A 4+1-dimensional discontinuous Galerkin numerical-relativity engine for White Hole Dynamics:
formulation, discretization, and validation

T. Scaramellino[1]

[1] *Independent Researcher*

(Dated: September 10, 2026 (**DRAFT / TEMPLATE** — methods paper, to be updated after production runs))

We describe the numerical-relativity engine developed for the White Hole Dynamics (WHD) research program: a 4+1-dimensional ($D = 5$) discontinuous Galerkin (DG) solver for the generalized-harmonic (GH) Einstein equations coupled to a field-theoretic matter sector consisting of an SU(2) Yang–Mills field, a complex SU(2)-doublet Higgs field, and two entropy currents. The spatial slices of the bulk are four-balls whose bounding three-spheres model our universe as a brane; the engine is built to evolve the nucleation and subsequent dynamics of a macroscopic white-hole-like source in this bulk. We detail the formulation (the 95 evolved GR fields, the matter action and its first-order reductions), the discretization (a hypercubed-sphere/4-ball domain, Legendre–Gauss–Lobatto collocation with spectral differentiation, upwind-penalty DG fluxes with a targeted Lax–Friedrichs stabilization of the marginally-propagating generalized-harmonic mode, fourth-order Adams–Bashforth time stepping, and a Hesthaven–Warburton modal filter), a curved-background extended conformal-thin-sandwich (XCTS) constraint re-solve used at mid-flight injection events, adaptive mesh refinement with a mortar treatment of non-conforming faces, sphere-based dynamical excision, and a single-GPU CUDA implementation with per-kernel CPU/GPU bit-equality. We report the validation that is in hand: vacuum Tangherlini (5D Schwarzschild) evolution, long-time stability of the multi-shell curved-initial-data configuration to $T = 35\, r_h$, CPU/GPU bit-equality across the kernel set, and the closure of the Layer-1 and Layer-2 matter gates. We are explicit about what remains: the headline production run (the "birth run") has not yet produced a validated THEORY-class result, and every place where a production number belongs is marked accordingly. This is a methods paper; the physics-results paper is separate.

I. INTRODUCTION

The White Hole Dynamics (WHD) program posits that our observable 3+1-dimensional universe is the boundary ("brane") of a 4+1-dimensional bulk spacetime [1]. Spatial slices of the bulk are four-balls; the bounding three-sphere is the brane. A macroscopic, long-lived, white-hole-like "engine" sits in the bulk and sources an outward radiation flux; that radiation interacts with a Higgs-like field, producing matter through Yang–Mills + Higgs dynamics; and localized black holes that form in the bulk act as "vents" through which bulk energy crosses onto the brane. The binding statement of the theory is an action-level commitment (the WHD action [1]): the bulk matter is an SU(2) Yang–Mills gauge field coupled to a complex SU(2)-doublet Higgs field, dressed by two entropy currents in the Andersson–Comer–Langlois (ACL) variational formalism, all minimally coupled to full 4+1-dimensional General Relativity.

Testing this program quantitatively requires a numerical-relativity (NR) engine that most existing codes are not built for: it must evolve Einstein's equations in five spacetime dimensions, on a domain that is a four-ball (not the exterior of a compact object), with a self-interacting non-Abelian gauge-plus-scalar matter sector, through the formation of horizons in the interior and the formation of a resolved domain-wall "brane" at the boundary. This paper documents the engine we built for that purpose. It is deliberately a *methods* paper in the sense that the SpECTRE and SXS collaborations write methods papers for their engines [2, 3]: its subject is the formulation, the discretization, the constraint and adaptivity machinery, the GPU implementation, and—above all—the validation. No WHD physics claim is made here; the physics-results paper is separate and will be written when the production runs land.

The engine is a nodal discontinuous Galerkin (DG) spectral-element code. It evolves the generalized-harmonic (GH) form of the $D = 5$ Einstein equations on a "hypercubed-sphere" domain (eight angular wedges per radial shell, stacked in N radial shells) using Legendre–Gauss–Lobatto (LGL) collocation and spectral differentiation matrices; there are no finite-difference stencils anywhere in the scheme. Inter-element coupling is by upwind-penalty numerical fluxes, with a targeted Lax–Friedrichs term added to the one GH characteristic family that becomes marginally propagating in the static gauge. Time integration is fourth-order Adams–Bashforth; a Hesthaven–Warburton modal exponential filter controls the top of the modal spectrum. The matter sector is evolved in first-order form (first-order Yang–Mills and first-order Higgs), a choice forced by the eigenvalue instability of the naive second-order schemes. The code runs on a single GPU (CUDA; tested on `sm_120` Blackwell and portable to `sm_80`/`sm_89`/`sm_90`), and every compute kernel is held to CPU/GPU bit-equality by a dedicated gate.

What is validated, and what is not. The vacuum GR sector is complete and validated at production scale (Sec. VII B); the multi-shell curved-initial-data configuration is stable to $T = 35\, r_h$ under the production recipe of Sec. III C–III E (Sec. VII C); CPU/GPU bit-equality holds across the kernel set (Sec. VII D); and the free-field (Layer 1) and spontaneously-broken-plus-dissipative (Layer 2) matter gates are closed (Sec. VII E). The headline production configuration—the "birth run," which evolves the nucleation of the engine and the formation of the brane from a de Sitter-throat seed—is under active bring-up: several of its pre-registered gates pass on the current build, but the energy-ledger and injection-event gates are not yet closed, and no THEORY-class production result yet exists (Sec. VII F). We are careful throughout to distinguish the two; the credibility of a methods paper is its validation section, and we do not overstate it.

Organization. Section II gives the continuum formulation: the $D = 5$ GH Einstein system, the 95 evolved GR fields, the Branch A matter action, and its first-order reductions. Section III describes the spatial and temporal discretization. Section IV describes the curved-background XCTS constraint re-solve. Section V describes adaptive mesh refinement and excision. Section VI describes the GPU implementation and the bit-equality discipline. Section VII is the validation. Section VIII states what the production runs will add.

II. CONTINUUM FORMULATION

A. The 4+1-dimensional generalized-harmonic Einstein system

We work in $D = 5$ spacetime dimensions with coordinates $x^a = (t, x, y, z, w)$, where (x, y, z, w) are the four spatial directions of the bulk, and metric signature $(-, +, +, +, +)$ ("mostly plus"). Indices $a, b, \ldots$ run over the five spacetime directions; spatial indices $i, j, \ldots$ over the four spatial directions. The gravity sector is full 4+1-dimensional Einstein–matter dynamics (the WHD action [1]); no dimensional reduction or symmetry ansatz is imposed.

The Einstein equations are evolved in generalized-harmonic (GH) form [4–6]. The gauge is controlled by a source vector H_a that prescribes the coordinate-wave operator acting on the coordinates, $\Box x^a = -H^a$; equivalently the reduced Einstein equations are written for g_{ab} with H_a (and its gradient) appearing as source terms, and the harmonic-constraint variable

$$\mathcal{C}_a \;\equiv\; H_a + \Gamma_a, \qquad \Gamma_a \equiv g^{bc}\Gamma_{abc}, \tag{1}$$

measures the departure from the prescribed gauge; $\mathcal{C}_a = 0$ on a solution. Damping of constraint violations follows the standard Gundlach–Pretorius prescription [7] with coefficients $(\gamma_0, \gamma_1, \gamma_2)$ multiplying the constraint-damping terms. For the vacuum configurations (Minkowski and Tangherlini) the code uses $(\gamma_0, \gamma_1, \gamma_2) = (0, -1, 1)$; $\gamma_1 = -1$ makes the fundamental variables propagate along the light cone and renders the principal system symmetric-hyperbolic. For the matter/brane configurations, γ_0 is promoted to a spatially varying profile that concentrates constraint damping in and around the resolved brane layer, and a separate gauge-field damping coefficient γ_A is introduced for the Yang–Mills Gauss/Lorenz constraint (Sec. II C); the specific profiles are configuration-dependent and set by the run's pre-registration.

B. Evolved GR variables: the 95 fields

The GH system is written in first-order-in-time, first-order-in-space form with 95 evolved real fields per quadrature point (`types.h`: `NF=95`; layout offsets in `gh_rhs.h`), organized as in Table I. The metric g_{ab} (symmetric 5×5: 15 components) is evolved together with a time-derivative variable $\Pi_{ab} \equiv -n^c\partial_c g_{ab}$ (where n^c is the future unit normal; 15 components) and a spatial-derivative variable $\Phi_{iab} \equiv \partial_i g_{ab}$ (four spatial derivatives of a symmetric 5×5: 60 components). The gauge-source vector H_a (5 components) completes the set. The auxiliary variables Π_{ab} and Φ_{iab} are what make the system first-order in space and time; the physical constraints $\mathcal{C}_{iab} \equiv \Phi_{iab} - \partial_i g_{ab}$ (the "three-index" constraint) and Eq. (1) are monitored during evolution and damped. (The contract's §4.1 counts 90 "GR DOFs"—metric, momentum, and spatial derivatives; the code additionally evolves the five gauge-source components H_a as part of the damped-harmonic gauge, for the 95 of Table I: $95 = 90 + 5$.)

C. The Branch A matter action

The matter content (the WHD action [1]; "Branch A") is a self-interacting field theory rather than a fluid. In addition to the Einstein–Hilbert term, the matter Lagrangian density is

$$\mathcal{L}_{\text{matter}} = -(D_a\Phi)^\dagger(D^a\Phi) - V_{\text{eff}}(\Phi) - \tfrac{1}{4}F^{\alpha}_{ab}F^{\alpha\,ab} + \mathcal{L}_{\text{entropy}} - \frac{\Lambda_{\cos}}{8\pi G}\sqrt{-g}, \tag{2}$$

TABLE I. The 95 evolved generalized-harmonic GR fields per quadrature point.

Symbol	Meaning	Components
g_{ab}	spacetime metric (symmetric 5×5)	15
Π_{ab}	time-derivative of g_{ab} (symmetric 5×5)	15
Φ_{iab}	spatial derivatives of g_{ab} ($4\times$ symmetric 5×5)	60
H_a	gauge source vector	5
	total	**95**

TABLE II. Branch A matter evolved variables, by implementation layer. Counts are per quadrature point; the "cumulative" column is the running total of matter degrees of freedom (Higgs + Yang–Mills + entropy).

Layer	Sector	Evolved variables	Components	Cumulative
1	Higgs (FO)	$\Phi + \Pi_\Phi$ (doublet + conjugate momentum)	$4 + 4 = 8$	8
1	Yang–Mills (FO)	$A^\alpha_a + E^\alpha_i$ (potential + electric field)	$15 + 12 = 27$	35
2	Entropy currents	$s^a_\Phi,\ s^a_A$	$5 + 5 = 10$	45

with the ingredients specified as follows.

Higgs sector. Φ is a complex SU(2) doublet (4 real components). The gauge-covariant derivative is $D_a\Phi = \partial_a\Phi - igA^\alpha_a T^\alpha \Phi$ with $T^\alpha = \sigma^\alpha/2$ the SU(2) generators and g the gauge coupling. The renormalized potential

$$V_{\text{eff}}(\Phi) = V_{L2}(\Phi) - V_{L2}(v) = -m^2|\Phi|^2 + \frac{2\lambda}{\Lambda}|\Phi|^4 + \frac{m^4\Lambda}{8\lambda}, \qquad |\Phi|^2 = \Phi^\dagger\Phi, \tag{3}$$

carries a constant counterterm so that the spontaneously-broken vacuum at $|\Phi| = v = \sqrt{m^2\Lambda/(4\lambda)}$ has $V_{\text{eff}}(v) = 0$: the vacuum contributes zero matter stress-energy, so no hidden vacuum energy gravitates and the bulk background curvature is controlled entirely by the explicit parameter Λ_{cos} (contract §4.2). The symmetric phase $|\Phi| = 0$ is tachyonically unstable; the physical Higgs excitation has mass $m_h = m\sqrt{2}$.

Yang–Mills sector. A^α_a is the SU(2) gauge potential ($\alpha = 1, 2, 3$; five spacetime components each). Its field strength is

$$F^\alpha_{ab} = \partial_a A^\alpha_b - \partial_b A^\alpha_a + g\,\epsilon^{\alpha\beta\gamma} A^\beta_a A^\gamma_b, \tag{4}$$

whose structure-constant term is the non-Abelian self-interaction that makes the radiation sector self-interacting. After spontaneous symmetry breaking the gauge bosons acquire mass $m_W = gv/\sqrt{2}$.

Entropy currents. $\mathcal{L}_{\text{entropy}}$ is the ACL variational Lagrangian for two entropy currents s^a_Φ and s^a_A (one five-vector each), tracking irreversible entropy production in the Higgs and Yang–Mills sectors respectively; a coupling term enforces global non-negativity of entropy production, $\nabla_a(s^a_\Phi + s^a_A) \geq 0$ (contract §4.1–4.2). These carry the sector's dissipative physics—friction, viscous transport, and radiation→matter conversion—which by the microphysics commitment (contract §4.5) must all trace to specific terms of the action rather than to hand-tuned limiters.

Parameters. Branch A has exactly four microphysics parameters—m (Higgs mass scale), λ (quartic coupling), g (gauge coupling), Λ (EFT cutoff)—and one background-geometry parameter, the bulk cosmological constant Λ_{cos} (default 0, an asymptotically-Minkowski bulk; static mode only, the dynamical mode having been retired for THEORY runs). No further parameters are admitted at the equation-of-motion level; viscosities, conversion rates, sound speeds, and threshold widths are *derived* from (m, λ, g, Λ) (contract §4.3–4.4).

D. First-order reductions of the matter sector

The matter fields are evolved in first-order (FO) form. The naive second-order ("D^2") Yang–Mills and Higgs schemes are eigenvalue-unstable on the curved, collapsing backgrounds of interest; the first-order reductions restore a well-behaved principal symbol and are the production schemes (research plan L1.3; the strong-form free Yang–Mills instability is documented, and the FO-Yang–Mills scheme was built in response). Table II summarizes the additional evolved matter degrees of freedom by implementation layer.

In the FO-Higgs reduction (`higgs_fo.h`) the doublet Φ is evolved together with a conjugate momentum Π_Φ and a spatial-gradient auxiliary $\Psi_i \equiv \partial_i\Phi$; in flat form $\partial_t\Phi = -\Pi_\Phi$, $\partial_t\Pi_\Phi = -\partial_i\Psi_i + V'(\Phi)$, $\partial_t\Psi_i = -\partial_i\Pi_\Phi$, and the reduction constraint $\Psi_i = \partial_i\Phi$ is monitored and re-established on refinement (Sec. V A). The curved-space right-hand side carries the full covariant d'Alembertian including the Hubble friction term $\propto \alpha K \Pi_\Phi$. In the FO-Yang–Mills reduction (`yangmills_fo.h`) the gauge potential A^α_a (15 components) and the electric field E^α_i (12 components) are evolved in temporal gauge, with the magnetic field B^α_{ij} obtained from F^α_{ij}; the non-Abelian Gauss and Lorenz constraints are controlled by a damping coefficient γ_A and a Dedner/GLM divergence-cleaning field. Both sectors use a local Lax–Friedrichs (Rusanov) interface flux. At Layer 2 the two entropy five-vectors are added, and the dissipative channels—the Closed-Time-Path friction $\Gamma(T)$, viscous transport, and the radiation→matter conversion—are sourced consistently so that the total matter stress-energy is conserved sector-by-sector and the second law holds at the point level. Care is taken that the stress tensor reads the first-order momentum as the Eulerian normal derivative it is (i.e. $\partial_0\Phi = -\alpha\Pi_\Phi + \beta^i\partial_i\Phi$ rather than $-\Pi_\Phi$), a subtlety that matters once the lapse departs from unity and the shift from zero.

E. The brane: $\mathbb{Z}_2$ mirror boundary and vents

The brane is realized numerically as a $\mathbb{Z}_2$ (reflection × SU(2)-center) identification of the doubled bulk (contract §2.4): the outer boundary of the four-ball domain is the fixed locus of the identification, and each field's even/odd parity at the boundary is *derived* from that identification rather than chosen field-by-field. In the code this is the `--gh-bc mirror` boundary condition. Rather than model the brane as an infinitely thin membrane with a prescribed tension and an imposed Israel junction condition, the production configuration *resolves* the brane as a kink layer in the Higgs field sitting at the equator of the birth slice; its effective tension and motion are then *measured* off the resolved layer and checked against the thin-shell (Israel) relations as residuals, not imposed as inputs. Localized black holes that form in the bulk are "vents": the only channel by which bulk energy crosses to the brane, whose energy is booked into a brane-matter density (contract §6). A vent is admitted only when an apparent horizon is actually located inside the resolved layer—never forced on, never suppressed—and the excision, horizon-finding, and energy-ledger machinery of Secs. IV–V exists to make that bookkeeping honest.

III. DISCRETIZATION

A. The hypercubed-sphere / 4-ball domain

The spatial domain is a four-ball (for matter/brane configurations, the full ball from the origin to the brane; for vacuum black-hole configurations, a shell whose inner radius stays outside the horizon, see Sec. VII B). It is tiled by a "hypercubed-sphere" construction that generalizes the six-face cubed sphere (which maps six tesseract-face charts to S^2) to *eight* tesseract cells mapped to S^3 (`hypercube_sphere_map.h`): each radial shell is divided into eight angular wedges, and N such shells are stacked radially, giving $8N$ spectral elements (`domain.h`, `build_spherical_domain`). The angular map is the equiangular (gnomonic) projection $a_i = \tan(\pi\xi_i/4)$ with $\sigma = (1 + a_1^2 + a_2^2 + a_3^2)^{-1/2}$; a smooth radial map places the shell breakpoints, with the inner shells linear in r and the outermost shell logarithmic (`radial_map.h`). For the full four-ball (`build_ball_domain`) the origin is an ordinary interior point: the ball is assembled from a tan-deformed central four-cube (containing $r = 0$), eight cube-to-sphere "blend" wedges, and the spherical shells. Refined ball elements use an exact (isoparametric-to-the-sphere) child map so that a refined face lies on the true spherical surface to round-off. Wedge connectivity and the inter-element face topology (including the face-point permutation tables) are precomputed (`wedge_connectivity.h`). Each element carries a tensor-product LGL grid (Sec. III B).

B. LGL collocation and spectral differentiation

Within each element the solution is represented nodally on a tensor product of N_p Legendre–Gauss–Lobatto (LGL) points per direction (`lgl.h`). The LGL nodes include the endpoints ± 1, so face data are simply the boundary volume nodes—no face interpolation is required. Spatial derivatives are computed by contracting with the precomputed LGL differentiation matrices (`lgl_dmatrix`) and chaining through the element inverse Jacobian to physical coordinates (`element.h`, `physical_derivatives`); integrals use the LGL quadrature weights (a diagonal, "mass-lumped" inner product, exact to polynomial degree $2N_p - 3$). **The engine has no finite-difference stencils of any kind**: all spatial derivatives—in the GH RHS, in the matter RHS, in the flux and lift operators, and in the diagnostics—are spectral. This is a hard design invariant of the code; there are no FD "fallback" paths and no hybrid spectral/FD schemes.

C. DG numerical fluxes: upwind penalty and the static-gauge Lax–Friedrichs term

Inter-element coupling uses the nodal DG penalty method: the shared face value is computed from the two element traces by a characteristic (upwind) penalty flux, lifted back into the volume by the DG lift operator (`flux.cpp`, `flux_point.h`, `face.h`). For the GH sector the flux is built from the GH characteristic decomposition (`gh_characteristics.h`, following [6, 9]): along a face normal n_i there are four families,

$$\begin{aligned} u^g_{ab} &= g_{ab}, & v_g &= -(1+\gamma_1)\,\beta_n, \\ u^0_{iab} &= P_i^k\,\Phi_{kab}\ \text{(tangential)}, & v_0 &= -\beta_n, \\ u^\pm_{ab} &= \Pi_{ab} \pm n^i\Phi_{iab} - \gamma_2\, g_{ab}, & v_\pm &= \pm\alpha - \beta_n, \end{aligned} \tag{5}$$

and the penalty upwinds each family according to its signed speed.

Two of these families require special treatment. The metric characteristic u^g propagates at $v_g = -(1+\gamma_1)\,\beta_n$ and the tangential/gauge family (together with the evolved sources H_a) at $v_0 = -\beta_n$. In the static, harmonic gauge used for the vacuum and quasi-static configurations ($\gamma_1 = -1$, shift $\beta = 0$), *both* of these speeds are identically zero: the standard upwind penalty for g_{ab} (and for H_a) across inter-element faces vanishes, and on a multi-shell domain with a curved ($g_{ab} \neq \delta_{ab}$) initial state, discretization disagreements between neighboring shells can accumulate undamped. The production remedy is a targeted Lax–Friedrichs (LF) stabilization applied to precisely the marginal u^g and H_a modes and *not* to the tangential u^0 mode (whose LF stabilization was tested and found to destabilize the scheme):

$$\texttt{GHParams.flux\_g\_visc} = 0.5 \quad \text{(production; default 0)}, \tag{6}$$

implemented as a central-difference LF penalty $\propto \texttt{flux\_g\_visc}\,(u^g_{\text{ext}} - u^g_{\text{int}})$ on the u^g and H_a modes (`flux_point.h`; the parameter default `flux_g_visc` $= 0$ is set in `gh_rhs.h` and makes the term a no-op on Minkowski data, preserving exact bit-equality of all prior single-shell regression tests). The LF term is needed only for multi-shell curved-initial-data runs; it introduces a parabolic contribution that tightens the time-step constraint to $\mathrm{d}t \lesssim 10^{-3}$ (Sec. III D), stricter than the hyperbolic CFL limit. This "L4-d" recipe is validated by the long-time closure test gate; single-shell runs need neither setting. A separate `flux_g_visc_subface` controls the LF strength on non-conforming AMR mortar faces (Sec. V A), defaulting to the conforming value.

D. Time integration

Time integration is by the fourth-order Adams–Bashforth (AB-4) linear multistep method (coefficients $[\frac{55}{24}, -\frac{59}{24}, \frac{37}{24}, -\frac{9}{24}]$; one right-hand-side evaluation per step; `adams_bashforth.h`, and the device integrator `ab_gpu.h`, `ABIntegratorGPU`). The first few steps are bootstrapped by a Runge–Kutta self-start. In a full matter run three AB-4 integrators—for the GH, Higgs, and Yang–Mills states—are advanced in lockstep on the shared orchestrator (`LayerOnePipeline`). The production time step for multi-shell curved-IC configurations is $\mathrm{d}t \le 10^{-3}$, set by the parabolic constraint of the LF face term rather than the hyperbolic CFL (Sec. III C). Because the AB start-up is not full order on its first interval, the energy-ledger diagnostics account for a bootstrap term that scales as $\mathcal{O}(\mathrm{d}t^2)$ on the first interval and judge closure from the second interval onward.

E. The Hesthaven–Warburton modal filter

High-order modal content is controlled by a Hesthaven–Warburton modal exponential filter [8] (`filter.h`, `apply_modal_filter`): the nodal solution is transformed to the modal (Legendre) basis by the element Vandermonde inverse, each mode n is attenuated by the tensor product of

$$\sigma(\eta_n) = \exp\Big[-\alpha\,\frac{(\eta_n-\eta_c)^p}{(1-\eta_c)^p}\Big] \ \ (\eta_n > \eta_c), \quad \sigma = 1\ (\eta_n \le \eta_c), \quad \eta_n = \tfrac{n}{N_p - 1}, \tag{7}$$

and the result is transformed back. The production parameters are the SpECTRE-typical $\alpha = 36$, $p = 16$, with the cutoff η_c at mode $\lfloor N_p/2 \rfloor$. The filter is applied to the GH state every 50 AB-4 steps in production (`set_filter_cadence(50)`), the AB history being restarted after each application. This cadence was fixed by a systematic sweep (the "L4-g" investigation): a subtle orchestrator bug had left the filter a silent no-op on vacuum-only multi-shell probes (the filter-enable flag was set inside the matter path rather than in the shared initialization), which is why an earlier cadence sweep produced bit-identical trajectories; once the filter initialization was hoisted into the shared `init()` path, a filter every step reduced the high-mode modal energy by roughly two orders of magnitude relative to no filter, and a cadence sweep (`filter_every` $\in \{10, 25, 50\}$) all survived to $T = 35\,r_h$ with negligible ($\sim$0.3%) Lax–Friedrichs decay of the metric amplitude and a few-percent overhead. The production operating point is `filter_every` $= 50$; single-shell runs do not require the filter for stability.

Production recipe (multi-shell, curved IC, $T \geq 22\,r_h$). For completeness, the validated recipe for multi-shell hypercubed-sphere domains with a curved GH initial state is: `flux_g_visc`$= 0.5$ (LF on the marginal u^g family), $\mathrm{d}t \leq 10^{-3}$ (parabolic CFL), and a modal filter every 50 AB-4 steps, with the vacuum constraint-damping choice $(\gamma_0, \gamma_1, \gamma_2) = (0, -1, 1)$ unchanged. This recipe was validated to $T = 35\,r_h$; see Sec. VII C.

IV. CONSTRAINT SOLVING: THE CURVED-BACKGROUND XCTS RE-SOLVE

The initial data for a configuration are constructed to satisfy the Einstein constraints; the harmonic-gauge and Lorenz-gauge initial conditions drive the GH and Yang–Mills constraint residuals to their truncation floor at $t = 0$. Beyond the initial slice, the engine also supports a *mid-flight* constraint re-solve on the curved, already-evolved slice. This is needed at the production run's mode-injection events: when a seed multiplet is injected into the evolving state (the classical stand-in for a mode becoming classical as it crosses the horizon), the injection perturbs the constraints, and rather than apply an ad-hoc "kick" the engine re-solves the constraints on the actual curved slice and *measures* the resulting jump (the WHD action [1]; the mid-flight ruling).

The re-solve is an extended conformal-thin-sandwich (XCTS) problem posed on the evolved slice. It is a coupled elliptic system for six fields—the conformal factor ψ (the Hamiltonian "$H4$" row), the product $u = \alpha\psi$ (the lapse "$L4$" row), and the four shift components $\beta^0 \ldots \beta^3$ (the weighted-York "$M4$" rows)—packed as $6n$ unknowns and solved by an inexact Newton iteration (the coupled curved-background XCTS solve). The linear systems are solved by right-preconditioned restarted GMRES (FGMRES with restart length 30; an Eisenstat–Walker forcing sequence η_k with $\eta_0 = \frac{1}{2}$, capped at 0.9, sets the inner tolerance). The preconditioner is one block Gauss–Seidel (block-Picard) sweep in the order $\psi \to u \to \beta$: each diagonal block is inverted by a few inner MINRES iterations to a loose relative residual (≈ 0.05), $\delta\psi$ feeding the u block's right-hand side and $\delta\psi, \delta u$ the β block's. As an independent check, a standalone block-Picard iteration is run from the same start and the two solutions' agreement is reported (a "two-derivation" cross-check). The elliptic operator is assembled in the same curvilinear continuous-Galerkin (CG-SEM) spectral-element discretization as the evolution (the re-solved data are re-assembled into the DG evolution state, with the CG$\leftrightarrow$DG transfer handled explicitly), and the solver runs on the run's actual AMR leaf mesh, including hanging nodes.

Two robustness features are worth recording because they are generic to elliptic re-solves on a near-de Sitter slice. First, the strong preconditioner is the *block* Gauss–Seidel structure, not a diagonal estimate: an earlier block-Jacobi *diagonal* preconditioner (a full element-by-element stiffness-diagonal probe every Newton step) dropped the long eigen-directions of the stiffness under GMRES restart and stagnated at $N_p = 6$ with $\sim 2 \times 10^5$ reduced degrees of freedom, because the restart length could not span them (RD-162); replacing it with the direct block solve fixed the stall. Second, the linearized operator on a near-time-symmetric curved slice has a near-null-space (the $\ell = 1$ Möbius near-kernel of a near-de Sitter slice) that is not deflated, so a naive GMRES stagnates or breaks down; the code deflates that mode symmetrically (PJP) and carries explicit "stall-stop" guards—a line-search stall is separated from a genuine round-off plateau (the iteration is declared converged only if the residual is within one decade of its measured round-off floor, otherwise it is flagged `STALL`), and a GMRES breakdown (including a NaN Arnoldi norm) is caught rather than dereferenced. The associated norm behavior is understood: a CG-SEM$\to$DG re-assembly leaves a fixed, resolution-independent elevated value in the pointwise (strong-form) Hamiltonian residual even when the weak-form residual is at round-off; this is a norm artifact of mixing the two discretizations, not a sign of an unconverged solve, and the injection-event gate is graded against a fixed absolute ceiling for exactly this reason rather than against a resolution-shrinking floor. **[in progress: the injection-event gate is pre-registered but not yet closed on the flown production binary; see Sec. VII F.]**

V. ADAPTIVITY AND EXCISION

A. Adaptive mesh refinement

The engine supports h-refinement (element subdivision into $2^4 = 16$ children in the four spatial directions, to a maximum level of 4) with a mortar treatment of the resulting non-conforming faces (`mortar.h`, `refinement_ops.h`). When two elements of different refinement level meet, the shared "mortar" face carries the flux exchange; the projection operators between a face and its mortar are L^2-optimal spectral interpolation/restriction (mass-weighted

Galerkin projection, following SpECTRE's mortar helpers and [8]), proven adjoint to round-off in the code's face-adjoint gate, so the non-conforming coupling is conservative and does not by itself seed constraint growth. Adaptation is driven by a modal-energy smoothness indicator, $\sqrt{E_{\rm high}/E_{\rm total}}$ with "high" modes those at order $\geq \lceil 2N_p/3 \rceil$ (`refinement_criterion.h`, after [10]), augmented by a field-magnitude sensor that catches tachyonic growth the spectral-fraction indicator misses. A "frontier" criterion force-refines the active elements bordering the excised region—expressed as a per-element mask over the excision boundary, so it works for point-based dynamical excision where there is no refinement sphere—concentrating resolution on the collapse front. Conservative density transfer is used when prolonging/restricting the conserved matter variables across a refinement level: the code transfers the reference density $q = \tilde{E}\,\det J$ rather than $\tilde{E}$ itself, so that the coordinate integral $\sum_{\rm nodes} w \det J\, \tilde{E}$ is conserved to round-off on curvilinear geometry (`refinement_transfer.h`).

A hard-won lesson recorded in the engine's history is worth stating in a methods paper because it is a general DG-NR pitfall: the mortar flux for the GH sector is *benign*; an early "GH wall" at refinement boundaries was misdiagnosed as a mortar problem, when the real cause was that the reduction constraints ($\mathcal{C}_{iab} = \Phi_{iab} - \partial_i g_{ab}$ and Eq. (1)) were not re-established on freshly-refined children. Recomputing the auxiliary variables from the metric on refinement ($\Phi_{iab} = \partial_i g_{ab}$ first, then $H_a = -\Gamma_a$) makes refinement constraint-transparent. Similarly, a matter blow-up at refinement under spontaneous symmetry breaking was cured by a broad matter modal filter (attenuating the cubic-nonlinearity aliasing of $|\Phi|^2\Phi$ into high-k modes), not by touching the mortar.

B. Excision and dynamical excision

Interior horizons are handled by excision. The engine locates apparent horizons with a Thornburg-style finder [11] (`horizon_finder.h`): it scans radial probes and bisects the outgoing-null expansion Θ_ℓ over an S^3 angular quadrature, and records the areal radius r_A defined by $A = 2\pi^2 r_A^3$; the sign-complementary search ($\Theta_{\rm in} > 0$) locates the anti-trapped, white-hole surfaces that are the object of interest for the engine. Quadrature points inside the located horizon are masked out (`excision.h`), and the excision-face boundary condition is pure characteristic outflow—no incoming characteristic families are imposed at a correctly-placed excision surface, and a per-face char-speed telemetry verifies that all incoming families vanish (the face is "outflow-cloaked," so no boundary condition is needed or imposed, as in SpEC/SpECTRE). An optional constraint-preserving Bjorhus correction on the excision-face ghost is available (`--exc-cp-bc`). Under AMR the excision detectors run on the orchestrator's current element set so that excision is AMR-aware (an early bug in which the excision mask was populated only in the no-refinement case, silently disabling excision under AMR, is fixed and gated).

Excision is *dynamical*: the excision margin is a fixed *fraction* of the located horizon radius rather than an absolute length, $r_{\rm exc} = (1-f)\,r_{\rm AH}$ with $f = 0.1$ pre-registered (the excision sphere sits at $0.9\,r_{\rm AH}$, inside the horizon; `--exc-sphere-dyn-frac 0.1`, `exc_sphere_policy.h`, `exc_sphere_growth_want_frac`), matching standard SpECTRE practice of an excision boundary a fixed 0.9–0.95 $r_{\rm AH}$ inside the horizon [12]. This scales correctly as the horizon grows and, critically, cannot cut outside the horizon; f is a resolution-independent method constant, never re-tuned against a result, and the absolute-margin form survives only as an INSTRUMENT flag. A companion rule forbids any excision cut before an apparent horizon is actually located (RD-108: in birth mode the growth policy yields zero excision radius until an areal horizon is found), which prevents a "phantom" excision of an untrapped region—a failure mode that had corrupted earlier runs. For curvature-driven excision in the engine configuration, a curvature indicator excises the collapse on its approach (before the trapped surface flips), which was validated to be load-bearing for late-time geometric quality rather than over-excision.

VI. GPU IMPLEMENTATION AND CPU/GPU BIT-EQUALITY

The production engine is a single-GPU CUDA code (C++17/CUDA; namespaces `whd::dg`; backend selected by `WHD_DG_BACKEND`$\in$ {`cpu`, `cuda`}). The full evolution—GH RHS, matter RHS, DG face/flux/lift, the AB-4 update, the modal filter, and the diagnostics—runs in-domain on one GPU (the orchestrator `MultiElementGPU`); the design deliberately does *not* treat "a finer grid on many GPUs" as a substitute for the full physical configuration on one GPU (a lesson from an earlier mis-scoped run). The build defaults to CUDA architectures `"120;90;80"`—`sm_120` (Blackwell, the primary development card, an RTX 5070 Ti), `sm_90` (Hopper, H100), and `sm_80` (Ampere, A100)—and a dedicated `sm_89` (Ada Lovelace, L4/L40S) build tree is maintained for cloud GPUs (`CMakeLists.txt`). A multi-rank/multi-GPU path (NCCL halo exchange) exists but is not the setting for the single-GPU in-domain production run.

Every compute kernel is held to CPU/GPU bit-equality by a dedicated test, at two tiers. Most kernels are held "bit-close"—a per-kernel single-evaluation tolerance of $\leq 10^{-13}$ in the max-norm, and $\leq 10^{-12}$ end-to-end (the

Layer-1 validation)—where the host and device differ only by fused-multiply-add and transcendental-library rounding. A stronger tier holds the flux and per-point kernels to *exact* bitwise equality (host-vs-device `memcmp`= 0): a single "`WHD_HD`" inline body is compiled for both host and device with fused-multiply-add disabled on the device (`--fmad=false`) and the x86-64 host built without FMA, and with `max`/`min` spelled out so no host-only library call can diverge; the non-conforming AMR sub-face mortar flux is gated against a frozen bitwise oracle. The engine's discipline is that bit-equality is necessary but *not sufficient*: every new kernel also carries a physical-correctness gate, because two implementations of a wrong formula can agree to round-off, and a multi-step (not merely single-step) CPU/GPU comparison is used, because a divergence in the multi-step evolution can hide below single-step agreement.

VII. VALIDATION

A. What "validated" means here

The program enforces a run-class discipline (CONTRACT_GATES.md §0) that we adopt in reporting validation. A *THEORY* run carries the full contract physics on ratified numerical machinery with zero configuration deviations; only a THEORY run can support a WHD physics claim. An *INSTRUMENT* run is a deliberately reduced configuration, declared in advance, used to isolate one mechanism (a bit-identity baseline, a scheme bring-up, a convergence probe); its result is a statement about the instrument, never about WHD physics. Every validation below is a statement about the *engine* (a gate on the numerics), which is exactly what a methods paper should establish; none is a physics claim. Where a gate result is quoted, it is the gate's own recorded outcome, not an estimate.

B. Vacuum: Tangherlini (5D Schwarzschild) evolution

The vacuum GR sector is complete and validated at production scale (Phase A). The $D = 5$ Tangherlini (5D Schwarzschild) black hole is set up in Schwarzschild form,

$$\gamma_{ij} = \delta_{ij} + \frac{\mu}{r^2-\mu}\, n_i n_j, \qquad \alpha = \sqrt{1-\frac{\mu}{r^2}}, \qquad \beta^i = 0, \qquad K_{ij} = 0, \tag{8}$$

(the metric and the $\partial_k g_{00}$ gradient are written analytically by `gh_rhs.cpp`, `set_tangherlini`; the remaining spatial-derivative auxiliaries Φ_{iab} and the gauge source H_a are completed by the harmonic-gauge initialization pass that sets $H_a = -\Gamma_a$, so that $\mathcal{C}_a = 0$ at $t = 0$). We note for the record two invariants that the code enforces and that are easy to get wrong in $D = 5$: (i) the spatial domain must satisfy $r_{\min} \geq 2\sqrt{\mu} = 2r_h$, because the $\mu/(r^2 - \mu)$ factor diverges at the horizon and regularizing near it manufactures artificial curvature (the hypercubed-sphere defaults to $r_{\min} = 2$ for $\mu = 1$); and (ii) one must *not* use the 4-dimensional conformal form $\gamma_{ij} = \psi^4\delta_{ij}$, which violates the $D = 5$ vacuum equations. The evolution of Eq. (8) under the full GH RHS is validated by the convergence gate `gh_full_rhs_tangherlini_convergence` (and the Phase A vacuum test suite `CUDA-0`...`CUDA-7`); the residual constraints converge under resolution refinement as expected for a spectral scheme. As an independent check of the mass diagnostics, the Misner–Sharp mass function (`misner_sharp.h`) reproduces its exact analytic anchors, $\mu_{\rm MS} \equiv 0$ on Minkowski and $\mu_{\rm MS} \equiv \mu$ on Tangherlini. Throughout, the monitored quantities are the ones the program's discipline requires—the GH constraint $\mathcal{C}_a = H_a + \Gamma_a$, the matter-aware Hamiltonian residual $H = R + K^2 - K_{ij}K^{ij} - 2\kappa\rho$, the Misner–Sharp mass profile, and integral fluxes—rather than pointwise field maxima. The measured spectral convergence is shown in Fig. 1: the GH constraint $\|\mathcal{C}_a\|_\infty$ on the vacuum Tangherlini state falls by roughly an order of magnitude per two nodes of polynomial order ($N_p = 6 \to 8 \to 10$ gives per-shell reduction factors ∼11 and ∼20), the geometric signature of a spectral scheme with no finite-difference floor.

C. Multi-shell curved-IC stability: the L4-d/L4-g closure

The nontrivial stability test of the discretization is a multi-shell hypercubed-sphere domain carrying a curved ($g_{ab} \neq \delta_{ab}$) GH initial state—the configuration in which the marginal u^g family (Sec. III C) can seed inter-shell drift. Under the production recipe (`flux_g_visc`= 0.5, d$t \leq 10^{-3}$, modal filter every 50 steps) this configuration is stable to $T = 35\,r_h$: the long-time closure test gate validates the face-flux recipe (growth factor 1.0000 over its window), and the The long-time stability gate validates survival to $T = 35\,r_h$ with the metric amplitude preserved to within the sub-percent Lax–Friedrichs decay quoted in Sec. III E. The investigation that produced this recipe is itself a validation of the diagnostic chain: a per-mode characteristic decomposition localized the late-time instability to the

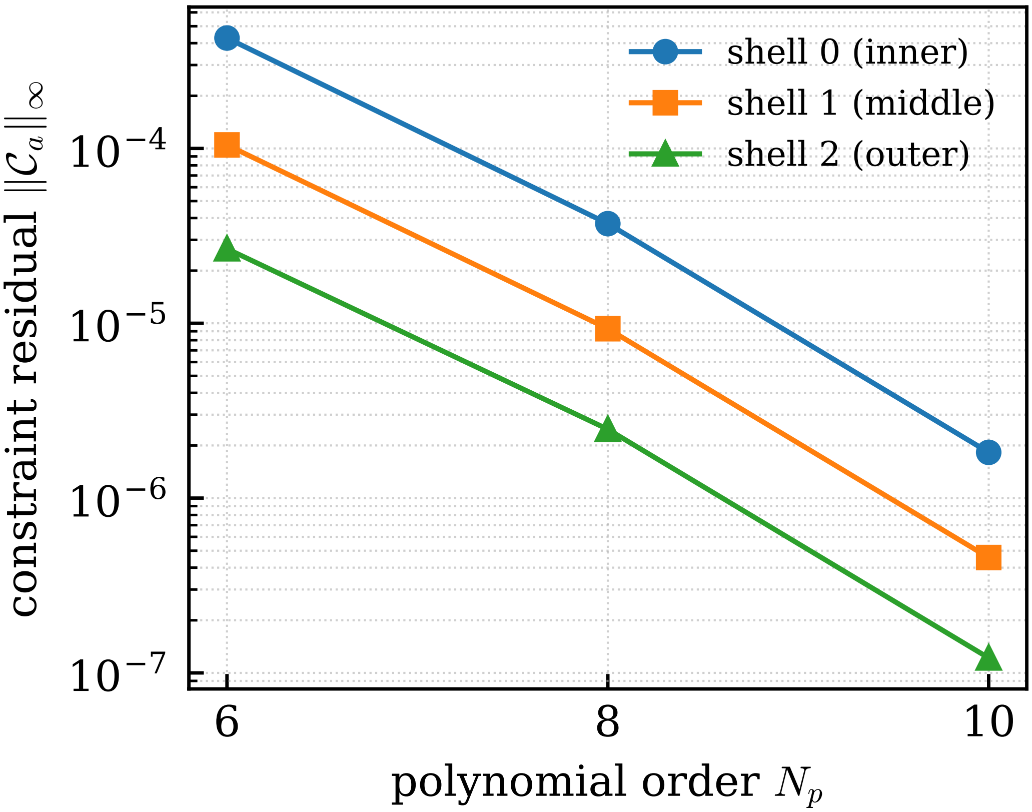

FIG. 1. Spectral convergence of the GH constraint on the vacuum Tangherlini ($D = 5$) state: $\|\mathcal{C}_a\|_\infty$ versus polynomial order N_p, per radial shell. The near-straight semilog trend is exponential (spectral) convergence.

TABLE III. Modal-filter cadence sweep for the multi-shell curved-IC configuration: every cadence survives to $T = 35\,r_h$ with the metric amplitude preserved (growth factor 1.0000) at only a few-percent overhead.

filter every (N steps)	survives $T = 35\,r_h$	metric growth	overhead
50 (production)	yes	1.0000	$\sim$2%
25	yes	1.0000	$\sim$4%
10	yes	1.0000	$\sim$10%

high-N modes (not the low-N structural modes first hypothesized), which correctly pointed at filter inadequacy and uncovered the silent-no-op filter bug of Sec. III E. The modal-energy spectrum that motivated the filter cadence is shown in Fig. 2: the per-mode energy E_N decays cleanly at $t = 0$, and by $t = 5$–$10\,r_h$ the highest modes ($N \gtrsim 18$) begin to lift while the low-N structural modes are unchanged—the signature of high-N noise, not a low-N instability. The cadence sweep (Table III) confirms that a modal filter applied every 10–50 steps suppresses this growth with negligible overhead; the production point is every 50 steps.

D. CPU/GPU bit-equality

CPU/GPU bit-equality holds across the kernel set at the tolerances of Sec. VI ($\leq 10^{-13}$ per single-step kernel, $\leq 10^{-12}$ end-to-end). This is checked kernel-by-kernel and per stage and is a standing regression gate; a large number of kernels are held bit-identical (host vs device `memcmp` $= 0$) where no transcendental-library difference intervenes. Table IV lists representative kernels and their measured host–device max-norm residuals; all sit at or below the 10^{-13} per-kernel tolerance, and the multi-step end-to-end pipeline stays below 10^{-12}.

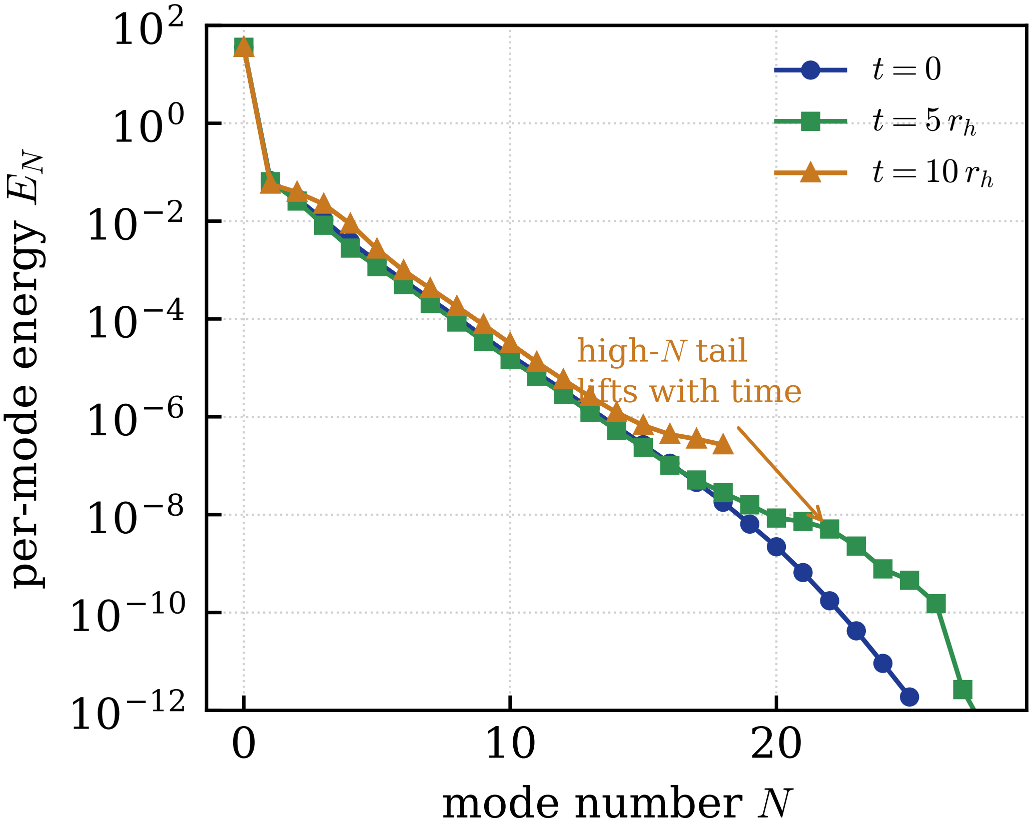

FIG. 2. Per-mode energy spectrum E_N of the multi-shell curved-IC configuration at $t = 0, 5, 10\,r_h$ (characteristic mode decomposition). The high-N tail lifts with time—the target of the modal filter—while the low-N modes are unaffected.

TABLE IV. Representative CPU/GPU bit-equality residuals (measured max |CPU − GPU|). Per-kernel tolerance 10^{-13}; end-to-end 10^{-12}.

Kernel / stage	max |CPU − GPU|
Spectral derivatives (poly / sin / random)	3.3×10^{-14}
GH volume RHS (Tangherlini)	4.5×10^{-15}
Full GH RHS (Kerr–Schild, driver+gauge)	1.8×10^{-15}
Face outward normals	2.2×10^{-16}
Lift / face correction	1.8×10^{-15}
Upwind penalty flux	1.1×10^{-16}
AB-4 update (200 steps)	1.4×10^{-15}
Two-element full pipeline	2.1×10^{-15}
End-to-end evolution (200 steps)	1.6×10^{-14}

E. The matter sector: Layers 1 and 2

The field-theoretic matter sector has been brought up and gated in layers.

Layer 1 (free fields + GR). The free Higgs doublet and the free SU(2) Yang–Mills field are implemented on the GPU in first-order form, coupled to the GH GR sector, with CPU/GPU bit-equality gates on all Layer-1 kernels (the Layer-1 bit-equality gates). The first-order substrate survives a long-time single-shell test to $T = 50$. Layer 1 closure is the Phase B closure criterion.

Layer 2 (SSB + entropy currents + dissipation). Higgs spontaneous symmetry breaking (tachyonic roll to $|\Phi| = v$) is demonstrated; the Closed-Time-Path friction $\Gamma(T)$ and the Higgs entropy current are implemented and match the 3+1-dimensional QFT benchmark to $\sim 10^{-16}$; the Yang–Mills entropy current and radiation friction are implemented with the second law gated at the kernel level; and a fully-coupled Layer-2 long-time gate (the coupled Layer-2 long-time test) sustains SSB Higgs + Yang–Mills + both entropy currents + both frictions to $T = 5$ (validated to $T = 10$). We state the honest scope of the Layer-2 gate as the code does: it runs with the metric frozen (`--freeze-gh`) to isolate the matter action from the gravitational-collapse confound, the first-order substrate runs the two sectors in their Abelian/free reduction (a fully gauge-covariant FO cross-coupling at $g \neq 0$ is a follow-on), and the global second-law integral is gated on the static metric where it is clean (the local production is non-negative at the kernel level regardless).

Engine and adaptivity infrastructure. The white-hole engine radiates outward and forms anti-trapped surfaces; dynamical excision, AMR↔excision coordination, and AMR on dynamic SSB matter are built and gated; and a long-time engine run with AMR plus curvature excision reaches a bounded stationary state to $T = 20$ (gate the engine+AMR+excision long-time test). These are engine-level validations of the adaptivity/excision machinery; per the run-class discipline (Sec. VII A) they are not, on their own, WHD physics claims.

F. The birth run (in progress)

The headline production configuration is the "birth run": the evolution, from a de Sitter-throat seed on a closed four-sphere birth slice, of the tachyonic growth of the engine, the formation of the resolved brane kink layer at the equator, and—if the physics produces them—the nucleation of vents. Under the current rules this is the one and only THEORY-stampable configuration, and it carries a specific pre-registered stack (the $\mathbb{Z}_2$ mirror boundary, a derived birth radius $r_{\max} = a^*$, a charge-projected seed, the full covariant first-order matter RHS, friction/entropy/conversion, the counterterm and areal-radius clock, and armed—never forced—vent machinery) with a set of required gates.

This configuration is under active bring-up. Several of its pre-registered layer gates pass on the current build at their pre-registered resolutions—the mirror parity floor, the kink-width convergence, the Israel-residual gate, and the layer-diagnostics printout gate. Others are not yet closed: the energy-ledger closure gate does not yet pass (the residual is understood—an $\mathcal{O}(\mathrm{d}t^2)$ Adams–Bashforth start-up term plus the entropy-transport bookkeeping that is itself still being brought up), and the measured-motion gate, the thin-limit reproduction gate, the vent-horizon gate, and the injection-event gate are pending. Consequently **no THEORY-class birth-run result yet exists**; the canonical production configuration presently stamps INSTRUMENT (with declared pending builds), which is the honest state.

[to add: This subsection is the one that the physics-results paper turns into a result. Once the birth run stamps THEORY—all required gates (all required gates) PASS on the flown binary with zero registry deviations—insert: (i) the passing gate ledger with resolutions; (ii) the constraint-monitor time series ($\mathcal{C}_a$, matter-aware Hamiltonian residual, Misner–Sharp mass, integral fluxes) for the production run; (iii) the resolved-layer diagnostics (u_{core}, σ_{eff}, δ_{layer}, H_b); and (iv) any convergence/self-convergence study. Until then, report only the engine methods above, not a birth-run outcome.]

G. Summary of validation status

Table V summarizes the state. The engine (vacuum GR, the multi-shell curved-IC discretization, CPU/GPU bit-equality, and the layered matter bring-up) is validated by named gates; the full production birth run is in progress.

VIII. CONCLUSION AND OUTLOOK

We have described a 4+1-dimensional discontinuous Galerkin numerical-relativity engine for the White Hole Dynamics program: the generalized-harmonic Einstein system in $D = 5$ with 95 evolved fields; a self-interacting SU(2) Yang–Mills + Higgs-doublet + entropy-current matter sector in first-order form; a hypercubed-sphere/4-ball domain with LGL collocation and strictly spectral differentiation; upwind-penalty DG fluxes with a targeted Lax–Friedrichs stabilization of the marginal generalized-harmonic family; AB-4 time stepping and a Hesthaven–Warburton modal filter; a curved-background XCTS re-solve for mid-flight injection events; mortar AMR and sphere-based dynamical excision; and a single-GPU CUDA implementation held to per-kernel CPU/GPU bit-equality.

The validation in hand establishes the engine, not the physics: vacuum Tangherlini evolution, long-time stability of the multi-shell curved-IC configuration to $T = 35\, r_h$, CPU/GPU bit-equality, and the closure of the Layer-1 and Layer-2 matter gates. The production birth run—the configuration that will turn these methods into a physics result—is under active bring-up: its layer, kink-width, Israel-residual, and diagnostics gates pass, while its energy-ledger, measured-motion, thin-limit, vent-horizon, and injection-event gates are not yet closed. When they are,

TABLE V. Validation status. "Gated" = validated by the named regression gate at the time of writing; "in progress" = built but not yet closed by its gate. Production numbers to be inserted after the runs (see the *[to add:]* markers in the text).

Component	Status	Evidence / gate
Vacuum Tangherlini ($D = 5$) evolution	Gated	constraint-norm convergence; CPU/GPU regression
Multi-shell curved-IC face-flux recipe (L4-d)	Gated	long-time closure test
Multi-shell long-time stability to $T = 35\,r_h$ (L4-g)	Gated	$T = 35\,r_h$ stability test
CPU/GPU bit-equality (per kernel / end-to-end)	Gated	per-kernel gates ($\leq 10^{-13}/\leq 10^{-12}$)
Layer 1 (free $\Phi + A^{\alpha}$+GR)	Gated	Layer-1 bit-equality; $T = 50$ FO survival
Layer 2 (SSB + entropy + friction)	Gated (scoped)	the coupled Layer-2 long-time test (frozen-GH, static-metric 2nd law)
Engine + AMR + excision infrastructure	Gated	the engine+AMR+excision long-time test ($T = 20$)
Curved-background XCTS re-solve	Built	unit gates; injection-event gate pending
Birth-run layer gates (parity, width, Israel-residual, refined-mesh)	Gated	pre-registered resolutions, current build
Birth-run ledger/motion/thin-limit/vent/injection	In progress	ledger gate open; motion, thin-limit, vent-horizon, injection gates pending
THEORY-class birth-run result	**Not yet**	no clean THEORY stamp yet

the physics-results paper will report (i) a fully-gated THEORY-class birth run, (ii) the constraint and conservation monitors through engine formation, brane formation, and venting, and (iii) the resolved-layer and vent diagnostics that make the WHD picture quantitative. This methods paper is the foundation for that report and the reference for the engine's numerics.

ACKNOWLEDGMENTS

The author thanks Amazon Web Services, and its representative Cristian Miranda, for the cloud computing resources that made the numerical experiments reported here possible.

[1] The White Hole Dynamics research program: the physics action and the companion papers in this series (the derived brane, the birth run, the dark-energy chain, and the string completion), in preparation.

[2] N. Deppe *et al.*, *SpECTRE: A task-based discontinuous Galerkin code for relativistic astrophysics*, spectre-code.org.

[3] The SXS Collaboration, *Spectral Einstein Code* (SpEC), black-holes.org.

[4] H. Friedrich, "On the hyperbolicity of Einstein's and other gauge field equations," Commun. Math. Phys. **100**, 525 (1985).

[5] D. Garfinkle, "Harmonic coordinate method for simulating generic singularities," Phys. Rev. D **65**, 044029 (2002).

[6] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, "A new generalized harmonic evolution system," Class. Quantum Grav. **23**, S447 (2006).

[7] C. Gundlach, J. M. Martín-García, G. Calabrese, and I. Hinder, "Constraint damping in the Z4 formulation and harmonic gauge," Class. Quantum Grav. **22**, 3767 (2005).

[8] J. S. Hesthaven and T. Warburton, *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications* (Springer, 2008).

[9] S. A. Teukolsky, "Formulation of discontinuous Galerkin methods for relativistic astrophysics," J. Comput. Phys. **312**, 333 (2016).

[10] C. Mavriplis, "Adaptive mesh strategies for the spectral element method," Comput. Methods Appl. Mech. Eng. **116**, 77 (1994); see also D. Houston and E. Süli, "A note on the design of *hp*-adaptive finite element methods," Comput. Methods Appl. Mech. Eng. **194**, 229 (2005).

[11] J. Thornburg, "A fast apparent-horizon finder for three-dimensional Cartesian grids in numerical relativity," Class. Quantum Grav. **21**, 743 (2004).

[12] D. A. Hemberger *et al.*, "Dynamical excision boundaries in spectral evolutions of binary black hole spacetimes," Class. Quantum Grav. **30**, 115001 (2013); M. A. Scheel *et al.*, Phys. Rev. D **79**, 024003 (2009).