

Particle Streams in a Four-Singularity Potential Well

Verification, GPU implementation, and a dimensional scan
Ongoing research log — rebuilt 2026-08-16 21:28 UTC

1. Scope, and one honest boundary

This document records an executable study of **high-density particle streams moving in the field of four fixed singularities**, with geodesic-style dynamics, ricochet near the singular surfaces, and long-time orbit classification. It covers verification of the physics, a GPU reimplement, measured performance, a scan over spatial dimension, and a time-reversibility experiment.

Before anything else, a boundary that has to be stated plainly, because everything downstream depends on it:

This simulation is Newtonian. It is not general relativity. It integrates $H = |p|^2/2m + m\Phi(r)$ in flat space with a softened multi-centre potential. There is no metric tensor, no light cone, no proper time, no frame dragging, no radiation reaction, and no event horizon in any relativistic sense. The "horizon" here is a **hard sphere**: the particle hits radius R and reflects specularly. That is a numerical regularization chosen to keep the ODE integrable near the centres — it is a modelling decision, not a physical result.

Everything measured below is a true statement about *this Newtonian system*. Section 9 addresses the relativistic questions separately and marks clearly where established physics ends and speculation begins. The two are not mixed.

1.1 What was inherited

A scaffold existed — `potential_fields.py`, `particle_dynamics.py`, `main_sim.py`, `verify.py` — written in a sandbox whose shell never started. **No line of it had ever been executed.** It was treated as unverified until proven otherwise, which turned out to be the right posture but a generous one: the CPU code was correct on first run.

1.2 The model

In normalized units ($G = 1$, test mass $m = 1$), the softened (Plummer-regularized) four-centre potential per unit mass is

```
Phi(r) = -G * sum_{i=1..4} q_i / sqrt(|r - r_i|^2 + eps^2)

dr/dt = v
dv/dt = -G sum_i q_i (r - r_i) / (|r - r_i|^2 + eps^2)^(3/2)
```

eps (softening) removes the $1/r$ divergence so the ODE stays integrable; the constructor enforces $\text{eps} \leq R_{\min}/2$ so the regularization never reaches into the sampled region. The physical surface $|r - r_i| = R_i$ is handled discretely: a terminal root-finding event stops integration, the velocity is specularly reflected,

```
n = (r - r_i) / |r - r_i|
v' = v - (1 + e)(v · n) n      applied only when v · n < 0
```

and integration restarts just outside. $e = 1$ is elastic and conserves energy exactly, which is what makes energy drift a clean diagnostic of integrator quality rather than a mix of physics and numerics.

Three geometries ship: **tetrahedron** (four $+q$ at tetrahedron vertices, monopole far field), **planar** (square, alternating $\pm q$, true quadrupole with $1/r^3$ far field), and **collinear** (four on the x-axis, an exact-symmetry control case).

2. Verification of the inherited physics

`verify.py` runs nine groups of checks — analytic gradient against central differences, tidal tensor against a numerical Hessian, quadrupole far-field exponent, the Kepler limit, energy conservation, elastic ricochet, and mirror symmetry. **All 20 assertions passed on the first execution**, with no tolerance loosened. Highlights:

Check	Result
<code>a == -∇Φ</code> (central differences)	max rel err 1.62e-10
Tidal tensor vs numerical Hessian	max rel err 4.31e-10 , asymmetry exactly 0
Planar quadrupole far field	fitted exponent 3.000
Kepler closure after 10 periods	4.71e-07 relative
Kepler energy drift	2.31e-11
Four-centre energy drift (t = 120)	2.68e-09
Elastic ricochet speed	6.167651 → 6.167651 (< 1e-9)
Mirror symmetry (collinear)	0.00e+00 — bit-exact

Two caveats worth recording rather than glossing:

- The mirror-symmetry result being *exactly* zero is not luck. Negating a float is exact in IEEE-754, and the error norms driving DOP853's step controller are identical for the mirrored problem, so the two integrations take bit-identical step sequences. It is a strong check of the field code and a weak check of the integrator.
- `test_bound_orbit_energy` recorded **0 bounces** and a minimum horizon gap of 5.08e-02. That orbit never approaches a horizon, so its "never penetrates a horizon" assertion is *vacuous*. Ricochet is genuinely exercised only by `test_ricochet` (5 bounces). This is a real gap in the inherited test suite.

3. The GPU implementation

3.1 Why `solve_ivp` had to go

`scipy.integrate.solve_ivp` is CPU-only and integrates exactly one particle per call. A stream of 10■ particles is 10■ serial Python-level ODE solves. But the field is static and identical for every particle, so the ensemble is embarrassingly parallel — the physics has no coupling between particles at all.

The replacement is a **batched Dormand-Prince 5(4)** in PyTorch where every particle carries its own adaptive step size and its own event state. The ensemble advances in lockstep at the level of *kernel launches*, not at the level of simulated time — particle A can be taking 1e-6 steps near a horizon while particle B coasts at 1e-1 in the far field, in the same kernel.

3.2 The parts that were not obvious

Event handling. The discontinuous ricochet is what makes this non-trivial. Four horizon surfaces and the escape sphere are folded into one continuous scalar

$$G(y) = \min(\min_i (|r - c_i| - R_i), R_{\text{escape}} - |r|)$$

positive in the physical region, negative once anything is crossed. A single root-finder then serves all five surfaces, and which one fired is decided afterwards from which term attained the minimum. The root is bracketed across the accepted step and refined with **Illinois regula falsi where every trial evaluation is a real RK5 step** of size `th` — not an interpolant. The located crossing therefore carries the integrator's own order of accuracy rather than a lower-order dense-output polynomial's.

Exact output-grid landing. `h` is clamped so a step never overshoots the next sample time. Particles land *exactly* on grid points, so recording needs no interpolation, and since at most one grid point can fall in a step (at its very end) an event resolved at `th < 1` can never skip a sample.

float64 throughout. The shadow-orbit separation starts at 1e-8 against coordinates of order 1. float32 carries ~1e-7 relative precision and would put the entire Lyapunov measurement below its own noise floor. On a consumer card FP64 runs at 1/64 the FP32 rate, but this kernel is launch- and bandwidth-bound rather than FLOP-bound, so the practical penalty is far smaller than that ratio implies.

3.3 Three bugs found by validation

These are recorded because the validation suite existing is the only reason they were caught.

- 1 **Conflated step-clamping semantics.** A single `clamped` flag meant "the step was shortened", which was true whether the *output grid* or the *proximity cap* did the shortening. When the cap shortened a step, the particle advanced its state by a tiny `h` but had its **clock jumped to the next grid point** — time ran ahead of the physics. Symptom: the GPU logged 3 ricochets where the CPU logged 5, because GPU particles covered less orbit per unit `t`. Fixed by separating `shortened` (don't grow `h`) from `grid_hit` (land on the grid and record).
- 1 **A Zeno trap in the proximity cap.** The safety cap limiting step size near a horizon was written as `h ∝ gap`. That makes the gap shrink geometrically and the surface is *never actually reached*. Replaced with a constant cap `h = 0.5 · R / |v|`, which still makes jumping clean over a horizon geometrically impossible but terminates.
- 1 **Retired particles re-flushed every iteration**, overwriting each particle's true `t_end` with `t_max`. Fixed with a freeze-then-retire scheme: a finished particle has its accepted-step mask cleared immediately, so its diagnostics are frozen at the instant it finished, and the bookkeeping flush happens lazily at compaction.

3.4 One deliberate divergence from the CPU reference

scipy's event machinery sees sign changes **only at step endpoints**. A step that dives through a horizon and back out in one bite is invisible to it — and would have been invisible to the GPU version too. The GPU path adds a proximity cap that makes that tunnelling geometrically impossible. This is an improvement, not a match, so **the number of times it binds is counted and reported** (`proximity_cap_bindings`) rather than left silent. In the production run it bound **0** times, so it changed nothing there — it is insurance, not an active correction.

4. Validating the GPU against the CPU

4.1 What can and cannot be compared

The four-centre problem is chaotic, with measured Lyapunov exponents around 0.1–0.3. Two integrators differing in the last bit of the initial state — let alone two different Runge-Kutta schemes — diverge exponentially and are uncorrelated by $t = 300$. **That divergence is physics, not a defect.** Demanding long-time trajectory agreement in the chaotic regime would be demanding the wrong thing, and any suite that "passed" such a test would be lying.

So the checks are layered by what is legitimately comparable:

Layer	What it tests	Result
A	Field kernels vs numpy	3.2e-16 max relative
B	Kepler (integrable), 10 periods	trajectories agree to 1.64e-08
C	Head-on ricochet, 5 bounces	bounce count 5 = 5 , trajectory 3.50e-06
D	Chaotic orbit at $t = 20$	6-D deviation 1.89e-10
E	Chaotic orbit, long time	divergence fitted at $\exp(+0.210\ t)$
F	Energy drift, 64-particle batch	max 1.32e-09 , median 4.14e-11
G	Batch invariance	1.22e-13 over the first quarter
H	Ensemble outcomes, $t = 40$	escape flags and ricochet counts agree particle-by-particle

Layer E is the one that matters conceptually. Rather than hiding the long-time disagreement, it is *measured*: the CPU/GPU separation grows as **$\exp(+0.210 \cdot t)$** , which is the system's own Lyapunov rate. The disagreement has the signature of chaos, not of a systematic error — and that is a stronger statement than agreement would have been.

Layer G was initially written as "must be bit-identical" and **failed at 2.22e-12**. That assertion was wrong, not the code: cuBLAS/ATen select different reduction kernels for different batch sizes, so a sum can differ in its last bit, and one ULP amplified by **$\exp(0.21 \cdot t)$** over $t = 60$ is exactly $2e-12$. The test now asserts round-off agreement *early*, before amplification acts, and reports the late-time value.

5. Measured performance

Both sides do identical work: same field, same tolerances, same `t_max`, same initial conditions. The CPU rate is measured on 24 particles and extrapolates linearly — which is exactly right for a serial per-particle loop, and is stated as an extrapolation, not a measurement.


CPU reference: 5.05 particles/s (scipy DOP853, serial).

N	wall	ms/iter	particles/s	speedup
1,024	32.9 s	10.2	31	6.2×
4,096	59.0 s	17.7	69	13.7×
16,384	119.7 s	33.6	137	27.1×
65,536	171.1 s	48.0	383	75.8×
262,144	206.4 s	49.8	1,270	251×


The shape of that table is the whole story. Going from 1k to 262k particles — **256× the work** — costs only **4.9× more time per iteration**. At small `N` the kernels are trivial and the loop is bound by launch latency and Python dispatch; the GPU sits ~38% utilized while the CPU issues work. Extra particles in that regime are nearly free, which is where the speedup comes from. The plateau between 65k and 262k (48.0 → 49.8 ms/iter) is the onset of memory-bandwidth saturation.

The honest reading: 251× is a throughput ratio at `N = 262,144`, not a universal speedup. At `N = 32` the GPU is *slower* than scipy (18.5 s vs 5.6 s), which the validation suite prints rather than hides. Break-even is around `N ≈ 150–200`.

6. What the trajectories actually did



Trajectories from the production run: 800 tracked particles of the 100,000-particle ensemble, coloured by classification, with the four singularities marked.



Ensemble diagnostics: escape fraction vs launch speed, Lyapunov spectrum, energy-drift distribution, orbit class vs speed, closest approach, and ricochet counts.

The stream is launched from $(3.0, 0.4, 0.6)$ into a cone of half-angle 0.30 rad about $(-1, -0.13, -0.2)$, with speeds swept across $0.55 \rightarrow 1.85$. The local escape speed at the launch point is **1.6097**, so the ensemble deliberately straddles it.

99,856 particles integrated to $t = 300$ in 1,108 s — 15,600 GPU loop iterations totalling **655,885,314 RK steps**, of which 1.66% were rejected by the error controller.

class	count	share
chaotic	64,631	64.7%
escape	22,120	22.2%
quasi-periodic	13,070	13.1%
periodic	35	0.035%

Four regimes appear, ordered by launch speed:

- **Chaotic scattering** — the dominant population. Particles fall into the four-centre region, undergo repeated close encounters, and wander with strongly positive Lyapunov exponents. This is a chaotic scattering system, which is exactly what four attracting centres should produce. The measured spectrum runs from $\lambda = 0.0058$ to 0.512 , median **0.0547**, with the 5th–95th percentile band spanning 0.0126 to 0.203.
- **Escape** — above threshold, particles leave on unbound paths. But see §8.4: 14.3% of these are not actually unbound, and that is a defect in the criterion rather than a property of the system.
- **Quasi-periodic** — wide rosette orbits, bounded, low λ , concentrated just below the escape threshold. Visible in the class-vs-speed panel as the blue wedge growing between $|v_{\text{esc}}| \approx 1.35$ and the escape edge.

- **Periodic — 35 of them.** The 16-particle CPU run found *none*, so this class only appears at scale. Their λ median is 0.0107, genuinely below the chaos threshold. But their recurrence values span **0.0233 to 0.0499** against a tolerance of **0.05** — every single one sits just inside the cutoff. Tighten the tolerance to 0.02 and this class vanishes entirely. They are best read as "the most nearly-recurrent orbits found", not as proven periodic orbits.

Two further readouts worth recording:

- **Closest approach.** A sharp spike sits exactly at the horizon radius $R = 0.05$ (the particles that bounced), with a spread out to $r \approx 1$. **Nothing goes below it.** The softening $\epsilon = 0.001$ is more than an order of magnitude below the minimum sampled distance, so the Plummer regularization never influenced the dynamics — which is precisely the design intent, now confirmed rather than assumed.
 - **Ricochets.** 81,614 horizon encounters total, distributed exponentially from 0 to a maximum of 12 per particle.
-

7. The dimensional scan

The request was for "all dimensions", not a 2-D or 3-D plane. Taken seriously, that cannot mean keeping a $1/r$ potential and adding coordinates — that would be the 3-D force law wearing extra axes. The physical generalization comes from the field equation: a point source in D spatial dimensions obeys Laplace's equation, whose Green's function gives

$$\begin{aligned}\Phi(r) &\sim -1 / r^{(D-2)} & (D \geq 3) \\ \Phi(r) &\sim +\ln r & (D = 2)\end{aligned}$$

so force falls off as $1/r^{(D-1)}$. Setting $D = 3$ reproduces the original module to **1.78e-15**, which is checked every run.

This makes a sharp, classical prediction. For a central force $F \sim r^{-n}$, circular orbits are linearly stable only when $n < 3$; in D dimensions $n = D - 1$, so stable bound orbits require $D < 4$. The effective potential $V_{\text{eff}} = L^2/2r^2 + \Phi(r)$ shows why: at $D = 4$ both terms scale as $1/r^2$ and there is no minimum at all, and for $D \geq 5$ the attraction wins at small r so the extremum is a *maximum*. This is Ehrenfest's 1917 argument that stable orbits single out three spatial dimensions.

Being a prediction, it was measured rather than asserted — 512 perturbed circular orbits per dimension:

D	force $\propto$	bound	stable annulus	absorbed	escaped	theory
2	r^{-1}	100.0%	100.0%	0.0%	0.0%	stable
3	r^{-2}	100.0%	100.0%	0.0%	0.0%	stable
4	r^{-3}	46.7%	2.7%	48.2%	5.1%	unstable
5	r^{-4}	0.0%	0.0%	47.3%	52.7%	unstable
6	r^{-5}	0.0%	0.0%	46.7%	53.3%	unstable
7	r^{-6}	0.0%	0.0%	47.7%	52.3%	unstable

The four-centre stream tells the same story. Bound fraction collapses from 75.3% ($D=2$) $\rightarrow$ 44.3% ($D=3$) $\rightarrow$ 2.3% ($D=4$) $\rightarrow$ 0.1% ($D=5$), while the absorbed fraction climbs to **99.8% by $D = 7$** : in high dimensions the short-range attraction $1/r^{(D-1)}$ overwhelms the centrifugal barrier $1/r^2$, so anything that comes near a centre is swallowed. The median r_{max} for $D \geq 5$ sits at 3.09, which is just the launch radius — those particles never went anywhere, they fell straight in.

These are extra **spatial** dimensions in a flat metric. They are not a second time dimension, not a curved manifold, and not compactified Kaluza-Klein circles. The result is a statement about Newtonian central forces, and it is a real one.

8. Where the model is weak

Stated bluntly, because a report that only lists successes is not useful.

8.1 The classification thresholds are heuristics, and they are marginal

`chaotic` is $\lambda > 0.02$ and `periodic` is a 6-D recurrence within 5%. Both are tunable constants, not derived quantities, and the 99,856-particle run shows both sitting inside populated regions of their own distributions rather than in empty gaps:

- The measured λ distribution has its **1st percentile at 0.0098 and its 5th at 0.0126** — the 0.02 chaos threshold cuts through live density, not through a trough. It lands between the low- λ peak and the main bulk, which is defensible, but moving it by a factor of two would reclassify thousands of particles.
- Every one of the 35 `periodic` particles has a recurrence between **0.0233 and 0.0499** against a **0.05** tolerance. The entire class is produced by where that cutoff sits.

The labels should not be trusted ahead of the raw `lyapunov` and `recurrence` values, which is why both are written to `particles.csv` for every particle. The histograms needed to retune them are in `diagnostics.png`.

Worse, λ **for an escaping particle is meaningless as a chaos measure**. Two nearly parallel escaping trajectories separate linearly, and the fit reports a small positive slope that has nothing to do with sensitive dependence. Escape is decided before λ is consulted, so no particle is mislabelled — but the reported λ values for escapers should be ignored, not interpreted.

8.2 The energy-drift metric is badly normalized

This one required correcting an initial diagnosis, so both the wrong answer and the right one are recorded.

The benchmark at $N = 262,144$ recorded a worst-case *relative* drift of **2.24e-04**, above the $1e-4$ threshold at which the README declares classifications untrustworthy. The obvious explanation is deep close encounters: particles passing very near a centre see enormous accelerations, and 5th-order DOPRI5 handles them less gracefully than the CPU's 8th-order DOP853. That was the initial diagnosis. **It is only half right.**

Testing it against the 99,856-particle run:

- Across the bulk, close encounters really do drive drift: `corr(log $\square$  relative drift, log $\square$  closest approach) = -0.622`. Closer approach, larger drift, as expected.
- But the **extreme tail is a normalization artifact**. Drift is reported as $|\Delta E| / |E|$, and the 481 particles with relative drift $> 1e-7$ have median $|E| = 0.00032$ — they are marginally bound, sitting almost exactly at $E = 0$. Their *absolute* energy error is **6.4e-11**, which is an order of magnitude **better** than the ensemble median of $1.0e-9$. They conserve energy beautifully; the denominator is what blows up.

The honest headline number is therefore the absolute one: **the worst absolute energy error anywhere in 99,856 particles is 1.51e-08**, with a median of $9.9e-10$.

So there are two separate defects. The real one is that `energy_drift` should be normalized by a *fixed* energy scale ($|\Phi(r)|$, or the initial kinetic energy) rather than by $|E|$, which is unstable precisely for the marginally-bound particles that matter most near the escape threshold. The secondary one is that a batched DOP853 would tighten the genuine close-encounter contribution. Until both are fixed, per-particle drift is written to `particles.csv` as a **quality flag** so the tail can be excluded rather than silently averaged in.

8.3 The inherited seeding correlates speed with direction

`seed_particle_stream` assigns particle `i` both the `i`-th direction on a Fibonacci spiral **and** the `i`-th speed of a linear sweep. Speed and direction are therefore perfectly correlated and their effects cannot be separated — fatal for any statistical readout like "escape fraction vs speed". The GPU driver defaults to a `grid` mode that takes the outer product instead. `stream` mode is retained for exact parity with the CPU reference.

8.4 The escape criterion does not measure escape

The single most consequential flaw found, and it only became visible at scale.

`escape` is triggered by crossing the sphere $|r| = R_{\text{escape}}$ (default 60). That is **not** the same as being gravitationally unbound. In the monopole far field $\Phi \approx -4/r$, a bound orbit with energy $E < 0$ has apoapsis $r_{\text{apo}} \approx -4/E$, so **any** bound particle with $|E| < 4/60 \approx 0.067$ reaches $r = 60$ on a very eccentric but perfectly closed orbit, gets stopped by the terminal event, and is recorded as having escaped.

Measured in the production run:

quantity	count
classified <code>escape</code>	22,120
genuinely unbound ($E > 0$)	18,960
bound, but labelled <code>escape</code>	3,160 (14.3% of all escapes)
unbound but not labelled escape	0

Those 3,160 particles have energies from -0.0655 to -0.0130, i.e. apoapsis distances from **61 to 308** — every one of them would have turned around and come back. The signature is visible directly in the escape-fraction curve: the transition to 100% escape sits at $|v_{\text{esc}}| \approx 1.5685$, while the true local escape speed is **1.6099**. The gap between those two numbers is exactly this artifact.

There are no false negatives, which is why the error is easy to miss — the criterion is one-sided. The fix is to require $E > 0$ for the `escape` label and treat the `R_escape` crossing as what it actually is: leaving the analysis volume. Until then, **the escape count in any run is an overestimate by roughly 14%**, and it grows as `R_escape` shrinks.

8.5 The ricochet is not horizon physics

Restated because it is the single most important caveat: specular reflection off a hard sphere is a **model choice**. There is no capture cross-section, no photon sphere, no innermost stable circular orbit, no frame dragging, no radiation reaction, no absorption. A real horizon absorbs; this one reflects perfectly. Every "bounce" statistic is a statistic about that modelling choice.

8.6 Other known limitations

- The singularities are **fixed**. They do not move, and they exert no force on each other — so this is a restricted problem, not a five-body problem.
- `t_max = 300` approximates $t \rightarrow \infty$. It is roughly 100–200 orbital times in the core region, which is adequate for classification but not for genuine asymptotic statements.
- The Lyapunov fit uses a coarse 256-point grid on the GPU path versus the full sample grid on the CPU path, so GPU λ values are slightly noisier.

9. On the relativistic questions

These questions were asked directly, and they deserve accurate answers rather than being quietly folded into simulation output that does not address them. **None of what follows is measured by this code.** It is established physics, marked where it becomes uncertain.

9.1 "How does time become space and space become time?"

This is real and specific. The Schwarzschild metric is

$$ds^2 = -(1 - r_s/r) c^2 dt^2 + (1 - r_s/r)^{-1} dr^2 + r^2 d\Omega^2$$

For $r > r_s$ the coefficient $(1 - r_s/r)$ is positive: dt is the timelike direction. For $r < r_s$ it turns **negative**, and the dr^2 coefficient turns negative with it. The roles swap — r becomes timelike and t becomes spacelike. The physical consequence is concrete: inside the horizon, moving toward $r = 0$ becomes as unavoidable as moving toward the future is outside. It is not that you *should not* escape; there is no future-directed path that increases r , in the same way there is no path that decreases your t .

9.2 "It never crosses the threshold"

Both halves of this are true, in different frames, and that is the resolution rather than a paradox:

- **Distant observer:** infalling matter is redshifted without limit and asymptotically freezes at the horizon. The crossing is never seen. This is the historical "frozen star" picture.
- **Infalling observer:** crosses in **finite proper time**, and nothing locally remarkable happens at r_s at all.

The key point is that the horizon is **not a curvature singularity**. The Kretschmann scalar $R_{abcd} R^{abcd} = 48G^2M^2/(c^4r^6)$ is perfectly finite at $r = r_s$ and diverges only at $r = 0$. The pathology at r_s in Schwarzschild coordinates is a *coordinate* artifact, removable by Eddington-Finkelstein, Kruskal-Szekeres or Painlevé-Gullstrand coordinates. For a supermassive black hole, tidal forces at the horizon ($\sim GM/r^3$, scaling as $1/M^2$) are small enough to cross unnoticed.

9.3 "Infinite curvature — does it bend all the way the other way?"

In **classical** GR, curvature genuinely diverges at $r = 0$; the singularity theorems (Penrose 1965, Hawking–Penrose 1970) show this is generic and not an artifact of symmetry. But a divergence is where the theory stops applying, not a physical infinity — the classical description fails once curvature reaches the Planck scale.

What happens instead is **unresolved**. Loop quantum cosmology proposes a bounce where the collapse reverses; various approaches propose regular ("non-singular") black holes; string-theoretic fuzzballs replace the interior entirely. **None of these is established.** Anyone who tells you the answer is known is overselling.

9.4 "When does the release happen?"

Hawking radiation: $T = \hbar c^3/(8\pi G M k_B)$, with evaporation time scaling as M^3 — about 10^{67} years for a solar-mass black hole, far longer than the current age of the universe. Whether information escapes with it is the **black hole information paradox**, still open.

9.5 "Space expanding faster than light"

True, and not a violation. Recession velocity follows $v = H_0 \cdot D$, so beyond the Hubble radius ($c/H_0 \approx 14.4 \text{ Gly}$) recession exceeds c . Special relativity's speed limit is **local** — it constrains relative velocities of objects at the same event, not the rate at which the space between distant objects grows. Objects beyond the Hubble sphere remain observable because the Hubble sphere itself expands.

Current observational state (checked August 2026):

- **Hubble tension, $\sim 5\sigma$ and not closing.** Local distance-ladder measurements give $\approx 73 \text{ km/s/Mpc}$; early-universe CMB/BAO inference gives $\approx 67\text{--}68$. JWST Cepheid observations (Riess et al. 2025) give 73.49 ± 0.93 , and a 2026 community distance network reports 73.5 ± 0.7 — a 1% local measurement. JWST reobserving the same Cepheids ruled out crowding and dust as the culprit, so this is no longer dismissible as an instrumental artifact.
- **DESI DR2 favours evolving dark energy.** BAO plus CMB plus supernovae prefer $w_0 > -1$ with $w_0 < 0$ at $2.8\text{--}4.2\sigma$ depending on which supernova compilation is used — i.e. dark energy that was *more* negative in the past and is weakening. That is below discovery threshold but has strengthened as data accumulated.

If that holds up, the "whole constant" is not constant, and the expansion history is not Λ CDM's.

9.6 "Does it flow backwards and forwards?"

This one the simulation can actually answer, and it was measured (§10) — because the Newtonian system here is exactly time-reversible.

10. Time reversibility, measured

The equations $\text{dr}/\text{dt} = \text{v}$, $\text{dv}/\text{dt} = \text{a}(\text{r})$ are invariant under $(\text{t} \rightarrow -\text{t}, \text{v} \rightarrow -\text{v})$: acceleration depends only on position, so reversing every velocity retraces the path exactly. Elastic specular reflection is its own inverse. There is no friction and no dissipation. **The microscopic laws here have no arrow of time.**

The experiment: integrate forward to T , negate every velocity, integrate forward another T , and measure how close particles come to their starting points. 256 particles, elastic horizons, $\text{rtol} = 1\text{e-}12$.

T	median error	p90 error	max error
5	1.52e-12	2.95e-12	7.53e-12
20	2.47e-11	1.34e-10	3.24e-09
40	2.90e-10	5.77e-09	7.38e-07
60	2.27e-09	2.77e-07	1.03e-04
80	1.92e-08	1.17e-05	7.81e-03
100	1.85e-07	3.60e-04	7.45e-01
140	1.25e-05	1.06e+00	6.28e+00
180	3.83e-04	4.21e+00	6.71e+00

At $\text{T} = 5$ the round trip returns every particle to within $1.5\text{e-}12$ — the system is reversible, as promised. By $\text{T} = 180$ the median particle misses its own starting point by $3.8\text{e-}04$ and the worst misses by **6.7 length units**, a distance comparable to the entire configuration.

The fitted growth is

$$\text{recovery error} \sim 2.05\text{e-}12 * \exp(0.0550 * 2\text{T})$$

The fitted amplification rate is 0.0550 per unit time. The independently measured median Lyapunov exponent of the 99,856-particle ensemble is 0.0547. Those two numbers come from completely different measurements — one from shadow-orbit divergence in the forward problem, the other from a forward-and-back round trip — and they agree to better than 1%. The loss of reversibility is chaotic amplification of round-off, and nothing else.

The **practical reversibility horizon** is $\text{T} \approx 182$: beyond it the past is arithmetically unrecoverable at double precision, even though it remains mathematically fully determined. Recovering a further factor of e in time costs a factor of $\text{e}^{0.11}$ in precision — buying reversibility back is exponentially expensive, which is why no amount of engineering makes this go away.

That is worth stating precisely, because it is easy to over-read. The arrow of time appearing here is **manufactured entirely by finite precision plus sensitive dependence**. It is a statement about numerics and classical chaos — *not* about thermodynamic entropy, CPT symmetry, or cosmological time asymmetry. Running with `--restitution 0.8` breaks the symmetry for a genuinely different reason: an inelastic horizon is really dissipative, and then the reversed run fails to return for reasons that have nothing to do with round-off.

11. The alternating quadrupole, and a mis-parameterized run

Every result up to this point used the **tetrahedron** geometry — four **+q** centres. The **planar** alternating configuration $q = (+1, -1, +1, -1)$, the only one with a true $1/r^3$ far field, had **never been run**. Running it exposed a seeding error worth recording in full, because the first run looked like a clean result and was not one.

First attempt, using the same beam as every previous run (speeds 0.55–1.85): 100% escape. All 59,902 particles. That is not a property of the geometry; it is a property of the beam. With alternating charges the monopole cancels, so

	tetrahedron	planar alternating
total charge	+4	0
$\Phi(r_{\blacksquare})$	-1.2955	-0.0984
local escape speed	1.6097	0.4437

The inherited beam therefore swept **1.2× to 4.2× the escape speed**. Every particle was unbound at launch. The run was a tautology.

Re-seeded across the real threshold (0.12–0.80), 59,902 particles to $t = 300$:

class	count	share
chaotic	28,852	48.2%
escape	27,393	45.7%
quasi-periodic	3,656	6.1%
periodic	1	0.002%

The striking difference is horizon traffic: **330,104 ricochets from 59,902 particles — 5.5 each, against 0.82 each for the tetrahedron**, nearly seven times higher. The mixed-sign configuration has repulsive centres that sling particles back into the attractive ones instead of letting them settle into wide orbits, and the lower escape speed keeps the survivors resident far longer. Energy conservation degrades correspondingly (median drift $5.4e-08$ against $1.7e-09$), which is the expected price of many more close encounters — consistent with the $corr = -0.622$ between drift and closest approach reported in §8.2.

The lesson is about method, not geometry: the beam was tuned once for the tetrahedron and then reused. Any seeding whose speed range is not derived from the *configuration's own* $escape_speed(r_{\blacksquare})$ will silently produce either 100% escape or 100% capture, and both look like findings.

12. Validating the asymptotic expansion

`verify.py`'s far-field test fits one number — the slope of $\log|\Phi|$ against $\log r$ — and asserts it is 3.000. That checks the leading *exponent* and no coefficient at all; it would pass identically if every term in the expansion were wrong.

The structure actually being claimed is an asymptotic series in $1/r$,

$$\Phi(r) = -G \sum_l M_l(n) / r^{l+1}, \quad M_l(n) = \sum_i q_i r_i^{l+1} P_l(n \cdot \hat{r}_i)$$

from the Legendre generating function. The testable consequence — the one Watson's lemma is about — is that **truncating after the leading term must leave a residual falling as $1/r^{l+2}$** , with the coefficient itself matching the analytic moment. `validate_multipole.py` checks this, and it found something the old test could not:

configuration	M_0	M_1	M_2	first correction	measured residual
planar alternating	0	0	$3a^2(n_x^2 - n_y^2)$	M_4 (hexadecapole)	r^{-6}
tetrahedron	4q	0	0	M_3 (octupole)	r^{-4}
collinear	4q	0	4.3182	M_2 (quadrupole)	r^{-3}

The tetrahedron has no quadrupole moment. Its four vertex vectors satisfy $\sum \mathbf{r}_i \otimes \mathbf{r}_i = (4/3)\mathbf{I}$, which is isotropic, so $\sum P_2(\mathbf{n} \cdot \mathbf{r}_i) = 0$ identically. It has no dipole either, so the leading correction to its monopole is the **octupole at $1/r$** — measured as a residual decaying at r^{-4} after monopole subtraction, with $M_0 = 2.1659$. That is a real property of the geometry this entire study has been built on, and nothing in the inherited test suite would ever have detected it.

The planar closed form $M_2 = 3a^2(n_x^2 - n_y^2)$ reproduces to **4.4e-16**, and all odd moments vanish to 1.1e-16 by inversion symmetry.

13. The relativistic module

Section 9 said that answering the horizon questions by simulation would need a separate program. That program now exists at [https://github.com/ricofour/four-singularity](#), and the claim in §9 — that this could not be bolted onto the Newtonian integrator — held: it shares no physics code at all.

13.1 What changed

	Newtonian module	Geodesic module
state	(r, v) in $\mathbb{R}^2$	(x^μ, u_μ) in $\mathbb{R}^4$
law	$dv/dt = -\nabla\Phi$, a force in flat space	$du^\mu/d\tau = -\Gamma^\mu_{\alpha\beta} u^\alpha u^\beta$; no force, curved geometry
conserved	$E = v^2/2 + \Phi$	$g_{\mu\nu} u^\mu u^\nu = -1$, plus $E = -u_t$, $L = u_\phi$
"horizon"	hard sphere with a restitution coefficient	a null surface of the geometry, imposed by nothing
time	a parameter	proper vs coordinate time — <i>the physics</i>

Only the numerics carry over: the same DOPRI5, per-particle adaptive stepping, and event root-finding.

Metrics implemented: Schwarzschild, Painlevé-Gullstrand and ingoing Eddington-Finkelstein (both horizon-regular), and Kerr in Boyer-Lindquist with frame dragging and an ergosphere. Christoffel symbols come two ways — closed-form for Schwarzschild, and a fully general numeric path from central differences of $g_{\mu\nu}$ — which agree to **2.97e-11**. That agreement is what makes the module extensible: any metric written as a function of position gets correct geodesics without hand-derived formulas.

13.2 The questions from §9, now answered by measurement

"It never crosses the threshold." The same infall, from rest at $r = 10M$, in two charts of the *same spacetime*:

- **Schwarzschild chart** (the distant static observer): $r_{\min} = 2.000000118$. It never crosses. Coordinate time diverges logarithmically — stopping successively closer gives $t = 47.17, 51.91, 56.53, 61.16, 65.76$ for $r - 2M = 1e-1 \dots 1e-5$, a fitted $dt/d[-\ln(r/2M-1)] = 2.0166$ against the predicted 2γ .
- **Painlevé-Gullstrand chart** (the infalling observer): crosses without incident and reaches $r = 9.98e-04$, in proper time **35.124058650** against the exact $\pi\sqrt{(r^3/8M)} = 35.124058748$ — relative error **2.79e-09**.

Both are correct. The horizon is a property of the *chart*, not of the spacetime, and the disagreement is the answer rather than a paradox.

"How does time become space?" Inside $r = 2M$, $g_{tt} = +1.000$ and $g_{rr} = -1.000$ at $r = M$: the signs have swapped, so r is timelike and t spacelike. The measured consequence: of 64 worldlines launched with a spread of angular momenta and **initial radial kicks in both directions**, 56 crossed the horizon, and every one of them had $dr/d\tau < 0$ **throughout the interior** (maximum $-8.96e-01$). Not one moved outward, including those fired outward. The algebra behind it: inside the horizon $(u^\mu)^2 \geq |f| > 0$, so u^μ can never reach zero and cannot change sign.

"Infinite curvature." The four-velocity norm holds to **8.66e-11** while crossing the horizon and stopping at $r_{\text{stop}} = 1.0$, and degrades to **2.62e-02** stopping at $r_{\text{stop}} = 1e-3$. The integration error is concentrated entirely at $r = 0$, not at $r = 2M$ — which is exactly the statement that the horizon is regular and the singularity is not. Where classical curvature diverges, the classical equations stop being integrable, and no integrator can fix that because it is not a numerical problem.

13.3 Your cubic

$P(y) = (2/3)M(2\cos(k/3) \pm \sqrt{3} \sin(k/3) + 1)$ is the trigonometric solution of the Schwarzschild turning-point cubic. With $u = 1/r$,

```
timelike: (du/dphi)^2 = (E^2-1)/L^2 + (2M/L^2) u - u^2 + 2M u^3
null:     (du/dphi)^2 = 1/b^2 - u^2 + 2M u^3
```

Depressing with $u = w + 1/(6M)$ gives roots $u = (1/6M)(1 + 2\cos((\theta-2\pi k)/3))$ — three roots 120° apart, which is where the $60^\circ/120^\circ$ reference comes from. (The factor is $2 \sin(2\pi/3) = \sqrt{3}$, not $\sqrt{2}$.) Implemented and verified in `orbit_cubic.py`:

- Given E and L , the cubic recovers $r_{apo} = 120$ and $r_{peri} = 80$ to $1e-9$.
- The roots in u sum to $1/(2M)$ to 12 digits, which **proves the third root always lies outside the horizon** ($u \blacksquare < 1/2M \blacksquare r \blacksquare > 2M$). An initial assertion that it lay inside was wrong, and the algebra corrected it.
- **The null double root is the photon sphere:** $b_{crit} = 3\sqrt{3} M = 5.196152422706632$ exactly, and integration measures the capture threshold at **5.19498** — relative error $2.26e-04$.
- ****The ISCO is a triple root**:** at $r = 6M$ all three coincide ($u = 1/6$ three times, summing to $1/2M$). A tidier characterisation than $d^2V_{eff}/dr^2 = 0$. Measured independently by perturbing circular orbits: $r = 5.90$ plunges to the horizon, $r = 6.10$ stays inside a band of 6.0915 – 6.1083 .

This also retires a caveat from §8.5. The Newtonian model has no capture cross-section because its horizon reflects; the geodesic module has one, and it is $27\pi M^2$ (the `b_crit` disc) — derived, not imposed.

13.4 Verification

Ten groups, all passing, each against a closed form that exists independently of this code — including perihelion precession at **3.0e-04** against the exact elliptic integral (the weak-field $6\pi M/p$ is genuinely 4.7% low at $p = 96$, so testing against *it* would have been testing the approximation), light deflection at $1.5e-02$ and $7.8e-03$ of $4M/b$, and the Newtonian limit recovering Kepler to $7.5e-04$ at $r = 2000M$. Full transcript in Appendix A.9.

What it still does not do. No radiation reaction, no self-force, no back-reaction on the metric, no matter fields, no quantum effects — so nothing here says anything about Hawking evaporation, the information paradox, or what replaces the singularity. Kerr is implemented but not yet exercised beyond the $a = 0$ consistency check, so frame dragging is available and unmeasured.

14. What to build next

Ordered by value per unit effort:

- 1 **Batched DOP853 on the GPU.** Directly addresses the energy-drift tail (§8.2), which is the main threat to the credibility of large-ensemble statistics.
 - 2 **Tune the classification thresholds against the measured distributions** rather than leaving them at inherited defaults — the raw λ and recurrence histograms are already written out for exactly this.
 - 3 **Absorbing horizons as the default** (`restitution = 0`) for any run meant to resemble a capture cross-section, with reflection kept as the energy-conservation diagnostic it is good at being.
 - 4 **Mobile singularities**, promoting this from a restricted problem to a genuine five-body problem.
 - 5 ~~If relativity is actually the goal, that is a different program.~~ **Done — see §13.**
integrates the geodesic equation in Schwarzschild, Painlevé-Gullstrand, Eddington-Finkelstein and Kerr, and all ten verification groups pass. Remaining work there: exercise Kerr (frame dragging, the ergosphere, the Penrose process), and add null-geodesic ray tracing to render what the four-centre configuration would actually *look like* if the centres were black holes rather than Newtonian wells.
 - 6 **Derive every beam from the configuration's own escape speed** rather than reusing a tuned constant — §11 shows how quietly that produces a 100% result that means nothing.
-

12. Reproducing every number in this document

```
python -m venv .venv
.\.venv\Scripts\Activate.ps1
pip install -r requirements.txt
pip install torch --index-url https://download.pytorch.org/whl/cu126
pip install reportlab

python verify.py                # 20 physics/numerics assertions
python validate_gpu.py          # GPU vs CPU, layered validation
python main_sim.py --outdir results_cpu
python bench.py                 # CPU vs GPU scaling
python gpu_sim.py --particles 100000 --t-max 300 --samples 3000 --track 800 \
    --outdir results_gpu_main
python hyper_dynamics.py        # dimensional scan
python reversibility.py         # time-reversal experiment
python validate_multipole.py    # asymptotic-expansion coefficients
python gpu_sim.py --config planar --particles 60000 --t-max 300 \
    --speed-min 0.12 --speed-max 0.80 --outdir results_planar_tuned
python make_report.py           # rebuild this PDF
```

The relativistic module lives in its own directory and uses the same venv:

```
cd ..\gr_geodesic_sim
python orbit_cubic.py          # turning-point cubic self-test
python verify_gr.py           # 10 groups against closed-form GR
```

`make_report.py --append "note"` timestamps a new entry into `RESEARCH_LOG.md` before rebuilding, so this document is meant to be appended to as the work continues.

Sources for §9.5: [DESI DR2 guide](#) · [On DESI's DR2 exclusion of \$\Lambda\$ CDM \(MNRAS Letters\)](#) · [Comparison of dynamical dark energy with \$\Lambda\$ CDM \(MNRAS\)](#) · [The Hubble Tension: A Decade Review](#) · [A Tale of Many \$H_0\$ \(Annual Reviews\)](#) · [The Hubble tension — CERN Courier](#)

Appendix A — Measured results (auto-generated)

Every number in this appendix is read directly from the run artifacts each time the report is rebuilt. Nothing here is typed by hand, so it cannot drift out of sync with what the code produced.

A.1 Physics and numerics verification (CPU)

23 passed, 0 failed. Full transcript of `verify.py`:

```
■
--- test_gradient ---
[PASS] acceleration == -grad Phi (max rel err 1.62e-10)

--- test_hessian ---
[PASS] tidal tensor == numerical Hessian (max rel err 4.31e-10)
[PASS] tidal tensor symmetric (asym 0.00e+00)

--- test_far_field_decay ---
[PASS] alternating config is charge-neutral
[PASS] far field ~ r^-3 (quadrupole) (fitted exponent 3.000)

--- test_rhs_shape ---
[PASS] rhs returns a length-6 vector ((6,))
[PASS] rhs velocity block matches state

--- test_kepler_limit ---
[PASS] Kepler orbit stays bound (ok)
[PASS] circular orbit closes after 10 periods (relative closure error 4.71e-07)
[PASS] radius stays constant (max dev 1.50e-08)
[PASS] energy conserved (Kepler) (relative drift 2.31e-11)

--- test_bound_orbit_energy ---
[PASS] four-center orbit integrates cleanly (ok)
[PASS] energy conserved (four-center) (relative drift 2.68e-09, bounces=0)
[PASS] particle never penetrates a horizon (min gap 5.08e-02)

--- test_ricochet ---
[PASS] head-on infall triggers a ricochet (5 bounce(s))
[PASS] elastic ricochet preserves speed (|v_in|=6.167651 |v_out|=6.167651)
[PASS] ricochet reverses the radial component
[PASS] energy conserved across ricochet (relative drift 3.47e-10)
[PASS] inelastic ricochet loses speed (6.1677 -> 3.0838)

--- test_symmetry ---
[PASS] mirror symmetry preserved (max deviation 0.00e+00)

--- test_stream_seeding ---
[PASS] stream speeds span the requested range ([0.500, 1.400])
[PASS] all directions inside the beam cone (min cos = 0.957197 vs 0.955336)
[PASS] launch points share the origin

=====
All checks passed.
```

A.2 GPU validation against the CPU reference

23 passed, 0 failed.

```
■device      : cuda
gpu          : a single consumer GPU
vram         : 24.0 GiB
torch       : 2.13.0+cu126 (cuda 12.6)

--- A field kernels ---
[PASS] GPU acceleration == numpy acceleration (rel 2.41e-16)
[PASS] GPU potential == numpy potential (rel 1.63e-16)
```

```

[PASS] GPU distances == numpy distances (rel 2.29e-16)
[PASS] GPU energy == numpy energy (rel 2.00e-16)

--- B Kepler limit ---
[PASS] Kepler run completes (ok)
[PASS] GPU circular orbit closes after 10 periods (relative closure error 4.80e-07)
[PASS] GPU radius stays constant (max dev 1.53e-08)
[PASS] GPU energy conserved (Kepler) (relative drift 1.81e-10)
[PASS] GPU trajectory == CPU trajectory (Kepler, 10 periods) (max abs deviation 1.64e-08)

--- C ricochet events ---
[PASS] GPU reproduces the CPU bounce count (GPU 5 vs CPU 5)
[PASS] GPU energy conserved across ricochets (relative drift 1.91e-09)
[PASS] GPU never penetrates a horizon (min gap 5.00e-14)
[PASS] GPU trajectory == CPU trajectory (5 ricochets) (max abs deviation 3.50e-06)
[PASS] inelastic horizon loses energy (relative energy change 1.941e+01)

--- DE chaotic orbit ---
[PASS] GPU == CPU on the chaotic orbit at t=20 (max 6-D deviation 1.91e-10)
      CPU/GPU separation grows at exp(+0.210 t); final |dY| = 1.078e+00
[PASS] long-time CPU/GPU divergence is exponential (chaos, not a bug) (fitted growth rate 0.210 per unit t)

--- FG batch invariance ---
[PASS] solo vs 64-batch agree to round-off over the first quarter (max abs deviation 0.00e+00 (whole run: 0.
[PASS] solo and batched bounce counts match (0 vs 0)
[PASS] GPU energy drift over the 64-particle batch (max 1.32e-09, median 4.14e-11)
[PASS] no particle penetrates a horizon (min gap 5.00e-14)

--- H ensemble outcomes ---
[PASS] escape flags agree particle-by-particle (GPU 0 escapes, CPU 0)
[PASS] ricochet counts agree particle-by-particle (GPU total 2, CPU total 2)
[PASS] escape/end times agree (max |dt_end| 0.00e+00)
      (this batch: GPU 16.94s vs CPU 5.59s -- 32 particles is far below the GPU's break-even point)

=====
All GPU validation checks passed.

```

A.3 CPU reference run

particles	16
integrator	scipy solve_ivp DOP853 (serial)
t_max	300.0
classifications	{'chaotic': 10, 'quasi-periodic': 2, 'escape': 4}
total ricochets	7
max energy drift	5.958e-09
median energy drift	3.060e-09
runtime (s)	30.29

A.4 GPU production run

particles	99,856
device	a single consumer GPU (not CPU)
integrator	batched Dormand-Prince 5(4), per-particle adaptive
t_max / samples	300.0 / 3000
tolerances	rtol=1e-10, atol=1e-12
runtime (s)	1108.1

throughput	90 particles/s
classifications	{'chaotic': 64631, 'quasi-periodic': 13070, 'escape': 22120, 'periodic': 35}
total ricochets	81,614
total RK steps	655,885,314
rejected steps	10,904,816
GPU loop iterations	15,600
max energy drift	3.104e-06
p99 energy drift	1.354e-08
median energy drift	1.724e-09
min horizon gap	4.990e-14
closest approach	0.0500
proximity-cap bindings	0
Lyapunov median	0.0547
Lyapunov 5th–95th pct	0.0126 ... 0.2028

A.4.1 Escape fraction vs launch speed

speed range	N	escape fraction
0.550 – 0.604	4,424	0.0%
0.604 – 0.658	4,424	0.0%
0.658 – 0.713	3,792	0.0%
0.713 – 0.767	4,424	0.0%
0.767 – 0.821	3,792	0.0%
0.821 – 0.875	4,424	0.0%
0.875 – 0.929	3,792	0.0%
0.929 – 0.983	4,424	0.0%
0.983 – 1.038	3,792	0.0%
1.038 – 1.092	4,424	0.0%
1.092 – 1.146	3,792	0.0%
1.146 – 1.200	4,424	0.0%
1.200 – 1.254	4,424	0.0%
1.254 – 1.308	3,792	0.0%
1.308 – 1.363	4,424	0.0%
1.363 – 1.417	3,792	0.0%
1.417 – 1.471	4,424	0.0%
1.471 – 1.525	3,792	0.0%
1.525 – 1.579	4,424	28.6%
1.579 – 1.633	3,792	100.0%
1.633 – 1.688	4,424	100.0%
1.688 – 1.742	3,792	100.0%
1.742 – 1.796	4,424	100.0%
1.796 – 1.850	4,424	100.0%

A.5 CPU vs GPU throughput

CPU reference: **5.05 particles/s** measured over 24 particles (scipy solve_ivp DOP853, one particle per call).
Larger counts extrapolate linearly because the loop is serial — that is an extrapolation, not a measurement.

N	wall s	iters	ms/iter	particles/s	speedup	max drift	VRAM MB
1,024	32.9	3,216	10.2	31	6×	2.5e-07	1
4,096	59.0	3,328	17.7	69	14×	2.0e-05	6
16,384	119.7	3,568	33.6	137	27×	1.9e-06	24
65,536	171.1	3,568	48.0	383	76×	8.2e-06	99
262,144	206.4	4,144	49.8	1,270	251×	2.2e-04	379

A.6 Dimensional scan

Consistency check — D=3 reproduces the 3-D field module: **PASS**.

A.6.1 Circular-orbit stability (Ehrenfest test)

D	force $\propto$	v_circ	bound	stable annulus	absorbed	escaped	theory
2	r^{-1}	1.0000	100.0%	100.0%	0.0%	0.0%	stable
3	r^{-2}	0.7071	100.0%	100.0%	0.0%	0.0%	stable
4	r^{-3}	0.7071	46.7%	2.7%	48.2%	5.1%	unstable
5	r^{-4}	0.6124	0.0%	0.0%	47.3%	52.7%	unstable
6	r^{-5}	0.5000	0.0%	0.0%	46.7%	53.3%	unstable
7	r^{-6}	0.3953	0.0%	0.0%	47.7%	52.3%	unstable

A.6.2 Four-center stream vs dimension

D	N	bound	escaped	absorbed	median r_max
2	4,096	75.3%	0.0%	24.7%	3.41
3	4,096	44.3%	21.0%	34.7%	6.60
4	4,096	2.3%	44.3%	53.4%	5.28
5	4,096	0.1%	10.7%	89.2%	3.09
6	4,096	0.0%	1.8%	98.2%	3.09
7	4,096	0.0%	0.2%	99.8%	3.09

A.8 Asymptotic (multipole) expansion

19 passed, 0 failed. Coefficient-level validation of the far field, not just the leading exponent.

```
■[PASS] planar: total charge (monopole M_0) vanishes (0.00e+00)
[PASS] planar: dipole M_1 vanishes in every direction (max |M_1| 1.11e-16)
[PASS] planar: M_2 == 3 a^2 (n_x^2 - n_y^2) [closed form] (max dev 4.44e-16)
[PASS] planar: all odd moments vanish (inversion symmetry) (max |M_odd| 1.11e-16)

--- planar ---
[PASS] planar: leading moment is l = 2
[PASS] planar: residual after truncation falls as r^-5
[PASS] planar: recovered leading coefficient matches M_2 (max rel err 1.96e-05)
[PASS] tetrahedron: monopole M_0 == 4q (4.0)
[PASS] tetrahedron: dipole M_1 vanishes (vertices sum to zero) (max |M_1| 1.11e-16)
[PASS] tetrahedron: quadrupole M_2 ALSO vanishes (isotropic 2nd moment) (max |M_2| 9.44e-16)
[PASS] tetrahedron: octupole M_3 is the first nonzero correction (max |M_3| 2.1659)
```

```

--- tetrahedron ---
[PASS] tetrahedron: leading moment is  $l = 0$ 
[PASS] tetrahedron: residual after truncation falls as  $r^{-4}$ 
[PASS] tetrahedron: recovered leading coefficient matches  $M_0$  (max rel err 2.01e-08)
[PASS] collinear: monopole  $M_0 == 4q$ 
[PASS] collinear: quadrupole  $M_2$  does NOT vanish (anisotropic by design) (max  $|M_2|$  4.3182)

--- collinear ---
[PASS] collinear: leading moment is  $l = 0$ 
[PASS] collinear: residual after truncation falls as  $r^{-3}$ 
[PASS] collinear: recovered leading coefficient matches  $M_0$  (max rel err 1.20e-05)

=====
All multipole / asymptotic-expansion checks passed.

```

A.9 Geodesic (general-relativistic) module

32 passed, 0 failed. Every check is against a closed-form GR result that exists independently of this code.

```

device : cuda
gpu     : a single consumer GPU
units   :  $G = c = 1$ ,  $M = 1$ , so  $r_s = 2M = 2$ 

--- A Christoffel symbols ---
[PASS] analytic == numeric Christoffels (Schwarzschild) (max rel err 2.97e-11)
[PASS] Christoffels symmetric in lower indices (max asym 0.00e+00)
[PASS] Kerr(a=0) metric == Schwarzschild (max |dg| 8.88e-16)
[PASS] Kerr(a=0) Christoffels == Schwarzschild (max rel err 2.97e-11)
[PASS]  $g g^{-1} == I$  (max dev 1.11e-16)
[PASS] Schwarzschild horizon at  $r = 2M$  ( $r_h = 2.0$ )
[PASS] Kerr(a=0.9) horizon at  $M + \sqrt{M^2 - a^2}$  ( $r_h = 1.435890$ )

--- B/C norm and Killing charges ---
[PASS] norm starts at -1 (max  $|N+1|$  2.22e-16)
[PASS] norm conserved over  $\tau = 4000$  (max drift 5.47e-14)
[PASS] energy  $E = -u_t$  conserved (max drift 1.61e-13)
[PASS] angular momentum  $L = u_\phi$  conserved (max drift 1.28e-11)

--- D radial infall proper time ---
[PASS] radial infall reaches the singularity (status 2)
[PASS] proper time ==  $\pi \sqrt{r_0^3/8M}$  [exact] (measured 35.124058650 vs exact 35.124058748, rel 2.79e-09)
norm drift vs where the integration is stopped:
  r_stop = 1      (inside the horizon)  norm drift = 8.662e-11
  r_stop = 0.1    (inside the horizon)  norm drift = 2.559e-08
  r_stop = 0.01   (inside the horizon)  norm drift = 2.554e-05
  r_stop = 0.001  (inside the horizon)  norm drift = 2.622e-02
[PASS] norm held to -1 while crossing the horizon ( $r_{\text{stop}} = 1.0 < 2M$ ) (drift 8.66e-11)
[PASS] norm error is concentrated at the curvature singularity, not the horizon (grows 8.7e-11 -> 2.6e-02 as  $r \rightarrow 0$ )

--- E the horizon in two charts ---
proper time to reach  $r=2M$  (exact): 33.700870
Schwarzschild chart :  $r_{\text{min}} = 2.000000118$ , coordinate  $t = 74.620$ 
Painleve chart      :  $r_{\text{min}} = 9.979e-04$ 
[PASS] Schwarzschild chart never crosses  $r = 2M$  ( $r_{\text{min}} = 2.000000118 > 2$ )
coordinate time as the horizon is approached:
  stop at  $r = 2M + 0.1$     ->  $t = 47.1669$ 
  stop at  $r = 2M + 0.01$    ->  $t = 51.9106$ 
  stop at  $r = 2M + 0.001$   ->  $t = 56.5291$ 
  stop at  $r = 2M + 0.0001$  ->  $t = 61.1577$ 
  stop at  $r = 2M + 1e-05$   ->  $t = 65.7598$ 
[PASS] Schwarzschild coordinate time diverges logarithmically at  $r = 2M$  (fitted  $dt/d[-\ln(r/2M-1)] = 2.0166$  v)
[PASS] Painleve chart crosses and reaches the singularity ( $r_{\text{min}} = 9.979e-04$ )
[PASS] Eddington-Finkelstein agrees with the exact proper time (rel err 1.28e-09)

--- F  $r \leftrightarrow t$  swap inside the horizon ---
[PASS] worldlines actually entered the horizon (56 of 64 crossed  $r = 2M$ )
[PASS]  $dr/d\tau < 0$  for every worldline inside  $r = 2M$  (max  $dr/d\tau$  inside = -8.961e-01 over 56 worldlines)

```

```

[PASS] g_tt > 0 and g_rr < 0 inside the horizon (roles swapped) (at r=1.0: g_tt = +1.000, g_rr = -1.000)

--- G photon capture cross-section ---
[PASS] photon capture threshold == 3 sqrt(3) M (measured 5.19498 vs exact 5.19615, rel 2.26e-04)
[PASS] photons start null (g_uv u^u u^v = 0) (max |N| 2.31e-16)

--- H ISCO at r = 6M ---
r = 4.00 r_min = 4.0000 r_max = 202.0501 -> UNSTABLE
r = 5.00 r_min = 2.0000 r_max = 10.0000 -> UNSTABLE
r = 5.50 r_min = 2.0000 r_max = 7.3333 -> UNSTABLE
r = 5.90 r_min = 2.0000 r_max = 6.2107 -> UNSTABLE
r = 6.10 r_min = 6.0915 r_max = 6.1083 -> stable
r = 7.00 r_min = 6.9963 r_max = 7.0037 -> stable
r = 9.00 r_min = 8.9962 r_max = 9.0038 -> stable
r = 15.00 r_min = 14.9933 r_max = 15.0067 -> stable
[PASS] no stable circular orbit inside r = 6M (ISCO) (0 of 4 sub-ISCO orbits stayed bounded)
[PASS] circular orbits stable outside r = 6M (4 of 4 supra-ISCO orbits stayed bounded)

--- I perihelion precession ---
[PASS] perihelion precession == exact elliptic-integral value (measured 0.206146752 vs exact 0.206084609 rad
(weak-field 6piM/p = 0.196349541 is 4.7% low at p = 96 -- the approximation, not the solver)
[PASS] norm conserved through the precessing orbit (drift 6.25e-14)
(same formula for Mercury: 43.0 arcsec/century -- observed 43)

--- J light deflection ---
[PASS] light deflection == 4M/b at b = 200 (measured 2.030192e-02 rad vs 4M/b = 2.000000e-02, rel 1.51e-02)
[PASS] light deflection == 4M/b at b = 400 (measured 1.007828e-02 rad vs 4M/b = 1.000000e-02, rel 7.83e-03)

--- K Newtonian limit ---
[PASS] circular orbit at r = 2000M closes after one Kepler period (swept 6.287903 rad vs 2pi, rel 7.51e-04)
[PASS] radius constant at r = 2000M (max dev 0.00e+00)

=====
All GR verification checks passed.

```

A.7 Time-reversal experiment

256 particles, restitution = 1.0 (elastic, so reversible in exact arithmetic), $\text{rtol} = 1\text{e-}12$. Each row integrates forward to T , negates every velocity, integrates forward another T , and measures the distance back to the starting state in 6-D.

T	recovered	median error	p90 error	max error	within 1e-6
5	256	1.52e-12	2.95e-12	7.53e-12	100%
10	256	3.47e-12	6.74e-12	1.13e-11	100%
20	256	2.46e-11	1.34e-10	3.24e-09	100%
40	256	2.90e-10	5.77e-09	7.38e-07	100%
60	256	2.27e-09	2.77e-07	1.03e-04	95%
80	256	1.92e-08	1.17e-05	7.81e-03	80%
100	256	1.85e-07	3.60e-04	7.45e-01	60%
140	256	1.24e-05	1.06e+00	6.28e+00	41%
180	256	3.83e-04	4.21e+00	6.71e+00	31%

Fitted growth: recovery error $\approx 2.05\text{e-}12 \cdot \exp(0.0550 \cdot 2T)$. The fitted amplification rate is directly comparable to the ensemble Lyapunov exponent — if they match, the loss of reversibility is chaotic amplification of round-off and nothing else. Practical reversibility horizon (median error reaches $1\text{e-}3$): $T \approx 182$.