The occurrence of event horizons

Daryl Janzen[1,*]

[1] *Department of Physics and Engineering Physics, University of Saskatchewan, 116 Science Place, Saskatoon, SK S7N 5E2, Canada*

(Dated: March 28, 2026)

It is shown that null geometry imposes a previously unrecognized constraint on the causal representation of black-hole horizons. Using only standard Lorentzian structure, we prove a Null–Temporal Degeneracy Theorem: along any smooth null geodesic segment, all admissible temporal functions can be freely deformed to assign arbitrarily small or large time differences between distinct null-related events. As a corollary, every null generator consists of manifold points that are topologically distinct and causally ordered yet metrically coincident. Specializing to the Schwarzschild event horizon, which is foliated by such generators, we find that any temporal slicing adapted to an exterior observer becomes asymptotically tangent to a single horizon generator and intersects $\mathcal{H}^+$ only in the limit of infinite exterior time. Within the causal domain of the external universe, the event horizon therefore occurs solely as a null future boundary approached asymptotically by all admissible "now" slicings. This causal–geometric structure clarifies the behavior of horizon neighborhoods in Eddington–Finkelstein diagrams and constrains the interpretation of horizon-proximal processes in gravitational-collapse scenarios.

A NULL–TEMPORAL DEGENERACY THEOREM IN LORENTZIAN GEOMETRY

According to general relativity, there is a fundamental distinction between the events contained within an observer's fixed causal past—defined by the interior of their past light cone—and events that are merely "past" according to some arbitrary coordinate slicing but are spacelike separated from the observer [1, 2]. Only the former constitute physically certain information; the latter do not admit invariant temporal status relative to the observer.

For example, consider M31 V1, a Cepheid variable in the Andromeda galaxy. In our fixed causal past, M31 V1 is observed as a bright, evolved supergiant. The fact that we see it as such today implies that the star must have previously formed and undergone stellar evolution to reach this phase before the light we now observe was emitted.

However, Andromeda lies roughly 2.5 million light-years away. What we observe "now" is M31 V1 as it was 2.5 million years ago. Stellar astrophysics strongly suggests that M31 V1 has already undergone a core-collapse supernova sometime within the past 2.5 million years; yet we do not possess causal knowledge of that event. The reason is formal: the putative M31 V1 supernova is spacelike separated from here-now. Therefore, no positive statement about whether that explosion has already occurred (in any absolute sense) is meaningful without adopting a specific relativistic coordinate convention.

Thus, an observer's fixed causal knowledge in relativity is constrained entirely by the past light cone. The physical state we ascribe to distant astronomical objects is necessarily based on the most recent event in our causal past, which depends on their lookback distance. No astronomer refers to M31 V1 as "a supernova remnant," although that is almost certainly its status on a natural cosmological time slicing; rather, we may only say with certainty what it was when it emitted the photons that currently reach Earth.

The Null–Temporal Degeneracy Theorem below is a precise mathematical articulation of a subtle and remarkable implication of this causal structure. Before presenting the theorem, it is helpful to consider a simplified thought experiment.

The Andromeda galaxy spans roughly 100 kly in the radial direction relative to Earth. Consequently, the image we observe is not Andromeda "as it is now," nor even "as it was exactly 2.5 million years ago." Instead, we receive light whose emission times vary by $\sim 10^5$ years across the radial extent: photons originating on the far side were emitted roughly $\sim 10^5$ years earlier than those from the near side. As light propagates along the null geodesic connecting those regions to us, photons emitted from progressively nearer regions join the same bundle of null generators; the image we observe is the accumulated result of this continual inflow of successively later emission events.

Now imagine a hypothetical scenario in which Andromeda contracts along the radial direction, bringing its near and far sides increasingly close together, asymptotically approaching coincidence in the infinite future. In that limit, photons from the near and far regions would ultimately arrive from exactly the same spatial location in our past light cone. Their emission events remain distinct and causally ordered—the far-side emission event precedes the near-side emission event along a null geodesic connecting the two—yet their spatial separation in our frame vanishes.

The theorem below demonstrates a key geometric subtlety: although these emission events remain well-defined, distinct points of the manifold with an unambiguous causal ordering, *there is no geometrically or physically meaningful temporal separation between them.* Null-separation carries no invariant metrical scale; distinct null-related events may be causally ordered, yet the metric fails to assign them any nonzero temporal distance. In the limit of vanishing spatial separation, the two emission events become coincident in both a practical and metrical sense, even though they remain topologically distinct and causally ordered.

Crucially, the theorem does not require that the worldlines hosting these emission events remain nonsingular in the future. Even if either or both regions of Andromeda encounter a curvature singularity at some later time, as long as the segment of spacetime in the causal past of the relevant null geodesic is smooth, the degeneracy result holds without modification.

Theorem 1 (Null–Temporal Degeneracy with Regular Past)**.** *Let (M, g) be a smooth, time-orientable Lorentzian manifold of dimension $n \geq 2$. Suppose there exists an open set $U \subset M$ containing a future-directed null geodesic segment from p to q, with*

$$q \in \partial J^+(p), \qquad p \neq q,$$

such that U contains no curvature singularities, metric degeneracies, or boundary pathologies (i.e., U is a smooth region of M). No assumptions are made about the global structure of M or the existence of singularities to the future of q.

Let $\tau : M \to \mathbb{R}$ be any smooth temporal function, meaning that its gradient is everywhere timelike [2, 3]:

$$g^{ab}(\nabla_a \tau)(\nabla_b \tau) < 0 \qquad \text{on } M.$$

Then the following hold:

(i) The quantity $\Delta\tau := \tau(q) - \tau(p)$ has no invariant geometric meaning; it is not determined by the Lorentzian metric or causal structure.

(ii) For every real number $\Delta \geq 0$ there exists another smooth temporal function τ_Δ such that

$$\tau_\Delta(q) - \tau_\Delta(p) = \Delta.$$

(iii) There exists a sequence of smooth temporal functions $\{\tau_n\}_{n=1}^{\infty}$ such that

$$\tau_n(q) - \tau_n(p) > 0, \qquad \lim_{n\to\infty} \big(\tau_n(q) - \tau_n(p)\big) = 0,$$

and each τ_n has everywhere timelike gradient.

(iv) The presence of singularities in the causal future of q does not affect any of these conclusions. Throughout the deformations described above:

- *the manifold remains smooth and Hausdorff in a neighborhood of the null segment,*
- *the points p and q remain distinct,*
- *the causal relation $p \prec_\partial q$ remains unchanged.*

Proof. By hypothesis, there exists an open, smooth region $U \subset M$ containing the null geodesic segment $\gamma : [0, 1] \to M$ joining p to q. In particular, U contains no singularities, so normal neighborhoods are well defined and standard differential-geometric constructions apply.

Let $\tau : M \to \mathbb{R}$ be a smooth temporal function. Since $\nabla\tau$ is everywhere timelike, it is nonvanishing and lies in the interior of the timelike cone in T^*M. Because this cone is open, any sufficiently small C^1 perturbation of τ preserves the timelike nature of its gradient.

Choose a convex normal neighborhood $V \subset U$ containing the image of γ [1]. Select smooth bump functions

$$\phi_p,\, \phi_q : M \to \mathbb{R},$$

supported in disjoint subsets of V, such that $\phi_p(p) > 0$ and $\phi_q(q) > 0$. For sufficiently small real parameters A, B, the function

$$\tau_{A,B} := \tau + A\phi_q - B\phi_p$$

remains a temporal function, as the perturbation may be taken arbitrarily small in the C^1 norm and is supported within U, where the manifold is smooth.

We compute

$$\begin{aligned} \tau_{A,B}(q) - \tau_{A,B}(p) &= \big[\tau(q) + A\phi_q(q)\big] - \big[\tau(p) - B\phi_p(p)\big] \\ &= \Delta\tau + A\phi_q(q) + B\phi_p(p). \end{aligned}$$

Because $\phi_q(q)$ and $\phi_p(p)$ are strictly positive, we may select A, B to obtain any prescribed value

$$\tau_{A,B}(q) - \tau_{A,B}(p) = \Delta,$$

which proves (ii).

If we choose $\Delta_n \to 0$ and set $\tau_n := \tau_{\Delta_n}$, then each τ_n is a temporal function and

$$\tau_n(q) - \tau_n(p) = \Delta_n \to 0,$$

establishing (iii).

Statement (i) follows: if a quantity may be altered arbitrarily by smooth local deformations supported in a region where the metric is nonsingular, it cannot be an invariant of the Lorentzian geometry.

Finally, any singularities lying strictly to the causal future of q do not affect the argument, since all constructions are confined to the smooth neighborhood U containing γ. The geometry and topology of U remain unchanged, and the causal relation $p \prec_\partial q$ is unaffected. This proves (iv), completing the proof. □

Remark 1. This theorem formalizes the fact that null-related events admit a well-defined causal ordering but no invariant temporal separation. Local freedom to deform temporal functions in a smooth neighborhood of a null segment is unaffected by the presence of singularities elsewhere in the spacetime, provided they lie outside the region in which the construction is performed.

TOPOLOGICALLY DISTINCT—CAUSALLY ORDERED—METRICALLY COINCIDENT

We now formalise a frequently underappreciated structural fact about null geometry, which follows immediately from Theorem 1: every null geodesic consists of manifold points that are topologically distinct and causally ordered, yet metrically coincident.

Corollary 1 (Topological Distinctness and Geometric Coincidence of Null Generators)**.** *Let (M, g), p, q, and U be as in Theorem 1, with $q \in \partial J^+(p)$ lying on a future-directed null geodesic segment γ joining p to q. Then the following statements hold:*

(i) ***(Topological distinction)*** *The events p and q are distinct points of the manifold M and are therefore topologically distinguishable.*

(ii) ***(Causal ordering)*** *The events satisfy the strict null ordering $p \prec_\partial q$ [4], meaning that q lies on the boundary of the causal future of p and no event strictly between them lies in $J^+(p)$.*

(iii) ***(Geometric coincidence)*** *The Lorentzian separation between p and q is identically zero:*

$$g(\dot{\gamma}, \dot{\gamma}) = 0$$

along the entire segment of γ [2]. Consequently, the metric assigns no nonzero temporal or spatial distance between p and q; they are geometrically coincident *in the sense that their separation is null.*

(iv) ***(Absence of invariant temporal separation)*** *For any smooth temporal function τ on M, the value $\tau(q) - \tau(p)$ has no invariant geometric meaning, and by Theorem 1, it may be altered arbitrarily by smooth local deformations of τ in a neighborhood of γ. Thus any temporal "distance" between p and q is coordinate-dependent and does not reflect their geometric relationship.*

Taken together, these facts imply that any null generator consists of a set of manifold points that are topologically distinct and causally ordered, yet metrically coincident: the Lorentzian geometry assigns no nonzero interval along a null geodesic.

This topological–geometric structure underlies the degeneracy of temporal assignments along null hypersurfaces, including but not limited to event horizons in black-hole spacetimes.

The event horizon of a black-hole spacetime is precisely such a null hypersurface, and therefore inherits this structure of topologically distinct yet metrically coincident generators. We now specialize these general results to the Schwarzschild geometry, where the metrical degeneracy becomes especially transparent in Eddington–Finkelstein coordinates.

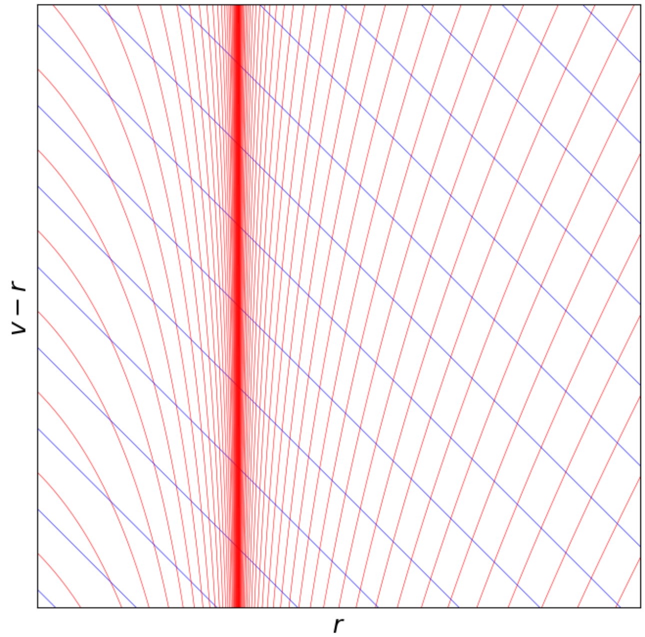

Figure 1. Null structure near the Schwarzschild event horizon in ingoing Eddington–Finkelstein coordinates (v, r). Outgoing null rays (red) and ingoing null rays (blue) are shown. The vertical line at $r = r_h$ represents a null generator of $\mathcal{H}^+$. Its apparent "length" in the diagram is *purely topological*: all points along the generator are metrically coincident, with zero invariant temporal or spatial separation along the null direction. Distinct horizon-crossing events correspond to distinct manifold points but share the same null geometric location; the vertical extent depicted here therefore does *not* represent metrical duration or physical extent in the Lorentzian geometry.

METRICAL DEGENERACY OF EVENT HORIZONS IN BLACK HOLE GEOMETRIES

The previous corollary established that any null generator consists of topologically distinct but metrically coincident events. The future event horizon $\mathcal{H}^+$ of a black-hole spacetime is a null hypersurface foliated by such generators, and therefore inherits this structure. In particular, horizon-crossing events along a single generator are causally ordered but admit no invariant temporal separation. Eddington–Finkelstein coordinates [5, 6] make this degeneracy visually explicit: the apparent "vertical" extension of the horizon at $r = r_h$ encodes only the ordering of null-related events, not any metrical duration. The following corollary formalizes this specialization for the Schwarzschild spacetime.

Corollary 2 (Degeneracy of Horizon-Crossing Events in EF Coordinates)**.** *In Schwarzschild spacetime expressed in ingoing Eddington–Finkelstein coordinates (v, r), let $\mathcal{H}^+$ denote the future event horizon $r = r_h$ [1], and let p and q be two distinct horizon-crossing events lying on the same null generator of $\mathcal{H}^+$, with q to the future of p*

along the generator.

For an exterior observer at fixed areal radius $r > r_h$, the events p and q lie on distinct points of the null boundary of the observer's causal past, and are therefore causally ordered but null-related.

Let Θ be any smooth temporal function adapted to the observer's worldline. Then the coordinate difference $\Theta(q) - \Theta(p)$ is non-invariant and can be altered arbitrarily by smooth deformations of Θ in a neighborhood of the generator. In particular, because $\mathcal{H}^+$ is a null hypersurface of metrically coincident generators, an exterior observer cannot assign any invariant temporal separation to horizon-crossing events: all such events occur at the same spatial location $r = r_h$ and admit no metrical duration between them.

Thus the apparent "vertical separation" of horizon points in an Eddington–Finkelstein diagram, as in Fig. 1, reflects only the causal ordering of distinct manifold events. It does not encode any geometric or physical temporal distance.

Having established the metrical degeneracy of horizon generators, we now turn to the causal structure of gravitational collapse and black-hole mergers, where the geometry of null boundaries determines the portion of the collapsing worldtube that can influence events in the external universe.

APPLICATION: GRAVITATIONAL COLLAPSE AND BLACK HOLE MERGER CAUSALITY

The causal structure of gravitational collapse, event horizon occurrence, and black hole mergers is most transparently described in the Schwarzschild geometry, with hypothetical processes modeled through purely radial motion. The event horizon at $r_h = 2GM/c^2$ has the same essential causal behavior as that of more general black hole solutions such the Kerr family; the rotational structure in such cases introduces additional features that do not alter the metrically null nature or causal occurrence of the horizon. For clarity of exposition, we therefore restrict attention to Schwarzschild spacetime.

Figure 2 illustrates a standard radial collapse scenario in ingoing Eddington–Finkelstein coordinates [7, 8]. The surface of the collapsing star follows a timelike worldline and crosses r_h in finite proper time, later forming a curvature singularity at $r = 0$. An external observer remains at fixed areal radius $r > r_h$ just beyond the star's original radius, and releases two test masses toward the collapsing object at different proper times. Test mass 1 intersects the stellar surface before the formation of the horizon, while Test mass 2 is released sufficiently late that its ingoing null ray at the moment of release fails to intersect the worldline of the collapsing surface at $r > r_h$.

Outgoing null rays in the exterior region, shown in red, all intersect the observer's worldline. Thus, regardless

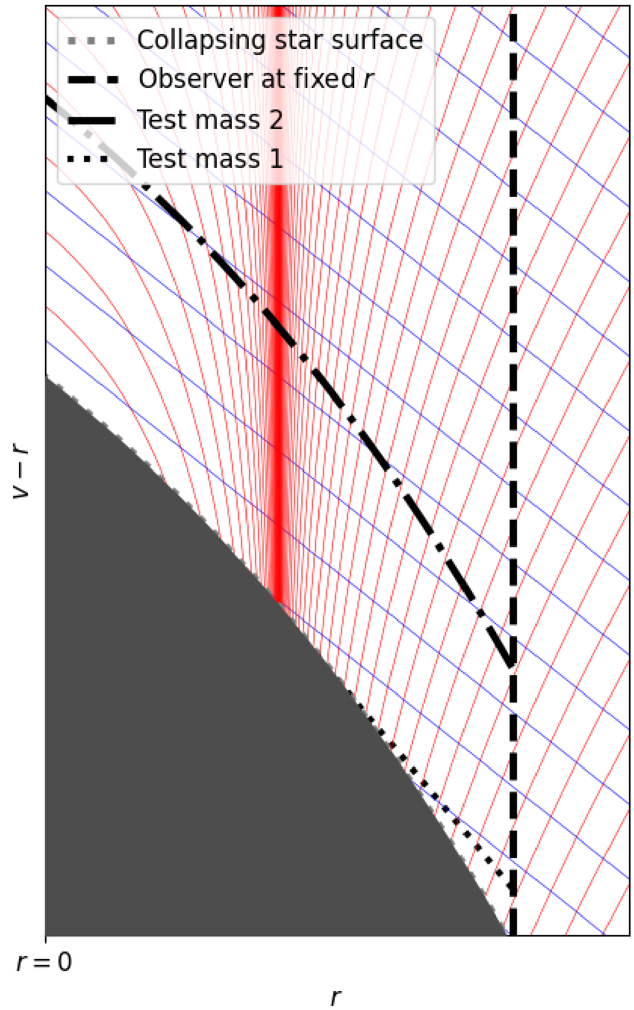

Figure 2. Radial gravitational collapse in ingoing Eddington–Finkelstein coordinates. The dotted grey curve represents the collapsing star's surface. The dashed curve is an external observer held at fixed radius $r > r_h$. Test mass 1 (dotted) reaches the stellar surface prior to horizon formation. Test mass 2 (dash–dot) is released later, and its ingoing null signal (blue grid) no longer intersects the collapsing surface. Outgoing null rays (red grid) emitted in the exterior region reach the observer, who therefore only ever sees the star and infalling masses asymptotically approach r_h. The shaded region indicates the stellar interior.

of how long the observer waits, the collapsing surface is never seen to reach r_h, since all (null-related) horizon events are spacelike separated from any finite-time exterior event. Even if the observer were equipped with an antenna capable of detecting signals of arbitrary wavelength, the portion of the stellar surface visible to them would only approach r_h as their own proper time tends to infinity.

Before turning to the behavior of the test masses and the application of the above theorem and its corollaries, it is helpful to summarize the observational appearance of mergers from the standpoint of such an exterior observer, at fixed $r > r_h$. If the observer is equipped with a perfect antenna and gravitational wave detector, they will continually see the collapsing surface at radii strictly larger than r_h. Therefore, any merger event that generates gravitational waves the observer later sees must have occurred at a time when the stellar surface remained visible in their causal past. Consequently, at the moment the merger signal arrives, the observer sees a consistent picture: the collapsing mass continues to asymptotically approach r_h, a test mass or additional object falls radially inward and meets it, a burst of gravitational radiation is emitted and later detected, and a post-merger object is subsequently observed to continue its approach toward a new effective radius r'_h determined by the total mass. Since the merger event lies within the observer's causal past, and since its outcome remains visible afterward, the entire process *must* have occurred in the region of the spacetime where both bodies remained larger than their respective horizon radii.

Remark 2. It is common in the literature to refer to compact objects that will develop event horizons as "collapsed objects," even in contexts where the horizon formation event has not yet occurred relative to any finite exterior time (e.g., Hawking [9], Penrose [10]). However, Hawking's own causal definition of the event horizon as $\mathcal{H}^+ = \partial J^-(\mathcal{I}^+)$ makes clear that such a horizon is a null future boundary rather than a structure occurring at any finite exterior time. Within the causal framework developed above, the two usages cannot be conflated.

The horizon-formation event for a collapsing object never lies within the past light cone of any gravitational-wave generation event that is itself visible to an exterior observer [11]. Since gravitational fields, light signals, and gravitational waves all propagate along null geodesics, no event outside an emission event's causal past can influence its generation. The situation is directly analogous to the M31 V1 example: we do not now observe its supernova remnant because that event does not lie within our causal past. Similarly, any gravitational-wave generation event accessible to an exterior observer must lie in a region where the relevant mass distributions were still larger than their horizon radii in the causal sense established above.

Just as referring to M31 V1 as a "supernova remnant" would implicitly place its supernova event in our causal past, Hawking's informal use of "collapsed object" implicitly places the horizon-formation event in the past of an exterior observer (whether the gravitational-wave generation event or our later observation of it)—a usage that contrasts with his own causal definition of the event horizon and is precisely the tension resolved by the framework developed here.

Having laid this conceptual groundwork, we can now proceed to interpret the causal generation of gravitational waves in the scenario shown in Figure 2, in the case of the two test-mass collisions.

Since test mass 1 encounters the collapsing surface while its radius is clearly larger than r_h, the interpretation is trivial. The wave generation occurs when the two masses are sufficiently close that the waves are generated in the strong field exterior. Those gravitational waves subsequently travel along an outgoing null line and are later observed by the observer at fixed r.

The case of test mass 2 is far more interesting and informative, and highlights the temporal degeneracy in the horizon vicinity remarkably well. Note that as the test mass approaches r_h, the most recent causal signal along the null connection to the collapsing star comes from ever closer in r. The situation is analogous to the "shrinking Andromeda" example in Section , where null-separated emission events become metrically coincident in the limit. One is tempted to interpret the graphical separation of these two events in Figure 2 as implying temporal distance, but null-degeneracy of the horizon events ensures this is merely an illusion due to the chosen coordinate system, which separates them by causal, not metrical order. Thus, the figure, accurately interpreted, does indeed illustrate that the distance in r between the collapsing star's surface and the test mass continually decreases when measured across their null connection. And if gravitational waves are generated, and later seen by the external observer, the wave generation event must have occurred at causal times preceding the horizon-formation event along both worldlines, relative to their shared null connection.

ASYMPTOTIC ALIGNMENT OF EXTERIOR TEMPORAL SLICINGS WITH THE HORIZON

In Schwarzschild spacetime, the future event horizon $\mathcal{H}^+$ is a null hypersurface [1] foliated by metrically coincident generators, as established in Sections –. Any temporal function adapted to an exterior observer must therefore accommodate this null–geometric structure: as one moves to increasingly late exterior times, the observer's level sets are forced to approach and ultimately "pile up" against a single generator of $\mathcal{H}^+$. The following lemma makes this geometric constraint precise.

Lemma 1 (Asymptotic Alignment of Exterior Temporal Slicings). *Let (M, g) be the exterior region of the Schwarzschild spacetime with areal radius $r > r_h$, and let*

$$O : \tau \mapsto \big(v(\tau), r_0\big)$$

be an observer at fixed radius $r_0 > r_h$, parametrized by proper time τ. Let $\Theta : M \to \mathbb{R}$ be any smooth temporal function adapted to O, meaning

$$\Theta\big(O(\tau)\big) = \tau + \text{const.}$$

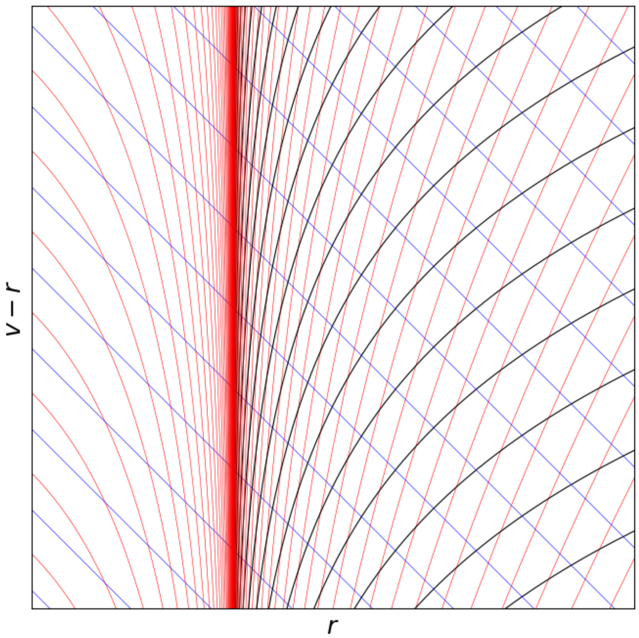

Figure 3. Level sets of the stationary Schwarzschild time coordinate t (black) plotted on an ingoing Eddington–Finkelstein diagram. Outgoing (red) and ingoing (blue) null rays are shown. The Schwarzschild slices t = const tilt upward and accumulate along a single generator of $\mathcal{H}^+$, becoming asymptotically tangent to it. Although this image uses the stationary t coordinate, the asymptotic piling-up is a generic consequence of the Null–Temporal Degeneracy Theorem and Lemma 1: any temporal function adapted to an exterior observer [3] must degenerate in exactly this way, with all late-time slicings converging on the same null generator.

For each τ, define the level set

$$S_\tau := \{\, x \in M : \Theta(x) = \tau \,\}.$$

Then:

(i) ***(Convergence to a unique generator)*** *The intersection sets $S_\tau \cap \mathcal{H}^+$ converge, as $\tau \to \infty$, to a single null generator of $\mathcal{H}^+$.*

(ii) ***(Asymptotic null tangency)*** *At the intersection points $S_\tau \cap \mathcal{H}^+$, the tangent hyperplanes of S_τ become asymptotically tangent to that generator. Equivalently,*

$$g^{ab}(\nabla_a\Theta)(\nabla_b\Theta) \longrightarrow 0 \quad \text{along } S_\tau \cap \mathcal{H}^+ \quad \text{as } \tau \to \infty,$$

so that $\nabla\Theta$ becomes asymptotically null in the horizon direction.

Proof. For each τ, the level set S_τ intersects the past light cone $J^-(O(\tau))$ in a smooth spacelike hypersurface. Since the Schwarzschild exterior is globally hyperbolic [2], the boundary $\partial J^-(O(\tau))$ is a smooth null hypersurface whose intersection with $\mathcal{H}^+$ is a single generator segment.

As $\tau \to \infty$, the sets $\partial J^-(O(\tau))$ monotonically approach $\mathcal{H}^+$, and by the defining property of the event horizon [9], every generator of $\mathcal{H}^+$ is the limit of such past-light-cone boundaries. Thus $S_\tau \cap \mathcal{H}^+$ converges to a single generator, establishing (i).

Since S_τ intersects $\partial J^-(O(\tau))$ transversely and $\partial J^-(O(\tau))$ approaches a null hypersurface, the tangent spaces $T(S_\tau)$ at intersection points become asymptotically tangent to the horizon generator. This implies that $\nabla\Theta$ becomes asymptotically null along $\mathcal{H}^+$, giving (ii). □

Remark 3. For any exterior observer—and, more generally, at all points in the exterior region $r > r_h$—the event horizon $\mathcal{H}^+$ lies permanently at the null future boundary of the observer's causal future. In particular, if Θ is any temporal function adapted to such an observer, then each finite-Θ slice S_Θ lies entirely in the region $r > r_h$ and therefore never intersects $\mathcal{H}^+$ (e.g. Figure 3). Only in the limit $\Theta \to +\infty$ do the slices approach and become tangent to a single horizon generator.

Thus, relative to any admissible exterior time coordinate, the event horizon does not occur at any finite time within the exterior domain. Within the causal domain of the external universe, the event horizon does not "exist" as a present structure—that is, it is not present across any sequence of finite exterior time-slices adapted to an exterior observer. Rather, it *occurs* only as the null future boundary approached in the infinite-time limit: a singular occurrence at the end of the universe. This conclusion concerns only the behavior of exterior-adapted slicings and does not address the global structure of the spacetime beyond that domain.

Remark 4. It is important to emphasize that Lemma 1 is not a statement about observational limitations or about the causal past of any particular observer. Rather, it characterizes the structure of every temporal slice S_Θ adapted to any exterior worldline. Such slices extend well beyond the observer's past light cone and include all events that may be taken as "out there now" relative to that observer, yet none of them—for any finite value of Θ—ever intersect the event horizon or contain any portion of a future singularity. Each finite-Θ slice contains only still-collapsing matter. The event horizon therefore does not occur as part of any finite exterior time-slice; it occurs only once, at the infinite-time boundary of the exterior domain. In maximal extensions, the apparent vertical "stack" of horizon points merely displays the causal ordering of distinct manifold events that are nevertheless metrically coincident, rather than a temporal extension of the horizon within finite exterior time.

* daryl.janzen@usask.ca

[1] Hawking S W and Ellis G F R 1973 *The large scale structure of space-time.* (Cambridge University Press)

[2] Wald R M 1984 *General relativity* (Chicago: The University of Chicago Press)

[3] Bernal A N and Sánchez M 2005 *Communications in Mathematical Physics* **257** 43–50

[4] Penrose R 1972 *Techniques in differential topology in relativity* (Society for Industrial and Applied Mathematics)

[5] Finkelstein D 1958 *Physical Review* **110** 965–967

[6] Misner C W, Thorne K S and Wheeler J A 1973 *Gravitation* (San Francisco: W. H. Freeman)

[7] Oppenheimer J R and Snyder H 1939 *Physical Review* **56** 455–459

[8] Penrose R 1965 *Physical Review Letters* **14** 57–59

[9] Hawking S W 1971 *Physical Review Letters* **26** 1344

[10] Penrose R 1969 *Nuovo Cim. Riv. Ser.* **1** 252–276

[11] Abbott B P, Abbott R, Abbott T D *et al.* 2016 *Phys. Rev. Lett.* **116** 061102