MeanField numerical research and consolidation handoff

Formulation, experiments, evidence, architecture, and remaining work

Research record prepared with Codex

Source and numerical evidence: 16–20 September 2026

20 September 2026

Experimental three-dimensional FEM stellar equilibria
Independent density; coupled free surface; polytropic EOS

The moderately rotating $n = 3$ and full virial targets remain open.

This document separates measured results from hypotheses, and retained research tools from promoted defaults.

Contents

1 Purpose, scope and current conclusion **5**

1.1 Physical and software objectives 5

1.2 What has actually been demonstrated 5

1.3 How to use this record 6

2 Definitions and evidence discipline **6**

2.1 Physical variables and virial diagnostics 6

2.2 Algebraic quantities and cost accounting 7

2.3 Controls, provenance and limits of inference 7

3 Starting point, problem definition, and architecture **8**

3.1 Continuum equations and independent discrete fields 9

3.2 Discretization and model compilation 9

3.3 Physical acceptance and measurement conventions 10

4 Initial diagnosis: several different plateaus **11**

4.1 Fresh control measurements 11

4.2 Normalization audit 12

4.3 Central condition: a useful example correction 12

4.4 Krylov window and warm-guess safeguard 13

4.5 Rejected one-way gravity completion 13

5 Separating pressure preparation from shape-related curvature **14**

5.1 The rotating pressure obstruction was localized and real 14

5.2 Higher-index merit rejection had a different mechanism 15

6 Three optional numerical strategies and their limits **16**

6.1 Polynomial enthalpy certification and anchor repair 16

6.2 Gravity-only curvature correction 16

6.3 Gravity elimination 17

6.4 Production outcomes of the three strategies 17

7 Preconditioner improvements measured at equal accuracy **18**

7.1 Baseline approximation and work accounting 18

7.2 Prepared cross-action optimization 18

7.3 Mass-consistent gravity Schur refinement 19

7.4 Nonlinear transfer limited the recommendation 19

8 Why better linear solves did not remove the nonlinear stall **20**

8.1 Tenfold tighter forcing did not improve the useful step 20

8.2 A large compensated direction amplified moderate curvature 21

9 Fully coupled second-order trial corrections **21**
9.1 Bounded correction against the complete residual . . . 21
9.2 Damping and work-aware eligibility . . . 22
9.3 Matched six-update live test . . . 22

10 Shape freedom, field transport, and the displacement-order control **23**
10.1 The isolated Q_4 displacement experiment . . . 24
10.2 Two structural cautions . . . 24

11 The bounded smooth-surface study **25**
11.1 Frozen late-state result . . . 25
11.2 Three cold reduced updates from the analytic seed . . . 26

12 Native spherical-harmonic surface prescription **27**
12.1 Architecture and normalization . . . 27
12.2 Native coupled results . . . 28

13 Negative density: a located discrete representation defect **29**
13.1 What the weak EOS equation determines . . . 29
13.2 Spatial location and root-based integration . . . 29
13.3 Mechanism controls and their limits . . . 30

14 Independent density-order study and logarithmic assessment **31**
14.1 Frozen approximation results . . . 31
14.2 Coupled controls and the stalled $n = 3$ trajectories . . . 31
14.3 What the logarithmic assessment established . . . 32

15 Detailed coupled-convergence campaign **33**
15.1 Initial nonlinear screen . . . 34
15.2 Matched gravity-curvature trial-path control . . . 34
15.3 Frozen inverse comparisons: rate, attained accuracy and total cost . . . 35
15.4 Weak coupled radial response . . . 36

16 A physically compatible rotating $n = 3$ family **37**
16.1 Matched live field-feedback result . . . 37

17 Virial consistency: pressure and gravity must be distinguished **38**
17.1 Pressure-moment identity and the density test space . . . 39

18 Physical two-dimensional model slices **39**
18.1 Display-refinement measurements . . . 40

19 What to carry forward and what remains unresolved **40**
19.1 Recent implementation verification and provenance . . . 43

20 Consolidated source, architecture, and handoff **43**
20.1 What the consolidation changes . . . 43
20.2 Architectural boundary between physics and numerics . . . 44
20.3 Fields, geometry, and the preferred surface representation . . . 44
20.4 Defaults, explicit controls, and experimental candidates . . . 45
20.5 Branch disposition and retained provenance . . . 47
20.6 Adoption, retention, and exclusion rationale . . . 49
20.7 A practical contract map for the rewrite . . . 50

21 Evidence organization and reproducibility after consolidation **53**
21.1 Rebuilding and reproducing results . . . 54

22 Consolidation validation on 20 September **55**
22.1 Fresh physical smoke model . . . 55
22.2 Artifact and document checks . . . 56

A Versioned study-report catalog **56**

1 Purpose, scope and current conclusion

MeanField is an experimental three-dimensional finite-element laboratory for steady rotating stellar equilibria. Its immediate purpose is to establish reliable numerical choices for a later implementation, rather than to present a finished stellar-evolution or production equilibrium code. The user is independently rewriting the implementation in order to own and understand the eventual shipped software. The useful output of this investigation is therefore the formulation, verified contracts, measured successes and failures, and reproducible evidence that can guide that rewrite.

This report consolidates the work recorded from 16 through 19 September 2026 and the source/evidence integration performed on 20 September. Earlier optimization and validation records are preserved as historical sources. The conversation is not treated as an experimental log: quantitative claims below come from committed reports, campaign manifests, saved vectors, diagnostic tables and recorded checks. The initial approximately tenfold performance gain and approximately 20-second Newton step for the nonrotating $n = 1$ model were the user's starting observation. They are not remeasured as a new speedup in this consolidation. At the outset, limited testing suggested plausible $n = 1$ models at very low rotation, described by the user as roughly 1.5 percent of critical, while increasing index or rotation produced residual plateaus. That rough starting description is not a calibrated critical-rotation measurement for the later campaign inputs.

1.1 Physical and software objectives

The immediate physical target is a moderately rotating $n = 3$ polytrope, with an eventual comparison against ESTER. The operational rotation goal is at least 50 percent of critical rotation. The more distant ambition of exceeding 99 percent of critical rotation is outside the demonstrated regime and would require renewed attention to surface regularity and the vanishing effective gravity near mass shedding. No value of the dimensionless angular momentum used in these campaigns has been calibrated to a critical-rotation fraction. Neither a visually oblate model nor a small J proves a particular fraction of critical rotation.

The key global milestone is a relative virial error no larger than 10^{-5}. The original July project README described an aspirational 10^{-6} target; the research acceptance target in this report is the subsequently specified 10^{-5}. Internal pointwise correctness remains necessary. Exterior accuracy is required insofar as it supports the correct interior and boundary response; an independently accurate external field everywhere is not the immediate scientific objective. Nevertheless, the exterior formulation can influence interior gravity and cannot be disregarded merely because its own residual is small.

The software objectives are equally explicit: retain a physics-oriented user API, express algorithm and field compatibility through types where practical, use runtime checks for mesh-dependent and state-dependent conditions, preserve independent density, and solve surface coefficients and interior fields in the same coupled system. Central density is an optional physical specification; it must not be silently imposed on every polytrope. Log density and new positivity enforcement are deferred. Existing optional feasibility experiments remain available for reproduction but are not enabled by consolidation.

1.2 What has actually been demonstrated

There is substantial progress, but the central rotating $n = 3$ goal is open. The strongest established results are:

1. The native all-m, degree-four real spherical-harmonic surface prescription is integrated and tested against the corresponding nodal geometry, physical residuals and Jacobian actions. It

reduces the surface to 25 independent coefficients without decoupling it from the interior solve.

2. Some low-index, slowly rotating native-harmonic cases meet both their algebraic stopping criterion and the global virial target. Higher-index representation errors and unresolved pointwise checks prevent generalizing this to every algebraically converged case.

3. At the same frozen state, gravity-Schur coupling reduced the measured total linear cost by 41.5 percent for $n = 3, J = 0$ and by 30.6 percent for $n = 1.5, J = 0.1$, at the stated true-defect gate. These are single matched comparisons, including setup, rather than statistical timing estimates.

4. In the compatible rotating $n = 3, J = 0.035, K/K_0 = 0.99$ family, stellar LDU plus gravity-Schur reached a 12.687-times lower residual after ten updates than gravity-Schur alone, at 14.593 percent greater Newton time. This is an equal-update endpoint comparison, not a speedup to the same converged model.

5. The best rotating $n = 3$ endpoint still fails the full virial target: the enthalpy-pressure energy error is 1.81141×10^{-5}, and the density-pressure energy error is 3.51558×10^{-5}. Passing one force virial is insufficient.

6. Physical-plane slices now show independent density and enthalpy on the actual deformed stellar geometry, with signed values and the undeformed outline available for comparison. Export refinement has been checked for selected $n = 2$ and rotating $n = 1.5$ states; it is not a mesh-convergence proof.

Several attractive paths did not justify promotion: simple one-way material-to-gravity completion, increasing displacement order alone, expensive coarse surface/scalar corrections, and universal use of second-order trial corrections. Some improve an isolated metric while increasing total work or worsening the accepted nonlinear trajectory. Those negative results are preserved because they constrain the next design, not because the algorithms are impossible to improve.

1.3 How to use this record

The report first reconstructs the baseline and sequence of hypotheses, then describes the current implementation, numerical evidence and interpretation. Later sections distinguish active defaults from selectable experimental policies and record the consolidation checks. The source reports under `experiments/` provide the most direct detailed companion records. `reports/archive/` contains branch and relocation inventories. `working/research/diagnostics/` contains the preserved raw campaigns. The large raw archive is local and intentionally not committed to Git; a source-only clone will not contain its saved models or copied executables. Versioned figure copies and this report remain usable without that archive.

2 Definitions and evidence discipline

2.1 Physical variables and virial diagnostics

For the polytropic experiments, pressure obeys

$$P(\rho) = K\rho^{1+1/n}, \qquad h(\rho) = (n+1)K\rho^{1/n}, \qquad \rho_\star(h) = \left(\frac{h}{(n+1)K}\right)^n \tag{1}$$

on their physical nonnegative domains. The discrete density and enthalpy are independent finite-element fields; the weak closure equation does not make them pointwise equal. Consequently an integral formed from $P(\rho_h)$ need not agree with one formed from $P(h_h)$ at the same discrete state.

Undefined fractional powers of negative density must be reported as undefined, not silently replaced with a positive or clipped value.

For a zero-pressure stellar boundary, define

$$T = \frac{1}{2}\int_{\mathcal{D}_\star} \rho|\boldsymbol{v}|^2\, dV, \qquad \Pi = \int_{\mathcal{D}_\star} P\, dV, \qquad W_\Phi = \frac{1}{2}\int_{\mathcal{D}_\star} \rho\Phi\, dV. \tag{2}$$

With the usual negative gravitational potential, $W_\Phi < 0$. The force binding integral W_f uses the corresponding gravitational force moment and agrees with W_Φ for a consistent isolated equilibrium. The campaign observers report energy and force virials using the common normalization $|W_\Phi|$:

$$v_\Phi = \frac{2T + 3\Pi + W_\Phi}{|W_\Phi|}, \qquad v_f = \frac{2T + 3\Pi + W_f}{|W_\Phi|}. \tag{3}$$

Tables generally show their absolute values unless signs are explicitly given. A nonzero external pressure requires the boundary-pressure term; the stated campaign comparisons use the zero-pressure surface convention. Both pressure reconstructions and both gravity forms are retained where they are defined. Agreement at quadrature orders 14 and 18 is an integration consistency check on a fixed discrete state, not evidence of spatial convergence or exact pointwise physics.

The common campaign scale is $M = R_{\text{reference}} = G = 1$ unless a manifest states otherwise. Angular momentum J, angular velocity, $T/|W|$, radius and critical-rotation fraction are different quantities. The physical-family compatibility study later explains why this distinction is particularly important at $n = 3$.

2.2 Algebraic quantities and cost accounting

Let x denote the complete normalized coupled coordinate vector and $F(x)$ its normalized residual. Write $A = DF(x)$ for its Jacobian; this notation avoids confusing the Jacobian with the stellar angular momentum J. A Newton correction p is checked through the fresh full-system defect

$$\eta_{\text{true}} = \frac{\|F(x) + Ap\|}{\|F(x)\|}. \tag{4}$$

An internal Krylov recurrence or a reduced-system residual cannot substitute for this full-system endpoint test. The nonlinear merit function is $\mathcal{M}(x) = \frac{1}{2}\|F(x)\|^2$; its directional derivative is $F(x)^T Ap$. Linear success and a negative initial merit derivative do not guarantee acceptance of a finite step on a curved coupled path.

The state and residual norms are produced by the configured physical Riesz normalization. They should not be compared across a change of basis, field order or normalization as though they were the same physical norm. The experiments therefore retain physical residual checks and integrals in addition to normalized algebraic quantities. A block's fraction of $\|F\|^2$ is an algebraic attribution; it is not an elementwise PDE error estimate.

Frozen comparisons hold the complete base state, Jacobian, right-hand side, mesh and coordinates fixed. Their relevant linear cost includes construction, binding, setup, solve and a fresh endpoint verification. Live comparisons follow evolving nonlinear states and include preparation, accepted/rejected trials and auxiliary solves in the recorded Newton time. Process wall time can additionally include observers and export overhead. Setup amortization must therefore be demonstrated on an actual repeated-use path rather than assumed from a smaller outer iteration count.

2.3 Controls, provenance and limits of inference

The adopted methodology used an isolated branch for each hypothesis or a clearly identified cumulative integration branch. A source revision and configuration were recorded before numerical

work; campaign runners copied the executable, recorded its digest and retained the exact commands, mesh digest, saved state and block diagnostics. Builds, FEM runs and timing trials were serialized within the later controlled studies. Read-only analysis and source review could proceed independently. Single-run timings are reported as observations; no confidence interval or cross-machine performance claim is inferred from them.

The pressure/shape investigations separated preparation feasibility, linear accuracy, initial descent and finite-step curvature. Derivative probes used directional or centered differences over a stated step grid and compared full versus component actions. Later snapshot loaders check dimensions, mesh, coordinate conventions and provenance rather than accepting a vector solely because its length matches. Unconverged restarts require an explicit diagnostic override. Saved-state initialization prepares and refreshes before committing, clears stale warm starts and rolls back failed initialization.

The evidence checkers reject corrupt or inconsistent fixtures as well as accepting valid ones. Passing an artifact audit establishes consistency of the recorded experiment under that checker's contract. It does not establish that the nonlinear model converged, that all interior samples are physical, or that a target virial was attained. Similarly, passing code assertions is not equivalent to obtaining a validated stellar model.

All finite iteration budgets and diagnostic ceilings remain explicit. A late small-step trajectory can be described as a stall observed over the available window. It cannot be promoted to a mathematically established convergence floor. Dense quadrature samples can locate negativity but cannot certify its absence everywhere; interval subdivision and polynomial bounds have their own scope and unresolved cases. Conclusions in this report retain those distinctions.

3 Starting point, problem definition, and architecture

The work began on 16 September 2026 with an already substantial experimental finite-element code. Its purpose was to discover a reliable numerical route to three-dimensional steady stellar equilibria before a later, independently written general-EOS implementation. The immediate physical model was a self-gravitating, rigidly rotating barotrope with a free stellar surface. The presence of older transport, Coriolis, viscosity, and continuity components did not mean that those equations participated in the active equilibrium solve. This distinction remains important when comparing the experimental code with the broader physics ultimately desired from a stellar structure application.

The incoming Git HEAD was `75cc63873992984966ea20011cc04d9984f54be4`. It did not uniquely describe the audited source: 49 tracked files already differed, with 3,626 additions and 267 deletions, and additional untracked experiments and results were present. The later baseline commit `80cc3d8` on `research/baseline-20260916` preserved that optimized source. The initial architecture audit was read-only. The subsequent diagnosis used fresh Release builds rather than assuming that existing executables represented the modified source. This provenance distinction is essential: historical binaries and historical passing tests are evidence about their own source, not automatically about the current checkout.

The user reported approximately a tenfold reduction in the cost of an ordinary nonrotating $n = 1$ Newton step, to about twenty seconds. Fresh measurements found 22–24 seconds for the first three accepted steps and were consistent with that operating point. The old pre-optimization source was not rerun in a matched comparison; the full historical factor of ten was therefore not independently remeasured. The inherited 9 September report had separately measured factors of 1.53 in whole-process time, 1.70 in Newton time, and 2.02 in linear-solve time for its particular batch. Later tensor, cache, and compact-action changes had different baselines. These gains must not be multiplied together.

3.1 Continuum equations and independent discrete fields

For a polytrope of index $n \geq 1$ and constant $K > 0$,

$$P = K\rho^{1+1/n}, \qquad h = (n+1)K\rho^{1/n}, \qquad \rho_{\mathrm{EOS}}(h) = \left[\frac{h}{(n+1)K}\right]^n, \qquad h \geq 0. \tag{5}$$

With positive centrifugal potential $\Psi = |\boldsymbol{\Omega} \times \boldsymbol{x}|^2/2$, the Bernoulli imbalance is

$$B = h + \Phi - \Psi - C. \tag{6}$$

The gravity flux represents $\boldsymbol{g} = \nabla\Phi$, so physical gravitational acceleration is $-\boldsymbol{g}$. The active system combines a mixed gradient–potential gravity formulation, a weak density closure, weak Bernoulli balance, reduced mechanical virtual work, and integral constraints for total mass and axial angular momentum. Density, enthalpy, potential, and the gravity flux are independent finite-element fields. In particular, the mechanical pressure is evaluated from enthalpy while gravity and centrifugal forces use independent density. Continuum equivalence of these quantities does not force their pointwise equality after discretization.

The surface condition replaces the relevant enthalpy test rows by the enthalpy value corresponding to prescribed surface pressure. At zero surface pressure this is an $h = 0$ trace condition. The mechanical residual is first formed against volume displacements and then pulled back by the transpose of the surface-to-volume deformation derivative. If q denotes independent surface parameters and $D(q)$ the generated displacement, its schematic form is

$$R_s(q, u) = D'(q)^T R_{\mathrm{mech}}(D(q), u). \tag{7}$$

The Jacobian includes both the derivative of the mechanical force and the derivative of this pullback. A zero R_s proves equilibrium only against the chosen extension-generated mechanical tests; it does not set every unrestricted volume mechanical row to zero.

The usual initial example additionally imposed a central-density specification. Its actual row was a central enthalpy point condition,

$$R_c = h(0) - h(\rho_{c,\mathrm{target}}), \qquad R_h \longleftarrow R_h + \lambda_c e_c. \tag{8}$$

It was not a direct point condition on the independent discontinuous density. This distinction motivated measuring the unbordered residual and artificial source amplitude. The user clarified that the central condition should remain optional, principally useful for the $n = 3$ homology, and should not be automatically inserted at other indices or hidden in a physics-facing API.

3.2 Discretization and model compilation

The common mesh had 1,216 hexahedra: 832 stellar and 384 vacuum elements. Its SHA-256 was `ce11ae99e21e6c3bbfbee402acd2e191c1da0d8261d2227b4203f73f2f337a74`. Reference geometry was order four. The ordinary spaces were RT index two for gravity, discontinuous Q_2 for potential and stellar density, and continuous Q_3 for enthalpy and displacement. The original scalar nodal radial surface had 866 parameters. With its three scalar constraints the root had 178,074 unknowns; removing the central specification removed one scalar, yielding 178,073. Mesh geometry order and the approximation order of an individual physical field were distinct controls.

Initial root block	Number of unknowns
Gravity gradient	98,928
Potential	32,832
Stellar density	22,464
Stellar enthalpy	22,981

Nodal radial surface	866
Generated scalar coordinates	3

The architecture was already strongly compositional. A compile-time model specification pack generated the block layout, scalar coordinates, Jacobian couplings, normalization metadata, and prepared specification runtimes. The active aggregate was the variadic prepared equilibrium operator. C++23 module interfaces under `libmeanfield/interface` were separated from implementations under `libmeanfield/impl`. MFEM/MPI spaces used explicit maps between supported solver coordinates and full true degrees of freedom. The nonlinear context owned accepted and trial states, normalization, inverse preparation, and linear backend; refresh, stale-view rejection, rollback, and observer lifetime were explicit contracts rather than informal conventions.

Prepared matrix-free actions cached reference basis tables, geometry, thermodynamics, and local contractions. Their dependency identities and revisions controlled invalidation. The numerical work therefore emphasized reusing established ownership and preparation interfaces rather than adding ad hoc mutable solver state. Static choices such as block identities, compatible policy combinations, and registered capabilities were checked at compile time. Mesh sizes, rank agreement, actual FE eligibility, values, normalization factors, and preparation generations necessarily required runtime validation.

The geometry map initially used nodal radial surface variation, a logical radius-squared interior extension, and a fixed-infinity vacuum extension. Kelvin compactification represented the exterior through infinity. The root geometry unknown was the surface, not the much larger volume displacement field. Surface parameterization, interpolation of displacement, extension, and the physical map were therefore distinct architectural responsibilities.

The seed initialized density and enthalpy independently by finite-element coefficient interpolation of a spherical Lane–Emden profile, using MFEM's `ProjectCoefficient`, and solved mixed gravity separately. This initialization was distinct from solving the weak closure's weighted L^2 projection relation; the method name must not be read as implying such a mass-matrix solve. It set shape to zero, initialized C from $-GM/R$, and initialized Ω from the target J divided by spherical inertia. A seed with nonzero target J was not an already rotating equilibrium. Throughout the early controlled campaign, $M = R_{\text{reference}} = G = 1$; at $n = 1$, $K = 2/\pi$ and $\rho_c = \pi/4$. Other indices used the production Lane–Emden integrator with step 10^{-3} to calibrate the seed. That was a consistent seed reference, not an independent high-precision continuum solution.

3.3 Physical acceptance and measurement conventions

The required milestone was virial balance within 10^{-5}, together with credible interior fields. Exterior pointwise accuracy was not an independent goal; exterior errors mattered through their effect on stellar equilibrium. The campaign measured both gravitational energies

$$W_E = \frac{1}{2}\int_\star \rho\Phi \, dV, \qquad W_F = -\int_\star \rho\, \boldsymbol{x} \cdot \nabla\Phi \, dV, \tag{9}$$

and, for the zero-pressure surface,

$$\epsilon_V(W, P) = \frac{|2T + W + 3\int_\star P \, dV|}{|W_E|}. \tag{10}$$

Pressure from h and pressure from ρ were evaluated separately, at independent diagnostic quadrature orders 14 and 18. Signed density was retained. When the EOS could not define pressure at a negative density sample, the full density-pressure diagnostic was reported undefined, rather than manufactured by clipping. Quadrature agreement tested integration sensitivity; it did not constitute spatial convergence or global positivity certification.

The first physical checker required recorded nonlinear convergence and all eight finite virial measurements to satisfy 10^{-5}. This conservative combined screen was stricter than a statement that a particular enthalpy-based virial passed. Neither passing the checker nor a small scalar virial alone was a pointwise continuum-accuracy certificate. Useful accompanying observations were mass and angular-momentum errors, Bernoulli variation, EOS mismatch, gravity consistency, geometry minima, density sign, and surface shape. No early value of J or reference-unit Ω was calibrated as a fraction of critical rotation.

The inherited saved $n = 1, J = 0$ benchmark already provided a useful physical reference beyond the scalar residual: density, enthalpy, interior potential, and gravity-gradient relative volume L^2 errors were approximately $4.70\,10^{-4}$, $3.27\,10^{-4}$, $1.26\,10^{-4}$, and $5.98\,10^{-4}$; surface-radius RMS error divided by reference radius was $2.75\,10^{-4}$. Its density-pressure energy virial was $3.10\,10^{-7}$ and mass error $2.36\,10^{-11}$. These were the particular saved coarse model's errors against its radial reference, not accuracy estimates for later rotating or higher-index cases. A one-level refined mesh and analytic geometry controls existed, but the corresponding directory did not contain completed refined equilibrium fields or an equilibrium convergence result. The audit did not infer completed mesh convergence from the presence of that directory.

The historical full-suite XML contained 592 cases, 588 passing and four failing: centrifugal-virial refinement geometry, an $n = 3$ projected-seed closure screen, and two- and four-rank MPI groups. Those failures had been reproduced in earlier preserved binaries. They were neither silently removed nor treated as fresh validation of the incoming modified source. At the initial audit, twelve production source files were newer than the available release sandbox binary and thirty newer than the test binary. This motivated the subsequent explicit rebuilds and targeted regression ledgers. A previous between-sample mesh fold was also documented, but the later saved $n = 1$ mesh had healthier sampled geometry; it was not assumed to suffer the same demonstrated folding mechanism.

Evidence. The original workspace reports were `reports/sources/original_audits/MeanField_audit_2026-09-16.md`, `reports/sources/original_audits/MeanField_solver_diagnosis_2026-09-16.md`, and `reports/sources/original_audits/formulation_audit_notes.md`. Their archival relocation is recorded in the consolidation manifest. Repository evidence includes `experiments/POLYTROPE_VALIDATION_FINDINGS_2026-09-09.md`, `experiments/POLYTROPE_H_CONVERGENCE_2026-09-09.md`, and `experiments/NUMERICAL_CAMPAIGN_2026-09-16.md`.

4 Initial diagnosis: several different plateaus

4.1 Fresh control measurements

The diagnostic rebuilt the incoming optimized source in Release using Homebrew Clang 22.1.8, with IPO enabled, profiling disabled, and uniform field order increment zero. Cases ran serially on one MPI rank. Initial live controls used FGMRES restart 40, at most 200 linear iterations, relative forcing tolerance $\eta = 0.03$, nonlinear relative tolerance 10^{-8} and zero absolute tolerance. The nonlinear threshold was relative to the initial residual, not an absolute 10^{-8} requirement.

Model	Accepted updates	Retained $\|F\|$	$P(h)$ energy/force virials
$n = 1, J = 0$	3 full	$6.6053\,10^{-9}$	$2.616\,10^{-7}/3.246\,10^{-7}$
$n = 1, J = 0.01$	4 full	$5.0898\,10^{-10}$	$2.610\,10^{-7}/3.246\,10^{-7}$
$n = 1, J = 0.05$	5 full	$6.0406\,10^{-9}$	$3.016\,10^{-7}/3.188\,10^{-7}$
$n = 1.5, J = 0$	$1, 1/16$	$1.6976\,10^{-4}$	$2.245\,10^{-6}/2.318\,10^{-6}$
$n = 3, J = 0$	0	$3.1006\,10^{-3}$	$5.326\,10^{-5}/8.341\,10^{-5}$

All these controls stopped at a failed linear gate. The $n = 3$ measurement was the unchanged seed.

At $n = 1, J = 0$ the strict target was approximately $1.81\,10^{-12}$, whereas an earlier validation driver accepted the same three-step state under an absolute 10^{-8} threshold. These descriptions were consistent; they did not imply that a previously converged model had become physically invalid. Conversely, the $n = 1$ seed itself already passed the virial threshold, directly demonstrating why virial alone was not a complete equilibrium test.

The initial audit separated three mechanisms: loss of useful Krylov information late in otherwise accurate $n = 1$ solves; stronger coupled linear difficulty at higher index; and severe nonlinear curvature or admissibility restrictions after a usable linear direction was obtained. It did not find evidence for one universal numerical floor.

4.2 Normalization audit

The state and residual registrations supplied their own physical scales and reference Gram diagonals. Writing these as $s_x, d_{x,i}$ for a value block and $s_F, d_{F,i}$ for its residual/test block, ordinary coefficient factors were

$$S_{ii} = \frac{\sqrt{d_{x,i}}}{s_x}, \qquad L_{ii} = \frac{1}{s_F \sqrt{d_{F,i}}}. \tag{11}$$

The Gram factors form paired primal and dual maps where the registered function-space topologies correspond, but the physical scales need not agree. For example, surface parameters use a length scale while mechanical shape residuals use a force scale; potential uses specific energy while its Poisson residual uses inverse time squared. Point replacement rows and scalar blocks have their declared special treatment. These distinctions are explicit in `libmeanfield/interface/normalization/physical_riesz.cppm`.

The reference factors were frozen through a solve. Their characteristic scales came from $M, R_{\text{reference}}, G$ and did not divide by the requested angular momentum J; zero rotation was therefore not a normalization singularity.

Let $\widehat{x} = Sx$ and $\widehat{F} = LF$. The consistent transformations were

$$\widehat{A} = LAS^{-1}, \qquad \widehat{P}^{-1} = SP^{-1}L^{-1}, \qquad \widehat{A}\widehat{P}^{-1} = LAP^{-1}L^{-1}. \tag{12}$$

Consequently, changing only state scales cancels from a consistently right-preconditioned system. Direct probes multiplying block state scales by 10^{-3} through 10^{3} preserved the physical correction and preconditioned action to approximately 10^{-16}. Wrapper/explicit actions and coordinate round trips agreed at roundoff. This ruled against a basic scaling-wrapper implementation error at the inspected states. It did not show that the chosen residual norm was optimal: changing L changes what both Krylov and line search minimize.

The same $M = R = G = 1$ factors were used across indices despite changing stellar concentration. Large factor ranges alone did not diagnose a bug. Similarly, small normalized border matrices had raw condition numbers $2.78\,10^{7}$ and $9.69\,10^{7}$, dominated by a central-column scale near $7.3\,10^{6}$; diagnostic equilibration reduced these to about 1.46 and 2.39. Those were border-matrix observations, not the condition number of the full stellar Jacobian or proof of a singular physical system.

4.3 Central condition: a useful example correction

At $n = 1, J = 0.05$, matched restart-40 runs with and without the central condition had identical first five accepted step lengths and linear iteration counts. Both then failed the sixth linear gate near relative defect 0.505. Removing the condition nevertheless markedly improved localized physical consistency:

Quadrature-18 measurement	With central condition	Without

Peak absolute EOS density mismatch	$7.30451\,10^{-3}$	$3.32568\,10^{-4}$
Peak mismatch / sampled density scale	$9.33951\,10^{-3}$	$4.30594\,10^{-4}$
Bernoulli sampled range	$1.45751\,10^{-2}$	$1.30398\,10^{-3}$
Unbordered residual	$4.54227\,10^{-6}$	$6.03609\,10^{-9}$
$P(h)$ energy virial	$3.01617\,10^{-7}$	$3.06913\,10^{-7}$

The formerly fixed central enthalpy changed from 1 to 0.983383508. Mass, angular momentum, angular velocity, and volume RMS errors changed little. The result supported removing that artificial restriction from the $n = 1$ example, while leaving explicit central constraints available. It did not establish that central constraints caused the general plateau.

4.4 Krylov window and warm-guess safeguard

Frozen cold comparisons retained the same state, full Jacobian, physical right-hand side, and inverse. Increasing only the restart window from 40 to 160 gave:

Frozen state	Window 40: iterations / defect	Window 160: iterations / defect
$n = 1, J = 0$ endpoint	200/0.04442	72/0.02914
$n = 1, J = 0.01$ endpoint	200/0.35468	125/0.02519
$n = 1.5, J = 0$ endpoint	200/0.06491	70/0.02694
$n = 3, J = 0$ seed	200/0.05631	200/0.04026

For the rotating $n = 1$ state, histories coincided through iteration 40 and separated immediately after restart. Independently recomputed full residuals agreed with reported residuals. This established actual loss of useful Krylov information, not false convergence reporting. In a live unbordered $n = 1, J = 0.05$ comparison, restart 160 reached strict convergence in seven full steps with $F = 3.38036\,10^{-11}$, versus five steps and $6.03609\,10^{-9}$ before linear failure with restart 40. Total linear iterations were 292 versus 370; complete diagnostic-process times were 271.95 versus 290.80 seconds. All eight virial checks were below $3.735\,10^{-7}$. These were single timing observations with increased Krylov storage, not a universal speedup claim.

The solver reused $(1 - \alpha)p$ after a damped update as the next initial linear guess. This was vector reuse, not Krylov recycling. At one $n = 1.5$ state that guess had a true defect larger than the zero guess. A small backend safeguard now checks the current true defect of a supplied nonzero guess and discards it only if the norm is nonfinite or exceeds $\|b\|$. Nonfinite input coefficients remain an input-contract error. The guard adds one action for a nonzero guess and none for a zero guess, without a new public option.

The matched $n = 1$ output was bitwise unchanged. At $n = 1.5$, the first shared guarded solve reduced the effective initial defect/RHS from 1.5097 to 1 and iterations from 71 to 65. Over eight attempts the control had seven successful linear solves followed by failure; the guard passed all eight. Total iterations fell from 672 to 561, mostly by avoiding that failed 200-iteration attempt. The endpoint changed only from $1.75135\,10^{-4}$ to $1.74946\,10^{-4}$, a 0.108 percent improvement, with final steps 1/2048 and 1/8192. This was a modest inner-solver robustness improvement, not a cure for nonlinear stagnation.

4.5 Rejected one-way gravity completion

The first coupling hypothesis changed the structure inverse from independent material/gravity blocks to material-then-gravity triangular completion. It retained the material inverse, order-five Chebyshev flux-mass inverse, and three AMG cycles, and rebuilt the scalar border consistently. At the same centrally constrained $n = 3$ seed, the restart-160 defect changed from 0.0402590 to 0.0381781 after 200 iterations; both failed $\eta = 0.03$. Time increased from 98.92 to 117.71 seconds,

correction norm from 0.75764 to 1.54083, and the sampled thermodynamic step limit worsened from 0.15244 to 0.09135. The candidate's $\alpha = 0.1$ trial was inadmissible. This particular one-way completion was not promoted.

The baseline restart-40 and restart-160 corrections had cosine -0.985 despite the same matrix and right-hand side. Their difference had a response bound $\|A(p_2 - p_1)\|/\|p_2 - p_1\| \leq 2.77\,10^{-4}$. This established a weak measured direction, not a singular value or an identification of homology, translation, or a surface mode. Later studies therefore examined the actual coupled directions instead of adding a gauge on analogy alone.

Evidence. `experiments/NUMERICAL_CAMPAIGN_2026-09-16.md`, `experiments/CENTRAL_CONSTRAINT_HYPOTHESIS_2026-09-16.md`, `experiments/KRYLOV_WINDOW_HYPOTHESIS_2026-09-16.md`, and `experiments/KRYLOV_HYPOTHESIS_2026-09-16.md`. The combined nonrotating $n = 1$ case subsequently reached $F = 1.65157\,10^{-13}$ in six full steps with worst virial $3.73092\,10^{-7}$. The transfer case $n = 1, J = 0.1$ passed all twelve linear gates but ended at $6.25389278\,10^{-3}$, with final $\alpha = 4.76837\,10^{-7}$ and failing enthalpy-based virials $4.75186\,10^{-4}/8.18118\,10^{-4}$. That transfer failure motivated the next investigation.

5 Separating pressure preparation from shape-related curvature

The focused 17 September campaign made no production numerical changes. It captured exact trial states, identified rejection stages, and split the same Newton correction into its actual shape block p_s and all remaining fields and scalars p_f. Frozen replays restored the base before each action; probing a trial never committed it as an accepted model.

5.1 The rotating pressure obstruction was localized and real

At $n = 1, J = 0.1$, exact replays of the first rejected candidates at zero-based iterations 2, 3, 4, 7, and 11 all returned the production code `thermodynamic_domain / pressure_force / outside_domain`. The pressure source takes this branch for finite negative enthalpy. Thus these were enthalpy-domain failures, not inverted geometry. The limiting pressure point remained at local stellar element 459, quadrature point 50; the closure's different rule limited at point 36 of the same element.

Iteration	Pressure-point h	δh	Pressure bound	Accepted α
2	$9.63674\,10^{-4}$	$-1.31441\,10^{-2}$	$7.33159\,10^{-2}$	1/16
3	$1.42165\,10^{-4}$	$-2.11119\,10^{-3}$	$6.73390\,10^{-2}$	1/16
4	$1.02159\,10^{-5}$	$-1.89704\,10^{-3}$	$5.38520\,10^{-3}$	1/256
7	$7.04574\,10^{-8}$	$-1.86569\,10^{-3}$	$3.77648\,10^{-5}$	1/32768
11	$1.07227\,10^{-9}$	$-1.86493\,10^{-3}$	$5.74963\,10^{-7}$	1/2097152

The pressure step budget fell by a factor of 127,514 while the late negative direction remained finite. The limiting element touched the stellar surface; the pressure point's reference Cartesian radius was 0.997553, near the negative rotation pole, approximately 8.66 degrees from the negative z axis at onset. It was an interior quadrature point, not a boundary row. Topological adjacency did not measure distance to the deformed surface.

At onset the closure bound, 0.0727489, was slightly tighter than the pressure bound; pressure was evaluated first and therefore reported first when both failed. A deliberately intermediate trial at $\alpha = 0.07303236911639208$ passed pressure preparation and failed closure, with predicted closure $h = -5.01979\,10^{-6}$ and limiting pressure $h = +3.72619\,10^{-6}$. This showed why a first rejection stage must not be confused with the tightest thermodynamic constraint.

All 866 surface enthalpy coefficients and their directions were exactly zero at onset and late capture. Boundary leakage was excluded for these trials. Nevertheless, positive nodal values did not imply a positive represented polynomial. A step just beyond the onset bound retained nonnegative values at all 64 local Q_3 nodes while making interior pressure and closure samples negative. At the late state even a unit step retained nonnegative nodes but gave $h = -1.86493\,10^{-3}$ at the limiting pressure point.

An independent reconstruction of the local tensor polynomial reproduced the nodes exactly and limiting-point values within $6.94\,10^{-18}$. A much denser grid found that the late accepted base already had negative enthalpy: a sampled minimum $-2.87114\,10^{-5}$ at reference coordinates $(0.0130467, 0.515625, 0.01)$. A separate 60-digit reconstruction agreed with the negative value at another selected interior point. These findings proved negative values existed in that element; the reported minimum was not a certified global minimum. Passing production quadrature preparation was thus not a pointwise positivity certificate. They also did not prove that every low-rotation or nonrotating converged control shared this defect.

5.2 Higher-index merit rejection had a different mechanism

For any directional piece p_d, the diagnostic measured

$$R_d(\alpha) = F(x + \alpha p_d) - F(x) - \alpha A p_d, \qquad R_{\text{mixed}} = R_{\text{full}} - R_s - R_f. \tag{13}$$

The mixed vector was formed before taking norms. All pieces used common step lengths and the original, unrenormalized block amplitudes. With $L = F + \alpha Ap$, the exact merit identity was

$$\|F(x + \alpha p)\|^2 = \|L\|^2 + 2\langle L, R_{\text{full}}\rangle + \|R_{\text{full}}\|^2. \tag{14}$$

At $n = 1.5, J = 0$, captured iteration 1, the raw full trial was admissible with minimum sampled Jacobian 0.896116. Its linear prediction was only 0.02576812 of the base norm; the actual residual was 12.243161 times the base norm. At $\alpha = 0.1$ the actual ratio remained 1.02052 rather than the predicted 0.900070; at 0.01 it decreased to 0.991209. This established nonlinear merit failure separately from the rotating thermodynamic wall.

Gravity-gradient actions from shape and fields had norms 0.115959 and 0.116055, cosine -0.999999831, and combined norm $1.16990\,10^{-4}$. Poisson actions also nearly cancelled, with cosine -0.999996055. Damping shape alone would have destroyed this measured first-order balance. The full-step second-order remainder was $2.18903\,10^{-3}$, with 99.7747 percent of its squared norm in gravity. Fixed-shape gravity is affine in its fields in the source; measured fields-only gravity remainders remained at arithmetic scale. Resolved remainders divided by α^2 stabilized, consistent with the directional Jacobian rather than a persistent first-order derivative error.

At the constrained $n = 3$ seed the original correction still failed the 0.03 forcing gate, attaining 0.0402590; it was not silently replaced by a better direction. At its common cap $\alpha = 0.12194828$, the actual norm was 1.0237652 of the base versus prediction 0.87826014. Gravity accounted for 90.7812 percent of the squared remainder. At one tenth of this cap the trial decreased the residual, and the net curvature contribution was slightly favorable. A large mixed term therefore did not imply that its signed merit projection was always adverse.

The consequence was to keep feasibility and nonlinear path curvature as separate hypotheses. A sampled enthalpy cap could avoid futile trials but would not remove a wall where $h \to 0$ and $\delta h < 0$. Nor could exterior weighting be justified from whole-gravity-block norms alone: the original probes had not yet localized stellar versus exterior contributions.

Evidence. `experiments/PRESSURE_AND_SHAPE_CAMPAIGN_2026-09-17.md`, `experiments/PRESSURE_OBSTRUCTION_METHODOLOGY_2026-09-17.md`, and `experiments/POLYNOMIAL_ENTHALPY_`

`ADMISSIBILITY_2026-09-17.md`. The eight focused and two supplemental cases restored their base residuals exactly; direction partitions and signed-merit identities passed independent checks. No additional virial-qualified model was produced by these probes.

6 Three optional numerical strategies and their limits

The next implementation added three interchangeable numerical approaches: polynomial enthalpy feasibility, gravity curvature correction, and gravity elimination. They were selected by concrete typed policies through `NewtonNumerics`, not added to the physics-facing model declarations. Unregistered or incompatible policies failed compile-time constraints; actual mesh and FE eligibility were checked collectively at runtime. The ordinary Newton path remained available and defaults were unchanged.

6.1 Polynomial enthalpy certification and anchor repair

The diagnostic converted the represented nodal polynomial to a tensor Bernstein representation with outward-rounded interval arithmetic. Adaptive subdivision returned a nonnegativity certificate, a negative-value witness, or unresolved bounds. A negative Bernstein coefficient by itself was not misclassified as a negative polynomial. Exact zero traces were retained. Certification used physical enthalpy coordinates, because general diagonal normalization does not preserve the polynomial represented by the coefficients.

For an admissible base, a fixed certified initial field supplied an anchor:

$$p_h(\beta) = (1-\beta)p_h + \beta(h_{\mathrm{anchor}} - h), \qquad 0 \le \beta \le 1. \tag{15}$$

Sufficient Bernstein inequalities constrained the repair fraction and the full linear residual model selected a fraction. The actual normalized-and-added candidate was certified again, and its complete merit derivative had to remain negative. This was a conservative feasible-direction family, not a constrained Newton optimizer, and it did not make the independent density positive.

At the rotating onset the base certified and the original full trial was proved negative. Whole-mesh certification cost about 3–4 ms. The sufficient repair fraction was 0.9152726503; repaired polynomials prepared successfully, but trial lengths 1, 0.5, and 0.125 all increased the full residual. The linear model's best step was near 0.00825, with only about 0.46 percent predicted residual reduction. The already negative late base correctly returned `invalid_base` and left accepted buffers unchanged. This path supplied useful diagnostic and admissibility machinery, but was not established as an efficient route through the nonlinear obstacle. The later user decision to defer positivity enforcement was consistent with this mixed result.

6.2 Gravity-only curvature correction

Partition the normalized state into retained variables u and gravity variables z, with gravity equations $G(u, z)$. At a raw trial $y = x + \alpha p$, the correction solved

$$G_z(y)\,\delta z = -\big[G(y) - G(x) - \alpha G'(x)p\big]. \tag{16}$$

Only gravity coordinates changed. This removes the nonlinear gravity remainder, not the entire gravity residual. It preserves the base tangent because the right-hand side is second order along a smooth path, and it vanishes on an affine gravity path even when the base residual is nonzero. The candidate inverse was refreshed, every residual was reevaluated, and ordinary full-merit acceptance remained authoritative.

The first frozen $n = 1.5$ comparison changed the full trial residual from $2.18913190375\,10^{-3}$ to $1.40702031008\,10^{-4}$ at base norm $1.78804469158\,10^{-4}$. The gravity remainder fell from

$2.18656200516\,10^{-3}$ to $1.37482284578\,10^{-12}$, with observed correction order 1.999898 on the smallest resolved steps. This was direct evidence that the identified curvature could be removed from gravity.

An initial production attempt was intentionally stopped after repeated auxiliary prepared-consistency rejection despite converged inner solves. The arithmetic allowance had omitted large cancelling gravity contributions. The revised check included separately measured gradient, potential, source, and correction actions; it used 256ϵ times their norm sum as a heuristic allowance, not a rigorous error bound. The true inner gate, lifted full-Jacobian gate, nonlinear tolerance, and physical criteria were not relaxed. The interrupted run and correction history were preserved.

6.3 Gravity elimination

Gravity elimination instead completed gravity at the initial state and each retained trial, then used the reduced Schur action

$$S = F_u - F_z G_z^{-1} G_u. \tag{17}$$

Right-hand-side condensation included the residual gravity defect, and lifting reconstructed a full correction that still had to satisfy the original full normalized linear gate. Tolerance-stopped inner Krylov solves only approximated this linear action; flexible outer preconditioning did not magically make it exact. Tight inner gates, finite-difference checks, and lifted verification were therefore retained.

The reference implementation restricted the full operator and inverse. A capability-checked compact path reused the existing prepared physical gravity actions with the same scaling. On three dense actual-FEM directions compact and full-restricted outputs agreed exactly. Gravity operator applications cost about 0.018 versus 0.244 seconds, around fourteen times faster, while inverse applications changed only from approximately 0.205 to 0.160 seconds. A matched tight inner solve retained forty iterations and the same result while falling from 18.36 to 7.25 seconds. The fourteenfold action gain was not a fourteenfold Newton speedup.

6.4 Production outcomes of the three strategies

Model and path	Accepted updates	Final $\|F\|$	Newton seconds / outcome
$n = 1, J = 0$, curvature	6	$6.51136\,10^{-14}$	273.71 / converged
$n = 1, J = 0$, eliminated	6	$8.98743\,10^{-14}$	869.67 / converged
$n = 1.5, J = 0$, curvature	6	$1.48794\,10^{-5}$	711.72 / outer gate failure
$n = 1.5, J = 0$, eliminated	5	$4.47888\,10^{-6}$	1821.79 / completed updates only
$n = 1, J = 0.1$, polynomial	3	$7.13059\,10^{-3}$	265.03 / outer gate failure
Polynomial plus curvature	1	$8.40387\,10^{-3}$	127.05 / outer gate failure
Polynomial plus elimination	1	$2.72634\,10^{-2}$	149.83 / non-descent repair

The two nonrotating $n = 1$ gravity paths and polynomial control reached the same discrete model to coefficient differences below $2.632\,10^{-11}$. All eight virials were below $3.73092\,10^{-7}$; independent enthalpy certificates covered all 832 stellar elements. The eliminated control used 107 inner solves and 3,875 inner iterations, versus six solves and 115 iterations for curvature correction. It was a verified alternative, but substantially more expensive.

The $n = 1.5$ curvature run reduced the initial norm about eighteenfold but failed its seventh outer solve at relative defect 0.059017 after 200 iterations. All fourteen auxiliary solves passed. The

eliminated run was externally interrupted at 2,400 seconds during its sixth attempt, after five accepted steps of $1, 1/2, 1/4, 1/128, 1/2048$. Its last accepted gravity residual was about $3.27\,10^{-14}$ while retained equations remained of order 10^{-6}; the last step improved the total residual only 0.00565 percent. This showed stagnation despite completed gravity. No final snapshot or complete physical diagnostics existed for the timed-out run, so these were not invented from its log.

Rotating runs remained both unconverged and virially unacceptable. The polynomial-only fourth linear attempt attained 0.0331634 rather than the required 0.03. Adding curvature improved the first step with the same tangent but the second linear solve failed at 0.0826534. Polynomial plus elimination passed its reduced and lifted linear gates before a repaired direction failed descent. A reporting replay later established that this failure occurred before another gravity relift, correcting an earlier ambiguous interpretation of counters. All these models retained negative independent density samples.

The implementation was covered by 43 targeted serial cases and 1,886 assertions at the final revision, with 37 supported cases on each of two MPI ranks. Full stellar multi-rank operation was not thereby certified: logical radial surface interpolation and per-rank material ownership restrictions remained. The frozen checks established algebra and lifecycle invariants; they did not constitute physical convergence.

Evidence. `experiments/NUMERICAL_STRATEGIES_2026-09-17.md`, `experiments/GRAVITY_SUBSYSTEM_IMPLEMENTATION_2026-09-17.md`, `experiments/GRAVITY_REDUCTION_IMPLEMENTATION_2026-09-17.md`, `experiments/GRAVITY_COMPACT_APPLICATIONS_2026-09-17.md`, and `experiments/GRAVITY_PREPARED_ROUNDOFF_2026-09-17.md` retain the complete policy contracts, intermediate failed builds/runs, and replay history.

7 Preconditioner improvements measured at equal accuracy

7.1 Baseline approximation and work accounting

The original inverse first treated surface and material through a triangular surface-then-enthalpy-then-density approximation. Surface used a reference mass diagonal; enthalpy used hydrostatic mass and identity treatment at replaced boundary rows; density included the live EOS enthalpy derivative. Material-to-surface feedback was omitted. Gravity used approximate mixed LDU with a degree-five Chebyshev mass inverse and three fixed AMG cycles for a potential surrogate. Structure-level material and gravity were independent, then a small dense scalar Schur complement completed specification borders relative to that approximate inverse. The live high-n EOS derivative was present; there was no evidence that the baseline simply omitted it.

Frozen benchmark roles each constructed fresh inverses and border responses at the same saved state, full Jacobian, normalization, and right-hand side. Reported complete verified cost included inverse construction, normalized binding, Krylov setup, solve, and an independent full-Jacobian endpoint action. Failed solves were retained but never used as finite baseline time-to-convergence measurements. Timing runs and large audits were serialized.

7.2 Prepared cross-action optimization

Existing stellar LDU was initially unexpectedly expensive: at an $n = 1.5$ plateau, 200 iterations cost 599.5473 seconds versus 87.9302 for the independent baseline, and both failed the gate. The gravity-to-material coupling used a legacy geometry-evaluation path. Replacing it with the already prepared complete gravity-force action, with zero density and displacement variations, preserved the same mathematical gravity-gradient column and subsequent deformation adjoint.

The cached LDU reduced complete cost to 117.4845 seconds, a factor 5.1032 relative to the old

LDU path, and mean inverse application time by 8.0425. The correction changed only $2.7051\,10^{-11}$ relatively and the gate still failed at 0.0600669. This was an execution optimization, not a conditioning improvement. It deserved retention independently of whether LDU was selected as a numerical policy.

7.3 Mass-consistent gravity Schur refinement

Writing the mixed gravity block as

$$D = \begin{bmatrix} M & B^T \\ B & 0 \end{bmatrix}, \qquad D \begin{bmatrix} g \\ \Phi \end{bmatrix} = \begin{bmatrix} f \\ r_g \end{bmatrix}, \tag{18}$$

let R be the same approximate mass inverse used for flux recovery and Q the existing AMG inverse on the diagonal-mass surrogate. The added policy used

$$S_R = BRB^T, \qquad r = BRf - r_g, \qquad \Phi_0 = Qr, \tag{19}$$

$$\Phi_{j+1} = \Phi_j + \omega Q(r - S_R\Phi_j), \qquad g = Rf - RB^T\Phi_k. \tag{20}$$

The Poisson defect was thereby measured against the same approximate mass action as recovery. This was still not an exact inverse of the mixed block. Fixed work preserved stationarity only when the child actions themselves declared stationary linear behavior. Positive relaxation alone did not prove contraction. Each extra sweep cost one application of R, Q, B, B^T; zero sweeps retained the old arithmetic path for bitwise baseline comparison.

One sweep with $\omega = 0.5$ first crossed the original 0.03 gate at the $n = 1.5$ saved state. An equal-accuracy confirmation gave:

Inverse, frozen $n = 1.5, J = 0$	Iterations	True relative defect	Complete cost (s)
Independent baseline	231	0.0290033753	100.9943
Mass-consistent gravity Schur	132	0.0292091550	72.3727

This was a measured 1.39548-fold speedup, 28.34 percent lower complete cost and 42.86 percent fewer iterations, without relaxing accuracy. At a saved $n = 1, J = 0.1$ system, the independent and gravity-only controls still failed after 600 iterations, at defects 0.0739197 and 0.0408506. Combining cached stellar LDU with gravity refinement passed in 142 iterations at defect 0.0299068971, complete cost 96.8236 seconds. It reduced Poisson and gradient defect norms by 63.33 and 94.47 percent relative to the matched 200-iteration independent control. Since the baseline had no finite observed successful solve, no rotating speedup ratio was asserted.

A nine-mode coupled shape coarse inverse was a negative control. All modes were active, overlap/rank checks passed, and there was no fallback, but the 200-iteration defect worsened to 0.0668658. Setup added 8.6012 seconds and persistent vectors occupied 28,491,680 bytes. That failed candidate remained experimental; it did not rule out every possible coarse space.

7.4 Nonlinear transfer limited the recommendation

For easy $n = 1, J = 0$, both independent and gravity-refined inverses converged in six full steps to the same physical state, with all virials below $3.74\,10^{-7}$. Linear iterations fell from 307 to 258, yet Newton time rose from 191.068 to 199.375 seconds, 4.35 percent. More expensive inverse applications outweighed iteration reduction.

For $n = 1.5, J = 0$ with gravity curvature correction, the stronger inverse passed all twelve permitted outer linear gates and all 31 auxiliary gravity solves. It removed the earlier immediate linear failure but ended at $1.3966991\,10^{-5}$, only 6.13 percent below the old six-update endpoint and far above target $2.7365922\,10^{-12}$. All nineteen rejections were merit failures, with late accepted steps 1/64

and 1/512. Newton time was 1075.368 seconds, including 415.038 in outer linear solves and 590.003 in preparation; nested auxiliary time must not be added again. Enthalpy-based virials passed, but negative density persisted.

The rotating combined-inverse follow-through accepted one repaired full step to 0.00685562, then failed the next linear gate at 0.0339685. Its enthalpy energy virial, 0.00712711, was worse than the historical independent endpoint's 0.000544539 despite its smaller nonlinear residual. Both failed physically. Frozen linear gains therefore did not justify claiming a solved nonlinear or physical problem, and the stronger policies remained optional.

Evidence. `experiments/PRECONDITIONER_IMPROVEMENT_2026-09-18.md`, `experiments/GRAVITY_CONSISTENT_SCHUR_2026-09-18.md`, and `experiments/PRECONDITIONER_CACHED_CROSS_ACTION_2026-09-18.md`. Focused serial coverage passed 20 cases and 441,978 assertions; two-rank algebra/coarse tests passed thirteen cases per rank. The actual multi-rank FEM test was blocked by a pre-existing partition where one rank had no vacuum element, required by mapped H(div) construction. Refining the mesh did not remove that ownership pattern. Serial FEM evidence must not be described as distributed FEM validation.

8 Why better linear solves did not remove the nonlinear stall

The next frozen investigation used the saved $n = 1.5, J = 0$ endpoint with $\|F\| = 1.3966991044\,10^{-5}$. All prior linear gates had passed; the final squared-residual shares were about 42.42 percent hydrostatic, 26.23 percent closure, 15.87 percent surface, 10.43 percent gravity, and 5.05 percent mass. It computed fresh cold directions rather than claiming to reconstruct raw directions that the earlier trajectory had not saved.

For $\phi = \|F\|^2/2$ and a forcing condition $\|Ap + F\| \le \eta\|F\|$, a consistent derivative gives

$$\phi'(0) = \langle F, Ap\rangle \le -(1-\eta)\|F\|^2. \tag{21}$$

Thus a gate-passing direction should be strongly locally descending, but this says nothing about how far the linear approximation remains useful.

8.1 Tenfold tighter forcing did not improve the useful step

Measurement	$\eta = 0.03$	$\eta = 0.003$
Iterations	84	143
Cold linear seconds	47.9374	81.6901
Verified relative defect	0.0295592	0.00294840
Direction norm	0.0794897	0.0796012
Corrected residual/base at 1/64	1.080323	1.080567
Corrected residual/base at 1/256	1.002025	1.002037
Corrected residual/base at 1/1024	0.999394139	0.999394272

The tenfold stricter gate cost about seventy percent more linear work and left the useful sampled step essentially unchanged. The directions differed by only 0.8885 percent in norm and had cosine 0.99996157. A supplementary minimum-norm calculation in their two-dimensional span, subject to the original forcing gate, gained only 0.00058549 percent beyond ordinary scalar damping of the tighter direction. This was a saved-vector algebra experiment, not a fresh FEM trial, and did not rule out richer subspaces.

Raw trials at $1, 1/4, 1/16$ failed pressure preparation for this new late direction. At 1/64, raw and gravity-corrected residual ratios were 1.807268 and 1.080323; at 1/256, 1.002510 and 1.002025. At 1/1024 both decreased, to 0.999170 and 0.999394. The remainders were resolved and approximately

quadratic over the small-step range. No persistent first-order error was detected in the tested tangent.

8.2 A large compensated direction amplified moderate curvature

The separate shape and field action norms were 0.7730368 and 0.7730377, while their sum had norm $1.396089\,10^{-5}$. The image cosine was -0.999999999838, and the cancellation factor was approximately 110,743. At the seed, the corresponding factor had been about 332.1. Normalization factors were byte-identical across these states.

The late direction was 17.35 times larger while the residual was 19.59 times smaller. The ratio $\|p\|/\|F\|$ grew from 16.7433 to 5691.26, about 340-fold. At the common step 1/64, corrected curvature per squared direction norm grew only from 0.656881 to 0.926040, about 1.41-fold; curvature relative to the base residual grew about 8,313-fold. This pointed toward weak coupled response and large corrections as the amplification mechanism, rather than orders-of-magnitude growth of curvature per unit direction. It did not identify a continuum symmetry or estimate the full Hessian.

Gravity correction reduced the small-step remainder norm by about 14.7-fold at 1/1024, but increased its adverse projection onto the linear predictor 2.88-fold. Hydrostatic, closure, surface, and mass rows contributed 43.52, 32.33, 19.53, and 4.63 percent of that signed projection. A smaller remainder therefore need not produce better merit. The observed full/shape gain ratio also gave a condition-number lower bound of about $9.07\,10^4$ if the finite-dimensional Jacobian was nonsingular; it was not a measured global condition number or a license to add a central constraint at $n = 1.5$.

Evidence. `experiments/NONLINEAR_STEP_INVESTIGATION_2026-09-18.md`. Three immutable probes covered 144 path observations and sixty successful auxiliary gravity solves; final saved-artifact audits passed 20,899 checks. All states, factors, and restored actions were preserved. Those checks established provenance and arithmetic identities, not a new model.

9 Fully coupled second-order trial corrections

9.1 Bounded correction against the complete residual

Because gravity-only correction moved adverse curvature into retained equations, a fully coupled trial correction was implemented. At a prepared trial y, it preserved the original affine target $L(\alpha) = F_0 + \alpha A_0 p$ and solved

$$A(y)c = L(\alpha) - F(y). \tag{22}$$

The auxiliary backend had its own true full-Jacobian defect check, finite budget, and report. The correction was capped by $\|c\| \le \alpha\|p\|$ by default. It was attempted only after a prepared candidate failed ordinary acceptance, with at most one solve per Newton iteration initially. A failed solve, cap, admissibility test, or merit comparison restored the original trial with fresh preparation. Accepted-state buffers were not mutated by the transaction. The policy composed statically with existing enthalpy/gravity choices and deliberately rejected the unimplemented eliminated-gravity combination at compile time.

At the common late state, the original direction still passed with defect 0.02955918315. Independent frozen trials showed:

Path and α	Before/base	After/base	Auxiliary iterations	Added seconds
Raw, 1/64	1.80726810	1.03360191	34	27.33
Raw, 1/256	1.00250966	0.99746464	52	37.59

Gravity first, 1/64	1.08032296	1.09566246	55	47.60
Gravity first, 1/256	1.00202482	0.99810801	62	43.53

All auxiliary solves passed their unchanged gate and cap. The 1/256 trials became acceptable, while the gravity-first 1/64 trial worsened and was restored. The already acceptable 1/1024 and seed full-step controls incurred no coupled solve. Raw timing came from an uncontended repeat after a mistakenly overlapping artifact audit was stopped and fixed; fifty saved vector files matched the first run byte-for-byte, and the original timing was excluded.

The correction's own curvature remained significant. For $b = L - F(y)$, $\delta = A(y)c - b$, and $E = F(y + c) - F(y) - A(y)c$,

$$F(y + c) - L = \delta + E. \tag{23}$$

At 1/256, $\|E\|$ exceeded the verified auxiliary defect by about 11.28 for the raw path and 254 for the gravity-first path. Correction/direction cosines were -0.9999330 and -0.9999942: much of the correction reversed the original step. Actual displacement increases relative to the old 1/1024 path were only 1.7596 and 1.2342, not the nominal factor four in α.

Fresh scalar controls matched projection along the original direction. They decreased residual 0.115623 and 0.096934 percent, whereas the coupled candidates decreased 0.253536 and 0.189199 percent. Small transverse components therefore had a measured benefit; the improvement was not just scalar damping. It remained a small absolute gain purchased with another coupled solve. Moreover, the independent frozen alphas each had a fresh correction budget, so this screen did not measure a live shared-budget line search.

9.2 Damping and work-aware eligibility

The next optional policy reused one computed c across $\beta = 1, 1/2, 1/4, 1/8$, always preparing $y + \beta c$ from the same y. The first ordinarily acceptable improvement ended the search; otherwise the best admissible improvement could be retained as a trial, with fresh replay. An eligibility rule reserved the single solve for a candidate near the affine prediction. For the configured deterministic merit M, it required positive finite $D = M(F_0) - M(L)$ and

$$\frac{M(F(y)) - M(F_0)}{D} \le k, \qquad k = 1 \text{ by default.} \tag{24}$$

This was a bounded-work heuristic, not a theorem of global convergence.

A sequential halving screen exposed an important correction to the earlier interpretation. Ordinary raw and gravity paths already accepted 1/512; the previous sparse grid had omitted it. The near-acceptance policy instead accepted 1/256, a twofold nominal increase. Without eligibility deferral, the sole solve was spent at an unsuccessful 1/32 and the final state was the baseline. Both successful frozen rescues used $\beta = 1$, so eligibility, not damping, explained that particular improvement.

In uncontended raw repeats, baseline search work cost 45.76 seconds for 0.116541 percent decrease; near-acceptance correction cost 74.90 seconds for 0.253536 percent decrease. The fractional decrease per measured operation-second was 1.329 times greater. Those timings excluded ordinary outer solve, commit refresh, initialization, and physical replay; they were not Newton speedups. A clock-anomalous original gravity/full run was retained and excluded, then repeated with independent UTC, monotonic, and child CPU measurements.

9.3 Matched six-update live test

The live pair kept the seed, mesh, $n = 1.5, J = 0$, absent central constraint, gravity-Schur inverse, restart 160, and $\eta = 0.03$ fixed. Only the trial policy changed. Both first updates reached $1.614841549\,10^{-5}$ with the same recorded diagnostics.

Endpoint after six updates	Gravity curvature control	Coupled policy
Final residual	$1.497197\,10^{-5}$	$1.245830\,10^{-5}$
Newton iteration seconds	502.21	613.09
Outer Krylov iterations	354	341
Auxiliary coupled iterations	0	163
$P(h)$ energy virial	$4.22590\,10^{-6}$	$3.78554\,10^{-6}$
$P(h)$ force virial	$2.71367\,10^{-6}$	$2.31894\,10^{-6}$
Minimum sampled density	$-7.73185\,10^{-4}$	$-7.85522\,10^{-4}$

The new endpoint residual was 16.79 percent lower for 22.08 percent more Newton time. Excluding the common first update, residual decreases were 22.85 versus 7.29 percent, giving 2.53 times more decrease per iteration-second over that bounded interval. Both stopped at the six-update cap; neither converged. The new method initially crossed several modest residual thresholds later than the control before eventually overtaking it.

Only one of three coupled searches was committed. At update two the full correction worsened merit but $\beta = 1/2$ passed, with a very small 0.01795 percent residual decrease. Later ordinary steps benefited from the changed trajectory. Two other kept improvements remained unacceptable and were followed by ordinary backtracking; they were not accumulated into the model. This result supported retaining the policy as an experiment, not promoting it or treating every rescued step as an efficient update. Density-pressure virials remained undefined and local EOS extrema did not uniformly improve.

Evidence. `experiments/COUPLED_TRIAL_CORRECTION_2026-09-18.md` and `experiments/DAMPED_COUPLED_TRIAL_2026-09-18.md`. The latter covered 58 serial numerical/solver cases, 23 correction cases on two MPI ranks, 56 Python checker tests, twelve frozen cases including repeats, and the matched live pair. Its tests covered fallback, exceptions, budgets, rank-local rejection, preserved outer reports, and typed composition.

10 Shape freedom, field transport, and the displacement-order control

The next investigation targeted the weak coupled response itself, preserving the equations and avoiding an assumed gauge interpretation. At independent quadrature orders 10 and 14, scalar pullback changes were decomposed as

$$\delta f_{\text{Eulerian}} = \delta \widehat{f} - \nabla f \cdot v, \tag{25}$$

where v was the generated physical displacement. Density and potential gradients were broken within-element gradients; facet jumps and RT transport were not included. Small Eulerian/pullback ratios therefore suggested transport compensation but could not certify a pure relabeling nullspace.

At the old twelve-update Q_3 endpoint, these ratios were 0.05769 for density, 0.18870 for enthalpy, and 0.01288 for potential, versus seed values 0.3243, 0.3080, and 0.1993. A six-block sampled singular-response minimum fell from 0.0580574 at seed to 0.000143543 late, with largest sampled response near 22.1–22.6. The normalized block direction span itself changed with the state, so this was not a full-Jacobian spectrum or a rank proof.

Despite an enthalpy energy virial of $4.06909\,10^{-6}$, the late nonrotating nodal surface radii spanned 0.997674–1.010093. Its direction was 99.9995107 percent cubic-invariant in sampled squared norm, while degree-eight spherical harmonics left 0.893679 of its nodal norm unexplained. This was fine angular structure consistent with mesh/discretization coupling, not evidence that all omitted modes were unphysical. The bulk and gravity discretizations also shared cubic symmetry.

10.1 The isolated Q_4 displacement experiment

The original Q_3 displacement could not exactly represent the order-four reference coordinate map: the affine identity control gave relative value error $9.01277\,10^{-5}$ and gradient error $1.23472\,10^{-3}$. On an isolated branch, only `libmeanfield/interface/field/field_registry.cppm`, `Displacement::vectorOrder`, changed from 3+ uniform increment to 4+ uniform increment, with the increment fixed at zero. The rationale and date were recorded in source. The shared scalar surface collection therefore also changed from Q_3 to Q_4, increasing surface unknowns from 866 to 1538. Density/potential Q_2, enthalpy Q_3, RT2 gravity, geometry, and the radius-squared lift were retained.

This was not a pure interpolation intervention: it added shape modes, changed surface normalization/preconditioning, increased derived pressure quadrature by one and gravity-force quadrature by two, and raised the geometry-inspection request from 18 to 24. Both variants started from analytic seeds; incompatible coefficient snapshots were not transferred.

Q_4 reduced affine value and gradient errors to $2.11629\,10^{-16}$ and $3.24701\,10^{-15}$. Nevertheless, seed physical displacement direction norm grew 31.63-fold, from 0.00121967 to 0.0385749, and seed linear iterations grew from 26 to 67. Both initial linear gates passed, and small trial remainders remained quadratic. The fixed six-attempt live comparison was unfavorable:

Measurement	Q_3 displacement/surface	Q_4 displacement/surface
Root unknowns	178,073	178,745
Accepted updates	6	3
First three accepted-step seconds	269.31	505.65
Residual after three / own initial	0.056755	0.906188
All attempted Newton seconds	547.34	638.42
Krylov iterations, including failure	354	492
$P(h)$ energy virial	$4.22590\,10^{-6}$	$8.89666\,10^{-5}$
$P(h)$ force virial	$2.71367\,10^{-6}$	$8.27467\,10^{-6}$
Hydrostatic-gradient RMS	0.0106811	0.325413
Nodal radius span / nodal mean	1.09849%	12.2559%

The fourth Q_4 solve and an independent cold endpoint solve both failed after 200 iterations at relative defect 0.0868564. Even its first three gate-passing corrections had already produced poor progress. The differing spaces and stopped endpoints prevented an equal-accuracy comparison; the result nevertheless ruled out matching displacement order alone as a sufficient cure in this configuration. The Q_4-only change was explicitly not recommended for promotion. It did not rule out compatible enrichment of multiple fields.

10.2 Two structural cautions

A naive replacement of mechanical shape rows by surface Bernoulli moments was identified as rank-deficient at an ideal $B = 0, J = 0$ root for the current trace spaces. On the actual closed surface, continuous Q_p had $96p^2 + 2$ degrees of freedom, versus $96p^2$ for broken Q_{p-1} potential traces. After enforcing the zero enthalpy trace, Q_3 surface supplied 866 rows against 864 DG Q_2 trace coefficients; C added no independent constant trace dimension and the direct shape-measure derivative vanished with B. This was an objection to that proposed replacement, not proof of rank deficiency in the existing mechanical operator. Any alternative boundary formulation needed compatible spaces and a stability analysis.

Second, correcting ambient interpolation did not make the radius-squared extension generate a physical dilation. If b was the boundary trace of reference position X, the existing lift $E_0(b)$ differed from X; at a logical radius 1/8 vertex it gave one eighth of the desired affine value. A proposed but unimplemented rank-one correction was

$$E(z) = E_0(z) + \big[X - E_0(b)\big]\ell(z), \qquad \ell(b) = 1. \tag{26}$$

It would give $E(b) = X$ while retaining the supplied trace and admit an exact transpose. This could place physical dilation in the mechanical test space, but would not automatically enforce the potential-energy virial, which uses the independent mixed potential. It was kept separate from surface-space reduction and was not claimed as an implemented remedy.

Evidence. `experiments/SHAPE_FIELD_WEAK_RESPONSE_2026-09-18.md`. The Q_4 production change was isolated at commit `4be1a23c3ec2650af6a2a49528a8f5248cb001c3`; 22 targeted cases and 380,856 assertions passed. That correctness coverage did not convert its negative numerical outcome into a recommendation.

11 The bounded smooth-surface study

The user authorized a one-hour study, recorded as 19:30:15–20:30:15 UTC on 18 September. It retained Q_3 displacement, Q_4 reference geometry, Q_2 density/potential, Q_3 enthalpy, RT2 gravity, the same extension, $n = 1.5$, $J = 0$, and no central constraint. It changed only independent surface variation through a frozen affine embedding

$$q = q_0 + Ba. \tag{27}$$

The existing base q_0 was never projected. Real spherical harmonics through degree L were evaluated at actual radial surface nodes and their interpolants orthonormalized in a quadrature-defined physical surface-area metric. All azimuthal components in that harmonic span were included.

Correct primal and dual normalized maps were essential. For physical state scales S and residual scales R, separate thin QR decompositions gave

$$SB = Q_x T_x, \qquad R^{-1}B = Q_f T_f, \tag{28}$$

so that

$$Q_f^T(RF_{\text{physical}}) = T_f^{-T} B^T F_{\text{physical}}. \tag{29}$$

The reduced operator was $Q_f^T \widehat{A} Q_x$, not a blind transpose of a normalized state basis. Other fields and scalar equations remained active. The projected existing inverse was an approximation, not an exact reduced inverse. Both retained and original full residuals, and excluded shape force, were always exported.

11.1 Frozen late-state result

At the common six-update Q_3 endpoint, $\|F\| = 1.4971971\,10^{-5}$, cold FGMRES used restart 160, cap 200, $\eta = 0.03$, and one gravity-Schur sweep with relaxation 0.5.

Shape space	Parameters	Iterations	Solve seconds	Unit-trial full norm
Full nodal trace	866	61	32.184	Pressure rejection
$L = 2$	9	12	6.447	$2.263903\,10^{-6}$
$L = 4$	25	21	11.249	$2.432960\,10^{-6}$
$L = 8$	81	29	15.745	$2.964119\,10^{-6}$

All reduced directions passed their retained linear gates, with relative defects 0.01719, 0.01608, and 0.01409. Their original full linear defects were 0.15121, 0.16187, and 0.19207: none passed the original full Newton gate. Nevertheless, all four sampled trial alphas prepared and decreased both norms, and full residual decreases at unit step were 84.88, 83.75, and 80.20 percent. This was a real improvement in the original residual, obtained by restricting the shape equations, not a more accurate solution of the original Newton system.

The $L = 4$ root had 177,232 unknowns, only 841 fewer than the full root; its less than 0.5 percent dimension reduction did not explain a factor 2–5 drop in cold solve time by itself. Fewer Krylov iterations supplied most of the observed gain. Physical displacement norm fell from $3.96028\,10^{-3}$ for the full direction to $9.14977\,10^{-5}$ at $L = 4$, 43.28-fold. The $L = 2$ direction was nearly field-only, while admitting higher modes increased coupled displacement and curvature.

After the $L = 4$ unit trial, excluded shape force accounted for 98.38 percent of the remaining original squared norm. The inherited fine-scale geometry also remained, since q_0 had not been projected. Enthalpy-based energy and force virials improved to $8.11913\,10^{-7}$ and $6.41398\,10^{-7}$, but local EOS/hydrostatic errors were almost unchanged and minimum sampled density remained $-7.27628\,10^{-4}$.

11.2 Three cold reduced updates from the analytic seed

A conditional extension, recorded before launching, chained three saved unit candidates from the analytic seed. Each stage passed preparation, the retained independent linear gate, and both recorded norm-decrease screens. This was an externally replayed cold reduced-Newton chain, not a production warm-start trajectory or a fully converged model.

Updates	Full norm	Retained norm	Excluded shape norm	Energy virial $P(h)$
0	$2.736592\,10^{-4}$	$2.590732\,10^{-4}$	$8.815010\,10^{-5}$	$2.410694\,10^{-5}$
1	$8.624827\,10^{-5}$	$9.387845\,10^{-6}$	$8.573583\,10^{-5}$	$2.180785\,10^{-6}$
2	$8.473526\,10^{-5}$	$3.147803\,10^{-6}$	$8.467677\,10^{-5}$	$8.660498\,10^{-7}$
3	$8.465191\,10^{-5}$	$7.357283\,10^{-8}$	$8.465188\,10^{-5}$	$6.769653\,10^{-7}$

The retained norm fell about 3521-fold but the original norm only 3.23-fold. At the final trial, excluded shape force represented 99.999924 percent of the original squared norm. Cold solves took 18, 20, and 18 iterations and 10.574, 11.341, and 9.594 seconds. Full relative linear defects were 31.45, 98.18, and 99.90 percent; these were not full-system Newton solutions.

The final force virial was $5.769064\,10^{-7}$ and mass error $3.79142\,10^{-9}$. Yet hydrostatic-gradient RMS rose from 0.001528098 at seed to 0.007255981, about 4.75-fold, while density remained negative at some samples. Nodal radius spread grew to 0.154699 percent, substantially less than unrestricted late distortion but not an exact spherical result. Small retained residual and excellent virial could therefore coexist with significant omitted mechanical force and imperfect interior fields.

The result supported decoupling smooth independent surface parameters from the ambient displacement carrier, with nested configurable spaces and consistent dual restriction. It did not establish that a fixed degree-four space was sufficient at all rotation rates, that excluded force was irrelevant, or that the complete original mechanical system had converged. The subsequent native spherical-harmonic implementation and higher-index studies should be read as later developments, not retroactively attributed to this bounded experiment.

All seven case audits passed 2,098 saved-artifact checks; a separate chain audit passed 107 checks with bitwise physical-state and residual handoffs. Actual builds, FEM cases, and expensive audits were serialized, and all seven FEM probes completed by 20:09:58 UTC. No production order, operator, or lift changed in this study. Optional $L = 6$, rotating cases, higher indices, mesh/enrichment sweeps, and production integration were outside its scope.

Evidence. `experiments/REDUCED_SURFACE_STUDY_2026-09-18.md`, producer commit `21cdbda22486e408adc52b65b25413f9380c7c30`, and the archived `working/research/diagnostics/reduced_surface_2026-09-18` campaigns. Together with the failed Q_4 experiment, this was the main reason to prefer a controlled smooth surface span over indiscriminate addition of nodal surface freedom.

12 Native spherical-harmonic surface prescription

The reduced-surface experiments motivated making spherical harmonics a native, selectable surface prescription. The implementation on `research/spherical-harmonic-surface-20260919` preserved the simultaneous Newton solve: surface coefficients, material fields, gravity fields and global multipliers remain coordinates of one coupled nonlinear system. This was not a separate surface-relaxation iteration. The preferred research candidate became maximum degree $L = 4$ with all real azimuthal orders, hence 25 independent coefficients. The unrestricted nodal radial prescription remained available as a control. The choice of $L = 4$ was an experimentally useful starting point, not a demonstrated resolution requirement for 50% critical rotation.

At reference boundary node i, with fixed radial direction $\widehat{r}_i$, the radial displacement is

$$q_i = \sum_{\ell=0}^{L} \sum_{m=-\min(\ell,M)}^{\min(\ell,M)} a_{\ell m} Y_{\ell m}(\widehat{r}_i).$$

Omitting the optional azimuthal cutoff retains $M \geq L$. The real harmonics are normalized to unit squared integral over solid angle, include the Condon–Shortley phase, and use negative m for sine and positive m for cosine. The catalog is ordered by increasing ℓ and then signed m. Coefficients have dimensions of length. Zero coefficients preserve the reference mesh; $\ell = 0$ is a genuine monopole rather than a removed gauge.

Configuration	Independent coefficients
$L = 4$, all m	25
$L = 4$, $\lvert m\rvert \leq 2$	19
$L = 4$, $m = 0$	5
$L = 8$, all m	81

No odd-degree or nonaxisymmetric modes are silently excluded. Choosing $m = 0$ constrains the shape about the fixed global z axis; it does not impose axisymmetry on the three-dimensional interior fields. The reference center is configurable, but the frame does not move as another Newton unknown.

12.1 Architecture and normalization

The prescription implements the existing forward, transpose and pullback contracts, with a cached reference basis and zero derivative of that fixed pullback. It composes with the existing interior and fixed-infinity vacuum extensions. The numerical discretization now owns typed surface, interior and vacuum prescriptions, instead of the modern equilibrium compiler hardcoding a nodal surface. Existing two-argument discretization calls preserve their nodal behavior. The physics-facing EOS, pressure boundary, mass and angular-momentum specification is unchanged.

Prescription/extension compatibility and coupled operator structure remain compile-time constrained. Requested degree, sampled rank and actual mesh resolution are runtime facts. Sampled aliasing is rejected rather than silently reducing the mode set. Modal normalization and the default material/surface inverse use a diagonal of the pulled-back reference vector trace mass matrix; state and residual factors retain their primal/dual relationship. Metric assembly checks FEM identity and a canonical trace-layout capability. The older nodal H^1 surface surrogate explicitly rejects harmonic coordinates.

Geometry remained Q4 and displacement remained Q3. The harmonic samples feed the existing finite-element trace, so the resulting surface between nodes is not an exact analytic spherical-harmonic graph. Reference-mesh interpolation error remains. The native implementation also differs from the earlier frozen reduced-space experiment, which retained an arbitrary nodal base

shape and used area whitening and QR maps. Physical residual and Jacobian equivalence is the appropriate comparison; identical normalized Krylov trajectories are not expected from those different coordinate metrics.

The implementation was single-rank, consistent with the radial volume extension. Successful two-rank rejection tests do not establish distributed harmonic-solver support.

12.2 Native coupled results

Five live cases used the same mesh, cold Lane–Emden initialization, no central constraint, a six-update budget, restart 160, at most 200 linear iterations and a true relative linear gate $\eta = 0.03$. Nonlinear convergence required $\|F\|/\|F_0\| \le 10^{-8}$, with zero absolute tolerance. All 29 accepted updates were full steps, with one trial each and passing linear gates.

Case	Updates	Final $\|F\|$	$\|F\|/\|F_0\|$	Newton s
$n = 1, J = 0$, all m	5	$1.52715\,10^{-12}$	$9.64952\,10^{-9}$	146.00
$n = 1.5, J = 0$, all m	6	$1.29965\,10^{-12}$	$5.01751\,10^{-9}$	221.40
$n = 1, J = .01$, all m	6	$3.86299\,10^{-13}$	$4.29274\,10^{-10}$	197.23
$n = 1, J = .01$, $m = 0$	6	$3.08442\,10^{-13}$	$3.42757\,10^{-10}$	164.90
$n = 1.5, J = 0$, gravity-Schur	6	$2.19823\,10^{-11}$	$8.48662\,10^{-8}$	189.54

The first four cases converged by their stated criterion. The gravity-Schur case exhausted its budget, although its last update still reduced the residual by approximately a factor of 40; calling that endpoint a plateau would be incorrect. Process wall times, including setup and observations, were 210.55, 290.70, 265.89, 232.54 and 257.49 seconds, respectively. Timings are single-run observations on the workstation, not replicated benchmark statistics.

At q18, absolute energy/force virial errors formed with $P(h)$ were $3.87036\,10^{-7}$ / $2.65938\,10^{-7}$ for the nonrotating $n = 1$ case, $6.55153\,10^{-7}$ / $6.00565\,10^{-7}$ for $n = 1.5$, and approximately $3.87\,10^{-7}$ / $2.66\,10^{-7}$ for both low-rotation $n = 1$ cases. They also passed 10^{-5} at q14. The corresponding density-pressure virials passed for the sampled $n = 1$ cases, but were undefined for $n = 1.5$, whose independent density was negative at some integration points. The later sign-region investigation also detected a thin $n = 1$ negative layer missed by ordinary volume quadrature.

Global balance did not establish pointwise accuracy. For nonrotating $n = 1.5$, the normalized EOS-density RMS was $1.25551\,10^{-4}$ and the hydrostatic-gradient RMS increased from $1.52810\,10^{-3}$ in the seed to $7.27475\,10^{-3}$ at the root, a factor of 4.76. This gradient diagnostic has acceleration units in the model normalization; it is not the normalized weak residual. The four converged cases had mass errors below $5.3\,10^{-14}$ and angular-momentum errors below $3.6\,10^{-16}$.

The full nodal mechanical force remained almost entirely outside the harmonic span: Euclidean nodal-dual norms were approximately $9.22\,10^{-6}$ for $n = 1, J = 0$ and $8.99\,10^{-6}$ for $n = 1.5, J = 0$. In the latter case the retained modal force norm was only $3.61\,10^{-14}$. This is convergence of the selected Galerkin shape problem, not convergence of the former unrestricted nodal shape problem. The excluded dimensional Euclidean norm cannot be directly compared with a normalized nonlinear norm or the dimensionless virial tolerance.

Nonaxisymmetric coefficients persisted in nominally spherical models. The $(\ell, m) = (4, 4)$ coefficients were $8.97263\,10^{-5}$ for $n = 1$ and $9.23867\,10^{-4}$ for $n = 1.5$, compared with $a_{40} = 1.09313\,10^{-3}$ in the latter. These are evidence to examine mesh/discretization imprint, not evidence of a physical nonaxisymmetric branch.

Total Krylov counts were 212, 335, 296, 219 and 206 in the table's order. The $m = 0$ rotating control used 26.0% fewer Krylov iterations and 16.4% less Newton time than the all-m control, but changes the admissible shape and is not a pure preconditioner speedup. All m remained the preferred three-dimensional candidate. Ten new harmonic tests and 64 selected existing regression cases passed; independent audits checked saved basis pullbacks, excluded force and physical/nodal

equivalence. Those implementation checks do not validate the higher-rotation target.

13 Negative density: a located discrete representation defect

The subsequent investigation was observational: it changed no equilibrium equation, field order, shape prescription, normalization, preconditioner or Newton policy, and solved no new equilibrium. It examined four frozen states: converged $n = 1$ and $n = 1.5$ nonrotating harmonic models, the original $n = 1.5$ seed, and the $n = 3$ nonrotating seed. Only the $n = 3$ case had a central condition; it was not a converged $n = 3$ model.

13.1 What the weak EOS equation determines

Density and enthalpy are independent unknowns. Baseline density is DGQ2 and enthalpy is continuous Q3. On a stellar element the closure is

$$R_i = \sum_q w_q \varphi_i(q) \left[\rho_h(q) - \rho_*(h_h(q))\right], \qquad \rho_*(h) = \left[\frac{h}{(n+1)K}\right]^n,$$

with physical mapped weights. At zero closure residual,

$$M\boldsymbol{\rho} = \boldsymbol{b}(h), \quad M_{ij} = \sum_q w_q \varphi_i \varphi_j, \quad b_i = \sum_q w_q \varphi_i \rho_*(h_h).$$

Thus the closure selects a quadrature-defined weighted L^2 projection, not pointwise EOS equality. Its unique polynomial can have negative regions even when the sampled EOS target is nonnegative.

Independently assembled element mass matrices and Cholesky solutions reproduced the converged $n = 1.5$ density coefficients within $1.40\,10^{-14}$ and the field within $2.25\,10^{-15}$ in physical L^2, while its negative density was of order 10^{-4}. Independent/production closure-residual differences were approximately $1.2\,10^{-17}$ to $2.1\,10^{-17}$ across the four states. Raising projection quadrature to q26 changed coefficients by at most $2.36\,10^{-7}$ and did not remove the undershoot. These comparisons identify the defect at the measured discrete root; it is not explained by incomplete closure convergence or the small measured quadrature difference.

The implemented dependency also corrects an earlier intuitive explanation: negative density does not directly generate the represented enthalpy. Pressure is evaluated from independently represented enthalpy; negative sampled enthalpy is rejected during preparation. All negative-density q18/q26 volume samples in this investigation had positive enthalpy. Some surface witnesses had exactly zero enthalpy. Density influences enthalpy indirectly through the coupled stellar equations.

13.2 Spatial location and root-based integration

All 96 surface-adjacent stellar cells contained measured negative regions; no negative intervals were found on the sampled lines through the other 736 stellar cells. All stored density nodal values nevertheless remained positive. The nodes are interior Gauss–Legendre points, and positivity of their values does not imply positivity between them.

Ordinary quadrature can miss thin sign regions: q18 and q26 missed the $n = 1$ volume layer, while its sampled boundary trace was negative. For the same unchanged $n = 1.5$ field, negative-mass estimates changed from $2.89\,10^{-5}$ at q18 to $2.18\,10^{-5}$ at q26 and the apparent negative volume nearly halved. The investigation therefore isolated zero crossings of the polynomial along reference-cell lines and integrated negative intervals. Define $M = \int \rho_h \, dV$ and $M_- = -\int \min(\rho_h, 0)\, dV$. The mass denominator below is $|M|$, not $\int |\rho_h| \, dV$.

Frozen state	$M_-/\lvert M\rvert$	V_-/V	Measured physical radius range
$n = 1$ equilibrium	$6.09585\,10^{-8}$	0.0327346%	.999643–.999886
$n = 1.5$ seed	$3.66771\,10^{-5}$	1.83795%	.993813–1.000019
$n = 1.5$ equilibrium	$2.24535\,10^{-5}$	1.49252%	.994917–1.001222
$n = 3$ seed	$2.88372\,10^{-5}$	12.0835%	.934409–.978469

Radii are in reference-radius units and are not distances from the local, slightly deformed surface. The first three cases have intervals reaching the surface; the $n = 3$ seed has an interior pocket in the same outer cell layer. Its surface trace is positive. Minimum-density witnesses were approximately $-1.14756\,10^{-4}$, $-9.26332\,10^{-4}$, $-7.21171\,10^{-4}$ and $-8.54870\,10^{-5}$, respectively; the $n = 3$ witness had enthalpy 0.0455581. These are sampled witnesses, not certified global minima.

Tangential rules with 8, 12 and 16 points per direction agreed on the negative integrals to roundoff in these cases. The finest two negative masses differed by at most 10^{-14} relative, and whole mass/volume agreed with separate observations within $1.77\,10^{-13}$ relative. Eight points integrated each root-partitioned interval. Analytic controls, decimal-arithmetic audits and state/coefficient hash checks support the measurements. They are not rigorous three-dimensional sign certificates: unsampled angular pockets or roots below reconstruction uncertainty remain possible.

13.3 Mechanism controls and their limits

The initial seed independently interpolates tabulated Lane–Emden density and enthalpy into their respective spaces. The converged closure instead projects the EOS image of represented enthalpy. Both operations can undershoot; neither Newton nor rotation is necessary. With inward surface-cell coordinate $t \in [0, 1]$, a linear enthalpy gives target density proportional to t^n. Exact quadratic L^2 projections in a flat one-dimensional cell are

$$\Pi_2 t^{3/2} = -\frac{2}{105} + \frac{16}{35}t + \frac{4}{7}t^2, \qquad \Pi_2 t^3 = \frac{1}{20} - \frac{3}{5}t + \frac{3}{2}t^2.$$

The fractional profile is negative for $0 \le t < 0.0396969$. The cubic profile has a negative pocket $0.118350 < t < 0.281650$, with minimum -0.01 at $t = .2$. Both preserve the exact cell mean and admit positive nodal samples. The cubic example demonstrates that a nonpolynomial target or fractional EOS is not necessary: insufficient degree alone can produce the pocket.

For $n = 1$ the leading linear surface model is exact in Q2, but the actual EOS image of Q3 enthalpy need not lie in Q2. Geometry weights do not destroy exact projection of a target already in the trial space. The measured $n = 1$ layer therefore cannot be attributed to curved weights alone. In the toy fixed-slope model, halving cell width reduces undershoot amplitude by 2^{-n} and negative mass per unit area by $2^{-(n+1)}$, without removing the fraction of that cell that is negative. These formulas explain a possible mechanism and scaling; they are not measured three-dimensional refinement rates.

The conclusions are specific. Tighter solves or better preconditioning cannot remove negativity while converging to this same discrete root. Exact algebraic elimination $\rho = M^{-1}b(h)$ reproduces the same polynomial; it does not turn weak closure into pointwise EOS density. A basis change alone also preserves the polynomial. Conversely, this investigation did not prove that all nonlinear plateaus, the hydrostatic-gradient discrepancy or external-domain errors share this cause.

An early diagnosis report proposed pointwise EOS-derived density as a possible barotropic remedy. The user explicitly rejected making density a derived field because the eventual code targets a general EOS, and subsequently deferred positivity intervention. That recommendation is historical and superseded. The current direction retains independent density and reports the defect without clipping or post-solve repair.

14 Independent density-order study and logarithmic assessment

The next controlled experiment varied density alone: DGQ2, DGQ3 and DGQ4 in the same Gauss–Legendre family. Enthalpy/displacement stayed Q3, gravity RT2/Q2, geometry Q4, and the surface all-m $L = 4$. There were 832 stellar elements, so the density block had 22,464, 53,248 or 104,000 coefficients. The default remained Q2. A compile-time registry option propagated order through typed field spaces, quadrature, adapters, normalization and block layouts; no extra physics-facing parameter was needed.

The closure rule must cover both its density mass form and EOS source form. The selected rule is the greater resolved request, preserving policy overrides. For these cases Q2/Q3 used identical achieved orders: 17, 19 and 23 for $n = 1, 1.5, 3$, with 729, 1000 and 1728 points per cell. Q4 used 19, 21 and 25, with 1000, 1331 and 2197 points. Thus the coupled Q4 comparison includes policy-selected quadrature enrichment; only the frozen common-rule comparison strictly isolates approximation order.

14.1 Frozen approximation results

Saved enthalpy and geometry were fixed for the $n = 1$ and $n = 1.5$ equilibria and $n = 3$ seed. Each density space received an independent physical-volume-weighted projection using common q26 quadrature. EOS error was checked with a separate q30 rule; no projection was installed in a solver state. Negative measures use the finest root-sliced tangential rule. The EOS L^2 error has model density and volume units, without further relative normalization.

Frozen target	Density	$M_-/\lvert M\rvert$	V_-/V	EOS L^2 error
$n = 1$ root	Q2	$6.09585\,10^{-8}$	$3.27346\,10^{-4}$	$1.25266\,10^{-4}$
	Q3	$1.05\,10^{-31}$	$2.49\,10^{-16}$	$1.56\,10^{-15}$
	Q4	$1.25\,10^{-31}$	$2.73\,10^{-16}$	$1.60\,10^{-15}$
$n = 1.5$ root	Q2	$2.24299\,10^{-5}$	.0149181	$3.67965\,10^{-4}$
	Q3	$9.33602\,10^{-6}$	.0104476	$1.12325\,10^{-4}$
	Q4	$3.18362\,10^{-6}$	.00683440	$4.80663\,10^{-5}$
$n = 3$ seed	Q2	$2.92063\,10^{-5}$	.121064	$1.06409\,10^{-2}$
	Q3	$2.63486\,10^{-6}$	.0373200	$6.86223\,10^{-4}$
	Q4	$1.02066\,10^{-8}$	.0169628	$3.84980\,10^{-5}$

The $n = 1$ Q3/Q4 sign measures are unresolved roundoff, not measured physical negative regions. Q3 contains the linear EOS image of Q3 enthalpy, explaining its roundoff-level mismatch. For $n = 3$, the generic target is Q9, and for $n = 1.5$ it is nonpolynomial. Q4 reduces the $n = 3$ projected negative mass by approximately 2862 times relative to Q2 but leaves measured negativity in about 1.70% of the volume. Enrichment is useful approximation evidence, not a positivity guarantee.

For $n = 1.5$ Q4, q26/q30 EOS norms differ by about 0.18%; the fixed-eight-point root-sliced norm differs from the unsliced q30 norm by about 2.2%. Stable negative masses under tangential refinement do not establish normal-direction integration accuracy for the fractional EOS error. This distinction motivated reporting the unsliced q30 norm.

14.2 Coupled controls and the stalled $n = 3$ trajectories

Twelve live cases retained the same default inverse, normalization, cold initialization, restart 160, 200-iteration linear cap, $\eta = .03$, eight-update budget and relative nonlinear tolerance 10^{-8}. Only the $n = 3$ cases used a central condition. All nine $n = 1$ and $n = 1.5$ cases converged with full steps and no rejected trials. All their $P(h)$ energy/force virials were below $6.6\,10^{-7}$.

Case	Density	Updates	Krylov iterations	Newton s
$n = 1, J = 0$	Q2	5	212	147.5
	Q3	6	286	204.3
	Q4	6	286	215.1
$n = 1, J = .01$	Q2	6	296	196.1
	Q3	6	305	217.8
	Q4	6	304	229.1
$n = 1.5, J = 0$	Q2	6	335	222.0
	Q3	6	329	241.7
	Q4	7	397	303.7

For $n = 1$, Q3/Q4 removed resolved undershoot and reduced EOS mismatch to roundoff. Nonrotating total costs were 39%/46% above Q2, but the enriched runs needed one extra update because their fifth residuals narrowly missed their own thresholds. The equal-six-update rotating case cost approximately 11%/17% more. Total costs therefore cannot be read as per-action order overhead.

For $n = 1.5$, Q3 reduced negative mass and EOS error by factors 2.40 and 3.28 for 8.9% more Newton time; Q4 reduced them by 7.01 and 7.66 for 36.8% more time. Negative volume fractions remained 1.493%, 1.046% and 0.685%. The hydrostatic-gradient RMS values were .00727475, .00731970 and .00729901; density enrichment did not appreciably improve that local balance measure.

No $n = 3$ case converged within eight updates:

Density	$\|F\|/\|F_0\|$	Min. α	Newton s	$\|v_E^h\|$	$\|v_F^h\|$
Q2	.430062	1/64	301.6	$1.77163\,10^{-4}$	$1.34597\,10^{-4}$
Q3	.985301	1/256	478.8	$8.07504\,10^{-6}$	$2.28169\,10^{-5}$
Q4	.987949	1/512	508.2	$1.37586\,10^{-5}$	$2.84372\,10^{-5}$

Krylov totals were 221, 452 and 458. All reported true linear gates passed. Q2 had four pressure and four insufficient-decrease rejections; Q3/Q4 each had five pressure and nine insufficient-decrease rejections. Density order does not enter the pressure-force sampling rule, whose operands remain enthalpy and displacement gradient. These differences reflect reached states/directions, not changed pressure sampling. Q3's passing energy virial alone does not make it successful: its force virial and nonlinear criterion fail.

Unbordered residuals were $1.33327\,10^{-3}$, $7.18640\,10^{-4}$ and $7.26009\,10^{-4}$; normalized central forcing was $3.17039\,10^{-5}$, $1.27669\,10^{-5}$ and $1.20966\,10^{-5}$. The unbordered diagnostic removes artificial central forcing but retains scalar constraint rows. Neither a physical root nor vanishing artificial force was demonstrated. The existing central condition pins central enthalpy to the EOS value at requested central density, not the independent DG density point value.

The production seed is $I_p(\rho_{\mathrm{LE,tabulated}})$, whereas the frozen controls are $I_p(\rho_*(h_h))$ and $\Pi_p(\rho_*(h_h))$. Nonlinear EOS evaluation does not commute with tabulation or interpolation. For example, the production Q4 $n = 3$ seed had approximately $5.26\,10^{-4}$ EOS error at q18, versus $3.85\,10^{-5}$ for the frozen Q4 projection at q30. These are different fields and quadratures, not a controlled improvement ratio. A closure-consistent initial guess remains a distinct hypothesis compatible with independent density.

Each order passed 48 focused contract tests (9766, 9822 and 9894 assertions), and the frozen/coupled artifact audits passed. Across-order residual norms belong to different discrete spaces and Riesz weights, so lower normalized residual does not rank physical accuracy. The endpoints are finite-budget outcomes, not proved floors or nonexistence results.

14.3 What the logarithmic assessment established

The separate log-density branch contained derivations and one-dimensional controls only. No stellar log-density solve, production representation change, speedup or virial improvement was established. Three constructions must remain distinct:

1. Exponentiating density nodal coefficients gives $\rho_h = \sum_j \rho_s e^{u_j}\varphi_j$. It changes coordinates on positive coefficient vectors but can retain the same between-node negative polynomial.
2. A pointwise exponential $\rho = \rho_s e^{u_h(\xi)}$ is positive in exact arithmetic for finite u_h, with $\delta\rho = \rho\,\delta u_h$. It changes the approximation family and every density-dependent physical evaluation.
3. A factored representation $\rho = \rho_{\text{ref}}(\xi, q)e^{u_h}$ can preserve an exact vacuum trace if the reference is positive inside and vanishes at the surface. Its variation includes $e^{u_h}\delta\rho_{\text{ref}}$, in addition to $\rho\delta u_h$.

Raw logarithmic density for $\rho \sim As^n$ behaves as $\log(A/\rho_s) + n\log s$ and is singular at an exact vacuum boundary. Interior nodes avoid evaluating $\log 0$ but do not give a finite polynomial log field the correct zero trace. A reference $\rho_{\text{ref}} \sim Bs^m$ makes the log amplitude bounded only when $m = n$. Such an exponent is a polytropic boundary assumption, not a general-EOS prescription; taking the reference from live enthalpy would introduce an additional dependency that is not required by independent density.

A Q2 pointwise exponential fitted to three exact moments of t^n gave:

n	Fitted surface density	Relative L^2 error	Tiny Newton updates
1	.0673707	.0578060	3
1.5	.0229264	.0517531	4
3	.00141091	.0333959	5

Verified moment defects were at most $2.38\,10^{-14}$. Thus positivity, mass conservation and weak closure can coexist with an incorrect positive vacuum trace. Exact Q2 polynomial projection errors for the same profiles were 0, .00952381 and .05, so the exponential family was not uniformly more accurate at the same degree. A matched factor $t^n e^0$ reproduces the toy target by construction, not by demonstrating a general reference-selection method.

With equal test/log bases the local closure derivative is the weighted mass matrix $\int \varphi_i\rho\varphi_j\,dV$. For $\rho = A\Delta^n t^n$ it scales as $A\Delta^{n+1}H_n$, against ΔI for an orthonormal polynomial density mass matrix. Cell-matrix condition numbers in the Q2 one-dimensional exact-profile controls were approximately 7.873, 16.852 and 94.561. These are not stellar Jacobian condition numbers. Small density reduces the response scale and can increase disparities across cells/blocks; logarithms do not automatically improve conditioning or enthalpy admissibility.

A future implementation would require a typed physical-density evaluator and exact directional action, consistent support in every operator, normalization, quadrature and cache. Zero inactive exterior coefficients must not be exponentiated into positive exterior density. Underflow/overflow needs an explicit policy rather than silent clipping. MESA's zonewise use of $\exp(\log\rho)$ supports investigating independent density coordinates, but its atmosphere boundary and zone representation do not establish correctness for this high-order exact-vacuum FEM problem. Following user direction, positivity work was deferred in favor of coupled convergence.

15 Detailed coupled-convergence campaign

The final campaign retained Q2 independent density, Q3 enthalpy/displacement, RT2/Q2 gravity, Q4 geometry, the same mesh/normalization and all-m $L = 4$ surface. No clipping, logarithmic field, derived density or positivity gate was introduced. Three kinds of evidence were deliberately separated: evolving nonlinear trajectories, same-state linear inverse comparisons, and unchanged saved-state visualization. Cost comparisons include the stated setup and verification work; a passed artifact audit is not equivalent to convergence.

15.1 Initial nonlinear screen

Five cases used 12 updates, restart 160 and a 300-iteration linear cap. Only $n = 3$ included the existing central phase condition. The two $n = 3$ cases differed in forcing tolerance. Every case exhausted its update budget.

Case	Final $\|F\|$	Newton s	Krylov total	$\|v_E^h\|/\|v_F^h\|$
$n = 2, J = 0, \eta = .03$	$8.85263\,10^{-6}$	742.74	1009	$3.25447\,10^{-7}$ / $2.24560\,10^{-6}$
$n = 3, J = 0, \eta = .03$	$8.24392\,10^{-4}$	499.89	429	$9.46671\,10^{-5}$ / $6.20929\,10^{-5}$
$n = 3, J = 0, \eta = .003$	$5.50731\,10^{-4}$	734.09	821	$2.23846\,10^{-4}$ / $1.98912\,10^{-4}$
$n = 1, J = .1, \eta = .03$	$5.83293\,10^{-3}$	310.43	226	$3.23282\,10^{-4}$ / $3.36216\,10^{-4}$
$n = 1.5, J = .1, \eta = .03$	$7.22567\,10^{-5}$	550.70	717	$5.10711\,10^{-5}$ / $5.20697\,10^{-5}$

The actual nonlinear thresholds were $3.81795\,10^{-12}$ for $n = 2$, $3.09931\,10^{-11}$ for $n = 3$, $8.57344\,10^{-10}$ for rotating $n = 1$ and $1.39488\,10^{-9}$ for rotating $n = 1.5$. q14/q18 enthalpy virial differences in all nonrotating cases were below $8\,10^{-14}$ absolute; this is a quadrature check, not mesh convergence or proof of equilibrium.

After $n = 2$'s second update, ten more updates reduced its residual by only 2.112%. Its final step was 1/256 and all rejections were insufficient decrease. Poisson and gravity-gradient blocks contributed 94.770% and 4.076% of the final squared normalized norm. Minimum q18 density was positive, $2.93907\,10^{-4}$, and density-pressure energy/force virials also passed: $2.49026\,10^{-6}$ / $5.70105\,10^{-7}$. This supplies a stalled control without an observed negative-density obstruction, although it is not a global positivity certificate.

Tenfold tighter $n = 3$ linear tolerance reduced its equal-update endpoint residual by 33.20%, but required 46.85% more Newton time and 91.38% more Krylov iterations, and worsened both virials. The final step decreased from 1/4 to 1/32. The ordinary run's squared residual was 79.927% Poisson, 13.854% gravity gradient and 6.144% closure; the tight run was 71.329% closure and 25.876% Poisson. These are different nonlinear trajectories, not same-state inverse comparisons. Both progressed beyond the earlier eight-step endpoints, showing why those endpoints must not be called demonstrated floors.

The rotating $n = 1$ case ended at $\alpha = 1/65536$, with 26 pressure rejections and no metric rejection. Rotating $n = 1.5$ had seven of each and ended at 1/64; four updates after its eighth reduced the residual only 4.97%. Both had negative sampled density, making density-pressure virials undefined.

Case	Scaled EOS RMS	Scaled Bernoulli RMS	Hydrostatic-gradient RMS
$n = 2, J = 0$	$1.87717\,10^{-4}$	$2.80766\,10^{-4}$	.0264341
$n = 3, \eta = .03$	$3.52140\,10^{-4}$	.00230017	.398054
$n = 3, \eta = .003$	$3.69501\,10^{-4}$	.00265289	.479759
$n = 1, J = .1$	.00126849	.00221774	.104394
$n = 1.5, J = .1$	$1.69365\,10^{-4}$	$6.32866\,10^{-4}$	.0366157

Scaled EOS and Bernoulli RMS use sampled density and enthalpy scales; hydrostatic-gradient RMS is unscaled in model acceleration units. These local observations remain distinct from weak residuals and virials. A source-derived row map placed 84.437% of saved Poisson norm squared inside the star for the $n = 2$ update-seven capture, 99.9768%/99.9995% for ordinary/tight $n = 3$, and over 99% for both rotating captures. The exterior rows alone therefore do not account for the stalled norm. Row shares are not pointwise PDE errors or proof that the exterior has no indirect influence on interior fields.

15.2 Matched gravity-curvature trial-path control

Raw and gravity-curvature Newton restarted from the identical $n = 2$ update-three capture, with initial $\|F\| = 9.0134584624\,10^{-6}$, six updates, $\eta = .03$ and common absolute target 10^{-12}. Every

reported full linear gate passed.

Trial path	Final $\|F\|$	Newton s	Final α
Raw Newton	$8.8762959137\,10^{-6}$	449.196	1/256
Gravity curvature	$8.6749533568\,10^{-6}$	887.543	1/2048

The corrected endpoint was only 2.2683% lower at 1.97585 times cost. Its first step accepted 1/16, against raw 1/128, but supplied 90.1683% of its eventual six-step gain; no durable recovery followed. Preparation time rose 3.72673-fold, with 22 auxiliary gravity solves and 866 auxiliary iterations. Their verified prepared defects passed, so the full stall was not a failed auxiliary solve. All rejections were residual-metric failures, not pressure failures. Both endpoints met the enthalpy virial target, but neither met the nonlinear target. This pair did not justify a default trial-path change.

15.3 Frozen inverse comparisons: rate, attained accuracy and total cost

Each case below compares a freshly constructed baseline and candidate on the identical saved physical state, Jacobian and right-hand side. The fresh full endpoint defect is $\|Ap + F\|/\|F\|$. Ordinary tolerance is .003; the marked tight case requests 10^{-7}. Every case had a 300-iteration cap.

State	Candidate	Iterations B/C	Total seconds B/C	Gate B/C
$n = 2$ late	gravity	105/83	51.1251/50.2194	pass/pass
$n = 2$ late	surface/scalars	105/300	51.6185/244.6878	pass/fail
$n = 2$ late	combined	105/109	51.4773/117.3489	pass/pass
$n = 2$ tight	combined	300/300	145.5541/283.3130	fail/fail
$n = 3$ late	combined	99/41	53.1688/65.6294	pass/pass
$n = 1, J = .1$ late	combined	63/28	29.9344/44.3827	pass/pass
$n = 1.5, J = .1$ late	combined	246/95	117.9250/101.7603	pass/pass
$n = 2$ late	LDU+gravity	105/78	51.4910/60.0450	pass/pass
$n = 3$ late	LDU+gravity	99/41	53.1904/35.2005	pass/pass
$n = 3$ late	gravity	99/46	53.3147/31.1652	pass/pass
$n = 1.5, J = .1$ late	gravity	246/137	118.6171/82.3113	pass/pass

"Gravity" uses the gravity-Schur refinement with independent subsystem structure. "Surface/scalars" wraps the independent base in the new approximate coarse Schur inverse. "Combined" wraps the gravity-refined base. LDU adds existing two-way material/gravity feedback. The total cost includes inverse construction, normalization binding, Krylov setup/solve and one fresh full Jacobian check. It excludes initial model/context setup, separate diagnostic action checks, trial-grid observations and output. Internal Schur Jacobian work is already included in application time; it must not be counted twice.

The strongest frozen improvements were gravity-only: 41.54% lower total cost for $n = 3$ and 30.61% for rotating $n = 1.5$. Their endpoint defects changed from .00282583089 to .00259308753 and from .00299226031 to .00199325392, respectively. The $n = 2$ reduction was only 1.77%, too small for a robust timing claim from one pair. There were no repeated timing samples or confidence intervals, and these are cold one-right-hand-side comparisons, not measured benefits from amortized setup reuse.

The surface-only $n = 2$ candidate ended at .00567029124 and failed its gate. The tight $n = 2$ pair attained $1.87757546\,10^{-4}$ and $1.21976551\,10^{-4}$, both failing 10^{-7}. Those are attained accuracies within a finite budget, not arithmetic or discretization floors. All $n = 2$ directions with raw-trial observations retained the same largest sampled acceptable step 1/256; approximately 90.18% of their small-step squared Taylor remainder was Poisson. Quadratic scaling over the sampled range supports a finite-curvature mechanism, without proving its global behavior.

The new coarse wrapper handles all 25 surface parameters and every generated scalar, with 27 or 28 coarse coordinates. It uses an explicitly unbordered field inverse, since restricting an

already scalar-bordered inverse could spuriously constrain the field response. Typed maps and stationary-application contracts are checked at compile time; generation, dimensions and rank are checked after preparation. The wrapper is frozen-state only and must be reconstructed when the Jacobian changes. Coarse construction uses $2k$ full Jacobian actions and k base-inverse actions; each application adds a full Jacobian action and base-inverse application. For the combined candidates, total preconditioner construction ranged from 22.311 to 27.717 seconds; the surface/scalar wrapper alone accounted for 20.515 to 25.578 seconds. These are distinct timing scopes. Reduced outer iterations therefore did not necessarily reduce total time, and the wrapper was not promoted to a live default.

15.4 Weak coupled radial response

For an exact derivative and a passing linear gate,

$$\|Ap + F\| \le \eta\|F\| \quad \Longrightarrow \quad F^T Ap \le -(1-\eta)\|F\|^2.$$

Hence $\eta = .03$ and $.003$ already guarantee infinitesimal descent in the stated norm. Tiny accepted steps motivate finite-step and inverse-response diagnostics, not the assumption that the direction fails to be a linear descent direction.

At the common centrally constrained, nonrotating $n = 3$ capture, baseline, combined and LDU directions all passed .003 and had merit slopes within $8\,10^{-6}$ relative of $-\|F\|^2$. Nevertheless:

Inverse	Normalized surface norm	Normalized signed monopole	Largest sampled α	Full-step $\|F_t\|/\|F\|$
Baseline	.0681146	−.057083	1/4	7.42157
Combined	.113440	−.104359	1/16	14.7443
LDU+gravity	.0287899	−.00681011	1/4	2.77979

The surface norm and signed monopole in this table use normalized solver coordinates, not physical displacement units. The baseline and LDU quarter-step residual ratios were .702686 and .674223. The complete combined-minus-baseline and LDU-minus-baseline directions had normalized norms 58.44% and 61.89% of baseline, but their A actions were only .3197% and .3364% of $\|F\|$. Response per unit difference norm was approximately 183–184 times smaller than for the baseline direction. This measures two weak directions, not a condition number, singular vector certificate or exact nullspace.

Approximately 43% of each squared normalized direction-difference norm was surface, 34% potential and 14% density. Within the squared normalized surface difference norm alone, 97.67% and 96.77% were monopole. These percentages are coordinate-metric contributions, not physical displacement or energy fractions. The effect is a coupled radial-size/field response, not primarily translation or high-m roughness. Freezing the monopole is unjustified: radius is physical, and this control already had a central phase condition. The combined inverse reduced the final surface/scalar defect from $4.64933\,10^{-4}$ to $1.07519\,10^{-6}$ relative to $\|F\|$, yet worsened finite trials. LDU permitted $9.45535\,10^{-4}$ there while improving the quarter-step trial. Coarse-row recovery alone is therefore not a reliable quality measure for a finite Krylov direction. The reported coarse reciprocal condition $1.31320\,10^{-9}$ belongs to the small approximate matrix, not full A.

If $e_i = Ap_i + F$, the common gate bounds $\|A(p_c - p_b)\| \le \|e_c\| + \|e_b\|$, but does not bound $\|p_c - p_b\|$ without inverse information. This supplies a mathematical reason why preconditioners can produce comparable linear defects and different nonlinear behavior. The weak-mode magnitude and composition must still be remeasured in the unbordered rotating family; they were not demonstrated there by this frozen control.

16 A physically compatible rotating $n = 3$ family

The driver originally calibrated K from a nonrotating Lane–Emden star with $M = R = 1$ and then retained $M = 1$ as J changed. At $n = 3$, nonrotating mass is independent of central density. This physical compatibility question is separate from the measured $n = 2$ numerical stall.

For nonnegative finite-energy density obeying the exact EOS $P = K\rho^{4/3}$, the sharp gravitational-energy inequality implies

$$|W| \le 3\Pi \left(\frac{M}{M_{\mathrm{LE}}(K)}\right)^{2/3}, \qquad \Pi = \int P\, dV.$$

Together with isolated steady virial balance $2T + W + 3\Pi = 0$, this gives the necessary condition

$$\beta = \frac{T}{|W|} \le \frac{1}{2}\left[1 - \left(\frac{M_{\mathrm{LE}}(K)}{M}\right)^{2/3}\right].$$

In particular, exact $M = M_{\mathrm{LE}}(K)$ cannot support nonzero T. Nonzero fixed J at finite inertia requires $T \ge J^2/(2I) > 0$. An extra central-density condition cannot repair incompatible mass/EOS inputs; artificial forcing could instead conceal the mismatch. The energy inequality and its Lane–Emden mass relation are documented in equations (1.14)–(1.15) and Theorem 3.1 of Cheng, Cheng and Lin, *Expanding solutions near unstable Lane–Emden stars*; the rotation bound above is our consequence of that inequality and the virial identity.

The numerical calibration is only approximately critical. Independent DOP853 Lane–Emden integrations using five tolerance/start/step controls found $M_{\mathrm{LE}} = .9999999918534539$ for recorded $K = .36393905995216302$, permitting only $\beta \lesssim 2.71552\,10^{-9}$. The control spread in mass was $1.95\,10^{-14}$, a numerical sensitivity observation rather than a rigorous interval enclosure. This near-critical input cannot represent the intended moderate rotational support under the continuum assumptions.

The discrete equations need not satisfy every assumption: independent $\rho/P(h)$, weak EOS closure, negative density, approximate gravity and central forcing all require separate scrutiny. The argument chooses a sensible target; it is not a density limiter or a proof that a discrete residual cannot vanish.

A separate experiment reduced K by 1%, kept target $M = 1$, prescribed $J = .035$ and omitted the central condition. Its necessary bound became $\beta \le .00500000269$. The seed's central density was multiplied by $.99^{3/2}$ to retain its reference radius, making its seed mass approximately $.99^{3/2}$ while target mass remained one. This is an explicitly off-constraint initial guess, not density rescaling during Newton. The seed Bernoulli scalar still used target mass; that mismatch remained in the initial residual.

Nonzero prescribed J already breaks the nonrotating homology, so an additional central condition was not added. At fixed M, K an $n = 3$ homology can change size and J without specifying a new dimensionless shape/rotation fraction. Neither input J nor an uncalibrated angular velocity should be called a fraction of critical rotation.

16.1 Matched live field-feedback result

Two cases shared $n = 3$, $J = .035$, K factor .99, no central condition, all-m L4/Q2, the gravity-curvature trial path, $\eta = .03$, ten updates and absolute target 10^{-12}. Their initial completed norm was .101333855231. Only the inverse differed. Both exhausted the budget.

Quantity	Gravity-Schur	Stellar LDU + gravity-Schur
Final normalized residual	$1.64730591\,10^{-4}$	$1.29842469\,10^{-5}$
Newton seconds	918.881	1052.972

Quantity	Gravity-Schur	Stellar LDU + gravity-Schur
Linear seconds	348.441	335.399
Krylov iterations	534	410
Trial preparation seconds	477.771	621.957
Final accepted α	1/16	1/64
Energy virial, $P(h)$	$7.18429\,10^{-5}$	$1.81141\,10^{-5}$
Force virial, $P(h)$	$5.23757\,10^{-5}$	$2.90028\,10^{-6}$
Energy virial, $P(\rho)$	$6.96335\,10^{-5}$	$3.51558\,10^{-5}$
Force virial, $P(\rho)$	$5.01663\,10^{-5}$	$1.41415\,10^{-5}$
Absolute mass error	$5.45536\,10^{-5}$	$3.37148\,10^{-6}$
EOS RMS / density scale	$3.99773\,10^{-4}$	$4.00529\,10^{-4}$
Bernoulli RMS / enthalpy scale	.00408345	.00372207
Hydrostatic-gradient RMS	.656446	.588702
$T/\|W\|$	.00497575	.00497421
Angular velocity, model units	.401764	.401388

LDU reached a 12.687-times lower residual after the same update count, at 14.593% more Newton time. Linear time fell 3.743%, while preparation rose 30.179%. These are evolving trajectories, not a speedup to the same converged solution. All full linear gates passed, and every rejected trial failed sufficient decrease; neither trajectory had pressure or geometry rejection. q18 density samples were positive without repair, not a global positivity certificate. Thus forcing positivity is not supported as the immediate remedy for this measured live obstruction.

The LDU trajectory reached $1.30261014\,10^{-5}$ after update seven with a full step; three subsequent 1/64 steps improved it only .3213%. Its final squared residual was 36.500% Poisson, 34.432% hydrostatic, 20.234% surface and 6.742% mass. Gravity-only ended 80.254% hydrostatic and 10.967% mass. Field feedback is a substantive improvement on this trial path but does not remove small-step stagnation. Its merit over raw-Newton policies has not been established by this pair, since both shared the gravity-curvature strategy.

17 Virial consistency: pressure and gravity must be distinguished

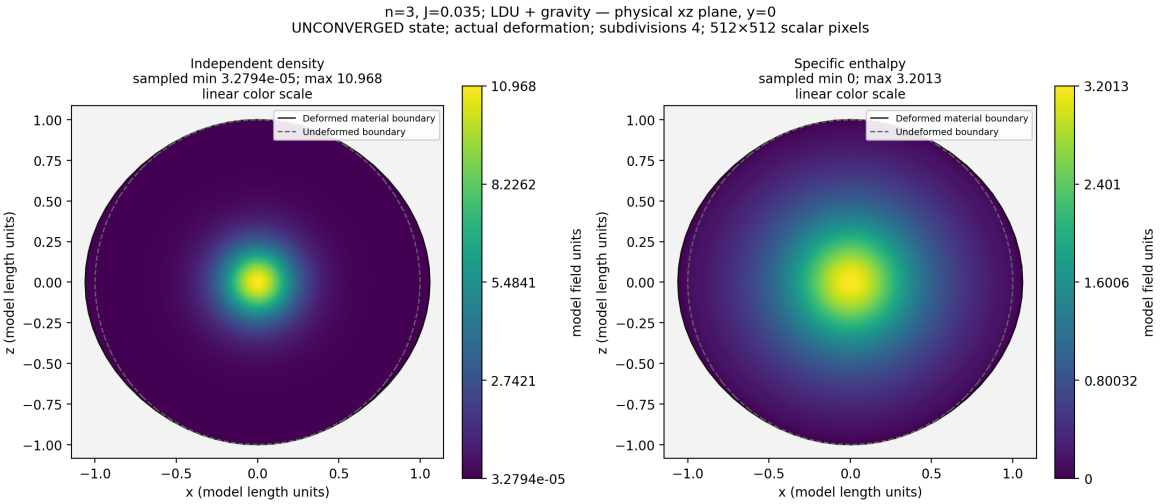


Figure 1: Physical meridional density and enthalpy slices of the final rotating $n = 3$, $J = .035$, K-factor .99 LDU plus gravity-Schur endpoint. The image shows the actual surface deformation without magnification. The model remains unconverged and fails the overall virial milestone; only four display subdivisions were used for this image.

The final rotating LDU model does not meet the overall 10^{-5} milestone. Only force virial formed with $P(h)$ passes. Its signed enthalpy virials are approximately

$$v_E^h = +1.811408\,10^{-5}, \qquad v_F^h = -2.900278\,10^{-6},$$

where both use $|W_E|$ as denominator. Their difference, approximately $2.10144\,10^{-5}$, agrees with the separately measured gravity-energy consistency discrepancy. Replacing $P(h)$ by $P(\rho)$ shifts both signed virials by approximately $+1.70418\,10^{-5}$, yielding $+3.515584\,10^{-5}$ and $+1.414148\,10^{-5}$.

At one fixed state, pressure changes both signed virials equally:

$$v_E^\rho - v_E^h = v_F^\rho - v_F^h = \frac{3(\Pi_\rho - \Pi_h)}{|W_E|}, \qquad v_E - v_F = \frac{W_E - W_F}{|W_E|}.$$

Consequently, even an arbitrary common pressure-only shift leaves at least one absolute virial at or above half their separation, approximately $1.05072\,10^{-5}$. Pressure adjustment alone cannot make both pass in this frozen state. This is a fixed-state arithmetic observation, not a bound on a newly solved or refined model.

q14/q18 energy virials agreed to below $8\,10^{-14}$ in both rotating runs. Quadrature stability must not be mistaken for mesh accuracy. The gravity gap remained about $2\,10^{-5}$, while EOS RMS stayed near $4\,10^{-4}$ despite the large algebraic gain. Since neither state solves the discrete equations and neither mesh nor order was varied, these observations do not establish a spatial error floor or identify an exterior-only cause.

17.1 Pressure-moment identity and the density test space

For nonnegative h, ρ, set $\rho_* = [h/((n+1)K)]^n$ and $\delta\rho = \rho - \rho_*$. Convexity of $P(\rho) = K\rho^{1+1/n}$ gives

$$P(\rho) - P(h) = \frac{h}{n}\delta\rho + D_P(\rho, \rho_*), \qquad D_P \geq 0,$$

because $P'(\rho_*) = h/n$. For a represented density test $v = \sum_i c_i\varphi_i$ and the same mapped quadrature Q used in closure,

$$Q\left[\frac{h}{n}\delta\rho\right] = \frac{1}{n}Q[(h-v)\delta\rho] + \frac{1}{n}\sum_i c_i R_i.$$

Current DGQ2 tests do not generally contain Q3 enthalpy. The first term can therefore remain even at zero weak closure. Compatible DGQ3 tests would permit $v = h$ elementwise, canceling the linear moment at exact closure under that same quadrature. Density would remain an independent field. At finite residual the remaining term is $c_h^T R/n$, which is not directly measured by a normalized Newton norm without undoing residual scaling.

The convex remainder persists, and is not uniformly quadratic near vacuum: for $n > 1$, $P''(\rho)$ diverges as density approaches zero; at $\rho_* = 0$, $D_P = K\rho^{1+1/n}$, with exponent $4/3$ for $n = 3$. Separate diagnostic quadrature can also retain a difference after assembled moment cancellation. The identity is conditional on nonnegative fields where evaluated and is specific to a constant-K polytrope; a general EOS uses its own thermodynamic derivatives. The contribution of each term has not yet been measured in the final models. This is a consistency hypothesis and a diagnostic to implement, not proof that raising order, clipping density or changing pressure will solve convergence.

18 Physical two-dimensional model slices

The added optional diagnostic exports physical xz and xy cuts, with yz and offsets available, showing independent density and enthalpy and the actual material boundary. It changes no equations or field representation. The important geometric contract is a cut by a physical plane, not a reference plane deformed afterward.

Each stellar hexahedron is subdivided in reference coordinates into a lattice and tetrahedra. Crossed edges solve the actual mapped plane equation using a safeguarded bracketed root calculation. Fields are evaluated at those reference coordinates on the owning element. Density traces are not averaged between DG elements; supplied connectivity prevents triangulation across vacuum. A plane coincident with an internal face uses the positive-normal-side trace. Boundary curves are separately evaluated from stellar/nonstellar faces, and the undeformed outline is independently cut by the same physical plane.

The plotter uses equal physical axis scales, no deformation magnification, and signed field values without clipping, logarithms or EOS replacement. Scalar values are rasterized on a 512-by-512 grid before color mapping, avoiding color-interpolation artifacts under triangle refinement. Exterior pixels stay masked. Images distinguish solid deformed and optional dashed reference outlines and mark known unconverged endpoints.

Straight triangles and boundary chords approximate the curved cut. Neighboring parent elements may choose different face diagonals; small gaps/overlaps can therefore occur on curved interfaces. A covered raster pixel uses the first supplied triangle, without cross-element averaging. Tangencies or multiple crossings between subdivision vertices can be missed. These limits are explicitly documented and prevent using the image as a pointwise positivity certificate or interpreting every seam as a physical feature.

18.1 Display-refinement measurements

Four unchanged-state replays compared four/eight subdivisions for final $n = 2$ and rotating $n = 1.5, J = .1$ states, in both planes. Source buffers remained bitwise unchanged, and both resolutions used common axes and color ranges.

Model/plane	Area change	$\int \rho \, dA$ change	$\int h \, dA$ change	Perimeter change
$n = 2$, xz	.12333%	.07667%	.07036%	.03065%
$n = 2$, xy	.12333%	.07667%	.07036%	.03065%
$n = 1.5, J = .1$, xz	.12495%	.00319%	.08797%	.03165%
$n = 1.5, J = .1$, xy	.12395%	.01249%	.07838%	.03088%

These are relative absolute changes of displayed straight-triangle measures, including overlaps; they are not curved-FEM integrals, three-dimensional mass, virial errors or a mesh convergence study. Sampled chord errors decreased about fourfold and axial extents agreed. The rotating meridional model was smoothly oblate, with equatorial extent approximately ± 1.12084 and polar extents $-.988807, +.988480$; the equatorial cut was approximately circular. Its sampled negative density minimum $-8.59\,10^{-4}$ was retained. The $n = 2$ outline was nearly spherical at radius 1.0049. Only the latter met its recorded virial milestone, and both remained unconverged algebraically.

Final $n = 3$ meridional/equatorial plots show strong central concentration and modest oblateness. Those exports used four subdivisions only. The separate $n = 2/n = 1.5$ display study must not be described as refinement verification of the $n = 3$ images. Visual plausibility supplies an additional diagnostic, not physical validation.

19 What to carry forward and what remains unresolved

The recent work establishes a usable native harmonic prescription, diagnoses a specific independent-density projection defect, measures order effects without changing density's physical status, adds reproducible physical slices, and shows real frozen-linear and live-nonlinear gains. It does not establish a converged moderately rotating $n = 3$ model, a calibrated 50%-critical sequence, or the ultimate 99%-critical ambition. No ESTER comparison was performed. Global virial balance at

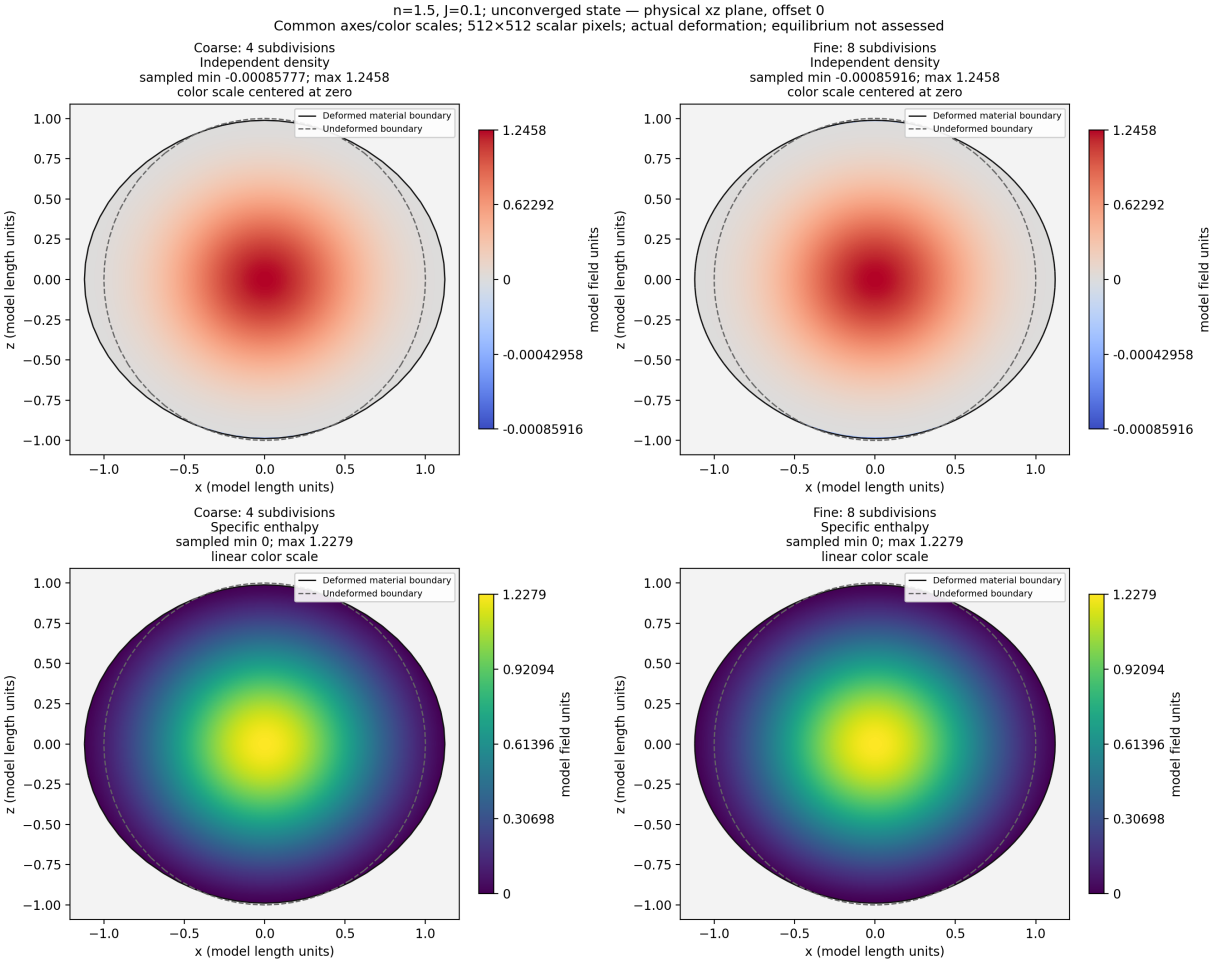


Figure 2: Same-state physical slice comparison for the unconverged rotating $n = 1.5$ endpoint at four and eight display subdivisions. Axes and color ranges are shared. This checks the displayed geometry and fields; it is not a change of the FEM mesh, a three-dimensional integral check or evidence of equilibrium.

10^{-5} remains a primary acceptance milestone, accompanied by internal EOS/hydrostatic checks; an attractive image or a small subset of residuals cannot replace it.

The evidence supports the following staged continuation.

1. Retain gravity-Schur as the strongest measured frozen linear-cost comparator and stellar LDU plus gravity as the better live candidate in the tested rotating $n = 3$ configuration. The latter is conditional on the shared gravity-curvature trial path. Compare complete time to a common convergence and physical target before selecting a production default. Do not promote the new coarse surface/scalar wrapper merely because it lowers iterations.

2. Remeasure coupled direction differences in the unbordered compatible rotating family, including true defects, field/modal composition and finite trial remainders. The centrally constrained nonrotating weak-monopole evidence is a motivated hypothesis, not a transferred diagnosis. A coupled trust-region or regularized step with a physical displacement scale is a reasonable swappable experiment. Radius must remain free, and surface/interior remain simultaneous.

3. Strengthen the material response in controlled stages if its frozen block defects warrant it. Existing material preconditioning already contains the exact prepared enthalpy-to-density derivative; it is incorrect to claim that cross derivative is missing. Exact local mapped DG density mass inversion is a narrow candidate (27 Q2 coefficients per hexahedron). A stronger stationary enthalpy mass inverse must respect H1 connectivity and replaced boundary rows. Only then consider a more costly EOS-aware material/gravity Schur approximation. Such algebraic eliminations belong inside the inverse and need not make density a derived nonlinear field.

4. Measure pressure-moment and gravity-energy consistency separately. Export the linear pressure moment, convex remainder, quadrature differences and both signed virials on identical states. Resolve the gravity relation/source and energy/force discrepancy with interior accuracy as the goal; the exterior need only be accurate insofar as it supports correct interior fields. The current observations do not justify allocating effort solely to exterior resolution.

5. Define a physically compatible $n = 3$ sequence in mass/EOS/rotation and calibrate a dimensionless rotation measure and critical reference. Input J alone is not that calibration. Keep central homology handling optional and separate from incompatible physical specifications.

6. After useful convergence is obtained, perform mesh and harmonic-degree studies, measure excluded surface force in a meaningful norm, and compare local profiles with analytic/nonrotating controls and eventually ESTER. Density-order results remain available as consistency controls; positivity intervention and logarithmic refactoring are deferred per the user's direction.

A closure-consistent seed and an affine-target closure correction are additional untested hypotheses, not completed remedies. At fixed admissible enthalpy and geometry, exact local density mass solves can populate independent density or correct its weak moments. Any trial correction must preserve the original Newton tangent and be freshly verified. If combined with gravity correction, density must be corrected first and gravity rebuilt against its original affine target afterward. That operation can change mass, inertia and mechanical balance, cannot repair unchanged inadmissible enthalpy, and must pass the full residual acceptance test. For fractional $1 < n < 2$, a uniform quadratic correction claim near vacuum would be unsound because the EOS derivative is only Hölder continuous there. No such closure trial policy was implemented in this phase.

A projected surface-equipotential equation and lower harmonic degree were also reviewed as future operator studies but not implemented or tested. They change the discrete problem. Any surface-potential derivative must respect the actual ALE pullback and field transport; substituting an effective-gravity derivative without that derivation would not be justified.

19.1 Recent implementation verification and provenance

Release builds and selected targets passed. The new coarse inverse passed 70 assertions in nine tests. The transactional diagnostic initializer passed 71 assertions in three real-FEM tests; it validates snapshot coordinates, prepares and refreshes before committing, clears stale warm starts and rolls back on failure. Snapshot compatibility checks include names, sizes, mesh and metadata, and unconverged captures require an explicit experimental override. Read-only observation and state installation remain separate operations.

Pure slice geometry controls passed, as did 14 plotting/comparison tests, 11 artifact-auditor controls and 11 tests for each convergence plotter including actual rendering. No unrelated full-suite claim is made. All 24 final-campaign cases passed retained-artifact audits: five nonlinear screens, two matched trial-path restarts, eleven frozen comparisons, four slice replays and two live $n = 3$ cases. All nine nonlinear runs in that final campaign were unconverged within their declared budgets. A passing audit can accurately certify records of a failed linear gate or unconverged model; it does not certify scientific success.

The Sept. 19 integration branch ended at `d7acb0d399aac5ac4c36d834a954c1c959042b81`. Producer revisions differ across adaptive experiments because hypotheses were declared and committed before execution. Campaigns retain private copied executables, source/mesh/input hashes, build settings, exact arguments, logs, states and output hashes. The study index binds six audit files and final report/figure hashes. Later report and plotting edits do not replace numerical producer provenance. Expensive builds, FEM runs, vector audits and rendering were serialized. Timing comparisons are workstation observations with the cost scopes stated above, not portable hardware benchmarks.

20 Consolidated source, architecture, and handoff

20.1 What the consolidation changes

The consolidation uses the cumulative research revision `d7acb0d399aac5ac4c36d834a954c1c959042b81` as its numerical source baseline. This revision descends from the former local main, `75cc638`, through the baseline snapshot `80cc3d8` and the successive controlled investigations. Most topic branches were incorporated by cherry-pick during those investigations. Their tips therefore need not appear as ancestors of the cumulative branch for their changes to be present. The consolidation checks both ancestry and final source content; it does not indiscriminately merge every experimental tip.

The integration branch is `integration/research-consolidation-20260920`. At the source checkpoint `b1fe359`, the additions after the numerical baseline are the preserved user formatting change, `6950dc6`, and six commits importing the historical log-density assessment. The latter begin at `549b1a7` and end at `b1fe359`; they add only `experiments/LOG_DENSITY_ASSESSMENT_2026-09-19.md` and `experiments/log_density_reference_model.py`. They do not implement log-density coordinates in the FEM solver. The consolidation itself adds no new numerical behavior. Its other changes organize evidence, document results, and retain useful source and tests. Local main is advanced to the validated consolidation descendant; there is no remote push and no deletion of historical research branches. The final integration commit and validation record identify the exact published local checkpoint. The final delivery is identified by the local Git tag `research-handoff-2026-09-20`; the tag records the complete report and validation as well as the integrated source.

The original programming checkout was initially on `research/numerical-strategies-20260917` at `10adf7e`, rather than on main. Its only tracked modification was declaration formatting in `libmeanfield/interface/quadrature/mfem.cppm`; that modification is preserved in `6950dc6`. This distinction matters because replacing the original checkout with a research tree without first

preserving the edit would lose user work. Research reports and raw evidence are also preserved, including unsuccessful experiments. Moving an artifact is not a revision of its original producer, parameters, or hash.

20.2 Architectural boundary between physics and numerics

The intended public construction remains physics forward: an EOS, a surface condition, and explicit integral specifications describe the stellar model. A separately owned discretization describes the finite elements, normalization, surface coordinates, and displacement extensions. An owning preconditioner prescription and an owning linear-backend configuration are prepared only after the physical problem exists. The solve context then owns the accepted state, prepared operators, prepared inverse, and nonlinear transaction machinery. This separation is valuable in the rewrite: changing a numerical approximation must not silently add a central-density constraint, change an EOS relation, or select a different physical family.

The current implementation uses C++23 modules, concepts, generated block manifests, and typed prescriptions. Compile-time checks cover structural facts: supported EOS relations, derivative availability, field and block roles, admissible prescription combinations, and backend application contracts. Mesh-dependent facts remain runtime checks: dimensions, rank of a sampled harmonic basis, actual topology, preparation generations, finite parameters, support layout, and distributed agreement. A type cannot prove that a sampled geometry is admissible or that a numerical inverse is accurate, so those checks remain explicit and measurable.

There are three distinct levels of composition:

1. A physical model compiles its field equations and specifications into value and residual blocks, including generated scalar coordinates, appended integral equations, and explicitly replaced carrier rows.

2. A discretization supplies coordinates and maps for these blocks. Physical operators are prepared on the accepted geometry and state; their actions and cross derivatives remain the source of truth.

3. A numerical recipe approximates the inverse or modifies a trial path, while the complete residual and physical diagnostics judge the resulting state. An exact coarse row in an approximate inverse is not an exact solution of the full problem.

Important source boundaries are `libmeanfield/interface/models/`, `libmeanfield/interface/equilibrium/stellar_discretization.cppm`, `libmeanfield/interface/operators/root_manifest.cppm`, `libmeanfield/interface/operators/stellar_equilibrium_compiler.cppm`, `libmeanfield/interface/preconditioning/stellar_recipe.cppm`, and `libmeanfield/interface/solver/stellar_context.cppm`. The experiments directory owns case selection, historical studies, optional probe recipes, artifact verification, and visualization. It should not be mistaken for a minimal production API.

20.3 Fields, geometry, and the preferred surface representation

The retained baseline has independent discontinuous Q2 density and independent continuous Q3 enthalpy. Gravity uses the existing RT2/Q2 mixed pair; displacement is Q3, while the reference mesh geometry is Q4. The density order is now independently configurable at compile time through `MEAN_FIELD_DENSITY_POLYNOMIAL_ORDER`; its default is two. A uniform order increment remains a separate existing option. The closure quadrature query now covers both density-mass and EOS-source contributions when density order changes. Neither this configuration mechanism nor the completed order study establishes that a higher density order cures the nonlinear stall.

The preferred research surface is the native real spherical-harmonic prescription with maximum

degree four and every azimuthal order:

$$q_i = \sum_{\ell=0}^{4} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\widehat{\boldsymbol{r}}_i), \qquad N_a = 25. \tag{30}$$

The coefficients have dimensions of length. Negative real orders encode sine modes and positive orders encode cosine modes; neither odd degree nor nonzero m is silently removed. A maximum absolute azimuthal order can be specified explicitly. For example, restricting the shape to $m = 0$ leaves five coefficients at degree four, but does not by itself make the interior three-dimensional fields axisymmetric.

The harmonic prescription supplies values at reference trace nodes. The existing Q3 trace and volume extension then realize the displacement in finite-element space. The resulting surface is therefore not an exact analytic harmonic graph between nodes. Q4 reference geometry and Q3 incremental displacement coexist deliberately; changing their orders is a separate approximation study. The higher displacement-order trial remains a recorded negative result, not the active default.

Surface coefficients and interior fields participate in the same Newton system with their cross derivatives. The reference trace metric provides compatible primal and dual coefficient scaling. It does not approximate the complete physical response of the interior to a surface change. The current native harmonic implementation requires one MPI rank, consistent with the retained radial volume extension. Tests of collective rejection on two ranks do not demonstrate distributed harmonic solving.

The sandbox selects the harmonic prescription explicitly. Existing *two-argument* calls to `makeStellarDiscretization` still select `NodalRadialSurface`; their behavior is preserved for compatibility. This is an intentional distinction between the preferred example and the backward-compatible factory default.

20.4 Defaults, explicit controls, and experimental candidates

The following table records the source checkpoint rather than proposing new settings. A capability entering main is not automatically promoted to a default.

Layer	Retained behavior and interpretation
Model example	The sandbox is $n = 1$, $J = 0$, with fixed total mass and angular momentum. Central density initializes the Lane–Emden seed; it is not an additional $n = 1$ equation. Central-density constraints remain explicit model specifications for studies that require them.
Surface	Sandbox: all-m, degree-four harmonics. Two-argument library factory: nodal radial surface. Both compose with the existing interior radial and fixed-infinity vacuum extensions.
Normalization	The sandbox selects `PhysicalRieszDiagonal` on reference geometry with a fixed branch mass/radius scale. Modal surface coordinates use their pulled-back reference trace metric.
Library linear options	`FGMRESOptions` has restart length 50. This default was not globally changed to 160.
Sandbox solve	Explicit restart 160; at most 30 Newton updates; relative nonlinear tolerance 10^{-8}, absolute tolerance zero; linear relative tolerance 0.03 and cap 200. Geometry validation/backtracking remain.

Layer	Retained behavior and interpretation
Diagnostic controls	Bare driver defaults retain nodal surface, restart 40, independent structure inverse, and baseline trial strategy. Campaign manifests explicitly override these. Reproduction requires the recorded manifest, not assumptions about the bare driver defaults.
Linear reuse	A nonzero incoming correction is checked with the current operator and RHS. It is reset if its true defect is nonfinite or larger than the cold-start RHS norm. Final convergence is checked against the fresh full-system true residual.
Default structure inverse	`IndependentStellarSubsystems` applies material/surface and gravity approximations independently at that level. Generated scalar-border completion is still present.
Default material inverse	`SurfaceThenMaterialTriangular` uses diagonal surface, enthalpy, and density inverses. The prepared EOS $h \to \rho$ derivative and surface-to-material cross actions are already included; the approximation is not fully field diagonal.
Default gravity inverse	`GravityApproximateLDU` uses a matrix-free Chebyshev mass inverse of order five with 20 power iterations, and three fixed AMG cycles for the potential surrogate.
Gravity refinement	`GravityConsistentSchurLDU` is an available typed policy. Its library default is one correction sweep with relaxation one; the diagnostic option defaults to relaxation 0.5. Exact run settings must come from the manifest.
Live research recipes	`gravity-schur`, `stellar-ldu`, and `stellar-ldu-gravity-schur` are explicit experiment recipes. The last is the stronger measured live candidate for the tested compatible rotating $n = 3$ family with a gravity-curvature trial path. It is not a demonstrated universal replacement for the default.
Default Newton path	`SampledEnthalpy`, `MonolithicGravity`, and `UncorrectedTrial`. No polynomial positivity repair, gravity elimination, or second-order correction is enabled implicitly.
Optional Newton paths	Polynomial admissibility/repair, gravity-curvature correction, reduced gravity, and coupled/damped trial correction remain explicit typed research choices with their checked contracts. Their presence does not imply successful high-rotation convergence.
Coarse surface/scalar inverse	Frozen experiment only. Its base inverse is explicitly unbordered and its state/generation binding is checked. Changed Jacobians require reconstruction. Setup and application expense prevent promotion on present evidence.
Continuation	Explicit diagnostic initialization validates coordinates, layout, mesh, and finite data; prepares and refreshes before commit; clears stale linear guesses; and rolls back on failure. Unconverged snapshots need an explicit experimental override.
Slices	Optional physical-plane output evaluates the independent fields inside their owning elements on the deformed stellar geometry. Signed values are retained. Display refinement is not mesh convergence.

The default material response can be written schematically, in physical block notation, as

$$\delta a = D_a^{-1} b_a, \tag{31}$$

$$\delta h = D_h^{-1}(b_h - A_{ha}\delta a), \tag{32}$$

$$\delta \rho = D_\rho^{-1}(b_\rho - A_{\rho a}\delta a - A_{\rho h}\delta h). \tag{33}$$

Here the diagonal inverses are approximate mass/trace inverses, whereas $A_{\rho h}$ is the actual prepared EOS cross action. At the next level the default structure inverse omits direct material–gravity feedback. The optional stellar LDU first obtains a material response, sends its source and geometry effects into gravity, and recovers gravity feedback into material. Gravity-only Schur refinement improves the gravity block; it does not itself add this EOS/material feedback. These distinctions explain why the later ablations compare multiple policies rather than treating all coupling changes as a single intervention.

20.5 Branch disposition and retained provenance

The table below refers to the research tips retained from the investigation. "Included" means ancestry or inspected patch/content equivalence, not necessarily a merge commit with that branch tip as a parent. The six log-assessment commits were imported during consolidation. All named historical branches are retained; no branch deletion is part of this work.

Branch	Tip	Disposition
`main`	`75cc638`	Former main checkpoint; ancestor of the consolidation baseline.
`research/baseline-20260916`	`80cc3d8`	Included baseline snapshot of user optimizations and initial diagnostics.
`research/central-constraint-20260916`	`2e19c2c`	Included equivalent commits: central density as seed data in the n1 example.
`research/combined-20260916`	`95d9fba`	Included cumulative first controlled campaign.
`research/experiment-harness-20260916`	`899cb18`	Included diagnostic and physical acceptance reporting.
`research/krylov-20260916`	`7127383`	Included warm-start true-residual safeguard and controls.
`research/krylov-window-20260916`	`1f9184b`	Included restart160 sandbox setting; report identical.
`research/gravity-coupling-20260916`	`db0ddd3`	Failed one-way completion remains historical; no old driver/default promotion.
`research/diagnostic-campaign-20260917`	`7a07cae`	Included pressure/shape attribution campaign.
`research/pressure-obstruction-20260917`	`c6572d3`	Included pressure, trace, and nodal diagnostic exports.
`research/shape-curvature-20260917`	`b24e6e0`	Included frozen shape/field curvature separation.
`research/numerical-strategies-20260917`	`10adf7e`	Included typed optional strategy integration and validation.
`research/polynomial-positive-newton-20260917`	`9aefb36`	Included explicit experimental certificate/repair; disabled by default.
`research/polynomial-endpoint-certification-20260917`	`cc6f4b0`	Included read-only polynomial endpoint certification.
`research/gravity-curvature-correction-20260917`	`3762c74`	Included checked auxiliary gravity correction; opt-in.
`research/gravity-reduced-newton-20260917`	`d25d141`	Included reduced formulation and full-system checks; opt-in.
`research/gravity-compact-application-20260917`	`ff1b889`	Included compact canonical gravity action and guarded fallback.

Branch	Tip	Disposition
research/gravity-prepared-roundoff-20260917	782827c	Included prepared consistency arithmetic scale.
research/gravity-trial-diagnostics-20260917	7f4b068	Included current-attempt diagnostics and verified records.
research/gravity-direction-diagnostics-20260917	47775a9	Included relift accounting before descent rejection.
research/preconditioner-benchmarks-20260918	dc1da62	Included benchmark methodology, implementation, and results.
research/gravity-consistent-schur-20260918	c10d019	Included fixed-work mass-consistent gravity Schur policy/tests.
research/preconditioner-cached-cross-actions-20260918	07769e7	Included reuse of prepared gravity force in cross preconditioning.
research/coupled-shape-coarse-20260918	4a6c1b5	Included optional frozen coarse experiment and tests; not default.
research/nonlinear-step-obstruction-20260918	2944dd5	Included nonlinear path, direction, and remainder diagnostics.
research/coupled-trial-correction-20260918	46249d0	Included bounded coupled correction and measured limitations.
research/damped-coupled-trial-20260918	db854c0	Included bounded damping/eligibility controls; no default promotion.
research/shape-field-weak-response-20260918	6ceda61	Included weak-response study, preserving Q3 baseline.
research/shape-field-displacement-q4-20260918	785a7ae	Q4 source change excluded; final negative report/comparison tools included.
research/reduced-smooth-surface-20260918	2171b8b	Included frozen reduced-space study and provenance audit.
research/spherical-harmonic-surface-20260919	1148c09	Included native coupled harmonic prescription and metrics.
research/negative-density-diagnosis-20260919	7ed4d23	Included localization, independent projection, and interval evidence.
research/density-order-study-20260919	ef76e0c	Included independent order configuration/closure quadrature and study.
research/log-density-assessment-20260919	90580b2	Imported as deferred assessment/reference script through b1fe359; no FEM change.
research/linear-coupling-study-20260919	0e6688f	Included frozen surface/scalar Schur and unbordered adapter tests.
research/model-slices-20260919	5da3f42	Included physical slice exporter, plots, bounds fixes, and tests.
research/n3-compatible-family-20260919	1d746b0	Included explicit mass-compatible physical-family control.
research/n3-linear-ablation-20260919	e78af98	Included predeclared matched gravity-only n3 comparison.
research/gravity-coupling-ablation-20260919	38c7184	Included gravity-only attribution for rotating higher-index state.
research/n3-live-field-feedback-20260919	1ed03f5	Included matched live field-feedback control.
research/coupled-convergence-study-20260919	d7acb0d	Cumulative numerical baseline for consolidation.

Branch	Tip	Disposition
integration/research-consolidation-20260920	b1fe359	Consolidation source checkpoint: preserved formatting and imported assessment; reports follow.

Several apparently unmerged commits are only isolated-branch registration changes. Commits 4a6c1b5, e599900, 2c64a92, 5952691, and a773c9a add test or module registrations already present in the cumulative CMake and module exports. Reapplying them would add no capability. Similarly, the surface/scalar Schur implementation from 1d95602 is included as 2398118, and the initial slice implementation from 966a4d4 is included as 9bac57a. Final topic and integration source were compared directly where patch IDs differed. These are documented equivalences, not assumed equivalences from similar commit titles.

20.6 Adoption, retention, and exclusion rationale

Adopted infrastructure. The consolidation retains the optimized prepared-operator baseline, reference-table and mapping caches, linear true-residual safeguards, checked operator/strategy interfaces, cache-aware preconditioner cross actions, consistent gravity Schur option, native harmonic surface, independent density-order configuration and corrected closure quadrature, transactional continuation, physical slice output, and artifact auditors. These are reusable capabilities with explicit contracts. In particular, the pre-existing performance work in 80cc3d8 predates many of the numerical experiments; the large diff from former main must not be attributed entirely to the later investigations.

Retained as experiments. Polynomial admissibility/repair, gravity elimination, auxiliary curvature paths, coupled/damped trial corrections, and coarse surface inverses remain available for controlled comparison. Their type-level integration and tests are useful even when a particular campaign did not justify default promotion. The current user preference is to improve broad convergence without forcing density positivity. The baseline path therefore remains unchanged, and the newly imported log-density document is historical analysis, not an active implementation direction.

Not adopted as active numerical behavior. The Q4 displacement change is excluded because the controlled trial did not supply a convergence improvement warranting the changed approximation and cost. Its report and analysis tools are retained with Q3 active. The early one-way material-to-gravity completion is also excluded from the current driver. At the matched $n = 3$, $J = 0$ seed with restart160 and cap200, it changed the true relative defect from 0.0402590 to 0.0381781, while both solves failed their 0.03 gate and solve time rose from 98.915 to 117.712 s. The correction norm grew from 0.757644 to 1.540825 and the sampled thermodynamic step limit decreased from 0.152435 to 0.0913461. That counterexample concerns one particular triangular composition; it does not invalidate later two-way LDU with a stronger gravity approximation.

Not implemented. A coupled trust region or regularized step, exact local DG density-mass inverse, stronger H1 enthalpy-mass inverse, EOS-aware field Schur inverse, and projected surface-equipotential equation are proposals, not completed solver features. Changing the physical $n = 3$ family or calibrating a critical-rotation fraction is likewise distinct from improving an inverse. No consolidation wording should turn these recommendations into results.

20.7 A practical contract map for the rewrite

The rewrite need not reproduce every experimental wrapper. It should reproduce the boundaries that made comparisons meaningful: explicit physical requirements, coordinates whose transformations are known, prepared operators whose validity is checkable, and a solver that reports what actually happened. The following contracts separate implementation responsibility from current algorithm choices. Source paths identify useful reference implementations, not a requirement to preserve every class name.

System	Contract to preserve and initial verification
EOS	Declare relations by input/output quantities and derivative support. An operator asks for the relation it needs rather than assuming every EOS has density-from-enthalpy. Type checks cover relation availability and units; evaluation checks cover finite inputs and the actual physical domain. Keep EOS variable choice separate from whether a field is independent. Verify derivatives on their valid domain and relation consistency for the polytrope. References: `interface/eos/` and `interface/material/thermodynamic_equations.cppm`.
Domain/support	Identify stellar, vacuum, interface, and outer-boundary support independently of coordinate storage. Density exists on stellar support; its zero exterior is not a vacuum density unknown. Generate restrictions/embeddings from support ownership, including shared trace DOFs. Verify coverage, no duplicated ownership, and interface consistency. References: `interface/utils/domain.cppm`, `interface/field/field_mfem.cppm`, and the deformation ownership layer.
Fields and blocks	Let typed field quantities carry approximation family, order, support, and units; let blocks distinguish value coordinates from residual duals. Generate layouts and equation contributions from these roles. Do not identify blocks only by historical numeric offsets or assume every field pair has equal dimensions. Verify rectangular coupling, restriction/injection identities, and generated specification borders. References: `interface/field/`, `interface/utils/blocks.cppm`, and `interface/operators/root_manifest.cppm`.
Quadrature	A form query declares the term, role, trial/test/coefficient orders, support, and mapping context. Runtime policy resolves actual rules and records them. A nonlinear EOS or curved map cannot be made polynomial merely by an order formula. Keep production, preconditioner, projection, and diagnostic roles distinct. Verify closure mass/source coverage and compare diagnostic rules independently. References: `interface/quadrature/` and field form descriptors.
Mapping	Consume mesh geometry, displacement data, and exterior-map policy; return physical coordinates, full map Jacobian/inverse/determinant, measures, normals, and coherent directional variations. Implement scalar pullback and H(div) Piola transformations consistently with the same map. Mapping must not depend on spherical harmonics. Verify identity/affine maps, transformation round trips, face orientation, flux/divergence identities, and finite-difference variations before caching. References: `interface/mapping/` and `impl/mapping/`.

System	Contract to preserve and initial verification
Deformation	Compose a surface prescription with independent interior and vacuum extensions. Supply displacement, tangent, transpose pullback, and derivative of the pullback. A harmonic parameter basis and a nodal basis should produce equal physical residual/Jacobian actions when they represent the same displacement. Check trace continuity, adjoint identities, mode rank, ownership, and stale mesh generations. References: `interface/deformation/` and `interface/equilibrium/stellar_discretization.cppm`.
Physical operator	Separate mathematical residual/Jacobian actions from preparation and caching. Cache by the complete set of dependencies, including geometry, density, enthalpy, gravity, rotation, and relevant specifications. Use the same discrete equations and rules for residual and linearization. Verify isolated blocks, cross derivatives, and the full coupled directional derivative; include physical boundary row replacement. References: `interface/operators/` and `impl/operators/`.
Normalization	Keep explicit maps for state coordinates and residual duals. Derive physical scales from declared reference quantities and Gram factors from a documented geometry/basis. Every normalized Jacobian, inverse, restriction, and diagnostic must use the corresponding maps. Verify round trips, physical/normalized action equivalence, and pullback metric consistency, including harmonic coefficients. References: `interface/normalization/`.
Preconditioner	Separate an owning unbound recipe from its prepared, problem-bound inverse. State linear/stationary application guarantees in the type contract where possible; check dimensions and generations at runtime. Distinguish unbordered field approximation from generated scalar completion. Measure setup, refresh, actions, fresh true defect, and complete nonlinear cost. References: `interface/preconditioning/` and the experiment recipe `experiments/plateau_newton_preconditioner.hpp`.
Solver	Keep accepted state authoritative. Trial preparation, evaluation, optional correction, acceptance, commit, and rollback form a transaction. Invalidate stale caches/guesses explicitly. A backend recurrence residual is not the full-system true residual; a reduced solve report must retain its coordinate-space identity. Report failure stages, update budgets, step lengths, all attempted work, and actual stopping thresholds. References: `interface/solver/`.
Validation and output	Preserve input/source/mesh hashes and actual commands; separate numerical convergence, artifact integrity, and physical acceptance. Evaluate both pressure routes, both virial forms, mass/angular momentum, EOS, hydrostatic/Bernoulli observations, and resolution trends. Slices show the actual deformed star with signed fields and explicitly label unconverged states. References: `experiments/` auditors, physical diagnostics, and slice exporter.

In the source references inside this table, `interface/` and `impl/` are relative to `libmeanfield/`. The table describes contracts to retain; a proposed verification is not a statement that every possible backend, EOS, geometry, or MPI configuration has already passed it.

20.7.1 Mapping before deformation-specific optimization

Given the systems already implemented in the rewrite—EOS, domains, fields/blocks, and quadrature configuration—mapping is a natural next foundation. Write it in terms of the complete element-to-physical map $\boldsymbol{x} = \mathcal{T}(\boldsymbol{\xi}; \boldsymbol{a})$, with $J = \partial\mathcal{T}/\partial\boldsymbol{\xi}$ and determinant j. The ownership of each map factor must be explicit: reference element, computational mesh, displacement, and exterior compactification must each be applied exactly once. The current slice exporter specifically avoids applying displacement a second time when evaluating the physical model. For the orientation convention under which the map is admissible, the standard transformations include

$$\nabla_x u = J^{-T}\nabla_\xi \widehat{u}, \qquad \boldsymbol{v}_x = j^{-1}J\widehat{\boldsymbol{v}}, \qquad \nabla_x \cdot \boldsymbol{v}_x = j^{-1}\nabla_\xi \cdot \widehat{\boldsymbol{v}}. \tag{34}$$

Measures and oriented face transforms must use the same determinant and orientation convention. The rewrite should make this convention explicit rather than mix absolute integration weights with signed Piola factors. Useful derivative identities are

$$\delta J^{-1} = -J^{-1}(\delta J)J^{-1}, \qquad \delta j = j\ \mathrm{tr}(J^{-1}\delta J). \tag{35}$$

These identities are small, independently testable building blocks for the shape derivatives of mass, gravity, pressure, and surface force. They are not a replacement for checking the derivative of the fully composed map.

The surface prescription should sit above this layer and supply a consistent displacement and its variations. It is then possible to change nodal parameters to harmonics, or change an extension, without rebuilding the meaning of every physical operator. A reference metric for harmonic normalization likewise belongs to the parameter-to-trace map, not to an arbitrary interior displacement norm.

Cache only after these contracts are established. The present mapper caches computational-mesh geometry separately from current displacement and compactification state. Mesh/node sequence changes invalidate cached mesh data; raw coordinate edits require the appropriate mesh-update signal or explicit clearing. Mesh lifetimes and thread ownership are part of this contract. Reusing a pointer or a quadrature index is not enough to prove that a prepared value is current.

20.7.2 Normalization and inverse composition

Let D_x map physical state coordinates to normalized coordinates and D_F map physical residual duals to normalized residuals. With these maps fixed during a linearization,

$$\widehat{x} = D_x x, \qquad \widehat{F}(\widehat{x}) = D_F F(D_x^{-1}\widehat{x}), \qquad \widehat{J} = D_F J D_x^{-1}. \tag{36}$$

A physical inverse approximation B correspondingly acts as

$$\widehat{B} = D_x B D_F^{-1}. \tag{37}$$

This explicit distinction prevents using the state scaling on residuals or accidentally applying a scalar-border completion twice inside a coarse Schur construction. If a future normalization changes with state, its lifecycle and any derivative interpretation need separate definition; the current fixed reference normalization should not be silently replaced by such a scheme.

For a nonlinear surface map $\boldsymbol{u} = E(\boldsymbol{a})$ and physical force dual $\boldsymbol{f}$, the parameter residual contains $E'(\boldsymbol{a})^T\boldsymbol{f}$. Its directional derivative contains both the pullback of $\delta\boldsymbol{f}$ and $\delta[E'(\boldsymbol{a})^T]\boldsymbol{f}$. The latter vanishes for a fixed linear parameterization but cannot be omitted from the general prescription contract. This is one reason the existing deformation interface exposes a pullback derivative separately from a forward Jacobian action.

20.7.3 Acceptance and evidence should remain separate

The rewrite should report at least three distinct outcomes: whether a linear equation met its true-residual target, whether the nonlinear iteration met its stated stopping criterion, and whether the model met physical validation requirements. A campaign artifact can be internally consistent while recording failure of any of these. A single passing virial number is not sufficient when the energy and force forms, or the enthalpy- and density-based pressure routes, disagree. The current target remains a well-resolved, physically credible, rotating $n = 3$ model with virial balance within 10^{-5}; the consolidated code does not yet claim that milestone or a calibrated 50-percent-critical rotation rate.

21 Evidence organization and reproducibility after consolidation

The consolidation preserved rather than deleted the accumulated working results. Seventy-one explicitly inventoried moves relocated 6,228 regular files containing 12,468,217,917 bytes. Every regular file was hashed before and after its move; symbolic links and directory structure were also compared. The complete inventory is `reports/archive/move_manifest.json`, with the original plan and organizing script preserved beside it. The verifier `reports/verify_archive.py` can recheck the archived contents on this machine or against a supplied repository location after transfer.

The principal directory layout is:

```
reports/
  MeanField_research_handoff_2026-09-20.tex
  MeanField_research_handoff_2026-09-20.md
  MeanField_research_handoff_2026-09-20.pdf
  branch_consolidation.md
  figures/             selected versioned scientific figures
  sources/             original audits copied without modification
  prior/               earlier optimization report and README
  archive/             inventories, provenance, historical experiments
working/
  research/diagnostics/ complete September research campaigns
  legacy_results/      earlier result directories and loose tables
  legacy_misc/         ZIPs, patches, scripts, old test output
  legacy_documents/    originals of earlier report files
  report_builds/       LaTeX intermediates and rendering checks
  consolidation_2026-09-20/  current validation records
```

The raw working tree is ignored by Git, except its directory guide. This is a storage policy, not evidence deletion: the complete raw files remain local, including old binaries and unsuccessful trials. The smaller source, reports, selected figures, case manifests, audit code and relocation inventory are versioned. A recipient needs an explicit copy of `working/` to rerun saved-state analyses from this archive; a Git clone alone is insufficient.

The original project-workspace `diagnostics` directory is now a symbolic link to `working/research/diagnostics/` in the main repository. Thus the absolute paths already recorded in case manifests and provenance continue to resolve on this machine. Original provenance was not rewritten. When moving the archive to another machine, produce a new execution manifest with rebased paths and retain the original manifest as provenance.

Existing CMake build trees were intentionally left at their original paths: their caches and C++ module artifacts contain absolute source/build identities. They are now ignored by Git and

documented as local caches. Validation used the isolated research worktree's rebuilt binaries; retained binaries in the original checkout can be stale and must be rebuilt before use. Moving them would create an avoidable rebuild or stale-cache ambiguity. Research worktrees and branch references are retained. Local IDE and tool configuration remains in place. No synced file under the ChatGPT project's `sources/` directory was modified or relocated.

The only originally uncommitted tracked source change was quadrature declaration formatting. Its exact patch is archived and its content is in commit `6950dc6`. The original checkout's edit is also retained in a named Git stash when that checkout is switched to main. The consolidated source already contains the edit, so replaying the stash on main is neither necessary nor advised.

21.1 Rebuilding and reproducing results

The compiled implementation requires a C++23 compiler with module support, CMake, Ninja, MPI, MFEM and the other dependencies identified by `CMakeLists.txt`. The checked local build uses Homebrew LLVM rather than Apple's default compiler, Release mode, interprocedural optimization, density order two and zero uniform order increment. Exact compiler/dependency paths remain in campaign provenance and the consolidation build record. A fresh build must be configured for its own source checkout; reusing a different checkout's CMake cache is not a valid shortcut.

The experiment entry point is `experiments/run_plateau_campaign.py`. Its arguments specify source, build, mesh, case manifest and a new output directory. It requires committed tracked source and refuses to overwrite an existing campaign directory. It copies the executable, records exact commands and input hashes, and runs cases serially. Case manifests state surface choice, constraint selection, tolerances, iteration budgets and preconditioner/trial policy. Starting the bare diagnostic executable without those options does not reproduce the preferred harmonic configuration.

The reference recipes include `experiments/spherical_harmonic_surface_cases_2026-09-19.json`, `experiments/convergence_linear_cases_2026-09-19.json`, `experiments/convergence_linear_followup_cases_2026-09-19.json`, `experiments/convergence_screen_cases_2026-09-19.json`, and `experiments/n3_compatible_family_cases_2026-09-19.json`. The last file specifies a compatible physical control, not an interchangeable numerical tuning of the original $n = 3$ model. Follow the saved-state chains and their exact coordinate metadata rather than manually transplanting a vector between discretizations.

Saved campaign artifacts can be checked with the study-specific auditors before any new solve. The final combined study uses `experiments/summarize_coupled_convergence.py`; earlier studies have separate pressure, coupled-trial, density-order and harmonic auditors with their own schemas. A checker should not be used on another schema merely because the files have similar names.

For visual output, use `experiments/MODEL_SLICES.md`. Export physical planes at the same saved state with two subdivision counts before interpreting small surface features. The existing comparison and convergence plot helpers require a fresh output path outside their active source checkout; render in external staging and then archive the outputs under `working/` with their sidecars. This consolidation preserves that output guard. The report builder is separate and deliberately writes its intermediates under `working/report_builds/research_handoff/`.

To regenerate the delivered report, run

```
python3 reports/build_report.py
```

The script invokes Pandoc directly on the authoritative LaTeX source to produce GitHub-flavored Markdown, then compiles the PDF with `latexmk`. It records tool commands, exit status, source/output hashes and build time. `-markdown-only` does not require a TeX installation. The Markdown is generated, not separately edited. A narrowly scoped Pandoc filter removes only a repeated first table-body row identical to the heading; explicit norm and absolute-value delimiters

avoid collision with Markdown table separators. Figure copies are bound to their original archived sources through `reports/archive/figure_sources.json`.

22 Consolidation validation on 20 September

The integration was validated in the existing isolated research worktree, using the numerical source checkpoint `b1fe359`. Its active numerical implementation differs from `d7acb0d` only by the preserved user quadrature declaration formatting; the additional log-density files are offline reference work. Documentation and archival changes are recorded in the final delivery tag. A 474-file source/fixture SHA-256 inventory binds the checks to their actual inputs and is verified against the original checkout after integration. This section describes new consolidation checks, separately from the earlier study validations reported above.

The Release build completed successfully for the main regression executable, sandbox, plateau diagnostic, density-order contracts, continuation contracts, surface/scalar Schur contracts and MPI test executable. Building the MPI test executable is not an execution of the distributed test suite. The checks below ran serially with the recorded local MPI shared-memory configuration; no new distributed harmonic-solver claim is made.

Check	Recorded result	Seconds
`mapping_surface_normalization`	150 cases; 7,224,720 assertions	794.18
`coupled_solver_preconditioner`	118 cases; 1,369,031 assertions	1432.04
`density_order_contract_tests`	52 cases; 722,405 assertions	199.70
`continuation_contract_tests`	8 cases; 515 assertions	610.43
`surface_border_schur_contract_tests`	9 cases; 70 assertions	0.44
`python_regressions`	146 tests; 2 initially skipped	12.32
`coupled_artifact_auditor_selftests`	11 tests	0.09
`slice_geometry_build`	completed	0.88
`slice_geometry`	completed	0.25
`n3_mass_compatibility`	completed	0.93
`n1_harmonic_smoke`	completed	210.08
`plot_render_contracts`	2 tests	11.04

The selected C++ checks account for 337 test-case executions and 9,316,741 assertions; the Python checks account for 157 tests. The 2 plot-render checks initially skipped by their environment gate were subsequently enabled and passed in the separate serial rendering check. The Python total counts executed tests, excluding the initial skips. These are executions, not a count of distinct cases across overlapping tag selections. Large assertion counts include array-element comparisons and do not indicate an equivalent number of independent scientific tests. The full unrelated regression suite was not run.

22.1 Fresh physical smoke model

The smoke case repeated the native all-m, $L = 4$, $n = 1, J = 0$ control with no central-density constraint, restart 160, linear cap 200, true linear target 0.03, and a six-update budget. The nonlinear relative target was 10^{-8} with zero absolute tolerance. It also exported physical model slices at four display subdivisions. This is a baseline regression check, not another high-index investigation or a claim about the unresolved rotating target.

The model converged in 5 accepted updates. Its final normalized residual was $1.5271452 \times 10^{-12}$. The four q18 absolute virials, in order energy/force from enthalpy pressure and energy/force from density pressure, were 3.8703561×10^{-7}, 2.6593819×10^{-7}, 4.2699365×10^{-7}, 3.0589623×10^{-7}.

All four also passed 10^{-5} at q14. The result reproduces the modest baseline capability; it does not resolve the rotating $n = 3$ deficit.

22.2 Artifact and document checks

All 6,228 relocated files were independently rechecked against their original content hashes with no mismatch. The report compiled to PDF and was converted from its authoritative LaTeX source with Pandoc. Conversion checks compared table cells, mathematical expressions, paths and figure captions, and verified the contents links. Repeated PDF table headings are removed from Markdown without removing data rows. PDF pages were rendered and inspected for table width, equation clipping, figure labels and page layout.

The full validation commands, durations, logs, source/fixture inventory and smoke artifacts are preserved under `working/consolidation_2026-09-20/validation/`. Compact versioned copies of the build record, validation results and source inventory are in `reports/validation/`. This provides both a readable handoff and a record that can be checked independently. Successful build, regression and artifact checks do not upgrade any unconverged historical model to a physical success.

A Versioned study-report catalog

The following reports preserve the detailed methods and original interpretation. Their source filenames are repository-relative; raw campaign locations are recorded within the reports and in the archive index.

```
experiments/CENTRAL_CONSTRAINT_HYPOTHESIS_2026-09-16.md
experiments/COUPLED_CONVERGENCE_STUDY_2026-09-19.md
experiments/COUPLED_TRIAL_CORRECTION_2026-09-18.md
experiments/DAMPED_COUPLED_TRIAL_2026-09-18.md
experiments/DENSITY_ORDER_STUDY_2026-09-19.md
experiments/FIELD_COUPLING_REVIEW_2026-09-19.md
experiments/GEOMETRY_QUALITY_EXPERIMENT.md
experiments/GEOMETRY_QUALITY_FINDINGS_2026-09-08.md
experiments/GRAVITY_COMPACT_APPLICATIONS_2026-09-17.md
experiments/GRAVITY_CONSISTENT_SCHUR_2026-09-18.md
experiments/GRAVITY_DIRECTION_DIAGNOSTICS_2026-09-17.md
experiments/GRAVITY_PREPARED_ROUNDOFF_2026-09-17.md
experiments/GRAVITY_REDUCTION_IMPLEMENTATION_2026-09-17.md
experiments/GRAVITY_SUBSYSTEM_IMPLEMENTATION_2026-09-17.md
experiments/KRYLOV_HYPOTHESIS_2026-09-16.md
experiments/KRYLOV_WINDOW_HYPOTHESIS_2026-09-16.md
experiments/LOG_DENSITY_ASSESSMENT_2026-09-19.md
experiments/MODEL_SLICES.md
experiments/N3_COMPATIBLE_FAMILY_2026-09-19.md
experiments/NEGATIVE_DENSITY_INVESTIGATION_2026-09-19.md
experiments/NONLINEAR_STEP_INVESTIGATION_2026-09-18.md
experiments/NUMERICAL_CAMPAIGN_2026-09-16.md
experiments/NUMERICAL_STRATEGIES_2026-09-17.md
experiments/POLYNOMIAL_ENTHALPY_ADMISSIBILITY_2026-09-17.md
experiments/POLYTROPE_H_CONVERGENCE_2026-09-09.md
experiments/POLYTROPE_VALIDATION.md
experiments/POLYTROPE_VALIDATION_FINDINGS_2026-09-09.md
```

experiments/PRECONDITIONER_CACHED_CROSS_ACTION_2026-09-18.md
experiments/PRECONDITIONER_IMPROVEMENT_2026-09-18.md
experiments/PRESSURE_AND_SHAPE_CAMPAIGN_2026-09-17.md
experiments/PRESSURE_OBSTRUCTION_METHODOLOGY_2026-09-17.md
experiments/README.md
experiments/REDUCED_SURFACE_STUDY_2026-09-18.md
experiments/SHAPE_FIELD_WEAK_RESPONSE_2026-09-18.md
experiments/SPHERICAL_HARMONIC_SURFACE_2026-09-19.md
experiments/SYSTEM_ISOLATION_PROFILE_HARNESS.md
experiments/check_numerical_strategies_README.md
experiments/check_plateau_physics_README.md
experiments/plateau_diagnostic_README.md
experiments/plateau_shape_curvature_README.md