

Quotient structure and on-generator diagnostics of vacuum traveling-wave time machines

Jonathan Schemoul

LibertAI, France

js@libertai.io

October 27, 2026

Abstract

Xavier’s metric admits a switchable quotient (the periodic train family); the paper’s own worked examples are helical and not switchable. In the train quotient a wave profile turns the compact z -circles from spacelike to timelike, so that closed timelike curves can be switched on, and a degenerate family is claimed in which the Hawking–Misner winding pile-up is structurally absent. We analyze the global and semiclassical structure of this core. First, we classify the quotient legality of the traveling-wave solution $h = P(z)Q(z - 2t)$ for two identifications of the family $(z, t) \sim (z + L, t + cL)$, the pure z -shift ($c = 0$) and the helical shift ($c = 1/2$). For general c the invariance condition admits further legal profiles (for example $c = 1/4$ with $P = 1$ and Q of period $L/2$); these lie outside the switchable machine class because the identification shifts t , and we do not classify them. For pure z -shift ($c = 0$) the general legal profile is *a train modulated by a pump*, $h = e^{2at}p(z)q(z - 2t)$ with p, q L -periodic: it is z -dependent generically, has boost $b = aL$, and carries the Misner pile-up for $a \neq 0$. The *trains* (the boost-degenerate subfamily $a = 0$, P and Q both L -periodic) are marginal; the *pumps* (the z -independent subfamily $p = q = 1$, $h = e^{2at}$) carry the pile-up. *Helical* cores ($c = 1/2$, P periodic and Q arbitrary) are a separate legal class, boost zero on the closed null generator, marginal, but not switchable. The paper’s own $\text{sech}^2(k(z - 2t))$ and exponential examples do not descend to the pure z -quotient the paper imposes; they are legal only under a helical identification, which has no global time and is not a switchable machine. Second, we measure the monodromy over 24 circuits: the train’s affine circuit length is exactly constant (ratio 1.0000000000; $\lambda = 2/3$ at $z_* = 0$), while the pump’s shrinks geometrically at $e^{-L/2\tau}$ (ratio 0.9512294245 for $\tau = 10$, $L = 1$). Third, we compute the truncated primary-branch deck image sum ($n \leq 12$) with Hadamard and Van Vleck weights. At fixed transverse distance ε the truncated train series is finite; it diverges as ε^{-2} for the σ^{-1} weight and ε^{-4} for the σ^{-2} weight as $\varepsilon \rightarrow 0$, the same powers as exact Misner space. This is the $\mathcal{N}$ -singularity of the closed null generator, not a core discriminator, and it falsifies the reading of the monodromy result as an ε -independent finite value. The genuine discriminator is *on the generator*: the train’s winding coefficients are bounded (measured log-rates $\equiv 0$ to 10^{-12}), while the pump’s screen coefficients grow at exactly $e^{+L/2\tau}$ and its spacing-weighted terms at $e^{+3L/2\tau}$ per circuit (measured rates $+0.05000000$ and $+0.15000000$). “Quantum silence” therefore means *no geometric pile-up*, not a finite on-generator two-point function. The off- $\mathcal{N}$ image sum/parametrix—unrenormalized and truncated at $n \leq 12$ —is computed at one point and is finite (at $x_* = 0.1$, $t_* = 0.2499204$: the truncated image sum is finite, its horizon limit is $\mathcal{O}(1)$ (degree zero in the horizon distance), and the near-axis scaling $S \sim C/x^2$ confines the divergence to the axis $x = 0$), while the train is classically neutral in the analyzed TT, profile, f -modulus, and vector-truncation sectors, with the full ten-component, longitudinal/constraint, $\mathcal{O}(A\rho k_\perp)$, $\mathcal{O}(A\rho^2)$, quotient-Cauchy, and nonlinear sectors open. All quantitative results use the core parameters $(A, L, \tau) = (0.10, 1, 10)$; the image sum uses $z_* = 0.25$ and the 24-circuit scan starts at $z_* = 0$.

1 Introduction

Hawking’s chronology protection conjecture proposes that the laws of physics prevent the appearance of closed timelike curves (CTCs) [3]. The proposed mechanism is semiclassical: on a chronology horizon, null geodesics winding around a closed null geodesic pile up, a quantum field acquires a boost holonomy at each circuit, and the renormalized stress-energy $\langle T_{\mu\nu} \rangle$ diverges. Fed back through the semiclassical Einstein equations, this is supposed to destroy the horizon before CTCs form. The mechanism is well established for compactly generated chronology horizons, where the Kay–Radzikowski–Wald theorem makes the two-point function ill-defined at the horizon [4], and it has been computed explicitly in Misner space and related models [3, 7, 8]; Misner space itself is due to [5]. It is not a theorem for non-compactly generated horizons, and the question of what happens there is model dependent [6].

The Ori (2005) vacuum core [2] is a leading non-compactly generated construction: a plane-wave core with boost-identified z -circles, surrounded by matter. Its chronology horizon carries a single closed null geodesic; the companion image-series analysis of this program [9], still in preparation, reports that the stress on approach follows the Misner law but is cut off by the horizon, leaving a finite stress at the horizon at its sampled point, and a true singularity only on the measure-zero generator. We cite that analysis as a qualitative precedent; its numerical normalization is not used here. The Ori core is permanent: the identification is part of the geometry and the CTCs are always present.

Xavier’s 2026 vacuum traveling-wave core [1] adds two features that the Ori core lacks. First, its metric admits a *switchable* train quotient: the z -circles are spacelike where $h > f$ and timelike where $h < f$, so a passing wave profile can turn the CTCs on. Second, it contains a degenerate-horizon family in which the Hawking–Misner winding compression is claimed to be absent. The core is therefore a candidate for a machine that can sit off and be armed on demand, and a candidate for a horizon that is quiet by construction rather than merely cut off.

This paper is an analysis of that core, not a construction of the machine. It answers four questions. (i) Which of the two identifications used by the Xavier mechanism and its worked examples, the pure z -shift ($c = 0$) and the helical shift ($c = 1/2$), are legal? (ii) What is the monodromy of the closed null generator, measured over many circuits rather than inferred from one? (iii) What does the truncated primary-branch deck image sum ($n \leq 12$) with Hadamard weights do at fixed and vanishing transverse distance? (iv) Where, if anywhere, is the train/pump discriminator?

The answers sharpen the Xavier proposal and correct two claims.

1. Neither the original dichotomy nor the first correction is exhaustive. For pure z -shift the general legal family is $h = e^{2at}p(z)q(z - 2t)$ with p, q L -periodic—a train modulated by a pump—with boost $b = aL$ and the Misner pile-up for $a \neq 0$. The trains ($a = 0$, P and Q both L -periodic) are its boost-degenerate subfamily and the pumps ($p = q = 1$, z -independent) its z -independent subfamily. The helical identification ($c = 1/2$) is a separate legal class. An earlier version of this analysis considered only the pure z -shift and concluded that there were exactly two legal families; that conclusion is retracted. The correction is recorded in the repository notes.
2. The paper’s own $\text{sech}^2(k(z - 2t))$ and degenerate exponential $e^{k\epsilon(z - 2t)}$ profiles do *not* descend to the pure z -quotient the paper itself imposes. They are legal under a helical identification $(z, t) \sim (z + L, t + L/2)$, but that quotient has no global time coordinate, and it is not a switchable machine: the identification vector is timelike wherever $f > 0$, so the helical core has permanent CTCs and no off state.

3. Measured over 24 circuits, the train's affine circuit length is exactly constant and the pump's shrinks at the exact Misner rate. This confirms the one-circuit spectrum and closes the prototype-inference gap.
4. The truncated image sum is finite at fixed ε but diverges as ε^{-2} (σ^{-1} weight) and ε^{-4} (σ^{-2} weight) as the generator is approached. The “ ε -independent finite value” reading is falsified. The surviving discriminator is on the generator: the train's spacing-weighted winding coefficients are bounded, while the pump's screen coefficients grow at $e^{+L/2\tau}$ and its spacing-weighted terms at $e^{+3L/2\tau}$.

The phrase “quantum silence,” retained here as the name of the geometric claim, is used in this precise sense: the geometric pile-up channel is absent for the train family. It does not mean that the stress is finite at the generator; no such claim is made. The off- $\mathcal{N}$ image sum/parametrix is computed at one point in Sec. 7 and is finite there; the train's linearized stability is neutral in the analyzed sectors, with the remaining sectors open as stated in Sec. 8.

The paper is organized as follows. Sec. 2 fixes the metric and the vacuum structure. Sec. 3 classifies the two identifications analysed. Sec. 4 measures the monodromy and the pile-up. Sec. 5 computes the Hadamard-weighted deck image sum. Sec. 6 isolates the on-generator discriminator. Sec. 7 computes the off- $\mathcal{N}$ image sum/parametrix at one horizon point, and Sec. 8 gives the linearized stability result for the analyzed sectors and lists the open ones. Sec. 9 discusses the result and Sec. 10 lists the open items. Sec. 11 concludes.

Conventions. We use signature $(-, +, +, +)$, units $8\pi G = c = \hbar = 1$ unless stated, and coordinates (t, x, y, z) . The coordinate z is periodic with period L ; the identification is specified in each section. A dot denotes a derivative with respect to an affine parameter; a comma denotes a partial derivative. The transverse saddle strength is A , the pump time scale is τ , and the reference core parameters are $(A, L, \tau) = (0.10, 1, 10)$; the image-sum reference point is $z_* = 0.25$. The transverse distance to the closed null generator $\mathcal{N}$ is ε .

2 The metric and vacuum structure

2.1 The metric

The Xavier core is

$$ds^2 = -2h(z, t) dt dz + [h(z, t) - f(x, y, z)] dz^2 + dx^2 + dy^2, \quad (1)$$

with $z \sim z + L$ at this stage only a formal periodicity of the profile. In matrix form, in the order (t, x, y, z) ,

$$g_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & -h \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -h & 0 & 0 & h - f \end{pmatrix}, \quad \det g = -h^2. \quad (2)$$

The inverse is

$$g^{\mu\nu} = \begin{pmatrix} -(h - f)/h^2 & 0 & 0 & -1/h \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -1/h & 0 & 0 & 0 \end{pmatrix}. \quad (3)$$

The vector ∂_t is null, and ∂_z has norm squared $h - f$. Thus the z -circles at fixed (t, x, y) are spacelike where $h > f$ and timelike where $h < f$. The locus $h = f$ is where they are null. This is the switch: a wave profile $h(z, t)$ passing through the core changes the causal character of the circles while f stays fixed.

2.2 Vacuum equations and general solution

A direct computation gives

$$G_{xx} = G_{yy} = -\frac{1}{2h} \partial_t \left(\frac{h_{,t} + 2h_{,z}}{h} \right), \quad G_{zz} = \frac{1}{2} (f_{,xx} + f_{,yy}), \quad (4)$$

and all other components vanish identically. Vacuum therefore requires

$$\partial_t \left(\frac{h_{,t} + 2h_{,z}}{h} \right) = 0, \quad f_{,xx} + f_{,yy} = 0. \quad (5)$$

The first equation is solved by

$$h(z, t) = P(z) Q(z - 2t) \quad (6)$$

for arbitrary nonvanishing P and Q . Indeed, the bracket in Eq. (5) is a function $C(z)$ of z alone, so $h_{,t} + 2h_{,z} = C(z)h$; along the characteristics $z - 2t = \text{const}$ (where $dz/dt = 2$) this is $d \ln h / dt = C(z)$, whose general solution is $h = P(z)Q(z - 2t)$. Substitution confirms Eq. (6) for arbitrary P, Q . The $\alpha = 2\epsilon$ exponential $h = e^{k\epsilon(z-2t)}$ is the special case $P = 1$, $Q = e^{k\epsilon(z-2t)}$; the general solution is not a special tuning but the full traveling-wave freedom. The second equation says that f is harmonic in the transverse plane at each z .

2.3 Geodesic identity and optical scalars

Two structural facts will be used below. First, the Christoffel symbol Γ^t_{zz} is proportional to $h - f$, so the z -circle tangent ∂_z is null where $h = f$. Being null is not sufficient for the circle to be geodesic: the transverse geodesic equation is $x'' + \Gamma^x_{zz}(z')^2 = 0$ with $\Gamma^x_{zz} = f_{,x}/2$ (and similarly for y). Thus a ∂_z orbit is a geodesic only where

$$h = f \quad \text{and} \quad f_{,x} = f_{,y} = 0. \quad (7)$$

For the saddle $f = P(z)^2 + (A/2)(x^2 - y^2)$ this is the axis $x = y = 0$. The horizon $h = f$ away from the axis is a locus of null circles, not of closed null geodesics; only the axis circle at the horizon epoch is the closed null generator $\mathcal{N}$.

Second, for transverse indices $a, b \in \{x, y\}$,

$$\Gamma^t_{ab} = \Gamma^z_{ab} = 0, \quad (8)$$

so the expansion, shear, and rotation of the z -congruence vanish identically. The transverse block is rigid. Hawking's per-circuit area factor is therefore

$$f_{\text{area}} = 0. \quad (9)$$

This matches the optical-scalar computation in the Xavier paper's appendix. The transverse directions do not focus the congruence per circuit; all of the interesting dynamics is in the null (t, z) block and in the transverse deviations of nearby geodesics, which we analyze in Secs. 4 and 5. (The mixed components Γ^t_{zx} and Γ^t_{zy} are nonzero and proportional to $f_{,x}$ and $f_{,y}$; they do not enter the transverse optical scalars.)

2.4 The locally flat option

If $f \rightarrow 0$, the metric (1) has Riemann tensor identically zero: the core is locally flat even though it is a quotient with nontrivial topology. This is the most benign member of the family and is useful as a validation limit. The flat core has no curvature-induced focusing and no local Hawking divergence; whatever global stress it carries is a Casimir-type effect of the identification. It is also the cleanest illustration of the fact that the interesting physics of the Xavier core is global, not local.

This locally flat $f \rightarrow 0$ limit is not the $h = f = 1$ limit used as the image-sum validation gate (g0). There h and f are both nonzero and equal, so ∂_z is null and every deck image is exactly null; here $f = 0$ and the z -circles are null only where h also vanishes. Both limits have vanishing curvature, but their causal character and their image sums are different.

3 Quotient legality

3.1 The identification and the invariance condition

A quotient is specified by an identification

$$(z, t) \sim (z + L, t + cL), \quad c \in \mathbb{R}. \quad (10)$$

We take f to be L -periodic in z throughout; the paper's saddle $f = (x^2 - y^2)/2$ is z -independent, so it is inert. For $h = P(z)Q(z - 2t)$, invariance requires

$$P(z + L)Q(z - 2t + L(1 - 2c)) = P(z)Q(z - 2t) \quad (11)$$

for all z, t . The two coordinate choices that matter are $c = 0$ (pure z -shift) and $c = 1/2$ (helical, which leaves $z - 2t$ invariant).

3.2 Pure z -shift: trains modulated by pumps

For $c = 0$, Eq. (11) reduces to $h(z + L, t) = h(z, t)$, or, with $w = z - 2t$,

$$P(z + L)Q(w + L) = P(z)Q(w). \quad (12)$$

Dividing by $P(z)Q(w + L)$ shows that $P(z + L)/P(z) = Q(w)/Q(w + L)$ is a constant, which we call e^{aL} with $a \in \mathbb{R}$. Hence the *general* legal pure- z profile is

$$h(z, t) = e^{2at}p(z)q(w), \quad w = z - 2t, \quad (13)$$

with p and q L -periodic; equivalently $P(z) = e^{az}p(z)$, $Q(w) = e^{-aw}q(w)$. The boost per circuit is

$$b = \ln \frac{P(z + L)}{P(z)} = aL. \quad (14)$$

For $a \neq 0$ with nonconstant p or q , this profile is z -dependent, is not a pump, and is not a train in the earlier sense (P and Q are not both periodic); it is a *train modulated by a pump* and it carries the Misner pile-up. The earlier two families are its special cases:

1. *Trains*: $a = 0$, so P and Q are both L -periodic. Then $h = p(z)q(w)$ is periodic in z and the quotient is legal. The closed generator is the orbit of ∂_z at fixed t at the horizon epoch. The boost per circuit is

$$b_{\text{train}} = \ln \frac{P(z + L)}{P(z)} = 0 \quad (15)$$

exactly, for every periodic P . This is the boost-degenerate subfamily.

2. *Pumps*: $p = q = 1$, so $h = e^{2at}$ is independent of z . The simplest example is

$$h(t) = e^{-t/\tau}, \quad (16)$$

which is the $\epsilon = 0$ member of the exponential family ($a = -1/2\tau$). The quotient is legal because h is z -independent. The closed generator is again the orbit of ∂_z at fixed t at the horizon epoch, but the profile now evolves in t ; the boost per circuit is

$$b_{\text{pump}} = -\frac{L}{2\tau} \neq 0 \quad (17)$$

for $h = e^{-t/\tau}$. This is the Misner-type family.

The Xavier paper's own worked examples are not in this list; they are the helical family discussed next. The train family is the boost-degenerate subfamily of the general pure- z family; the generic member and the pumps carry the pile-up.

3.3 The helical identification

For $c = 1/2$, the argument of Q in Eq. (11) is invariant, so the condition reduces to $P(z + L) = P(z)$: P must be L -periodic, while Q is *arbitrary*. This is a separate legal class. In the invariant coordinates

$$v = z - 2t, \quad w = z, \quad (18)$$

the identification is a pure w -translation, and the metric becomes

$$ds^2 = P(w)Q(v) dv dw - f(x, y, w) dw^2 + dx^2 + dy^2. \quad (19)$$

The closed null generator is the orbit of the identification vector

$$\xi = \partial_z + \frac{1}{2}\partial_t = \partial_w, \quad (20)$$

which is null where $f = 0$. On the axis $f = 0$ and $\partial_w f = 0$, the orbit is geodesic:

$$\Gamma^w_{ww} = \frac{P'(w)}{P(w)} \implies (P\dot{w})' = 0, \quad (21)$$

and the affine tangent rescaling per circuit is

$$\frac{P(w)}{P(w + L)} = 1 \quad (22)$$

exactly, because P is periodic. Thus the helical family is boost-degenerate for the whole family. Numerically, the helical sech^2 core gives a tangent rescaling 1.0000000000 per circuit, with no boost pair.

The helical family includes the paper's own worked examples: $\text{sech}^2(k(z - 2t))$ is $Q(v)$ with $P = 1$, and the degenerate exponential $e^{k\epsilon(z-2t)}$ is also $Q(v)$ with $P = 1$. Both are legal quotient spacetimes under $c = 1/2$. This is the mathematical content of the correction to the earlier dichotomy.

3.4 The detuned exponential and the paper’s boost formula

For the detuned exponential

$$h = e^{k(\epsilon z - \alpha t)}, \quad (23)$$

the combination $\epsilon z - \alpha t$ is invariant under the helical identification with

$$c = \frac{\epsilon}{\alpha}, \quad (24)$$

for any tuning. In invariant coordinates the boost per circuit is

$$b = \Gamma^w_{ww} L = \frac{k(2\epsilon - \alpha)L}{2}, \quad (25)$$

which is exactly the Xavier paper’s h_{boost} formula, their Eq. (10). For $\alpha \neq 2\epsilon$ the boost is nonzero and the generator carries the Misner pile-up; for $\alpha = 2\epsilon$ it is zero and the generator is marginal. The paper’s formula is therefore the holonomy of a helical identification that the paper never imposes. Under the paper’s pure z -shift the detuned exponential is not even a legal quotient for $\epsilon \neq 0$ (it shifts by $e^{k\epsilon L}$), so its published h_{boost} is a cover-space statement, not a quotient statement.

3.5 What descends to the paper’s quotient

The Xavier paper specifies a pure z -identification with t a global coordinate [1, §II and Appendix A]: $t, x, y \in \mathbb{R}$, z periodic with $0 \leq z \leq L$ and the endpoints identified. No t -shift is mentioned. In that quotient the CTCs are the closed curves of constant x, y, t —the z -circles—and the generator is the orbit of ∂_z at fixed t . This is the mechanism the paper uses.

Under that identification, the paper’s own profiles do not descend:

$$\left. \text{sech}^2(k(z - 2t)) \right|_{z \rightarrow z+L} - \text{sech}^2(k(z - 2t)) \neq 0, \quad e^{k(\epsilon(z+L) - \alpha t)} - e^{k(\epsilon z - \alpha t)} = e^{k\epsilon L} - 1 \neq 0. \quad (26)$$

The metric with the paper’s worked examples is therefore not invariant under the paper’s own identification. The cover spacetimes are unchanged—local curvature statements remain true—but the CTC mechanism, which is a statement about closed curves in the quotient, is not defined for those examples as written. The examples become legal only under the helical identification, whose global structure is different.

3.6 Why the helical cores are not switchable machines

The helical quotient is a legal spacetime, but it does not have the two properties that make the Xavier core interesting. First, t is only defined up to $L/2$ (or cL): the identification shifts the time coordinate, so there is no global time function on which the wave can evolve. The paper’s fixed- t z -circles are not closed curves in the quotient, and the switch mechanism—CTCs developing on a Cauchy surface S_0 when the wave arrives—is inexpressible. There is no off state.

Second, for the $c = 1/2$ identification—the one that contains the paper’s sech^2 and degenerate exponential—the identification vector $\xi = \partial_z + \frac{1}{2}\partial_t$ has norm

$$g(\xi, \xi) = -f. \quad (27)$$

It is timelike wherever $f > 0$. For the paper’s saddle $f = (x^2 - y^2)/2$, this is the entire $f > 0$ half-plane, at all “times”: the helical quotient has closed timelike curves through every point of that region, permanently. It is a Gödel-type geometry, not a chronology horizon with an off phase. There

is no compact regular initial-data story with an off state and no trigger. The detuned exponential under $c = \epsilon/\alpha$ has the same no-global-time obstruction; its boost is nonzero unless $\alpha = 2\epsilon$, and it is not a switchable machine either.

The flat helical core $f = 0$ is a useful check. For any P, Q the $f = 0$ metric (19) has Riemann tensor identically zero: $K = -(1/A)\partial_v\partial_w \ln A = 0$ with $A = P(w)Q(v)$. For the exponential profile, with $U = e^{kev}/(2k\epsilon)$, $V = w$, it is Minkowski space modulo a null translation—the flat null cylinder. It has a finite Casimir-type stress and no Hawking divergence, but it also has no f to cross and no switch. The flat helical core is benign but not a machine.

3.7 The classification

The legal quotients of the two identifications analysed here fall into two classes:

1. *Pure z -shift* ($c = 0$): the general legal family is (13), $h = e^{2at}p(z)q(z - 2t)$ with p, q L -periodic and boost $b = aL$. It is a train modulated by a pump; for $a \neq 0$ with nonconstant p or q it is z -dependent and carries the Misner pile-up. Its subfamilies are the *trains* ($a = 0$, P and Q both L -periodic, boost 0 exactly, marginal) and the *pumps* ($p = q = 1$, z -independent, boost $aL \neq 0$, pile-up).
2. *Helical cores* ($c = 1/2$): mixed identification, P periodic and Q arbitrary, boost 0 on the closed null generator, marginal, but no global time and not switchable.

The earlier list of three (pumps, trains, helical) was therefore not exhaustive: it omitted the generic modulated pure- z family. Nor is it a classification for all c : the invariance condition (11) admits further legal profiles when $c \notin \{0, 1/2\}$, for example $c = 1/4$ with $P = 1$ and $Q(w) = 1 + \frac{1}{2}\cos(4\pi w/L)$ (period $L/2$). We do not analyse these; because the identification shifts t , they are not switchable machines in the sense used here. A general- c classification is left open. Within every legal class, “degenerate boost” is equivalent to “no geometric pile-up.” The boost is

$$b = \begin{cases} aL, & \text{pure } z\text{-shift } (aL = 0: \text{trains}), \\ 0, & \text{helical.} \end{cases} \quad (28)$$

Our monodromy measurement reproduces the train and pump values ($aL = 0$ and $aL = -L/2\tau$) exactly and the helical value numerically to 10^{-10} .

Within the two identifications analysed, the machine-relevant legal cores are exactly the trains. The generic modulated pure- z family and the pumps carry the pile-up, and helical cores cannot be switched. The exact, globally identified, degenerate, switchable cores are the trains, i.e. the $a = 0$ subfamily. This is the sharpened form of the Xavier degeneracy: it is a whole subfamily, not a razor’s edge, but it is not the full $c = 0$ family and not the paper’s helical examples. Rieger’s isocausality and deck-equivariance classification of Misner-type spacetimes [10] is the closest external treatment of quotient legality; the classification above is specific to the Xavier metric and does not derive from it.

4 Monodromy and the pump/train pile-up

4.1 The winding-pile-up mechanism

The Hawking–Misner mechanism lives on the closed null generator $\mathcal{N}$. Let λ_k be the affine length of the k -th circuit of $\mathcal{N}$ and let T_k be the affine tangent at the reclosure. The boost per circuit is

core	tangent/circuit	length/circuit	generator
pump ($\tau = 10, L = 1$)	$\times e^{L/2\tau} = 1.0512711$	geometric shrink	incomplete (pile-up)
train (P, Q periodic)	$\times 1.0000000$	equal	complete (no pile-up)
helical sech^2	$\times 1.0000000$	equal	complete (no pile-up)

Table 1: One-circuit monodromy of the closed null generator. The pump’s tangent grows at the exact Misner rate; the train’s is marginal. The helical row is the corrected-classification addition: despite the expectation of a Misner-type boost, the helical identification acts by a null translation in the invariant frame and its closed generator is affinely complete.

core	λ_k/λ_{k-1}	$v_z/\text{circuit}$	λ_k
pump ($\tau = 10, L = 1$)	0.9512294245	$\times e^{L/2\tau} = 1.0512711$	geometric shrink
train (P, Q periodic)	1.0000000000	1.0000000000	2/3 exactly

Table 2: Measured affine circuit lengths over 24 successive circuits. The pump’s geometric law is exactly $e^{-L/2\tau}$; the train’s is exactly constant. The train value $\lambda = 2/3$ is at the starting point $z_* = 0$, where $P(z_*) = 3/2$; the general law is $\lambda_{\text{circ}} = 1/P(z_*)$ in units $L = 1$. At $z_* = 0.25$ used in the image sum, $P = 1$ and $\lambda = 1$.

$b = \ln |T_{k+1}/T_k|$. If λ_k shrinks geometrically, successive self-images of a quantum field crowd onto the generator exponentially; if the affine length is constant, the images sit at algebraic spacing and the geometric crowding channel is absent. The sign convention is the one used in the image-series analysis: for the pump, the affine length ratio is $e^{-L/2\tau}$ while the tangent ratio is $e^{+L/2\tau}$. The generic pure- z family (13) has $b = aL$ and hence the same pile-up for $a \neq 0$; only its $a = 0$ train subfamily is neutral.

4.2 The one-circuit spectrum

The exact geodesic flow and its variational equations were integrated for one circuit around the closed generator, for the pump $h = e^{-t/\tau}$ ($\tau = 10, L = 1$) and the train $h = P(z)Q(z - 2t)$ with

$$P(z) = 1 + \frac{1}{2} \cos \frac{2\pi z}{L}, \quad Q(w) = 1 + \frac{1}{2} \cos \frac{2\pi w}{L}, \quad (29)$$

and a transverse saddle $f = f_0(z) + A(x^2 - y^2)/2$, with f_0 tracking the axis profile at the horizon epoch so that the whole z -circle is null there. The monodromy spectrum is shown in Table 1.

The pump row is the known-divergent case, and its exact rate validates the machinery. The train row is the degenerate case. The helical row is the correction: the degenerate tuning is forced there too, but the helical quotient is not a machine (Sec. 3).

4.3 The 24-circuit measurement

The one-circuit spectrum is a linearized statement. The quantity the pile-up mechanism actually depends on is the affine circuit length over many successive circuits. The exact orbit was integrated circuit-by-circuit for $n = 1, \dots, 24$ (event $z = kL$, relative tolerance 10^{-11} , nonlinear, no linearization beyond the variational eigenvectors used to aim the probes). The results are in Table 2.

The pump cross-check is exact: the measured geometric law $e^{-0.05} = 0.9512294245\dots$ is the Misner-type crowding, demonstrated over 24 circuits rather than inferred from one. The train is exactly marginal. The train’s value $\lambda = 2/3$ is not a special tuning: it is $1/P(z_*)$ at $z_* = 0$, where $P = 1 + \frac{1}{2} \cos 0 = 3/2$; the per-circuit ratio is exactly 1 either way.

4.4 The stable-direction probe

A natural worry is that the train’s transverse saddle might supply a geometrically shrinking direction even though the boost is zero. Neighbors were launched along the transverse eigendirections of the one-circuit map and integrated exactly (nonlinear pairs, 24 circuits). For the train, the stable direction has $|\text{ev}| = 0.7865$ and the deviation shrinks at the measured rate $d \ln(\text{dev})/dk = -0.2402$; the unstable direction has $+0.2402$ (the symplectic partner, determinant 1). For the pump, there is also a shrinking direction, with $|\text{ev}| = 0.9512$ and rate -0.0507 , *including in the known-divergent case*.

The structural lesson is that a stable transverse direction is present in both cores, so its presence cannot be what distinguishes divergent from benign. The distinguishing observable measured here is the affine circuit length law: geometric versus exactly equal. This is the quantity the image spacing directly depends on, and it is the one we carry into the image-sum analysis.

4.5 What the monodromy settles

The monodromy settles the spacing premise. The train’s generator winds with exactly equal affine circuits, so the n -wound images sit at algebraic spacing, while the pump’s crowd at exactly $e^{-L/2\tau}$ per circuit. The Hawking pile-up mechanism, which lives on geometric shrinkage, has no geometric shrinkage to feed on in the train. What this does *not* settle is the full image sum: the transverse hyperbolicity enters the two-point function through the Van Vleck/focusing factors, and the sum over deck images must be computed with those factors included. That is the subject of Sec. 5.

5 The deck image sum with Hadamard weights

5.1 Construction

The image-series approximation to the two-point function on a quotient spacetime sums the Hadamard parametrix over the deck images of a point. Let P have coordinates $(t_*, x_*, y_*, z_*) = (0, \varepsilon, 0, 0.25)$, and let the deck images be

$$P_n = P + nL \hat{z}, \quad (30)$$

which is an exact isometry of both cores: for the train because the profile is L -periodic, for the pump because it is z -independent. For each n we solve the full 4D boundary-value problem for the connecting geodesic $P \rightarrow P_n$: the initial tangent is $v_0 = (v_t, v_x, 0, 1)$ and the endpoint must return all of (t, x, y) , not only x . The $y = 0$ plane is invariant, so this is a two-dimensional Newton shoot in (v_t, v_x) with the exact fixed- z endpoint Jacobian from the variational flow. The earlier one-dimensional shoot fixed $v_t = 0$ and imposed only $x_{\text{end}} = x_*$; it leaves $t_{\text{end}} \neq t_*$ and is a projection, not a deck-connecting geodesic. The corrected coefficient is derived in Sec. 7.

The world function is

$$\sigma_n = \frac{1}{2} \kappa \lambda_n^2, \quad \kappa = g_{\mu\nu}(P) v_0^\mu v_0^\nu, \quad (31)$$

with λ_n the total affine time at the z -event. This is Synge’s exact affine-segment identity, identical to the convention used in the Ori-core image-series analysis. The Van Vleck factor is

$$|\Delta_n|^{1/2} = \lambda_n^2 \left(\frac{|\det g(P_n)|}{|\det g(P)|} \right)^{1/4} |\det J_v|^{-1/2}, \quad (32)$$

where $J_v = \partial x_{\text{end}}/\partial v_0$ is the fixed-affine-time differential of the initial-velocity-to-endpoint map, taken from the exact 8×8 variational flow. The derivation is the standard one: for the affine-time-one map, $\sigma_{;\nu'} = -u_{\nu'}$ at the start, so $-\sigma_{;\mu\nu'} = (J_{\text{cov}}^{-1})^{\nu'}_{\mu}$ with $J_{\text{cov}} = J_u g'^{-1}$, $J_u = J_v/\lambda$, hence Eq. (32). In flat space $J_v = \lambda I$ and $|\Delta|^{1/2} = 1$ exactly.

We use two weights:

$$w_{1,n} = \frac{|\Delta_n|^{1/2}}{|\sigma_n|}, \quad w_{2,n} = \frac{|\Delta_n|^{1/2}}{\sigma_n^2}. \quad (33)$$

The first is the Hadamard σ^{-1} weight; the second is the σ^{-2} weight used in the Ori-core stress analysis. The directional $\cos \theta$ factor is $\mathcal{O}(1)$ here and is not the subject of this paper.

5.2 Validation gates

Four gates were passed.

- (g0) Null-circle limit ($h = f = 1$; $A = 0$, $\tau \rightarrow \infty$): $J_v = \lambda I$ exactly, so $|\Delta|^{1/2} = 1.000000000000$ for every n , and every image is exactly null ($\sigma = 0$). This is not the locally flat $f \rightarrow 0$ core of Sec. 2: here $h = f$ and ∂_z is null.
- (g1) Exact Misner reference: the machinery reproduces $\sigma_n = 4|T_0|(\cosh(nL/2) - 1)$ to a relative 3×10^{-12} for $n = 1, \dots, 4$.
- (g2) On-generator ($\varepsilon = 0$): the generator tangent hits P_n exactly (miss $< 10^{-9}$) with $\kappa = 0$ (null) and $\lambda_n = n \cdot 1.0000000000$ for all n .
- (g3) Pump cross-check: the off-generator σ -series terms decay like the train's, so the anticipated off-generator divergence does *not* appear; the pump divergence is instead measured on-generator, machine-exactly (Sec. 6). The gate passes in that corrected form.

Gate (g3) is itself a result: the naive expectation that the pump's off-generator image series would blow up geometrically is wrong. The pump's divergence is an on-generator phenomenon.

5.3 World functions: $\sigma_n = \varepsilon^2 s_n$

For the train, the deck-connecting world functions obey

$$\sigma_n = \varepsilon^2 s_n \quad (34)$$

to a maximum relative deviation 6.3×10^{-7} along the horizon family at $x = 0.2, 0.1, 0.05$ ($n \leq 4$); s_n is x -independent to that accuracy. The first values on the deck-connecting branch are shown in Table 3. Every image hugs the generator at $\mathcal{O}(\varepsilon^2)$.

The structure is still reproduced from a linear transverse problem, now evaluated on the geodesics that actually return to the deck image. The world functions are $\mathcal{O}(\varepsilon^2)$ and monotone in n : the deck-connecting branch has no near-null recurrences, unlike the one-dimensional projection. The transverse caustics of the connecting family remain at $n = 13.06k$ (π/ω with $\omega = 0.2404$ per circuit), but on this branch the Van Vleck weight is suppressed there rather than enhanced.

n	s_n	n	s_n
1	-0.0290	7	-1.880
2	-0.1177	8	-2.753
3	-0.2716	9	-4.060
4	-0.5010	10	-6.211
5	-0.8230	11	-10.40
6	-1.267	12	-22.29

Table 3: Measured $s_n = \sigma_n/\varepsilon^2$ on the deck-connecting (4D) branch at $\varepsilon = 0.1$, $n = 1, \dots, 12$, for the off- $\mathcal{N}$ horizon point of Sec. 7; the t -independence established there makes the values at $t = 0$ agree. The world functions are monotone. The first transverse caustic is at $n \approx 13.06$, where the primary weight is suppressed. The earlier one-dimensional projected branch has near-null recurrences and different s_n ; those are projection artifacts (Sec. 7).

5.4 Truncation at fixed ε

On the deck-connecting branch the weights decrease monotonically from $w_{1,1} = 3.45 \times 10^3$ to $w_{1,12} = 8.60$. The truncated sum is finite:

$$\begin{aligned} S_1(12) &= 5237.5, & C_1 &= \varepsilon^2 S_1(12) = 52.4, \\ S_2(12) &= 1.282 \times 10^7, & C_2 &= \varepsilon^4 S_2(12) = 1.282 \times 10^3, \end{aligned} \quad (35)$$

at the cutoff $n \leq 12$; at the smaller cutoff $n \leq 6$, $C_1 = 50.7139$ (Eq. 37). The last weight is 8.60 and the observed weight ratios (0.62–0.72) give a geometric tail $\sum_{n>12} w_{1,n} \approx 10$ –20: a few tenths of a percent of $S_1(12)$, but 23–47% of the last three included terms ($w_{1,10} + w_{1,11} + w_{1,12} = 42.7$). Since $C_1 = \varepsilon^2 S_1$ with $\varepsilon = 0.1$, the corresponding tail in C_1 is about 0.1–0.2; a conservative reading that the tail is comparable to the last included terms would leave C_1 uncertain at the several-percent level. Decreasing increments do not by themselves imply convergence ($w_n \sim 1/n$ is a counterexample): the truncated sum is finite, the tail is estimated, but convergence of the n -series is not proved. A converged coefficient requires summing past the $n \approx 13.06$ caustic with the Maslov dressing and a rigorous tail bound, neither of which is done here. These are the corrected deck-connecting values; the one-dimensional projection gives different ones (below).

5.5 The ε -scaling and the falsified ε -independence

The truncated sums are not ε -independent. Along the horizon family, $\sigma_n = \varepsilon^2 s_n$ gives the exact ε^{-2} and ε^{-4} powers,

$$S_1 \sim C_1/\varepsilon^2, \quad S_2 \sim C_2/\varepsilon^4. \quad (36)$$

The coefficients $\varepsilon^2 S_1$ measured at $\varepsilon = x = 0.2, 0.1, 0.05, 0.025$ (each at its own horizon $t_*(x)$, $n \leq 6$) are

$$50.713855, 50.713857, 50.713866, 50.713894, \quad \frac{d \ln S_1}{d \ln x} = -2.000000, \quad (37)$$

constant to $\mathcal{O}(\varepsilon^2)$. With $n \leq 12$ at $\varepsilon = 0.1$, $C_1 = 52.4$ and $C_2 = 1.282 \times 10^3$. The coefficient is cutoff-dependent: $C_1 = 50.71$ at $n \leq 6$ and 52.4 at $n \leq 12$, with the tail estimated but not bounded. The divergence is therefore the Misner-type power law with a cutoff-dependent coefficient; the earlier one-dimensional projection gave a fitted exponent -2.006 , a transient of the projected branch.

This falsifies the reading of the monodromy result as a sum that converges to an ε -independent finite value. The ε^{-2} (σ^{-1}) and ε^{-4} (σ^{-2}) behavior is the generic approach-to-a-closed-null-curve

light-cone divergence; it is present for the train, for the pump, and for exact Misner space. It is not a core discriminator.

For reference, the same weights give

$$C_1^{\text{Misner}} = 2.81, \quad C_2^{\text{Misner}} = 4.09. \quad (38)$$

The corrected deck-connecting train coefficient is $C_1 \simeq 52.4$ at the $n \leq 12$ cutoff (50.71 at $n \leq 6$; tail estimated, not bounded). Misner's $C_1 = 2.81$ is evaluated at $A = 0$, whereas the train has $A = 0.10$, and C_1 diverges as $A \rightarrow 0$ (the train's $A \rightarrow 0$ limit is the $h = f = 1$ null-circle limit, where all images are exactly null). The ratio $\simeq 18.6$ is therefore cutoff- and parameter-dependent, and it is largely set by the transverse saddle ($s_1 \approx -A/4$); we quote it only as an indicative number at this cutoff, not as a converged coefficient. The C_2 comparison is separate: $C_2 = 1.282 \times 10^3$ is about 313 times Misner's $C_2 = 4.09$; it does not share the C_1 ratio and is equally indicative. The train's deck images all sit at the same transverse distance ε from the generator (the deck is a pure z -translation), so every winding geodesic hugs the generator and every term is near-light-cone. The equal-affine-spacing premise does *not* make the two-point function's approach singularity milder than Misner's. The transverse saddle strength A sets the coefficient ($s_1 \approx -A/4$); the null-circle limit ($h = f = 1$, $A = 0$, $\tau \rightarrow \infty$) has all images exactly null and is infinitely singular. At the same indicative normalization, the train's $\mathcal{N}$ is more singular per unit distance than Misner's—the price of the switchable wave profile at the generator.

5.6 Why the focusing factors are load-bearing

In the $t \equiv 0$ class the transverse dynamics is exactly linear:

$$x'' = -\frac{A}{2}v_z^2(\lambda)x, \quad y'' = +\frac{A}{2}v_z^2(\lambda)y, \quad (39)$$

a confining rotation with monodromy pair $e^{\pm i0.2404}$ and a hyperbolic pair $e^{\pm 0.2402}$ (the measured $e^{\pm 0.24}$). The rotation creates the caustics of the connecting family; the hyperbolic partner enters $\det J_v$ with a growing weight $e^{+0.24n}$, which suppresses the Van Vleck factor $|\det J_v|^{-1/2}$; together with the growth of the world functions this makes the weights decrease, and it tames the weight near the caustics. The transverse hyperbolicity suppresses the weights and keeps the truncated sum finite; it is not a proof of convergence. The screen optical determinant on the generator is 0.99997, the symplectic “no net area growth” measured directly.

The 4D deck-connecting branch is the one the earlier one-dimensional projection called the “second branch” ($v_t \propto \varepsilon \neq 0$, a small t -excursion). The x -only projection is not a connecting geodesic: it leaves $t_{\text{end}} \neq t_*$ by an amount that grows with n (up to -2.97×10^{-1} at $n = 12$ at the off- $\mathcal{N}$ point), so its s_n and its $\sim 39\times$ coefficient are projection artifacts. The 4D branch reproduces the projected second branch's $s'_3 = -0.271612$ and $s'_8 = -2.752609$ to 5×10^{-7} and gives the corrected $C_1 \simeq 52.4$. All structural verdicts— ε^2 world functions, Van Vleck suppression, $\varepsilon^{-2}/\varepsilon^{-4}$ approach scaling—are unchanged.

5.7 Relation to the $\mathcal{N}$ -singularity

The ε in this section is the transverse distance to the closed null generator $\mathcal{N}$ at the horizon epoch: $P = (0, \varepsilon, 0, 0.25)$, and at $t_* = 0$ the null-circle locus is exactly the axis circle. The $\varepsilon^{-2}/\varepsilon^{-4}$ divergence is therefore the $\mathcal{N}$ -singularity—the measure-zero divergence every closed-null-curve machine carries, including the Ori 2005 core, where the return distance vanishes only on $\mathcal{N}$

and the program accepts it as benign. It is *not* the stress at horizon points off $\mathcal{N}$. This distinction is the subject of Sec. 7.

The correction stands against the literal “ ε -independent finite value” phrasing: nothing with a closed null geodesic can be ε -independent as $\mathcal{N}$ is approached. The truncated n -series at fixed ε is finite (measured), and the on-generator coefficients are bounded (measured, Sec. 6); but these statements do not control the approach at off- $\mathcal{N}$ horizon points. The truncated $\mathcal{N}$ coefficients are $C_1 \simeq 52.4$ at $n \leq 12$ (50.71 at $n \leq 6$; tail estimated, not bounded) and, separately, $C_2 = 1.282 \times 10^3$. The Misner comparison is at $A = 0$ while the train is $A = 0.10$, and C_1 diverges as $A \rightarrow 0$, so the ratios ~ 18.6 and ~ 313 are cutoff- and parameter-dependent indicative numbers, largely set by the saddle, not converged coefficients. They are the $\mathcal{N}$ -singularity strength, relevant only to back-reaction at $\mathcal{N}$ itself. The closed-null-geodesic census is not a reliable guide to the approach (the companion Ori analysis [9] found it superseded); it tells us where a closed null geodesic exists, not whether $\sigma_n \rightarrow 0$. The off- $\mathcal{N}$ image sum/parametrix is computed at one point in Sec. 7 and is finite there; the full off- $\mathcal{N}$ scan remains open.

6 The on-generator discriminator

6.1 Why the off-generator series cannot discriminate

Gates (g1) and (g3) and the ε -scaling make the point precisely: the off-generator image series has the same $\varepsilon^{-2}/\varepsilon^{-4}$ approach law for train, pump, and Misner. The difference between the train and the pump is not in the off-generator series. It must be sought where the two-point function is actually singular: on the generator.

At a point $P \in \mathcal{N}$, every winding returns to P with $\sigma = 0$, so the two-point function is $\delta(\sigma)$ -supported. The meaningful Hadamard data are the *winding coefficients*. We measure the screen optical Van Vleck coefficient

$$\Delta_{s,k}^{1/2} = |\det S_k|^{-1/2}, \quad S_k = \frac{\partial(x, y)_{\text{end}}}{\partial(v_x, v_y)_0} \Big|_{\text{event slice}}, \quad (40)$$

and the spacing-weighted term

$$w_{\text{on},k} = \frac{\Delta_{s,k}^{1/2}}{\lambda_k^2}. \quad (41)$$

The winding coefficient is $\Delta_{s,k}^{1/2}$; $w_{\text{on},k}$ is the term that enters the δ -comb sum.

6.2 Measured coefficients

The results are in Table 4.

The sums over the first 16 windings are

$$\sum_{k=1}^{16} w_{\text{on},k} \Big|_{\text{train}} = 16.000244, \quad \sum_{k=1}^{16} w_{\text{on},k} \Big|_{\text{pump}} = 66.7386. \quad (42)$$

The train’s first five coefficients are all 1.0000153; the pump’s are 1.0776, 1.252, 1.455, 1.690, 1.963. The pump’s screen coefficients grow at exactly $e^{+L/2\tau}$ per circuit, machine-exactly, and its spacing-weighted terms at $e^{+3L/2\tau}$; the train’s rates are zero to 10^{-12} .

core	λ_k per circuit	$\Delta_{s,k}^{1/2}$ per circuit	$d \ln w_{\text{on},k}/dk$
train	1.0000000000	1.000015	-4.8×10^{-12}
pump	$e^{-L/2\tau} = 0.9512294$	$e^{+L/2\tau} = 1.0512711$	+0.15000000

Table 4: On-generator winding coefficients over $k = 1, \dots, 16$ circuits. The train’s coefficients are bounded; the pump’s grow at the exact Misner rate. The pump’s spacing-weighted term grows at $e^{+3L/2\tau}$ per circuit: the affine crowding $e^{-L/2\tau}$ plus an equal-rate focusing growth. The measured log-rates for λ_k and $\Delta_{s,k}^{1/2}$ are $+1.6 \times 10^{-12}$ and -1.6×10^{-12} (train) and -0.05000000 and $+0.05000000$ (pump).

6.3 Interpretation: uniform comb versus geometric comb

The train’s winding terms form a *uniform comb* of bounded weight, with spacing $\lambda_1 = 1$ at $z_* = 0.25$. The comb is a regular distribution: its partial sum grows linearly with the winding cutoff, as it must on any complete closed null curve because there are infinitely many δ -terms at the same argument. The local weight density is bounded; this is the regular-comb case, but a finite renormalized stress is not established here. The coefficient sum still diverges with the winding cutoff, and no full renormalized stress tensor is computed here.

The pump’s comb *crowds geometrically*: the spacing shrinks by 0.9512 per circuit, the screen coefficients grow by $e^{+L/2\tau}$, and the spacing-weighted terms by $e^{+3L/2\tau}$, so the local weight density diverges. That is the Hawking divergence.

This is the genuine discriminator, and it is on-generator. The surviving content of the “quantum silence” claim is: *the train’s winding coefficients are bounded; the pump’s grow at the Misner rate*. “Silent” can therefore only mean *no geometric pile-up*, not a finite on-generator two-point function. The coefficient sum still diverges with the winding cutoff on any complete closed null curve; the question is whether the divergence is the regular comb divergence (whose removal by standard renormalization is not established here) or the geometric weight-density divergence (the Hawking mechanism).

7 Off- $\mathcal{N}$ image sum and parametrix at one point

The image-sum computation of Sec. 5 approaches the generator $\mathcal{N}$ at the horizon epoch $t_* = 0$. The horizon at other epochs is not the axis circle. Because the wave profile shifts, the locus $h = f$ at $t_* \neq 0$ contains points with $x \neq 0$. At such a point $\partial_x f \neq 0$, so the fixed- t z -circle is *not* geodesic, and no closed null geodesic was found through it. The closed-null-geodesic census is not a reliable guide to the approach (the companion Ori analysis [9] found it superseded); it tells us where a closed null geodesic exists, not whether $\sigma_n \rightarrow 0$. The returning-distance heuristic therefore does not predict bounded σ_n here. The off- $\mathcal{N}$ image sum is a different quantity from the $\mathcal{N}$ -singularity, and it is the one that decides whether the train is physically benign at the horizon.

We computed it at one point with the full 4D deck-connecting boundary-value problem. The profile is the train $P(u) = Q(u) = 1 + \frac{1}{2} \cos(2\pi u/L)$ with $f = P(z)^2 + (A/2)(x^2 - y^2)$ and $(A, L, \tau) = (0.10, 1, 10)$, at

$$x_* = 0.10, \quad y_* = 0, \quad z_* = 0.25, \quad t_* = 0.2499204225\dots, \quad (43)$$

where $Q(w_*) = P(z_*) + Ax_*^2/(2P(z_*))$ fixes $w_* = -0.2498408450\dots$ and $t_* = (z_* - w_*)/2$. At this point $h = f = 1.0005$ and $g_{zz} = 0$, but the circle is not geodesic ($\Gamma_{zz}^x = Ax_*/2 = 5 \times 10^{-3}$, $\Gamma_{zz}^z = P'/P = -\pi$).

For each deck image $P_n = P + nL\hat{z}$ we solve the exact affine geodesic and its 8×8 variational flow, imposing the return of all of (t, x, y) : $t_{\text{end}} = t_*$, $x_{\text{end}} = x_*$, $y_{\text{end}} = 0$. The t -return condition is essential. The earlier one-dimensional shoot of Sec. 5 fixed $v_t = 0$ and imposed only $x_{\text{end}} = x_*$; it leaves $t_{\text{end}} \neq t_*$ by an amount that grows with n (from -4.7×10^{-5} at $n = 1$ to -2.97×10^{-1} at $n = 12$), so it is not a deck-connecting geodesic. The genuine 4D branch is the one the earlier image-sum note called the second branch ($v_t \neq 0$); it reproduces that branch's $s'_3 = -0.271612$ and $s'_8 = -2.752609$ to 5×10^{-7} .

Five validation gates pass:

- (v0) On-generator $\mathcal{N}$ limit ($x = 0, t = 0$): the tangent $(0, 0, 0, 1)$ hits P_n at machine precision with $\kappa = 0$ and $\lambda_n = n$.
- (v1) t -independence: σ_n and Δ_n at t_* agree with their $t = 0$ values to 2.1×10^{-7} and 3×10^{-9} for $n = 1, \dots, 4$; moving off the horizon by δt (with $|h - f|$ up to 0.294) leaves them unchanged to the same accuracy.
- (v2) $\sigma_n = x^2 s_n$: σ_n/x^2 is independent of $x = 0.2, 0.1, 0.05$ along the horizon family to 6.3×10^{-7} .
- (v3) Second-branch cross-check: $s_3 = -0.2716115$ and $s_8 = -2.7526086$ reproduce the earlier projected second branch ($s'_3 = -0.271612$, $s'_8 = -2.752609$) to 5×10^{-7} .
- (v4) Projection diagnostic: the one-dimensional shoot is reproduced together with its nonzero t -miss; its near-null recurrences do not appear on the 4D branch.

On the primary (generator-connected) branch, $n = 1, \dots, 12$ before the first transverse caustic at $n \simeq 13.06$, the world functions are monotone and the weights decrease from 3.45×10^3 to 8.60. The partial sums are

$$\begin{aligned} S_1(12) &= 5237.5, & C_1 &= x_*^2 S_1(12) = 52.4, \\ S_2(12) &= 1.282 \times 10^7, & C_2 &= x_*^4 S_2(12) = 1.282 \times 10^3. \end{aligned} \quad (44)$$

The last weight is 8.60; as in Sec. 5 the observed weight ratios give a geometric tail of about 10–20 in the weight sum, a few tenths of a percent of $S_1(12)$ (and 23–47% of the last three included terms). The truncated off- $\mathcal{N}$ image sum is therefore finite at this point; the tail is estimated, not bounded, and convergence is not proved.

Two scalings characterize the result. First, at fixed $x = 0.1$ the horizon approach is $\mathcal{O}(1)$ (degree zero in the horizon distance): varying t off the horizon (with $|h - f|$ up to 0.294) changes σ_n by at most 2.1×10^{-7} and Δ_n by at most 3×10^{-9} . This is the analytic t -independence of the train structure, not a numerical accident; the horizon condition only selects t_* , and the image sum is blind to it. Second, along the horizon family $\sigma_n = x^2 s_n$ exactly, so $S_1 \sim C_1/x^2$; measured at $x = 0.2, 0.1, 0.05, 0.025$ (each at its own $t_*(x)$, $n \leq 6$),

$$x^2 S_1 = 50.713855, 50.713857, 50.713866, 50.713894, \quad \frac{d \ln S_1}{d \ln x} = -2.000000. \quad (45)$$

The near-null enhancement is therefore real but lives in the transverse coordinate: the off- $\mathcal{N}$ image sum is finite at $x \neq 0$ and diverges only on the axis $x = 0$, where closed null geodesics exist. Its coefficient is the truncated $\mathcal{N}$ coefficient $C_1 \simeq 52.4$ at $n \leq 12$ (50.71 at $n \leq 6$; tail estimated, not bounded), not the projected $\sim 39 \times / 111$ value. The Misner comparison is at $A = 0$ while the train is $A = 0.10$ and C_1 diverges as $A \rightarrow 0$; the ratio ~ 18.6 is cutoff- and parameter-dependent and largely set by the saddle, so it is indicative only.

The result is a parametrix-level statement with the following caveats. It is one point, not a scan; only the primary (generator-connected) branch is summed, and only $n \leq 12$, before the first transverse caustic, with the tail estimated but not bounded; the higher transverse branches are not summed; no subtraction is performed, so this is not a renormalized $\langle T_{\mu\nu} \rangle$ and no back-reaction is computed; and no classical stability claim is made in this section (the linearized analysis is in Sec. 8). The full off- $\mathcal{N}$ scan, the higher branches, the Maslov/caustic dressing, and the full renormalized tensor remain open (Sec. 10).

The companion analysis of the Ori 2005 core [9] reports a finite parametrix-level off- $\mathcal{N}$ image sum at its sampled chronology-horizon point; we do not rely on its numerical normalization here. The present train result is finite at the truncation level at $x_* = 0.1$, the same qualitative parametrix-level outcome, but we do not attach a physical stress normalization to it here.

8 Classical stability: analyzed sectors

The linearized stability of the train is now computed in several tractable sectors at the closed null generator: the TT tensor sector at $k_\perp = 0$ and $k_\perp \neq 0$, the vector truncation, the leading off-axis $\mathcal{O}(A\rho)$ TT problem, the transversely homogeneous profile/scalar h -mode, and the f -modulus. In each analyzed sector the Floquet spectrum is neutral. The per-circuit multipliers lie on the unit circle (the deck multipliers are 1 for the reference train) and the measured per-circuit log-rates are 0.0. The computation is the linearization of the exact vacuum Ricci tensor about the periodic train; its symbolic and finite-difference validation gates are listed in the reproducibility appendix.

TT tensor sector. For the perturbations $\delta g_{xx} = a$, $\delta g_{yy} = -a$, $\delta g_{xy} = b$, with the remaining components zero, the linearized equations at the axis are

$$\delta R_{zz} = -Aa, \quad \delta R_{xy} = -\frac{2}{Q^2} \left(b_{,vv} - \frac{Q'}{Q} b_{,v} + \frac{Q}{P} b_{,vw} \right), \quad (46)$$

where $v = z - 2t$ and $w = z$. The first equation forces $a = 0$ on the axis for any $A > 0$: the $h_{xx} - h_{yy}$ polarization is *expelled from the generator* by the transverse curvature, the same mechanism as the Ori A4 A -mode exclusion. Only at $A = 0$ is it a second neutral TT mode. The second equation is $\square_2 b = 0$ on the flat null cylinder $P(w) dW dw$, with $dW = Q dv - P dw$, and its general solution is

$$b = F(W) + G(w). \quad (47)$$

The h_{xy} mode is therefore a two-dimensional massless scalar; it is independent of A at the axis and its neutrality is not tuned to the reference value $A = 0.10$. The deck map $(t, z) \rightarrow (t, z + L)$ leaves it with multiplier 1, and the time isometry $(t, z) \rightarrow (t + L/2, z)$ multiplies a Fourier mode e^{ikW} by e^{-ikL} , again of unit modulus.

TT at $k_\perp \neq 0$. In the light-cone gauge $h_{u\mu} = 0$ with trace $\eta = h_{xx} + h_{yy} = 0$, the two tensor polarizations $S = h_{xx} - h_{yy}$ and $B = h_{xy}$ obey the common principal equation

$$L_k[h] \equiv h_{,uu} - \frac{Q'}{Q} h_{,u} + \frac{Q}{P} h_{,uz} - \frac{Q^2 k_\perp^2}{4} h = 0. \quad (48)$$

In the null coordinates $dW = Q du - P dz$ this is $h_{,Wz} = (Pk_\perp^2/4)h$, with the exact chirped solution

$$h = e^{i\nu W} \exp \left[-\frac{ik_\perp^2 P_{\text{int}}(z)}{4\nu} \right], \quad P'_{\text{int}} = P, \quad (49)$$

and deck multiplier $e^{-ik_\perp^2 L/(4\nu)}$, of unit modulus. The $k_\perp$ correction is a pure phase, not an amplifier, and the neutral B mode survives it. The $h_{xx} - h_{yy}$ expulsion is a $k_\perp = 0$ statement: at $k_\perp \neq 0$ the curvature term $-AS/2$ in δR_{zz} is absorbed by $(k_\perp^2/2)E$, so $E \sim AS/k_\perp^2$. The expulsion is therefore not gauge-invariant, but E is constraint-determined and no new propagating tensor mode appears at this order.

Vector sector. For $\delta g_{ui} = V_i$, $\delta g_{zi} = W_i$ ($i = x, y$) at $k_\perp = 0$, the two polarizations decouple and the linearized equations are first order in $D_i = V_{i,z} - W_{i,u}$. The general solution is the constraint family $D_i = Q(u)G_i(z)$ with G_i an arbitrary periodic function, V_i algebraically fixed, and $W_i = Q_{\text{int}}(u)K_i(z) + w_i(z)$, where $w_i(z)$ is pure gauge. The deck multiplier is 1; the apparent linear growth of W_i is a longitudinal response whose deck shift depends on z only and is therefore pure gauge. There is no second-order propagating vector mode and no per-circuit vector amplifier in this truncation.

Off-axis $\mathcal{O}(A\rho)$. For pure TT at $k_\perp = 0$, the principal equations are unchanged at $\mathcal{O}(A\rho)$; the first correction to the principal operator is $\mathcal{O}(A\rho^2)$. The mixed components

$$\delta R_{xz} = \frac{A(xa_{,u} - yb_{,u})}{PQ}, \quad \delta R_{yz} = \frac{A(xb_{,u} + ya_{,u})}{PQ} \quad (50)$$

do not vanish away from the axis, so pure-TT localization fails at $\mathcal{O}(A\rho)$ and the longitudinal components $C = h_{xz}$, $D = h_{yz}$, $E = h_{zz}$ are forced. This is gauge relaxation, not a new leading-order amplifier. The cross term $\mathcal{O}(A\rho k_\perp)$ is not computed.

Profile and f -modulus. The profile perturbation $h \rightarrow h + \phi$, i.e. $\delta g_{tz} = -\phi$ and $\delta g_{zz} = +\phi$, has the general solution

$$\phi = P(z)B(v) + Q(v)H(z), \quad (51)$$

which is the linearization of the vacuum family $P \rightarrow P + H$, $Q \rightarrow Q + B$. It is an in-family modulus, not a gauge artifact, and its multipliers are 1 (deck) and e^{-ikL} (time isometry). The f -modulus, $\delta g_{zz} = \Phi$, has the general solution $\Phi = f(z)q(v) + g(z)$ with $q'(v) = Q(v)$. Under $v \rightarrow v + L$ one has $q(v + L) = q(v) + L\langle Q \rangle$, so the $f(z)q(v)$ branch is not single-valued and is excluded; the surviving periodic $g(z)$ branch is neutral. In all three sectors there is no exponential per-circuit amplification of the metric.

Geodesic saddle as ejection. The transverse saddle is not a metric mode. In the $t \equiv 0$ class the transverse equations are Eq. (39): a rotation in x with one-circuit pair $e^{\pm i0.2404}$ and a hyperbolic pair in y with $e^{\pm 0.2402}$. Both maps have determinant 1. The hyperbolic pair is the Jacobi field of the geodesic congruence: nearby null geodesics converge exponentially along the stable direction and are ejected along the unstable one. It is a symplectic $\det = 1$ shape deformation, not a growing perturbation of the metric, and it is not registered by the area factor f_{area} .

The $f_{\text{area}} = 2h_{\text{boost}}$ criterion. For the train, $h_{\text{boost}} = \ln[P(z + L)/P(z)] = 0$ exactly because P is periodic, and $f_{\text{area}} = 0$ at the axis (Eq. 9), so Xavier's criterion $f_{\text{area}} > 2h_{\text{boost}}$ reads $0 > 0$: it is *marginal*. The criterion is a semiclassical optical balance between area focusing and boost; it is not a linearized-Einstein mode-stability test, and a marginal value can certify neither stability nor instability. The direct mode computation above is the stronger statement for the sectors it covers.

What is and is not established. The train is classically neutral in the analyzed TT (including the $k_\perp$ phase and the leading off-axis problem), profile, f -modulus, and vector-truncation sectors: no per-circuit metric-perturbation amplifier was found there. This is not a general stability theorem, and several sectors remain explicitly open.

1. *Full ten-component system:* the gauge-complete linearized system at $k_\perp \neq 0$, including the vector-tensor mixing found in the $k_\perp$ vector equations, a complete gauge choice, and a constraint/Bianchi reduction.
2. *Longitudinal/constraint propagation:* the exact C, D, E system, the integration constants $c(z), d(z), e(z)$, the non-elementary E integral, and Bianchi consistency.
3. *Off-axis beyond leading order:* $\mathcal{O}(A\rho k_\perp)$, $\mathcal{O}(A\rho^2)$, and the full off-axis gauge-complete system.
4. *Quotient Cauchy problem:* the quotient is not globally hyperbolic (the horizon is a chronology horizon); the modes were solved on the cover by characteristics and by the deck/time isometries, and no initial-value problem on the quotient was posed.
5. *Nonlinear stability:* untouched.

The numerical rates are computed at one profile ($P = Q$, $A = 0.10$, $L = 1$); the neutrality structure is profile-independent for periodic P, Q , but the coefficients are for this profile. A scan of the Floquet spectrum over the parameter point is not performed.

Related gauge-complete check (Ori background). A companion analysis of the gauge-complete linearized problem on the Ori 2005 core—not the Xavier train—finds no per-wrap exponential amplifier in any sector. There the full ten-component system in the light-cone gauge $h_{T\mu} = 0$ reduces, after the trace mode $h_{xx} + h_{yy}$ is eliminated as a transverse dilation, to a common principal equation for $S = h_{xx} - h_{yy}$ and $B = h_{xy}$,

$$T h'' + (1 - 2ik_z)h' - k_\perp^2 h = 0, \quad (52)$$

with T the Ori null coordinate, exact Bessel solutions, and indicial exponents $\{0, 2ik_z\}$. At $k_\perp = 0$ it recovers the TT result (B neutral, S expelled for positive transverse curvature). Because ∂_z is Killing, every single-valued mode on $z \sim z + L$ has $k_z = 2\pi n/L$ and multiplier $e^{ik_z L} = 1$, so the absence of a per-wrap amplifier is structural rather than mode-by-mode. The companion quotient-Cauchy analysis of the same Ori system shows that the natural $T = \text{const}$ slice is spacelike only for $T < f$, that the quotient is not globally hyperbolic, and that the exact constraints propagate by the contracted Bianchi identity. It also identifies an $\mathcal{O}(a)$ obstruction to the frozen-WKB core truncation: the $\nu = z$ Bianchi component is $\Delta_z = -aS'/2$, so the truncated system forces $aS' = 0$, while the $B = h_{xy}$ sector remains a clean propagating mode obeying Eq. (52). This is a methodological companion, not a result for the Xavier train: it does not establish Xavier-train stability, and the Xavier full ten-component, longitudinal/constraint, $\mathcal{O}(A\rho k_\perp)$, $\mathcal{O}(A\rho^2)$, quotient-Cauchy, and nonlinear sectors remain open.

Within these limits, the correct summary is: the train is classically neutral (marginally stable) in the analyzed sectors, not unstable there and not proven stable in general.

9 Discussion

9.1 Mechanism inapplicable versus mechanism evaded

For the Ori 2005 core, the Hawking mechanism *fires* on approach to the non-compactly generated horizon and is then cut off by the horizon itself. For the Xavier train, the mechanism is not merely evaded or cut off: the geometric pile-up channel is absent by construction. The affine circuit length is exactly constant, the on-generator winding coefficients are bounded, and the boost is zero for the whole train subfamily. Within the two identifications analysed, this is forced by periodicity, not achieved by tuning. The degeneracy is a whole subfamily, not a razor’s edge.

The price is visible in the same computation. At the same indicative normalization, the train’s generator is more singular per unit distance than Misner’s: the deck images all hug the generator. The truncated $\mathcal{N}$ -singularity coefficient is $C_1 \simeq 52.4$ at $n \leq 12$ (50.71 at $n \leq 6$; tail estimated, not bounded). Misner’s $C_1 = 2.81$ is evaluated at $A = 0$ whereas the train has $A = 0.10$, and C_1 diverges as $A \rightarrow 0$, so the ratio ~ 18.6 is cutoff- and parameter-dependent and largely set by the saddle; it is indicative only. The earlier $\sim 39\times/111$ number was the one-dimensional projection, not the deck-connecting result; the C_2 ratio is separate (1.282×10^3 versus Misner’s 4.09). The off- $\mathcal{N}$ stress is finite at the computed horizon point; the train is classically neutral in the analyzed TT (including the $k_\perp$ phase and the leading off-axis problem), profile, f -modulus, and vector-truncation sectors, with the full ten-component, longitudinal/constraint, $\mathcal{O}(A\rho k_\perp)$, $\mathcal{O}(A\rho^2)$, quotient-Cauchy, and nonlinear sectors open. “Quantum silence” is therefore a statement about the *mechanism*, not about the full stress tensor: the train has no geometric pile-up, but it does not have a finite stress at $\mathcal{N}$.

9.2 Which cores are machines

Within the two identifications analysed, the classification separates the legal spacetimes from the machine-relevant ones. The generic modulated pure- z family and the pumps are legal and carry the Misner divergence. Helical cores are legal and boost-degenerate, but they have no global time, no off state, and permanent CTCs where $f > 0$; they are not switchable machines. Trains ($a = 0$) are legal, boost-degenerate, switchable, and globally identified. They are the only machine-relevant family among the two identifications analysed.

The paper’s own worked examples— sech^2 and the degenerate exponential—fall in the helical family, not the train family, when made legal. This is the sharpened correction: the mathematics of the examples is right, but the global structure they require is not the one the paper’s switch mechanism uses. The switchable cores are the periodic trains, and their profiles are periodic in both P and Q .

9.3 Relation to the Ori core and to the machine program

The Xavier train and the Ori 2005 core are complementary. Ori’s core has permanent CTCs, a cut-off stress at the chronology horizon, and a well-developed envelope and stability analysis; it is the better understood object. Xavier’s train is switchable and has no geometric pile-up, and its off- $\mathcal{N}$ image sum is finite at the one horizon point computed; it is classically neutral in the analyzed TT, profile, f -modulus, and vector-truncation sectors, with the full ten-component, longitudinal/constraint, $\mathcal{O}(A\rho k_\perp)$, $\mathcal{O}(A\rho^2)$, quotient-Cauchy, and nonlinear sectors open. It is the more attractive machine core if those open sectors close favorably.

The identification is still assumed. This paper classifies the quotients of the two identifications analysed and computes the semiclassical image series on them; it does not construct the identifica-

tion, estimate the ignition energy, or propose a laboratory realization. The topology is global, and the question of how a boost identification could be acquired is outside this analysis. The switch is a wave profile on a fixed identification, not a local manipulation of the topology.

A quiet train relic would be difficult to see: with $f \approx 0$ the core is locally flat, and the lensing signature is degenerate. The detectable epoch is the trigger, when the wave profile turns the circles timelike. The signature computation is open and is not attempted here.

9.4 Limitations

The quantitative results use the core parameters $(A, L, \tau) = (0.10, 1, 10)$ at single points, with $t_* = 0$ for the on-generator image sum and $t_* = 0.2499204$ for the off- $\mathcal{N}$ point; the image sum uses $z_* = 0.25$ and the 24-circuit scan starts at $z_* = 0$. The structural laws ($\sigma_n = \varepsilon^2 s_n$, the caustics at π/ω per cycle, the exact affine circuit ratios, the on-generator rates) are mechanism-derived and are not expected to change qualitatively, but the coefficients are point values. The ε^{-2} coefficient is dominated by the first few images ($n = 1, 2, 3$); a parameter point placing a caustic or otherwise changing the near-null weights would change it at the $\mathcal{O}(1)$ level.

The image sum counts the winding (generator-attached) geodesic family only. Other connecting geodesics (bulk chords) have $\mathcal{O}(1)$ σ_n and join the series with bounded, ε -independent terms; they do not affect the ε -scaling. The Van Vleck convention was validated in flat space and exact Misner space, both machine-exactly; the 4×4 detector definition used here is the standard one. The off- $\mathcal{N}$ image sum/parametrix is computed at one point (Sec. 7); a full off- $\mathcal{N}$ scan is not computed, and the stability result is restricted to the analyzed TT (including the $k_\perp$ phase and the leading off-axis problem), profile, f -modulus, and vector-truncation sectors (Sec. 8); the full ten-component, longitudinal/constraint, $\mathcal{O}(A\rho k_\perp)$, $\mathcal{O}(A\rho^2)$, quotient-Cauchy, and nonlinear sectors are not computed. No claim is made about the full renormalized stress tensor.

10 Open problems

Five problems are posed precisely by the present analysis. The first four are the items recorded in the image-sum note; the fifth records the completion of the stability analysis.

1. *Renormalized stress from the parametrix.* The train's ε^{-4} approach scaling matches Hawking's Misner law, so chronology protection at the stress level is not discharged by the spacing argument. The off- $\mathcal{N}$ image sum is now computed at one point (Sec. 7) and is finite; what remains is to build $\langle T_{\mu\nu} \rangle$ from the computed Hadamard parametrix with the standard renormalization, to extend the off- $\mathcal{N}$ evaluation to a scan and to the higher transverse branches, and to determine whether the renormalized stress is finite and small.
2. *Maslov/caustic dressing.* The windings past $n \approx 13.06k$ cross caustics of the connecting family. The caustics are benign for the truncated sum (the Van Vleck suppression keeps it finite), but a full parametrix requires the Maslov phase and the caustic dressing. This is a technical but necessary completion.
3. *The t -direction approach.* The present ε is the transverse/spatial distance to $\mathcal{N}$ at the horizon epoch. The approach to the generator in the timelike direction is a different limit and has not been computed. It is the natural counterpart of the near-horizon limit in the Ori-core analysis.

4. *Parameter scan.* The measurements use $(A, L, \tau, z_*) = (0.10, 1, 10, 0.25)$ and the off- $\mathcal{N}$ point $(x_*, t_*) = (0.1, 0.2499204)$. A scan in A, L, τ, z_* , and the horizon point is needed to confirm that the coefficients (not the structural laws) vary as expected and to identify parameter points where the near-null weights are enhanced.
5. *Completion of the stability analysis.* The linearized analysis now establishes neutrality in the TT h_{xy} mode (both $k_\perp = 0$ and $k_\perp \neq 0$), the leading off-axis $\mathcal{O}(A\rho)$ TT problem, the profile h -mode, the periodic f -modulus, and the vector truncation; the $h_{xx} - h_{yy}$ polarization is expelled at $k_\perp = 0$ and absorbed by E at $k_\perp \neq 0$ (Sec. 8). The full ten-component system, the longitudinal/constraint propagation, the $\mathcal{O}(A\rho k_\perp)$ cross terms and higher off-axis orders, the quotient Cauchy problem, and nonlinear stability remain open, as does a scan of the Floquet spectrum over the parameter point. The companion gauge-complete and quotient-Cauchy checks on the Ori background (Sec. 8) narrow the methodological question—the 10-component reduction, the absence of per-wrap amplifiers, and the Bianchi propagation of the constraints can be carried through there—but they are not Xavier-train computations and do not close these Xavier sectors. No general stability theorem is claimed.

The deck-connecting recalculation also replaces the projected $\sim 39\times$ Misner ($C_1 \sim 111$) by $C_1 \simeq 52.4$ at $n \leq 12$ (50.71 at $n \leq 6$; tail estimated, not bounded), with $C_2 = 1.282 \times 10^3$ compared separately to Misner's $C_2 = 4.09$. The ratios are cutoff- and A -dependent and indicative only. The state-choice problem is also open: Melton, Strominger and Wang construct an invariant vacuum and Fock space on the time-periodic CTC quotient AdS_3/Z [11], the closest external construction, but that is a different quotient, so an adapted state for the Xavier train is not yet supplied. In addition, the acquisition/ignition questions of the machine program are outside the scope of this paper.

11 Conclusion

Among the two identifications analysed here ($c = 0$ and $c = 1/2$), the Xavier metric admits a genuine switchable quotient—the trains, the $a = 0$ subfamily of the pure- z -shift class—and its degeneracy is a whole subfamily, not a tuning. But the global structure is more constrained than either the original dichotomy or the first correction suggested. For pure z -shift the general legal profile is $h = e^{2at}p(z)q(z - 2t)$, a train modulated by a pump with boost $b = aL$ and the Misner pile-up for $a \neq 0$; the trains are its boost-degenerate subfamily and the pumps its z -independent subfamily. The helical identification ($c = 1/2$) is a separate legal class; for general c further legal profiles exist (for example $c = 1/4$), which we do not classify and which are not switchable machines because the identification shifts t . The paper's own sech^2 and exponential examples do not descend to the pure z -quotient the paper imposes; they are legal only under a helical identification, which has no global time and is not a switchable machine. Within the two identifications analysed, the machine-relevant legal cores are the trains.

For the trains, the monodromy is exactly marginal: the affine circuit length is constant over 24 circuits, and the on-generator winding coefficients are bounded. The pump's affine length shrinks at $e^{-L/2\tau}$, its screen coefficients grow at $e^{+L/2\tau}$, and its spacing-weighted terms grow at $e^{+3L/2\tau}$. The truncated primary-branch deck image sum ($n \leq 12$) with Hadamard and Van Vleck weights is finite at fixed ε but diverges as ε^{-2} (σ^{-1}) and ε^{-4} (σ^{-2}) as the generator is approached; this is the $\mathcal{N}$ -singularity, common to Misner, pump, and train, and it is not the discriminator. The discriminator is on-generator: the train's spacing-weighted coefficients are bounded, while the pump's screen coefficients and spacing-weighted terms grow at the Misner rates quoted above.

The correct statement of “quantum silence” is therefore: the geometric pile-up channel is absent for the train family. It is not a statement that the stress is finite at the generator. The off- $\mathcal{N}$ image sum/parametrix is computed at one point ($x_* = 0.1$, $t_* = 0.2499204$) and is finite there, with the divergence confined to the axis; the train is classically neutral in the analyzed TT (including the $k_\perp$ phase and the leading off-axis problem), profile, f -modulus, and vector-truncation sectors, while the full ten-component, longitudinal/constraint, $\mathcal{O}(A\rho k_\perp)$, $\mathcal{O}(A\rho^2)$, quotient-Cauchy, and nonlinear sectors remain open. A full off- $\mathcal{N}$ scan and the full renormalized tensor remain open; together with those open stability sectors, they decide whether the train is a physically viable machine core. The deck-connecting calculation also replaces the projected $\sim 39\times$ Misner by $C_1 \simeq 52.4$ at $n \leq 12$ (50.71 at $n \leq 6$; tail estimated, not bounded), with $C_2 = 1.282 \times 10^3$ separately about 313 times Misner’s $C_2 = 4.09$; the ratios are cutoff- and A -dependent and indicative only. The quotient structure of the two identifications analysed, the monodromy, and the on-generator discriminator are settled, and they are the results this paper reports.

Provenance and reproducibility

Every quantitative statement from this paper’s computations is reproduced by the scripts `simulations/xavier_core.py`, `xavier_monodromy.py`, `xavier_legality.py`, `xavier_image_series.py`, `xavier_image_sum.py`, `xavier_image_sum_analysis.py`, `xavier_offn_stress.py`, `xavier_stability.py`, `xavier_vector_offaxis.py`, and `ori_gauge_complete.py`; the results and validation records are included with the manuscript for review (Ref. [12]); no public repository URL or DOI is assigned. The off- $\mathcal{N}$ result of Sec. 7 is reproduced by `simulations/xavier_offn_stress.py`; the full run takes about four minutes, `--quick` uses $n \leq 6$, and both write `simulations/xavier_offn_stress.png` (world functions 4D versus the 1D projection, Hadamard weights, partial sums and the x -family, near-axis log-log law). The companion Ori analysis [9] is cited only qualitatively; its numerical normalization is not used here. The linearized stability result of Sec. 8 is reproduced by `simulations/xavier_stability.py`; the run takes about 12 s and writes `simulations/xavier_stability.png` (amplitude versus circuit, the Floquet spectrum, finite-difference validation residuals, and the transverse Jacobi integration). The script uses a five-point finite-difference linearization of the exact nonlinear Ricci tensor and passes the following gates: the symbolic linearized Ricci factorization, the exact-vacuum and $f = 0$ (locally flat) limits, a positive control (a deliberately nonzero perturbation, with $\max|\delta R| > 0.1$), and the one-circuit Jacobi map (eigenvalues with determinant 1). The analyzed sectors have measured residuals below 10^{-5} . The companion gauge-complete check quoted in Sec. 8 is reproduced separately by `simulations/ori_gauge_complete.py`; it analyzes the Ori 2005 background, *not* the Xavier train, runs in about 3.5 s, and writes `simulations/ori_gauge_complete.png`. Its ten-component symbolic linearization is validated against a finite-difference linearization of the exact Ricci tensor at relative error 1.7×10^{-5} , and its scope and open sectors are recorded in the project notes. The symbolic vacuum and legality results are exact; the monodromy, image-sum, and stability results are numerical at the stated tolerances. The parameter point and the validation gates are those listed in Secs. 4, 5, 7, and 8.

References

- [1] S. Xavier, “Closed Timelike Curves from a Vacuum Traveling Wave,” arXiv:2607.00788 (2026).

- [2] A. Ori, A class of time-machine solutions with a compact vacuum core, *Phys. Rev. Lett.* **95**, 021101 (2005); arXiv:gr-qc/0503077.
- [3] S. W. Hawking, Chronology protection conjecture, *Phys. Rev. D* **46**, 603 (1992).
- [4] 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.
- [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] M. Visser, “The quantum physics of chronology protection,” in *The Future of Theoretical Physics and Cosmology* (Cambridge University Press, Cambridge, 2003), pp. 161–176; arXiv:gr-qc/0204022 (2002).
- [7] G. Klinkhammer, Vacuum polarization of scalar and spinor fields near closed null geodesics, *Phys. Rev. D* **46**, 3388 (1992).
- [8] W. A. Hiscock and D. A. Konkowski, Quantum vacuum energy in Taub-NUT type cosmologies, *Phys. Rev. D* **26**, 1225 (1982).
- [9] J. Schemoul, “Protection without teeth: semiclassical chronology protection on the Ori time machine,” in preparation (2026); cited here only qualitatively.
- [10] 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.
- [11] W. Melton, A. Strominger, and T. Wang, “Quantum Fields on Time-Periodic AdS_3/Z ,” arXiv:2510.15036 (2025).
- [12] J. Schemoul, Analysis scripts and validation records for this paper (2026), included with the manuscript for review; no public repository URL or DOI is assigned.