
Particle Streams in a Four-Singularity Potential Well

A Newtonian GPU study: verification, numerical method, a dimensional scan, and measured time reversibility

Thomas Garren

SOPHIA XT LLC

Run of 2026-08-16. Retypeset 2026-09-19. · rico_inf_curve, research log rev. 2

sophiaxt.com

CC BY 4.0. Code referenced herein is Apache 2.0.

Abstract

This report records an executable study of high-density particle streams moving in the field of four fixed singularities. The solver is Newtonian and nothing in it is relativistic. It integrates the Hamiltonian $H = |p|^2/2m + m\Phi(r)$ in flat space with a softened four-centre potential, in normalised units with $G = 1$ and test mass $m = 1$. There is no metric tensor, no proper time, no light cone, no frame dragging, no radiation reaction and no event horizon in any relativistic sense. What the code calls a horizon is a hard sphere of radius $R = 0.05$ off which the particle reflects specularly. That reflection is a numerical regularisation chosen to keep the ordinary differential equation integrable near the centres. It is a modelling decision and not a physical result, and it means the trajectory figures in this report say nothing about black holes even though they look as though they might. Within that Newtonian system the results are measured rather than asserted. A batched Dormand-Prince 5(4) integrator in float64 on one consumer GPU advanced 99,856 particles to $t = 300$ in 1,108 s, taking 655,885,314 Runge-Kutta steps of which 1.66% were rejected by the error controller, at $rtol = 1e-10$ and $atol = 1e-12$. The report covers verification of the inherited physics against closed forms, validation of the GPU path against a serial SciPy DOP853 reference, measured throughput scaling, a scan over spatial dimension that reproduces Ehrenfest's stability argument, coefficient-level validation of the multipole expansion, and a time-reversal experiment whose measured amplification rate agrees with the ensemble Lyapunov exponent to better than one percent. A separate program, sharing no physics code with the four-centre solver, integrates geodesics in Schwarzschild, Painlevé-Gullstrand, Eddington-Finkelstein and Kerr geometries; its results are reported here as a distinct study and they do not transfer to the Newtonian ensemble. Every number in this document is simulated or derived. Nothing was measured in a laboratory.

1. What this study is, and the boundary that defines it

The object of study is a stream of non-interacting test particles moving in the static field of four fixed point sources. Particles do not act on each other and do not act on the sources, so the ensemble is a collection of independent single-particle problems that happen to share a field. That is what makes the problem worth putting on a GPU, and it is also what limits what the results mean.

Before any result, the boundary, because everything downstream depends on it. This simulation is Newtonian. It is not general relativity. The solver advances a position and a velocity in flat three-dimensional space under a gradient of a scalar potential. There is no metric, no line element, no causal structure, no proper time and no curvature. The four centres are called singularities because the unsoftened potential diverges at them, not because anything in the model resembles a spacetime singularity.

The consequence is specific rather than rhetorical. Questions about horizons, about time and space exchanging roles, about capture cross-sections or about redshift are questions about a metric, and this solver has no metric. Section 12 describes a separate program that does integrate geodesics and can address some of them.

1.1 What was inherited, and its status

A scaffold existed at the start of this work: `potential_fields.py`, `particle_dynamics.py`, `main_sim.py` and `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. That turned out to be a generous posture: the CPU code was correct on first run, and section 4 records the checks that established it.

2. The model and the equations of motion

In normalised units with $G = 1$ and test mass $m = 1$, the softened Plummer-regularised four-centre potential per unit mass and the equations of motion are

```
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)
```

The softening length ε removes the $1/r$ divergence so the ordinary differential equation stays integrable. The constructor enforces $\varepsilon \leq R_{\min}/2$ so the regularisation never reaches into the sampled region. In the production run $\varepsilon = 0.001$, more than an order of magnitude below the minimum sampled distance of 0.0500, so the Plummer term never influenced the

$$\begin{aligned} n &= (r - r_i) / |r - r_i| \\ v' &= v - (1 + e)(v \cdot n) n \end{aligned} \quad \text{applied only when } v \cdot n < 0$$

The restitution coefficient $e = 1$ is elastic and conserves energy exactly. That choice is what makes energy drift a clean diagnostic of integrator quality rather than a mixture of physics and numerics. Running with $e = 0.8$ breaks the symmetry for a genuinely different reason, because an inelastic surface really is dissipative.

Three geometries ship. The tetrahedron places four sources of charge $+q$ at the vertices of a regular tetrahedron and has a monopole far field. The planar configuration places four sources of alternating sign at the corners of a square, cancels the monopole and gives a true quadrupole with a $1/r^3$ far field. The collinear configuration places four sources on the x-axis and serves as an exact-symmetry control case.

3. Numerical method

Two integrators are used and the report is careful about which produced which number. The CPU reference is `scipy.integrate.solve_ivp` with the DOP853 method, one particle per call, serial. The production path is a batched Dormand-Prince 5(4) written in PyTorch in which every particle carries its own adaptive step size and its own event

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

One root-finder therefore serves all five surfaces, and which surface fired is decided afterwards from which term attained the minimum. The root is bracketed across the accepted step and refined with Illinois regula falsi in which every trial evaluation is a real fifth-order Runge-Kutta step of size θh rather than an interpolant, so the located crossing carries the integrator's own order of accuracy instead of a lower-order dense output polynomial's. The step size h is additionally clamped so a step never overshoots the next sample time. Particles therefore land exactly on output grid points and recording needs no interpolation, and because at most one grid point can fall inside a step, an event resolved at $\theta < 1$ can never skip a sample.

3.2 Three defects the validation suite caught

- Conflated step-clamping semantics. A single 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, so time ran ahead of the physics. The symptom was the GPU logging 3 ricochets where the CPU logged 5. Fixed by separating the shortened flag from the grid-landing flag.

dynamics. That was the design intent and section 7 confirms it from the closest-approach histogram rather than assuming it.

The physical surface $|r - r_i| = R_i$ is handled discretely. A terminal root-finding event stops the integration, the velocity is reflected specularly, and integration restarts just outside:

state. The ensemble advances in lockstep at the level of kernel launches rather than at the level of simulated time, so one particle can take steps of $1e-6$ near a surface while another coasts at $1e-1$ in the far field within the same kernel.

Tolerances for the production run were $\text{rtol} = 1e-10$ and $\text{atol} = 1e-12$, $t_{\max} = 300$ with 3,000 output samples. The time-reversal experiment of section 10 used $\text{rtol} = 1e-12$. All arithmetic is float64. The shadow-orbit separation used for the Lyapunov measurement starts at $1e-8$ against coordinates of order 1, and float32 carries roughly $1e-7$ relative precision, which would put the entire Lyapunov measurement below its own noise floor. FP64 runs at 1/64 the FP32 rate on a consumer card, but this kernel is launch-bound and bandwidth-bound rather than FLOP-bound, so the practical penalty is far smaller than that ratio implies.

3.1 Event handling

The discontinuous ricochet is what makes the batched integrator non-trivial. Four horizon surfaces and one escape sphere are folded into a single continuous scalar, positive in the physical region and negative once anything has been crossed:

- A Zeno trap in the proximity cap. The safety cap limiting step size near a surface was written as h proportional to the gap, which makes the gap shrink geometrically so the surface is never reached. Replaced with a constant cap $h = 0.5 R / |v|$, which still makes jumping clean over a surface geometrically impossible but terminates.
- Retired particles were re-flushed every iteration, overwriting each particle's true end time with $t_{\max}$. Fixed by clearing a finished particle's accepted-step mask immediately so its diagnostics freeze at the instant it finished.

One divergence from the CPU reference is deliberate. SciPy's event machinery sees sign changes only at step endpoints, so a step that dives through a surface and back out in one bite is invisible to it. The GPU path's proximity cap makes that geometrically impossible. That is an improvement rather than a match, so the number of times the cap binds is counted and reported. In the production run it bound zero times, so it changed nothing there. It is insurance, not an active correction.

4. Verification of the inherited physics

`verify.py` runs nine groups of checks against analytic results that exist independently of this code. Twenty-three assertions

passed on first execution with no tolerance loosened. The narrative section of the source document says twenty; the auto-generated transcript says twenty-three, and the transcript

is the figure used here because it is read from the run rather than typed.

Table 1. Verification of the CPU reference against closed forms and numerical derivatives. All values are computed, not measured.

Check	Result
$a \equiv -\nabla\Phi$ against central differences	max relative error 1.62e-10
Tidal tensor against a numerical Hessian	max relative error 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 at $t = 120$	2.68e-09
Elastic ricochet speed	6.167651 to 6.167651, below 1e-9
Mirror symmetry, collinear geometry	0.00e+00, bit-exact

Two caveats are recorded rather than glossed. 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. Separately, the bound-orbit test recorded zero bounces and a minimum surface gap of 5.08e-02, so its assertion that the orbit never penetrates a horizon is vacuous. Ricochet is genuinely exercised only by the dedicated ricochet test, which produced five bounces. That is a real gap in the inherited test suite.

5. Validating the GPU path against the CPU reference

The four-centre problem is chaotic, with measured Lyapunov exponents in the range 0.1 to 0.3 for the orbits used in validation. 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 rather than a defect. Demanding long-time trajectory agreement in the chaotic regime would be demanding the wrong thing, and a suite that passed such a test would be lying. The checks are therefore layered by what is legitimately comparable.

Table 2. Layered GPU-against-CPU validation. Twenty-three assertions, none failing.

Layer	What it tests	Result
A	Field kernels against numpy	3.2e-16 max relative
B	Kepler limit, 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 $t = 20$	six-dimensional 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 at $t = 40$	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 and GPU separation grows as $\exp(+0.210 t)$, which is the system's own Lyapunov rate for that orbit. The disagreement has the signature of chaos rather than of a systematic error, and that is a stronger statement than agreement would have been. Layer G was originally an assertion of bit-identity and failed at 2.22e-12. The assertion was wrong, not the code: different batch sizes select different reduction kernels, so a

sum can differ in its last bit, and one unit in the last place amplified by $\exp(0.21 t)$ over $t = 60$ is 2e-12.

6. Measured throughput

Both sides do identical work: same field, same tolerances, same $t_{\max}$, same initial conditions. The CPU rate of 5.05 particles per second was measured over 24 particles and larger counts extrapolate linearly, which is exactly right for a serial per-particle loop and is stated as an extrapolation rather than a measurement.

Table 3. Measured GPU scaling on one consumer card against the serial SciPy reference. Speedup is a throughput ratio at the stated N, not a universal figure.

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

Going from 1k to 262k particles is 256 times the work and costs only 4.9 times more per iteration. At small N the kernels are trivial and the loop is bound by launch latency and dispatch, with the GPU roughly 38% utilised, so extra particles in that regime are nearly free. The plateau between 65k and 262k is the onset of memory-bandwidth saturation. The honest reading is that 251× is a throughput ratio at N = 262,144: at N = 32 the GPU is slower than SciPy, 18.5 s against 5.6 s, and break-even sits near N = 150 to 200.

7. The production ensemble

The stream is launched from (3.0, 0.4, 0.6) into a cone of half-angle 0.30 rad about the direction $(-1, -0.13, -0.2)$, with speeds swept from 0.55 to 1.85. The local escape speed at the launch point is 1.6097, so the ensemble deliberately straddles it. 99,856 particles were integrated to $t = 300$ in 1,108 s across 15,600 GPU loop iterations totalling 655,885,314 Runge-Kutta steps, of which 10,904,816 or 1.66% were rejected by the error controller.




Figure 1. Trajectories in (x, y, z) for 800 tracked particles drawn from the 99,856-particle tetrahedron ensemble, coloured by orbit class and with the four singularity positions marked. Axes are position in normalised length units. Escaping particles leave as near-straight rays; chaotic and quasi-periodic particles remain inside the core region. Recovered from the source PDF; it is matplotlib output from the run described in this section.

Ensemble diagnostics: tetrahedron, N = 99,856, $t_{\text{max}} = 300.0$




Figure 2. Six ensemble diagnostics from the same run. Top row: escape fraction against launch speed with the local escape speed marked; the finite-time Lyapunov spectrum with the 0.02 classification threshold and the 0.0547 median marked; the distribution of $\log_{10}$ relative energy drift against the $1e-4$ trust threshold. Bottom row: orbit class against launch speed; $\log_{10}$ closest approach to any centre against the horizon radius 0.05 and the softening length 0.001; ricochets per particle on a log count axis, 81,614 in total. Redrawn as vector art from the run's own diagnostics.npz by scripts/papers/make_rico_diagnostics.py.

Table 4. Classification of the tetrahedron production ensemble. The class labels are produced by threshold heuristics, not by derived criteria; section 13 gives the thresholds and shows where they are marginal.

Orbit 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 is the dominant population: particles fall into the four-centre region, undergo repeated close encounters and wander with strongly positive Lyapunov exponents, which is what four attracting centres should produce. The measured spectrum runs from $\lambda = 0.0058$ to 0.512 with median 0.0547 and a fifth to ninety-fifth percentile band of 0.0126 to 0.2028 . Escaping particles leave on paths that cross the analysis boundary, but 14.3% of them are not actually unbound, which section 13 treats as a defect in the criterion rather than a property of the system. Quasi-periodic orbits are wide bounded rosettes with low λ , concentrated just below the escape threshold. The periodic class holds 35 particles; the 16-particle CPU run found none, so the class only appears at scale.

Two further readouts are worth recording. The closest-approach distribution has a sharp spike exactly at the surface

radius $R = 0.05$, which is the particles that bounced, with a spread out to $r \approx 1$ and nothing below it. Since the softening length is 0.001 , the Plummer regularisation never reached the sampled region. Horizon encounters total 81,614, distributed exponentially from 0 to a maximum of 12 per particle.

8. The alternating quadrupole, and a mis-parameterised run

Every result to this point used the tetrahedron. The planar alternating configuration with charges $(+1, -1, +1, -1)$, the only geometry 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 attempt looked like a clean result and was not one. Reusing the tetrahedron beam, speeds 0.55 to 1.85 , produced 100% escape across all 59,902 particles. That is a property of the beam rather than of the geometry.

Table 5. Why the inherited beam was a tautology for the planar geometry: it swept 1.2 to 4.2 times the local escape speed, so every particle was unbound at launch.

Quantity	Tetrahedron	Planar alternating
Total charge		+4
$\Phi(r_0)$ at the launch point	-1.2955	-0.0984
Local escape speed		1.6097
		0.4437

Re-seeded across the real threshold at speeds 0.12 to 0.80 , 59,902 particles were integrated to $t = 300$ with the same tolerances. 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 relative drift $5.4e-08$ against $1.7e-09$, which is the expected price of many more close encounters and is consistent with the correlation of -0.622 between drift and closest approach reported in section 13.

Table 6. Classification of the re-seeded planar quadrupole ensemble, 59,902 particles to $t = 300$.

Orbit class	Count	Share
chaotic	28,852	48.2%
escape	27,393	45.7%
quasi-periodic	3,656	6.1%
periodic	1	0.002%




Figure 3. Trajectories from the re-seeded planar quadrupole run, same axes and same colouring as Figure 1. The source document described this run but did not illustrate it; the plot is matplotlib output taken from that run's own output directory.

Ensemble diagnostics: planar_quadrupole, N = 59,902, t_max = 300.0




Figure 4. Ensemble diagnostics for the re-seeded planar run, panels as in Figure 2. The escape-fraction panel now shows the transition sitting on the planar geometry's own escape speed of 0.444 rather than on a reused constant, the energy-drift distribution is visibly broader and shifted toward larger drift, and the ricochet histogram extends to about 30 encounters per particle against 12 for the tetrahedron. Redrawn as vector art from that run's own diagnostics.npz.

The lesson is about method rather than 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 at the launch point will silently produce either total escape or total capture, and both look like findings.

9. The dimensional scan

Generalising to D dimensions cannot mean keeping a $1/r$ potential and adding coordinates, which would be the three-dimensional force law wearing extra axes. The generalisation 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 on every run. This makes a sharp classical prediction. For a central force $F \sim r^{-n}$, circular orbits are linearly stable only when $n < 3$, and 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, with 512 perturbed circular orbits per dimension.

Table 7. Circular-orbit stability against spatial dimension, 512 perturbed orbits per row. The measured transition sits between $D = 3$ and $D = 4$ where Ehrenfest's argument puts it.

D	force	v_{circ}	bound	stable annulus	absorbed	escaped	theory
2	r^{-1}	1.0000	100.0%	100.0%	100.0%	0.0%	0.0% stable
3	r^{-2}	0.7071	100.0%	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

The four-centre stream tells the same story. The bound fraction collapses from 75.3% at $D = 2$ to 44.3% at $D = 3$, 2.3% at $D = 4$ and 0.1% at $D = 5$, 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 maximum radius 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.

10. Time reversibility, measured

The equations $dr/dt = v$ and $dv/dt = a(r)$ are invariant under $t \rightarrow -t$ with $v \rightarrow -v$, because the acceleration depends only on position. Elastic specular reflection is its own inverse. There is no friction and no dissipation, so the microscopic laws here have no arrow of time. The experiment integrates forward to T , negates every velocity, integrates forward another T , and measures how close each particle comes to its starting state in six dimensions. 256 particles, elastic surfaces, $\text{rtol} = 1\text{e-}12$.

Table 8. Round-trip recovery error against half-period T . Errors are distances in the six-dimensional state space in normalised units.

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%

At $T = 5$ the round trip returns every particle to within $1.5\text{e-}12$, so the system is reversible as promised. By $T = 180$ 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

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

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 and 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, where the median error reaches $1\text{e-}3$, is $T \approx 182$. Beyond it the past is arithmetically unrecoverable at double precision even though

it remains mathematically fully determined, and recovering a further factor of e in time costs a factor of $e^{0.11}$ in precision.

The arrow of time appearing here is manufactured entirely by finite precision plus sensitive dependence. It is a statement about numerics and classical chaos. It is not a statement about thermodynamic entropy, CPT symmetry or cosmological time asymmetry.

11. Coefficient-level validation of the multipole expansion

The inherited 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 $1/r$ from the Legendre generating function:

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

The testable consequence 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. Nineteen checks of that form pass, and they found something the old test could not.

Table 9. Leading moments and the first surviving correction for each geometry, measured as the decay rate of the residual after subtracting the leading term.

Configuration	M_0	M_2	First correction	Measured residual
planar alternating		$0 \ 3a^2(n_x^2 - n_y^2)$	M_4 , hexadecapole	r^{-5}
tetrahedron	$4q$		$0 \ M_3$, octupole	r^{-4}
collinear	$4q$		$4.3182 \ M_2$, quadrupole	r^{-3}

The tetrahedron has no quadrupole moment. Its four vertex vectors satisfy a second moment proportional to the identity, which is isotropic, so the $l = 2$ Legendre sum vanishes identically. It has no dipole either, so the leading correction to its monopole is the octupole at $1/r^4$, measured as a residual decaying at r^{-4} after monopole subtraction, with $M_3 = 2.1659$. That is a real property of the geometry this entire study is built on, and nothing in the inherited test suite would have detected it. The planar closed form reproduces to $4.4e-16$ and all odd moments vanish to $1.1e-16$ by inversion symmetry.

12. The separate geodesic program, and why it does not transfer

A second program integrates the geodesic equation and shares no physics code with the four-centre solver. Only the numerics carry over: the same Dormand-Prince 5(4), the same per-particle adaptive stepping, the same event root-finding. It is reported here because the two programs are often confused, and the confusion is exactly the error this document exists to prevent.

Table 10. The two programs are different physics with shared numerics. Nothing in the right column applies to any result in sections 4 through 11.

	Newtonian four-centre module	Geodesic module
State	(r, v) in R^6	(x^μ, u^μ) in R^8
Law	$dv/dt = -\nabla\Phi$, a force in flat space	$du/d\tau = -\Gamma u u$, no force, curved geometry
Conserved	$E = v^2/2 + \Phi$	$g_{\mu\nu}u^\mu u^\nu = -1$, $E = -u_t$, $L = u_\phi$
Horizon	hard sphere with a restitution coefficient	a null surface of the geometry, imposed by nothing
Time	a parameter	proper against coordinate time, which is the physics

Metrics implemented are Schwarzschild, Painlevé-Gullstrand and ingoing Eddington-Finkelstein, both of the latter 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 the metric, and they 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.

Thirty-two checks pass across ten groups, each against a closed form that exists independently of this code.

Table 11. Geodesic module against closed-form general relativity, in units $G = c = 1$ and $M = 1$ so that $r_s = 2M = 2$. The weak-field precession formula $6\pi M/p$ is itself 4.7% low at $p = 96$, so the comparison is against the exact elliptic integral.

Check	Measured	Exact	Relative error
Radial infall proper time from $r_0 = 10M$	35.124058650	35.124058748	$2.79e-09$
Schwarzschild coordinate-time divergence rate	2.0166	$2M = 2.0$	$8.3e-03$
Photon capture impact parameter	5.19498	$3\sqrt{3} M = 5.196152$	$2.26e-04$
Perihelion precession, rad/orbit	0.206146752	0.206084609	$3.02e-04$
Light deflection at $b = 200$	$2.030192e-02$	$4M/b = 2.0e-02$	$1.51e-02$
Light deflection at $b = 400$	$1.007828e-02$	$4M/b = 1.0e-02$	$7.83e-03$
Newtonian limit, one Kepler period at $r = 2000M$	6.287903 rad	2π	$7.51e-04$

Two results bear directly on the Newtonian model's weakest assumption. The same infall from rest at $r_0 = 10M$, computed in two charts of the same spacetime, gives $r_{\min} = 2.000000118$ in the Schwarzschild chart, which never crosses, and $r_{\min} =$

$9.979e-04$ in the Painlevé-Gullstrand chart, which crosses without incident; the horizon is a property of the chart rather than of the spacetime. And this module has a capture cross-section of $27\pi M^2$, derived from the critical impact parameter

rather than imposed. The Newtonian model has none at all, because its surface reflects.

What the geodesic module still does not do: no radiation reaction, no self-force, no back-reaction on the metric, no matter fields, no quantum effects. Kerr is implemented but exercised only as far as the $a = 0$ consistency check, so frame dragging is available and unmeasured.

13. Where the Newtonian model is weak

13.1 The classification thresholds are marginal heuristics

The chaotic label is $\lambda > 0.02$ and the periodic label is a six-dimensional recurrence within 5%. Both are tunable constants rather than 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 first percentile at 0.0098 and its fifth at 0.0126, so the 0.02 threshold cuts through live density. 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, so the entire class is produced by where that cutoff sits; tighten the tolerance to 0.02 and the class vanishes. They are best read as the most nearly recurrent orbits found rather than as proven periodic orbits. Separately, λ for an escaping particle is meaningless as a chaos measure, because two nearly parallel escaping trajectories separate linearly. Escape is decided before λ is consulted so no particle is mislabelled,

but the reported λ values for escapers should be ignored rather than interpreted.

13.2 The energy-drift metric is badly normalised

The benchmark at $N = 262,144$ recorded a worst-case relative drift of $2.24\text{e-}04$, above the $1\text{e-}4$ threshold at which the project declares classifications untrustworthy. The obvious explanation is deep close encounters, and it is half right. Across the bulk, close encounters really do drive drift: the correlation between log relative drift and log closest approach is -0.622 . But the extreme tail is a normalisation artifact. Drift is reported as $|\Delta E| / |E_0|$, and the 481 particles with relative drift above $1\text{e-}7$ have median $|E_0| = 0.00032$, sitting almost exactly at $E = 0$. Their absolute energy error is $6.4\text{e-}11$, an order of magnitude better than the ensemble median of $1.0\text{e-}9$. They conserve energy well; 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 $1.51\text{e-}08$, with a median of $9.9\text{e-}10$.

13.3 The escape criterion does not measure escape

This is the most consequential flaw found and it only became visible at scale. The escape label 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 $E < 0$ has apoapsis $r \approx -4/E$, so any bound particle with $|E| < 4/60 \approx 0.067$ reaches $r = 60$ on a very eccentric but perfectly closed orbit, is stopped by the terminal event, and is recorded as having escaped.

Table 12. The escape label is one-sided: it has no false negatives, which is why the error is easy to miss.

Quantity	Count
Classified escape	22,120
Genuinely unbound, $E > 0$	18,960
Bound, but labelled escape	3,160, or 14.3% of all escapes
Unbound but not labelled escape	0

Those 3,160 particles have energies from -0.0655 to -0.0130 , giving apoapsis distances from 61 to 308, so every one of them would have turned around and come back. The signature is visible directly in the escape-fraction curve of Figure 2: the transition to 100% escape sits at launch speed 1.5685 while the true local escape speed is 1.6097, and the gap between those two numbers is exactly this artifact. The fix is to require $E > 0$ for the escape label and to treat the boundary 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 the overestimate grows as R_{escape} shrinks.

13.4 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 and no absorption. A real horizon absorbs; this one reflects perfectly. Every bounce statistic in this report, including the 81,614 encounters of Figure 2 and the 330,104 of Figure 4, is

a statistic about that modelling choice rather than about gravity.

13.5 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 rather than a five-body problem.
- $t_{\text{max}} = 300$ stands in for $t \rightarrow \infty$. It is roughly 100 to 200 orbital times in the core region, adequate for classification but not for genuine asymptotic statements.
- The Lyapunov fit uses a coarse 256-point grid on the GPU path against the full sample grid on the CPU path, so GPU λ values are slightly noisier.
- The seeding routine inherited from the scaffold gives particle i both the i -th direction on a Fibonacci spiral and the i -th speed of a linear sweep, so speed and direction are perfectly correlated and cannot be separated in any statistical readout. The GPU driver defaults to a grid mode that takes the outer product instead; the correlated mode is retained only for parity with the CPU reference.

14. What these results establish

- The inherited field and dynamics code is correct. Gradients match central differences to $1.6\text{e-}10$, the tidal tensor matches a numerical Hessian to $4.3\text{e-}10$ with exactly zero asymmetry, and the Kepler limit closes to $4.7\text{e-}07$ over ten periods.
- A batched adaptive Runge-Kutta integrator with per-particle event handling reproduces a serial reference implementation everywhere the comparison is legitimate, and where it stops reproducing it, the divergence is measured to grow at the system's own Lyapunov rate rather than at some other rate.
- Throughput on one consumer GPU reaches 1,270 particles per second at $N = 262,144$ against a serial reference of 5.05, with break-even near $N = 150$ to 200 and a measured onset of bandwidth saturation between 65k and 262k particles.
- The four-centre Newtonian system is a chaotic scattering system. In the tetrahedron geometry, 64.7% of a 99,856-particle ensemble is chaotic by the stated threshold, with a finite-time Lyapunov spectrum of median 0.0547.
- Ehrenfest's argument that stable bound orbits require fewer than four spatial dimensions is reproduced by measurement: 100% stable annuli at $D = 2$ and $D = 3$, 2.7% at $D = 4$, and zero at $D \geq 5$, over 512 perturbed circular orbits per dimension.
- The loss of numerical time reversibility in this system is chaotic amplification of round-off and nothing else. The independently fitted amplification rate of 0.0550 and the independently measured median Lyapunov exponent of 0.0547 agree to better than 1%.
- The tetrahedron geometry has no dipole and no quadrupole moment; its leading correction is the octupole with $M_3 = 2.1659$ and a residual decaying as r^{-4} .
- The geodesic program reproduces seven independent closed-form results of general relativity, the tightest being radial infall proper time to $2.8\text{e-}09$.

15. What these results do not establish, and honest scope

The four-centre study is Newtonian. It integrates $dv/dt = -\nabla\Phi$ in flat space with a Plummer-softened potential and a specular reflecting sphere at $r = R$. It therefore says nothing whatever about the following, and no amount of resemblance between its trajectory plots and relativistic imagery changes that:

- Anything about an event horizon. The surface at $R = 0.05$ reflects; a horizon absorbs. There is no capture cross-section in this model, no photon sphere, no innermost stable circular orbit and no one-way causal boundary.
- Anything about time. There is one universal parameter t . There is no proper time, no coordinate time, no redshift, no

time dilation and no exchange of timelike and spacelike roles.

- Anything about curvature. The space is flat by construction. The word singularity here denotes a divergence of a Newtonian scalar potential, not a curvature singularity, and the softening removes even that.
- Anything about gravitational radiation, self-force, back-reaction or mass loss. The four sources are fixed, infinitely rigid and eternal.
- Anything about real astrophysical systems. Test particles do not interact, the sources do not move, and there is no gas, no magnetic field, no relativistic correction and no observational comparison.
- The extra dimensions of section 9 are additional flat spatial coordinates. They are not a second time dimension, not a curved manifold and not compactified extra dimensions of the kind that appear in higher-dimensional gravity.
- The arrow of time in section 10 is numerical. It is finite precision amplified by sensitive dependence. It is not thermodynamic entropy, not CPT violation and not cosmological time asymmetry.

Within the Newtonian system, several reported quantities carry known defects that are not yet fixed. The escape count is an overestimate by roughly 14% because the criterion measures a boundary crossing rather than unboundedness. The orbit-class labels depend on two threshold constants that sit inside populated regions of their own distributions, so the periodic class in particular is an artifact of where a tolerance was placed. The relative energy-drift metric is unstable for marginally bound particles, which are precisely the particles that matter most near the escape threshold. Long-time statements rest on $t_{\text{max}} = 300$, which is 100 to 200 core orbital times and not an asymptotic limit.

The geodesic results of section 12 are a separate study. They are correct general relativity for test particles and photons on a fixed background, and they do not license any relativistic reading of sections 4 through 11. Within their own scope they exclude radiation reaction, self-force, back-reaction, matter fields and every quantum effect, so they say nothing about Hawking evaporation, the information paradox or what replaces a classical singularity. The Kerr metric is implemented and has been exercised only at $a = 0$, so frame dragging and the ergosphere are available and unmeasured.

Finally, provenance. Every number in this document is simulated or derived. No apparatus exists, nothing was measured in a laboratory, and no comparison against observation has been made. The hardware is described as one consumer GPU because the specific card is not load-bearing for any result; the numerical method, the tolerances and the step counts are, and they are stated.

16. Reproducing every number in this document

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

python verify.py          # 23 physics and numerics assertions
python validate_gpu.py    # layered GPU-against-CPU validation
python main_sim.py        # CPU reference run, 16 particles
python bench.py           # CPU against GPU scaling
python gpu_sim.py --particles 100000 --t-max 300 --samples 3000 --track 800
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

# the geodesic module is a separate program in its own directory
python orbit_cubic.py     # turning-point cubic self-test
python verify_gr.py       # 10 groups against closed-form GR
```

Per-particle Lyapunov exponents, recurrence values and energy drifts are written to particles.csv for every particle in every run, so the classification thresholds discussed in section 13 can be retuned against the raw distributions rather than trusted as labels.

References

- P. Ehrenfest, Proceedings of the Amsterdam Academy 20, 200 (1917). The dimensional stability argument reproduced in section 9.
- J. R. Dormand and P. J. Prince, Journal of Computational and Applied Mathematics 6, 19 (1980). The 5(4) pair used on the GPU path.
- E. Hairer, S. P. Nørsett and G. Wanner, Solving Ordinary Differential Equations I, 2nd ed., Springer, 1993. The DOP853 method used by the SciPy reference.
- H. C. Plummer, Monthly Notices of the Royal Astronomical Society 71, 460 (1911). The softening used to regularise the four-centre potential.
- C. W. Misner, K. S. Thorne and J. A. Wheeler, Gravitation, Freeman, 1973. Closed forms against which the geodesic module of section 12 is checked.