From One Clock to Programmable Coordinates: A Material Layer for Coordinate Freedom in General Relativity

Hermann Borotschnig

May, 2, 2026

Abstract

In general relativity, any valid coordinate chart may describe a given spacetime region, yet material reference frames used to visualize this freedom—such as Einstein's "reference mollusk," a deformable reference body equipped with arbitrarily running clocks—can directly display only charts adapted to their built-in protocol. This contribution introduces an elementary enhancement that lifts that restriction: each observer in a seed congruence carries a programmable device that computes and displays four coordinate values as functions of the four seed coordinates. Several consequences follow. A single observer swarm can directly display charts not adapted to its underlying congruence, including light-cone coordinates with two null and two spacelike basis directions, Painlevé–Gullstrand coordinates inside the Schwarzschild horizon with four spacelike basis directions, and Kruskal–Szekeres coordinates with two null directions. Coordinate directions need not mirror observer motion: observers travel on timelike worldlines while the displayed coordinates may have null or spacelike basis directions. The same physical swarm can display multiple charts during deployment. Standard coordinates and special constructions can be programmed, including harmonic coordinates or coordinates built from curvature invariants. The metric is then reconstructed in the displayed coordinates by standard local-inertial measurements, keeping three layers manifestly distinct: the material reference apparatus, the coordinate labels it displays, and the metric subsequently reconstructed in those labels.

1 Introduction

"How is one to drive a nail into spacetime to mark a point?" With this vivid question, Misner, Thorne, and Wheeler highlight a fundamental challenge in the teaching of general relativity (GR): spacetime has no natural coordinate grid waiting to be discovered [1, 2]. Their answer—characterize points by "what happens there"—suggests using events themselves as the markers. But this immediately raises the next question: who observes these events, and how could they record the observations together with the coordinates they assign to what they observe [3]?

One could answer: in addition to the events of interest, let a swarm of local observers also "happen there"—a continuous distribution of observers filling spacetime, each recording local events and assigning coordinate labels.[1] This idea, as naive as it may appear at first sight, has deep roots: Einstein himself envisioned a "reference mollusk" in 1917, a single deformable reference body populated by observers with clocks [4, 5]. A closely related construct is described by Landau and Lifshitz as "bodies which fill all the space like some sort of 'medium' ... with arbitrarily running clocks fixed on them." [6] Modern formulations in canonical and quantum gravity employ dust-based reference frames [7, 8], and philosophical and pedagogical discus-

[1] Dense coverage conflicts with event-observability and test-matter passivity: the frame must not source the gravitational field. These and other standing idealizations are mutually competing and can only be approximately fulfilled. In practice one would need limited observation regions and a compromise finite (near zero) density. As density decreases, a sparse swarm connected via information-carrying signals could emerge; similarly, any physical implementation of programmable devices would operate at finite readout resolution. Even under such test-matter assumptions, one must still additionally restrict the spacetime region $\mathcal{U}$ so as to avoid singularities. We also exclude both caustics arising from the spacetime geometry (e.g., geodesic focusing) and caustics resulting from inadequate steering of accelerated observers. See also Suppl. Mat., Assumptions.

sions sometimes speak of "observer dust" [9, 10, 11, 12].[2] [3]

Yet a subtler pedagogical gap persists between the full mathematical freedom of choosing any valid coordinate charts and the narrower class of raw coordinates directly supplied by these material constructions. The form of the equations of GR remains invariant under *any* valid coordinate transformation on a chosen domain [3]. By contrast, a reference frame in which each observer in a congruence carries one clock and fixed worldline labels imposes additional conditions on its coordinatization. While this does not affect the symmetry content of the theory itself [12, 9], the built-in protocol adaptation limits which coordinate systems can be *directly* displayed during actual deployment without further processing.

[2]For entry points into the literature on material reference frames, intrinsic coordinates, and relational observables in modern theoretical physics, including canonical and quantum gravity, see Bergmann and Komar; Brown and K. V. Kuchař; Rovelli; Giesel and Thiemann; Dittrich; Pons; Pitts; Bamonti; and historical reviews by Tambornino and Salisbury (Refs. [13, 7, 8, 14, 15, 16, 17, 18, 19, 20]). The present proposal pursues an aim complementary to that of canonical approaches, which often explicitly provide physical models for specific classes of reference frames as a tool for *reducing* coordinate freedom, obtaining observables and evolution. The programmable layer instead seeks to *lift* gauge adaptations in order to provide a tangible pedagogical analogy between programming freedom and chart freedom on top of a seed frame. It does not change the gauge adaptation of the underlying seed, but allows the enhanced reference frame to directly display (multiple different) coordinate systems that each lie outside the adapted class entirely. While the discussed construction operates in the test-matter limit throughout (see also Ref. [3], § 4.3.1) a functional programmable layer would, by definition, continue to apply a passive coordinate transformation to the seed coordinates regardless of coupling regime.

[3]For concise entry points into philosophical discussions of reference frames, general covariance, diffeomorphism invariance, the hole argument, and relational versus substantival readings of spacetime, see Norton, Earman, Callender and Hoefer, Pooley, Stachel, Weatherall, Mozota Frauca, and Giovanelli (Refs. [12, 21, 22, 23, 24, 25, 26, 27]). Specifically, general covariance is often discussed in two related senses. *Formal* general covariance concerns writing laws in geometric form so that their component expressions transform appropriately under passive coordinate changes. Though Einstein was influenced by this (see also Ref. [27]), any theory can be formally recast in this manner (see e.g., Ref. [1], § 12). *Substantive general covariance* typically concerns whether models related by *active* diffeomorphisms—acting on all dynamical fields including g_{ab}—represent the same physics, especially in the absence of fixed background structures (see Ref. [12]). The passive/active distinction is subtle: mathematically equivalent; the gauge reading arises only without fixed background structures (e.g., no absolute–time τ_a or spatial metric h^{ab}). We focus on passive transformations performed on the spot via programmable devices. For a careful discussion of the relation between general invariance and general covariance, see Ref. [3], § 2.

The limitation can be stated concretely. The original purpose of coordinates is to label individual *events*, not observers or worldlines. Yet in a mollusk or congruence-adapted protocol, three displayed labels remain constant along each observer worldline and therefore play a double role: they identify the observer who recorded the event, and they also serve as three of the four coordinates assigned to that event. Operationally, the event label is encoded by answering two questions: which observer $\mathbf{y}$ detected the event, and what clock reading λ that observer recorded. Even before any metric has been measured, each observer can verify from successive readings that the three displayed values $\mathbf{y}$ remain fixed along its own worldline.

Once a metric is introduced, this protocol is seen to impose a gauge adaptation: one coordinate-line family coincides with a timelike congruence of observers, while the remaining three coordinates are comoving labels. A general chart need not respect that structure. Light-cone coordinates, for example, have two null and two spacelike coordinate basis directions, while Painlevé–Gullstrand coordinates inside the Schwarzschild horizon have four spacelike coordinate basis directions, with timelike directions appearing only as linear combinations of the basis vectors.[4] Such charts remain admissible mathematically, but they cannot be displayed directly by the raw congruence-adapted protocol during deployment; they can only be obtained by an additional coordinate transformation. Yet direct display is what preserves the locality that motivates the material swarm: observers are placed throughout the region precisely so that they can label events on the spot. Moreover, on-the-fly metric reconstruction yields metric components in whichever coordinates are available during deployment. The pedagogical question is therefore whether one can retain this locality while freeing the directly displayed coordinates from the protocol-induced restrictions.

Before addressing this question, let us first distinguish the underlying material reference-frame apparatus—which includes physical devices for local event detection, labeling by displayed coordinate values, and record collection—from the mathematical chart that the displayed coordinate values approximate: the devices (including the executed programs, if present) carry operational content, whereas the chart and the numerical values of coordinates

[4]The Supplementary Material examines which coordinates may be directly displayed by mollusk and congruence-adapted seed frames (Appendix A) and shows by explicit examples how a programmable layer can lift or switch the gauge adaptation through programmable transforms in SR and in GR (Appendix B), including interior Painlevé–Gullstrand coordinates (B.4), and double null Kruskal–Szekeres coordinates (B.5).

remain conventional [10, 3].[5] Where no confusion can arise, we refer to the physically displayed values also simply as "coordinates."

A congruence-based construct is the starting point: a seed reference frame that provides gauge-adapted seed coordinates on a spacetime region $\mathcal{U}$. It reflects the kinematic structure of a smooth timelike congruence of observers with one clock and fixed labels per worldline. No specific (e.g., dust) matter model is assumed, however, and the seed frame observers may also be accelerated as long as they conceptually remain test-matter.

We now equip each observer in the seed frame with a programmable device that physically evaluates transformations of the seed coordinates during deployment and displays the resulting coordinate labels on its comoving screen. Operationally, one may load any program one wishes. In the present context, however, our interest lies in coherently programmed transformations that constitute valid coordinate transforms, both at the level of the associated continuum chart and at the level of the displayed values under the idealizations adopted here. Hence, whenever the corresponding coordinate transform is computed by a valid program, the displayed values implement a valid chart during deployment. In this sense, programming freedom serves as a pedagogical analogy for the mathematical freedom to choose any valid chart, while the program and the chart remain conceptually distinct. We may for now leave open how far this analogy extends as long as we can use it as a pedagogical tool.[6] [7]

[5] Suppl. Mat., Appendix B.1 provides details for programmable linear transforms of standard Minkowski coordinates that lead to coordinates with double-timelike, double-null, and four spacelike basis directions already in SR. Appendix B.2 presents a rotating cylinder or disk example in SR. It emphasizes the decoupling of (accelerated) observer motion and displayed (inertial) coordinate *values*, while highlighting the physical nature of the programmable layer.

[6] Any chart $\tilde{X}$ on $\mathcal{U}$ compatible with a given seed chart X is related to it by the usual transition-map structure of an atlas. In standard manifold terms (see also Ref. [28]), two charts $X, \tilde{X} : \mathcal{U} \to \mathbb{R}^4$ on the same open set are related by $\Phi = \tilde{X} \circ X^{-1}$, so that $\tilde{X} = \Phi \circ X$, where Φ is here taken to be a C^k diffeomorphism between open subsets of $\mathbb{R}^4$ mapping $X(\mathcal{U})$ onto $\tilde{X}(\mathcal{U})$, with $k \geq 1$. Thus every compatible chart $\tilde{X}$ arises from some valid Φ defined on the seed chart image. The pedagogical analogy uses this mathematical structure when it can be mirrored by valid programmed coordinate transformations $\tilde{X}^\mu(X)$ that play the role of Φ.

[7] Once the seed reference frame has supplied seed coordinates on a working region $\mathcal{U}$, one may load an effectively specified program that converts those seed labels into new displayed labels. The programmable layer therefore enlarges what can be displayed, but it inherits the familiar GR requirement that chart validity be verified case by case. We

The underlying point of this approach is not specific to congruence-adapted seed frames: a freely programmable computational layer may also decouple the displayed coordinates of other protocol-induced adaptations applying to seed frames of a different nature [14].[8]

But in what follows, the idea is applied in a beginner-friendly setting by attaching an explicit programmable layer to elementary seed frames based on timelike observer congruences, including mollusk-type frames. In this context, the co-deployed computing devices constitute a genuine material enhancement of the seed reference-frame apparatus and a physical program upgrade that enables the devices to compute and display the result of valid programmed coordinate transformations away from the seed coordinates. For pedagogical vividness, we may imagine this added material layer in the form

restrict attention to programmed transformations that satisfy the usual chart conditions both at the level of the displayed values and at the level of the associated continuum map on $\mathcal{U}$. In this sense the analogy with mathematical coordinate freedom implies free choice only among those programmed transformations that are valid on the region under study. Compared with the original mollusk, one replaces the single clock function $x^0 = \lambda(\tau, \mathbf{y})$ by four functions $X^0(\tau, \mathbf{y}), \ldots, X^3(\tau, \mathbf{y})$, which need no longer vary monotonically in τ. The relevant validity requirements are structurally the same type of conditions that already apply to the seed frame $(\lambda(\tau, \mathbf{y}), \mathbf{y})$ itself. We need not commit to any specific computational paradigm (digital, analog, hybrid, etc.). In any realistic scenario one would work not with ideal real-valued continuum inputs and outputs, but with a discrete set of observers, a finite set of recorded events, and finite readout resolution. The implementation-level computational requirement then becomes distinguishability of finitely many discrete values at the adopted numerical tolerance: different recorded events should not display the same finite-precision coordinate quadruple. Likewise, when local Jacobians are reconstructed, one works in practice at a chosen resolution with finite differences between nearby events. These are computable counterparts to continuum chart-validity criteria, whose finite precision may be extended as required. On a conceptual level, one may still ask how far the analogy between programming freedom and chart freedom can be pushed; a related mathematical question is isolated for guided exploration in Suppl. Mat., Appendix C. The pedagogical point in the main text does not depend on any stronger claim beyond explicit standard textbook examples.

[8] Rovelli studies four emitters which broadcast their proper-time readings. The coordinates s^α of an event are the four emission times encoded in the signals received at that event. The corresponding coordinate one-forms are null ($g^{\alpha\alpha} = 0$, no summation), so the raw coordinates are adapted to the emission protocol. Rovelli then applies a fixed linear transform to the raw coordinates and reconstructs the metric by additional operational means. The shared logic—protocol-dependent seed coordinates, transformation, metric reconstruction—makes these papers natural follow-up reading for readers who have absorbed the elementary framework presented here. For a contrasting perspective and an elementary visualization, see also Ref. [26] and Fig. 1 therein.

of "smartphones" replacing the "arbitrarily running clocks" in a congruence-adapted seed frame.

The proposal is illustrated first through light-cone coordinates in special relativity (SR), then extended to curved spacetimes, and finally applied in standard metric reconstruction via the equivalence principle, offering a coherent route from SR to the threshold at which the Einstein field equations could be introduced.

2 Reference Frames in Special Relativity

To prepare for GR, we begin by revisiting and slightly reframing the implementation of physical reference frames in SR. A central operational guideline in relativistic theories is that physically invariant statements can be phrased in terms of local coincidences of events. This *locality principle* will also serve as the operational starting point for the formulation of reference frames.[9]

Consequently, even in SR, an "observer" is best understood not as a single point-like abstraction but as a *swarm of local observers*, each equipped with measuring devices and capable of recording events encountered along their respective worldlines. A spacetime diagram, then, may be interpreted as the operational representation of the totality of recorded measurements: the collective event data gathered by all these assistant observers who conceptually fill spacetime. Each local device assigns coordinates—its own clock reading and spatial position—to the events it directly registers.

Mathematically, these coordinates are typically denoted by x^μ with $\mu = 0, 1, 2, 3$, such that a spacetime event is labeled by $x = (x^0, x^1, x^2, x^3)$. The spatial coordinates are given by (x^1, x^2, x^3). The temporal coordinate is given by $x^0 \equiv ct$, where c is the speed of light and t is the local clock reading.

Some introductory texts visualize the underlying reference frame as a regular three-dimensional lattice, composed of idealized rigid rods and locally placed synchronized clocks [30, 32, 31]. These setups are designed to

[9] We adopt standard idealizations (see also Endnote [**?**]): (i) a conceptually dense observer swarm; (ii) point-coincidence between observer and event; (iii) ex-post-facto (delayed) data analysis. These follow established precedent, e.g. Ref. [29] (p. 109), Ref. [30] (§ 2–8), and Ref. [31] (§ 1.13). Extended observers could be modeled via worldtubes or local comoving frames, but point–coincidence suffices here. On-the-fly signal exchanges (with delay $d\tau$ to probe local four-dimensional neighborhoods) are also conceivable in addition to post-deployment evaluation. For a list of physical and computational assumptions, see also Suppl. Mat., Assumptions.

illustrate the operational assumptions of inertial frames in SR—such as the Euclidean geometry of space, global clock synchronization, and the coincidence idealization between the event and its recording observer.

An equivalent and more locality-focused depiction is shown in Fig. 1. Here, spacetime events are assigned coordinates by a distributed network of observers, each using their locally maintained synchronized clocks and pre-established positions relative to others. When needed, spatial alignment across the network can be re-established by activating orthogonal laser pulses, which form a transient coordinate grid. This operational picture emphasizes that it is the observers—and not the coordinate mesh—that serve as the physical basis of the reference frame.

Once set up, a reference frame allows for the measurement of spatial distances, time intervals, and kinematic quantities by combining purely local observations, even when the relevant events are separated in space and/or time. For instance, measuring the length of an idealized rod requires two simultaneous events at different locations: one recorded by the observer who happens to be at one end of the rod at the agreed time, the other by the observer at the opposite end. Both rely on their synchronized clocks to ensure simultaneity in their frame. The length is then the Euclidean distance between their spatial positions.

The configuration of free-floating observers shown in Fig. 1 naturally generalizes to multiple overlapping inertial reference frames, such as frame A and frame B, covering the same region of spacetime. Each frame is assembled independently, using only its own network of observers, synchronized clocks, and internal measurement protocols.

As a result, any given spacetime event can be assigned two sets of coordinates—x and $\tilde{x}$—depending on which frame's observers perform the local assignment. Because the two networks are coextensive, observers in one frame can also access the measurements made by their local counterparts in the other. This operational comparison allows one to empirically determine a transformation function:

$$\tilde{x} = \Lambda_{A\to B}(x). \tag{1}$$

A striking empirical result emerges from this experiment: when two inertial frames A and B are realized independently—each relying solely on its own local observers and synchronization protocol—the transformation $\Lambda_{A\to B}$ turns out not to be Galilean, as would be expected in Newtonian mechanics. Instead, it corresponds to a Lorentz transformation (up to translations), in

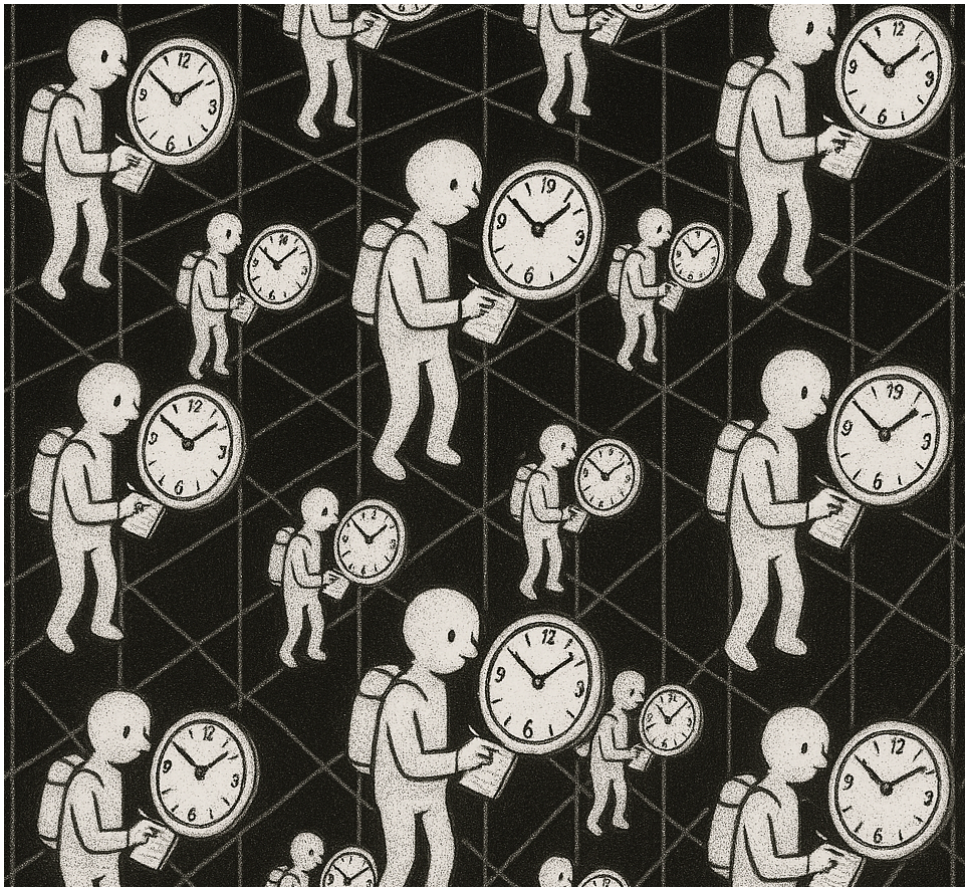

Figure 1: An inertial orthonormal reference frame in special relativity, realized by distributed observers with synchronized clocks. The laser grid serves only as an alignment aid; the observers constitute the physical reference frame.

accordance with the Poincaré symmetry that preserves the form of Maxwell's equations. This result is theoretically equivalent to the constancy of the speed of light as measured in any inertial reference frame—the principle originally postulated by Einstein and used to derive the Lorentz transformation. Equivalently, SR is Poincaré invariant: the spacetime interval between any two events remains the same in all inertial frames.

Expressing the Poincaré transformation in component form with Einstein notation as

$$\tilde{x}^\mu = \Lambda^\mu{}_\nu x^\nu + a^\mu, \tag{2}$$

where a^μ accounts for arbitrary (but physically irrelevant) offsets in spacetime origins between frames, and introducing the standard Minkowski metric

$$\eta_{\mu\nu} = \mathrm{diag}(-1, +1, +1, +1), \tag{3}$$

the invariance of the infinitesimal spacetime interval ds^2 is given by

$$ds^2 = \eta_{\mu\nu}\, dx^\mu dx^\nu = \eta_{\mu\nu}\, d\tilde{x}^\mu d\tilde{x}^\nu. \tag{4}$$

Importantly, Eq. (4) concerns coordinate differences between events. In the operational setup, these differences are not read off by a single observer at one location; they are assembled from local records distributed across the frame, much as in the rod-length measurement above. When both frames record the same events, each independently evaluates the spacetime interval from its coordinate differences between those events, and the equality across frames is an empirical fact.

The spacetime interval also indicates the causal character of infinitesimal displacements through each event. For example, consider the tangent directions to the coordinate curves obtained by varying one coordinate while holding the other three fixed. For an infinitesimal displacement with components dx^μ, the invariant interval satisfies

$$ds^2 \begin{cases} < 0 & \text{timelike,} \\ = 0 & \text{null or lightlike,} \\ > 0 & \text{spacelike.} \end{cases} \tag{5}$$

In standard Minkowski coordinates, the displacement for which only $dx^0 \neq 0$ is timelike, whereas the displacements for which only $dx^i \neq 0$ are spacelike. For the subsequent sections, we simplify notation by adopting units in which $c = 1$.[10]

[10] In an orthonormal local inertial frame, length and time are measured using idealized rods and clocks (see also § 3.1 in Ref. [3]). While physical interpretations are conceptually reduced to measurements in such frames, nothing here depends on a specific choice of units. A coordinate chart is a map $x\colon U \to \mathbb{R}^4$, in which $\mathbb{R}^4$ is unit-free; changing units corresponds to constant rescalings with induced metric component transformations, keeping $ds^2 = g_{\mu\nu}\, dx^\mu dx^\nu$ invariant. The computed and displayed coordinates are just numerical labels (see Fig. 3); physical interpretation arises only through the metric. Correspondingly, a natural assignment would be $[ds^2] = L^2$, $[g_{\mu\nu}] = [\eta_{\mu\nu}] = L^2$, $[x^\mu] = 1$, as in Eq. (1) of R. H. Dicke, "Mach's Principle and Invariance under Transformation of Units," *Physical*

3 General Reference Frames

In curved spacetimes, the notion of a global inertial frame must be replaced by more flexible constructions. Einstein—writing for general audiences—offered the vivid image of a "reference mollusk" as an intuitive way to think about curvilinear reference frames [4]:

> *What does it mean to assign to an event the definite coordinates x_1, x_2, x_3, x_4, if these coordinates have no significance in themselves? Closer reflection shows, however, that this concern is unfounded [. . .] One therefore uses non-rigid reference bodies which are not only in arbitrary motion as a whole, but also undergo arbitrary changes of shape during their motion. For the definition of time one uses clocks with arbitrary, however irregular, rates, which one is to imagine as fixed, each at a point of the non-rigid reference body, and which satisfy only the single condition that the simultaneously observable indications of spatially neighboring clocks differ only infinitesimally from one another. This non-rigid reference body, which one might not inappropriately call a "reference mollusk," is essentially equivalent to an arbitrary four-dimensional Gauss coordinate system. What gives the "mollusk," in comparison with the Gauss coordinate system, a certain intuitive vividness is the (actually unjustified) formal preservation of the separate existence of the spatial coordinates as opposed to the time coordinate. Every point of the mollusk is treated as a spatial point, and every material point at rest relative to it simply as at rest, so long as the mollusk is treated as a reference body. The general principle of relativity requires that all such mollusks can be used with equal right and equal success as reference bodies for the formulation of the general laws of nature; the laws themselves*

Review, vol. 125, no. 6, pp. 2163–2167, 1962. However, any other consistent choice of units is equally permissible. For example, consider SI units for Schwarzschild in standard coordinates (t, r, θ, ϕ): metric components must have varying units so that $[ds^2] = L^2$. Crucially, units on coordinates do *not* imply invariant measurement interpretations: the dimensions of t, r, and metric components do not change when crossing the horizon, even though the radial direction changes from spacelike to timelike. This highlights the difference between coordinates as labels—a role unaffected by unit choice—and attaching units to them as scale factors, which is permitted as long as physical relationships remain invariant. See also B. Crowell, *General Relativity*, Light and Matter, Fullerton, 2009, § 5.11.

are to be entirely independent of the choice of mollusk.

(author's translation from the German original)

A visual representation of this idea is shown in Fig. 2, which generalizes the familiar spatial lattice of SR.[11] The mollusk offers considerable flexibility through deformable spatial labeling and nonlinear clock functions—allowing independent warping of spatial and temporal coordinate directions.[12]

3.1 One Clock per Observer

In conceptual continuity with the earlier transition from a reference body to a distributed network of local observers in SR (see Fig. 1), we shift the focus from coordinate labels attached to a deformable body to the comoving observers themselves, modeled as a smooth timelike congruence covering the region of interest, a common model in GR [6, 8, 12, 7, 10, 11, 1].

For present purposes only the kinematic structure matters: each observer follows a distinct worldline, carries a clock, and is assigned fixed labels $\mathbf{y} = (y^1, y^2, y^3)$ that vary smoothly over the congruence so as to reflect the local neighborhood structure of the observers. For concreteness, let each observer carry an idealized standard clock whose mechanism is insensitive to acceleration—for example, an atomic clock. We use the reading τ of this clock as the seed time parameter along the worldline. Once the metric has been reconstructed below, one may identify the same empirical quantity with proper time in the usual way; for accelerated observers this is the standard clock-hypothesis idealization.[13] [14] In the spirit of Einstein's mollusk, whose clocks are not required to run at any prescribed rate, the displayed clock

[11] One may visualize the three-dimensional deformable reference body as nested, curved surfaces—like matryoshka shells. Alternatively, picture a deformable dough with no preferred foliation.

[12] Einstein's "Gauss coordinate system" (*Gaußsches Koordinatensystem*) refers to general curvilinear coordinates in the sense of Riemann's generalization of Gauss's surface theory—what would today be called a chart—not the more restrictive Gaussian normal (synchronous) form $ds^2 = -d\tau^2 + g_{ij}\, dx^i dx^j$.

[13] Strictly speaking, the seed frame's construction requires only that each observer's time label λ be a monotonic parameter along the worldline that varies coherently across the congruence. Since the programmed transformation can also absorb any reparametrization of λ, no generality is lost by simplifying the seed coordinates to $X^\mu_{\text{seed}} \equiv (\tau, \mathbf{y})$ instead of $(\lambda, \mathbf{y})$. The role of τ is then solely to provide a concrete physical prescription for how the carried clock could be operationally realized, via the clock hypothesis.

[14] Physics produces empirical facts for the swarm: zero accelerometer readings indicate

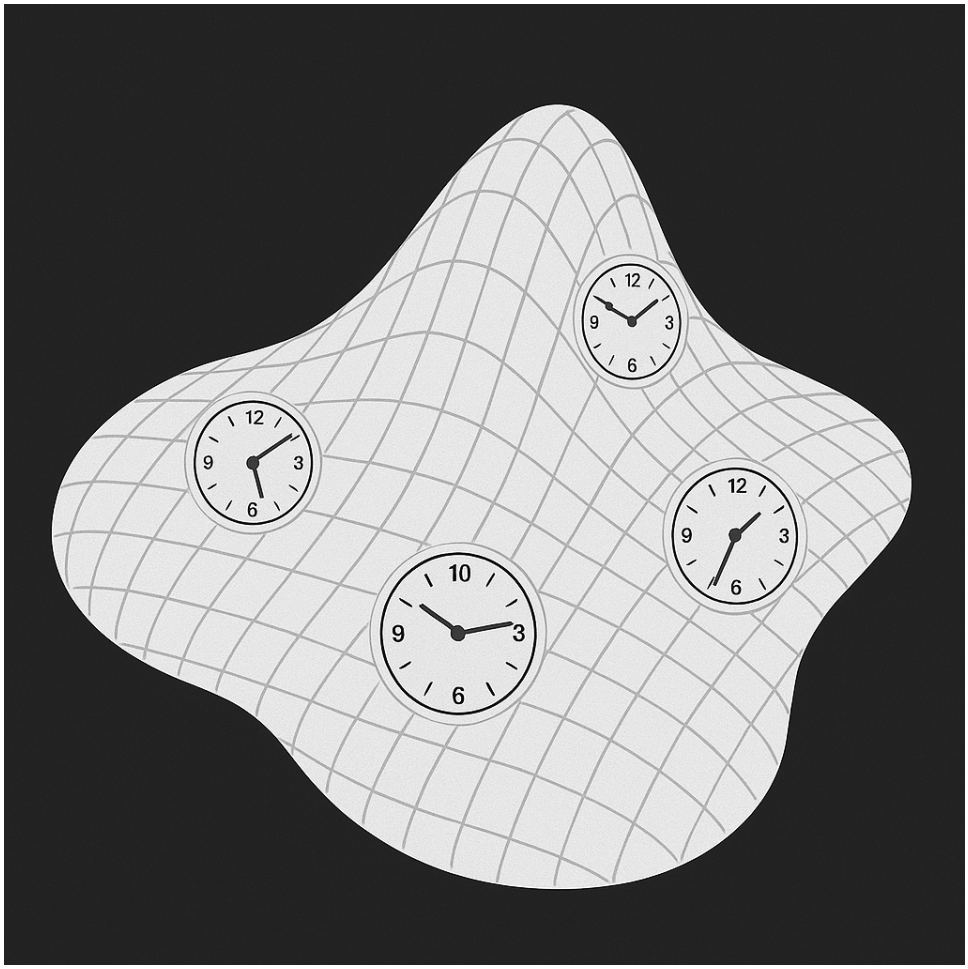

Figure 2: A two-dimensional slice through a deformable Einstein mollusk, representing a malleable reference body. Clocks with variable readings are positioned along warped spatial coordinates. This surface represents one of many ways to visualize the continuously deformable reference frame.

free fall, enabling local inertial frames; light propagation reflects causal structure; atomic clocks can operationalize proper time before metric knowledge (for details see section Assumptions in the Suppl. Mat.). Mathematically, in the geometrized viewpoint, the full theory encodes gravity through geometric notions, such as curvature. The adopted empirical approach begins by conceptually equipping spacetime only with material coordinates—observers must still subsequently measure the metric. Even given the constraints of this physical setting, programming freedom lets us imitate differential geometry's progression

label λ is taken to be an adjustable strictly monotonic function of τ along each worldline (since distinct events on a given worldline must receive distinct coordinates). In the explicit examples below we will simplify to the special case $\lambda = \tau$.

An ansatz for comoving coordinates established by these observers is thus

$$(x^0, x^1, x^2, x^3) = \big(\lambda(\tau, y^1, y^2, y^3),\, y^1,\, y^2,\, y^3\big). \tag{6}$$

As a whole, the setup must ensure that x^μ defines a valid coordinate system on the four-dimensional spacetime region $\mathcal{U}$ covered by the congruence. By itself, this ansatz reflects only the chosen timelike congruence of observers. If one interprets Einstein's mollusk more specifically as a reference *body*, then one adds the stronger requirement that equal-λ displacements $d\mathbf{y}$ between neighboring material points express locally spatial separations. The more general congruence-adapted form suffices here, however: each comoving clock rests at $dy^i = 0$, and only $d\lambda \neq 0$ along the worldline of the respective observer $\mathbf{y}$. The coordinate line of λ therefore coincides with the observer worldline.[15]

Can such a material reference frame *directly* provide *every* admissible coordinate system on the chosen region? As indicated in the introduction, the answer is no. The issue is not whether compatible charts would be mathematically related by coordinate transformations. The issue is whether the physically realized frame can by itself and on the spot reflect that chart freedom already during deployment. This question is central in GR pedagogy, but it arises equally in SR.

---

by a real-world construction—topological spaces to charts and atlases to metric–equipped manifolds.

[15] Eq. (6) defines coordinates adapted to a timelike congruence: one coordinate-line family, namely the curves $y^i = \text{const}$, coincides with the observer worldlines and is therefore timelike, with $ds^2 = g_{00} d\lambda^2 < 0$ for $dy^i = 0$; but this form does not by itself enforce any causal character for the equal-λ hypersurfaces $d\lambda = 0$. Reading Einstein's mollusk as a reference *body* with locally *spatially* separated points $\mathbf{y}$ and $\mathbf{y} + d\mathbf{y}$ results in the additional, stronger requirement that those leaves be locally spacelike: $ds^2 = g_{ij}\, dy^i dy^j > 0$ for nonzero spatial displacements $d\mathbf{y}$ at $d\lambda = 0$. In the standard matter-model sense, by contrast, "dust" means a pressureless perfect fluid with stress tensor $T_{ab} = \rho u_a u_b$, whose flow is geodesic when $\nabla_a T^{ab} = 0$ and $\rho \neq 0$. Here we assume only the weaker kinematic structure of a smooth, non-crossing timelike congruence of test observers with one clock and fixed labels per worldline; the congruence may consist of accelerated observers. For details and examples, see Suppl. Mat., Appendix A.

For an elementary example, let us therefore briefly return to Minkowski spacetime, where mollusk-type frames can of course also be constructed; in fact, the standard frame presented in Fig. 1 is a special instance of an Einstein mollusk. Consider now light-cone coordinates, which align with the paths of light rays:

$$u = \frac{x^0 + x^1}{\sqrt{2}}, \qquad v = \frac{x^0 - x^1}{\sqrt{2}}.$$

The coordinate basis vectors are null. Equivalently, infinitesimal displacements with $du \neq 0$, $dv = dx^2 = dx^3 = 0$ or with $dv \neq 0$, $du = dx^2 = dx^3 = 0$ satisfy $ds^2 = 0$, whereas the directions obtained by varying only x^2 or only x^3 remain spacelike. The coordinate system therefore has two null and two spacelike coordinate-basis directions. Such coordinates have no timelike coordinate-basis direction and therefore cannot be directly expressed by any congruence-adapted seed frame of the form (6). In such frames, the worldline of each massive observer is timelike, and the coinciding coordinate curve along which only the clock label changes must likewise be timelike. In other words, displacements with $dy^i = 0$ and $d\lambda \neq 0$ must satisfy $ds^2 < 0$. These incompatible coordinate conditions make it impossible that any such seed frame could directly display light-cone coordinates. This exemplifies the additional reference-frame-induced restrictions that Einstein himself called "actually unjustified" in his original mollusk description and motivates the generalization below.

3.2 From Fixed Labels to Programmable Coordinates

To make alternative chart choices directly available on a material swarm, we physically add a programmable layer to the labeling mechanism. The construction still rests on the congruence-adapted seed coordinates X^μ_{seed} defined by Eq. (6) on a spacetime region $\mathcal{U}$. The new ingredient is a co-deployed computational layer: each observer $\mathbf{y}$ carries a freely programmable portable device that computes and displays four coordinate values $X^\mu(X_{\text{seed}})$ from its own seed values alone. For the simplest seed ansatz $X^\mu_{\text{seed}} \equiv (\tau, \mathbf{y})$, this reduces to $X^\mu(\tau, \mathbf{y})$, as depicted in Fig. 3. Whenever the transformation $X^\mu(X_{\text{seed}})$ is a valid coordinate transformation on the chosen region $\mathcal{U}$—that is, a sufficiently smooth map with correspondingly smooth inverse onto its image and nonsingular local Jacobians—the displayed labels realize, in

the chosen idealizations, the numerical values of a valid chart on $\mathcal{U}$.[16] [17]

Additional coordinate systems can also be obtained by composing transformations: $\tilde{X}^\mu = \tilde{X}^\mu(X)$. Each observer's device evaluates this transform from its own current $X^\mu(\tau, \mathbf{y})$; since the same function $\tilde{X}^\mu(X)$ is applied across the entire congruence, the swarm as a whole realizes the four-dimensional coordinate transformation $(\tau, \mathbf{y}) \mapsto \tilde{X}^\mu\big(X(\tau, \mathbf{y})\big)$. In this scenario, knowledge of the programmed transform also allows each device to compute the local Jacobian on the spot.[18]

Operationally, the construction has three stages. *Pre-deployment* arranges a seed reference frame in which the fixed labels $\mathbf{y}$ reflect the physical neighborhood structure of observers and vary smoothly across the congruence—much like a dance ensemble whose performers learn their steps in advance so that, during the performance, each dancer maintains contact with the same

[16] Although $(\tau, \mathbf{y}) \mapsto X^\mu(\tau, \mathbf{y})$ may recall comoving cosmological coordinates, we assume no homogeneity, isotropy, or staticity (see Fig. 27.1 and Exercise 27.2 in Ref. [1]). The proper–time origin may be shifted $\tau \to \tau + \tau_0(\mathbf{y})$, fixed by any protocol—e.g., $\tau = 0$ when a reference signal reaches each observer (the dancer's loudspeaker cue). Such shifts are absorbed in $X^\mu(\tau, \mathbf{y})$.

[17] The programmable layer may also emulate scalar-field-based coordinate constructions. For example, if a metric hypothesis is available, the devices may be pre-loaded with values of four scalar fields satisfying the harmonic (de Donder) gauge condition (see Ref. [9], § 7.4; Ref. [3], § 4.3.1), or with intrinsic coordinates built from curvature invariants (see Ref. [13]). The latter would amount to an operational realization of such coordinates on a physical reference-frame apparatus. In such cases the programmable layer allows one to replace the gauge adaptation of the mollusk by a different gauge. This may highlight the nature of coordinate conditions: for example, even in Minkowski spacetime, standard spherical coordinates are not harmonic, illustrating that the harmonic gauge excludes a different class of useful coordinate systems than congruence-adapted seed frames, but for the same structural reason: a specific gauge adaptation constrains which coordinates can be directly displayed.

[18] C^k-diffeomorphisms locally imply nonzero Jacobian determinants: $\det\Big(\frac{\partial \tilde{X}^\mu}{\partial X^\nu}\Big) \neq 0$. For a consistency check, $\tilde{X}(X)$ is invertible with inverse $X(\tilde{X})$; we thus have $X(\tilde{X}(X)) = X$. Differentiating (chain rule) gives $\frac{\partial X^\lambda}{\partial \tilde{X}^\mu}\,\frac{\partial \tilde{X}^\mu}{\partial X^\nu} = \delta^\lambda{}_\nu$. Taking determinants and using $\det(AB) = \det(A)\det(B)$, we obtain $\det\Big(\frac{\partial X^\lambda}{\partial \tilde{X}^\mu}\Big)\det\Big(\frac{\partial \tilde{X}^\mu}{\partial X^\nu}\Big) = 1$, so neither vanishes. Likewise, if $(\tau, \mathbf{y}) \mapsto X^\mu(\tau, \mathbf{y})$ defines a valid chart, its Jacobian must be nonsingular. At each point, the Jacobian is the invertible 4×4 map relating coordinate differentials: $d\tilde{X}^\mu = \frac{\partial \tilde{X}^\mu}{\partial X^\nu}\, dX^\nu$. It can be *measured* from coordinate differentials in a four-dimensional neighborhood, see also Eq. (7). For pre-computed transforms, each observer can alternatively *directly compute* $\tilde{X}^\mu(X)$ and the Jacobian along their worldline $\mathbf{y}$. Since all devices implement the same $\tilde{X}^\mu(X)$, each observer computes intrinsically 4-D quantities.

local partners while the ensemble as a whole executes complex collective motion patterns. At the same stage, each observer receives a device pre-loaded with the coordinate functions $X^\mu(\tau, \mathbf{y})$. During *deployment*, the seed protocol instantiates the base coordinates $(\tau, \mathbf{y})$: the smooth, non-crossing timelike congruence provides dense coverage of the observation region $\mathcal{U}$; the seed congruence need not be geodesic; any such family of timelike worldlines covering $\mathcal{U}$ suffices. Each encountered event is detected, labeled with the computed coordinates $X^\mu(\tau, \mathbf{y})$ displayed at the moment of detection, and recorded. *Post-deployment*—possibly supplemented by *on-the-fly* information-carrying signal exchanges during deployment—consists of compiling and evaluating the recorded data across the swarm.

We can now put this "computationally enhanced swarm" to work in our SR example and see if it can display light-cone coordinates. We first set up the standard observers of Fig. 1 as our seed mollusk-like reference frame and equip each observer with a device programmed to initially display the respective standard SR coordinates x^μ on the screens: $X^0(\tau) = x^0 = \tau$, $X^i(\tau) = x^i = \text{const.}$, with $i = 1, 2, 3$. Consider now a secondary program which computes the coordinate transform: $\tilde{X}^0 := (X^0 + X^1)/\sqrt{2}$, $\tilde{X}^1 := (X^0 - X^1)/\sqrt{2}$, $\tilde{X}^2 := X^2$, $\tilde{X}^3 := X^3$. The very same swarm can thus directly display the previously inaccessible light-cone coordinates $\tilde{X}^\mu(X)$.

Crucially, this does not require superluminal signaling. In standard Minkowski language, the observer worldlines remain vertical and timelike, while the light-cone chart corresponds to a 45° tilt of the coordinate lines in the (X^0, X^1)-plane. Each lightlike coordinate line is therefore not traced by a single observer; it is *created* as a collectively preprogrammed pattern: different observers display the appropriate coordinate values at different proper times τ, without any signal traveling across the swarm during deployment.

The device screens may also simultaneously display the two distinct material coordinate systems $\tilde{X}^\mu$ and X^μ in different areas, thereby allowing for direct pointwise comparison. By collecting all these comparisons across a single swarm, the observers can also empirically reconstruct the coordinate transformation. For example, the coordinate curve with $d\tilde{X}^0 \neq 0$, $d\tilde{X}^i = 0$ can be reconstructed offline from a corresponding set of event records labeled by different observers at different proper times.

The same construction also works for the one-parameter family $\tilde{X}^0 = (\alpha\, X^0 + \alpha^{-1} X^1)/\sqrt{2}$, $\tilde{X}^1 = (\alpha\, X^0 - \alpha^{-1} X^1)/\sqrt{2}$, $\tilde{X}^2 = X^2$, $\tilde{X}^3 = X^3$. The light-cone case above corresponds to $\alpha = 1$. For $\alpha > 1$, a step in either $\tilde{X}^0$ or $\tilde{X}^1$ has a larger spatial than temporal component and therefore lies

outside the light cone, giving four spacelike coordinate directions. For $\alpha < 1$ both directions are timelike. All three regimes are implemented by the same swarm via a single programmed transformation with different parameter values. Notably, this entire range of causal characters for coordinate directions arises within SR alone, showing that the lifted coordinate conditions concern the adaptations implied by the mollusk's coordinatization protocol, not a property of spacetime or gravity.

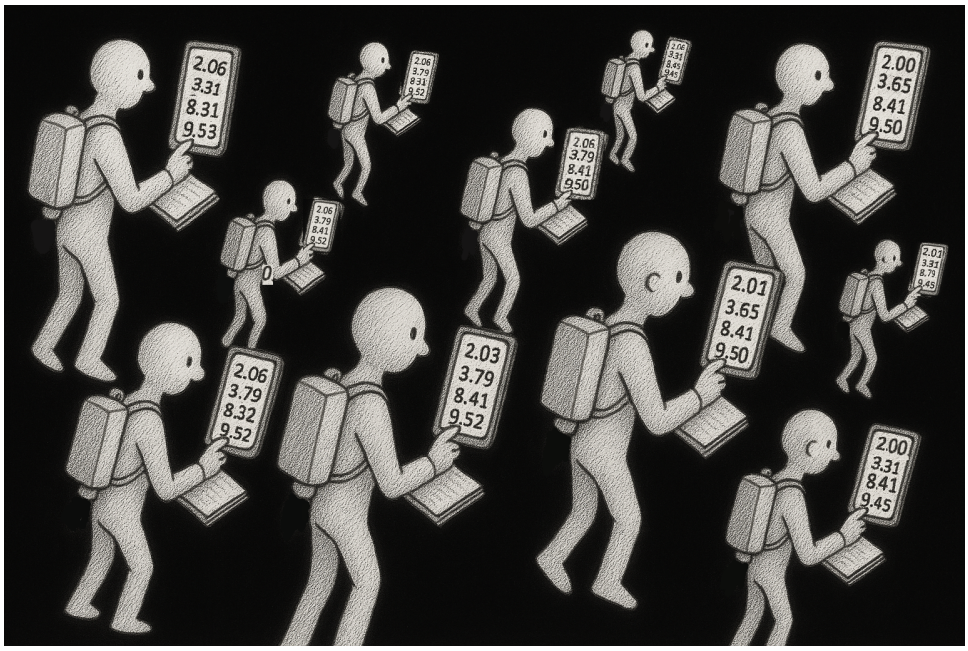

Figure 3: An observer swarm with programmable portable devices. Each device separately computes and displays four coordinate values X^μ as functions of its own seed coordinates $(\tau, \mathbf{y})$. Starting from these seed coordinates, the same swarm can display many alternative coordinate systems during deployment, obtained by evaluating any pre-programmed valid coordinate transformation on the seed region.

These examples also highlight a general point: each coordinate curve passing through a point in $\mathcal{U}$—obtained by varying one coordinate while holding the other three fixed—generically crosses multiple observer worldlines rather than following any single one. These curves must therefore be reconstructed from observations collected from different observers, typically at different proper times.

The flexibility of the programmable layer moves the enhanced apparatus one step beyond the traditional notion of a single reference frame, often tied to a single gauge adaptation. Operationally, the apparatus can mimic the effect of moving between charts of an atlas, even though the underlying material setup of observers remains unchanged.[19] Indeed, a single physical swarm can optionally be programmed to concurrently or separately compute and display n different coordinate systems $\tilde{X}^\mu_k$ $(k = 1, \ldots, n)$ in n separate screen windows.[20]

In this manner, the operational setup retraces differential geometry's logical progression even within a physical deployment. The displayed numbers collectively visualize coordinate systems before explicitly introducing further geometric structure. They carry no intrinsic metrical or causal meaning by themselves—they are coordinate labels shown on screens [3]. In particular, unlike the mollusk clock reading λ, an individual displayed component $X^\mu(\tau, \mathbf{y})$ need not vary monotonically with τ along a given worldline. What matters is only that the full map $(\tau, \mathbf{y}) \mapsto X^\mu(\tau, \mathbf{y})$ define a valid chart on the region under consideration.

Since the observer-assigned numbers X^μ are now displayed coordinate values that can serve as idealized operational stand-ins for mathematical coordinates x^μ, we henceforth adopt again lowercase notation x^μ for both.

The construction extends naturally to multiple swarms. Once coordinate values are computed and displayed on the devices of a single observer swarm, we may also deploy multiple independent swarms concurrently within the same observation region. In this setting, the coordinate transformation between reference frames is always empirically obtained: at each spacetime point, an observer in one swarm also records the coordinate values displayed by the coincident observer from the other swarm. The full transformation $\tilde{x}^\mu = \tilde{x}^\mu(x)$ can be operationally reconstructed by compiling such pointwise

[19] The transition-map structure is usually defined for atlases and uses overlapping regions $\mathcal{U}_i \cap \mathcal{U}_j$; here we restrict attention to a single patch $\mathcal{U}$ for simplicity. Extending the approach to multiple overlapping patches would require compatible programming on the overlaps, just as ordinary atlases require compatible transition maps. These concepts are commonly visualized for surfaces via maps to $\mathbb{R}^2$, analogous to cartography (for an alternative 2D surveyor visualization, see Ref. [5], § VII.4).

[20] Although the seed congruence fixes a particular gauge, the enhanced apparatus can concurrently display coordinates from differently gauge-adapted classes (e.g., congruence-adapted, harmonic, null) in separate screen windows during a single deployment. Gauge adaptations that may otherwise require a different seed reference frame can here be changed by (re-)programming.

comparisons across the region $\mathcal{U}$.

To determine empirical derivatives at any point, one assembles data from its (infinitesimally) small local four-dimensional neighborhood—either in the post-deployment evaluation phase or during deployment via local signal exchange with corresponding small temporal delays. The gathered data suffices to determine the Jacobian:

$$d\tilde{x}^\mu = \frac{\partial \tilde{x}^\mu}{\partial x^\nu}\, dx^\nu, \tag{7}$$

where both $d\tilde{x}^\mu$ and dx^ν represent coordinate differences between neighboring events as recorded by the respective observers. The Jacobian can be obtained either by differentiating the compiled transformation $\tilde{x}^\mu(x)$, or equivalently, by measuring coordinate differences and applying linear algebra—provided both swarms implement well-defined coordinate systems in the evaluated neighborhood. At finite resolution this would yield a local numerical approximation; the exact derivative is recovered in the dense-swarm idealization.

3.3 Measuring the Metric in GR

At this stage, the operational framework provides a natural starting point for introducing the general relativistic metric $g_{\mu\nu}(x)$ in the observer-swarm coordinates. Importantly, when following the standard metric reconstruction procedure, students are no longer tied to the seed protocol's coordinate conditions; they may work with any chosen programmed coordinate system whose validity has been established on the relevant region. The following steps mirror the reconstruction of a metric on curved surfaces by comparing orthonormal coordinates on locally flat patches to general curvilinear charts. In the present operational setting, a structurally analogous approach is implemented in spacetime itself by comparing local, inertial coordinates (defined via the equivalence principle) to physically realized programmable coordinates. This picture goes back to Einstein [33], was explained again by Born [5], and most succinctly summarized by Misner, Thorne and Wheeler [1], p. 400:

> *For example, to "measure" the metric near a given event, one typically lays out a latticework of rods and clocks (local orthonormal frame, small enough that curvature effects are negligible), and uses it to determine the interval between neighboring events.*

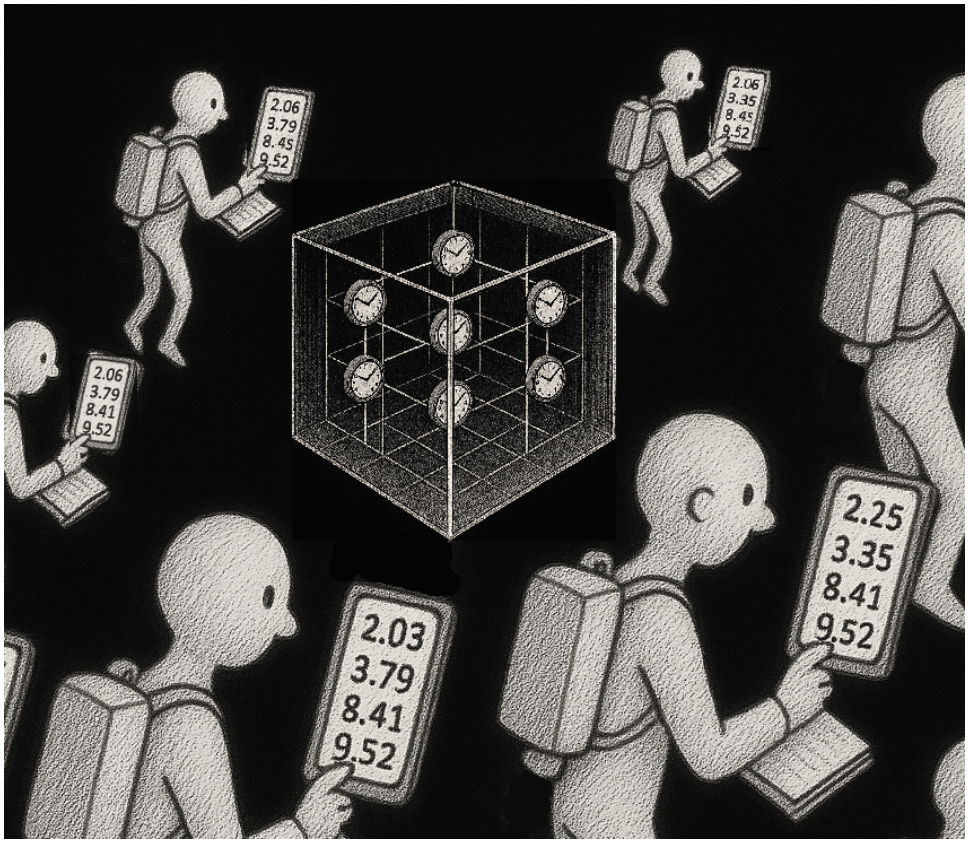

Figure 4: Two coexisting reference frames: a GR observer swarm whose programmable devices display computed coordinates, and a freely falling Einstein elevator featuring local inertial orthonormal coordinates. By comparing the coordinate values assigned by the two frames to the same events, one can reconstruct the metric components in the programmable coordinates.

Indeed, at any event on an observer's worldline one can introduce a momentarily comoving freely falling frame—the local inertial frame of Einstein's equivalence principle, often pictured as his "elevator," but imagined without physical walls so that observers from both frames are interspersed throughout the same small region (Fig. 4). The equivalence principle[21] ensures that,

[21]Contemporary accounts distinguish between multiple formulations of the equivalence principle, which differ in empirical scope and conceptual function. We here employ the Einstein Equivalence Principle (EEP), which ensures that freely falling frames are locally indistinguishable from inertial frames in special relativity, see D. Lehmkuhl, "The

at that event and to first order in sufficiently small displacements, physics in this tiny transient frame is indistinguishable from special relativity. One can therefore set up a local inertial orthonormal coordinate system $\tilde{x}^\mu$ inside the elevator, using a small latticework of observers as in Fig. 1, centered on the event and valid only within the elevator's extent.

Because both sets of observers—those belonging to the general relativistic swarm and those belonging to the local inertial frame—occupy the same small region, they can record the same local occurrences. At each such occurrence, the corresponding GR observer assigns displayed coordinate values x^μ while the elevator observer assigns local inertial coordinates $\tilde{x}^\mu$, and each reads off the values displayed by their counterpart. Compiling these pointwise comparisons yields a local coordinate transformation $\tilde{x}^\mu = \tilde{x}^\mu(x)$, valid only throughout the elevator but otherwise an ordinary coordinate transformation.

Since the $\tilde{x}^\mu$ coordinates belong to an orthonormal local inertial frame, the interval written in these coordinates has direct operational meaning. In particular, $ds^2 = \eta_{\mu\nu}\, d\tilde{x}^\mu\, d\tilde{x}^\nu$ remains valid at the considered event and to excellent approximation within the elevator. Substituting the coordinate-differential relation (7) then expresses the interval in the general coordinates:

$$ds^2 = \eta_{\mu\nu} \frac{\partial \tilde{x}^\mu}{\partial x^\alpha} \frac{\partial \tilde{x}^\nu}{\partial x^\beta}\, dx^\alpha\, dx^\beta \equiv g_{\alpha\beta}(x)\, dx^\alpha\, dx^\beta, \tag{8}$$

where we define the metric components in the swarm coordinates:

$$g_{\alpha\beta}(x) = \frac{\partial \tilde{x}^\mu}{\partial x^\alpha} \frac{\partial \tilde{x}^\nu}{\partial x^\beta}\, \eta_{\mu\nu}. \tag{9}$$

The general coordinates x^μ alone carry no immediate physical meaning—they are the numbers shown on the swarm's screens. Physical content enters through the metric $g_{\alpha\beta}(x)$, which is reconstructed from local inertial measurements together with the Jacobian of the transformation between the two coordinate systems. Carrying out this procedure at each event throughout the observation region $\mathcal{U}$ determines the metric components in the displayed coordinates everywhere on $\mathcal{U}$. Equivalently, one may bypass the explicit construction of the Jacobian (7): sufficiently small displacements $d\tilde{x}^\mu$ in the elevator frame and the corresponding coordinate differences dx^μ in the general frame, both locally approximating the infinitesimal limit, can be plugged

equivalence principle(s)," in *The Routledge Companion to Philosophy of Physics*, edited by E. Knox and A. Wilson (Routledge, London, 2021), pp. 125–144.

directly into the defining relation $g_{\alpha\beta}\,dx^\alpha\,dx^\beta \equiv \eta_{\mu\nu}\,d\tilde{x}^\mu\,d\tilde{x}^\nu$, and standard linear algebra yields the metric components $g_{\alpha\beta}$ [33, 1].

Since the local GR metric $g_{\alpha\beta}(x)$ has so far been reconstructed from the local SR coordinates via Eq. (9), the causal classification is inherited from the local inertial frame. The SR observers classify infinitesimal displacements $d\tilde{x}^\mu$ as timelike, null, or spacelike by evaluating Eq. (5). Because the interval ds^2 is invariant under the transformation, the same value of ds^2 applies to the corresponding displacement dx^α expressed in the general coordinates. The causal structure in the GR coordinates can therefore be operationally obtained by transferring the local SR classification through the invariant interval.[22] This provides a useful practical insight, particularly in spacetime regions with potentially misleading coordinate labels (such as the standard Schwarzschild interior). Regardless of naming conventions, the causal character of any direction dx^α can be determined locally from the measured GR metric $g_{\alpha\beta}(x)$ via Eq. (8).

Consider now two different GR coordinate systems describing the same spacetime region, related by $x^\mu = x^\mu(x')$. At any event, each can measure the metric by constructing a local "elevator." Since all freely falling frames at a point are locally inertial, their local orthonormal coordinate descriptions, at that event and to first order, are related by Lorentz transformations, which preserve the spacetime interval ds^2. Therefore, both GR frames will also yield the same invariant:

$$g_{\mu\nu}(x)\,dx^\mu dx^\nu = g'_{\alpha\beta}(x')\,dx'^\alpha dx'^\beta. \tag{10}$$

Since the coordinate differentials transform via the Jacobian (Eq. (7)), this invariance requires:

$$g'_{\alpha\beta}(x') = \frac{\partial x^\mu}{\partial x'^\alpha}\frac{\partial x^\nu}{\partial x'^\beta}\,g_{\mu\nu}\big(x\big), \tag{11}$$

which is precisely the component transformation law of a (0,2) tensor, Eq. (9) being a special case. Thus, the metric's tensorial nature emerges from a physical requirement: all coordinate systems must agree on spacetime intervals measured in local inertial frames.

[22] An alternative elementary route to local metric components uses light-cone structure: observing light rays and their coordinate directions dx^α. Since null directions satisfy $ds^2 = g_{\alpha\beta}(x)\,dx^\alpha dx^\beta = 0$, measuring sufficiently many independent directions dx^α determines $g_{\alpha\beta}(x)$ up to a conformal factor; see Ref. [34]. The remaining conformal ambiguity must be fixed; e.g., by measuring a proper time interval $d\tau$. This is operationally simpler but obscures the relation to local inertial frames. The programmable layer remains applicable in this scenario as well.

It is worth emphasizing, however, that we have so far only scratched the surface of differential geometry. Physics—not abstract preference—motivates further mathematical development: local metric measurement employs the equivalence principle, leading to a procedure conceptually analogous to using small, flat patches on curved surfaces to measure the metric locally. In this sense, physics has led us naturally to the tools of differential geometry, not the other way around.

Moreover, it was gravity that motivated us to engage in local metric reconstruction. It may therefore come as no surprise that local variations in the metric tensor $g_{\alpha\beta}(x)$ also play a central role in the gravitational field equations themselves. It is, however, instructive to contrast this intuition with curvilinear coordinates in special relativity, which also yield locally varying metric components but a vanishing Riemann curvature tensor [31, 9]. Hence, within the geometrized viewpoint adopted here, the operational framework naturally suggests exploring field equations in which gravitating sources give rise to nonzero spacetime curvature, the empirically successful route historically followed by Einstein [33].[23] [24]

Conclusion

We have presented a framework that connects general coordinate systems to idealizations of physically realizable operations by equipping local observers with portable devices that display four programmable numbers per chart. This setup is intended to help build intuition for general coordinate freedom. Several implications—though straightforward in hindsight—may initially contrast with acquired intuitions:

[23] The same reference-frame and local metric-reconstruction procedure would also work in the absence of gravitating sources. Moreover, whether gravity is described geometrically or as fields or other fundamental objects on a fixed background is a separate issue. In standard GR pedagogy, all freely falling bodies in the presence of gravitating sources follow geodesics of a dynamical metric with torsion-free Levi-Civita connection. Within that geometrized viewpoint, a vanishing Riemann tensor implies (locally, and globally under mild topological assumptions) equivalence to flat Minkowski spacetime; nonzero spacetime curvature is then the hallmark of gravitating sources. In theories with independent torsion or additional structure, seed frames may still be enhanced through a programmable layer, but the detailed relation between free fall, connection, and metric would differ from theory to theory.

[24] Gravitational effects of the dense reference frame itself were ignored, per the zero-density limit: it represents only passive test matter, not a source (see also [3], § 4.3.1).

- *Programming freedom can overcome seed-frame constraints*: for each programmed transformation that is a valid chart transition on $\mathcal{U}$, the displayed coordinates implement valid coordinates that need not inherit the protocol-dependent constraints of the seed reference frame.

- *A material computational layer genuinely enlarges the frame*: the co-deployed computing and display devices are physical parts of our reference-frame apparatus, while the numerical values shown on their screens remain conventional coordinate labels. In this way, a material frame may by itself emulate the chart freedom of differential geometry as closely as one may reasonably expect in a physical implementation.

- *Charts precede the metric*: whenever the seed frame can be specified before metric measurement, adding a computational layer does not destroy this property. The enhanced frame can then also emulate chart freedom before metric measurement, mirroring the logical progression of differential geometry.

- *Coordinate directions need not mirror observer motion*: although the observers travel on timelike worldlines, the displayed coordinates may have null or spacelike coordinate directions.

- *One swarm suffices*: a single observer swarm can compute and display coordinates of several coordinate systems in parallel, physically mimicking the mathematical operation of moving between charts.

- *Programmable devices reproduce alternative coordinate prescriptions*: familiar coordinate choices may arise as particular programming options on the same seed frame, as may more specialized constructions such as harmonic coordinates or coordinates built from curvature invariants when a metric hypothesis is assumed.

The framework adds an explicit computational layer to an elementary material seed frame, thereby making pedagogically transparent the distinction between raw operational labels and the freely chosen coordinate labels permitted by the coordinate freedom of GR.

In this sense, the "computationally enhanced mollusk" is best viewed as an enhanced teaching device that makes standard chart transitions operationally vivid. Its value lies in connecting a concrete swarm of observers with

the abstract atlas viewpoint, thereby offering another tool for bridging mathematical formalism and physical measurement procedures in the teaching of coordinate freedom in GR. In the chosen setting (dust-frames enhanced by "smartphones"), even beginners need only associate the freedom to program a familiar device with the freedom to choose valid coordinates—an association that, once grasped, may later help them interpret more sophisticated reference-frame constructions as well.

References

[1] Misner, Ch. W. and Thorne, K. S. and Wheeler, J. A. *Gravitation*. W. H. Freeman, San Francisco, CA, 1973.

[2] James B. Hartle. *Gravity: An Introduction to Einstein's General Relativity*. Addison-Wesley, San Francisco, 2003.

[3] Westman, H. and Sonego, S. Coordinates, Observables and Symmetry in Relativity. *Annals of Physics*, 324(8):1585–1611, 2009.

[4] Einstein, A. *Über die spezielle und die allgemeine Relativitätstheorie (Gemeinverständlich)*. Vieweg & Sohn, Braunschweig, 1917. English translation: *Relativity: The Special and the General Theory — A Popular Exposition*, Methuen & Co. Ltd., 1920.

[5] Max Born. *Die Relativitätstheorie Einsteins*. Springer, Berlin, 1920. English transl.: *Einstein's Theory of Relativity* (E. P. Dutton & Co., New York, 1922).

[6] L. D. Landau and E. M. Lifshitz. *The Classical Theory of Fields*, volume 2 of *Course of Theoretical Physics*. Pergamon Press, Oxford, 3rd rev. ed. edition, 1971.

[7] Brown, J. D. and Kuchař, K. V. Dust as a standard of space and time in canonical quantum gravity. *Physical Review D*, 51(10):5600–5629, 1995.

[8] Rovelli, C. What is observable in classical and quantum gravity? *Classical and Quantum Gravity*, 8(2):297–316, 1991.

[9] Weinberg, Steven. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. John Wiley & Sons, New York, NY, 1972.

[10] Petr Hájíček. *An Introduction to the Relativistic Theory of Gravitation*, volume 750 of *Lecture Notes in Physics*. Springer, Berlin, Heidelberg, 2008.

[11] R. K. Sachs and H. Wu. *General Relativity for Mathematicians*. Springer-Verlag, New York, 1977.

[12] Norton, J. D. General covariance and the foundations of general relativity: Eight decades of dispute. *Reports on Progress in Physics*, 56:791–858, 1993.

[13] Peter G. Bergmann and Arthur Komar. Poisson brackets between locally defined observables in general relativity. *Physical Review Letters*, 4(8):432–433, 1960.

[14] Carlo Rovelli. GPS observables in general relativity. *Physical Review D*, 65:044017, 2002.

[15] Kristina Giesel and Thomas Thiemann. Scalar material reference systems and loop quantum gravity. *Classical and Quantum Gravity*, 32(13):135015, 2015.

[16] Bianca Dittrich. Partial and complete observables for canonical general relativity. *Classical and Quantum Gravity*, 23(22):6155–6184, 2006.

[17] Josep M. Pons, Donald C. Salisbury, and Kurt A. Sundermeyer. Observables in classical canonical gravity: Folklore demystified. In *Journal of Physics: Conference Series*, volume 222, page 012018, 2010.

[18] J. Brian Pitts. Equivalent theories redefine hamiltonian observables to exhibit change in general relativity. *Classical and Quantum Gravity*, 34(5):055008, 2017.

[19] Johannes Tambornino. Relational observables in gravity: A review. *SIGMA*, 8:017, 2012.

[20] Donald Salisbury. A history of observables and hamilton–jacobi approaches to general relativity. *The European Physical Journal H*, 47(1):7, 2022.

[21] John Earman. *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*. MIT Press, Cambridge, MA, 1989.

[22] Craig Callender and Carl Hoefer. Philosophy of space-time physics. In Peter Machamer and Michael Silberstein, editors, *The Blackwell Guide to the Philosophy of Science*, pages 173–198. Blackwell / John Wiley & Sons, 2002.

[23] Oliver Pooley. Substantivalist and relationalist approaches to spacetime. In Robert Batterman, editor, *The Oxford Handbook of Philosophy of Physics*, pages 522–586. Oxford University Press, Oxford, 2013.

[24] John Stachel. The hole argument and some physical and philosophical implications. *Living Reviews in Relativity*, 17(1):1, 2014.

[25] James Owen Weatherall. Regarding the hole argument. *The British Journal for the Philosophy of Science*, 69(2):329–350, 2018.

[26] Á. Mozota Frauca. GPS observables in Newtonian spacetime or why we do not need 'physical' coordinate systems. *European Journal for Philosophy of Science*, 14:51, 2024.

[27] Giovanelli, M. Nothing but coincidences: the point-coincidence and Einstein's struggle with the meaning of coordinates in physics. *European Journal for Philosophy of Science*, 11(2):45, 2021.

[28] Wald, Robert M. *General Relativity*. University of Chicago Press, Chicago, USA, 1984. See especially Ch. 8.

[29] Robert Geroch. *General Relativity from A to B*. University of Chicago Press, Chicago, 1978.

[30] Edwin F. Taylor and John Archibald Wheeler. *Spacetime Physics: Introduction to Special Relativity*. W. H. Freeman, New York, 2 edition, 1992.

[31] Bernard F. Schutz. *A First Course in General Relativity*. Cambridge University Press, Cambridge, 2 edition, 2009.

[32] Wolfgang Rindler. *Relativity: Special, General, and Cosmological*. Oxford University Press, Oxford, 2001.

[33] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7):769–822, 1916.

[34] Stephen W. Hawking and George F. R. Ellis. *The Large Scale Structure of Space-Time*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 1973.

Supplementary Material for:

From One Clock to Programmable Coordinates: An Elementary Material Reference Frame for General Relativity

This supplementary document contains additional technical material—including derivations and detailed arguments—that support but are not essential to the main pedagogical narrative. Instructors may find concrete examples for advanced undergraduate or graduate courses. Specific references to this document are made in the endnotes of the main text. Throughout these appendices, unless otherwise noted, we adopt geometrized units with $c = G = 1$, as in the main text (from section 3 onwards). The sign convention remains $(-+++)$. For readers accustomed to differential-geometric language, the following correspondences may help clarify the physics notation used here: the line element $ds^2 = g_{\mu\nu}\, dx^\mu\, dx^\nu$ encodes the metric tensor $\mathbf{g} = g_{\mu\nu}\, dx^\mu \otimes dx^\nu$; "setting all differentials to zero except dx^μ" stands for evaluating $\mathbf{g}(\partial_\mu, \partial_\mu) = g_{\mu\mu}$, with no sum on μ. The causal character of a coordinate direction is diagnosed via the covariant diagonal component $g_{\mu\mu}$; the contravariant component $g^{\mu\mu}$ carries different geometric information whenever off-diagonal terms are present. Appendix B.4.1 discusses this distinction.

Contents

Assumptions **2**

Appendix A: Congruence-adapted frames vs. programming **4**
A.1 Seed frame: mollusk and timelike congruence interpretation 4
A.2 Double-null Kruskal–Szekeres and seed-frame limits 9
A.3 ADM charts with timelike grid lines can be instantiated by a single mollusk 10

Appendix B: Implementations of programmable layers **11**
B.1 Linear coordinate transformations in SR 11
B.2 Swapping the seed: rotating frames showing inertial coordinates in SR 15
B.3 GR Scenarios: with and without metric hypothesis 18
B.4 Interior Painlevé–Gullstrand Coordinates: four spacelike directions 20
B.5 Kruskal–Szekeres Coordinates: two lightlike coordinate basis directions 26

Appendix C: Mathematical questions for guided exploration **28**

Assumptions

Valid seed frame

We assume the existence of a *valid seed frame* on a spacetime region $\mathcal{U}$: a dense congruence of smooth, non-crossing, timelike worldlines, each carrying a clock, together with labels $\mathbf{y}$ that vary across the congruence, such that the seed coordinates $(\tau, \mathbf{y})$ define a valid chart on $\mathcal{U}$. This entails the following subsidiary idealizations, which are standard for such idealized observer-swarm constructions and mutually competing in any attempted real-world implementation.

- *Test matter.* The observer swarm is treated in the zero-density limit: it detects and labels events but does not source the gravitational field ([1], § 4.3.1). In practice one needs a compromise between sufficient density for coverage and sufficient diluteness for passivity, with correspondingly limited observation regions. This assumption also excludes self-generated gravitational effects of the swarm itself, such as self-collapse, self-induced caustics, or gravitational radiation emitted by the swarm. It does not exclude the possibility that the background spacetime being probed already contains such features.
- *Dense coverage and event observability.* Every event in $\mathcal{U}$ can be detected and labeled by an observer via point-coincidence along that observer's worldline. Strictly speaking, this idealization is self-contradictory: a fully dense frame would leave no room for other events to occur. A realistic implementation would use a sparse swarm—much as the SR lattice of rods and clocks leaves room for the events it is meant to record—with coordinates being reconstructed by interpolation and local signal exchanges. We adopt the dense idealization because it allows us to imagine observers present directly at each event, avoiding the need to reconstruct events from indirect traces. Note that we still conceptually restrict attention to events occurring within the observation region $\mathcal{U}$ covered by the congruence, unlike typical astrophysical observations where events of interest occur far from the observers.
- *Sufficiently smooth congruence and labeling.* The worldlines are sufficiently smooth and timelike, and the labels $\mathbf{y}$ vary sufficiently smoothly across the congruence. Distinct worldlines do not pass through the same event (no caustics). The seed congruence need not be geodesic: accelerated observers are permitted, provided they remain test matter and the worldlines remain timelike, sufficiently smooth, and non-crossing on $\mathcal{U}$, ensuring full coverage of the observation region. We exclude both caustics arising from the spacetime geometry (e.g., geodesic focusing) and caustics resulting from inadequate steering of accelerated observers.
- *Avoidance of singularities.* The region $\mathcal{U}$ is restricted to avoid singular boundaries where unbounded tidal forces destroy any physical apparatus.
- *Data collection.* Post-deployment compilation of recorded data across the swarm is assumed feasible, possibly supplemented by on-the-fly signal exchanges with small temporal delays.
- *Horizon crossings.* A horizon crossing is locally undetectable and harmless provided post-deployment data collection inside the horizon can be assured—for example, by choosing a compilation event suitably in the future of $\mathcal{U}$ and sufficiently remote from any singularity. Appendix B.4 illustrates this.
- *Atlas structure.* Larger regions may be covered by overlapping swarms on separate patches with compatible programming on the overlaps. We do not develop such multi-patch scenarios in detail: the text requires only a single working region $\mathcal{U}$.

- *Proper-time clock and clock hypothesis.* Each observer carries an idealized standard clock (e.g., an atomic clock) with identical calibration conventions across the swarm. The reading τ is taken operationally as proper time along that observer's worldline; this does not presuppose knowledge of the metric tensor. For accelerated observers, we adopt the standard clock hypothesis: the carried clock agrees, to first order, with an auxiliary standard clock at rest in a momentarily comoving freely falling frame. Local inertial comparison frames are assumed available in principle by the equivalence principle, also in the presence of gravity. Only after the metric has been reconstructed is the same empirical quantity represented by $d\tau^2 = -g_{\mu\nu}\, dx^\mu dx^\nu = -ds^2$.

Valid programmable layer

Whenever the transformation $X^\mu(X_{\text{seed}})$ is a valid coordinate transformation on the chosen region $\mathcal{U}$—that is, a sufficiently smooth map with correspondingly smooth inverse onto its image and nonsingular local Jacobians—the displayed labels realize, in the chosen idealizations, the numerical values of a valid chart on $\mathcal{U}$. The programmable layer assumes that each observer's device can evaluate such a transformation from its own seed values and display the result. This entails:

- *Test matter.* The co-deployed computing and display devices are likewise treated as test matter: they add no gravitational source term beyond the seed frame's zero-density idealization.

- *Local computation.* Each device computes $X^\mu(\tau)\big|_{\mathbf{y}}$ from its own seed values alone; no inter-observer communication is required during deployment. Pre-deployment ensures consistency (e.g., identical programs loaded onto all devices).

- *Effective specifiability.* The programmed transformation is effectively specified (e.g., by a closed formula or a well-defined algorithm). We need not commit to a specific computational paradigm (digital, analog, hybrid, etc.).

- *Chart validity, case by case.* Validity of the displayed coordinates must be established for each programmed transformation, just as for any coordinate construction in GR. For present purposes, any chart $\tilde{X}$ on $\mathcal{U}$ compatible with a given seed chart X is related to it by the usual transition-map structure of an atlas: $\tilde{X} = \Phi \circ X$, where Φ is a C^k diffeomorphism between open subsets of $\mathbb{R}^4$ mapping $X(\mathcal{U})$ onto $\tilde{X}(\mathcal{U})$, with $k \geq 1$. Thus every compatible chart arises from some valid Φ over the seed chart image. In the present programmable framework, one analogously chooses within the class of programmed transition functions from the coordinates provided by a seed reference frame, with chart validity checked for the case at hand, as exemplified in Appendix B. Any finite-precision discrete implementation only approximates these continuum conditions at the adopted numerical resolution, which may be increased as required for the cases studied. The pedagogical analogy therefore concerns the practically relevant class of programmed valid coordinate transforms, not an unrestricted notion of arbitrary programmability (for a mathematical exploration of possible limits of this analogy, see Appendix C).

  Moreover, the working region $\mathcal{U}$ is understood to lie in a single chart patch for simplicity; for example, in a single angular chart patch whenever spherical or cylindrical angular coordinates are used.

Appendix A: Congruence-adapted frames vs. programming

This appendix analyzes which coordinate systems can be *directly* produced by three successively more flexible material reference-frame constructions: Einstein's physical mollusk, a congruence-adapted seed frame, and a seed reference frame enhanced by a programmable computational layer.

A.1 Seed frame: mollusk and timelike congruence interpretation

All seed constructions considered here share a common kinematic starting point: a dense congruence of non-intersecting timelike worldlines, each carrying a clock that displays a value λ monotonically changing with proper time τ. The remaining three labels $\mathbf{y} = (y^1, y^2, y^3)$ vary smoothly across the congruence but are constant along each worldline, and thereby identify the worldline. The adapted coordinates are

$$(x^0, x^1, x^2, x^3) = (\lambda(\tau, \mathbf{y}),\ y^1, y^2, y^3) \tag{1}$$

assumed to define a valid base coordinate system. Once the metric $g_{\mu\nu}$ has been measured, we can write down a line element

$$ds^2 = g_{\mu\nu}(x) dx^\mu dx^\nu,$$

knowing that ds^2 could also be measured in each momentarily comoving local inertial frame. By a *coordinate line* (or, infinitesimally speaking, a *coordinate (basis) direction*) we mean the curve obtained by varying one coordinate x^α while holding the other three fixed ($x^\beta = const., \beta \neq \alpha$). Its causal character at any point is determined by the sign of the quadratic form $g_{\mu\nu} dx^\mu dx^\nu$ evaluated on the corresponding coordinate basis vector, i.e. with $dx^\alpha \neq 0$ and $dx^\beta = 0$ for all $\beta \neq \alpha$. Any actual motion of a massive material object would have to follow a timelike worldline, $ds^2 < 0$. The coordinate-line "steps" discussed below need not themselves be physically traversable; they are infinitesimal bookkeeping displacements reconstructed from neighboring recorded events across the observation region $\mathcal{U}$. Spacelike and null coordinate basis directions can be analyzed even though no massive observer travels along them.

- **Timelike time direction (observer worldlines):**

  Because the mollusk and congruence-adapted seed frames are conceptually made of real (massive) observers carrying real clocks (albeit under the test-matter assumption), they impose a simple physical restriction on the causal nature of the coordinate basis directions that can occur. Observers sit at fixed labels $\mathbf{y} = (y^1, y^2, y^3)$, so along any one observer worldline we have

  $$dy^1 = dy^2 = dy^3 = 0,$$

  and only the time label changes. For $d\lambda \neq 0$ along such a worldline the line element reduces to[1]

  $$ds^2 = g_{00}\,(d\lambda)^2, \tag{2}$$

  because all terms involving dy^i vanish on that step. These temporal steps along the worldline of a massive observer must be timelike in the comoving local inertial frame $ds^2 < 0$ and, given that ds^2 is invariant, satisfy $ds^2 < 0$ also in the GR coordinates and their metric. Therefore we must have

  $$g_{00} < 0. \tag{3}$$

[1] This "small-step" test—setting all differentials to zero except one and reading off the causal character from the sign of the corresponding diagonal metric component—is used throughout these appendices. It probes the covariant diagonal component $g_{\mu\mu}$, which characterizes the coordinate basis vector. See the footnote in Appendix B.4.1 for a discussion in the context of Painlevé–Gullstrand coordinates.

In fact, we can say more. Along a fixed observer worldline $\mathbf{y} = \text{const}$, the clock relation $\lambda = \lambda(\tau, \mathbf{y})$ implies $d\tau = \left(\frac{\partial \tau}{\partial \lambda}\right)_{\mathbf{y}} d\lambda = \left(\frac{\partial \lambda}{\partial \tau}\right)_{\mathbf{y}}^{-1} d\lambda$. Since $ds^2 = -d\tau^2 = g_{00}\, d\lambda^2$ on that worldline, we obtain

$$g_{00}(\lambda, \mathbf{y}) = -\left(\frac{\partial \tau}{\partial \lambda}\right)_{\mathbf{y}}^{2} = -\left(\frac{\partial \lambda}{\partial \tau}\right)_{\mathbf{y}}^{-2} < 0.$$

Strictly speaking, this is the *only* causal-structure consequence of the ansatz (1). The most general metric compatible with the adapted-coordinate ansatz therefore takes the form

$$ds^2 = -\left(\frac{\partial \lambda}{\partial \tau}\right)_{\mathbf{y}}^{-2} d\lambda^2 + 2\, g_{0i}\, d\lambda\, dy^i + g_{ij}\, dy^i\, dy^j. \tag{4}$$

The corresponding four-velocity of the observer on worldline $\mathbf{y}$ is

$$u^\mu(\tau, \mathbf{y}) = \left.\frac{dx^\mu}{d\tau}\right|_{\mathbf{y}} = \left(\frac{\partial \lambda}{\partial \tau}, 0, 0, 0\right), \tag{5}$$

so $u^0 = (\partial\lambda/\partial\tau)_{\mathbf{y}}$ along that worldline. In these equations, we do not yet impose *a priori* restrictions on $g_{0i}(\lambda, \mathbf{y})$ or $g_{ij}(\lambda, \mathbf{y})$. The off-diagonal terms g_{0i} encode the tilt between the observer worldlines and the equal-λ surfaces, and g_{ij} is the induced metric on those leaves.[2]

For the following discussion we explicitly distinguish two readings of Einstein's mollusk passage quoted in the main text. The weaker reading yields the most general congruence-adapted seed frame. The stronger reading treats the mollusk literally as a reference body whose points are spatial points. The programmable layer is applicable to both interpretations as possible seed frames.

• **Weak reading (congruence-adapted reference frame).**

A congruence of clock-carrying observers, each labeled by $\mathbf{y}$, supplies only a timelike threading: the curves $\mathbf{y} = \text{const}$ are worldlines of massive observers, so only Eq. (4) holds. This is the minimal kinematic structure common to all congruence-adapted seed-frame constructions. The labels $\mathbf{y}$ are conventional names and $\lambda(\tau, \mathbf{y})$ is the single programmable clock per observer $\mathbf{y}$. But at this stage nothing would require the equal-λ leaves to have any particular causal character.

Indeed, a 1+1 Minkowski illustration quickly shows that the ansatz $(\lambda(\tau, y), y)$ requires at least one timelike coordinate basis direction but does not force a space versus time split by itself across the congruence. With the standard Minkowski line element $ds^2 = -d\tau^2 + dy^2$ and $\lambda := \tau + \beta y$, one finds $ds^2 = -d\lambda^2 + 2\beta\, d\lambda\, dy + (1 - \beta^2)\, dy^2$. Hence, at fixed λ, $ds^2 = (1 - \beta^2)\, dy^2$. Therefore the equal-$\lambda$ direction dy with $d\lambda = 0$ is spacelike for $|\beta| < 1$, null for $|\beta| = 1$, and timelike for $|\beta| > 1$. For example, $\beta = 0, 1, 2$ give spacelike, null, and timelike cases, respectively. In all cases, λ remains a valid clock along each worldline ($d\lambda = d\tau$, monotonically increasing); what changes with the parameterization β is not the clock's behavior along any single worldline but the intra-leaf causal character for each equal-λ leaf across neighboring worldlines.

Hence, the general metric takes the form (4) with no restriction on g_{ij} beyond smoothness and Lorentzian signature of the full $g_{\mu\nu}$.

Type of small step	**Operational requirement**	**Metric condition (weak reading)**
Along one observer	$dy^i = 0 \Rightarrow ds^2 < 0$	$ds^2 = g_{00}(d\lambda)^2 = -(\partial\lambda/\partial\tau)^{-2}(d\lambda)^2$

[2] Permuting coordinate labels corresponds to a relabeling of metric indices, leaving the geometry and the logical structure of the adaptation unchanged.

• **Strong reading (the mollusk as a reference body).**

Einstein's own description goes further. He characterizes the mollusk by the "formal preservation of the separate existence of the *spatial* coordinates as opposed to the *time* coordinate" such that "every point of the mollusk is treated as a *spatial* point, and every material point at rest relative to it simply as at rest, so long as the mollusk is treated as a reference *body*". In this reading, the mollusk is not merely a congruence of worldlines but a material reference *body* whose *spatial* points have *spatial* coordinates. Consider the spacetime region $\mathcal{U}$ covered by the mollusk in motion over time. In particular, take two neighboring material points on the mollusk $\mathbf{y}$ and $\mathbf{y}+d\mathbf{y}$ at the same clock reading λ_0. This fixes two events, and we may ask how we can interpret these infinitesimally separated two events in terms of a spatial reference body. If a light signal emitted by material point $\mathbf{y}$ at its local time λ_0 can reach the neighboring material point $\mathbf{y}+d\mathbf{y}$ at that neighbor's time λ_0, then the second event lies in the causal future of the first. The separation is null and the equal-λ set would locally follow a light front instead of representing a material body. If the displacement between them were timelike, even a massive particle could travel from one to the other, so the equal-λ leaf would locally contain a piece of history rather than a body. Thus, for configurations to locally represent reference *bodies* at any fixed $\lambda=\lambda_0$, the infinitesimal separation of nearby spatial points at fixed λ should be spacelike.

• **Spacelike spatial grid (steps between neighboring observers):**

To conceptually "move" *through the mollusk* from one observer to a nearby observer at the *same clock reading*, we consider small coordinate "steps"[3] with

$$d\lambda = 0, \qquad (dy^1, dy^2, dy^3) \neq (0,0,0).$$

For such purely spatial "steps" (e.g., performed in an offline phase) the line element reduces to

$$ds^2 = g_{ij}\, dy^i\, dy^j. \tag{6}$$

Requiring these inter-observer "steps" to be spacelike means $ds^2 > 0$ for every nonzero infinitesimal displacement in any direction inside the body, i.e.

$$g_{ij}\, dy^i\, dy^j > 0 \quad \text{for all } (dy^1, dy^2, dy^3) \neq (0,0,0). \tag{7}$$

In other words, the 3×3 matrix g_{ij} is positive definite: it assigns positive squared-distance to every nonzero displacement at fixed λ. The following table summarizes the mollusk conditions using the line element $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$ evaluated on simple coordinate "steps".

Type of small step	**Operational requirement**	**Metric condition (mollusk)**
Along one observer	$dy^i = 0 \Rightarrow ds^2 < 0$	$ds^2 = g_{00}(d\lambda)^2 = -(\partial\lambda/\partial\tau)^{-2}(d\lambda)^2$
Between neighbors	$d\lambda = 0,\ (dy^i) \neq 0 \Rightarrow ds^2 > 0$	$ds^2 = g_{ij}dy^i dy^j > 0$ for all nonzero (dy^i)

The second line is not presented as a historical claim that Einstein implied the modern criterion in these terms; rather, it is our proposed interpretation of his reference body language.

[3] We remind the reader that dy^i and $d\lambda$ are infinitesimal displacements, "moves" or "steps", not physically moving observers. These conceptual "moves" are made on the gathered event data, e.g., in an offline processing phase.

Table 1 collects the results. Key observations:

- The **mollusk's** local 3+1 structure of coordinate basis directions results from interpreting Einstein's physical requirement that the reference frame be a spatial *body*.
- **Congruence-adapted reference frames** are in principle more flexible: there is necessarily one coordinate-line family tangent to the congruence's worldlines, hence timelike, while no additional spatial foliation assumption is yet imposed.
- The **programmable layer** imposes no intrinsic causal restriction on the displayed coordinate values anymore, no matter whether it starts from the weaker congruence-adapted reading or the stronger mollusk reading: coordinates are conventional numbers on screens. It can therefore display any valid programmable coordinates, including those with no timelike directions.

Construction framework	**Built-in gauge adaptation**	**Compatible coordinate-line patterns**	**Patterns excluded without extra processing**
Einstein's mollusk (reference body + clocks)	$g_{00} < 0$ and g_{ij} positive definite; structurally: timelike threading plus ADM-like spacelike leaves	$1\,\mathrm{T} + 3\,\mathrm{S}$ only	All other patterns: all $a\,\mathrm{T} + b\,\mathrm{N} + c\,\mathrm{S}$, differing from $a = 1, b = 0, c = 3$
Observer congruence (timelike congruence, leaves free)	$g_{00} < 0$ only; g_{ij} unrestricted	$a\,\mathrm{T} + b\,\mathrm{N} + c\,\mathrm{S}$ $a + b + c = 4$, with $a \neq 0$	zero-T patterns: $b\,\mathrm{N} + c\,\mathrm{S}$, $b + c = 4$
ADM decomposition (spacelike foliation, not a material frame)	g_{ij} positive definite, lapse function $N > 0$; g_{00} unrestricted	$1\,\mathrm{T} + 3\,\mathrm{S}$ $1\,\mathrm{N} + 3\,\mathrm{S}$ $4\,\mathrm{S}$	Fewer than 3 S: $a\,\mathrm{T} + b\,\mathrm{N} + c\,\mathrm{S}$, $a+b+c = 4, c < 3$
Programmable layer (seed + computation)	no built-in causal restriction on displayed values beyond validity of the coordinates	programmable coordinate-line causal patterns	none beyond programmability and chart validity

Table 1: Coordinate-line causal patterns for four frameworks. T, N, and S denote the local causal character (timelike, null, spacelike) of the four coordinate lines obtained by varying one coordinate while holding the other three fixed; $a, b, c \in \{0, 1, 2, 3, 4\}$ with $a + b + c = 4$. For the three material frameworks, the last two columns indicate which patterns can or cannot be directly displayed from the raw protocol. ADM is included for comparison as a foliation formalism rather than a material reference frame; see A.3 below. The programmable layer (last row) supplies the extra processing needed to reach patterns outside the seed frame's adapted class.

Three levels of spatial geometry in congruence-adapted coordinates

We conclude this introductory section by an optional note. Much of the literature on comoving coordinates and dust reference frames writes the comoving metric in forms that implicitly allow for a mollusk under the strong reading. To see exactly which assumptions enter where, it is useful to distinguish three levels.

1. The spatial projector (Lorentzian spacetime)

Given any timelike congruence with four-velocity u^μ (normalized so that $g_{\mu\nu}u^\mu u^\nu = -1$), the tensor

$$h_{\mu\nu} = g_{\mu\nu} + u_\mu u_\nu \tag{8}$$

projects orthogonally to u^μ. It measures the spatial geometry of the three-dimensional subspace perpendicular to the flow at each event. In comoving coordinates ($u^i = 0$), the normalization gives $u^0 = 1/\sqrt{-g_{00}}$, and the projector satisfies $h_{0\mu} = 0$ identically. Its purely spatial components are[4]

$$h^{\text{proj}}_{ij} = g_{ij} + \frac{g_{0i}\, g_{0j}}{-g_{00}}. \tag{9}$$

This object is *always* positive definite for any timelike congruence.[5]

2. Positive definiteness of the leaf metric (spatial body assumption)

The adopted interpretation of the mollusk as a reference *body* requires that the tangent directions within the equal-λ leaves be spacelike in the local sense relevant here (Section A.1, strong reading). This is the condition that the *induced metric on the leaf*—which in adapted coordinates is simply g_{ij}—be positive definite:

$$g_{ij}\, dy^i\, dy^j > 0 \quad \text{for all } d\mathbf{y} \neq 0 \quad \text{at } d\lambda = 0. \tag{10}$$

This is a genuinely stronger assumption (for a reference *body*), not an automatic consequence of the congruence being timelike. It is the content of the strong reading of Einstein's mollusk. When $g_{0i} \neq 0$, the projector h^{proj}_{ij} and the leaf metric g_{ij} are *different objects*. The projector measures geometry orthogonal to the flow; the leaf metric g_{ij} measures geometry within the equal-λ surface.

3. Orthogonal leaves (irrotationality)

If the congruence is irrotational—i.e. satisfies the Frobenius integrability condition $u_{[\mu}\nabla_\nu u_{\rho]} = 0$, equivalently vanishing vorticity $\omega_{\mu\nu} = 0$—then there exist hypersurfaces orthogonal to u^μ, and one can choose the equal-λ surfaces to coincide with them. In this case $g_{0i} = 0$, and the projector coincides with the leaf metric: $h^{\text{proj}}_{ij} = g_{ij}$. The metric takes the block-diagonal form

$$ds^2 = g_{00}\, d\lambda^2 + g_{ij}\, dy^i\, dy^j, \tag{11}$$

with $g_{00} = -(\partial\lambda/\partial\tau)^{-2} < 0$ and g_{ij} positive definite. For proper-time clocks ($\lambda = \tau$, $g_{\tau\tau} = -1$) this becomes the *synchronous* or *Gaussian normal* form (for an example, see Eq. (51)):

$$ds^2 = -d\tau^2 + g_{ij}\, dy^i\, dy^j. \tag{12}$$

[4] This coincides with the spatial metric of [2], § 84, for comoving observers.

[5] The reason is a basic fact of Lorentzian geometry: any nonzero vector orthogonal to a timelike vector is spacelike. The projector $h_{\mu\nu}$ extracts the component of any vector orthogonal to u^μ; for a spacetime vector $V^\mu = (0, v^i)$ with $v^i \neq 0$, this component is nonzero (since V^μ is not parallel to u^μ) and therefore spacelike: $h_{\mu\nu}V^\mu V^\nu > 0$.

A.2 Double-null Kruskal–Szekeres and seed-frame limits

We consider the Schwarzschild metric in standard coordinates t, r, θ, ϕ:

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dt^2 + \left(1 - \frac{2M}{r}\right)^{-1} dr^2 + r^2 d\Omega^2, \tag{13}$$

where $d\Omega^2 = d\theta^2 + \sin^2\theta\, d\phi^2$. Here and below, (θ, ϕ) are understood on a local angular patch of S^2; the poles and the ϕ-branch cut are excluded from the working region. This chart and its associated metric satisfy the mollusk conditions outside the horizon ($r > 2M$): a mollusk can be used as a physical reference frame in this situation. Note that this particular mollusk is not based on a free-falling congruence: static observers at fixed r in Schwarzschild spacetime are accelerated.

Moreover, inside the horizon ($r < 2M$), the same chart still satisfies the mollusk conditions, but with reassigned roles: $g_{rr} < 0$ so that dr serves as the timelike coordinate basis direction, while dt becomes spacelike. Hence, a mollusk is possible in principle (see also Appendix B.4) but we require a different congruence of observers—at fixed (t, θ, ϕ) advancing in r—since the exterior static congruence (observers at fixed r) cannot be continued past the horizon (timelike worldlines at constant r cease to exist there).

Now consider the exterior Schwarzschild geometry expressed in Kruskal–Szekeres coordinates in their double-null form (we follow the coordinate convention of Poisson, Eqs. (5.7)–(5.9), [3]). These are defined via null coordinates

$$U = -e^{-u/4M}, \quad V = e^{v/4M}, \tag{14}$$

where $u = t - r^*$ and $v = t + r^*$ are the usual retarded and advanced times, and r^* is the tortoise coordinate,

$$r^* = r + 2M \ln\left|\frac{r}{2M} - 1\right|. \tag{15}$$

In these coordinates, the Schwarzschild metric (line–element) takes the form:

$$ds^2 = -\frac{32M^3}{r} e^{-r/2M} dU dV + r^2 d\Omega^2, \tag{16}$$

where r is implicitly a function of UV determined by the relation

$$e^{r/2M}\left(\frac{r}{2M} - 1\right) = -UV. \tag{17}$$

The key operational feature of the double-null chart is easy to read directly from the line element. If we take a small "step" with $dV = d\theta = d\phi = 0$ (i.e. we change only U), then $ds^2 = 0$, so the coordinate line obtained by varying only U is *null*. Likewise, if we take $dU = d\theta = d\phi = 0$ and change only V, we again find $ds^2 = 0$. This is equivalent to the statement that $g_{UU} = g_{VV} = 0$ with only a cross-term $g_{UV} \neq 0$. In other words: mollusk and congruence-adapted seed frames require one coordinate-line family tangent to the observer worldlines to be timelike (see Table 1), but the coordinate curves when changing only U (or only V) are *null*, not timelike. So these coordinate lines cannot coincide with the worldlines of massive, clock-carrying observers. On the chosen angular patch, the θ- and ϕ-coordinate lines are spacelike as well. Hence double-null KS coordinates are not available as the raw comoving coordinates of either a mollusk (3+1) or the weaker one-clock-plus-fixed-label congruence-adapted construction. They can nevertheless be displayed after a further programmable transformation from a valid timelike seed frame, see Appendix B.5.

A.3 ADM charts with timelike grid lines can be instantiated by a single mollusk

The Arnowitt–Deser–Misner (ADM) formalism expresses the spacetime metric using a time coordinate t, a lapse function N, a shift vector N^i, and a spatial metric h_{ij}. The metric takes the form:

$$ds^2 = -N^2 dt^2 + h_{ij}(dx^i + N^i dt)(dx^j + N^j dt). \tag{18}$$

To determine at the local level whether an ADM chart can be physically instantiated by a mollusk, we apply the same *small-step test with the line element* used above.[6]

1. **Timelike "grid-observer" time steps:** A mollusk ties its observers to fixed spatial labels. In an ADM chart that means we consider the coordinate lines with $x^i = \text{const}$, i.e. small steps with $dx^i = 0$, and only t changes. Plugging this into the ADM line element gives

$$ds^2 = g_{tt}\, dt^2 = \left(-N^2 + h_{ij}N^i N^j\right) dt^2. \tag{19}$$

   For these $x^i = \text{const}$ lines to represent massive observers carrying clocks, we must have $ds^2 < 0$ for such "steps". Therefore the mollusk condition in an ADM chart is

$$g_{tt} = -N^2 + h_{ij}N^i N^j < 0 \tag{20}$$

   throughout the region of interest.

2. **Spacelike "steps" within a $t = \text{const}$ slice:** To conceptually "move" (during the offline evaluation) between neighboring grid points at the same coordinate time, we set

$$dt = 0, \qquad (dx^1, dx^2, dx^3) \neq (0, 0, 0).$$

   The line element becomes

$$ds^2 = h_{ij}\, dx^i dx^j, \tag{21}$$

   which is positive for all nonzero spatial steps because h_{ij} is assumed Riemannian (positive definite) in ADM.

3. **Non-crossing / well-defined labeling:** Finally, the coordinate lines $x^i = \text{const}$ should form a smooth, non-intersecting family of curves.

The second and third conditions are largely built into the ADM setup: h_{ij} is assumed Riemannian, and the coordinate grid (t, x^i) defines smooth families of curves locally, and assuming the chart is well-defined on the region of interest. The central condition is thus the first: $g_{tt} < 0$. The lapse N controls how much proper time separates neighboring $t = \text{const}$ slices, while the shift N^i describes how the spatial coordinate grid slides sideways from one slice to the next. If the sideways slide is too large compared to the lapse, then the coordinate lines $x^i = \text{const}$ become lightlike or spacelike rather than timelike. The inequality condition $g_{tt} < 0$ is exactly the statement that fixed-x^i lines can serve as worldlines of massive observers. Hence, in practice, one often *chooses* lapse and shift so that $g_{tt} < 0$ on the region of physical interest when one wants the coordinate grid points $x^i = \text{const}$ to behave like sensible "stand-in observers". But even when $g_{tt} \geq 0$, the ADM decomposition remains mathematically valid—only the physical mollusk interpretation fails, see Appendix B.4.

[6] Note that the h_{ij} appearing in the ADM line element is the induced metric on the $t = \text{const}$ leaves, i.e. g_{ij} in adapted coordinates. It coincides with the flow projector h_{ij}^{proj} of Section A.1 only when the shift vanishes, $N^i = 0$.

Appendix B: Implementations of programmable layers

To illustrate the operational reference frame model developed in the main text, we now consider examples of observer swarms with programmable, coordinate-outputting devices with displays. We begin with two examples in SR (B.1–B.2), then give a brief GR framing discussion (B.3), followed by two GR examples (B.4–B.5).

B.1 Linear coordinate transformations in SR

Even in flat Minkowski spacetime with the standard inertial coordinates $X^\mu = (X^0, X^1, X^2, X^3)$ and line element

$$ds^2 = -(dX^0)^2 + (dX^1)^2 + (dX^2)^2 + (dX^3)^2,$$

a simple linear coordinate transformation can change the causal character of the coordinate basis directions.

A one-parameter family of linear charts. For any $\alpha > 0$, define new coordinates

$$\tilde{X}^0 = \frac{\alpha\, X^0 + \alpha^{-1} X^1}{\sqrt{2}}, \qquad \tilde{X}^1 = \frac{\alpha\, X^0 - \alpha^{-1} X^1}{\sqrt{2}}, \qquad \tilde{X}^2 = X^2, \qquad \tilde{X}^3 = X^3. \tag{22}$$

The Jacobian matrix of the transformation $X^\mu \mapsto \tilde{X}^\mu$ is

$$\frac{\partial \tilde{X}^\mu}{\partial X^\nu} = \begin{pmatrix} \alpha/\sqrt{2} & \alpha^{-1}/\sqrt{2} & 0 & 0 \\ \alpha/\sqrt{2} & -\alpha^{-1}/\sqrt{2} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \qquad \det = -1, \tag{23}$$

which is nonzero for all $\alpha > 0$, confirming that (22) is a valid coordinate transformation.

The Minkowski metric in the transformed coordinates. Inverting (22) gives

$$dX^0 = \frac{d\tilde{X}^0 + d\tilde{X}^1}{\sqrt{2}\,\alpha}, \qquad dX^1 = \frac{\alpha\,(d\tilde{X}^0 - d\tilde{X}^1)}{\sqrt{2}}. \tag{24}$$

Substituting into the Minkowski line element yields

$$ds^2 = \tilde{g}_{\mu\nu} d\tilde{X}^\mu d\tilde{X}^\nu = \frac{\alpha^4 - 1}{2\alpha^2}\left[(d\tilde{X}^0)^2 + (d\tilde{X}^1)^2\right] - \frac{\alpha^4 + 1}{\alpha^2}\, d\tilde{X}^0\, d\tilde{X}^1 + (d\tilde{X}^2)^2 + (d\tilde{X}^3)^2, \tag{25}$$

with metric tensor

$$\tilde{g}_{\mu\nu} = \begin{pmatrix} \frac{\alpha^4 - 1}{2\alpha^2} & -\frac{\alpha^4 + 1}{2\alpha^2} & 0 & 0 \\ -\frac{\alpha^4 + 1}{2\alpha^2} & \frac{\alpha^4 - 1}{2\alpha^2} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{26}$$

At $\alpha = 1$ this reduces to

$$ds^2 = -2\, d\tilde{X}^0\, d\tilde{X}^1 + (d\tilde{X}^2)^2 + (d\tilde{X}^3)^2, \tag{27}$$

the standard light-cone line element, with $\tilde{X}^0 = (X^0 + X^1)/\sqrt{2}$ and $\tilde{X}^1 = (X^0 - X^1)/\sqrt{2}$.

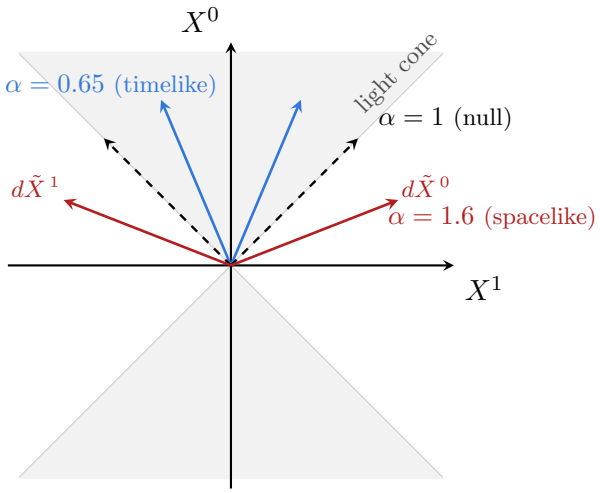

Figure B.1: Tangent directions to coordinate curves in the (X^0, X^1) plane for the one-parameter family (22). Each pair of arrows shows the directions obtained by varying only $\tilde{X}^0$ (right-leaning) or only $\tilde{X}^1$ (left-leaning) while holding the other transformed coordinates fixed. The two directions are mirror images about the X^0 axis. The shaded region is the interior of the light cone. Observer worldlines of the seed swarm (not shown) run vertically (constant X^1, varying X^0). Steep pair (blue): $\alpha < 1$, both directions timelike. Dashed 45° pair (black): $\alpha = 1$, both null. Shallow pair (red): $\alpha > 1$, both spacelike.

A linear transform that is *not* Lorentz. The determinant of the Jacobian matrix is -1, so the transformation is invertible for every $\alpha > 0$. The negative sign means that the linear map reverses orientation in the (X^0, X^1) plane. Chart validity requires only $\det \neq 0$; orientation-reversing charts ($\det < 0$) are perfectly admissible coordinate descriptions of the same spacetime region. Without the $1/\sqrt{2}$ factors, the determinant would be -2, as in Dirac's original unnormalized definition of light-cone coordinates [4]. Flipping the sign of one null coordinate would change the determinant to $+1$, corresponding to an orientation-preserving convention. This would *not* make the transformation Lorentz, however. In the present linear setting, a Lorentz transformation is a linear map Λ satisfying $\Lambda^{\mathrm{T}} \eta \Lambda = \eta$, so that the Minkowski metric retains the same component matrix $\eta_{\mu\nu} = \mathrm{diag}(-1, 1, 1, 1)$ in the transformed coordinates. By contrast, the transformation (22) is invertible but does *not* satisfy that condition for generic α; it rewrites the same Minkowski spacetime in coordinates with different, generally non-diagonal components. This makes the distinction between freely programmable valid coordinates and actual physical content manifest already in SR. A programmable layer can display such transformed coordinates: the observers are equipped with devices that directly display the transformed values.

Small-step causal-character test. Following the same method used for the mollusk analysis above, we examine the causal character of each coordinate basis direction in the $(\tilde{X}^0, \tilde{X}^1)$ plane.

- *Vary only* $\tilde{X}^0$: set $d\tilde{X}^1 = d\tilde{X}^2 = d\tilde{X}^3 = 0$. Then $ds^2 = \tilde{g}_{00}\,(d\tilde{X}^0)^2$.
- *Vary only* $\tilde{X}^1$: set $d\tilde{X}^0 = d\tilde{X}^2 = d\tilde{X}^3 = 0$. Then $ds^2 = \tilde{g}_{11}\,(d\tilde{X}^1)^2$.

Since $\tilde{g}_{00} = \tilde{g}_{11} = (\alpha^4 - 1)/(2\alpha^2)$, the sign of $\alpha^4 - 1$ determines three regimes:

Parameter	$\tilde{g}_{00} = \tilde{g}_{11}$	**Causal character when varying only $\tilde{X}^0$ or $\tilde{X}^1$**
$\alpha > 1$	> 0	both spacelike
$\alpha = 1$	$= 0$	both null
$\alpha < 1$	< 0	both timelike

In every case, $d\tilde{X}^2$ and $d\tilde{X}^3$ remain spacelike. A mollusk requires exactly one timelike and three spacelike coordinate basis directions. A congruence-adapted seed frame requires at least one timelike coordinate-line family (see Table 1). These conditions fail for $\alpha \geq 1$ for all transformed charts in the family (22), and, for the mollusk, also for $\alpha < 1$, even though the underlying spacetime remains flat Minkowski spacetime throughout. Figure B.1 illustrates the three regimes. Because $\tilde{g}_{00} = \tilde{g}_{11}$ for all $\alpha > 0$, the two directions in the $(\tilde{X}^0, \tilde{X}^1)$ plane always share the same causal character, so no member of the family provides the one-timelike-plus-three-spacelike structure required by a mollusk.

Congruence-adapted seed frame for $\alpha < 1$. The $\tilde{X}^0$-coordinate lines—curves of constant $\tilde{X}^1, \tilde{X}^2, \tilde{X}^3$—are timelike. They could therefore serve as worldlines of an observer congruence, with $\tilde{X}^0$ as the clock parameter and $(\tilde{X}^1, \tilde{X}^2, \tilde{X}^3)$ as comoving labels. Which congruence? The inverse transformation gives

$$X^0 = \frac{\tilde{X}^0 + \tilde{X}^1}{\sqrt{2}\,\alpha}, \qquad X^1 = \frac{\alpha\,(\tilde{X}^0 - \tilde{X}^1)}{\sqrt{2}}, \tag{28}$$

so along a $\tilde{X}^0$-line (fixed $\tilde{X}^1$): $dX^1/dX^0 = \alpha^2$. These are inertial observers uniformly boosted to velocity $v = \alpha^2 < 1$ in the X^1-direction.

A congruence-adapted reference frame based on this boosted congruence, with comoving coordinates $(\tilde{X}^0, \tilde{X}^1, \tilde{X}^2, \tilde{X}^3)$, would directly produce these coordinate values from its clocks and labels. However, since $\tilde{g}_{11} < 0$ (the $\tilde{X}^1$ basis direction is timelike), the equal-$\tilde{X}^0$ leaves are not spacelike: two neighboring "spatial" labels at the same clock reading can be connected by a massive particle. This is not a body at an instant; it is not a mollusk. It is an example of a congruence-adapted comoving frame that is more general than a mollusk. It is, however, still otherwise adapted: each target coordinate system requires a matching congruence whose worldlines coincide with one of the target's timelike coordinate lines.

In the above example, the underlying congruence is the *same* set of worldlines that a standard Lorentz boost to velocity $v = \alpha^2$ would single out. We may ask again, why, then, is the metric (26) not simply $\eta_{\mu\nu}$ in this case? Because the α-transformation and a Lorentz boost, although they can be adapted to the same timelike congruence, assign *different* time labels and *different* equal-time hypersurfaces to that congruence. A Lorentz boost is chosen so that the boosted observers' equal-time hypersurfaces are orthogonal with respect to $\eta_{\mu\nu}$ to their worldlines (i.e., tilted symmetrically about the light cone in Fig. B.1), and in those coordinates the metric retains the standard Minkowski form $\eta_{\mu\nu}$. The α-transformation, by contrast, tilts the equal-$\tilde{X}^0$ hypersurfaces away from orthogonality and therefore produces off-diagonal terms such as $\tilde{g}_{01} \neq 0$.

Two timelike directions still span the full tangent space. For $\alpha < 1$, both $d\tilde{X}^0$ and $d\tilde{X}^1$ are timelike. One might wonder whether such a coordinate basis can reach spacelike or lightlike directions at all. It can: the coordinate basis vectors in the (X^0, X^1) plane are

$$e_{\tilde{0}} = \frac{1}{\sqrt{2}}(\alpha^{-1},\ \alpha), \qquad e_{\tilde{1}} = \frac{1}{\sqrt{2}}(\alpha^{-1},\ -\alpha),$$

which are linearly independent (they differ in the sign of the spatial component). The purely spatial direction $(0, 1)$ is reached by the combination $e_{\tilde{0}} - e_{\tilde{1}}$, i.e. a step with $d\tilde{X}^0 = +1$, $d\tilde{X}^1 = -1$. The purely temporal direction $(1, 0)$ is reached by $e_{\tilde{0}} + e_{\tilde{1}}$, i.e. $d\tilde{X}^0 = +1$, $d\tilde{X}^1 = +1$. For the *mathematical* validity of a coordinate chart, the causal character of the individual coordinate basis directions is irrelevant; what matters is that the four coordinate basis vectors remain locally linearly independent, equivalently that the Jacobian be nonsingular. Causal character is central, however, for *direct material instantiation* by a mollusk or congruence-adapted seed frame: one coordinate-line family must coincide with massive observer worldlines.

Lessons. This entirely special-relativistic example demonstrates multiple points:

- The causal-character restriction is a property of the *display protocol*.
- The three rows of Table 2 map onto levels of a hierarchy: in our example, the $\alpha < 1$ regime is the unique window accessible to a congruence-adapted seed frame but not to any mollusk; the $\alpha \geq 1$ regimes are inaccessible to the congruence-adapted constructions devoid of a programmable layer.
- A single fixed seed swarm—the standard inertial observers of Fig. 1 of the main text—equipped with a programmable layer can display all three regimes by varying a single parameter α in the same programmed transformation.

α	Directions	Mollusk?	Congruence-adapted frame?	Programmable layer?
$\alpha < 1$	2 timelike + 2 spacelike	no (equal-time leaf not spacelike)	yes (but boosted congruence)	yes
$\alpha = 1$	2 null + 2 spacelike	no	no	yes
$\alpha > 1$	4 spacelike	no	no	yes

Table 2: The α-family of coordinates. Each row is excluded at a different level. "Congruence-adapted frame" means: realizable as the adapted (comoving) coordinates of some timelike congruence, without requiring the equal-clock leaves to be spacelike. The programmable layer operates on a fixed seed swarm (e.g., the standard inertial observers) and displays all three regimes by programming alone.

B.2 Swapping the seed: rotating frames showing inertial coordinates in SR

The previous examples all started from a "standard" seed frame and programmed more "exotic" coordinates. This section illustrates that the seed and target roles can be interchanged: the same two coordinate systems can be related by programming in either direction, depending on which frame is physically realized as the seed. The example uses only special relativity.[7]

Consider two coordinate systems on a local cylindrical chart patch of Minkowski spacetime with $0 < R_{\min} \le R \le R_0 < 1/\omega$, $T \in [0, T_{\rm obs}]$, and $Z \in [Z_{\min}, Z_{\max}]$, where Φ and ϕ' are understood as local angular coordinates on a chosen branch of width exceeding the angular sweep $\omega\, T_{\rm obs}$, centered so that both Φ and $\phi' = \Phi - \omega\, T$ remain within the branch throughout the deployment.[8]

- *Inertial coordinates* (T, R, Φ, Z): standard cylindrical Minkowski coordinates with line element
$$ds^2 = -dT^2 + dR^2 + R^2\, d\Phi^2 + dZ^2. \tag{29}$$

- *Corotating coordinates* (t, r, ϕ', z): adapted to observers distributed throughout the chosen angular patch, corotating rigidly[9] with angular velocity ω, defined by
$$t = T, \quad r = R, \quad \phi' = \Phi - \omega\, T, \quad z = Z. \tag{30}$$
On any chosen cylindrical chart patch away from the rotation axis, this transformation has Jacobian determinant 1 and is a smooth local diffeomorphism there. For the *rotating-seed* interpretation (used in Direction 2 below) we additionally restrict to $R_{\min} \le r \le R_0 < 1/\omega$ (all material points lying inside the light cylinder) and assume Born rigidity[10]. The line element becomes
$$ds^2 = -(1 - \omega^2 r^2)\, dt^2 + 2\,\omega\, r^2\, d\phi'\, dt + dr^2 + r^2\, d\phi'^2 + dz^2. \tag{31}$$

[7] The spacetime is flat throughout (backreaction ignored; see Assumptions section). The off-diagonal terms and position-dependent coefficients in the corotating metric are pure coordinate artifacts.

[8] Thus we work away from the rotation axis. A branch of angular width $\Delta\Phi$ must accommodate the angular sweep $\omega\, T_{\rm obs}$ with margin for the initial positions of both Φ and ϕ'; for longer observation times one may instead use the Cartesian reformulation noted below, which avoids angular branch bookkeeping altogether. Moreover, the z-direction decouples entirely from the rotation ($g_{zz} = 1$, no cross-terms), so every $z = $ const slice is equivalent. We retain z only so that all four coordinates are displayed. The restriction $0 < R_{\min} \le R \le R_0 < 1/\omega$ specifies the common region in which both programming directions are physically realizable. In Direction 1 we will also note, as an extension, that the inertial seed can display the same corotating coordinates on a larger region, including $R > 1/\omega$, even though the coordinate curves with $dt \neq 0$ and $dr = d\phi' = dz = 0$ are then no longer timelike.

[9] The rotating disk played a pivotal role in Einstein's path to general relativity: in [5], § 3, he used it to argue that coordinates on the disk lose their direct metrical significance, motivating the need for general coordinates and non-Euclidean geometry.

[10] For the kinematic framework used in the example, Born rigidity is actually not required: a smooth, non-crossing timelike rotating congruence with fixed labels (r, ϕ', z) already suffices as a seed frame on the chosen region. We impose Born rigidity only if the rotating seed is to be interpreted literally as a rigid material disk. In that stronger reading, neighboring proper distances measured orthogonally to the worldlines remain constant. This is the standard relativistic notion of rigidity. It does not imply Euclidean spatial geometry: for a rotating disk, the spatial geometry measured by the corotating observers is non-Euclidean, as in the Ehrenfest discussion.

Direction 1: inertial seed → corotating display. The seed swarm consists of standard inertial observers at rest in (R, Φ, Z), each carrying a programmable device. The program evaluates Eq. (30) at each observer's proper time $\tau_M = T$:

$$X^t = T, \quad X^r = R, \quad X^{\phi'} = \Phi - \omega T, \quad X^z = Z.$$

Every device displays corotating coordinates on the fly. Since R, Φ, and Z are constant along each inertial worldline, only X^t and $X^{\phi'}$ vary with τ_M—both updated locally by each observer.

The displayed corotating coordinates have $g_{tt} = -(1-\omega^2 r^2) < 0$ (within the light cylinder) and g_{ij} positive definite: all conditions for a mollusk are satisfied. The off-diagonal term $g_{t\phi'} = \omega r^2 \neq 0$ indicates that the equal-t surfaces, while spacelike, are not orthogonal to the corotating worldlines.

Observers actually sitting still in the inertial frame could in principle display corotating coordinates also beyond the light cylinder ($R > 1/\omega$). On any such local cylindrical patch the chart remains mathematically valid, only g_{tt} changes sign ($g_{tt} > 0$), so these displayed coordinates cease to be mollusk-type and instead have four spacelike coordinate basis directions.

Direction 2: rotating seed → inertial display. Now reverse the roles. The seed swarm consists of the material points of a physically rotating configuration: observers with fixed (r, ϕ', z) following helical worldlines in the inertial frame, each carrying a programmable device. Along each helical worldline, r, ϕ', and z are constant by construction. The physical clock carried by each disk observer rotating with angular velocity ω at radius r around the z-axis measures proper time τ_D, not the coordinate time t. For an observer at fixed $r < 1/\omega$, Eq. (31) gives

$$d\tau_D = \sqrt{1-\omega^2 r^2}\, dt,$$

since $ds^2 = -d\tau_D^2$ along the worldline. After fixing clock origins (e.g., $t_0 = \tau_0 = 0$), the device computes

$$t = \frac{\tau_D}{\sqrt{1-\omega^2 r^2}}. \tag{32}$$

Using this computed t together with the fixed labels (r, ϕ', z), the device evaluates

$$X^T = t, \qquad X^R = r, \qquad X^\Phi = \phi' + \omega t, \qquad X^Z = z.$$

Again, no inter-observer communication is required. The metric expressed in the displayed inertial coordinates (X^T, X^R, X^Φ, X^Z) is the standard Minkowski metric of Eq. (29). Yet the observers producing these coordinates travel on helical worldlines—they are *not* at rest in the coordinate system they display. This is precisely the decoupling between observer motion and displayed coordinate structure that the programmable layer makes possible.

Lessons: Note the asymmetry:

- In Direction 1 the inertial seed is not limited by the light cylinder even though the observers move on timelike worldlines (by sitting still in the inertial seed frame). On any chosen cylindrical patch away from the axis, the same programmed transformation can be used also for regions with $R > 1/\omega$; no superluminal signaling is required, only local evaluation of the programmed formulas.

- In Direction 2 the massive observers must move on timelike worldlines, but now they rotate around the z-axis. Hence, the rotating seed frame itself cannot physically exist at or beyond $R = 1/\omega$: equivalently, the proper acceleration of each corotating observer,

$$a = \sqrt{a_\mu a^\mu} = \frac{\omega^2 r}{1 - \omega^2 r^2},$$

diverges as $r \to 1/\omega$, so no finite force can maintain the congruence at the light cylinder. Hence the displayed coordinate *values* can still be decoupled from the physical devices used to construct them (and coordinates for standard inertial Minkowski space can be displayed by accelerated observers), but the spacetime region in which a physically *rotating* programmable layer can operate remains limited by the seed frame's physical constraints.

This also illustrates the distinction between the mathematical transformation and the physical devices that are programmed to evaluate it. The coordinate transformation $\phi' = \Phi - \omega T$ is well-defined on any chosen cylindrical chart patch away from the rotation axis (and in the equivalent Cartesian form discussed below it is even global on Minkowski spacetime). The devices co-moving with a rotating seed, by contrast, are physical objects attached to physical observers: they share the observers' worldlines, are subject to accelerations, and cannot exist where the rotating seed congruence breaks down. The programmable layer inherits the seed frame's physical constraints precisely because it consists of material devices—flexible ones with computational capabilities, but material nonetheless. Its flexibility lies in what it computes, not in escaping the physical conditions under which it operates.

Coordinate conventions and long observation times. The cylindrical presentation above is the standard disk-adapted one: fixed (r, ϕ', z) label the rotating material points directly, and the Born metric follows immediately. The only extra bookkeeping issue is the angular variable.

1. *Local angular branch.* In the main discussion above we work on a chosen angular branch and a finite observation time $T_{\rm obs}$, with branch width large enough that both Φ and $\phi' = \Phi - \omega T$ remain within the branch throughout the deployment.
2. *Long runs or multiple revolutions.* If one wants to follow the same observers for arbitrarily many revolutions, one may instead reformulate the same example in Cartesian coordinates. Introduce inertial Cartesian coordinates (X^T, X^X, X^Y, X^Z) and corotating Cartesian coordinates (t, x', y', z) related by

$$X^T = t, \qquad X^X = x' \cos\omega t - y' \sin\omega t, \qquad X^Y = x' \sin\omega t + y' \cos\omega t, \qquad X^Z = z,$$

whose full Jacobian determinant is 1. The corotating metric then takes the Cartesian Born form

$$ds^2 = -(1 - \omega^2\rho^2)\, dt^2 + 2\omega(-y'\, dx' + x'\, dy')\, dt + dx'^2 + dy'^2 + dz^2,$$

where $\rho^2 = x'^2 + y'^2$. In this Cartesian reformulation no angular branch cut arises, so multiple revolutions require no additional angular bookkeeping.

These alternatives do not change the physical events, the observer congruence, or the measurement protocol. They change only the software convention by which the devices present the coordinate labels. This again illustrates the programmable layer's central point: coordinate conventionality made operationally tangible as a programming choice.

B.3 GR Scenarios: with and without metric hypothesis

We now consider general-relativistic examples. For explanatory purposes, we distinguish two scenarios, depending on whether a metric hypothesis serves as the starting point. In both scenarios, coordinates are "the work of man" (see [6], pg. 409), but Scenario 1 remains the chosen pedagogical route in the main text because it operates directly at the coordinate level and preserves the emphasized logical structure: first establish charts, then determine the metric. Scenario 2 becomes more natural once the metric has been introduced and a physical congruence can also be described theoretically.

Scenario 1: Creating coordinates without a metric hypothesis

- The spacetime metric is unknown.
- We freely choose a program to display valid coordinate values $X^\mu(\tau, \mathbf{y})$ across the congruence for each observer $\mathbf{y}$. For example, we may begin from a mollusk-type seed reference frame and then apply any valid coordinate transformation.
- We measure the metric in the chosen coordinates.
- This is the route emphasized in the main text: it follows the logical progression of differential geometry by establishing the coordinate description first and determining the metric afterward.

Note that in this scenario we are genuinely agnostic as to which metric tensor we may find. The operational procedure has merely introduced valid coordinates X^μ on the measured spacetime region, and no assumption about the metric tensor has been made. This illuminates a central aspect of GR: established coordinates together with measured metric components suffice to provide a valid description.

Incidentally, an *a posteriori* mismatch between measured components and a particular canonical form of a metric tensor would not by itself rule out any geometry: the same geometry can be expressed in different coordinates $\tilde{X}^\mu = \tilde{X}^\mu(X)$ with different component values of the metric. If one wants to test subsequently whether the measured metric is equivalent to a known spacetime geometry one therefore has to solve the equivalence problem, for example using the Cartan–Karlhede algorithm.

Scenario 2: Testing a specific coordinate representation of a hypothesized metric

- We hypothesize both a geometry (e.g., Schwarzschild) and its representation in specific coordinates x^μ. The same physical spacetime admits many coordinate representations; here we select one and test the combined package of coordinates and line element.
- We program the displays to show $X^\mu := x^\mu$. One convenient way to do this is to choose a suitable timelike congruence suggested by the hypothesis—in our example below, a geodesic congruence—together with suitable initial data, and then integrate along the observer worldlines to obtain $X^\mu(\tau, \mathbf{y})$.
- Upon measurement, we test whether the empirically obtained metric components take the predicted theoretical values in these coordinates, which were constructed using expected values for those metric components.

- If confirmed, we recover the hypothesized line element (e.g., Eq. (33)).
- If not confirmed, the displayed coordinates X^μ still correspond to valid chart patches in the regions where the reference frame could operate successfully and label all encountered events; only the measured metric components differ from the hypothesized ones. In that case, the combination of operationally established coordinates and actually measured metric still correctly describes the physical spacetime in the actual swarm deployment region. Chart freedom then permits subsequent offline coordinate transformations to coordinates better adapted to the measured metric, if such coordinates can be found.

A useful lesson emerges: a signal of genuine physical disagreement arises only when invariant relationships differ, whereas coordinate-dependent discrepancies merely reflect chart choices.

While idealized for pedagogical purposes, Scenario 2 also provides a toy model for actual GR research practice, introducing students to the methodology of starting with a combined set of theoretical assumptions and connecting theoretical predictions with experimental observations. Combined hypotheses encountered in practical situations would, however, typically also employ different coordinatizations; for example by extracting distance measures from observed signals received from remote events, see e.g., [7], Ch. 14. Moreover, even in our toy example, physical geometries would usually be parameterized (e.g., by mass, charge, or angular momentum), and failure to measure a particular line element would only directly exclude a single set of parameter values; Scenario 2 would therefore still require solving the equivalence problem if specific line-elements for geometric hypotheses were to be tested across a whole range of parameters.

In the following examples in Appendices B.4 and B.5 we remain within Scenario 2 and assume that a metric hypothesis underlies the programming to be applied to the devices carried by the observer swarm.

B.4 Interior Painlevé–Gullstrand Coordinates: four spacelike directions

We now work in Scenario 2, assume a Schwarzschild spacetime, and describe it in ingoing Painlevé–Gullstrand (PG) coordinates $x^\mu = (T, r, \theta, \phi)$, with line element

$$ds^2 = -\left(1 - \frac{2M}{r}\right) dT^2 + 2\sqrt{\frac{2M}{r}}\, dT\, dr + dr^2 + r^2 d\Omega^2, \qquad d\Omega^2 = d\theta^2 + \sin^2\theta\, d\phi^2, \tag{33}$$

where dT denotes the coordinate differential associated with the coordinate T, equivalently dx^T in index notation; no distinction between these notational conventions is intended here.

The structure of this appendix is as follows. First, we show that inside the horizon all four coordinate curves of the PG chart are spacelike on a local angular patch, so the chart cannot be directly realized by a mollusk-type seed frame. Second, we identify a timelike infalling congruence whose 4-velocity is a linear combination of those spacelike coordinate basis vectors. Third, we use that congruence together with suitable initial data to program the swarm so that it displays the PG coordinates. Finally, we check that the resulting displayed coordinates are valid on the chosen observation region but are no longer congruence-adapted.

B.4.1 The small-step (coordinate-line) causality test in the interior. The observer swarm works in both the exterior and interior of Schwarzschild spacetime. We first focus on the interior, where the Painlevé–Gullstrand chart has all four coordinate basis directions spacelike on a local angular patch away from the poles. Thus the chart admits no direct mollusk realization.

Fix an event with $0 < r < 2M$ (we may imagine all phases of deployment to happen inside a very large horizon). Consider the four coordinate curves through that event, i.e. vary one coordinate while holding the other three fixed.

- Vary only T: set $dr = d\theta = d\phi = 0$. Then

$$ds^2 = g_{TT}\, dT^2 = \left(\frac{2M}{r} - 1\right) dT^2 > 0, \tag{34}$$

  so the T-coordinate curve is *spacelike* inside the horizon.

- Vary only r: set $dT = d\theta = d\phi = 0$. Then

$$ds^2 = dr^2 > 0, \tag{35}$$

  so the r-coordinate curve is spacelike.

- Vary only θ or only ϕ (with the other differentials zero). Then

$$ds^2 = r^2 d\theta^2 > 0, \qquad ds^2 = r^2 \sin^2\theta\, d\phi^2 > 0, \tag{36}$$

  with the ϕ-statement understood on the chosen angular patch away from the poles. Thus the angular coordinate curves are spacelike.

Thus, in the Schwarzschild interior the PG chart has the striking property: *on any local angular patch away from the poles, all four coordinate lines obtained by holding the other three coordinates fixed are spacelike.*[11]

[11] The coordinate-line test probes g_{TT}, the covariant diagonal component, which determines the causal character of the basis vector ∂_T. By contrast, the contravariant component g^{TT} determines the causal character of the normal to the $T = \text{const}$ hypersurface. In PG coordinates $g^{TT} = -1$ everywhere, so the $T = \text{const}$ hypersurfaces are spacelike; but $g_{TT} > 0$ for $r < 2M$, so the basis vector ∂_T is spacelike there. These are different geometric statements. For a tabulation of both sets of components in various Schwarzschild slicings, see [8], § 6, Tables 3 and 4.

This does *not* mean there are no timelike directions: the metric is still Lorentzian. It means only that no coordinate basis direction is timelike in this chart:

Direction	**Metric component**	**Interior** $(0 < r < 2M)$
vary T only	$g_{TT} = -(1 - 2M/r)$	> 0 (spacelike)
vary r only	$g_{rr} = 1$	> 0 (spacelike)
vary θ only	$g_{\theta\theta} = r^2$	> 0 (spacelike)
vary ϕ only	$g_{\phi\phi} = r^2 \sin^2\theta$	> 0 (spacelike)

B.4.2 Timelike motion exists as a linear combination of spacelike coordinate basis vectors. A concrete timelike direction is given by the infalling "rain" 4-velocity field

$$u^\mu = u^\mu_{\mathrm{PG}} = \left(1, -\sqrt{\frac{2M}{r}}, 0, 0\right). \tag{37}$$

Even though the basis directions are spacelike inside, the 4-vector (37) is timelike:

$$\begin{aligned} g_{\mu\nu}u^\mu u^\nu &= g_{TT}(u^T)^2 + 2g_{Tr}u^T u^r + g_{rr}(u^r)^2 \\ &= \left(\frac{2M}{r} - 1\right) + 2\sqrt{\frac{2M}{r}}\,(1)\left(-\sqrt{\frac{2M}{r}}\right) + (1)\left(\frac{2M}{r}\right) = -1. \end{aligned} \tag{38}$$

This is also the standard geodesic "rain" congruence of ingoing PG coordinates. For a quick check (sidestepping the explicit evaluation of Christoffel symbols in PG coordinates), lower the index using Eq. (33) and obtain $u_\mu = g_{\mu\nu}u^\nu = (-1, 0, 0, 0) = -\partial_\mu T$. Hence $\partial_{[\mu}u_{\nu]} = 0$, and because the connection is torsion-free this implies $\nabla_{[\mu}u_{\nu]} = 0$, i.e. $\nabla_\mu u_\nu = \nabla_\nu u_\mu$. Using metric compatibility $\nabla_\rho g_{\mu\nu} = 0$ together with Eq. (38), we then obtain the four-acceleration

$$a_\mu := u^\nu \nabla_\nu u_\mu = u^\nu \nabla_\mu u_\nu = \tfrac{1}{2}\nabla_\mu(u^\nu u_\nu) = 0,$$

equivalently $a^\mu = 0$. The 4-velocity field thus represents observers in free fall towards the singularity in PG coordinates. The experimental protocol (Section B.4.4) will thus use these hypothesized u^μ for freely falling observers (zero accelerometer readings) to obtain the displayed coordinates X^μ by integration along each worldline.

B.4.3 A programmable swarm to display the interior PG coordinates. In our framework, observers follow timelike worldlines (here, free fall), while the displayed coordinates arise from applying an explicitly specified program to the seed-frame coordinates. The resulting map must then be checked to define a valid chart on the working region. We use the rain congruence (37) as a suitable example: together with the initial data specified below, it yields a map $(\tau, \mathbf{y}) \mapsto X^\mu(\tau, \mathbf{y})$ that is locally invertible on the region of interest. Each device is programmed to output the PG coordinates by integrating

$$\frac{dX^\mu}{d\tau} = u^\mu_{\mathrm{PG}}\big(X(\tau, \mathbf{y})\big), \qquad X^\mu(\tau, \mathbf{y}) = (X^T, X^r, X^\theta, X^\phi). \tag{39}$$

Explicitly,

$$\begin{aligned}
\frac{dX^T}{d\tau} &= 1, \\
\frac{dX^r}{d\tau} &= -\sqrt{\frac{2M}{X^r}}, \\
\frac{dX^\theta}{d\tau} &= 0, \qquad \frac{dX^\phi}{d\tau} = 0.
\end{aligned} \tag{40}$$

B.4.4 Physical setup. Choose an interior observation region (bounded away from the singularity and contained in a single angular chart patch) such as

$$\mathcal{U}: \quad r \in [r_{\min}, r_{\max}], \qquad 0 < r_{\min} < r_{\max} < 2M, \qquad T \in [0, T_{\text{obs}}], \tag{41}$$

together with

$$\theta \in (\theta_{\min}, \theta_{\max}) \subset (0, \pi), \qquad \phi \in (\phi_{\min}, \phi_{\max}), \qquad \phi_{\max} - \phi_{\min} < 2\pi.$$

This excludes the poles and any ϕ-branch cut from the working patch; a global treatment would use multiple angular patches.

Pre-deployment initializes a dense set of observers on the slice $T = 0$ with labels $\mathbf{y} = (r_0, \theta_0, \phi_0)$, where

$$r_0 \in [r_{0,\min}, r_{0,\max}], \qquad \theta_0 \in (\theta_{\min}, \theta_{\max}), \qquad \phi_0 \in (\phi_{\min}, \phi_{\max}).$$

Although the observation region $\mathcal{U}$ is entirely interior, the seed congruence need not be: for sufficiently large T_{obs}, some observers that later enter $\mathcal{U}$ can satisfy $r_0 > 2M$ at $T = 0$. For the congruence to cover $\mathcal{U}$, the initialization interval must be chosen so that the inverse map sends every $(T, r) \in [0, T_{\text{obs}}] \times [r_{\min}, r_{\max}]$ into the allowed initial range.[12]

We thus establish the following initial conditions:

$$X^T(0, \mathbf{y}) = 0, \quad X^r(0, \mathbf{y}) = r_0, \quad X^\theta(0, \mathbf{y}) = \theta_0, \quad X^\phi(0, \mathbf{y}) = \phi_0. \tag{42}$$

The simplest physical picture for the congruence may begin outside the horizon. Observers are released at staggered times, each with the inward velocity $dr/d\tau = -\sqrt{2M/r}$ of the PG rain congruence (corresponding to free fall from rest at infinity), and subsequently fall freely. The release times are coordinated so that the required interval of r_0-values is occupied at $T = 0$; depending on T_{obs}, that interval may lie wholly inside the horizon or may straddle it. This is not the only possibility, but it provides a simple scenario: pre-deployment programming can be established before horizon crossing, while the actual data-gathering region remains the interior patch $\mathcal{U}$.

[12] Equivalently, $r_{0,\min} \le \left(r^{3/2} + \frac{3}{2}\sqrt{2M}\,T\right)^{2/3} \le r_{0,\max}$ for all $(T, r) \in [0, T_{\text{obs}}] \times [r_{\min}, r_{\max}]$, see Eq. (50) below. A sufficient choice is $r_{0,\max} \ge \left(r_{\max}^{3/2} + \frac{3}{2}\sqrt{2M}\,T_{\text{obs}}\right)^{2/3}$ and $r_{0,\min} \le r_{\min}$. If one insists that all seed observers already lie inside the horizon at $T = 0$, one must additionally require $\left(r_{\max}^{3/2} + \frac{3}{2}\sqrt{2M}\,T_{\text{obs}}\right)^{2/3} < 2M$, equivalently $T_{\text{obs}} < \frac{2}{3\sqrt{2M}}\left((2M)^{3/2} - r_{\max}^{3/2}\right)$.

B.4.5 The seed is a mollusk. In the seed frame, the initial data are used to define the coordinates

$$X^\mu_{\rm seed} = (\tau, \mathbf{y}) = (\tau, r_0, \theta_0, \phi_0). \tag{43}$$

Each observer sits at fixed labels while only τ advances. The four-velocity takes the congruence-adapted form

$$u^\mu_{\rm seed} = \frac{dX^\mu_{\rm seed}}{d\tau} = (1, 0, 0, 0).$$

In seed coordinates $\lambda(\tau, r_0, \theta_0, \phi_0) = \tau$, so $g_{\tau\tau} = -1$, see Eq. (4). Because the rain congruence is geodesic and irrotational, the seed metric can be written in Gaussian normal form, Eq. (12). Moreover, Eq. (44) shows that the equal-τ leaves coincide with the $T = $ const hypersurfaces of the PG chart, whose induced spatial metric is positive definite (B.4.10). The explicit pullback in B.4.12, Eq. (51), confirms this directly: the cross terms vanish and the spatial part is positive definite for $r > 0$. Hence both mollusk conditions are satisfied on the chosen region.

B.4.6 Coordinates by integration. On top of these seed coordinates, the programmed coordinates $X^\mu(X_{\rm seed})$ are obtained by integrating (40) using the initial data shared with the seed. This gives

$$X^T(\tau, \mathbf{y}) = \tau, \tag{44}$$

$$X^r(\tau, \mathbf{y}) = \left(r_0^{3/2} - \frac{3}{2}\sqrt{2M}\,\tau\right)^{2/3}, \tag{45}$$

$$X^\theta(\tau, \mathbf{y}) = \theta_0, \tag{46}$$

$$X^\phi(\tau, \mathbf{y}) = \phi_0, \tag{47}$$

valid until the singularity at $r = 0$ is reached at finite proper time

$$\tau_{\max}(r_0) = \frac{2}{3\sqrt{2M}}\, r_0^{3/2}. \tag{48}$$

From a Lagrangian viewpoint, Eq. (45) amounts to displaying the infalling r-coordinate along each worldline. From an Eulerian viewpoint, the congruence of worldlines labeled by $\mathbf{y} = (r_0, \theta_0, \phi_0)$ covers $\mathcal{U}$ by sweeping through it during a finite time interval. The map $X^\mu(X_{\rm seed})$, i.e.

$$X^\mu_{\rm seed} = (\tau, r_0, \theta_0, \phi_0) \mapsto (X^T, X^r, X^\theta, X^\phi)$$

has Jacobian $J^\mu_\nu = \partial X^\mu / \partial X^\nu_{\rm seed}$ (we here write $r = X^r$, $T = X^T$)

$$\frac{\partial X^\mu}{\partial(\tau, r_0, \theta_0, \phi_0)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -\sqrt{\dfrac{2M}{r}} & \sqrt{\dfrac{r_0}{r}} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \qquad \det\left(\frac{\partial X^\mu}{\partial(\tau, r_0, \theta_0, \phi_0)}\right) = \sqrt{\frac{r_0}{r}}, \tag{49}$$

which is nonzero for $r > 0$.

We can also write the inverse map $X^\mu_{\rm seed}(X)$ explicitly:

$$\tau = T, \qquad r_0 = \left(r^{3/2} + \frac{3}{2}\sqrt{2M}\,T\right)^{2/3}, \qquad \theta_0 = \theta, \qquad \phi_0 = \phi. \tag{50}$$

Thus the map is not only locally invertible but globally one-to-one on the chosen observation region $\mathcal{U}$, provided $\mathcal{U}$ is chosen away from $r = 0$ and within the finite lifetime of the swarm.

B.4.7 The displayed coordinates are not congruence-adapted. $dX^\mu/d\tau = u^\mu_{\rm PG}$ converts seed labels into PG coordinate values via integration, with the initial conditions identifying the two systems at $\tau = 0$, see Eq. (42). As τ advances, the seed keeps its spatial labels fixed while the displayed PG values evolve. Using the above Jacobian $J^\mu_{\ \nu} = \partial X^\mu/\partial X^\nu_{\rm seed}$, the four-velocity transforms as

$$u^\mu_{\rm PG} = J^\mu_{\ \nu}\, u^\nu_{\rm seed} = J^\mu_{\ 0} = \left(1,\, -\sqrt{2M/r},\, 0,\, 0\right).$$

The restricted seed form $(1,0,0,0)$ and the two-component PG form $(1,-\sqrt{2M/r},0,0)$ describe the same worldlines in different coordinates, cf. again Eq. (37). The underlying motion is identical; what changes is which numbers the devices display: proper time and initial congruence conditions versus actual PG coordinates with four spatial coordinate basis directions (B.4.1). It is instructive to compare the four-velocity components in the two frameworks. In congruence-adapted coordinates, Eq. (5) always forces $u^\nu_{\rm seed}$ to have only one nonzero component: only the clock reading changes along each worldline. In the programmable-swarm framework, by contrast, all four components of u^μ may vary. Eq. (37) for $u^\mu_{\rm PG}$ provides a concrete example with two nonzero components.

B.4.8 PG coordinates in the exterior ($r > 2M$). The worldline equations in B.4.3 and their solutions in B.4.6 are identical in the exterior and interior regions—only the causal character of the coordinate basis directions changes at $r = 2M$. Hence, the same programming allows the enhanced swarm to display PG coordinates both outside the horizon (where a direct mollusk realization is also available) and inside (where the PG chart itself has no direct mollusk realization).

B.4.9 A Mollusk in Schwarzschild interior coordinates ($r < 2M$). Another mollusk can operate inside the horizon—but again not in PG coordinates. In standard Schwarzschild interior coordinates, $g_{rr} < 0$ for $r < 2M$, so the r-direction is timelike; a mollusk adapted to those coordinates would use observers at fixed (t,θ,ϕ) evolving along r. The programmable swarm could also start from such an interior mollusk and transform on the fly to PG coordinates. The key point is that PG coordinates inside the horizon are not available for any mollusk or congruence-adapted seed frame as the raw comoving coordinates; reaching them always requires a further coordinate transformation, which the programmable devices perform operationally on the fly.

B.4.10 ADM formulation in the interior ($r < 2M$). Painlevé–Gullstrand coordinates admit a valid 3+1 (ADM) decomposition. Written as $ds^2 = -dT^2 + \left(dr + \sqrt{2M/r}\, dT\right)^2 + r^2 d\Omega^2$, the ADM fields are $N = 1$, $N^r = \sqrt{2M/r}$ (with $N^\theta = N^\phi = 0$), and $h_{ij} = \mathrm{diag}(1, r^2, r^2\sin^2\theta)$, i.e. intrinsically flat Euclidean 3-space. These conditions remain well-behaved both outside and inside the horizon ($0 < r < 2M$): h_{ij} is positive definite, the lapse is positive, and the unit normal $n^\mu = (1, -\sqrt{2M/r}, 0, 0)$ to the $T = \text{const}$ hypersurfaces is timelike with $g_{\mu\nu}n^\mu n^\nu = -1$. Thus, the foliation stays regular and horizon-penetrating (up to the physical singularity at $r = 0$) in the region covered by ingoing PG coordinates. What fails inside the horizon is the mollusk and congruence-adapted seed-frame condition: $g_{TT} > 0$ makes the coordinate-time vector spacelike, so no massive observer can remain at fixed (r,θ,ϕ) while "advancing in T". The unit normal n^μ remains of course timelike throughout, as guaranteed by the ADM construction. This illustrates that a valid ADM foliation does not by itself guarantee a mollusk realization, and that the programmable coordinates can mirror that foliation nonetheless.

B.4.11 A lesson about coordinates. A mollusk-type frame attempts to realize coordinates by assigning fixed spatial labels to observer worldlines. To "instantiate" the PG chart in that sense, observers would have to sit at, for example, fixed (r, θ, ϕ) and advance in T. But (34) shows those T-lines are spacelike for $r < 2M$. Nevertheless, along each infalling worldline, the PG coordinate T increases monotonically: since $dT/d\tau = 1$, each observer's device displays a steadily growing value of T as proper time elapses. One might therefore expect T to behave like a time coordinate in every sense. There is no contradiction: the infalling observers do not move in the direction of T. Their worldlines have $dr/d\tau \neq 0$, so they move through a *combination* of coordinate basis directions—precisely the timelike combination given by Eq. (37). A coordinate can thus increase monotonically along timelike worldlines even when its associated coordinate basis direction is spacelike. These are distinct notions: the displayed coordinate values are generated by each observer individually along their own worldline, whereas a coordinate curve (varying one coordinate while holding the others fixed) cuts across multiple worldlines and must be reconstructed from the full congruence.

B.4.12 Testing the metric hypothesis. The programmed functions $X^\mu(\tau, \mathbf{y})$ reflect coordinates chosen to reproduce the form of the metric underlying the metric hypothesis. The operational test comes when one measures the physical metric, for example using local Einstein elevators and the procedures described in the main text. If empirical reconstruction yields Eq. (33), the hypothesis is supported. Even if the measured metric differs from the hypothesized one, the combination of reference frame plus measured metric still provides a valid description of reality. This follows from ordinary chart freedom.

However, these statements only apply on those regions where the displayed coordinates remain valid. Since the actual physical situation differs from the situation assumed when establishing the coordinate system based on a metric hypothesis, validity failures of the coordinates may become apparent during deployment. For example, neighboring observers may pile up (indicating the presence of caustics). On those regions that remain valid, however, the collected data may be analyzed offline and subjected to further coordinate transformations, if such better adapted coordinates can be established for the actually measured metric.

If each observer additionally displays the seed coordinates, the metric can be reconstructed in those coordinates as well. We obtain the metric by pulling back the PG line element (33) through the Jacobian (49). Substituting $dT = d\tau$ and $dr = -\sqrt{2M/r}\, d\tau + \sqrt{r_0/r}\, dr_0$ yields

$$ds^2 = -d\tau^2 + \frac{r_0}{r(\tau, r_0)}\, dr_0^2 + r^2(\tau, r_0)\big(d\theta_0^2 + \sin^2\theta_0\, d\phi_0^2\big), \tag{51}$$

where $r(\tau, r_0) = \left(r_0^{3/2} - \frac{3}{2}\sqrt{2M}\,\tau\right)^{2/3}$ is given by Eq. (45). This is the Gaussian normal form anticipated in B.4.5: the cross terms cancel identically, confirming $g_{\tau i} = 0$, and the spatial part is positive definite for $r > 0$. With

$$R = \frac{2}{3}\frac{r_0^{3/2}}{\sqrt{2M}}, \qquad r = (2M)^{1/3}\left[\frac{3}{2}(R - \tau)\right]^{2/3},$$

Eq. (51) can be rewritten as

$$ds^2 = -d\tau^2 + \frac{2M}{r}\, dR^2 + r^2\big(d\theta_0^2 + \sin^2\theta_0\, d\phi_0^2\big), \tag{52}$$

i.e. the metric in standard Lemaître coordinates in mostly-plus signature.

Note that in line elements (51) and (52), τ is proper time for the radially infalling observers at fixed θ_0 and ϕ_0; for an arbitrary worldline through the spacetime, τ, r_0, θ_0, and ϕ_0 are coordinates that vary along the worldline in line element (51), while τ, R, θ_0, and ϕ_0 are the coordinates of (52).

B.5 Kruskal–Szekeres Coordinates: two lightlike coordinate basis directions

While the PG example focused on assuming and testing a metric hypothesis in a scenario in which all coordinate basis directions are spacelike, the present example highlights a different capability: once valid seed coordinates are available on a chosen region, a single physical observer swarm can be (re)programmed to display other compatible coordinates on that same region, including charts with null coordinate basis directions. Such coordinates are not available as the raw output of traditional congruence-adapted seed frames; they require a further coordinate transformation.

For the subsequent formulas we assume the exterior scenario $r > 2M$ (see B.4.8) to further demonstrate the versatility of the approach.[13]

An observer swarm thus populates a region $\mathcal{U} \subset \{r > 2M\}$ of Schwarzschild spacetime and, on a fixed local angular patch, displays the PG coordinate values $x^\mu = (T, r, \theta, \phi)$ along its timelike worldlines via a valid coordinate transform

$$(\tau, \mathbf{y}) \;\longmapsto\; X^\mu(\tau, \mathbf{y}). \tag{53}$$

The very same swarm can then display an alternative coordinate system by applying another transformation $\tilde{X}^\mu(X)$. The double null Kruskal–Szekeres (KS) coordinates provide a useful contrast due to their incompatibility with mollusk or congruence-adapted seed frames (see Appendix A.2). The enhanced swarm implements them without changing its motion—observers on timelike worldlines display coordinate values aligned with null directions in spacetime.

Coordinate transformation $(T, r) \mapsto (U, V)$, $r > 2M$. Horizons were excluded in the main text in part because of complications that may arise with measurement collection, even though coordinate labels could still be assigned for the present example (for $r < 2M$ by adjusting the sign conventions in the KS definitions).

Since both systems share the angular coordinates (θ, ϕ), we restrict to the $(T, r) \leftrightarrow (U, V)$ sector. The Schwarzschild tortoise coordinate is (see [3], Eq. (5.5))

$$r^* = r + 2M \ln\left|\frac{r}{2M} - 1\right|, \tag{54}$$

and the null coordinates are $u = t - r^*$, $v = t + r^*$, where t is Schwarzschild time. Double null Kruskal–Szekeres coordinates are (see [3], Eqs. (5.7)–(5.9)):

$$U = -\exp\left(-\frac{u}{4M}\right), \quad V = \exp\left(\frac{v}{4M}\right). \tag{55}$$

The relation between Schwarzschild time t and PG time T is (see [9], Eq. (2.5)):

$$T = t + 4M\left(\sqrt{\frac{r}{2M}} + \frac{1}{2}\ln\left|\frac{\sqrt{r/2M} - 1}{\sqrt{r/2M} + 1}\right|\right) =: t + F(r). \tag{56}$$

Combining the relations $t = T - F(r)$ and $r^*(r)$ defined above, the Kruskal–Szekeres coordinates (U, V) can be written directly as smooth functions of the PG coordinates (T, r).[14]

[13] The interior KS regions are obtained by the corresponding region-dependent sign choices for U and V; the implicit relation between r and UV is unchanged.

[14] Here the additive constant in the PG time T is fixed to zero. A shift $T \mapsto T + T_0$ would rescale the KS coordinates as $U \mapsto e^{-T_0/(4M)}U$ and $V \mapsto e^{T_0/(4M)}V$, leaving the invariant product $-UV$ unchanged.

From PG coordinates (T, r) to Kruskal–Szekeres coordinates (U, V), $r > 2M$

For $r > 2M$ the Kruskal–Szekeres coordinates can be written directly as functions of the Painlevé–Gullstrand coordinates (T, r):

$$s(r) = \sqrt{\frac{r}{2M}}$$

$$U(T,r) = -\,(s(r)-1)\,\exp\!\left(-\frac{T}{4M} + s(r)\left(\frac{1}{2}s(r)+1\right)\right), \tag{57}$$

$$V(T,r) = \quad (s(r)+1)\,\exp\!\left(\frac{T}{4M} + s(r)\left(\frac{1}{2}s(r)-1\right)\right). \tag{58}$$

These define the usual exterior Kruskal patch and satisfy

$$-\,U(T,r)\,V(T,r) = \left(\tfrac{r}{2M}-1\right)\exp\!\left(\tfrac{r}{2M}\right) = (s^2-1)e^{s^2}.$$

These expressions determine (U, V) as smooth functions of (T, r), defining a coordinate transformation $\tilde{X}^\mu = \tilde{X}^\mu(X)$ with $\tilde{X}^\mu = (U, V, \theta, \phi)$ and $X^\mu = (T, r, \theta, \phi)$. Correspondingly, the Jacobian has nonzero determinant

$$\det\!\left(\frac{\partial(U,V,\theta,\phi)}{\partial(T,r,\theta,\phi)}\right) = \frac{r}{16M^3}\,\exp\!\left(\frac{r}{2M}\right) = \frac{s^2 e^{s^2}}{8M^2}. \tag{59}$$

Programming of the swarm. Each observer programs its device to compute KS values alongside PG values. At each proper time reading τ along the worldline $\mathbf{y}$, the device obtains the current PG coordinates $X^\mu(\tau, \mathbf{y})$, computes $\tilde{X}^\mu\big(X(\tau, \mathbf{y})\big)$, and displays both in separate windows. Because every observer $\mathbf{y}$ uses the same transformation, the swarm collectively implements the four-dimensional KS chart over $\mathcal{U}$.

Null directions from timelike worldlines. The KS chart has two null coordinate basis directions (obtained by varying only U or only V) and two spacelike ones (obtained by varying only θ or only ϕ)—a structure unavailable as the raw comoving output of mollusk- or congruence-adapted seed frames (see Table 1). Yet the swarm enhanced by a computational layer displays KS coordinates on the fly on the same region: observers still follow timelike geodesics with zero accelerometer readings, but the numbers displayed at each event now correspond to a chart with null directions. The timelike motion of observers and the causal character of displayed coordinate directions are decoupled: a single observer swarm can display coordinates that such seed frames do not directly realize on their own, including coordinates with null coordinate basis directions, provided that the programmed transformation satisfies the usual chart validity conditions on the region of interest.

Summary: The fundamental lesson remains the same as in the main text: a single conceptual framework—a material seed swarm supplemented by programmable devices—covers a variety of scenarios without requiring students to switch between widely different mental models, thus providing a single pedagogical picture. Within the idealizations adopted here, the same computationally enhanced swarm can display different compatible coordinate systems on the same region. For conceptual discussion and visualization, this provides a concrete operational picture to which students may return when interpreting unfamiliar coordinate systems.

Appendix C: Mathematical questions for guided exploration

Note to instructors: This appendix provides optional exploratory material for instructors and readers interested in theoretical limits of the proposed concept. It is not needed for the main pedagogical argument, which rests only on explicit programmed transformations whose validity is checked directly, case by case, on the chosen region. For the standing idealizations on the computational layer, see Suppl. Mat., Assumptions.

- **The question**

A stronger mathematical question suggested by the pedagogical analogy between coordinate freedom and programming freedom is the following. Let $k \geq 1$ be a finite integer. Let $S \subset \mathbb{R}^4$ be open and bounded and with compact closure $\overline{S}$. Let

$$\Phi : S \to \Phi(S) \subset \mathbb{R}^4$$

be a valid chart transition that extends to a C^k diffeomorphism on a slightly larger open neighborhood $\mathcal{W}$ of $\overline{S}$.[15] One may then ask whether, under suitable additional hypotheses, for every $\varepsilon > 0$ there exists a rational-coefficient polynomial map

$$F : \mathbb{R}^4 \to \mathbb{R}^4$$

such that

$$\|F - \Phi\|_{C^k(\overline{S})} < \varepsilon$$

and such that the restriction

$$F|_S : S \to F(S)$$

is still a C^k diffeomorphism onto its image. Here $C^k(\overline{S})$-closeness means the usual uniform closeness on $\overline{S}$ of the map values and of all partial derivatives up to order k, a natural measure of closeness for this approximation problem. Working with a compact closure $\overline{S}$ is natural for operationalizations and for simplifying the mathematical question.

- **Interpretation**

Proposition 1 below answers this question positively under the stated hypotheses. In the above operationally restricted sense, every valid chart transition Φ on a fixed working region S lies arbitrarily close to a rational-coefficient polynomial surrogate F that remains a valid chart transition on that region. Since rational-coefficient polynomial maps are forward-computable, we may conceptually restrict each observer's device to compute only such maps from the seed coordinates during deployment. Such observers would lose no expressivity in the following approximation-theoretic sense: the class of such maps that are valid chart transitions on the working region is dense in the space of all valid chart transitions. This result supplies one precise sense in which an idealized programmable layer approximates ordinary chart freedom on a fixed region under the stated additional assumptions.

The result may be of general interest: since any physical device operates at finite precision, similar approximations are inherent in any real-world implementation and not a property specific to the programmable layer; the present construction merely sharpens the operational question.

Before turning to the solution, we stress that chart validity must still be verified on a case-by-case basis, just as for continuum transition maps.

[15] Such technical assumptions facilitate the approach. Alternatively, S may be covered from inside by subsets.

- **Hints**

A possible route to the stronger mathematical question may combine two standard ingredients. The first is approximation by rational-coefficient multivariate Bernstein polynomials (e.g., [10], Thm. 6.7 plus truncation to rationals, or Thm. 6.10). However, this would not yet ensure that $F|_S$ is itself a diffeomorphism: one would still have to verify that F has everywhere non-vanishing Jacobian $\det(DF) \neq 0$ and is injective on S, that is $F(x) \neq F(y)$ for distinct $x, y \in S$. Both properties hold for Φ; whether they transfer to a sufficiently C^k-close F under the assumptions stated above may be addressed via standard techniques in differential topology, where the set of diffeomorphisms (and the set of embeddings) is known to be open under appropriate conditions (see e.g., [11], Ch. 2).

- **Details**

The proof presented below extends Φ to all of $\mathbb{R}^4$ via a smooth cutoff, approximates the extension in the C^k norm by rational-coefficient Bernstein polynomials on a slightly larger compact set, and then uses the compact-set openness of embeddings in the C^1 topology to ensure that the polynomial approximant remains an embedding on S. Since source and target have the same dimension, the resulting embedding is a diffeomorphism onto its image. For mathematically inclined readers, the details are given in Proposition 1 below.

The proposition is stated for arbitrary finite $k \geq 1$. For basic chart validity and for the pointwise transformation law of the metric, $k = 1$ suffices. In GR applications, however, $k \geq 2$ is the natural minimum if one wants the transformed metric to be at least C^1 and the Christoffel symbols to be classically well defined. Indeed,

$$g'_{\alpha\beta} = \frac{\partial x^\mu}{\partial x'^\alpha} \frac{\partial x^\nu}{\partial x'^\beta} g_{\mu\nu}$$

uses first derivatives of the coordinate change, whereas the transformation law for Christoffel symbols contains a term of the form

$$\frac{\partial x'^\alpha}{\partial x^\mu} \frac{\partial^2 x^\mu}{\partial x'^\beta \partial x'^\gamma},$$

which requires second derivatives. For curvature or Einstein-equation manipulations one usually assumes whatever higher regularity is needed for the specific classical formulas being used.

As mentioned, the technical details provided below may help guided-exploration tasks proposed by interested instructors, combining two known results from approximation theory (rational-coefficient polynomials) and differential topology (openness of embeddings).[16] These details are intended only as optional mathematical orientation; the pedagogical claims do not depend on it; they use the analogy as a pedagogical tool exemplified through the different scenarios presented in Appendix B and the main text.

[16] One could use any other type of effectively specifiable functions (e.g., triangulations or Fourier coefficients); the question is formulated here using polynomials for simplicity. Also, the set of computable input values (computable reals or rationals) is dense in $\mathbb{R}^4$, thus preserving any established approximation properties. Moreover, because the computation always starts from the seed reference frame coordinates, one would also only require the forward map F to be computable and to restrict to a diffeomorphism on S. The global map $F : \mathbb{R}^4 \to \mathbb{R}^4$ need not be bijective, nor is it assumed here that the inverse of $F|_S$ is polynomial (usually not the case). In fact, since we do not attempt to compute the inverse (e.g., by Newton iteration), we need not even assume it is computable. The question, as posed, concerns chart transitions as maps between subsets of $\mathbb{R}^4$ and is independent of the spacetime metric and signature. Extending any such result to a larger region covered by multiple overlapping patches would require verifying the usual atlas conditions: notably, additional chart compatibility conditions on the overlaps; when modeling finite operational tolerance, one might instead ask for approximate compatibility.

Proposition 1. *Let $k \geq 1$ be finite. Assume that $S \subset \mathbb{R}^4$ is open and bounded with compact closure $\overline{S}$, and that*

$$\Phi : \mathcal{W} \to \Phi(\mathcal{W}) \subset \mathbb{R}^4$$

is a C^k diffeomorphism on an open neighborhood $\mathcal{W}$ of $\overline{S}$. Then for every $\varepsilon > 0$ there exists a rational-coefficient polynomial map $F : \mathbb{R}^4 \to \mathbb{R}^4$ such that

$$\|F - \Phi\|_{C^k(\overline{S})} < \varepsilon$$

and such that

$$F|_S : S \to F(S)$$

is a C^k diffeomorphism onto its image. In fact, $F|_S$ is smooth.

Proof. Choose a bounded open set V such that $\overline{S} \subset V \subset \overline{V} \subset \mathcal{W}$. Choose $\chi \in C_c^\infty(\mathcal{W})$ such that $\chi \equiv 1$ on $\overline{V}$. For each component Φ^μ, extend the domain from $\mathcal{W}$ to all of $\mathbb{R}^4$ using the smooth cutoff χ:

$$G^\mu(x) = \begin{cases} \chi(x)\Phi^\mu(x), & x \in \mathcal{W}, \\ 0, & x \notin \mathcal{W}. \end{cases}$$

Then $G = (G^0, G^1, G^2, G^3) \in C^k(\mathbb{R}^4, \mathbb{R}^4)$ and $G = \Phi$ on $\overline{V}$. Note that $\operatorname{supp}\chi$ is a compact subset of $\mathcal{W}$, i.e. χ vanishes outside a compact set contained in $\mathcal{W}$, so the two branches match with all derivatives.

• *Part 1. Approximation.*

Choose a closed box $Q = [-N, N]^4$ with $N \in \mathbb{N}$ and $\overline{V} \subset \operatorname{Int} Q$. After an affine rescaling between Q and $[0,1]^4$, approximation by multivariate Bernstein polynomials ([10, Thm. 6.7]) yields, for each $n \in \mathbb{N}$, a real-coefficient polynomial map $P_n : \mathbb{R}^4 \to \mathbb{R}^4$ such that

$$\|P_n - G\|_{C^k(Q)} < \frac{1}{2n}.$$

Indeed, one applies the scalar theorem to each component of G and to the finite set of multiindices β with $|\beta| \leq k$, and then chooses one polynomial index large enough for all components and all such β simultaneously. Since the space of polynomial maps of degree at most $\deg P_n$ is finite-dimensional, perturbing the coefficients slightly changes the $C^k(Q)$-norm continuously. Hence we may perturb the coefficients of P_n to rational numbers (rationals being dense in the reals) and obtain a rational-coefficient polynomial map $F_n : \mathbb{R}^4 \to \mathbb{R}^4$ with

$$\|F_n - P_n\|_{C^k(Q)} < \frac{1}{2n}.$$

Therefore, by the triangle inequality,

$$\|F_n - G\|_{C^k(Q)} \leq \|F_n - P_n\|_{C^k(Q)} + \|P_n - G\|_{C^k(Q)} < \frac{1}{n},$$

and in particular

$$\|F_n - \Phi\|_{C^k(\overline{V})} = \|F_n - G\|_{C^k(\overline{V})} \to 0.$$

• *Part 2. Diffeomorphism.*

Now $\Phi|_V$ is a C^k embedding, and $\overline{S} \subset V$ is compact. By the compact-set openness of embeddings in the C^1 topology, every sufficiently small C^1 perturbation of $\Phi|_V$ is still an embedding on some open neighborhood of $\overline{S}$ in V; see [11], Ch. 2, Exercise 14 (with V as M, $\mathbb{R}^4$ as N, $\overline{S}$ as K, and $\Phi|_V$ as f). Since $F_n|_V \to \Phi|_V$ in $C^1(\overline{V})$, it follows that for all sufficiently large n there exists an open neighborhood $U_n \subset V$ of $\overline{S}$ such that $F_n|_{U_n}$ is an embedding. In particular, $F_n|_S$ is an embedding (since $S \subset U_n$ for each such n). Now choose an n also large enough that $\|F_n - \Phi\|_{C^k(\overline{S})} < \varepsilon$ (which is possible since $\|F_n - \Phi\|_{C^k(\overline{S})} \leq \|F_n - \Phi\|_{C^k(\overline{V})} \to 0$).

Because F_n is polynomial, it is smooth, i.e. C^∞ (note here that the C^k diffeomorphism property of the conclusion will follow from the smoothness of F_n itself, not from the C^1 topology used). Since $F_n|_S$ is an embedding, it is an injective immersion that is a homeomorphism onto its image. So $D(F_n)_x$ is injective for every $x \in S$. As $\dim S = \dim \mathbb{R}^4 = 4$, $D(F_n)_x$ is therefore an isomorphism. Hence, by the inverse function theorem, $F_n|_S$ has a smooth local inverse at every point (the inverse inherits the smoothness of F_n). Being also globally injective as an embedding, $F_n|_S$ is a smooth diffeomorphism from S onto the open set $F_n(S)$. Setting $F := F_n$ proves the proposition. □

References

[1] Westman, H. and Sonego, S. Coordinates, Observables and Symmetry in Relativity. *Annals of Physics*, 324(8):1585–1611, 2009.

[2] L. D. Landau and E. M. Lifshitz. *The Classical Theory of Fields*, volume 2 of *Course of Theoretical Physics*. Pergamon Press, Oxford, 3rd rev. ed., 1971.

[3] Poisson, E. *A Relativist's Toolkit: The Mathematics of Black-Hole Mechanics*. Cambridge University Press, 2004.

[4] P. A. M. Dirac. Forms of relativistic dynamics. *Reviews of Modern Physics*, 21:392–399, 1949.

[5] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7):769–822, 1916.

[6] Misner, Ch. W. and Thorne, K. S. and Wheeler, J. A. *Gravitation*. W. H. Freeman, San Francisco, CA, 1973.

[7] Weinberg, Steven. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. John Wiley & Sons, New York, NY, 1972.

[8] Colin MacLaurin. Schwarzschild spacetime under generalised Gullstrand–Painlevé slicing. In S. Cacciatori, B. Güneysu, and S. Pigola, editors, *Einstein Equations: Physical and Mathematical Aspects of General Relativity*, Tutorials, Schools, and Workshops in the Mathematical Sciences, pages 267–287. Birkhäuser, Cham, 2019.

[9] Martel, K. and Poisson, E. Regular coordinate systems for Schwarzschild and other spherical spacetimes. *American Journal of Physics*, 69(4):476–480, April 2001.

[10] Adrian Fellhauer. Approximation of smooth functions using Bernstein polynomials in multiple variables. 2016. arXiv:1609.01940 [math.CA].

[11] Morris W. Hirsch. *Differential Topology*, volume 33 of *Graduate Texts in Mathematics*. Springer, 1976.