

Protection without teeth: semiclassical chronology protection on the Ori time machine

Jonathan Schemoul

LibertAI, France

js@libertai.io

October 27, 2026

Abstract

Hawking’s chronology protection conjecture rests on divergent vacuum polarization at chronology horizons, established for spacetimes whose horizons are hypersurfaces of closed null geodesics (Misner [5], Grant [6], and wormhole machines) and on the Kay–Radzikowski–Wald theorem for compactly generated horizons. Both are silent for machines such as the Ori time machine [Phys. Rev. Lett. **95**, 021101 (2005)], whose horizon is probably (but not proven) non-compactly generated; Ori notes that this must be verified by an explicit study of the Cauchy horizon. We compute the image-series energy density of a massless conformal scalar on the Ori geometry, and assemble the full image-series stress tensor at the anchor and on the local eikonal chronology-horizon branch, conditional on the Huygens truncation of the Hadamard form suggested by the exact Brinkmann pp-wave structure. Each image is transported by the exact bitensor; the full 4×4 Jacobi system gives a longitudinal block with determinant exactly λ^2 , so the Van Vleck block is not assumed affine. At the anchor and on the local eikonal chronology-horizon branch, the frame-invariant eigenvalues of $T^\mu{}_\nu$ are $(-\frac{3}{2}, +\frac{1}{2}, +\frac{1}{2}, +\frac{1}{2})\mathcal{E}$ up to the $O(a)$ transverse splitting. The transport tangent is spacelike at both sampled points ($\kappa = u^2 > 0$), so the image-direction contraction $T_{\mu\nu}l^\mu l^\nu = -\frac{3}{2}\mathcal{E}$ with $l = u/\sqrt{|\kappa|}$ is a normalized spacelike eigenvalue, not a null contraction; its sign is negative, opposite to the positive energy-density proxy $\mathcal{E}$, and we make no null-energy (NEC) claim from it. On approach to the horizon the stress follows the Misner law $\mathcal{E} \propto \epsilon^{-2}$ ($\min_n |\sigma_n| \approx 0.51 \epsilon$). The surface $T = f$ is the timelike CTC boundary, not the horizon; the local eikonal horizon $T = F$ lies below it at $\epsilon_{\text{CH}} = a^2 \rho^2 + O(a^3)$ and is not proven to be the global Cauchy horizon. At ϵ_{CH} the energy series is within 8–9% of its asymptotic constant (91–93%), and the sharp turnover lies below the CH. The largest full-tensor chart component at the CH is $1.78 \times 10^{-13} \text{ J/m}^3$ for a metre-scale machine at $a = 0.02$, while the energy-density proxy is $5.3 \times 10^{-18} \text{ J/m}^3$; the often-quoted 6.6×10^{-5} is the ratio of the latter to its formal past-horizon on- ∂V extrapolation, not evidence of a cutoff. A nine-point scan over $\rho = 0.10\text{--}0.50$ (one transverse direction, $y = 0$) is a consistency check: ϵ_{CH} is defined as $a^2 \rho^2$, so the fit ρ^{-4} follows from the ϵ^{-2} law, and the CH/ ∂V ratio is a - and ρ -independent. A separate $n_{\text{max}}/\text{no-}\Delta$ transport check gives 0.870; the full bitensor gives 0.856 at the anchor and 0.999 at the CH. The gauge-complete ten-component analysis finds no per-wrap exponential amplifier in any sector, and the TT expulsion is a $k_\perp = 0$ statement only. Two-dimensional backreaction is not exactly zero in general; what is exact there is the horizon shift, which vanishes because 2D gravity is topological and the CFT stress is traceless, so the literal semiclassical equation reduces to a state constraint ($\langle T_{ab} \rangle = 0$), not a metric equation. Four-dimensional backreaction and a closed-form adapted state remain open (the conical anomaly is computed, and the conical piece and quotient Casimir do not cancel automatically). Thus, if the horizon is non-compactly

generated, semiclassical protection operates in form but is quantitatively negligible at the sampled point: the proxy is about sixty orders below the unit-curvature stress, and the full-tensor chart components about 55.4–55.7 orders (the largest, $1.78 \times 10^{-13} \text{ J/m}^3$, is 55.4 orders below). This is a test-field statement; survival of a macroscopic machine requires the uncomputed 4D backreaction. If the horizon is compactly generated, the KRW theorem applies and this finite test-field result is not the relevant statement.

1 Introduction

Hawking’s chronology protection conjecture proposes that “the laws of physics prevent the appearance of closed timelike curves” [4]. The proposed enforcement mechanism is semiclassical: on a chronology horizon, null geodesics winding around a closed null geodesic (the “fountain”) pile up, virtual particles acquire a boost holonomy at each circuit, and the renormalized stress-energy operator $\langle T_{\mu\nu} \rangle$ of quantum fields diverges; fed back through the semiclassical Einstein equations, this is supposed to destroy the horizon before closed timelike curves can form [4]. Explicit computations supported the mechanism: vacuum polarization diverges on the chronology horizons of Taub–NUT space [7], of wormhole time machines [8], and of generic compactly generated horizons [9]. The mechanism was eventually put on a rigorous footing — of a sort — by Kay, Radzikowski and Wald (KRW) [10], who proved, by microlocal analysis, that on a spacetime with a *compactly generated* Cauchy horizon no quantum state can be Hadamard everywhere: the two-point function fails to be Hadamard at points of the horizon itself, and $\langle T_{\mu\nu} \rangle$ is there singular or undefinable.

Every step from these results to “time machines are impossible” crosses a known crack; the two that matter here are:

(i) *The theorem’s scope.* KRW is a statement about distributions on a fixed background, not about back-reaction, and it applies only to *compactly generated* horizons [10, 17]. Non-compactly generated machines — rotating cylinder machines, Krasnikov’s constructions, the Ori–Soen family — are untouched by both the classical (Tipler–Hawking) theorems and KRW. Visser’s review emphasizes that the failure can be localized to points or parts of the horizon, and that the renormalized stress tensor can remain finite up to the horizon in special models; the measure-zero characterization is Li and Gott’s, not Visser’s [17, 13].

(ii) *The state-choice loophole.* For Misner space, self-consistent Hadamard states with $\langle T_{\mu\nu} \rangle = 0$ everywhere off the horizon have been constructed [11, 12, 13], meeting the KRW-type reply that on the horizon itself the operator remains undefinable [16]. Hawking himself found propagators with vanishing stress tensor and dismissed them on positivity grounds: “I am almost sure there is no quantum state on Misner space for which $\langle T^{\mu\nu} \rangle$ is finite on the horizon, but I do not have a rigorous proof” [4].

Against this background, the sharpest open question is the non-compactly generated case: we are not aware of a computation of the renormalized $\langle T_{\mu\nu} \rangle$ on the best available constructive models, the Ori time machines. The 2005 model of Ori [18] is a class of curved, asymptotically flat, topologically trivial *vacuum* solutions that develop closed timelike curves inside an empty toroidal region bounded by a chronology horizon that is probably (but not proven) non-compactly generated; as Ori notes, this must be verified by an explicit study of the Cauchy horizon. The surrounding matter satisfies the weak, dominant and strong energy conditions. Ori’s own verdict: “Stability remains the main open question.” And, more pointedly: “Perhaps we shall have to await the formulation of the full theory of quantum gravity before we know whether quantum instabilities provide chronology protection or not” [18].

Related work. Geometric analyses of the Ori family have concentrated on its causal and curvature structure [21]; here we instead compute the renormalized stress on that structure. Rieger’s

isocausality and deck-equivariance classification of Misner-type spacetimes [22] provides a causal-structure framework for quotient legality rather than a stress computation. On the constructive side, Xavier has exhibited CTCs in a vacuum traveling wave [28], whose train family lacks the winding-compression mechanism; that geometry is the natural next test of whether “protection without teeth” is generic to non-compactly generated cores. Analogue-gravity proposals now aim to test chronology protection directly, with Bermudez and Leonhardt proposing an optical time-travel analogue of Hawking’s conjecture [29]. On the no-go side, Lee, Nilsson and Thakur derive global causality constraints for rotating scalar-tensor spacetimes [30], and Ahmed, Al-Badawi and Sakalli ask whether Petrov type-N and D spacetimes admitting CTCs remain valid in $f(R, L_m, \Phi, X)$ gravity [31]; both extend the classical constraints beyond the Einstein–vacuum setting of the Ori model, the latter directly relevant because the Ori core is type N. Senjaya probes the conjecture through scalar quasinormal modes [32], a complementary scattering-based diagnostic to the local stress computation performed here.

This paper is the computation Ori’s remark invites. We compute the image-series energy density of a massless conformal scalar on the Ori geometry, and the full image-series tensor at the anchor and on the local eikonal chronology horizon: to our knowledge, the first image-series energy-density computation of the conformal scalar on the Ori 2005 core in GR (literature search 2026-10-26). The core is a leading constructive candidate, whose horizon is probably (but not proven) non-compactly generated and must be verified by an explicit study of the Cauchy horizon. The investigation corrected itself three times along the way, and we document all three corrections openly, since they are diagnostics of the method’s honesty: (1) our initial closed-null-geodesic census suggested the divergence was localized to a single closed geodesic, and the image-series stress computation *overturned* that verdict (Sec. 3); (2) an apparent disagreement between our ϵ^{-2} approach law and Hawking’s t^{-4} was resolved from the primary source as a coordinate-convention difference (Sec. 6); and (3) the surface $T = f$ was corrected from “the chronology horizon” to the timelike boundary of the CTC region, with the local eikonal horizon below it (Sec. 5).

Our conclusion is a hybrid neither camp of the chronology-protection debate expected. The full image-series tensor at the anchor and at the local eikonal horizon has frame-invariant eigenvalues $(-\frac{3}{2}, +\frac{1}{2}, +\frac{1}{2}, +\frac{1}{2})\mathcal{E}$ up to the $O(a)$ splitting, so the image-direction contraction is $-\frac{3}{2}\mathcal{E}$: the transport tangent is spacelike at both sampled points, so this is a normalized spacelike eigenvalue, not a null contraction. Its sign is negative, opposite to the positive energy-density proxy $\mathcal{E}$, but it carries no null-energy (NEC) claim. The energy series near the local eikonal horizon follows the Misner law ϵ^{-2} — and never gets the chance to leave it. A locus correction (Sec. 5) shows the surface $T = f$ is not the horizon but the timelike CTC boundary; the local eikonal horizon lies $\epsilon_{\text{CH}} \propto a^2 \rho^2$ below it (its identification with the global Cauchy horizon is not proven), where the energy series is within 8–9% of the asymptotic constant (91–93%) and finite. The energy-density series there is a factor 6.6×10^{-5} of its formal past-horizon on- ∂V extrapolation — a ratio between two surfaces with the same $(a^2 \rho^2)^{-2}$ scaling, not evidence of a cutoff. The largest full-tensor chart component at the CH is $1.78 \times 10^{-13} \text{ J/m}^3$ for a metre-scale machine at $a = 0.02$, and the proxy is $5.3 \times 10^{-18} \text{ J/m}^3$. The only genuine singularity is on the single closed null geodesic $\mathcal{N}$ on ∂V , and by inference on the CH. The CH series is verified at $x_0 = 0.3$ and scanned over $\rho = 0.10$ – 0.50 (nine values, $y = 0$, $a = 0.02$, $L = 1$) as a consistency check: it is finite, at 91–93% of the asymptotic constant ($\epsilon_{\text{CH}}^2 |\text{series}| = 0.2192$ at every ρ ; $\min_n |\sigma_n| = 0.5000 \epsilon_{\text{CH}} = 0.9804$ of the 0.51ϵ law), and scales as $(a^2 \rho^2)^{-2}$ (fitted slope -4.000). Since ϵ_{CH} is defined as $a^2 \rho^2$, the ρ^{-4} fit follows from the ϵ^{-2} law by construction; the caveats are one transverse direction, the local $O(a^2)$ eikonal identification, and the test-field series. The standard Hawking stress mechanism is quantitatively negligible at the sampled point: the energy-density proxy is about sixty orders below the unit-curvature stress at $a = 0.02$, and the full-tensor chart components about 55.4–55.7 orders (the

largest, $1.78 \times 10^{-13} \text{ J/m}^3$, is 55.4 orders below). This is protection in form, not in force — for the test field, and *protection without teeth* at the local eikonal horizon, without a divergence away from $\mathcal{N}$. Survival of a macroscopic machine requires the uncomputed 4D backreaction. This conclusion is conditional on non-compact generation: if the horizon is compactly generated, the KRW theorem applies and this finite test-field result is not the relevant statement.

The plan of the paper follows. Sec. 2 introduces the metric and its chronology horizon. Sec. 3 presents the closed-null-geodesic census and its supersession. Sec. 4 describes the image-series method; Sec. 5 the results; Sec. 6 the parallel-transport check and the Hawking exponent reconciliation; Sec. 7 the classical stability analysis of $\mathcal{N}$; Sec. 8 the discussion; Sec. 9 the limitations; Sec. 10 the conclusions. An appendix collects numerical methods and the reproducibility statement.

2 The Ori metric and its chronology horizon

We work with the metric of Ref. [18] (their Eq. (1)),

$$ds^2 = dx^2 + dy^2 - 2dzdT + [f(x, y) - T]dz^2, \quad f(x, y) = \frac{a}{2}(x^2 - y^2), \quad (1)$$

with z cyclic, $z \sim z + L$; we sample $a = 0.02$, $L = 1$ throughout. The determinant is $\det g = -1$ exactly (a useful numerical gate), and the metric is Ricci-flat, a fact the method of Sec. 4 exploits. The surface

$$T = f(x, y), \quad (2)$$

where the metric coefficient $g_{zz} = f - T$ vanishes and the identification vector ∂_z becomes null, is the boundary ∂V of the causality-violating region: crossing it, the circles of constant (x, y, T) change from spacelike to timelike and closed timelike curves appear. For $a \neq 0$ this surface is *timelike* (its normal has norm $a^2\rho^2$ there), like the CTC boundary of the Tipler–Van-Stockum family — and unlike a chronology horizon, which is null wherever smooth. The chronology horizon — more precisely, the local eikonal branch through $\mathcal{N}$ that we use in this paper, since a global Cauchy-horizon identification is not proven — is the null eikonal surface

$$T = F(x, y), \quad |\nabla F|^2 = f - F, \quad F = f - a^2\rho^2 + 4a^3(x^2 - y^2) + O(a^4), \quad (3)$$

displaced *below* ∂V by $\delta_{\text{CH}} = a^2\rho^2 + O(a^3)$; the two surfaces coincide only in the degenerate Misner case $a = 0$. The eikonal solution is established at quadratic order for z -invariant leaves and is local; far from the core the global Cauchy horizon may have other leaves, so Eq. (3) is a local branch, not a global theorem. (A correction we owe to the locus analysis of Sec. 5; an earlier draft called $T = f$ “the chronology horizon”.) Inside the horizon region the geometry is vacuum; in Ori’s full solutions the region exterior to the core is supported by matter obeying the weak, dominant and strong energy conditions. Ori suggests that the horizon generators may emanate from the non-compact matter region and that the horizon is probably non-compactly generated, a scenario he says must be verified by an explicit study of the Cauchy horizon; if confirmed, it places the model beyond the reach of the Tipler–Hawking classical theorems and the KRW quantum theorem [18, 10].

Two exact structures of the geodesics of Eq. (1) anchor everything that follows.

(a) *Affine incompleteness.* The z -geodesic equation closes on itself,

$$\ddot{z} = \frac{1}{2}\dot{z}^2 \implies \dot{z}(\lambda) = \frac{\dot{z}_0}{1 - \dot{z}_0\lambda/2}, \quad z(\lambda) = -2\ln(1 - \tfrac{1}{2}\dot{z}_0\lambda), \quad (4)$$

so every wrapping null geodesic is future-incomplete, with affine blow-up at $\lambda_* = 2/\dot{z}_0$: an wrapping trajectory completes infinitely many circuits ($z \rightarrow \infty$) within finite affine parameter. This

rediscovers, independently and exactly, Ori’s own remark that the closed null geodesic at the core is future-incomplete [18]. Numerically, starting with $\dot{z}_0 = 1$ the z -velocity after one wrap is $1.6487 = e^{1/2}$, and the z -circle through the origin ($x = y = T = 0$) closes to machine zero.

(b) *Transverse dynamics.* In z -parameterization (primes denoting d/dz) the transverse equation is exact and universal, independent of whether the geodesic is timelike or null (Sec. 7):

$$x'' + \frac{1}{2}x' = \frac{a}{2}x. \quad (5)$$

The closed null geodesic $\mathcal{N}$ sits at $x = y = T = 0$; it is the unique locus where the z -circle closes exactly.

3 The closed-null-geodesic census (and its supersession)

Why count closed null geodesics? In the image-sum representation of the two-point function (Sec. 4) each image term is singular when a connecting geodesic becomes null, $\sigma_n \rightarrow 0$. Exact closed null geodesics — points where a geodesic returns to its own light cone — are therefore the natural first diagnostic for where the semiclassical stress tensor diverges. This is also the geometric core of the KRW microlocal statement: for compactly generated horizons the divergence lives on the whole horizon because the closed null geodesics fill a hypersurface [10].

We performed the census as follows. From each sample point on the horizon $T = f(x, y)$ we integrate null geodesics in the z -circle direction (and, for the direction search, over the full three-parameter family of null directions), and measure the one-wrap return distance to the initial point. Validation gates: $\det g = -1$ exactly; the null condition conserved to 4.6×10^{-12} after three wraps; $\mathcal{N}$ closure to machine zero; and the Misner limit $f = 0$, which must reproduce the known result. It does: for $f = 0$ (four-dimensional Misner) the return distance is identically zero *everywhere* on the horizon — the entire horizon is a fountain, which is the geometric reason KRW divergence fills it. For the Ori case $f \neq 0$ the census gives *exactly one* closed null geodesic, $\mathcal{N}$: the return distance vanishes only at $x = y = 0$; the z -circle census distance grows as ρ away from it, while the direction-minimized return distance δ_1 grows as ρ^2 (here ρ is transverse distance from $\mathcal{N}$). A direction search over all null directions, $n = 1$ to 8 wraps at $\rho = 0.15$, shows the minimum return distance δ_n *growing* with n ($7.3 \times 10^{-8} \rightarrow 5.9 \times 10^{-6}$ for $n = 1 \rightarrow 8$), stable across integrator tolerances 10^{-9} – 10^{-13} (at $n = 1$: 7.28×10^{-8} at every tolerance), so the growth is a real signal, not noise. Moreover the per-wrap transverse dynamics (5) is contracting, factor $\sim e^{-L/4} \cosh \mu L \approx 0.8$: $\mathcal{N}$ is an attracting closed null geodesic, and images of any point are bounded away from it. The census verdict was that the semiclassical divergence collapses from a hypersurface to the set $\{\mathcal{N}\}$, suggesting $\langle T_{\mu\nu} \rangle$ regular on the horizon away from $\mathcal{N}$ (Fig. 1).

Correction during the investigation. The census verdict was **superseded** by the image-series stress computation of Secs. 4–5, and we document the supersession openly because it identifies the correct diagnostic. The stress tensor does not require exact closure: it requires $\sigma_n \rightarrow 0$, and that occurs wherever the identification vector ∂_z becomes null — which is the *entire* horizon $T = f(x, y)$, not only $\mathcal{N}$. Exact closed null geodesics are the locus where the divergence survives *on* the horizon; they say nothing about the approach to it. The image-series computation below bears this out: the stress grows on approach to the Ori horizon exactly as in Misner space. The census machinery, validation gates, and the $f = 0$ cross-check all remain useful; the inference drawn from them did not survive contact with the actual operator.




Figure 1: The census. Left: one-wrap return distance on the Ori horizon $T = f(x, y)$: it vanishes only at the single closed null geodesic $\mathcal{N}$ (star). Middle: the contrast with the four-dimensional Misner limit $f = 0$, where the return distance vanishes identically. Right: minimum return distance δ_n over all null directions grows with wrap number n — no image pile-up off $\mathcal{N}$.

4 Method: the image series on a plane-wave quotient

For a massless conformally coupled scalar, the Hadamard parametrix of the two-point function on a generic spacetime is $G(x, x') = \Delta^{1/2}/4\pi^2\sigma + (\ln \sigma)$ terms + $\dots$. **Which terms survive is a question about the Riemann, not the Ricci, tensor:** Ricci-flatness alone does *not* truncate the recursion — Schwarzschild is Ricci-flat and its scalar Green function has an inside-the-light-cone tail. What does truncate it on the Ori background is the plane-fronted-wave structure: in the coordinates $(u, v, x, y) = (T, z, x, y)$ the metric is exactly of Brinkmann form

$$ds^2 = -2 du dv + [f(x, y) - u] dv^2 + dx^2 + dy^2, \quad (6)$$

a plane-fronted wave with parallel rays, of Weyl type N, with *all* polynomial curvature invariants vanishing (we verified $R_{\mu\nu} = 0$, $R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma} = 0$, $\Psi_{1,2,3} = 0$ symbolically). The relevant classification is Carminati and McLenaghan's: for the conformally invariant scalar wave equation on a Petrov type-N background, Huygens' principle holds if and only if the spacetime is conformally related to a plane wave (the type-N Huygens class), not for all type-N metrics; Künzle treats Maxwell fields and Harte's plane-wave analysis concerns Maxwell and gravitational perturbations [2, 1, 3]. Our symbolic checks verify type N and the vanishing invariants; they do not establish conformal equivalence of Eq. (6) to a plane wave, and therefore do not by themselves prove Huygens. The exact Brinkmann form is the local structure we rely on; we use the Huygens truncation as a working hypothesis, supported by that form and by the vanishing invariants, not as a theorem applied verbatim. It is a local statement. Under this hypothesis the scalar propagator is tail-free — and since the background is vacuum ($R = 0$), conformal and minimal coupling coincide — so the two-point function is an image sum over the quotient $z \rightarrow z + nL$,

$$G(x, x') = \sum_{n=-\infty}^{\infty} \frac{\Delta_n^{1/2}}{4\pi^2\sigma_n} + (\text{Hadamard subtraction}), \quad (7)$$

where for the n -th image σ_n is the Synge world function along the geodesic connecting x to $\Lambda^n x$ (the n -th identified point), and Δ_n the Van Vleck determinant. Point-splitting — differentiate, subtract the trivial $n = 0$ geodesic, and take $x' \rightarrow x$ — turns each image into a stress-tensor contribution.

Transporting the flat Hadamard form by the exact bitensor gives, per image,

$$T_{\mu\nu}^{(n)} = \frac{\Delta_n^{1/2}}{4\pi^2\sigma_n^2} \left(g_{\mu\nu} - \frac{4u_\mu u_\nu}{\kappa_n} \right), \quad \kappa_n = u^2, \quad \sigma_n = \frac{1}{2}\kappa_n\lambda^2, \quad (8)$$

with u the tangent of the connecting geodesic. Its eigenvalues are $(-\frac{3}{2}, +\frac{1}{2}, +\frac{1}{2}, +\frac{1}{2})\mathcal{E}_n$ with $\mathcal{E}_n = \Delta_n^{1/2}/(2\pi^2\sigma_n^2)$. The unit vector $l = u/\sqrt{|\kappa_n|}$ is spacelike at both evaluated points ($\kappa_n > 0$), so $T_{\mu\nu}l^\mu l^\nu = -\frac{3}{2}\mathcal{E}_n$ is a normalized spacelike eigenvalue along the image direction, not a null contraction; the leading-order trace vanishes identically. The script also reports a genuine null-tetrad component for the null pair $l = \partial_T$, $n = \partial_z + \frac{1}{2}(f-T)\partial_T$ (with $l \cdot n = -1$): $T_{ll} = -5.63 \times 10^{12}$ series units at the CH, equal to the chart component T_{TT} . That component is chart-enhanced ($\propto \epsilon^{-3}$), is not an eigenvalue of $T^\mu{}_\nu$, and is not aligned with the image tangent, so we do not identify it with the image-direction quantity. The energy-density series used for the earlier scan and for the nine-point ρ scan is

$$\mathcal{E}(x) = \frac{1}{2\pi^2} \sum_{n \neq 0} \frac{\Delta_n^{1/2} \cos \theta_n}{\sigma_n^2}, \quad (9)$$

where θ_n is the start-point approximation to the transported tangent angle used in those runs; its normalization differs from the exact tensor (8) by an $O(1)$ transport factor discussed in Sec. 6. The full tensor is evaluated at the anchor and on the local eikonal chronology horizon in Sec. 5. The divergence structure is controlled by (i) the light-cone relation $\sigma_n \rightarrow 0$ and (ii) the large- n convergence of the series. We note the scope of the Huygens hypothesis honestly: it fixes the *leading* light-cone structure that Eq. (9) keeps; should a subleading tail term exist on the quotient (e.g. from the identification or from any departure of the exact quotient geometry from the plane-wave structure), it would be smooth inside the light cone and would contribute regular, subleading image terms — it cannot change the $\sigma_n \rightarrow 0$ singularity structure that our conclusions rest on, but the coefficients would acquire corrections. The series is thus exact in its leading, divergence-relevant content, and approximate at subleading order.

The geodesic boundary-value problem $x \rightarrow \Lambda^n x$ (the connecting geodesic must end exactly at the identified point) is solved by multi-start root finding over the full three-parameter initial-direction family (v_x, v_y, v_T) at fixed $v_z = \pm 1$: shoot, integrate the geodesic equations by an explicit Runge–Kutta–Dormand–Prince scheme with event detection at $z = nL$, and iterate on the initial direction until the endpoint residual vanishes (tolerance 10^{-5}). Along each connecting geodesic we evaluate $\kappa = g_{\mu\nu}u^\mu u^\nu|_0$, $\sigma = \frac{1}{2}\kappa\lambda^2$, and the exact bitensor transport of Eq. (8). The Van Vleck determinant comes from the full 4×4 Jacobi system: the longitudinal 2×2 block is not affine individually ($J_{Tz} \neq 0$), but its determinant is exactly λ^2 , so that $\Delta_n = \lambda^2/|J_x J_y|$ is exact (we verify $\det J/(J_x J_y \lambda^2) = 1 + 1.2 \times 10^{-10}$). The n -th term uses the connecting geodesic of smallest $|\sigma_n|$ among those found. All numbers below are reproducible from the scripts in the repository (Appendix A).

5 Results

5.1 Validation: the Misner benchmark

Setting $f = 0$ recovers four-dimensional Misner space, for which the world function of the n -th image is available analytically from the Rindler dictionary $U = -4Te^{z/2}$, $V = e^{-z/2}$ (the identification




Figure 2: The image-series energy density. Left: validation on four-dimensional Misner ($f = 0$), where the series realizes the $|T_0|^{-2}$ law with the analytic world function $\sigma_n = 4|T_0|(\cosh \frac{nL}{2} - 1)$. Center-left: the Ori horizon approach off $\mathcal{N}$: the same ϵ^{-2} law, $\min |\sigma_n| = 0.51 \epsilon$. Center-right: convergence — $\sigma_n \propto n^4$, Δ_n bounded. Right: on the horizon, the series blows up only approaching $\mathcal{N}$, as ρ^{-4} .

$z \rightarrow z + L$ is a boost of rapidity $L/2$):

$$\sigma_n = 4|T_0| \left(\cosh \frac{nL}{2} - 1 \right), \quad (10)$$

with T_0 the (negative) time coordinate of the field point. At $n = 1$, $L = 1$ this gives $\sigma_1 = 0.5105 |T_0|$, and the measured value agrees to six decimals ($\sigma_1/|T_0| = 0.510504$ at $|T_0| = 0.5$). The series diverges as $|T_0|^{-2}$: measured power-law slope -2.00 . An earlier draft quoted the incorrect form $|T_0|(\cosh nL - 1)$ and a “6% discrepancy” against it; both were wrong (the rapidity is $L/2$, not L), and the corrected Eq. (10) removes the discrepancy entirely — the machinery is exact against the benchmark, not approximate.

5.2 Approach to ∂V : cut off at the local eikonal horizon, not turned over

Off $\mathcal{N}$, approaching the surface $T = f(x, y)$ from below at fixed transverse position (here $x_0 = 0.3$), the dominant image obeys, to numerical precision,

$$\min_n |\sigma_n| \approx 0.51 \epsilon, \quad \epsilon \equiv f(x_0, 0) - T, \quad (11)$$

with the coefficient $0.51 (= 4(\cosh \frac{L}{2} - 1))$, the exact Misner value) constant to 0.1% over $\epsilon = 0.5 \rightarrow 10^{-3}$ and within 1% down to 10^{-4} : the approach is locally Misner-like, and the energy density follows $\mathcal{E} \propto \epsilon^{-2}$ in this range. **A locus correction now fixes where this law is physically meaningful.** As established in Sec. 2 (Eq. 3), $T = f$ is not the chronology horizon for $a \neq 0$ but the timelike CTC boundary ∂V ; the local eikonal horizon — the null surface $T = F$, not proven to be the global Cauchy horizon — lies below it at

$$\epsilon_{\text{CH}} = f(x_0, 0) - F(x_0, 0) = a^2 \rho^2 + O(a^3) = 3.3\text{--}3.6 \times 10^{-5} \quad (12)$$

at our sampling point ($a = 0.02$, $x_0 = 0.3$). Re-evaluating the images there: $\min_n |\sigma_n| = 1.80 \times 10^{-5} = 0.98 \times (0.51 \epsilon_{\text{CH}})$, so the dominant image is still on the 0.51ϵ law. The energy series at ϵ_{CH} evaluates to 1.69×10^8 in series units ($\epsilon^2 |\text{series}| = 0.2192$), which is 91–93% of the asymptotic constant 0.236–0.242: **the stress at the leading-order CH endpoint is within 8–9% of the asymptotic constant**, and the sharp turnover lies below the CH. **The horizon, not a curvature**

ρ	ϵ_{CH}	$\min_n \sigma_n $	$ \text{series} _{\text{CH}}$	$\epsilon_{\text{CH}}^2 \text{series} $	$\text{CH}/\partial V$
0.10	4.000×10^{-6}	2.000×10^{-6}	1.3698×10^{10}	0.2192	6.589×10^{-5}
0.15	9.000×10^{-6}	4.500×10^{-6}	2.7058×10^9	0.2192	6.589×10^{-5}
0.20	1.600×10^{-5}	8.000×10^{-6}	8.5614×10^8	0.2192	6.589×10^{-5}
0.25	2.500×10^{-5}	1.250×10^{-5}	3.5067×10^8	0.2192	6.589×10^{-5}
0.30	3.600×10^{-5}	1.800×10^{-5}	1.6911×10^8	0.2192	6.589×10^{-5}
0.35	4.900×10^{-5}	2.450×10^{-5}	9.1283×10^7	0.2192	6.589×10^{-5}
0.40	6.400×10^{-5}	3.200×10^{-5}	5.3509×10^7	0.2192	6.589×10^{-5}
0.45	8.100×10^{-5}	4.050×10^{-5}	3.3405×10^7	0.2192	6.589×10^{-5}
0.50	1.000×10^{-4}	5.000×10^{-5}	2.1917×10^7	0.2192	6.589×10^{-5}

Table 1: Transverse scan of the CH image series on $T = F$ ($y = 0$, $a = 0.02$, $L = 1$, $n_{\text{max}} = 5$; series units L^{-4}). The anchor row is $\rho = 0.30$. The $O(a^3)$ -refined CH gives $\epsilon_{\text{CH}}^2 |\text{series}| = 0.2164$ (90.6% of the 0.239 midpoint asymptote) and $\text{CH}/\partial V = 7.688 \times 10^{-5}$ at every scanned ρ .

turnover, ends the rising regime. The on- ∂V value is 2.57×10^{12} : the CH energy-density series is a factor 6.6×10^{-5} of that formal past-horizon on- ∂V extrapolation at the leading-order endpoint (7.7×10^{-5} with the $O(a^3)$ correction). This ratio is not evidence of a cutoff: both surfaces scale as $(a^2 \rho^2)^{-2}$, so the ratio is a - and ρ -independent.

A nine-point scan over $\rho = 0.10$ – 0.50 ($y = 0$, $a = 0.02$, $L = 1$, $n_{\text{max}} = 5$) is a consistency check, not an independent confirmation. At every ρ , $\epsilon_{\text{CH}}^2 |\text{series}| = 0.2192$ (91–93% of the asymptotic constant), $\min_n |\sigma_n| = 0.5000 \epsilon_{\text{CH}}$ (0.9804 of the 0.51ϵ law), and the CH series fits ρ^{-4} with slope -4.000 . The ρ^{-4} fit follows by construction, because ϵ_{CH} is defined as $a^2 \rho^2$; what the scan checks is the $(a^2 \rho^2)^{-2}$ law over a decade in ρ . The $\text{CH}/\partial V$ ratio is constant at 6.589×10^{-5} (7.688×10^{-5} for the $O(a^3)$ -refined CH) and is a - and ρ -independent (Table 1). The provenance of the quoted asymptotic band is the plateau of $\epsilon^2 |\text{series}|$ in the intermediate range $\epsilon = 10^{-3}$ – 10^{-1} , recomputed here as 0.2398 at $n_{\text{max}} = 3$ and 0.2421 at $n_{\text{max}} = 5$; we use the conservative envelope 0.236–0.242. The caveats are that the scan is one-dimensional ($y = 0$) and local (the $O(a^2)$ eikonal CH identification), and that the series remains a test-field evaluation at and beyond the local eikonal horizon.

What lies between the horizon and ∂V (the “pre-CTC shell”, $0 < \epsilon < \epsilon_{\text{CH}}$, of width $a^2 \rho^2$)? Sampling into it is formally possible — $\min |\sigma_n|$ departs from 0.51ϵ below $\epsilon \sim 10^{-5}$, passes through a shallow minimum (1.3×10^{-7} at $\epsilon = 10^{-6}$) and settles onto 3.78×10^{-7} as $\epsilon \rightarrow 0$ — but the region is already past the local eikonal horizon: no causal curve determination from the pre-chronology data reaches it, and the image ansatz there no longer represents the response to any state. The turnover point $\epsilon_c \simeq 7 \times 10^{-7}$ and the “20 $\times$ peak” measured in an earlier draft live in this shell and carry no physical on-horizon interpretation. (That draft called $T = f$ the horizon and read the turnover as the curvature capping the divergence; both readings are corrected here.) The physically meaningful statement is the one above: **the energy-density series at the local eikonal horizon is finite, within 8–9% of the asymptotic constant, and 6.6×10^{-5} of the formal on- ∂V extrapolation** — the proxy is $\sim 5.3 \times 10^{-18} \text{ J/m}^3$ at $a = 0.02$ for a metre-scale machine (the full-tensor chart component is treated below). Hawking’s mechanism, which on Misner space produces a genuine divergence, here is interrupted by the horizon while still in its intermediate-asymptote regime.

5.3 The full image-series tensor

At the anchor ($x_0 = 0.3$, $\epsilon = f - T = 0.2$) and on the leading-order CH ($\epsilon_{\text{CH}} = 3.6 \times 10^{-5}$) we evaluate the exact bitensor expression (8) with the full 4×4 Jacobi block. The frame-invariant content is the same up to the $O(a)$ splitting at both points: the eigenvalues of $T^\mu{}_\nu$ are $(-1.5145, 0.5000, 0.5000, 0.5145)\mathcal{E}$ at the anchor and $(-1.5004, 0.4999, 0.5000, 0.5005)\mathcal{E}$ at the CH, i.e. $(-\frac{3}{2}, +\frac{1}{2}, +\frac{1}{2}, +\frac{1}{2})\mathcal{E}$ up to the $O(a)$ transverse splitting. The image-direction (unit-spacelike) contraction is $-1.500000\mathcal{E}$ at both points; because $\kappa_n > 0$, this is a normalized spacelike eigenvalue, not a null contraction. The trace is zero to relative 7.3×10^{-17} (anchor) and 1.1×10^{-20} (CH); on this VSI pp-wave the local trace anomaly vanishes identically, so the numerical zero matches the exact anomaly.

The energy-density proxy is positive, $\mathcal{E} = 5.34 \times 10^{-18} \text{ J/m}^3$ at the CH for $L = 1 \text{ m}$, but the image-direction eigenvalue is $-3\mathcal{E}/2 = -8.0 \times 10^{-18} \text{ J/m}^3$ (a spacelike direction): the full tensor flips the sign and multiplies the magnitude by $3/2$. A genuine null-tetrad contraction is reported by the script for the null pair $l = \partial_T$, $n = \partial_z + \frac{1}{2}(f - T)\partial_T$: $T_{ll} = -5.63 \times 10^{12}$ series units at the CH, equal to the chart component below. The (z, T) -chart components are larger: $T_{TT} = -5.63 \times 10^{12}$ and $T_{uu} = -2.82 \times 10^{12}$ in series units at the CH, i.e. 1.78×10^{-13} and $8.9 \times 10^{-14} \text{ J/m}^3$. These carry an extra $1/\sigma$ because the chart degenerates at the horizon and the stress is null-dominated; they scale as ϵ^{-3} while the eigenvalues and the image-direction eigenvalue scale as ϵ^{-2} . This is a chart effect, not a second physical divergence. The $O(a^3)$ -refined CH gives $\mathcal{E} = 2.01 \times 10^8$ ($6.3 \times 10^{-18} \text{ J/m}^3$) and $T_{TT} = -7.28 \times 10^{12}$ ($2.3 \times 10^{-13} \text{ J/m}^3$), within the $O(a^3)$ uncertainty of the CH locus. The full-tensor and energy-density normalizations differ by 0.856 at the anchor and 0.999 at the CH; that is the transport factor discussed in Sec. 6.

5.4 On the CTC boundary: the curvature cap (formal, past-the-horizon values)

Everything in this subsection is computed on ∂V itself — at $\epsilon = 0$, i.e. *past* the local eikonal horizon $\epsilon = \epsilon_{\text{CH}}$, inside the pre-CTC shell where the image ansatz no longer represents the response to any state. These are formal properties of the image series (they bound the shell's interior structure), not physical stress measurements; the physical horizon values are those of the preceding subsection. At fixed ρ on ∂V itself the situation changes completely. In Misner space the on-surface σ_n vanishes identically (the census: the whole surface is a fountain) and the series is infinite everywhere on it. On the Ori ∂V the $f \neq 0$ curvature keeps every image off the light cone:

$$\sigma_1 = 4.2 \times 10^{-6} \rho^2 \quad (\text{exact } \rho^2 \text{ scaling, } a = 0.02), \quad (13)$$

measured at $\rho = 0.3, 0.15, 0.08, 0.04, 0.02$ (e.g. 3.78×10^{-7} at $\rho = 0.3$, 1.68×10^{-9} at $\rho = 0.02$). The series is therefore *finite* on ∂V away from $\mathcal{N}$, with a capped value scaling as $(a^2 \rho^2)^{-2} \propto \rho^{-4}$; approaching $\mathcal{N}$ along the surface we measure the slope -4.00 (Fig. 2, right) — again a formal, on- ∂V value. The barrier survives only as a genuine singularity on the single closed null geodesic $\mathcal{N}$.

5.5 Convergence of the image series

At a fixed off- $\mathcal{N}$ horizon point the dominant $|\sigma_n|$ grows with n approximately as n^4 ($3.78 \times 10^{-7} \rightarrow 6.40 \times 10^{-4}$ for $n = 1 \rightarrow 6$, local log-log slopes 4.0–4.5), so the series (9) converges at least as fast as a $\zeta(8)$ -type tail $\sum_n n^{-8}$. The Van Vleck determinants stay bounded, $\Delta_n = 1.000\text{--}1.001$ over the measured range: there is no focusing blow-up, and partial sums converge ($21.9 \rightarrow 23.6 \rightarrow 24.0$ for $n_{\text{max}} = 1, 2, 3$ at the sampled point; the remaining tail is dominated by the $\zeta(8)$ estimate).




Figure 3: Rigor checks. Left: the Rindler dictionary ($U = -4Te^{z/2}$, $V = e^{-z/2}$) is exact, giving $r^2 = |T_0|$. Center: replotted against r , the series realizes Hawking's t^{-4} law with slope -4.00 . Right: convergence in $n_{\max}$.

5.6 Physical magnitude: test-field values

Restoring units, with L the periodicity scale and $\tilde{\rho}$ the transverse distance from $\mathcal{N}$ in units of L , the energy-density proxy at the local eikonal horizon ($\epsilon = \epsilon_{\text{CH}} = a^2 \rho^2$) scales as

$$|\mathcal{E}|_{\text{CH}} \sim \frac{\hbar c}{L^4} \frac{1}{a^4 \tilde{\rho}^4}, \quad (14)$$

and the formal past-horizon on- ∂V series scales as $(\hbar c/L^4)(a^4 \tilde{\rho}^4)^{-1}$, the same law; their ratio is therefore a - and ρ -independent. At the sampled point $\tilde{\rho} = 0.3$ the proxy evaluates to $1.69 \times 10^8 L^{-4}$ at the local eikonal horizon (and 2.6×10^{12} on ∂V), i.e. $\mathcal{E} \approx 5.3 \times 10^{-18} \text{ J/m}^3$ for a machine with $L = 1 \text{ m}$ and $a = 0.02$ (the canonical design $a = 10^{-4}$ gives $\sim 8 \times 10^{-9} \text{ J/m}^3$ from the a^{-4} scaling). The full tensor gives an image-direction eigenvalue $-3\mathcal{E}/2 \approx -8.0 \times 10^{-18} \text{ J/m}^3$ (a spacelike direction) and a largest chart component $1.78 \times 10^{-13} \text{ J/m}^3$ at the same point. The density that curves a metre-scale region by order unity is $c^4/8\pi GL^2 \approx 4.8 \times 10^{42} \text{ J/m}^3$. The proxy is about sixty orders below it at $a = 0.02$, and the full-tensor chart components about 55.4–55.7 orders below it (the largest is 55.4 orders below). These are test-field magnitudes at the sampled point: they show that the stress is negligible there, not that the machine survives. Survival requires the uncomputed 4D backreaction; the ρ^{-4} growth is cut off only at scales where the semiclassical approximation has itself broken down.

6 Rigor checks: parallel transport, and reconciling the exponent with Hawking

6.1 Transport checks

The endpoint-tangent angle θ_n was initially computed with both tangents evaluated at the start point, an $O(\text{curvature} \times \text{length}^2)$ approximation. A separate energy-density transport check (`a2_rigor.py`; fixed $n_{\max} = 3$, no Δ factor) replaces this by a parallel-transported tangent and shifts the series coefficients by a *stable* factor 0.870 at both sampled ϵ ($\epsilon = 0.2$ and 0.05). That check is not the full bitensor: it keeps the energy-density normalization and the old $n_{\max}$, and omits Δ . The exact bitensor transport of Eq. (8), with the full 4×4 Jacobi block, gives instead a ratio $\mathcal{E}_{\text{exact}}/\mathcal{E}_{\text{approx}}$ of 0.856 at the anchor ($\epsilon = 0.2$) and 0.999 at the CH. The difference is the

$O(1)$ transport/Van Vleck normalization, not a change of the power law; the full-tensor numbers of Sec. 5.3 supersede the 0.870 check for the anchor and the CH.

6.2 The exponent convention, resolved from the primary source

Our series diverges as $|T_0|^{-2}$ with $\sigma_n = 4|T_0|(\cosh \frac{nL}{2} - 1)$; Hawking’s Misner computation gives $K = Bt^{-4}$, where t is a distance-type function to the closed null geodesic, and an observer crossing the horizon measures $\mathcal{E} \sim Bd/s^4$ [4]. Apparent disagreement: exponent -2 or -4 ?

Resolved from the primary source. Two-dimensional Misner space is exactly Rindler under the dictionary

$$U = -4T e^{z/2}, \quad V = e^{-z/2}, \quad -dU dV = -2 dz dT - T dz^2, \quad (15)$$

verified numerically to a maximum relative residual of 2×10^{-15} over 200 random samples. Hence $UV = -4T$ exactly, and the distance scale to the closed null geodesic is $r \equiv \frac{1}{2}\sqrt{|UV|} = \sqrt{|T_0|}$, i.e. Hawking’s t is our r with $t^2 = |T_0|$. Our $|T_0|^{-2}$ is his t^{-4} . The content of this reconciliation is the *dictionary* — the exact identification of his distance function with our coordinate distance, established by Eq. (15). Once $r \equiv \sqrt{|T_0|}$ is fixed, the slope of the series against r follows algebraically from the $|T_0|^{-2}$ law; we plot it (measured -4.00 , Fig. 3) as an illustration of the identification, not as an independent numerical confirmation. The two computations agree exactly; the exponents differed only by coordinate convention.

Correction during the investigation. We record openly that our first write-up treated the ϵ^{-2} series law and Hawking’s t^{-4} as an unresolved discrepancy. It is not a discrepancy: $|T_0| = t^2$, and the reconciliation above — a convention difference, not a physics difference — was confirmed against the primary source rather than argued away. (For completeness: Hiscock and Konkowski’s Taub–NUT computation has $B < 0$ [7], and Hawking’s tensorial form is $\text{diag}(K, 3K, -K, -K)$ [4] — the sign structure our series also produces.)

7 Classical stability of the closed geodesic

Ori’s open question — “Stability remains the main open question” [18] — has a classical-geodesic part, which we answer quantitatively. In z -parameterization the transverse dynamics (5) is exact and universal (timelike or null), with solutions

$$x(z) \sim e^{s_{\pm}z}, \quad s_{\pm} = -\frac{1}{4} \pm \mu, \quad \mu = \sqrt{\frac{a}{2} + \frac{1}{16}}, \quad (16)$$

so the per-wrap transverse log-growth is $\mu - \frac{1}{4}$ per wrap of period L . At $a = 0$ (four-dimensional Misner) $\mu = 1/4$ exactly: the exponents are $\{0, -\frac{1}{2}\}$ and the transverse mode is marginal — wrapping trajectories are not ejected and the boost pile-up accumulates energy near the horizon plane (the Misner pile-up). For any $a > 0$ the sign flips: the growing mode $e^{(\mu - \frac{1}{4})z}$ wins and wrapping geodesics are *ejected transversely*. Full numerical integration of the geodesics confirms the analytic rate (Table 2; Fig. 4): the fitted per-wrap growth matches $\mu - 1/4$ across the curvature scan.

The Ori curvature thus reverses Misner’s test-geodesic pile-up into slow transverse ejection — the machine’s core does not accumulate test geodesics. Dust, the field stress, and the 4D backreaction are not included in this statement. The ejection scale $z_{\text{esc}} = \ln(X/x_0)/(\mu - \frac{1}{4})$ diverges as $a \rightarrow 0$: the Misner limit is the accumulation boundary. Combined with (4) — every wrapping geodesic completes infinitely many wraps within finite affine parameter, consistent with Ori’s future-incompleteness remark — the ejection is reached within finite proper time.

a	fitted per-wrap log-growth	analytic $\mu - 1/4$
0 (Misner)	+0.0022	0
0.005	+0.0068	+0.0050
0.010	+0.0114	+0.0098
0.020 (Ori)	+0.0203	+0.0193
0.050	+0.0457	+0.0458
0.100	+0.0844	+0.0854

Table 2: Transverse stability of $\mathcal{N}$: fitted per-wrap log-growth of the transverse offset vs. the analytic $\mu - 1/4$ (full timelike geodesic integration, 12 wraps, three initial conditions per value). The sign flip at exactly $a = 0$ is the Misner marginality.

7.1 Tensor modes: no amplifier, but horizon pile-up

The metric-perturbation analogue — stability of the geometry rather than of the geodesics — we analyze on the transverse-traceless ansatz $h_{xx} = -h_{yy} = A$, $h_{xy} = B$, all other components zero. The linearized Einstein system on this background was derived with the covariant machinery $\delta\Gamma_{\mu\nu}^\rho = \frac{1}{2}g^{\rho\sigma}(\nabla_\mu h_{\nu\sigma} + \nabla_\nu h_{\mu\sigma} - \nabla_\sigma h_{\mu\nu})$ and validated end-to-end against a finite-difference linearization of the exact nonlinear Ricci tensor (the Ricci-linearization checks pass at 10^{-5} or better and the core-mode evolution matches the exact solution to 3.6×10^{-4} after four wraps). Three results:

1. **No per-wrap tensor amplifier.** At the core the surviving h_{xy} polarization is exactly solvable: $B = F(z) + G(S)$ with $S = e^z(T - f_0)^2$; per wrap the profile *shifts*, $G(S) \mapsto G(e^L S)$, preserving the sup-norm (measured per-wrap log-growth -0.0002 over six wraps). Single-valuedness on the quotient $z \sim z + L$ forces $G \propto S^{2\pi i n/L}$: the Floquet spectrum lies on the unit circle — neutrally stable, for every a . The tensor sector does not share the geodesic ejection law $\mu - \frac{1}{4}$.
2. **One polarization is expelled from $\mathcal{N}$ at $k_\perp = 0$.** The $h_{xx}-h_{yy}$ polarization obeys $\delta R_{zz} = aA$ at the core in the $k_\perp = 0$ TT ansatz: for any $a > 0$ the linearized system forces $A = 0$ on the closed geodesic itself — the same curvature that ejects geodesics removes this polarization from $\mathcal{N}$ entirely (at $a = 0$ only, an extra non-periodic, horizon-singular family appears; the $a \rightarrow 0$ limit is discontinuous for this sector). This is a $k_\perp = 0$ statement, not a gauge-invariant one; at $k_\perp \neq 0$ the longitudinal component $E = h_{zz}$ can absorb the curvature term $aA/2$.
3. **Gradient pile-up onto the CTC boundary.** The h_{xy} packet's T -gradient grows $\times e^{L/2}$ per wrap (measured $+0.5000$) while its T -location compresses by $e^{-L/2}$ per wrap (measured -0.5001), asymptoting onto the surface $T = f_0 = a\rho^2/2$ — the CTC boundary ∂V of Eq. 2, which for $a \neq 0$ sits $a^2\rho^2$ above the local eikonal chronology horizon. The classical pile-up thus passes through the same region in which the semiclassical images of Sec. 5 compress; the stress scale $(a^2\rho^2)^{-2}$ keeps both small at the sampled point.

Caveats: the TT analysis is the $k_\perp = 0$ sector. The gauge-complete ten-component linearized system has now been derived and validated against a finite-difference linearization of the exact Ricci tensor (see the project notes [25]); it finds no per-wrap exponential amplifier in any sector, because ∂_z is Killing and every single-valued mode has multiplier $e^{ik_z L} = 1$. The expulsion of $h_{xx}-h_{yy}$ is a $k_\perp = 0$ statement only; at $k_\perp \neq 0$ the E component absorbs the curvature term and $S = h_{xx} - h_{yy}$ satisfies the same Bessel mode equation as B . The quotient Cauchy problem, constraint propagation, off-core $O(a\rho)$ terms, the invariant interpretation of the T -growth,

A4: classical stability of N — Ori's open question, answered quantitatively:
the curvature reverses Misner's matter accumulation into slow transverse ejection




Figure 4: The Misner–Ori flip. Left: per-wrap transverse offset for $a = 0$ (collapse toward the horizon plane) vs. $a = 0.02$ (ejection). Center: fitted growth rates vs. analytic $\mu - 1/4$ (Table 2). Right: the ejection scale diverges as $a \rightarrow 0$.

vector/scalar physicality, and nonlinear stability remain open. Within the analyzed tensor sector, there is no per-wrap instability mechanism that would shred the machine by its own radiation.

8 Discussion

8.1 Protection without teeth

The result is sharper than either the naive “protection fails” or the standard “protection holds”. Hawking’s mechanism is real near the local eikonal horizon: the ϵ^{-2} approach law holds as in Misner space, precisely as the KRW-type expectations suggest — but the full image-series tensor has eigenvalues $(-\frac{3}{2}, +\frac{1}{2}, +\frac{1}{2}, +\frac{1}{2})\mathcal{E}$ up to the $O(a)$ splitting, so the image-direction contraction is negative and $3/2$ times the magnitude of the positive proxy $\mathcal{E}$. The transport tangent is spacelike at both sampled points, so this contraction is a normalized spacelike eigenvalue, not a null contraction, and carries no NEC claim. The locus analysis of Sec. 5 shows the mechanism is *interrupted before it can do anything*: the local eikonal chronology horizon (the null surface $T = F$, not the timelike CTC boundary $T = f$, and not proven to be the global Cauchy horizon) lies at $\epsilon_{\text{CH}} = a^2 \rho^2 + O(a^3)$, where the energy series is within 8–9% of the asymptotic constant (91–93%; the dominant image is at 0.98 of the 0.51ϵ law) and finite. The CH value is verified at $x_0 = 0.3$ and scanned over $\rho = 0.10$ – 0.50 (nine values, $y = 0$, $a = 0.02$, $L = 1$): it is finite, at 91–93% of the asymptotic constant ($\epsilon_{\text{CH}}^2|_{\text{series}} = 0.2192$ at every ρ ; $\min_n |\sigma_n| = 0.5000 \epsilon_{\text{CH}} = 0.9804$ of the 0.51ϵ law) and scales as $(a^2 \rho^2)^{-2}$ (fitted slope -4.000). Since ϵ_{CH} is defined as $a^2 \rho^2$, the ρ^{-4} fit is a consistency check of the ϵ^{-2} law, not an independent measurement; the caveats are one transverse direction, the local $O(a^2)$ eikonal identification, and the test-field series. The CH energy-density series is a factor 6.6×10^{-5} of its formal past-horizon on- ∂V extrapolation at the leading-order endpoint (7.7×10^{-5} with the $O(a^3)$ correction); both surfaces scale as $(a^2 \rho^2)^{-2}$, so this ratio is a - and ρ -independent and is not evidence of a cutoff. In physical units the proxy is $\sim 5.3 \times 10^{-18} \text{ J/m}^3$, the image-direction eigenvalue is $-8.0 \times 10^{-18} \text{ J/m}^3$ (a spacelike direction), and the largest chart component is $1.78 \times 10^{-13} \text{ J/m}^3$ at $a = 0.02$: the proxy is about sixty orders below the unit-curvature stress, and the chart components about 55.4–55.7 orders (the largest, 55.4). The canonical design $a = 10^{-4}$ gives a proxy $\sim 8 \times 10^{-9} \text{ J/m}^3$. These are test-field magnitudes, not a survival proof. The only genuine singularity is on the single closed null geodesic $\mathcal{N}$, a set of measure zero. Whether one reads that as protection depends on which side of the interpretive fork one stands: Visser’s points-

or-parts reading — the pathology occupies a set too localized to enforce anything — or Hawking’s barrier reading [17, 4]. Our computation sharpens the fork: on a leading constructive candidate whose horizon is probably non-compactly generated, the barrier reading is quantitatively empty *except on $\mathcal{N}$* . This is conditional: if the horizon is compactly generated, KRW applies and this finite test-field result is not the relevant statement.

8.2 The state-choice band

What we computed is the in-vacuum image series; its singular structure is state-independent within the Hadamard class (the image sum’s light-cone singularities are fixed by the geometry, and the state-dependent coefficient is B -like, negative in the Taub–NUT computation of Hiscock and Konkowski [7, 4]). The state-choice literature is the second battleground. Krasnikov constructed states with bounded $\langle T_{\mu\nu} \rangle$ in two-dimensional CTC spacetimes [11]; Li and Gott constructed the four-dimensional adapted Rindler state for Misner space, with $\langle T_{\mu\nu} \rangle = 0$ everywhere in the initial globally hyperbolic region and self-consistent at $b = 2\pi$ under the semiclassical Einstein equations, and Cassidy’s effective-action analysis supports the $b = 2\pi$ zero [13, 12]. The KRW/Cramer–Kay reply is that in those states the stress tensor is nevertheless undefinable on the horizon itself [16]. Hawking’s own admission is the sharpest statement of the open issue: he exhibited propagators with vanishing stress tensor and dismissed them on positivity grounds — “I am almost sure there is no quantum state on Misner space for which $\langle T^{\mu\nu} \rangle$ is finite on the horizon, but I do not have a rigorous proof” [4].

A two-dimensional prototype computation performed for this project (on the sectional reduction of the Misner geometry, with the pipeline anchored to the Fulling result and validated against the in-vacuum benchmark) sharpens the band considerably, and corrects a common shorthand. The natural “adapted Rindler vacuum” — the quotient cylinder/Milne vacuum, i.e. Rindler vacuum with boost images — does *not* have vanishing stress: it carries $\langle T_{UU} \rangle = -(1/4\pi)(\pi^2/3\Lambda^2 + 1/12)/U^2$ (with $\Lambda = L/2$ the identification rapidity), a quotient-circle Casimir contribution plus the Fulling residue, which is *larger* in magnitude than the in-vacuum divergence for every Λ (ratio $R(\Lambda) = (\pi^2/3\Lambda^2 + 1/12)/S(\Lambda) > 1$). An exactly-stress-free state does exist, but it is a *tuned thermal state* — KMS in the cylinder time, at inverse temperature $\beta^*(\Lambda)$ with $2\pi\beta^*/\Lambda \approx 3.3$ — whose thermal excitations cancel the Casimir-plus-anomaly piece identically. This is a separate 2D mechanism from the self-consistency point of Li and Gott [13]: read from the primary source, their $\langle T_{ab} \rangle \propto [(2\pi/b)^4 - 1]$ vanishes at $b = 2\pi$; they note that the other zero, $b = i2\pi$, “is just the usual Minkowski vacuum expressed as a thermal state relative to a Rindler observer,” with periodicity in imaginary time “a characteristic of a thermal Green function.” The 2D thermal tune is a fine-tuning, not a generic property of adapted vacua; in 4D the $b = 2\pi$ zero is instead geometric, as we now describe. The Cramer–Kay objection (stress undefinable on the horizon itself, where the boost-image singularities sit) survives undiminished. On the Ori background this dispute becomes quantitative: the band between (i) the in-vacuum series computed here (rises on approach, cut off and finite at the local eikonal horizon, singular only on $\mathcal{N}$) and (ii) the fine-tuned thermal zero-state (if its analogue exists on this quotient) is the honest state-choice uncertainty — and the horizon divergence is, in this sense, entirely state-dependent in the 2D prototype: whether an analogous zero-stress state exists on the 4D Ori quotient is the open A3 task, and in the 2D prototype the dispute reduces to whether nature supplies the fine-tuning.

A four-dimensional algebraic check performed for this project (roadmap item A3; App. A) verifies the Li–Gott formula; the full 4D point-split has not been re-done. For 4D Misner space

($f = 0$), Li and Gott’s image-summed Rindler vacuum has the exact stress

$$\langle T_{ab} \rangle = \frac{1}{1440\pi^2\xi^4} \left[\left(\frac{2\pi}{b} \right)^4 - 1 \right] (4\xi^2 d\eta_a d\eta_b + g_{ab}), \quad (17)$$

which is traceless and *exactly zero at $b = 2\pi$* [13]; we verify its trace, the $b = 2\pi$ zero, the $b \rightarrow \infty$ Fulling limit $\rho = -1/(480\pi^2\xi^4)$, and the coefficient against the Brown–Ottewill anomaly result as quoted by Michel [14]. The 2D statement that the adapted vacuum is worse for every Λ is therefore 2D-specific: the relative sign of the image and Fulling terms is dimension-dependent, and the 2D thermal β^* zero is a separate mechanism. In 4D the $b = 2\pi$ zero is a *geometric* condition (the smooth Euclidean cone, $L = 4\pi$), not a thermal fine-tuning. For the Ori $f \neq 0$ core the situation differs again: exact symbolic evaluation shows the core is a VSI pp-wave ($R_{\mu\nu} = R = 0$, $\text{Riemann}^2 = 0$, Euler density = 0), so the local conformal trace anomaly vanishes identically and cannot enter the bulk balance; the metric has no flat covering or boost isometry, the Rindler modes are not solutions, and no closed-form adapted vacuum exists (a restricted no-go: no flat covering and no real Euclidean section, since $\det g = -1$ and the required Wick rotations are complex). The quotient/conical anomaly on $T = f$ is now computed [26]: the surface is timelike (null only at $\mathcal{N}$), the local cone has rapidity $\Lambda = L/2$ and deficit $\delta = 2\pi - L/2$, with the smooth cone at $L = 4\pi$; the local 4D invariants vanish, but the conical Li–Gott/Misner piece and the independent quotient-circle Casimir $\rho_{\text{Cas}} = -1/(1440\pi^2\tau^4)$ at $L = 4\pi$ do not cancel automatically, so a state tune is still required. A zero-stress state on the Ori quotient is therefore neither constructed nor excluded.

The two-dimensional backreaction companion is exact and narrow: in 2D the Einstein tensor vanishes identically, every CFT state on the flat Misner background is traceless, and the conformal mode is unsourced, so the horizon does not move at any order for the in-vacuum, adapted, or thermal states [27]. That is a no-go for the 2D model, not an answer to the 4D question: the 4D Weyl sector is dynamical and is sourced by the traceless stress. The 4D backreaction remains open.

On the Ori quotient the substantive QFT tasks are thus the construction of an adapted state and a state-selection principle (the quotient anomaly is now computed; Sec. 9). The closest 2025 construction of an invariant state on a CTC quotient is that of Melton, Strominger and Wang, who build a time-periodic AdS_3/Z mode basis with an $\text{SO}(2,2)$ -invariant vacuum and Fock space [23]; it is a different quotient, so the 4D Ori choice with $f \neq 0$ remains open.

8.3 Comparison with the holographic program

Empanan and Tomašević, in controlled holographic models, find that strongly coupled fields decouple the chronology-violating region entirely [19]; in a companion backreaction analysis the null chronology horizon turns into a strong spacelike curvature singularity [20]. Protection, if it holds, is by a different and weaker mechanism than Hawking’s conjecture, and does not involve Planck-scale physics. Our result is the weakly-coupled, Einstein-gravity counterpart of the same message: on a leading constructive candidate, the standard Hawking stress mechanism is negligible off the core and cannot enforce protection, so if protection holds there, it must come from elsewhere — from the holographic mechanism, from the back-reaction endpoint, or from full quantum gravity, exactly as Ori anticipated [18].

8.4 Ori’s open question, partially answered

Combining the pieces: no accumulation of test geodesics (ejection at rate $\mu - 1/4$, Sec. 7; dust and backreaction are not computed), no effective quantum barrier off $\mathcal{N}$ (finite at the local eikonal

horizon, Sec. 5), and every wrapping geodesic future-incomplete (the affine law, Sec. 2 — consistent with Ori’s own remark). The classical-geodesic part of “stability remains the main open question” now has a quantitative answer, and it points toward viability; survival still requires the uncomputed 4D backreaction.

9 Limitations

We state the caveats explicitly.

(1) *Image series; full tensor at two points.* We assemble the full image-series tensor, all independent components, at the anchor and on the leading-order chronology horizon, with the exact bitensor transport and the full 4×4 Jacobi block. The statements that are exact in that computation are the bitensor transport, the longitudinal Van Vleck determinant λ^2 , tracelessness, and the vanishing anomaly. The approximations are: leading order in $1/\sigma$ (estimated relative $O(R\sigma) < 10^{-6}$ at the CH, not implemented); one dominant connecting geodesic per image (completeness of the connecting set is not proved); in-vacuum series only; fixed background and test field. The nine-point ρ scan and the 0.870 transport check are energy-density series evaluated with the earlier start-point tangent, fixed $n_{\max}$, and no Δ factor; they are consistency checks of the power law, not full-tensor evaluations, and the full-tensor normalization supersedes 0.870 at the anchor and the CH. The scan is one-dimensional ($y = 0$), uses the local $O(a^2)$ eikonal CH identification, and remains a test-field evaluation at and beyond the local eikonal horizon. The open items are the subleading $O(1/\sigma)$ finite part, the completeness of the connecting-geodesic set, and the 4D backreaction, not the tensor assembly.

(2) *In-vacuum only.* All computed stresses are for the in-vacuum image series. The state-choice map (the “band” of Sec. 8) is partly resolved in 4D: the Li–Gott formula is verified algebraically (trace, the $b = 2\pi$ zero, the Fulling limit), but the full 4D point-split has not been re-done. The adapted Misner ($f = 0$) state is exactly stress-free at $b = 2\pi$ (Li–Gott Eq. (12); Sec. 8.2); the 2D “adapted is worse” result is 2D-specific and does not generalize. On the Ori $f \neq 0$ core the adapted state is not constructed: the metric has no flat covering or boost isometry, the Rindler modes are not solutions, and there is no real Euclidean section (a restricted no-go). The exact core is a VSI pp-wave, so its local conformal trace anomaly vanishes identically; the quotient/conical anomaly on the timelike $T = f$ surface is now computed, and the conical and quotient-Casimir pieces do not cancel automatically, so a state tune is still required. What remains open is therefore state selection on the Ori quotient: a zero-stress state there is neither constructed nor excluded. The benchmark state-choice results (Krasnikov and our 2D prototype) are two-dimensional, while Li and Gott’s construction is four-dimensional; our computation is four-dimensional on the $(3 + 1)$ Ori geometry, and two-dimensional lessons about adapted states do not automatically transfer.

(3) *Fixed background, test field.* We compute $\langle T_{\mu\nu} \rangle$ of a test field on the fixed geometry; no back-reaction is included, and the geometry’s own matter support (outside the vacuum core, in the full Ori solutions) is not part of the quantum computation. In two dimensions the horizon shift is exactly zero and structurally so: 2D gravity is topological, the CFT stress is traceless, and the horizon locus is conformally invariant, so the literal 2D semiclassical equation reduces to a state constraint ($\langle T_{ab} \rangle = 0$), not a metric equation. That is a no-go for the 2D model, not a statement that backreaction vanishes in general: the 4D Weyl sector is dynamical and is sourced by the traceless stress. The 4D backreaction endpoint is therefore open, holography suggests a spacelike singularity forms in strongly-coupled models [20], and the semiclassical approximation itself is unreliable in an arbitrarily small neighborhood of $\mathcal{N}$ [17]. Easson’s state-dependent Cauchy-horizon flux coefficients in the reduced Polyakov model [24] are a useful methodological comparator for this back-reaction

question, though the setting is a charged black hole rather than a CTC core.

(4) *The z -parameterization subtlety.* The transverse stability law is exact in z -parameterization, but z is not an affine parameter: wrapping geodesics reach $z \rightarrow \infty$ within finite affine parameter ($\lambda_* = 2/\dot{z}_0$). “Per-wrap” rates are rates in z ; the picture is nevertheless coherent, since the ejection scale is reached at finite affine parameter.

(5) *Mass.* We treated a massless conformal scalar. Massive fields are known to regularize chronology-horizon polarization in related spacetimes [15]; a massive-field scan on the Ori geometry is straightforward but not performed here.

10 Conclusions

We computed the renormalized image-series energy density of a massless conformal scalar on the Ori (2005) time machine, and the full image-series tensor at the anchor and on the local eikonal chronology horizon: to our knowledge, the first image-series energy-density computation of the conformal scalar on the Ori 2005 core in GR (literature search 2026-10-26). The mechanism Hawking proposed for chronology protection operates: the in-vacuum stress follows the ϵ^{-2} Misner law on approach to the CTC boundary ∂V , with $\min_n |\sigma_n| \approx 0.51\epsilon$ over the asymptotic range, and the $|T_0|^{-2}$ series law is Hawking’s own t^{-4} law in disguise (slope -4.00 against his distance function; Rindler dictionary exact to 2×10^{-15}). But the full tensor changes the sign: the eigenvalues are $(-\frac{3}{2}, +\frac{1}{2}, +\frac{1}{2}, +\frac{1}{2})\mathcal{E}$ up to the $O(a)$ splitting, and the image-direction contraction is $-3\mathcal{E}/2$, negative and opposite to the positive proxy. The transport tangent is spacelike at both sampled points, so this is a normalized spacelike eigenvalue, not a null contraction; we draw no NEC conclusion from it. The same curvature that displaces the local eikonal horizon below ∂V ends the rising regime at $\epsilon_{\text{CH}} = a^2\rho^2 + O(a^3)$, where the energy series is within 8–9% of the asymptotic constant (91–93%) and finite. At $a = 0.02$ for a metre-scale machine the proxy is $5.3 \times 10^{-18} \text{ J/m}^3$, the image-direction eigenvalue $-8.0 \times 10^{-18} \text{ J/m}^3$ (a spacelike direction), and the largest full-tensor chart component $1.78 \times 10^{-13} \text{ J/m}^3$; the energy-density series is 6.6×10^{-5} of its formal past-horizon on- ∂V extrapolation, a ratio between two surfaces with the same $(a^2\rho^2)^{-2}$ scaling, not evidence of a cutoff. At the canonical design $a = 10^{-4}$ the proxy is $\sim 8 \times 10^{-9} \text{ J/m}^3$. On the CH the only genuine singularity remains on the single closed null geodesic $\mathcal{N}$; the CH series is verified at $x_0 = 0.3$ and scanned over $\rho = 0.10\text{--}0.50$ (nine values, $y = 0$, $a = 0.02$, $L = 1$) as a consistency check of the ϵ^{-2} law ($\epsilon_{\text{CH}}^2|\text{series}| = 0.2192$ at every ρ , 91–93% of the asymptotic constant; $\min_n |\sigma_n| = 0.5000\epsilon_{\text{CH}} = 0.9804$ of the 0.51ϵ law; fitted slope -4.000 because $\epsilon_{\text{CH}} \equiv a^2\rho^2$); the caveats are one transverse direction, the local $O(a^2)$ eikonal identification, and the test-field series. The classical-geodesic part of Ori’s stability question answers the same way: no accumulation of test geodesics, transverse ejection at rate $\mu - 1/4$ instead (dust and backreaction are not computed). Semiclassical chronology protection thus operates on a leading constructive time machine candidate, but the standard Hawking stress mechanism is quantitatively negligible at the sampled point: the proxy is about sixty orders below the unit-curvature stress, and the full-tensor chart components about 55.4–55.7 orders (the largest, 55.4). This is protection in form, not in force — for the test field. Survival of a macroscopic machine requires the uncomputed 4D backreaction. The conclusion is conditional on non-compact generation: if the horizon is compactly generated, KRW applies and this finite test-field result is not the relevant statement. What remains is the construction of an adapted state and a state-selection principle on the Ori quotient (the conical anomaly is computed; a zero-stress state is neither constructed nor excluded), the subleading tensor terms and completeness of the connecting-geodesic set, and the 4D backreaction endpoint — and, if protection is real, the expectation that it comes from beyond the semiclassical mechanism

[18, 19]. The gauge-complete ten-component analysis narrows the metric-stability front: there is no per-wrap exponential amplifier in any sector, and the TT expulsion is a $k_{\perp} = 0$ statement only.

The paper documented three mid-investigation corrections openly — the supersession of the closed-null-geodesic census by the image-series stress computation, the coordinate-convention reconciliation with Hawking’s exponent, and the CH-locus correction (the surface $T = f$ is the timelike CTC boundary, with the local eikonal horizon below it) — and we suggest they be read as what they are: evidence that the computation, and its error bars, are honest.

A Numerical methods and reproducibility

Geodesics are integrated with the explicit Runge–Kutta Dormand–Prince 8(5,3) scheme (tolerances 10^{-9} – 10^{-13}), with event detection at $z = nL$; the image boundary-value problems are solved by multi-start root finding (hybrid Powell method, 4^3 -point direction grid, endpoint tolerance 10^{-5}). Validation gates, all passed: $\det g = -1$ exactly; null condition conserved to 4.6×10^{-12} after three wraps; $\mathcal{N}$ closure to machine zero; the Misner limit $f = 0$ reproducing the known fountain structure; the direction-search result stable across integrator tolerances 10^{-9} – 10^{-13} ; and the transport checks (Sec. 6). All figures regenerate from the repository scripts (`simulations/ori_horizon.py`, `ori_tmn_full.py`, `a2_rigor.py`, `ori_full_tensor.py`, `ori_stability.py`, `horizon_locus.py`, `a4_tensor_stability.py`, `ori_gauge_complete.py`, `a3_adapted_state.py`, `ori_adapted_4d.py`, `ori_adapted_state.py`, `ori_backreaction_2d.py`, `ch_scan.py`), which run on a standard scientific Python installation; every simulation number is reproduced by one of them (`ch_scan.py` reproduces the $x_0 = 0.3$ anchor and the nine-point $\rho = 0.10$ – 0.50 CH scan). The full-tensor script reproduces the geometry gates ($\det g = -1$, $|R_{\mu\nu}| = 2 \times 10^{-13}$, Kretschmann = 0), the full 4×4 Jacobi check ($\det J/(J_x J_y \lambda^2) = 1 + 1.2 \times 10^{-10}$, $J_{Tz} = 0.112$), the Misner benchmark slopes (-2 for the eigenvalues and -3 for the chart components), the zero trace, the $n_{\max}$ convergence, and the ρ scan of the eigenvalue ratios. The 4D state-choice algebra check is `ori_adapted_4d.py` (runtime ≈ 1.5 s; figure `simulations/ori_adapted_4d.png`); its validation gates are the exact symbolic Ori invariants ($\det g = -1$, $R_{\mu\nu} = R = 0$, Riemann² = 0, Euler density = 0), the exact Li–Gott algebra (tracelessness, the $b = 2\pi$ zero, the $b \rightarrow \infty$ Fulling value $-1/(480\pi^2\xi^4)$), the 2D identity $C_{\text{ad}}(\Lambda) - C_{\text{ad}}(2\pi) = (1/48\pi)[(2\pi/\Lambda)^2 - 1]$ and thermal β^* gates at $\Lambda = 0.5, 1, 2$, the 2D ratio $R(\Lambda) > 1$, and the $S_4(\Lambda)$ image-sum asymptotics (within 1% at small Λ , 2% at large Λ). The $f \neq 0$ adapted-state programme is `ori_adapted_state.py`: its gates are the exact $f = 0$ Hankel basis (1.2×10^{-15}), the timelike conical geometry on $T = f$ (rapidity $L/2$, smooth cone $L = 4\pi$), the traceless Li–Gott algebra at $L = 4\pi$, and the nonzero quotient Casimir. The 2D backreaction is `ori_backreaction_2d.py` (Rindler dictionary 6.6×10^{-15} , the closed-form stresses, the trace 2.9×10^{-16} , and the horizon-locus invariance). Unit conversions and the $O(a^3)$ endpoint correction are derived in the text.

References

- [1] H. P. Künzle, “Maxwell fields satisfying Huygens’s principle,” *Math. Proc. Camb. Phil. Soc.* **64**, 779 (1968).
- [2] J. Carminati and R. G. McLenaghan, “Determination of all Petrov type-N space-times on which the conformally invariant scalar wave equation satisfies Huygens’ principle,” *Phys. Lett. A* **105**, 351 (1984). doi:10.1016/0375-9601(84)90278-0.

- [3] A. I. Harte, “Tails of plane wave spacetimes: Wave-wave scattering in general relativity,” *Phys. Rev. D* **88**, 084059 (2013).
- [4] S. W. Hawking, Chronology protection conjecture, *Phys. Rev. D* **46**, 603 (1992).
- [5] C. W. Misner, “Taub-NUT space as a counterexample to almost anything,” in *Relativity Theory and Astrophysics I: Relativity and Cosmology*, Lectures in Applied Mathematics, Vol. 8, edited by J. Ehlers (American Mathematical Society, Providence, RI, 1967), pp. 160–169.
- [6] J. D. E. Grant, “Cosmic strings and chronology protection,” *Phys. Rev. D* **47**, 2388 (1993); arXiv:hep-th/9209102.
- [7] W. A. Hiscock and D. A. Konkowski, “Quantum vacuum energy in Taub-NUT (Newman-Unti-Tamburino)-type cosmologies,” *Phys. Rev. D* **26**, 1225 (1982). doi:10.1103/PhysRevD.26.1225.
- [8] S.-W. Kim and K. S. Thorne, Do vacuum fluctuations prevent the creation of closed timelike curves?, *Phys. Rev. D* **43**, 3929 (1991).
- [9] G. Klinkhammer, Vacuum polarization of scalar and spinor fields near closed null geodesics, *Phys. Rev. D* **46**, 3388 (1992).
- [10] B. S. Kay, M. J. Radzikowski, and R. M. Wald, Quantum field theory on space-times with a compactly generated Cauchy horizon, *Commun. Math. Phys.* **183**, 533 (1997); arXiv:gr-qc/9603012.
- [11] S. V. Krasnikov, “Quantum stability of the time machine,” *Phys. Rev. D* **54**, 7322 (1996); arXiv:gr-qc/9508038. doi:10.1103/PhysRevD.54.7322.
- [12] M. J. Cassidy, “Divergences in the effective action for acausal spacetimes,” *Class. Quantum Grav.* **14**, 3031 (1997); arXiv:gr-qc/9705075. doi:10.1088/0264-9381/14/11/006.
- [13] L.-X. Li and J. R. Gott, “A self-consistent vacuum for Misner space and the chronology protection conjecture,” *Phys. Rev. Lett.* **80**, 2980 (1998); arXiv:gr-qc/9711074. doi:10.1103/PhysRevLett.80.2980.
- [14] B. Michel, “Remarks on Rindler quantization,” arXiv:1612.03158 (2016).
- [15] T. Tanaka and W. A. Hiscock, Chronology protection and quantized fields: nonconformal and massive scalar fields in Misner space, *Phys. Rev. D* **49**, 5240 (1994).
- [16] C. R. Cramer and B. S. Kay, “The thermal and two-particle stress-energy must be ill-defined on the 2-d Misner space chronology horizon,” *Phys. Rev. D* **57**, 1052 (1998); arXiv:gr-qc/9708028.
- [17] M. Visser, “The quantum physics of chronology protection,” in *The Future of Theoretical Physics and Cosmology*, edited by G. W. Gibbons, E. P. S. Shellard, and S. J. Rankin (Cambridge University Press, Cambridge, 2003), pp. 161–176; arXiv:gr-qc/0204022.
- [18] A. Ori, A class of time-machine solutions with a compact vacuum core, *Phys. Rev. Lett.* **95**, 021101 (2005); arXiv:gr-qc/0503077.
- [19] R. Emparan and M. Tomašević, “Holography of time machines,” *J. High Energy Phys.* **03**, 212 (2022); arXiv:2107.14200. doi:10.1007/JHEP03(2022)212.

- [20] R. Emparan and M. Tomašević, “Quantum backreaction on chronology horizons,” *J. High Energy Phys.* **02**, 182 (2022); arXiv:2109.03611. doi:10.1007/JHEP02(2022)182.
- [21] J. Dietz, A. Dirmeier, and M. Scherfner, “Geometric Analysis of Ori-type Spacetimes,” arXiv:1201.1929 (2012).
- [22] N. E. Rieger, “Causal classification of pathological Misner-type spacetimes,” *Annals of Physics* **489**, 170455 (2026); arXiv:2501.00841. doi:10.1016/j.aop.2026.170455.
- [23] W. Melton, A. Strominger, and T. Wang, “Quantum Fields on Time-Periodic $\text{AdS}_3/\mathbb{Z}$,” arXiv:2510.15036 (2025).
- [24] D. A. Easson, “Cauchy-horizon flux coefficients in the reduced Polyakov model,” *Class. Quantum Grav.* (2026), published online; arXiv:2511.05656. doi:10.1088/1361-6382/aeb78e.
- [25] J. Schemoul, “Gauge-complete linearized stability of the Ori time machine,” project notes (in preparation, 2026).
- [26] J. Schemoul, “Adapted states and the conical anomaly on the Ori core,” project notes (in preparation, 2026).
- [27] J. Schemoul, “Two-dimensional backreaction on the Misner/Ori quotient,” project notes (in preparation, 2026).
- [28] S. Xavier, “Closed Timelike Curves from a Vacuum Traveling Wave,” arXiv:2607.00788 (2026).
- [29] D. Bermudez and U. Leonhardt, “Optical time travel: proposal for testing Hawking’s Chronology Protection Conjecture in an optical analogue,” *Class. Quantum Grav.* **43**, 14LT01 (2026); arXiv:2607.22056.
- [30] B.-H. Lee, N. A. Nilsson, and S. Thakur, “Global causality constraints in rotating scalar-tensor spacetimes,” *Phys. Rev. D* **114**, 084027 (2026); arXiv:2602.17724. doi:10.1103/3wgy-tzpl.
- [31] F. Ahmed, A. Al-Badawi, and I. Sakalli, “Are Petrov type-N and D spacetimes admitting CTCs valid in $f(R, L_m, \Phi, X)$ gravity?,” *Eur. Phys. J. C* **86**, 792 (2026). doi:10.1140/epjc/s10052-026-16044-w.
- [32] D. Senjaya, “Probing the chronology protection conjecture via scalar quasisresonance,” *Eur. Phys. J. C* **86**, 34 (2026). doi:10.1140/epjc/s10052-026-15299-7.