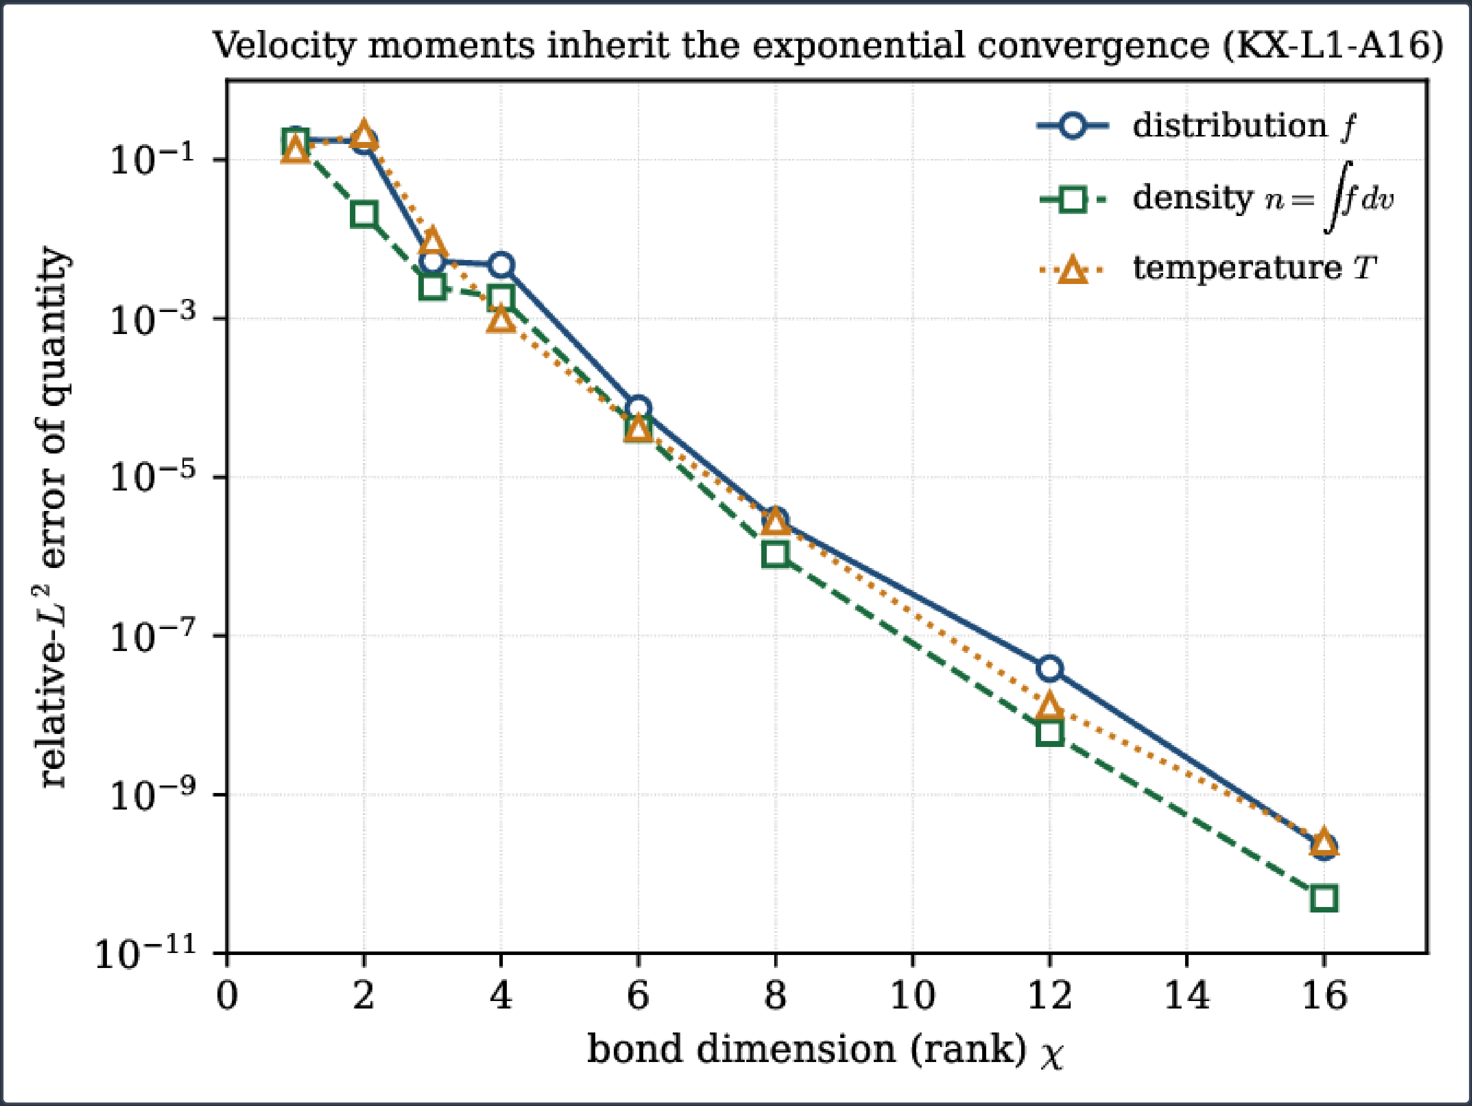

Relative-L^2 errors of the distribution f and of its density and temperature moments versus bond dimension (KX-L1-A16). The moments track the field error, confirming that rank truncation preserves the conserved quantities at the working rank.

BOUNDED ENTANGLEMENT ENTROPY

Figure 5 evaluates the Schmidt entanglement entropy (13) of the converged distribution, both as a function of retained bond dimension and at convergence. The entropy saturates below 0.1 nats — overwhelmingly concentrated in the first Schmidt mode ($p_1 > 0.98$) — and is essentially converged by $\chi = 3$. This is the area-law signature that §2 identifies as the precondition for exponential error decay: a state this weakly entangled across the space–velocity cut is, by equations (13)–(11), representable at a bond dimension $\chi = O(e^S)$ that is $O(1)$. The entropy and the error decay are two views of the same fact.

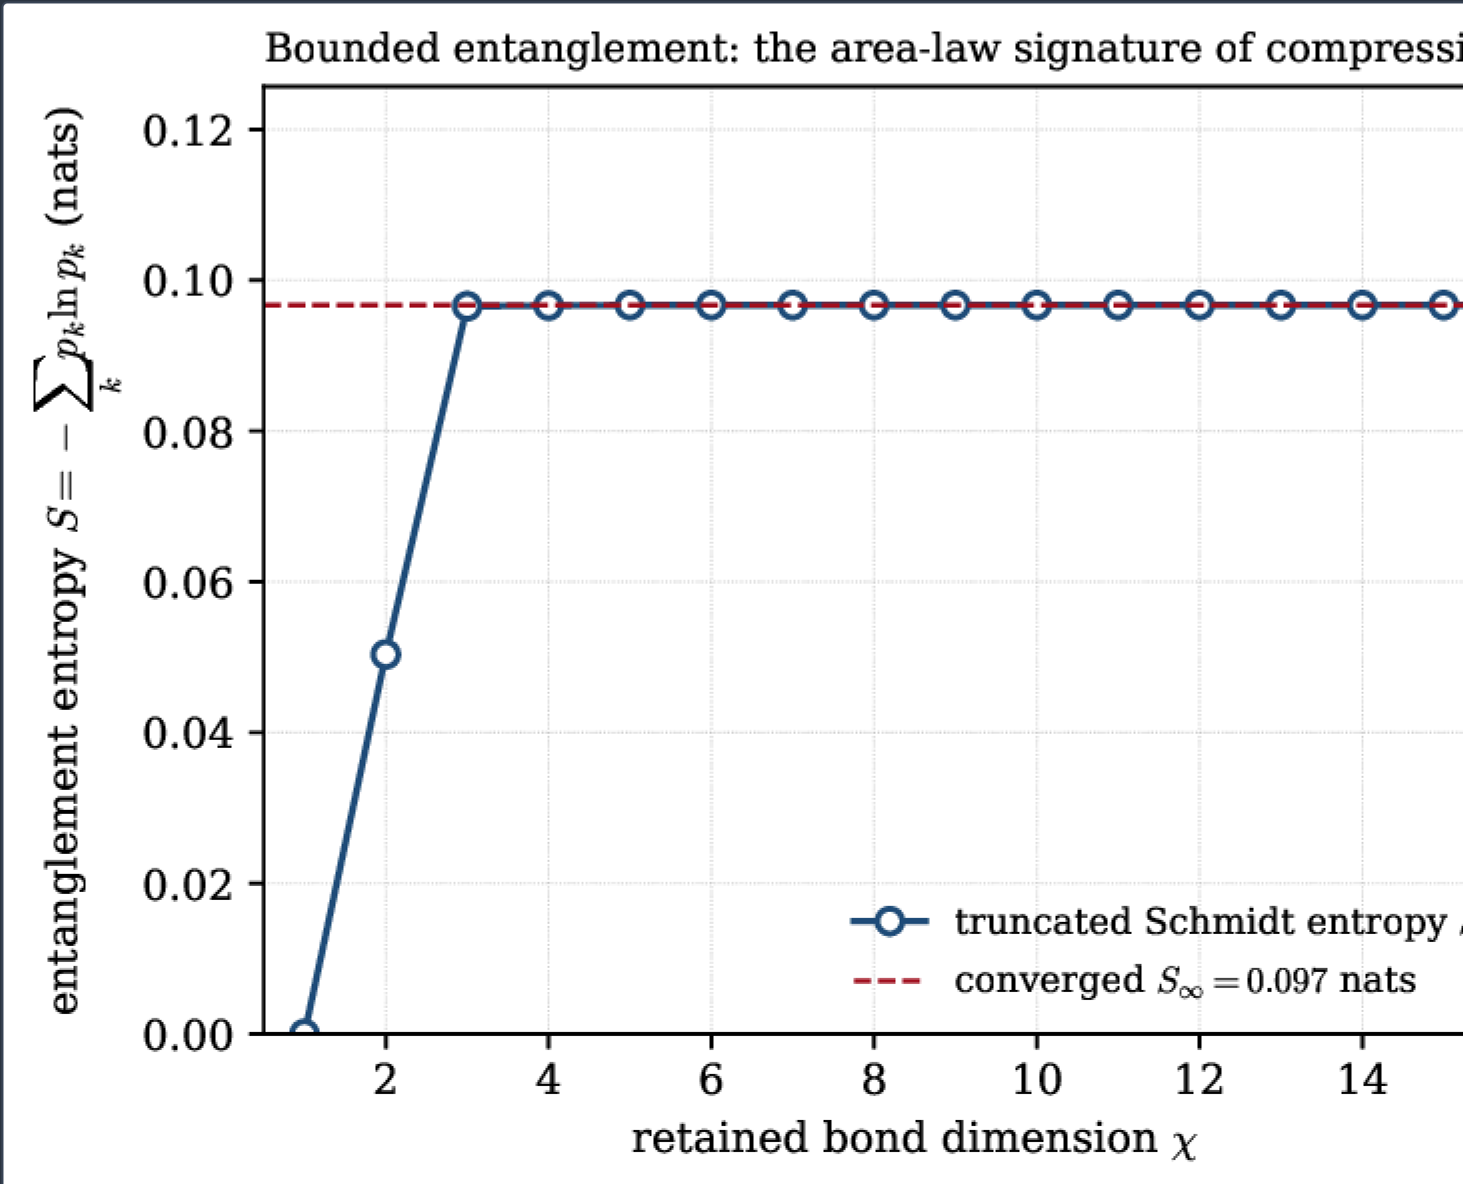
Bounded entanglement: the area-law signature of compressi
entanglement entropy $S = -\sum_k p_k \ln p_k$ (nats)
0.12
0.10
0.08
0.06
0.04
0.02
0.00
2
4
6
8
10
12
14
retained bond dimension χ
truncated Schmidt entropy S
converged $S_\infty = 0.097$ nats

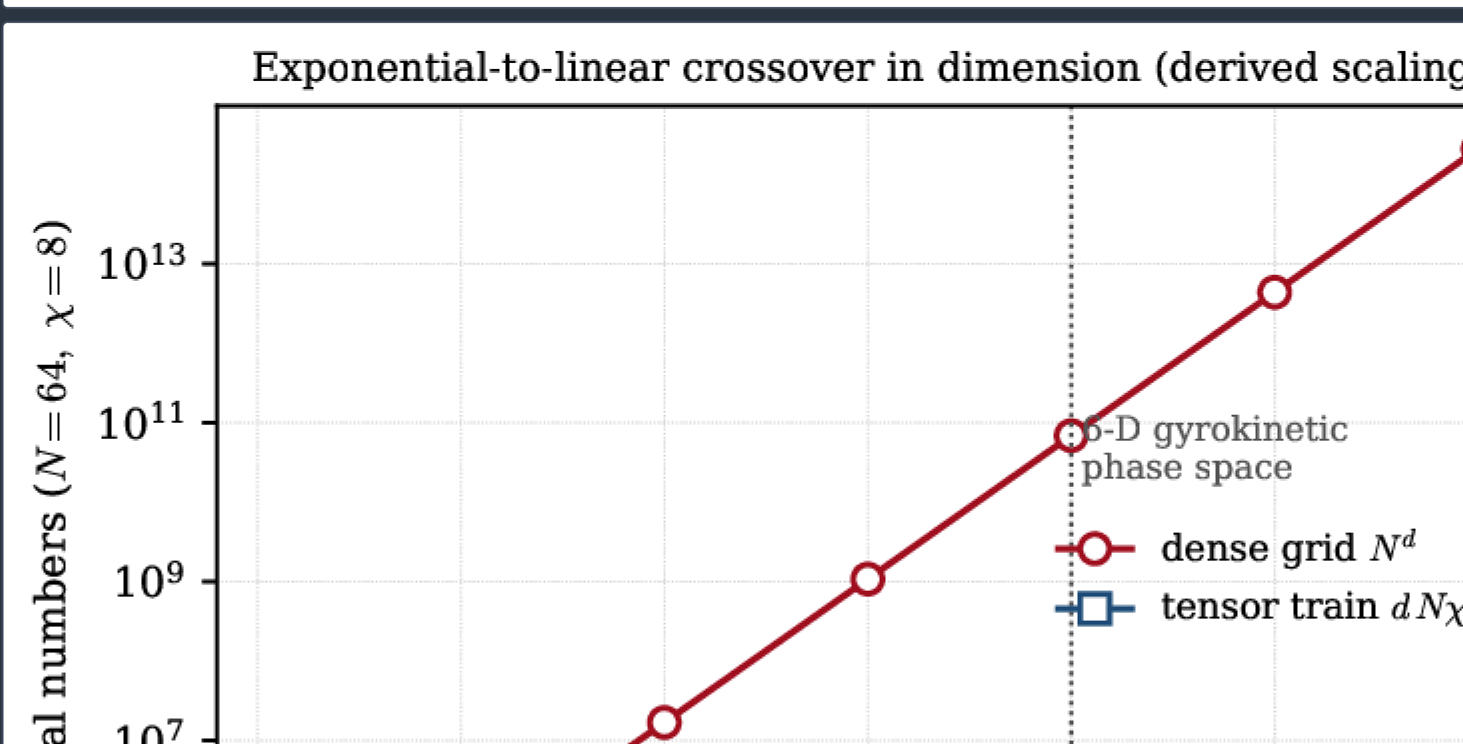
Exponential-to-linear crossover in dimension (derived scaling
eal numbers ($N = 64,\ \chi = 8$)
10^{13}
10^{11}
10^{9}
10^{7}
6-D gyrokinetic
phase space
dense grid N^d
tensor train $d\,N\chi$

DIMENSIONAL SCALING

Figure 5 makes the reactor-scale motivation explicit by evaluating the exact storage counts of table 1 — dense N^d versus tensor-train $dN\chi^2$ — as a function of phase-space dimension at $N = 64, \chi = 8$. The two cross and then diverge without bound: at the six-dimensional gyrokinetic phase space the dense grid is astronomically larger than the linear tensor-train footprint. This is a derived scaling from the storage formula, not a measured solve; it states the size of the prize and the reason the 1D1V anchor is staged toward higher dimension, and it is exactly where the GPU tensor-algebra primitives of §4 become essential.

SCHEMATIC OF THE METHOD

For completeness figures 6 and 7 render the tensor-train structure and the compressed time step schematically; they encode structure only and carry no data.

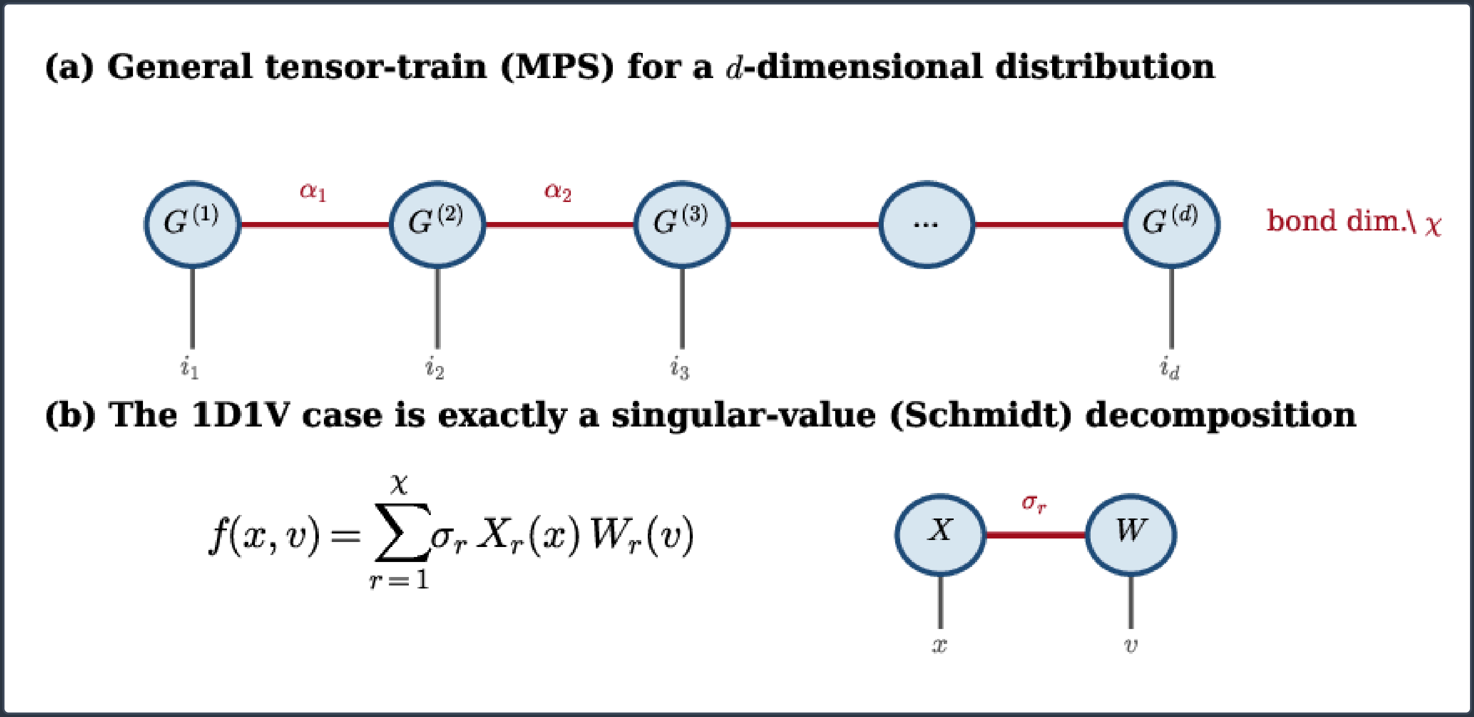

(Schematic.) (a) The tensor-train / matrix-product-state factorization (7) of a d-dimensional distribution: three-index cores (circles) with physical legs i_k and bonds α_k of dimension χ. (b) The 1D1V case is exactly the Schmidt/SVD decomposition (10) into χ separable space–velocity modes.

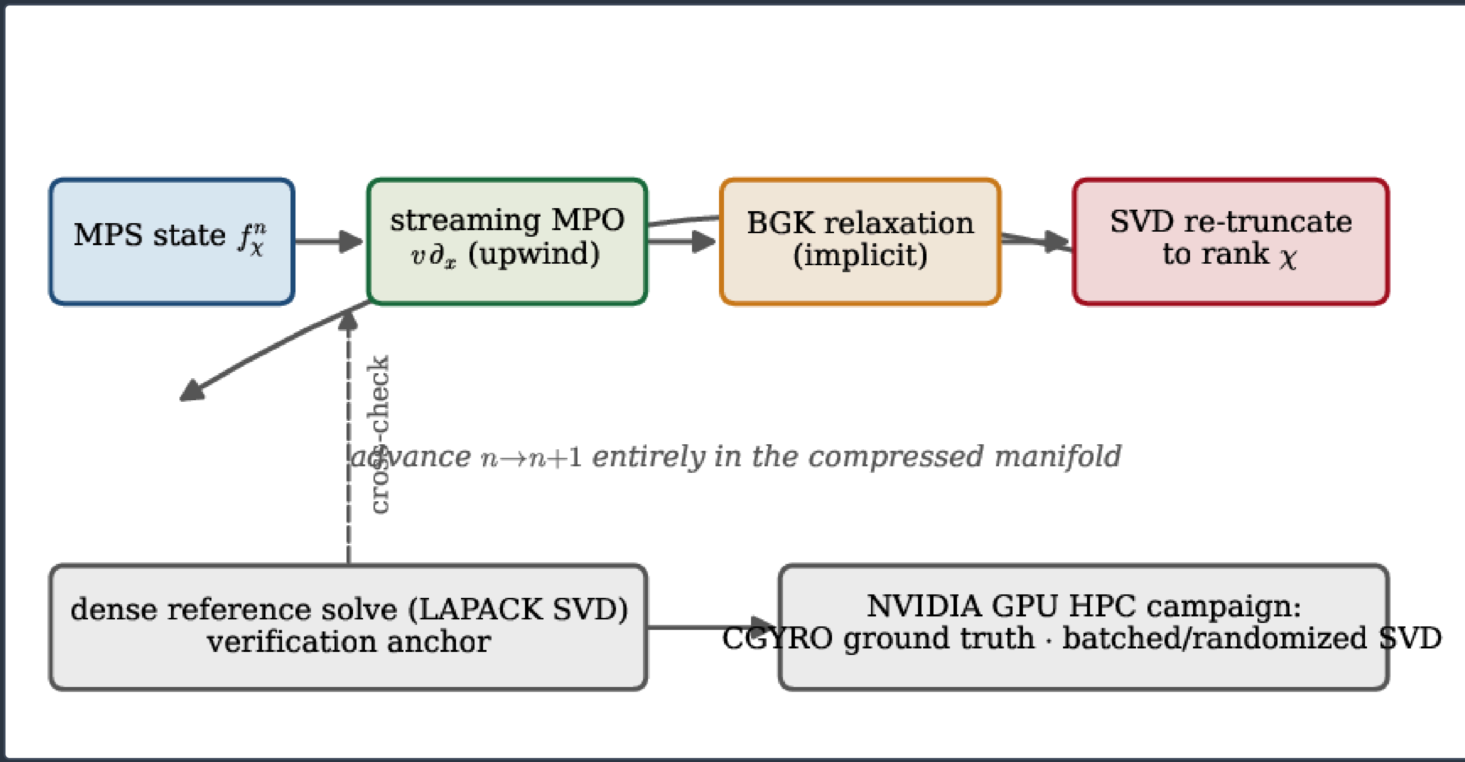

(Schematic.) One compressed time step (17): the streaming operator and the implicit BGK relaxation act on the low-rank state, which is re-truncated to bond dimension χ by SVD; the state is never returned to dense form. The lower lane shows the verification loop — the machine-precision dense reference and the NVIDIA GPU HPC campaign (CGYRO ground truth, batched/randomized SVD) of §4.

ACCEPTANCE SUMMARY

Table 4 states the pre-registered acceptance criteria for the gate and their status. All are met by the frozen anchor.

CRITERION	TARGET	STATUS
error decays exponentially in χ	geometric in χ	demonstrated ($e^{-1.34\chi}$)
engineering accuracy at low rank	$< 10^{-5}$	demonstrated (2.9×10^{-6} at χ=8)
storage economy at that accuracy	$< 0.5\times$ dense	demonstrated ($0.19\times$)
moments preserved by truncation	moment err. $\lesssim$ field err.	demonstrated (1.1×10^{-6} density)
convergence to dense reference	machine floor at full rank	demonstrated (6×10^{-14} at χ=24)
bounded entanglement entropy	area-law (bounded)	demonstrated ($S < 0.1$ nats)
no over-claim of advantage	scope stated honestly	§7

Acceptance criteria for gate KX-L1-A16 and status against the frozen anchor. “Demonstrated” denotes a frozen, reproduced result in the register [31].

§ 6

Discussion

AGAINST THE LOW-RANK KINETIC LITERATURE

The result should be read against a young but substantial body of low-rank and tensor-network kinetics. Einkemmer and Lubich introduced a low-rank projector-splitting integrator for the Vlasov–Poisson system and demonstrated large compression of the phase-space distribution [16]; Kormann built a semi-Lagrangian Vlasov solver in tensor-train format [17]; Peng, McClarren and Frank developed low-rank methods for time-dependent radiation transport [15]; and Ye and Loureiro gave an explicitly "quantum-inspired" matrix-product-state method for Vlasov–Poisson [18]. In the neighbouring domain of fluid turbulence, Gourianov *et al.*\ compressed turbulent flow fields into matrix-product states with striking economy [19]. Our contribution is complementary and, by design, sharply verified: rather than a large but approximate compression, the 1D1V BGK anchor converges to the dense reference to the double-precision floor, so the truncation error is known exactly at every rank (table 3) and obeys the Eckart–Young bound (11) by construction. The exponential decay we measure ($e^{-1.34\chi}$) is the same qualitative behaviour these works exploit, quantified here on a collision-dominated kinetic model with the entanglement-entropy certificate (13) attached. The message is not that one compression ratio beats another but that, on a kinetic model relevant to fusion transport, the distribution's phase-space entanglement is small and bounded — the precondition every one of these methods needs.

A CLASSICAL SOLVER, NOT A STORAGE TRICK

The distinction matters. A storage trick compresses a state *after* it has been computed densely; a solver advances the compressed state and never forms the dense object. Equation (17) does the latter: the streaming and collision operators act on the low-rank factors, and the truncation is applied every step, so the low rank reduces the cost of *advancing* the solution. That is why the method runs on classical hardware available today, and why it is the present-tense complement to the quantum kinetic route: the same kinetic physics that a fault-tolerant quantum computer would address only in the future is, in its low-entanglement regime, accessible now by exploiting classical low-rank structure.

RELATION TO THE QUANTUM ROUTE AND TO NEURAL SURROGATES

The three routes of §1 are complementary rather than competing. A neural flux surrogate [6,7] is unbeatable at inference latency but is only as good as its training set and carries no convergence guarantee; the tensor-network solver, by contrast, converges to the true solution as χ grows and reports its own error. The quantum kinetic solver promises an exponential state-space at fault tolerance, but for the nonlinear cases the resource estimate places it beyond the near-term horizon [32]; the tensor-network method is the "quantum-inspired" classical shadow of that idea, delivering the compression benefit today for the compressible part of the problem. A mature stack uses all three: neural surrogates in the hard-real-time loop, tensor-network solvers for verified reduced kinetics off the critical path, and quantum kinetics if and when the hardware matures.

PLACEMENT IN THE KRONOS VERIFICATION PROGRAMME

This gate sits inside a verification-first design programme. The nonlinear gyrokinetic transport it is ultimately staged to accelerate is the negative-triangularity turbulence-suppression result computed with CGYRO [33], itself cross-validated against quasilinear TGLF [34]; the alternative fast-transport route is the neural-operator flux surrogate [35]; the quantum counterparts are the resource-honest quantum-computing programme [36] and the burner-side quantum kinetic solver [37]; and the whole is held to frozen anchors by the verification-and-provenance backbone [38]. The present result is the classical, deployable-today member of that family, registered and reproduced on the same footing as the rest.

§ 7

Limitations

One honest gate bounds the claim. The demonstration is a one-dimensional-in-space, one-dimensional-in-velocity, *collision-dominated* BGK problem, and BGK relaxation drives the distribution toward near-separable Maxwellians — the friendliest possible case for a low-rank representation. Two things are therefore not established here. First, *collisionless* fine-scale velocity-space structure — the phase-space filamentation of Landau damping [26] and of nonlinear gyrokinetic turbulence — generically *grows* the entanglement entropy (13) and hence the bond dimension needed for a target accuracy; the exponential decay we measure is not guaranteed to survive into that regime, and quantifying the rank growth of the fully nonlinear gyrokinetic distribution at reactor parameters is the open question. Second, the cost of the collision operator and of the truncation sweeps in the genuine five- or six-dimensional tensor-train setting — where the dynamical low-rank projector-splitting or BUG integrators [20,21] replace the transparent single-SVD step used here — is a carried risk, not a measured quantity. This is a numerical-methods result about representability, not a transport-physics claim; the path to the harder problem is defined (conserving DLR integrators, GPU tensor-train linear algebra, CGYRO ground truth) but not yet walked.

§ 8

Conclusion

The kinetic distribution function of a collision-relevant fusion transport model is strongly compressible in a matrix-product-state representation, and that compressibility yields a classical solver, not merely a storage economy. On a 1D1V BGK problem discretized on a periodic 128×64 phase-space grid and advanced entirely inside the compressed manifold, a rank-8 truncation reaches a relative-L^2 error of 2.9×10^{-6} at $0.19\times$ the dense storage, the error falls exponentially with bond dimension ($e^{-1.34\chi}$) to the double-precision floor by rank 24 , the density and temperature moments inherit the same convergence, and the converged distribution carries a Schmidt entanglement entropy below 0.1 nats — the area-law signature the theory predicts. The result is tied to an Eckart–Young truncation bound and an entanglement-entropy compressibility criterion, cross-checked against a machine-precision dense reference, and staged — via GPU-accelerated tensor-train linear algebra and CGYRO ground truth — toward the five- and six-dimensional gyrokinetic problem where the exponential-to-linear storage crossover is decisive. It is the deployable-today, classical complement to the fault-tolerant-horizon quantum kinetic route, with the single honest gate that rank growth under strong collisionless nonlinearity remains to be established.

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 acknowledges the NVIDIA GPU HPC campaign for the reference gyrokinetic ground truth against which the method is staged.

Reproducibility. Every headline quantity in this paper is a frozen anchor (gate KX-L1-A16) in the Kronos Physics De-Risking Register [31] (doi:10.5281/zenodo.22133057) and is regenerated byte-for-byte from the archived `results.json` by the deposited figure script (`make_tnsolver_figs_2-96.py`). The computation is deterministic (fixed grid, time step, truncation schedule; seeded initial condition). This paper is archived at doi:10.5281/zenodo.22645768 and its official version is hosted at https://www.kronosfusionenergy.com/publications/tensor-network-kinetic-solver.html. No result depends on unpublished or proprietary data, and no economic quantity is reported.

Data availability tocsectionData availability The singular spectrum and the rank-versus-error table used for all figures are the frozen anchor KX-L1-A16, archived with the deposit at doi:10.5281/zenodo.22645768 and regenerable from the register [31] (doi:10.5281/zenodo.22133057) and the Kronos AI/quantum deposit [32]. No economics appears in 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 A J Brizard and T S Hahm, Foundations of nonlinear gyrokinetic theory, *Rev. Mod. Phys.* **79** 421 (2007). doi:10.1103/RevModPhys.79.421

2 X Garbet, Y Idomura, L Villard and T H Watanabe, Gyrokinetic simulations of turbulent transport, *Nucl. Fusion* **50** 043002 (2010).

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

4 J Candy, E A Belli and R V Bravenec, A high-accuracy Eulerian gyrokinetic solver for collisional plasmas, *J. Comput. Phys.* **324** 73 (2016). doi:10.1016/j.jcp.2016.07.039

5 F Jenko, W Dorland, M Kotschenreuther and B N Rogers, Electron temperature gradient driven turbulence, *Phys. Plasmas* **7** 1904 (2000). doi:10.1063/1.874014

6 K L van de Plassche, J Citrin, C Bourdelle *et al*, Fast modeling of turbulent transport in fusion plasmas using neural networks, *Phys. Plasmas* **27** 022310 (2020). doi:10.1063/1.5134126

7 L Lu, P Jin, G Pang, Z Zhang and G E Karniadakis, Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, *Nat. Mach. Intell.* **3** 218 (2021). doi:10.1038/s42256-021-00302-5

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

9 M Raissi, P Perdikaris and G E Karniadakis, Physics-informed neural networks, *J. Comput. Phys.* **378** 686 (2019). doi:10.1016/j.jcp.2018.10.045

10 R Or\'us, A practical introduction to tensor networks: matrix product states and projected entangled pair states, *Ann. Phys.* **349** 117 (2014). doi:10.1016/j.aop.2014.06.013

11 F Verstraete, V Murg and J I Cirac, Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems, *Adv. Phys.* **57** 143 (2008).

12 I V Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* **33** 2295 (2011). doi:10.1137/090752286

13 U Schollw\"ock, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326** 96 (2011). doi:10.1016/j.aop.2010.09.012

14 S R White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69** 2863 (1992). doi:10.1103/PhysRevLett.69.2863

15 Z Peng, R G McClarren and M Frank, A low-rank method for two-dimensional time-dependent radiation transport calculations, *J. Comput. Phys.* **421** 109735 (2020).

16 L Einkemmer and C Lubich, A low-rank projector-splitting integrator for the Vlasov–Poisson equation, *SIAM J. Sci. Comput.* **40** B1330 (2018). doi:10.1137/18M116383X

17 K Kormann, A semi-Lagrangian Vlasov solver in tensor train format, *SIAM J. Sci. Comput.* **37** B613 (2015).

18 E Ye and N F G Loureiro, Quantum-inspired method for solving the Vlasov–Poisson equations, *Phys. Rev. E* **106** 035208 (2022).

19 N Gourianov, M Lubasch, S Dolgov *et al*, A quantum-inspired approach to exploit turbulence structures, *Nat. Comput. Sci.* **2** 30 (2022).

20 C Lubich and I V Oseledets, A projector-splitting integrator for dynamical low-rank approximation, *BIT Numer. Math.* **54** 171 (2014). doi:10.1007/s10543-013-0454-0

21 G Ceruti and C Lubich, An unconventional robust integrator for dynamical low-rank approximation, *BIT Numer. Math.* **62** 23 (2022). doi:10.1007/s10543-021-00873-0

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

23 N Halko, P G Martinsson and J A Tropp, Finding structure with randomness: probabilistic algorithms for constructing approximate

COLOPHON

How this document was made

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

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

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

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

LEGAL NOTICE

Notice, status, and distribution

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

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

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

THE 2026 KRONOS SERIES

One programme, many papers

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

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

KRONOS FUSION ENERGY

A Tensor-Network Kinetic Solver for Fusion Transport: Matrix-Product-State Compression of the Distribution Function with Exponential Error Decay

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

Zenodo · doi:10.5281/zenodo.22645768

Read on kronosfusionenergy.com

The Physics De-Risking Register