From Measure to Flux: Energy-Conjugate Bookkeeping and Exact Semidiscrete Scattering in Finite Radial Wave Models

Jeffery Huckstead
Cerebral Graphix ORCID 0009-0007-0234-2177

Working draft v0.3 (production), 2026-07-13

Abstract

We study a finite spatial window on the radial wave equation—the setting in which a bounded observer accounts for wave energy on a Schwarzschild background—and develop its discrete bookkeeping to the point of exact solvability. For the standard second-order semidiscretization with trapezoid endpoint weights, we construct the boundary flux conjugate to the discrete energy, extend it to a one-parameter characteristic reflection family with an exact energy identity, and derive the exact per-mode reflection coefficient of the family: $r(k) = (\cos(kh/2)-\alpha)/(\cos(kh/2)+\alpha)$. The discrete boundary is therefore the continuum characteristic boundary with a secant-dressed impedance, $\alpha_{\text{eff}}(k) = \alpha \sec(kh/2)$, equivalently a Möbius composition of the target reflection coefficient with one additional factor $\beta_k = \tan^2(kh/4)$. From the exact symbol we obtain, for Gaussian packets: an exact odd reflected-power law whose first two defect orders share the cubic kernel $\mathcal{R}(1-\mathcal{R}^2)$, with extremum at $1/\sqrt{3}$; the packet impedance coefficient $\kappa = \frac{3}{16}(h/\sigma)^2[1-\frac{5}{48}(h/\sigma)^2]$; the continuum-transparent target member's reflected-power floor $\rho_0 = \frac{15}{1024}(h/\sigma)^4$; for the leading continuum Gaussian energy weight, a spectral mean–variance decomposition that is exactly 3/5–2/5; and a $1/|\mathcal{R}|$ reconstruction law near the transparent point. The leading packet coefficients κ and ρ_0 close against frozen solver receipts at relative discrepancies 4.71×10^{-9} and 1.41×10^{-8}; higher-order and reconstruction-law checks are reported at their own independently tabulated accuracies. Finally, we close the gap between the semidiscrete theory and the fully discrete solver: a boundary-parameter-dependent timestep policy converts the classical dissipation of fourth-order Runge–Kutta integration into a signed, odd-in-$\mathcal{R}$ reflected-power bias, whose scaling law $h^5\sigma^{-6}$ we derive, and which a pre-registered timestep-halving falsifier removes to 96.407% of the observed residual, meeting a locked target to 0.41%. The complete attribution chain—policy → branch-dependent dissipation → odd bias → matched-pair closure → locked falsifier—was not found in the sources examined.

1 Introduction and contribution boundary

A finite observer of a radial wave field owns a window $[x_L, x_R]$, an energy functional on that window, and two boundaries through which energy leaves or returns. Whether the observer is a numerical relativist truncating a Schwarzschild exterior or an accountant of any bounded wave system, the same discipline applies: every claimed transfer must be measured at a declared boundary by a flux that is *conjugate* to the energy being evolved. This paper carries that discipline through a complete, exactly solved example and reports what it caught.

The model is the 1+1 radial wave equation with a potential, semidiscretized by the standard three-point Laplacian with trapezoid endpoint weights. Our contributions divide cleanly into what is inherited and what is new. Inherited, and used as foundation: the summation-by-parts energy identity for characteristic boundary closures and the reflection-coefficient parametrization $\alpha = (1 -$

$\mathcal{R})/(1+\mathcal{R})$, both due to Erickson, Kozdon and Harvey [1] within the SBP–SAT tradition [2, 3, 4]; the fact that spurious numerical reflection is analyzable mode by mode, from the group-velocity and boundary-reflection literature [5, 6, 7]; and the fact that classical RK4 is mildly dissipative for oscillatory modes, a textbook property [8, 9, 10].

New, in the sense that we did not find them in the sources examined: the exact per-mode reflection symbol of this closure (Theorem 1); its reading as an exact multiplicative, α-independent impedance dressing $\alpha\sec(kh/2)$ and as a Möbius composition; the exact odd reflected-power law (Theorem 2) with its cubic kernel; the coefficient-level Gaussian packet laws κ, ρ_0, the 3/5–2/5 split, and the near-zero reconstruction law; and the fully discrete attribution chain closed by a locked falsifier. We claim no element of standard numerical analysis as new; we claim the chain, and we state every claim with its hypotheses and its receipt.

2 The finite-observer radial model

Let $\psi(x,t)$ satisfy

$$\partial_t^2\psi \;=\; \partial_x^2\psi \;-\; V(x)\,\psi, \qquad x\in[x_L,x_R], \tag{1}$$

where x is the tortoise coordinate $x = r + 2M\ln(r/2M-1)$ on a Schwarzschild exterior and, for the scalar $(s=0)$ master field of multipole ℓ,

$$V(r) \;=\; \Big(1-\tfrac{2M}{r}\Big)\Big(\tfrac{\ell(\ell+1)}{r^2}+\tfrac{2M}{r^3}\Big), \tag{2}$$

the $s=0$ member of the Regge–Wheeler family [11]. The potential is a barrier peaked near $r\approx 3M$ and decaying toward both ends of the window. Two regimes organize the paper: the *flat laboratory* $V\equiv 0$, in which every boundary statement below is exact and certifiable in isolation, and the *barrier testbed*, in which the certified boundaries meter a genuine scattering problem. All exact boundary results are proved in the flat laboratory and applied to the testbed only through the energy identity, whose validity does not require $V=0$.

3 Geometry-weighted discrete energy

On the uniform grid $x_j = x_L + jh$, $j=0,\dots,N$, with $v_j=\dot\psi_j$, define

$$E_h \;=\; \frac{h}{2}\sum_{j=0}^{N} w_j v_j^2 \;+\; \frac{1}{2h}\sum_{j=0}^{N-1}\big(\psi_{j+1}-\psi_j\big)^2 \;+\; \frac{h}{2}\sum_{j=0}^{N} w_j V_j\psi_j^2, \qquad w_0=w_N=\tfrac12,\ w_j=1 \text{ otherwise.} \tag{3}$$

The endpoint half-weights are not cosmetic. They are the trapezoid quadrature of the continuum energy, and—as the next section shows—they are precisely what makes a factor-of-two endpoint closure the conjugate partner of the edge-gradient energy. Replacing the trapezoid endpoint half-weights by unit weights breaks the exact edge-flux closure used here; the resulting energy is not conjugate to the same endpoint flux identity for this family.

4 Conjugate boundary flux

Evolve the interior by $\ddot\psi_j = (\psi_{j+1}-2\psi_j+\psi_{j-1})/h^2 - V_j\psi_j$ and close the endpoints conservatively by

$$\ddot\psi_0 \;=\; \frac{2(\psi_1-\psi_0)}{h^2} - V_0\psi_0, \qquad \ddot\psi_N \;=\; -\frac{2(\psi_N-\psi_{N-1})}{h^2} - V_N\psi_N. \tag{4}$$

Differentiating E_h and telescoping the gradient sum leaves boundary residues $-v_0 g_0$ and $+v_N g_{N-1}$, where $g_j = (\psi_{j+1} - \psi_j)/h$ is the one-sided edge gradient; the half-weighted endpoint equations produce exactly the cancelling terms. The conservative closure conserves E_h identically, and the *conjugate boundary fluxes* are the one-sided edge expressions

$$F_a = -v_0\, g_0, \qquad F_b = -v_N\, g_{N-1}, \qquad \frac{dE_h}{dt} = F_a - F_b \quad \text{exactly.} \tag{5}$$

This identity is the paper's accounting primitive, and it already caught one error worth recording. The centered flux $F = -v_j(\psi_{j+1} - \psi_{j-1})/2h$, the "obvious" discretization, is not conjugate: at the left face $F_{\text{centered}} - F_{\text{conj}} = -\frac{h}{2} v_j\, \psi_{xx} + O(h^3)$, with the sign reversed at the right face, and in practice the mismatch degrades the ledger's closure order from $p \approx 2$ to $p \approx 1$. A flux that looks consistent is not thereby conjugate; the energy chooses its flux.

5 The characteristic reflection family

Following [1], parametrize boundary impedance by a target reflection coefficient $\mathcal{R} \in (-1, 1]$ through $\alpha = (1 - \mathcal{R})/(1 + \mathcal{R})$, and damp the conservative closure:

$$\ddot{\psi}_0 \;=\; \frac{2(\psi_1 - \psi_0)}{h^2} - V_0\psi_0 - \frac{2\alpha_L}{h}\, v_0, \tag{6}$$

mirrored at the right face. The same telescoping now yields the exact, sign-definite energy law

$$\frac{dE_h}{dt} \;=\; -\,\alpha_L\, v_0^2 \;-\; \alpha_R\, v_N^2, \tag{7}$$

the standard SBP–SAT identity for this family [1, 2]. The family's anchors are exact members: $\mathcal{R} = +1$ ($\alpha = 0$) is the bare conservative closure and is Neumann-type; $\mathcal{R} = 0$ ($\alpha = 1$) is the continuum-transparent target member; $\mathcal{R} = -1$ is Dirichlet, implemented strongly because the α-chart is singular there—the family requires two coordinate charts, each endpoint anchor strong in its own chart. Measured anchor receipts: reflected energy fractions 1.000000, 2.759×10^{-9}, and 1.000000 with the correct parities, at time-ledger closures of order 10^{-14}.

6 The exact semidiscrete reflection symbol

Theorem 1 (exact boundary symbol)**.** *Let the boundary lie in a constant-coefficient region in which $V = 0$ over the neighborhood required by the discrete plane-wave ansatz, and let the face carry the closure of Section 5 with parameter $\alpha > 0$. For discrete plane waves with dispersion $\omega(k) = \frac{2}{h}\sin(kh/2)$, the reflection amplitude of the face is exactly*

$$r(k) \;=\; \frac{\cos(kh/2) - \alpha}{\cos(kh/2) + \alpha}. \tag{8}$$

Proof. Insert $\psi_j = e^{-i\omega t}(e^{-ikjh} + r\, e^{ikjh})$ into the endpoint equation. The conservative terms cancel identically against the dispersion relation, leaving $\sin(kh)\,(1 - r) = \alpha h\omega\,(1 + r)$; since $\sin(kh) = h\omega\cos(kh/2)$, the closed form follows. □

The anchors are immediate: $\alpha = 0$ gives $r \equiv 1$; $\alpha \to \infty$ gives $r \equiv -1$; and every k-dependence enters through the single function $\cos(kh/2)$. The hypotheses are part of the claim: the theorem is semidiscrete (time is continuous), local ($V = 0$ over the modal stencil), and modal (propagating discrete plane waves); the fully discrete solver is treated in Section 11, and the theorem's own energy identity forbids growing boundary modes for $\alpha > 0$.

7 Effective impedance defect and Möbius representation

Theorem 1 admits two exact readings. First, since $r(k) = (1-\alpha_{\text{eff}})/(1+\alpha_{\text{eff}})$ with

$$\alpha_{\text{eff}}(k) \;=\; \alpha \,\sec(kh/2) \;=\; \alpha\Big(1 + \tfrac{(kh)^2}{8} + \tfrac{5(kh)^4}{384} + \cdots\Big), \tag{9}$$

the discrete face is the continuum characteristic boundary with a *multiplicative, α-independent, frequency-dependent* impedance dressing. The defect is a property of the discretization alone; the target $\mathcal{R}$ never enters it. Second, writing $\beta_k = \tan^2(kh/4)$, the symbol is the Möbius composition

$$r(\mathcal{R}, k) \;=\; \frac{\mathcal{R} - \beta_k}{1 - \beta_k\,\mathcal{R}}, \tag{10}$$

i.e. a consistent discretization composes the target reflection coefficient with one additional factor. The bilinear impedance–reflection map and its composition algebra are classical in transmission-line theory and Schur analysis; we offer this as a Möbius composition with a Blaschke-type scattering interpretation, not as new algebra.

Remark 1. The per-mode odd-kernel coefficient is $2\beta_k = (kh)^2/8 + (kh)^4/192 + O((kh)^6)$. To make the distinct quartic coefficients explicit, set $\delta = \sec(kh/2) - 1$. Since $\beta_k = \delta/(2+\delta)$,

$$2\beta_k = \frac{2\delta}{2+\delta} = \delta - \frac{\delta^2}{2} + O(\delta^3).$$

With $\delta = (kh)^2/8 + 5(kh)^4/384 + O((kh)^6)$, this gives $2\beta_k = (kh)^2/8 + (5/384 - 1/128)(kh)^4 + O((kh)^6) = (kh)^2/8 + (kh)^4/192 + O((kh)^6)$. Thus the impedance-defect quartic coefficient $5/384$ and the kernel quartic coefficient $1/192$ differ exactly by the map's second-order term.

8 The exact odd reflected-power law

Theorem 2 (odd reflected-power law)**.** *For the symbol of Theorem 1, the odd-in-$\mathcal{R}$ part of the reflected-power error $\rho - \mathcal{R}^2$, per mode, is exactly*

$$\delta_{\text{odd}}(\mathcal{R}, k) \;=\; -\,\frac{2\beta_k\,(1-\beta_k^2)}{\big(1 - \beta_k^2\mathcal{R}^2\big)^2}\,\mathcal{R}\,(1-\mathcal{R}^2). \tag{11}$$

Proof. With $A = (\mathcal{R} - \beta)(1 + \beta\mathcal{R})$ and $B = (\mathcal{R} + \beta)(1 - \beta\mathcal{R})$, one finds $A - B = -2\beta(1-\mathcal{R}^2)$ and $A + B = 2\mathcal{R}(1-\beta^2)$, whence $r(\mathcal{R})^2 - r(-\mathcal{R})^2 = (A^2 - B^2)/(1-\beta^2\mathcal{R}^2)^2$ gives the claim. ☐

Three consequences carry the section. (i) To leading order in the defect the law is a cubic kernel, $\delta_{\text{odd}} \approx -\kappa\,\mathcal{R}(1-\mathcal{R}^2)$ with $\kappa = \langle 2\beta\rangle_W$ a packet average, and the kernel shape is preserved exactly through the first two defect orders (claim 8); beyond that, the Möbius denominator supplies the corrections, and we claim no all-orders purity. (ii) The kernel's extremum sits at $|\mathcal{R}| = 1/\sqrt{3}$, a statement about the map $\mathcal{R}(1-\mathcal{R}^2)$, not about geometry (claim 8). (iii) The kernel coefficient measured across 950 symmetric pairs of solver runs is constant to 1.9×10^{-7} relative—the numerical face of the α-independence proved in Section 7. Figure 1 shows the executed notebook comparison: the measured odd component, the first-order cubic law, and the exact effective-impedance Cayley map on the same axes, with the $1/\sqrt{3}$ extremum marked explicitly.

9 Gaussian packet asymptotics

Certification runs initialize a Gaussian $\psi_j = A\, e^{-(x_j - x_0)^2/2\sigma^2}$ with centered-gradient velocity data $v_j = -(\psi_{j+1} - \psi_{j-1})/2h$. Decomposing this data into exact discrete movers gives the measured energy weight

$$W(k) \;\propto\; \sin^2(kh/2)\,\big(1 + \cos(kh/2)\big)^2\, e^{-\sigma^2 k^2}, \tag{12}$$

the $\sin^2$ from the discrete gradient energy and the $(1 + \cos)^2/4$ from the mover decomposition of centered-gradient initialization. With $u = (h/\sigma)^2$ and Gaussian energy moments $\langle k^2\rangle_E = 3/2\sigma^2$, $\langle k^4\rangle_E = 15/4\sigma^4$:

Proposition 1 (packet impedance coefficient)**.** *For the stated family and centered-gradient Gaussian initialization,*

$$\kappa \;=\; \langle 2\beta\rangle_W \;=\; \frac{3}{16}\, u\Big[1 - \frac{5}{48}\, u + O(u^2)\Big]. \tag{13}$$

The $5/48$ arises from a two-term budget, each term one integral: the per-mode kernel quartic contributes $+\frac{5}{48}u$ through $\langle k^4\rangle_E$, and the discrete weight's covariance contributes $-\frac{10}{48}u$. The hypotheses are named because they are load-bearing: the coefficient is Gaussian-specific and *initialization-specific*—exact discrete right-mover data would alter the weight factor and hence the coefficient. Neglected terms are itemized in Appendix A: left-mover denominator content at relative $O(u^2)$, sampling aliasing at $O(e^{-(\pi\sigma/h)^2})$, domain tails exponentially small, and no time-discretization error at the semidiscrete level. Receipt: at $(n, \sigma) = (1601, 3)$ the frozen closure ledger records analytic $8.137652908 \times 10^{-5}$ against solver $8.137652946 \times 10^{-5}$, relative discrepancy 4.707×10^{-9}; the next-order coefficient is measured as 0.104132 against $5/48 = 0.104167$.

10 Continuum-transparent member, spectral split, and signed reconstruction

At $\mathcal{R} = 0$ the symbol gives $r = -\beta_k$, so the continuum-transparent target member reflects, on the finite grid, the packet-averaged square of the defect:

$$\rho_0 \;=\; \langle \beta^2\rangle_W \;=\; \frac{15}{1024}\, u^2 + O(u^3), \tag{14}$$

verified across a seven-octave resolution ladder at observed order 4.000 down to machine precision, with the frozen ledger recording relative discrepancy 1.410×10^{-8} at the certification point (claim 10). The transparent power decomposes by Cauchy–Schwarz into a mean-power fraction $\langle\beta\rangle^2/\langle\beta^2\rangle = \langle k^2\rangle_E^2/\langle k^4\rangle_E \le 1$ and a variance remainder; for the leading continuum Gaussian energy weight the mean-power and variance fractions are exactly $3/5$ and $2/5$, and the frozen finite-grid receipts approach these values with the separately reported discrepancies (0.599984 and 0.40007; claim 10). The variance floor then controls signed reconstruction near the transparent point: with $\mathcal{R}_{\rm rec} = s\sqrt{\rho}$ and parity supplying the sign s,

$$\mathcal{R}_{\rm rec} - \mathcal{R}_{\rm eff} \;\approx\; \frac{(2/5)\,\rho_0}{2\,\mathcal{R}}, \tag{15}$$

a $1/|\mathcal{R}|$ law with measured local log-slope -1.00116 (claim 10). Figures 2 and 3 show the actual near-zero reconstruction runs and the corresponding local scaling rate approaching the predicted -1 law. This is a *reconstruction singularity of the estimator*, not a physical divergence; below the declared sign threshold the reconstruction is undefined by policy. Finally, the extremum of Section 8 acquires a derived displacement

$$\Delta\mathcal{R}_* \;=\; \frac{4\sqrt{3}}{27}\,\frac{\langle\beta^3\rangle}{\langle\beta\rangle} \;=\; \frac{35\sqrt{3}}{6912}\Big(\frac{h}{\sigma}\Big)^4, \tag{16}$$

equal to $+1.652\times10^{-9}$ at the certification point. The frozen ledger records the corresponding solver-estimator offset as 7.148×10^{-9} with the status annotation adopted here verbatim: the estimator value is sign-consistent but includes known floor and systematic contributions and is not a direct resolution of the symbol shift.

11 Fully discrete attribution

Everything to this point is semidiscrete; the solver is not. The production integrator is classical RK4 under the timestep policy $\Delta t = \min(0.2h,\, 0.35h/2\alpha_{\max})$, and the policy has a consequence the semidiscrete theory cannot see: runs at $+\mathcal{R}$ (where $\alpha < 1$) and $-\mathcal{R}$ (where $\alpha > 1$) execute at *different absolute timesteps.* Classical RK4 is mildly dissipative for oscillatory modes—$|R_4(ix)|^2 = 1 - x^6/72 + x^8/576$ exactly, a textbook fact [8, 9, 10] rederived in Appendix B—and over a run of duration T the modal loss accumulates to $D = (T/72)\langle\omega^6\rangle_E\,\Delta t^5$, with $\langle k^6\rangle_E = 105/8\sigma^6$ for the Gaussian weight. Branch-dependent timesteps therefore inject a *signed, odd-in-$\mathcal{R}$* reflected-power bias, $\epsilon_{\rm odd} \approx -\frac{\rho}{2}(D_+ - D_-)$, scaling as $h^5\sigma^{-6}$ under $\Delta t \propto h$—so that at fixed h/σ the bias scales as $1/h$. A deliberately matched pair, $(n,\sigma) = (801, 3)$ and $(1601, 1.5)$ with $h/\sigma = 1/24$ in both, made the mechanism unmistakable: the exact-symbol predictions of the two configurations are identical, their negative-branch solver outputs agree with maximum absolute difference 6.03×10^{-14} (2.01×10^{-13} relative, across 31 pairs), and their estimator residuals differ by a factor of two—4.8151×10^{-8} against 9.6311×10^{-8}—exactly the $1/h$ law. A parameter-free exact-mode amplification model, using only $|R_4|^{2N}$ per mode at the recorded per-row timesteps, reproduces the residual curves pointwise to 0.1–0.4% and was independently re-derived, cross-implementation and cross-agent, on hash-pinned inputs.

The chain was then closed by a pre-registered falsifier. The leading prediction—halve the positive-branch cap, collapse the positive channel by $2^5 = 32$—was locked, together with a stronger branch-aware target of $2.9775682586\times10^{-8}$ for the new estimator offset, before any rerun. The realized offset was $2.9897167741\times10^{-8}$: 0.408% from the locked target, removing 96.407% of the prior fully-discrete residual, with the small remainder recorded rather than repaired. We classify the mechanism as *numerically certified*: derived in form, reproduced across implementations and agents, and falsifier-tested, for this integrator and this policy—not as a universal theorem about time integration. The episode's role in this paper is exactly that of a validation success for the accounting discipline of Sections 3–5: the ledger opened the fully discrete discrepancy into named channels, and the channels emptied into one derived mechanism. A method with $|R(ix)| = 1$ on the relevant imaginary-axis linear modes would remove the pure modal amplitude-loss component of this RK4 channel, while leaving phase error and the semidiscrete boundary-impedance defect untouched.

12 Numerical certification and the locked falsifier

The paper's evidentiary architecture is uniform across claims. Every analytic coefficient was locked in a pre-PDE prediction ledger, hashed and timestamped, before the corresponding solver campaign; every solver receipt is a configuration-stamped file whose SHA-256 appears in its parent receipt; the

Table 1: Locked timestep-halving falsifier settlement. All values from the frozen falsifier receipt (`fz/...falsifier_receipt.json`), locked 2026-07-12T23:12Z.

quantity	value
old estimator offset / residual	$1.2274770156 \times 10^{-7}$ / $9.6310748932 \times 10^{-8}$
locked leading 1/32 target	$2.9446663531 \times 10^{-8}$
locked branch-aware target (pre-registered)	$2.9775682586 \times 10^{-8}$
realized new offset	$2.9897167741 \times 10^{-8}$
closure vs target / residual removed	0.408% / 96.407%

fully discrete comparisons run on byte-identical inputs across independent implementations. Under the project's certification wall—derived $\neq$ numerically reproduced $\neq$ cross-implemented $\neq$ cross-agent reproduced $\neq$ externally independently validated $\neq$ physically certified—the exact results of Sections 4–8 are derived and numerically reproduced; the packet laws of Sections 9–10 are derived and numerically reproduced, with the leading coefficients κ and ρ_0 closing at relative discrepancies 4.71×10^{-9} and 1.41×10^{-8} and the higher-order and reconstruction-law checks at their own tabulated accuracies; and the attribution chain of Section 11 is additionally cross-implemented and cross-agent reproduced. Nothing in this paper claims the last two grades: external validation and physical certification are open by construction, and the finite-radius export quantities used in the barrier testbed are declared proxies, not horizon or null-infinity fluxes.

13 Scope, limitations, and nonclaims

The fences of this paper, stated as claims about what has *not* been shown: (1) semidiscrete $\neq$ fully discrete—Theorems 1–2 are semidiscrete, and the fully discrete solver matches them only after the Section 11 attribution; (2) exact per-mode $\neq$ exact packet-level at all orders—the cubic kernel is exact through two defect orders only; (3) Gaussian-specific coefficients $\neq$ universal spectral laws—5/48, 15/1024, 3/5–2/5 carry their weight and initialization hypotheses; (4) reconstruction singularity $\neq$ physical divergence; (5) standard RK4 facts $\neq$ new contributions—the unit-circle identity and the $T\omega^6\Delta t^5$ scaling are inherited [8, 9, 10]. Our application-specific claim is bounded. In the sources examined, we did not find a prior analysis that closes the complete chain from a boundary-parameter-dependent timestep policy to branch-dependent RK4 dissipation, odd-in-$\mathcal{R}$ reflected-power bias, matched-pair $h^5\sigma^{-6}$ closure, extremum-shift attribution, and a locked timestep-halving falsifier. Additionally: ledger closure is not physical boundary certification; the barrier bins are packet-spectrum-weighted transmission and reflection fractions of the declared packet, not universal barrier constants; and internal cross-agent reproduction is not external validation.

14 Conclusion

A finite window on a wave equation becomes trustworthy accounting only when the energy chooses its flux, the boundary declares its family, and the family is solved rather than approximated. For the trapezoid-weighted second-order semidiscretization, that program completes: the boundary's exact symbol is two lines long, its defect is one secant, its packet consequences are rational coefficients, and the one place the solver disagreed with the theory was traced to its timestep policy and removed by a falsifier written down in advance. The same discipline—declared boundary, conjugate flux, named channels, locked prediction, receipt—is the subject of the companion architecture paper, where the present model serves as the executable lead example.

A Derivations

A.1 Conjugate flux and the centered-flux mismatch. Telescoping $\frac{d}{dt}\frac{1}{2h}\sum(\psi_{j+1}-\psi_j)^2 = \sum_k g_k(v_{k+1}-v_k)$ gives $g_{N-1}v_N - g_0v_0 - h\sum_{j=1}^{N-1} L_j v_j$ with L_j the discrete Laplacian; the interior equations cancel the sum and the half-weighted endpoint equations cancel the residues. At the left face, centered minus one-sided flux equals $-\frac{h}{2}v_j\psi_{xx}+O(h^3)$ by Taylor expansion of $(\psi_{j+1}-\psi_{j-1})/2h-(\psi_j-\psi_{j-1})/h$; the sign reverses at the right face.

A.2 The κ budget. $2\beta_k = 2\tan^2(kh/4) = (kh)^2/8+(kh)^4/192+O((kh)^6)$. Under the weight W: $\langle 2\beta\rangle_0 = (3/16)u[1+(5/48)u]$ from the moments; the weight expansion $W\propto k^2e^{-\sigma^2k^2}[1-\frac{5}{24}(kh)^2]$ (with $\frac{1}{12}$ from $\sin^2$ and $\frac{1}{8}$ from the mover factor) contributes the covariance $-\frac{5}{24}h^2\,\mathrm{Cov}_0(k^2,2\beta) = -\frac{5}{128}u^2$, i.e. $-\frac{10}{48}u$ relative. Net: $\kappa = (3/16)u[1-(5/48)u+O(u^2)]$. Neglected: left-mover denominator content, relative $O(u^2)$ on κ; Gaussian sampling aliasing $O(e^{-(\pi\sigma/h)^2})$; domain tails; time discretization (none, semidiscrete).

A.3 Extremum displacement. With $\delta_{\mathrm{odd}} = -K(\mathcal{R})\,\mathcal{R}(1-\mathcal{R}^2)$, $K(\mathcal{R}) = \langle 2\beta(1-\beta^2)/(1-\beta^2\mathcal{R}^2)^2\rangle_W = M_1+2\mathcal{R}^2M_3+O(\beta^5)$, stationarity at $\mathcal{R}^2=\frac{1}{3}+\delta$ gives $\delta = \frac{8}{27}M_3/M_1$, hence $\Delta\mathcal{R}_* = \frac{4\sqrt{3}}{27}\langle\beta^3\rangle/\langle\beta\rangle = \frac{35\sqrt{3}}{6912}(h/\sigma)^4$ using $\langle k^6\rangle = 105/8\sigma^6$. The exact map preserves the kernel shape through second order, $\delta_{\mathrm{odd}} = -\kappa(1-\kappa/2)\mathcal{R}(1-\mathcal{R}^2)+O(\kappa^3)$, so the extremum is displaced only at the quartic order above.

B RK4 oscillatory defect and the exact repeated-mode model

For $y' = i\omega y$, $x=\omega\Delta t$: $R_4(ix) = \left(1-\frac{x^2}{2}+\frac{x^4}{24}\right)+i\left(x-\frac{x^3}{6}\right)$—the truncated cosine–sine pair—and exactly $|R_4(ix)|^2 = 1-\frac{x^6}{72}+\frac{x^8}{576}$: the truncated pair fails to remain on the unit circle, which is the geometric reading of RK4's order-five dissipation [8, 9]. The ladder: $|R_4| = 1-\frac{x^6}{144}+\frac{x^8}{1152}-\frac{x^{12}}{41472}+\cdots$; $\log|R_4|^2 = -\frac{x^6}{72}+\frac{x^8}{576}-\frac{x^{12}}{10368}+\cdots$; $\arg R_4 = x-\frac{x^5}{120}+\frac{x^7}{336}-\frac{x^9}{5184}+\cdots$, separating amplitude from phase error. Four tiers are kept distinct: the exact one-step identity; the exact repeated-mode recurrence $G_N(k) = |R_4(i\omega\Delta t)|^{2N}$ at the recorded timestep; the packet asymptotic $\log G_N = -\frac{T}{72}\omega^6\Delta t^5+\frac{T}{576}\omega^8\Delta t^7-\cdots$ with $D = (T/72)\langle\omega^6\rangle_E\Delta t^5$; and the full PDE, for which no theorem is claimed. Under the Section 11 policy the branch difference gives $\epsilon_{\mathrm{odd}} \approx -\frac{\varrho}{2}(D_+-D_-)$ and an estimator displacement $\Delta\mathcal{R}_{\mathrm{RK4}} \approx 1.995\times10^{-3}\,h^3/\sigma^4$ at $T=75$. The exact-mode audit of the falsifier run reproduces the realized estimator offset to 0.06% with pointwise curve error 3.6×10^{-14} RMS (3.4% of curve scale); estimator-level closure, not pointwise perfection, is the claim. Unit-modulus methods ($|R(ix)|=1$: implicit midpoint, Gauss–Legendre; Störmer–Verlet preserves a modified energy) remove the amplitude channel only.

Table 2: Claim-to-receipt seed. Status families: T exact theorem; I exact identity under stated closure; A asymptotic corollary; N numerically certified mechanism. Receipt paths refer to the frozen v0.6.11 package.

ID	Claim	Status	Receipt anchor
FO-A-001	conjugate flux identity	I	gates; Sec. 4
FO-A-002	$\mathcal{R} \leftrightarrow \alpha$ family + anchors	I	anchor receipts
FO-A-003	exact symbol $r(k)$	T	`exact_symbol_predictions.csv`; ledger κ row
FO-A-004	$\alpha_{\text{eff}} = \alpha \sec(kh/2)$	T	ibid.
FO-A-005	Möbius representation	I	`rk4_receipt_audit_receipt.json`
FO-A-006	per-mode odd law	T	overlay receipts
FO-A-007	cubic kernel, two-order fence	T	950-pair receipts; Fig. 1
FO-A-008	$1/\sqrt{3}$ extremum	T	ledger R_* row; Fig. 1
FO-A-009	$\Delta\mathcal{R}_*$ formula	A	root-solve receipts
FO-A-010	κ, 5/48	A	ledger κ, c rows
FO-A-011	$\rho_0 = 15/1024\, u^2$	A	ledger ρ_0 row; ladder
FO-A-012	3/5–2/5 split	A	ledger rows (one provisional)
FO-A-013	$1/\lvert\mathcal{R}\rvert$ knot	A	ledger slope row; Figs. 2–3
FO-A-014	RK4 unit-circle identity	I	App. B; [8, 9]
FO-A-015	repeated-mode model	I	`exact_mode_receipt.json`
FO-A-016	branch-dt odd bias	N	`bk/` audit; matched receipts
FO-A-017	matched-pair closure	N	matched finals; prediction ledger
FO-A-018	dt-halving falsifier	N	`falsifier_receipt.json`

C Claim-to-receipt seed table

D Figure provenance

Figures bind to `img/image_manifest.csv` rows (notebook, cell index, SHA-256). Previously bound: falsifier figure = `runtime_figures/...falsifier.pdf` (sha `1890cb0e...`); exact-mode closure = `...exact_mode_closure.pdf` (sha `39b5b0a3...`). The three selected human-facing figures included below are copied byte-for-byte from the frozen notebook image export:

- Fig. 1: notebook cell 97, sha `cdc8b06f1dce6f75ce07c6db52c4dee4 542a0cb25962cc614ce939f48e86922c`.
- Fig. 2: notebook cell 128, sha `6b968de119e3b59f91427c3087c27b1e 71186956cc30e98da8a4ab64ef8c5d7c`.
- Fig. 3: notebook cell 128, sha `6fed5163e1ed10a142e3c0f39a649008 54e2073ec59c474b243cafcb368e01f6`.

E Selected visual evidence

The archive carries the complete image export and receipts. The three plots below are included directly so that a human reader can inspect the central mathematical signatures without traversing nested archives.

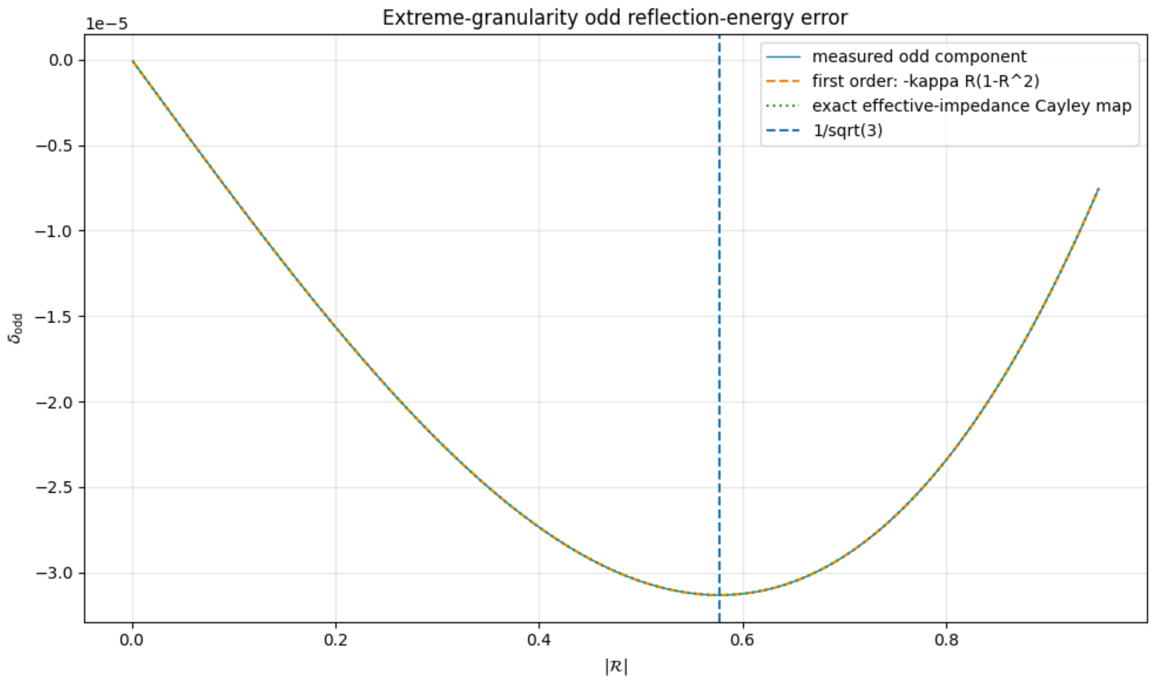

Figure 1: Executed odd reflected-power comparison at extreme granularity. The measured odd component lies on the first-order cubic law $-\kappa\mathcal{R}(1-\mathcal{R}^2)$ and the exact effective-impedance Cayley map at the scale shown; the dashed line marks $|\mathcal{R}| = 1/\sqrt{3}$. This is the visual counterpart of Theorem 2 and claims FO-A-007–FO-A-008.

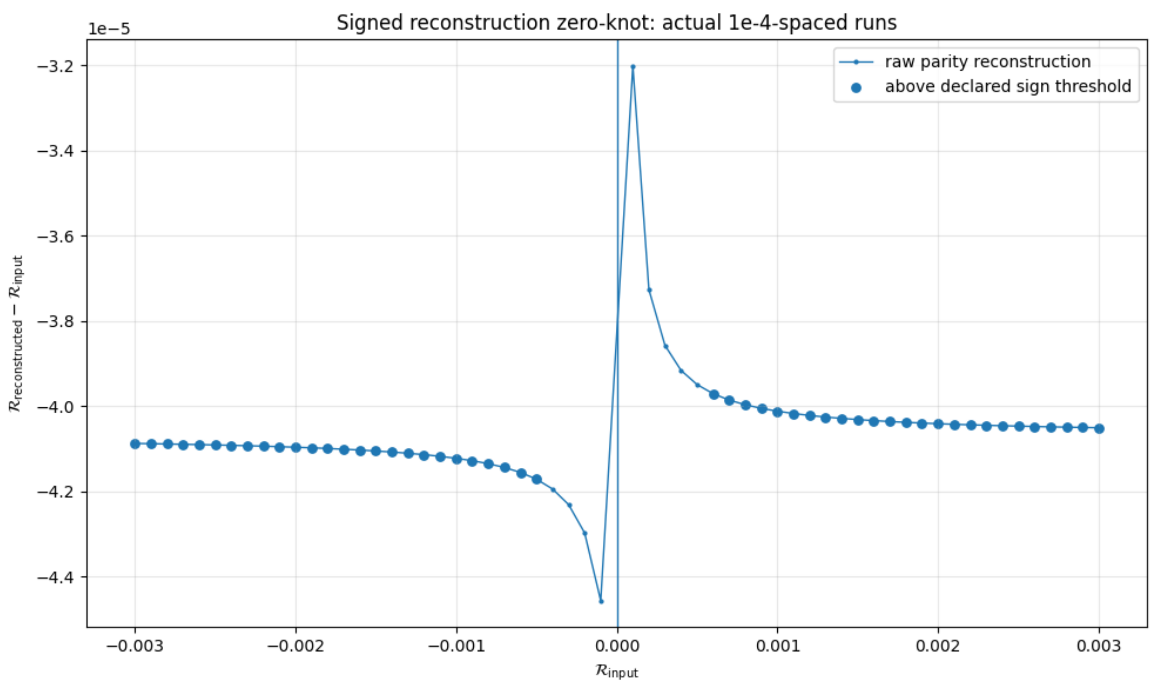

Figure 2: Actual 10^{-4}-spaced signed-reconstruction runs through the zero-knot neighborhood. The estimator's correction becomes singular near the transparent point, while the declared sign-threshold policy determines where a signed reconstruction is reported.

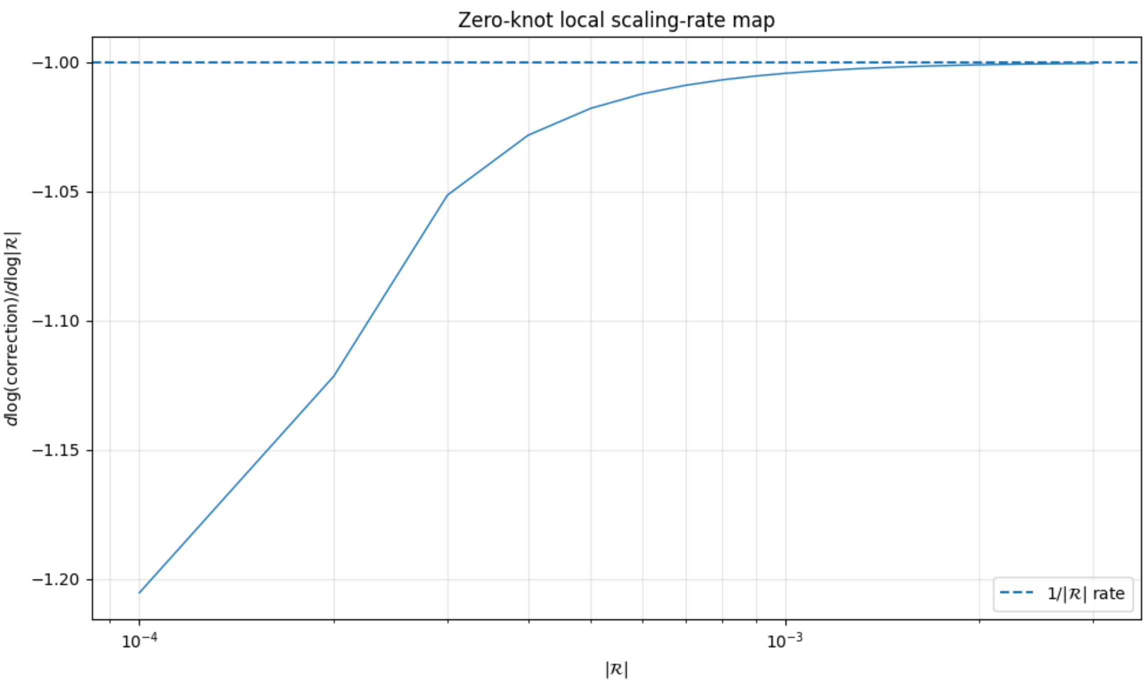

Figure 3: Local logarithmic scaling rate of the zero-knot correction. The rate approaches -1, directly visualizing the $1/|\mathcal{R}|$ reconstruction law of claim FO-A-013; the dashed line is the predicted asymptotic rate.

References

[1] Brittany A. Erickson, Jeremy E. Kozdon, and Tobias W. Harvey. A non-stiff summation-by-parts finite difference method for the scalar wave equation in second order form: Characteristic boundary conditions and nonlinear interfaces. *Journal of Scientific Computing*, 93:17, 2022.

[2] Siyang Wang, Daniel Appelö, and Gunilla Kreiss. An energy-based summation-by-parts finite difference method for the wave equation in second order form. *Journal of Scientific Computing*, 91:52, 2022.

[3] Ken Mattsson and Jan Nordström. Summation by parts operators for finite difference approximations of second derivatives. *Journal of Computational Physics*, 199:503–540, 2004.

[4] Magnus Svärd and Jan Nordström. Review of summation-by-parts schemes for initial-boundary-value problems. *Journal of Computational Physics*, 268:17–38, 2014.

[5] Lloyd N. Trefethen. Group velocity in finite difference schemes. *SIAM Review*, 24(2):113–136, 1982. [publisher-check].

[6] Robert Vichnevetsky and John B. Bowles. *Fourier Analysis of Numerical Approximations of Hyperbolic Equations*. SIAM, Philadelphia, 1982.

[7] Heinz-Otto Kreiss, N. Anders Petersson, and Jacob Yström. Difference approximations for the second order wave equation. *SIAM Journal on Numerical Analysis*, 40:1940–1967, 2002. [publisher-check].

[8] P. J. van der Houwen and B. P. Sommeijer. Explicit Runge–Kutta (–Nyström) methods with reduced phase errors for computing oscillating solutions. *SIAM Journal on Numerical Analysis*, 24(3):595–617, 1987. [publisher-check].

[9] F. Q. Hu, M. Y. Hussaini, and J. L. Manthey. Low-dissipation and low-dispersion Runge–Kutta schemes for computational acoustics. *Journal of Computational Physics*, 124:177–191, 1996. [publisher-check].

[10] Ernst Hairer, Christian Lubich, and Gerhard Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer, 2nd edition, 2006.

[11] Karl Martel and Eric Poisson. Gravitational perturbations of the Schwarzschild spacetime: A practical covariant and gauge-invariant formalism. *Physical Review D*, 71:104003, 2005.