Cosmological Simulations in Pathological Geometries

And How To Make Them

Balázs Pál, Gábor Rácz, István Csabai, István Szapudi

GPU Days 2026

ELTE Department of Physics of Complex Systems
HUN-REN Wigner Research Centre for Physics

The basics of $\mathbb{T}^3$ cosmological simulations

Why simulate the Universe at all?

We only have one Universe.

All galaxies, clusters, and filaments we observe are one realisation of a stochastic process: the formation of large-scale structure from initial conditions. We see one Hubble constant, one CMB, one expansion history. This irreducible uncertainty is **cosmic variance**.

And the non-linear regime where halos form, merge, and the cosmic web is built is analytically intractable. Perturbation theory does not get us there.

Simulations let us:

- Generate ensembles of possible universes
- Bridge linear theory and non-linear observations
- Build mock catalogues for survey design
- Test alternative cosmologies

Gravity on discretised matter

Strip away the details and every cosmological simulation does one thing: track discrete chunks of matter under gravity. The comoving equation of motion has *three independent pieces*.

$$m_i\,\ddot{\boldsymbol{x}}_i \;=\; \underbrace{\sum_{j\neq i} \frac{m_i\,m_j\,\boldsymbol{F}(\boldsymbol{x}_i-\boldsymbol{x}_j,\;h_i+h_j)}{a(t)^3}}_{\text{law of interaction}} \;\underbrace{-\,2\,m_i\,H(t)\,\dot{\boldsymbol{x}}_i}_{\text{background expansion}} \;+\; \underbrace{\text{boundary term}}_{\text{spatial topology}}$$

Law of interaction
What gravity *is*. Newtonian $1/r^2$, MOND, $f(R)$, screened fifth forces, plus your softening kernel and force solver.

Background expansion
How the Universe expands. $H(t)$ from Friedmann or any tabulated $a(t)$ (ΛCDM, wCDM, early dark energy, ...).

Spatial topology
Geometry of the manifold. The *boundary condition* of your simulation: $\mathbb{T}^3$, $\mathbb{R}^3$, $\mathbb{S}^1\times\mathbb{R}^2$. We will return to this.

Structure formation in $\mathbb{T}^3$

From nearly homogeneous initial conditions, gravity sculpts the cosmic web of voids, filaments, walls, and dark matter halos. A standard periodic cube, $L = 500$ Mpc, $N = 256^3$ particles.

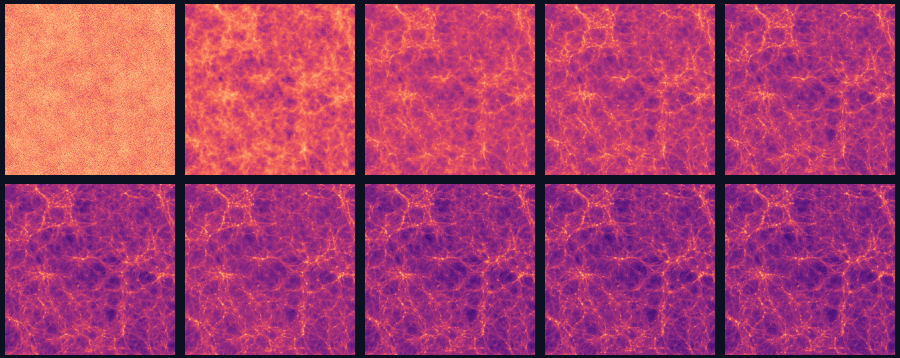

Figure 1: $z = 63$ (top-left) $\rightarrow$ $z = 0$ (bottom-right).

The problems of $\mathbb{T}^3$ topology

What the periodic cube quietly breaks

When you make the box periodic, you also make a physical claim. The simulation lives on a manifold whose properties are *not* the properties of our Universe. [See Rácz et al., 2019].

Broken rot. invariance

The cube has only the 48-element octahedral group $\mathrm{O_h}$, not $\mathrm{SO}(3)$. Angular momentum is *not conserved*. The force depends on orientation relative to the box axes.

Modified force law

Ewald summation over periodic images alters gravity at scales near L_{box}. The force deviates from $1/r^2$ with direction-dependent corrections. *Such simulations are not even Newtonian anymore!*

No observational support

Planck and SDSS find no signs of periodicity in the CMB or in galaxy surveys. This means that $\mathbb{T}^3$ is not an observed physical property of the Universe, but rather a computational choice!

Small boxes ($L < 50$ Mpc) amplify all of these: missing large-scale power, halos seeing their own periodic images, invalidated sub-box physics.

But the cube is also extraordinary

In fairness, the cube has been the standard for fifty years for excellent reasons.

- FFT-based force solvers (PM, TreePM) are extremely efficient on periodic domains.
- Translational invariance ensures linear momentum is exactly conserved.
- Statistical homogeneity is automatic.
- Decades of validated infrastructure: GADGET, AREPO, PKDGRAV, ...
- Enormous scientific output: Millennium, FLAMINGO, the whole canon.

The question is not whether $\mathbb{T}^3$ is useful

The usefulness of $\mathbb{T}^3$ simulations cannot be questioned. The question is whether their hidden topological assumptions matter for the physics we want to study.

Resolution: $\mathbb{R}^3$ and $\mathbb{S}^1 \times \mathbb{R}^2$

What would an ideal geometry look like?

1. **Isotropic boundaries** — no preferred directions; the box looks the same from every direction.
2. **Topology matching observations** — either fully open $\mathbb{R}^3$, or a single periodic direction $\mathbb{S}^1 \times \mathbb{R}^2$.
3. **Angular momentum conservation** — full $\mathrm{SO}(3)$, or at minimum $\mathrm{SO}(2)$ about a physically motivated axis.
4. **Physically motivated force law** — pure Newtonian $1/r^2$ in open geometries, or the natural Green's function of the manifold. No spurious periodic-image contributions.
5. **Natural zoom-in capability** — high resolution where you need it, smooth degradation outward.

The StePS framework achieves all five by construction.

[See Rácz et al., 2018, 2019]

Stereographic compactification

Instead of tiling space periodically, map $\mathbb{R}^3$ onto a finite sphere via inverse stereographic projection. *One extra term in the EOM — Newton's shell theorem — and the topology is fixed.*

The stereographic map

$$\omega = 2 \arctan\left(\frac{r}{2R_{\mathrm{S}}}\right)$$

compactifies $r \in [0, \infty) \;\rightarrow\; \omega \in [0, \pi)$. The outer shell becomes an **isotropic boundary** via the shell theorem applied to homogeneous matter outside R_{sim}.

Equal angular steps $\Delta\omega$ on the sphere $\mathcal{S}$ become exponentially growing radial steps Δr on the plane $\mathcal{P}$.

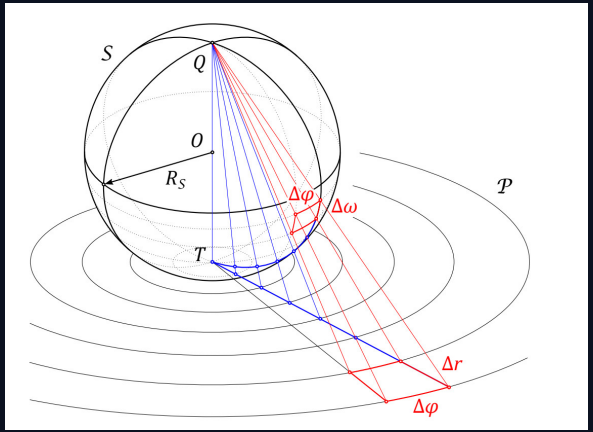

The first cosmological simulation in $\mathbb{S}^1 \times \mathbb{R}^2$ topology

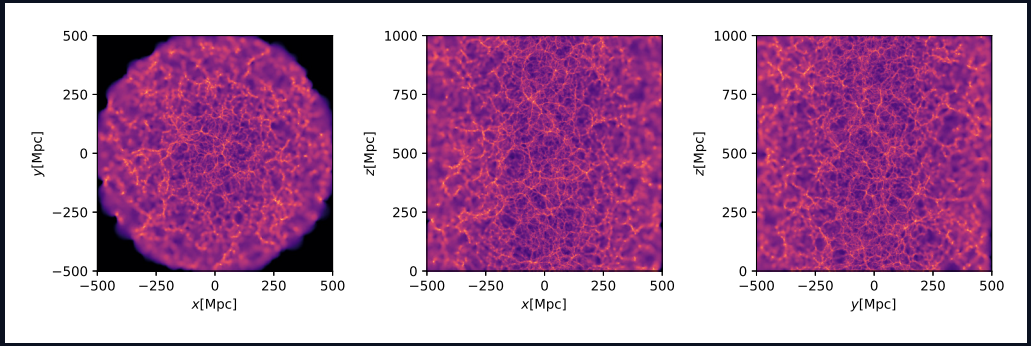

Figure 2: Dark matter density in a cylindrical $\mathbb{S}^1 \times \mathbb{R}^2$ simulation. 2.4×10^7 particles, $R_{\rm sim} = 500$ Mpc, $L_z = 1$ Gpc, Planck 2018 ΛCDM, 640 CPU cores. The non-linear cosmic web, i.e., voids, filaments, halos form, exactly as they do in a cube.

[See Rácz et al. 2026 and Pál et al. 2026]

A topology test, and a null result

Shamir (2025) reported a galaxy-spin asymmetry in JWST/JADES data. Our cylindrical gravity is anisotropic by construction (weaker axially, stronger radially). Does the topology fake the signal?

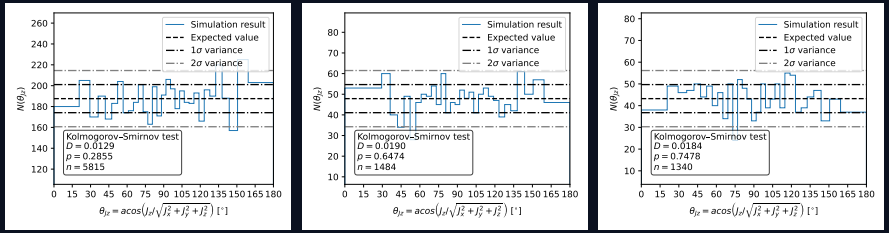

Figure 3: Halo spin orientation θ_{J_z} relative to the periodic axis, three mass bins at $z = 0.5$.

Result: null

The cylinder topology does *not* fake the JWST asymmetry: one class of explanations is ruled out.

The takeaway

Did the $\mathbb{T}^3$ topology become dogmatism?

The most dangerous conventions are the ones that have grown so familiar they have stopped looking like conventions.

1. From convenience to assumption

$\mathbb{T}^3$ was chosen because FFTs are fast on a cube. Fifty years later it had stopped looking like a choice and now it just became the norm *how the Universe is simulated*.

2. Every field has its own $\mathbb{T}^3$

In ML it is the canonical benchmark. In molecular dynamics it is *the* force field. In climate it is the parametrisation that ages into orthodoxy. Always: a tool too useful to question.

3. Our first job: keep asking

Which of our assumptions are *physics*, and which are merely *convenience*? Keep the distinction alive long after it has stopped being obvious.

Final takeaway

The standard periodic cube is an extraordinary tool. But it is not the only geometry the Universe permits us to simulate.

To build on the work of giants who came before is to inherit not just their answers, but their questions too. Take the time to learn why every assumption was made and never be afraid to play with them, to probe their limits, to ask whether they still deserve to stand.

Cautious and rigorous scientific curiosity is the only thing that keeps every field honest.

Thank you for your attention!

Backup slides

The structure of scientific research programmes

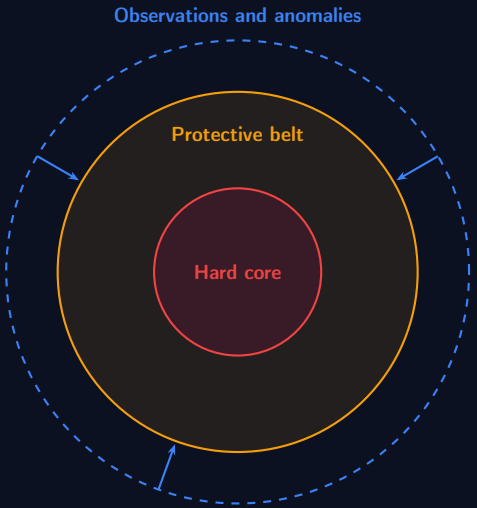

Kuhn (1962, 1970): Science operates within a *disciplinary matrix*, which is a shared framework of exemplars, methods, and commitments. Day-to-day work is *puzzle-solving* within the paradigm. Anomalies accumulate until a *crisis* triggers a paradigm shift.

Popper (1970): Warns against *methodological dogmatism*. A paradigm that immunizes itself against refutation risks ceasing to be scientific. Healthy science requires that auxiliary hypotheses remain genuinely testable.

Lakatos (1970): A *research programme* has an irrefutable *hard core* shielded by a *protective belt* of adjustable auxiliary hypotheses. *Progressive* if belt modifications predict novel facts; *degenerative* if they only accommodate post hoc.

Where does $\mathbb{T}^3$ live in the Lakatosian structure?

Hard core

- General Relativity
- Cold dark matter
- Cosmological constant Λ
- Gaussian, adiabatic initial conditions

Foundational commitments. Questioning them requires overwhelming anomalies.

Protective belt

- Specific inflation model
- Dark energy equation of state
- Baryonic feedback physics
- Neutrino mass hierarchy
- Reionization history
- **Simulation boundary conditions**

Adjustable choices that shield the core from direct falsification.

The periodic cube is not part of the hard core of cosmology. It is an auxiliary hypothesis in the protective belt — a numerical convenience, not a physical law. Replacing it with $\mathbb{R}^3$ or $\mathbb{S}^1 \times \mathbb{R}^2$ is the kind of belt modification Lakatos considers *progressive*: it yields testable predictions (angular momentum statistics, topological anisotropy) that $\mathbb{T}^3$ cannot make.

Conversely, refusing to test alternative geometries is what Popper calls methodological dogmatism.

From primordial fluctuations to the cosmic web

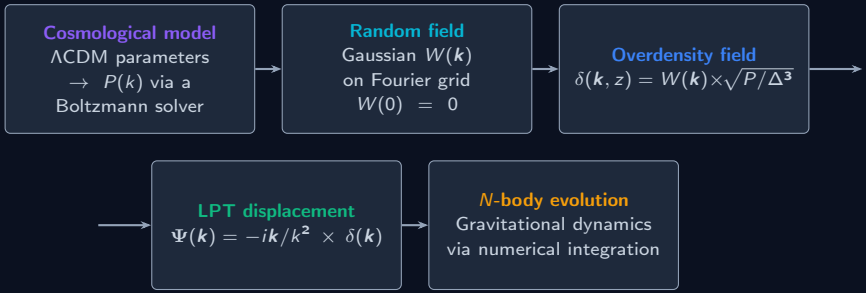

Primordial linear power spectrum

$$P(k, z) = A_{\mathrm{s}} \cdot \left(\frac{k}{k_{\mathrm{p}}}\right)^{n_{\mathrm{s}}-1} \cdot T^2(k) \cdot D_1^2(z)\,,$$

where $T(k)$ is the transfer function and $D_1(z)$ is the linear growth factor.

Lagrangian perturbation theory

LPT displaces particles from Lagrangian positions $\boldsymbol{q}$ to Eulerian positions $\boldsymbol{x}$:

$$\boldsymbol{x}(\boldsymbol{q}, z) \;=\; \boldsymbol{q} \;+\; D_1(z)\,\boldsymbol{\Psi}^{(1)}(\boldsymbol{q}) \;+\; D_2(z)\,\boldsymbol{\Psi}^{(2)}(\boldsymbol{q}) \;+\; \cdots$$

1st order: Zel'dovich approximation

$$\boldsymbol{\Psi}^{(1)}(\boldsymbol{k}) = -\frac{i\boldsymbol{k}}{k^2}\,\delta(\boldsymbol{k})$$

Solves $\nabla^2\Phi = \delta$, then $\boldsymbol{\Psi} = -\nabla\Phi$ in Fourier space. Particles stream along straight trajectories until shell crossing.

2nd order: tidal field correction

$$S(\boldsymbol{q}) = \sum_{i<j}\left[\Phi_{,ii}\,\Phi_{,jj} - (\Phi_{,ij})^2\right]$$

$$\boldsymbol{\Psi}^{(2)}(\boldsymbol{k}) = -\frac{i\boldsymbol{k}}{k^2}\,S(\boldsymbol{k})$$

The quadratic tidal tensor source captures anisotropic collapse, i.e., the initial skeleton of filaments and walls.

Conservation laws by topology

	Lin. mom.	Ang. mom.	Symmetry	Periodic correction
$\mathbb{T}^3$ (periodic cube)	✓ all 3	×	$\mathrm{O_h}$	3D Ewald
$\mathbb{R}^3$ (spherical StePS)	×	✓ full $\mathrm{SO}(3)$	$\mathrm{SO}(3)$	none (shell thm.)
$\mathbb{S}^1\times\mathbb{R}^2$ (cylindrical)	✓ z-axis	✓ z-comp.	$\mathrm{SO}(2)\times$transl.	1D Ewald

$\mathbb{S}^1\times\mathbb{R}^2$ is a natural hybrid: periodic along one physically motivated axis (e.g., a filament spine) plus isotropic boundaries and angular momentum conservation in the transverse plane. Ideal for filamentary structures, anisotropic cosmologies (Bianchi VII), and topology tests.

Computational advantage of the zoom-in geometry

The radially decreasing resolution keeps particle counts at $\sim 10^6$–10^7 even for Gpc-scale volumes, making $\mathcal{O}(N^2)$ direct summation on GPUs feasible. The StePS Millennium-equivalent: 1.2×10^7 particles, 1.35 Gpc^3, 0.34 GB memory, 106 h on 12 GPUs.

The stereographic map visually

Points on the hypersphere $\mathcal{S}$ are projected from the pole Q through the tangent plane $\mathcal{P}$ at the antipode T. Equal angular steps $\Delta\omega$ on $\mathcal{S}$ become exponentially growing radial steps Δr on $\mathcal{P}$.

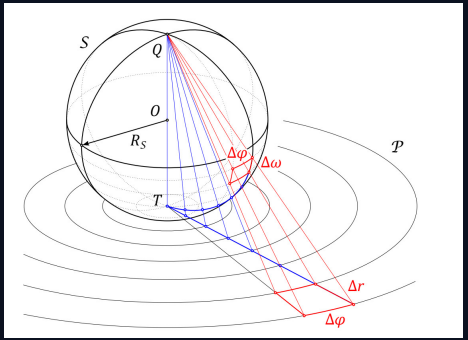

Figure 4: 3D: $\mathcal{S}^2 \to \mathcal{P}$. Uniform shells $\to$ concentric rings of growing Δr.

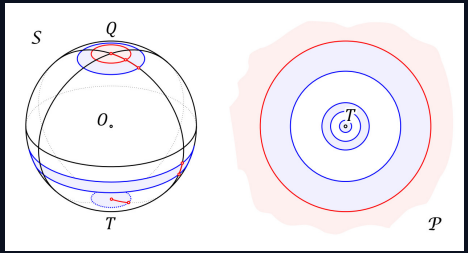

Figure 5: Circles on the sphere map to circles on the plane. Area near Q inflates; area near T is nearly isometric.

`stepsic`: initial conditions for non-standard geometries

The `stepsic` code extends the standard IC pipeline to $\mathbb{R}^3$ and $\mathbb{S}^1 \times \mathbb{R}^2$ geometries, as well as non-cubical $\mathbb{T}^3$ boxes:

Step	Standard (periodic)	StePS geometries
Particle load	Regular grid or periodic glass, uniform mass	Concentric shells with variable mass resolution, or pre-relaxed glass. $m_i = \bar{\varrho}\, V_{\mathrm{shell}}/N_{\mathrm{shell}}$
Fourier grid	Cubic grid, isotropic $k_{\mathrm{f}} = 2\pi/L$	Equal voxel spacing Δ, but anisotropic IR cutoffs $k_{\mathrm{f},\alpha} = 2\pi/L_\alpha$ for non-cubical boxes.
Interpolation	NGP/CIC/TSC, grid-coincident particles	Same kernels + compensation $W^{-1}(\boldsymbol{k})$ for off-grid particles. Multiresolution interpolation for zoom-in.

Verification of StePS initial conditions

Strategy: generate density fields from a known $P(k)$, apply LPT to obtain a particle distribution, measure the power spectrum back, and verify $P_{\mathrm{meas}}/P_{\mathrm{ref}} \approx 1$.

Result

Sub-percent agreement ($|P_{\mathrm{meas}}/P_{\mathrm{ref}} - 1| < 0.5\%$) up to $\sim k_{\mathrm{Ny}}/2$, across all tested resolutions (128^3–512^3), redshifts ($z = 15$–63), LPT orders (1st and 2nd), mass assignment schemes (NGP, CIC, TSC), and box aspect ratios (1:1 to 10:1).

The cubic-voxel Fourier grid introduces **no geometry-dependent systematic bias**.

External benchmark: `monofonic`

White noise generated in `stepsic` and fed directly into `monofonic`. Both produce density and displacement fields that agree at the level expected from floating-point discretization across the full resolved k-range.

Verification of StePS initial conditions (visual)

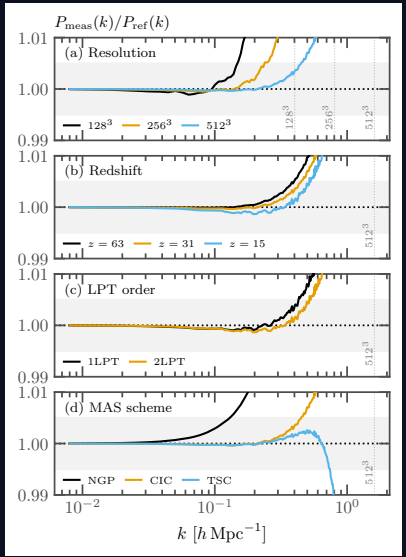

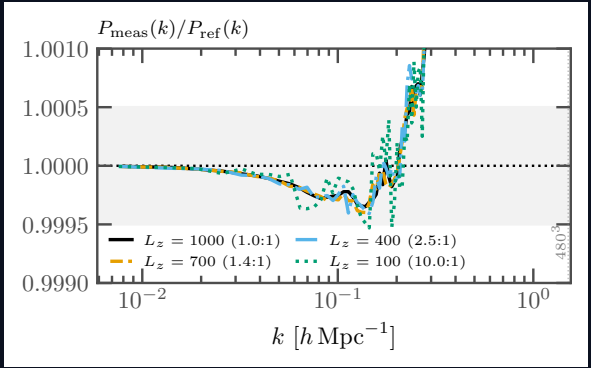

Figure 6: Aspect ratios from 1:1 to 10:1. The cubic-voxel Fourier grid introduces **no geometry-dependent systematic bias**. Flatter slabs show a slight accuracy trade-off from reduced volume, not from the grid construction itself.

The first cosmological simulation in $\mathbb{S}^1 \times \mathbb{R}^2$ — run parameters

Simulation setup

Cosmology	Planck 2018 ΛCDM
H_0	67.66 km/s/Mpc
Ω_m	0.3111
σ_8	0.8102
R_sim	500 Mpc
L_z	1.0 Gpc
Particles	2.4×10^7
Mass range	10^{11}–10^{13} $M_\odot$
z_init	49 (2LPT)
Timesteps	3085
Hardware	640 CPU cores

Key results

- Cosmic web forms naturally: voids, filaments, walls, halos. Visually identical to $\mathbb{T}^3$ and $\mathbb{R}^3$.
- $P(k)$ follows linear CAMB on large scales and non-linear HaloFit on small scales.
- 1.1×10^5 DM halos at $z = 0$, identified with a custom SO halo finder ($\textsc{StePS\_HF}$).
- FKP estimator used for $P(k)$, designed for non-uniform selection functions.

Testing for topological anisotropy — motivation

If $\mathbb{S}^1 \times \mathbb{R}^2$ gravity (anisotropic: weaker axially, stronger radially) imprints a physical signal, it should appear in the angular momentum distribution of dark matter halos.

$$\theta_{J_z} = \arccos\left(\frac{J_z}{|\boldsymbol{J}|}\right) \qquad \text{— angle between halo spin and the periodic } z\text{-axis}$$

Motivation: Shamir (2025) reported a statistically significant deviation from isotropy in galaxy spin directions from JWST/JADES data. Since gravity is inherently anisotropic in $\mathbb{S}^1 \times \mathbb{R}^2$, we tested whether the topology induces a corresponding signal.

Result: null

The distribution of $\cos\theta_{J_z}$ is consistent with uniformity in all mass bins at all output redshifts. One-sample KS test: $D < 0.02$, $p > 0.28$ ($N_{\mathrm{part}} > 100$ per halo).

The $\mathbb{S}^1 \times \mathbb{R}^2$ topology with $L_z = 1$ Gpc does not provide a natural explanation for the JWST spin asymmetry. Smaller L_z could in principle enhance anisotropy, but such values are strongly constrained by Planck and SDSS.

Future directions

3rd-order LPT	Further reduces transient errors at moderate initial redshifts; minimal additional computational cost.
Primordial non-Gaussianity	Local-type $f_{\rm NL}$ modifications to the initial density field — broadens the range of testable cosmological models.
Rotating cosmologies	Shear-free Heckmann–Schücking models in $\mathbb{S}^1 \times \mathbb{R}^2$ and $\mathbb{R}^3$ topologies. CMB isotropy is preserved at all ω; searching for late-time dynamical signatures (source counts, polarisation rotation) that could independently constrain $\omega_0/H_0 \lesssim \mathcal{O}(10^{-2})$.
GPU-accelerated octree	OpenMP-CUDA Barnes–Hut implementation for $> 10^8$ particle counts in cylindrical topology.

All codes are open source: `github.com/eltevo/StePS` · `github.com/eltevo/stepsic`

Structure formation in the $\mathbb{S}^1 \times \mathbb{R}^2$ StePS geometry with rotation

Particle load configurations in stepsic

Grid Rectilinear mesh, cubic voxels. $\Delta = \min(L_i)/N_{\text{input}}$. Cuboid geometries only.

Random Homogeneous Poisson process. Primarily for testing and glass generation seed.

Shells Concentric spherical shells or cylindrical annuli, N_{shell} particles each. Variable mass: $m_i = \bar{\varrho}\, V_{\text{shell},i}/N_{\text{shell}}$.

Glass Pre-relaxed via repulsive gravity (StePS or GADGET). Automatic coordinate/mass rescaling.

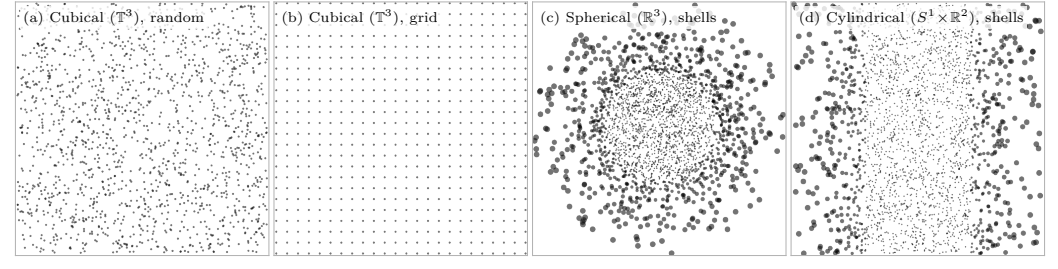

Anisotropic displacement in slab geometries

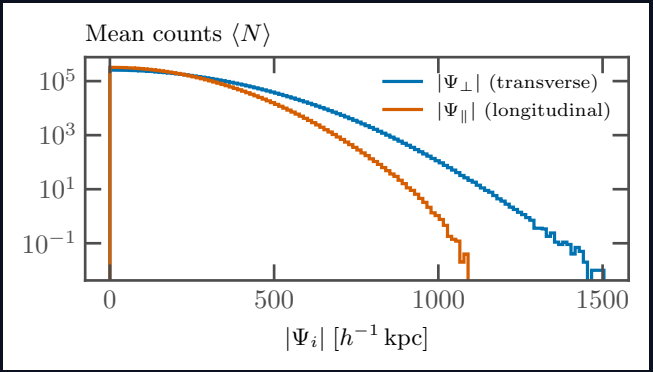

Per-component absolute displacements in a 5:1 periodic slab, averaged over 200 independent realizations.

Setup: $L_x = L_y = 1000$, $L_z = 200\ h^{-1}\mathrm{Mpc}$; equal physical voxel spacing Δ along every axis.

Mode selection: the short axis imposes $k_{\mathrm{f},z} = 2\pi/L_z = 5\,k_{\mathrm{f},\perp}$.

Result: missing long-wavelength z-modes suppress $|\Psi_\parallel| = |\Psi_z|$ relative to $|\Psi_\perp| = \frac{1}{2}(|\Psi_x| + |\Psi_y|)$.

Interpretation: geometric conditioning of the allowed Fourier modes, not an anisotropic-resolution artifact or an isotropic crop.

Multiresolution interpolation scheme

Heavy outer shells span large effective volumes and resolve only long-wavelength modes; light inner particles need a fine mesh to see small-scale structure. A single resolution either starves the centre or injects unresolved noise into the periphery.

Stack N_{grid} grids, one per distinct mass:

$$m_j \longmapsto N_{\mathrm{vox}}^{(j)} = \frac{L}{\left(V\, m_j / M_{\mathrm{box}}\right)^{1/3}},$$

all covering the same volume. Each grid yields its own displacement and velocity field; for every particle we then *linearly interpolate across the stack as a function of its mass*.

The non-obvious step

All resolution levels are seeded from the *same* white noise realization.

Modes resolved on multiple grids share **identical amplitudes and phases**; coarser grids simply truncate the high-k tail.

Heavy particles never receive modes their volume cannot resolve $\Rightarrow$ no noise injection. Light particles see the full small-scale field. The stack is internally consistent by construction.

Mass assignment and compensation kernels

Order-p B-spline assignment functions ($p \in \{1, 2, 3\}$ for NGP, CIC, TSC) have Fourier transforms:

$$W(\boldsymbol{k}) = \prod_{i=x,y,z} \left[\frac{\sin(\pi k_i / 2k_{\mathrm{Ny}})}{\pi k_i / 2k_{\mathrm{Ny}}} \right]^p$$

The interpolation acts as a low-pass filter. To compensate when interpolating grid-based fields onto off-grid particle positions, a deconvolution factor is applied *before* the inverse FFT:

$$W^{-1}(\boldsymbol{k}) = \prod_{i=x,y,z} \left[\mathrm{sinc}\left(\frac{k_i}{2k_{\mathrm{Ny}}} \right) \right]^{-p}$$

This restores unsmoothed field amplitudes for modes below k_{Ny}. The correction is unnecessary when particles lie exactly on the grid (standard periodic case with grid IC).

Variance reduction: paired and fixed amplitudes

`stepsic` supports two variance-reduction techniques for precision validation:

Phase-shifted (paired)

A global phase offset $\varphi_0 = \pi$ is applied: $\delta(\boldsymbol{k}) \to \delta(\boldsymbol{k})\, e^{i\varphi_0}$. Phases are exactly anti-correlated $\to$ sample variance partially cancels when two runs are averaged.

Fixed-amplitude

Stochastic Rayleigh amplitudes $|W(\boldsymbol{k})|$ are replaced by deterministic target amplitudes from $P(|\boldsymbol{k}|)$. Only random phases are retained. Combined with pairing: *paired-and-fixed*.

Caveat

Fixed amplitudes artificially reduce power spectrum variance $\to$ bias higher-order statistics (e.g. covariance). Only suitable for first-order comparisons.

Ewald summation in cosmological simulations

In a periodic cosmological box, one particle represents an infinite lattice of copies. The force is therefore not just the nearest $1/r^2$ pair force, but the conditionally convergent sum over all periodic images:

$$\boldsymbol{F}_{\mathrm{per}}(\boldsymbol{r}) = -Gm \sum_{\boldsymbol{n}\in\mathbb{Z}^d} \frac{\boldsymbol{r}+\boldsymbol{n}L}{|\boldsymbol{r}+\boldsymbol{n}L|^3}$$ with the homogeneous background treated separately.

Real-space part

A Gaussian screening makes nearby image forces rapidly convergent. This gives a short-range sum in real space.

Fourier-space part

The smooth long-range remainder is summed over reciprocal modes. The $k = 0$ mode is omitted or absorbed into the mean-density/background term.

In StePS this is used as a *force correction*: compute the ordinary nearest-image softened Newtonian force, then add a precomputed Ewald correction table for the chosen topology.

Ewald summation in $\mathbb{S}^1 \times \mathbb{R}^2$ topology

StePS uses a singly-periodic Ewald sum to tabulate the correction to the nearest-image Newtonian force. The full pair force is split into real-space, non-zero Fourier, and $m=0$ pieces:

$$\boldsymbol{F}_{\mathrm{1P}} = \boldsymbol{F}^{\mathrm{real}} + \boldsymbol{F}^{m\neq 0} + \boldsymbol{F}^{m=0}, \qquad \boldsymbol{D}(\varrho, z) = \boldsymbol{F}_{\mathrm{1P}} - \boldsymbol{F}_{\mathrm{NI}}$$

Real Erfc-damped direct sum over the image chain only: $\boldsymbol{r}_n = \boldsymbol{r} + n\, L_z\, \hat{\boldsymbol{e}}_z$.

$m \neq 0$ Discrete axial modes $k_m = 2\pi m / L_z$; the transverse integrals are done analytically. The code evaluates the erfc-based kernel with stable $\exp(\pm k_m \varrho)\,\mathrm{erfc}(\cdots)$ products, so no numerical Bessel evaluation is needed.

$m=0$ Zero mode contributes only to the radial component: $F_{\varrho}^{m=0} = -\frac{2}{L_z\,\varrho}(1 - e^{-\alpha^2 \varrho^2})$.

Runtime stores only (D_ϱ, D_z) on an $N_\varrho \times N_z$ table ($\varrho_{\max} = 2.25\, R_{\mathrm{sim}}$), interpolates it with NGP/CIC/TSC, and rotates D_ϱ back into (x, y).

Force field comparison: $\mathbb{T}^3$ vs. $\mathbb{S}^1 \times \mathbb{R}^2$ vs. $\mathbb{R}^3$

$\mathbb{T}^3$ **(periodic cube):**
Nearest-image softened Newtonian force plus a 3D Ewald correction table $\boldsymbol{D}(x, y, z)$. Only $\mathrm{O_h}$ symmetry remains, so the correction is orientation-dependent, not an isotropic rescaling of $1/r^2$.

$\mathbb{S}^1 \times \mathbb{R}^2$ **(cylinder):**
Only z is wrapped. The table stores (D_ϱ, D_z) and rotates D_ϱ into (x, y). The axial piece carries the image-chain periodicity; the transverse piece carries the line-image and zero-mode response.

$\mathbb{R}^3$ **(free boundary conditions):**
Direct softened Newtonian pair force. No image correction table.

Background compensation:
Comoving cylindrical runs subtract the homogeneous cylinder's radial collapse with a transverse linear term:

$$\boldsymbol{a}_{\mathrm{bg}} = 2\pi G\, \bar{\varrho}\, \varrho\, \hat{\boldsymbol{e}}_\varrho \qquad (a_z = 0).$$

Full Ewald uses this exact factor, the nearest-image cylindrical path uses a finite-cylinder radial lookup.