


Starting Cosmological Simulations from the Big Bang

Florian List¹,* Oliver Hahn², and Cornelius Rampf¹

¹Department of Astrophysics, University of Vienna, Türkenschanzstraße 17, 1180 Vienna, Austria
²and Department of Mathematics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria

 (Received 26 September 2023; revised 4 February 2024; accepted 29 February 2024; published 26 March 2024)

The cosmic large-scale structure (LSS) provides a unique testing ground for connecting fundamental physics to astronomical observations. Modeling the LSS requires numerical N -body simulations or perturbative techniques that both come with distinct shortcomings. Here we present the first unified numerical approach, enabled by new time integration and discreteness reduction schemes, and demonstrate its convergence at the field level. In particular, we show that our simulations (i) can be initialized directly at time zero, and (ii) can be made to agree with high-order Lagrangian perturbation theory in the fluid limit. This enables fast, self-consistent, and UV-complete forward modeling of LSS observables.

DOI: 10.1103/PhysRevLett.132.131003

Introduction.—The gravitational evolution of collisionless matter is governed by the cosmological Vlasov-Poisson system (VP) [1–4], which describes how the phase-space distribution $f = f(t, \mathbf{x}, \mathbf{p})$ of a continuous medium evolves,

$$\frac{df}{dt} = \frac{\partial f}{\partial t} + \frac{\mathbf{p}}{a^2} \cdot \nabla_{\mathbf{x}} f - \nabla_{\mathbf{x}} \varphi \cdot \nabla_{\mathbf{p}} f = 0, \quad (1)$$

where the gravitational potential is subject to Poisson's equation $\nabla_{\mathbf{x}}^2 \varphi = 3/(2a)H_0^2 \Omega_m \delta$. Here, (\mathbf{x}, \mathbf{p}) is the canonical position-momentum pair, a is the scale factor, H_0 is the Hubble constant, Ω_m is today's density parameter, and $\delta = \rho/\bar{\rho} - 1 = \int_{\mathbb{R}^3} f d^3 p - 1$ is the density contrast.

At sufficiently early times and assuming that matter is perfectly cold, the first two kinetic moments of Eq. (1) form a closed set of fluid equations. This resulting Euler-Poisson system is the starting point for perturbative approaches to structure formation, which form the basic *theoretical* class of methods for studying the large-scale structure of the Universe: in Eulerian (standard) perturbation theory (e.g., [2]), the density contrast δ is expanded in a Taylor series, and a hierarchy of recursion relations for δ is derived. However, as density fluctuations grow and $\delta \sim 1$, this technique breaks down.

An alternative approach is given by Lagrangian perturbation theory (LPT) [5–8], where instead a series ansatz is used for the displacement field $\Psi(\mathbf{q}) = \mathbf{x}(\mathbf{q}) - \mathbf{q}$, i.e., the vector pointing from each Lagrangian position \mathbf{q} to the currently associated Eulerian position $\mathbf{x}(\mathbf{q})$ when moving along the fluid characteristics. All-order recursive solutions for Ψ are available [9–11], with the exact solution of the VP system arising in the limit of infinite order [12]. Although converging significantly faster than Eulerian perturbation theory, LPT eventually also breaks down,

namely, at the first shell crossing, i.e., when particle trajectories cross for the first time. Then, the (Eulerian) velocity field becomes multivalued, and the fluid description ceases to be valid as the Vlasov hierarchy can no longer be truncated at first order. Analytical post-shell-crossing approaches exist (e.g., [13–16]); however, they do not (yet) extend into the strongly nonlinear regime and are therefore not mature enough to be useful in practice. An alternative to this is, e.g., “effective field theory of large-scale structure” [17–19], which, however, relies on matching free parameters to a UV-complete approach, typically provided by simulations.

Hence, resolving the nonlinear late-time dynamics in a UV-complete manner requires *numerical* methods, with the most prominent technique given by N -body simulations. Here, the continuous phase-space distribution f is represented by a set of N discrete tracer particles with canonical positions and momenta $(\mathbf{X}_i, \mathbf{P}_i)$, for $i = 1, \dots, N$. Requiring $df(t, \mathbf{X}_i, \mathbf{P}_i)/dt = 0$ leads to the Hamiltonian equations of motion $\dot{\mathbf{X}}_i = \mathbf{P}_i/a^2$ and $\dot{\mathbf{P}}_i = -\nabla_{\mathbf{x}} \varphi|_{\mathbf{X}_i}$. Note that if one had access to the exact (continuous) potential φ , the particles would move *exactly* according to the characteristics of the underlying continuous system. However, since the true density contrast δ in the Poisson equation can only be approximated based on the positions of the N particles, an estimate $\delta_N \approx \delta$ sources the Poisson equation, resulting in an approximate potential $\varphi_N \approx \varphi$. This is the crucial approximation made by N -body simulations and, as we will see later, carefully designed techniques to improve the match $\varphi_N \rightarrow \varphi$ are therefore key in suppressing discreteness effects at early times and accessing the fluid limit with N -body simulations.

Although perturbation theory and N -body simulations are the theoretical and numerical pillars of modeling cosmological structure formation, there are only very

few studies on their agreement in the fluid limit at early times when perturbation theory is still valid. Comparison studies in this regime are hampered by the fact that spurious discreteness effects become significant at early times as the N -body system quickly deviates from the underlying continuous dynamics [20,21]. Techniques for correcting at least the linear discreteness error of the particle lattice exist [22], but are not widely employed. Despite discreteness errors, N -body simulations have traditionally been initialized using first-order LPT (the Zel'dovich approximation [5]) or, more recently, second-order LPT (2LPT) [6,23] at early times (redshift $z = a^{-1} - 1 \gtrsim 100$), to avoid truncation errors arising from the residual between LPT and the true solution, which ultimately bias the statistics of the simulated fields. In Ref. [24], it was recently shown that a more favorable trade-off between numerical discreteness errors and LPT truncation errors is achieved by initializing cosmological simulations at rather *late* times (e.g., $z \approx 15$ –40), by employing higher-order LPT, namely, 3LPT [25,26].

In this Letter, we bridge the gap between the analytical and numerical descriptions of cosmic structure formation in the fluid limit at early times. Specifically, we show for the first time that by applying an array of discreteness reduction techniques, together with a time integrator that has the correct asymptotic behavior for $a \rightarrow 0$, one obtains excellent agreement between N -body dynamics and perturbation theory. The choice of appropriate initial conditions (ICs) and time variable allows us to initialize N -body simulations at $a = 0$, enabling a clean comparison between N -body and LPT dynamics. Remarkably, a *single* N -body drift-kick-drift (DKD) step from $a = 0$ to a “typical” 3LPT initialization time for cosmological simulations yields a displacement field at ≈ 3 LPT accuracy. This effectively renders moot the LPT-based initialization of cosmological N -body simulations and demonstrates that starting them directly at $a = 0$ is a promising alternative.

The structure of this Letter is as follows. First, we briefly review the time integrator `PowerFrog`, which we recently introduced in Ref. [27]. This integrator is asymptotically consistent with 2LPT for $a \rightarrow 0$ and a crucial ingredient for achieving agreement between LPT and the N -body dynamics. Next, we describe the discreteness suppression techniques that enable us to achieve extremely low-noise results in the fluid limit at early times. Then, we present and discuss our results for a single N -body simulation step from $z = \infty$ to $z = 18$ (shortly before the time of the first shell crossing). Finally, we comment on the present-day (i.e., $z = 0$) statistics of N -body simulations initialized either directly at $z = \infty$ or with LPT. We find that while the power spectra match to within 1% even without applying any discreteness suppression techniques, these techniques are necessary in order to obtain the correct cross-power spectrum with $z = \infty$ -initialized simulations.

Π integrators.—The leapfrog integrator is ubiquitous in cosmological simulations thanks to its simplicity,

and suitability for individual time steps for different particles [28]. While it converges at second order towards the correct solution as the time step decreases, it does not exploit the fact that before shell-crossing the displacement field Ψ can be expressed analytically in the form of a series in the linear growth-time D of the Λ CDM concordance model, namely the LPT series $\Psi(\mathbf{q}, D) = \sum_{n=1}^{\infty} \psi^{(n)}(\mathbf{q}) D^n$. (We only consider growing-mode solutions and neglect higher-order LPT corrections stemming from the cosmological constant Λ ; see Refs. [29,30].)

In Ref. [27], we introduced a class of integrators, which we named Π integrators in view of the momentum variable $\mathbf{\Pi} = d\mathbf{X}/dD$ with respect to which they are formulated. Expressing the integrator in terms of momentum $\mathbf{\Pi}$ w.r.t. growth-factor time enables the construction of second-order accurate integration schemes which, when performing only few time steps, mimic LPT dynamics.

The only previously existing representative of this class is the popular `FastPM` scheme by Ref. [31], which was constructed to match the dynamics of the Zel'dovich approximation on large scales. One of our new integrators, which we named `PowerFrog`, further matches the 2LPT asymptote at early times $a \rightarrow 0$, which turns out to be essential for initializing simulations at $a = 0$, as we will see later.

As usual, we choose the ICs to be $\delta(D = 0) = 0$ and $\mathbf{\Pi}(D = 0) = -\nabla_{\mathbf{q}}\phi_{\text{ini}}$, which implicitly selects the growing-mode solution and ensures that the initial momentum is curl-free [32–34]; see, e.g., Ref. [24] for details on how ϕ_{ini} can be obtained from a standard Boltzmann code employing a standard backscaling approach. Notice that the canonical variables (\mathbf{X}, \mathbf{P}) are incompatible with these ICs: due to Liouville’s theorem for Hamiltonian mechanics, the contraction of the positions to a single point in the limit $a \rightarrow 0$ leads to the divergence of the momenta. This is not so, however, for the coordinates $(\mathbf{X}, \mathbf{\Pi})$, which are employed by Π integrators. In fact, the transformation from $(\mathbf{X}, \mathbf{P}) \mapsto (\mathbf{X}, \mathbf{\Pi})$ is noncanonical (but rather “contact,” see Refs. [35,36]), for which reason these new variables are not subject to Liouville’s theorem, and it is easy to see that the contact Hamiltonian for $(\mathbf{X}, \mathbf{\Pi})$ remains bounded for $a \rightarrow 0$ (subject to suitable ICs [37]).

Equipped with an integrator that works in terms of these variables and, by construction, is consistent with the 2LPT trajectory at early times, we will demonstrate that it is possible to start cosmological simulations at $a = 0$, with the particles placed on an unperturbed homogeneous grid [which approximates $\delta(D = 0) = 0$], and the growth-factor “Zel’dovich” momentum initialized as $\mathbf{\Pi}_i = -\nabla_{\mathbf{q}}\phi_{\text{ini}}|_{X_i}$.

We emphasize that—in contrast to LPT—the time integration of cosmological N -body systems using Π integrators is UV complete in that the N -body dynamics should converge towards the VP solution in the limit of infinitely many particles and time steps, even in the highly nonlinear multistreaming regime (albeit the mathematical proof thereof is still missing; but see, e.g., [38,39]).

Towards the fluid limit: Suppressing particle noise with sheet-based interpolation.—In this Letter, we focus on the particle-mesh (PM) method for the force computation, but we also briefly consider tree-PM [40,41] and the nonuniform fast Fourier transform (FFT). To control particle discreteness effects at the required level, we apply *four* important steps: (1) Increasing the number of gravity source particles (“resampling”) by *sheet interpolation* during the force calculation (extending the quadratic interpolation of Ref. [42] to Fourier interpolation, see also Ref. [43]) (2) Using *higher-order mass-assignment schemes* to represent particle positions more accurately [44] and deconvolving the density field on the grid with the mass assignment kernel (3) Using *grid interlacing* to suppress low-order aliases [28] (4) Using the *exact gradient kernel* $\mathbf{i}\mathbf{k}$ for the force calculation, rather than a finite difference gradient kernel. The sheet interpolation has by far the largest effect in terms of suppressing discreteness. It harnesses the fact that for cold ICs, the phase-space density $f(t, \mathbf{x}, \mathbf{p})$ in the VP equation (1) occupies a three-dimensional manifold in the six-dimensional phase space at all times, the Lagrangian submanifold. Hence, to increase the spatial resolution of the gravitational potential, we can “spawn” new N -body gravity source particles in Lagrangian space on a finer grid, determine their displacement by Fourier interpolation, and compute the resulting force on the N particles using this refined potential field.

For illustration, Fig. 1 shows a plot of the $z = 0$ density field with $N = 512^3$ particles, evaluated on a grid with $M = 1024^3$ cells, from a standard N -body simulation (*left half*) and a simulation with 5^3 -fold resampling of each particle for the density computation and the other discreteness reduction techniques applied in each simulation step (*right half*). The density field in the standard simulation is poorly sampled, particularly in underdense regions, with many cells containing no particles and hence $\delta_N = -1$. The resampling evidently suppresses discreteness and leads to a




FIG. 1. Slice through the overdensity field at $z = 0$ of a standard N -body simulation (*left half*) and a simulation with discreteness reduction techniques applied (in each time step, as well as for the computation of the plotted density slice, *right half*). The particle and grid resolutions for both cases are $N = 512^3$ and $M = 1024^3$.

much more continuous density field. (Because of increasing complexity, the sheet-based interpolation [42,45] should not be applied in halo regions without using any refinements. Figure 1 is intended for illustrative purposes; here, we employ resampling only in the fluid regime at early times when it is well suited to suppress discreteness.) For a detailed explanation of each of these techniques, we refer to the Supplemental Material [46].

Initializing simulations without LPT.—We will now perform a single PowerFrog DKD time step starting from $z = \infty$ (i.e., $a = 0$) to a redshift where one would typically initialize a cosmological simulation with 3LPT, namely, $z = 18$. We checked that the Jacobi determinant $\det(d\mathbf{x}^{\text{1LPT}}/d\mathbf{q}) > 0$ for all particles at that time, and the standard deviation of the density field $\sigma(\delta^{\text{1LPT}}) = 0.30$. Hence, the entire simulation box is still in the single-stream regime, for which there is strong evidence that LPT converges [12,15].

We consider the evolution of $N = 512^3$ particles in a periodic simulation box of edge length $L = 100 \text{ Mpc}/h$ subject to a flat Λ CDM cosmology with $\Omega_m = 0.3$, $H_0 = 67.11 \text{ km/s/Mpc}$, $n_s = 0.9624$, $\sigma_8 = 0.8$. We perform our computations on a single GPU, computing the forces with the PM method at grid resolution $M = 1024^3$.

Figure 2 depicts the residual between the 1-step N -body result and different LPT orders at $z = 18$; specifically, we show a slice of the displacement component Ψ_x . In view of PowerFrog being designed to only match the 2LPT asymptote for $a \rightarrow 0$, it might surprise that the 1-step N -body displacement lies even closer to 3LPT and 4LPT than to the 2LPT result. Intuitively, this can be understood by noting that the LPT terms are computed at the Lagrangian particle positions, that is, by pulling back the evolution of each particle to its initial location, while the kick in the N -body step updates the velocities at growth-factor time $\Delta D/2$ directly based on the potential that solves the Poisson equation at that time, which excites higher-order LPT terms; we leave a detailed investigation on this for future work.

We remark that also the velocity field is in good agreement with its LPT counterpart (see the Supplemental Material [46]). The excellent match between the positions and momenta of a single PowerFrog step and high-order LPT makes the initialization of cosmological simulations directly at the origin of time at $a = 0$ with PowerFrog (or another integrator that follows the ≥ 2 LPT asymptotic behavior for $a \rightarrow 0$) an attractive alternative to the traditional LPT-based ICs.

We now study how the different discreteness reduction techniques affect the numerical solution of the N -body simulation. Figure 3 shows the relative root-mean-square (rms) error of the displacements between a single N -body step from $z = \infty$ to $z = 18$ using all discreteness suppression methods discussed above, together with the results when omitting one of these techniques at a time.




FIG. 2. Residuals of the x displacement Ψ_x between our results with a *single* PowerFrog N -body step from redshift $z = \infty$ to $z = 18$ and the corresponding LPT fields at $z = 18$ for different LPT orders. Shown is a slice in the Lagrangian y - z coordinate plane.

Clearly, the sheet-based resampling of the density is crucial for achieving convergence between N -body and LPT: without it, the residual towards LPT is entirely dominated by errors due to the particle-based approximation of the continuous density at a level of $\sim 50\%$, and no differences between the different LPT orders are visible. The second-most important technique is the deconvolution of the density with the mass assignment kernel, whose absence results in significant high-frequency noise that conceals the

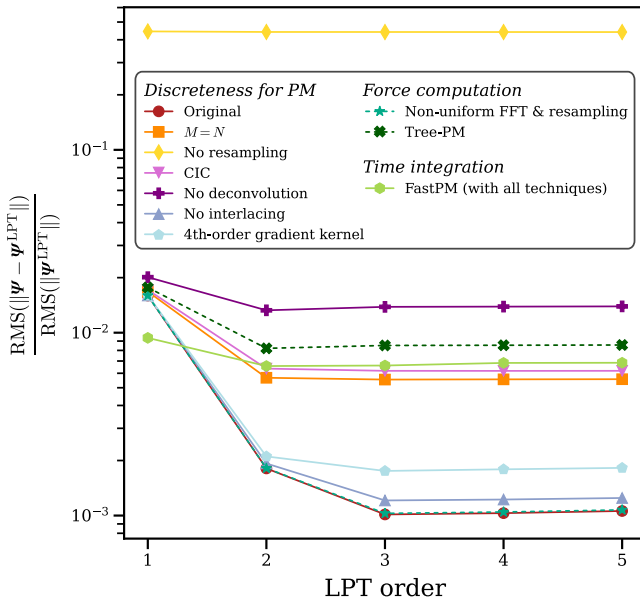


FIG. 3. Relative rms displacement error between the 1-step N -body simulation and different LPT orders at $z = 18$ when using the PowerFrog integrator and applying all discreteness reduction techniques (“original”), when omitting one technique at a time, and when performing a FastPM DKD step instead of a PowerFrog step (applying all discreteness reduction techniques). Evidently, a carefully designed time integrator, resampling, deconvolution, higher-order mass assignment, and a fine PM grid (e.g., $M = 2^3 N$) are all key ingredients to access the 3LPT regime. We also show the residuals when using the nonuniform FFT instead of local mass assignment (with resampling), and with tree-PM force computation (without resampling).

3–4LPT contributions in the residual. The residual also increases substantially when reducing the number of PM grid cells from $M = 2^3 N$ to $M = N$ or when using cloud-in-cell (CIC) instead of piecewise-cubic spline (PCS) mass assignment. The impact of dealiasing the density by means of interlaced grids, and of using the exact Fourier gradient kernel $i\mathbf{k}$ instead of a 4th-order finite-difference gradient kernel is much more modest; however, leaving out any of these methods imprints a characteristic grainy structure in the 3–4LPT residuals. With all techniques active, the 3LPT vs 1-step PowerFrog residual is only 0.1%. For completeness, we also show the residuals when replacing the local PM mass assignment by Ref. [66]’s implementation of the nonuniform FFT [67–69] together with resampling, which yields the same residuals as our PM baseline and could be an exciting avenue for future exploration. Using tree-PM for the force computation (without resampling) rather than PM also somewhat reduces discreteness, but much less than in our discreteness-suppressed PM baseline.

Finally, the green hexagons show the residual when performing a single step with the (DKD variant of the) FastPM stepper instead of PowerFrog which, recall, is consistent with the Zel’dovich approximation, but whose asymptotic behavior for $a \rightarrow 0$ differs from 2LPT. Clearly, there is a significant 2LPT contribution in the residual, which prevails in the residuals with respect to higher LPT orders. A plot of the residual fields and their power spectra can be found in the Supplemental Material [46]. FastPM is therefore not suitable as a 1-step initializer.

In principle, it should be possible to construct integrators that match even higher LPT orders with a single time step by composing each step out of more than three drift or kick components, but the gain from going beyond 3LPT can be expected to be relatively small in practical applications. After the first time step, the assumption that the time step starts in the asymptotic regime at $a \approx 0$ is no longer exactly valid, and second-order-in- a residuals arise with PowerFrog; we will explore potential improvements in this regard in future work.

Analysis at $z = 0$.—Finally, let us comment on the results one obtains when using the positions and momenta

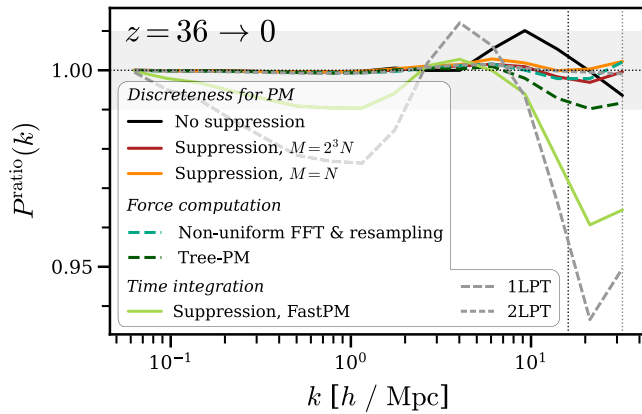


FIG. 4. Density power spectrum ratio at $z = 0$ between cosmological simulations with 1LPT, 2LPT, or 1-step N -body ICs and 3LPT ICs. Even without any discreteness suppression, starting cosmological simulations with a single PowerFrog step from $z = \infty$ to 36 and then continuing in standard N -body fashion leads to power spectrum errors $< 1\%$ towards 3LPT on all scales. The dotted vertical lines show the Nyquist modes for $N = 512^3$ and $M = 1024^3$, respectively.

computed with a single PowerFrog step as the ICs for a (standard) cosmological simulation down to $z = 0$. We take the same cosmology as in the previous section, $N = 512^3$ particles, and initialize the simulations either with 1, 2, or 3LPT, or with a single N -body time step (that starts from $z = \infty$) at $z = 36$; then, we run a cosmological simulation with the industry standard code Gadget-4 [70]. For the single N -body step, we consider PowerFrog (i) with discreteness suppression and grid sizes $M = 512^3$ and 1024^3 , (ii) without any discreteness suppression and $M = 512^3$, (iii) a FastPM DKD step with discreteness suppression, (iv) non-uniform FFT forces with resampling, and (v) tree-PM forces (without resampling).

Figure 4 shows the power spectrum ratio at $z = 0$ w.r.t. the 3LPT ICs, which we take as our reference. Interestingly, even without any discreteness suppression, the residual between the power spectra with PowerFrog and 3LPT ICs is $\leq 1\%$ on all scales. Also, for the equilateral bispectrum, we find excellent agreement; however, the cross-spectrum drops significantly when omitting the discreteness reduction (e.g., from 99% to 90% at $k = 21 h/\text{Mpc}$, see the Supplemental Material [46]). This implies that, in principle, standard N -body simulations can be started with a PowerFrog-like stepper without any discreteness suppression, and the resulting $z = 0$ density field will be correct in terms of power spectrum and bispectrum, but its phases will be somewhat corrupted due to the discreteness.

Discussion.—In this Letter, we have provided the first demonstration of the field-level agreement between high-order LPT and cosmological N -body simulations in the single-stream regime. Choosing kinematic variables in which the solution remains regular for $a \rightarrow 0$ allowed us to initialize simulations at the origin of time, making the

customary LPT-based computation of the initial conditions at some scale factor $a > 0$ obsolete—provided discreteness artifacts are sufficiently suppressed. Remarkably, the use of an LPT-informed time integrator implies that a *single* N -body step starting from $a = 0$ yields more accurate results than 2LPT, which is the established technique for initializing cosmological simulations.

From a practical point of view, this opens up a wide range of applications: the computational cost of the discreteness-suppressed step we applied to obtain the close match with LPT at early times shown in Fig. 2 is similar to that of 3LPT, but already applying very few or even none of these techniques might give sufficiently accurate results in fast simulations and for analyses focused on late times (see the highly accurate power spectrum for “No suppression” in Fig. 4). Also, an N -body initialization step from $a = 0$ might be superior in terms of memory requirements—no matter how fine the grid in Lagrangian space used for the resampling—as no large arrays need to be stored for each LPT order. Another interesting scope of application is given by zoom simulations, where the intricacies in the (usually FFT-based, but cf. [22] for 2LPT computed in configuration space) LPT computation arising from different resolutions can be circumvented.

In the era of precision cosmology, it is crucial to thoroughly test the agreement of complementary approaches to structure formation such as perturbative techniques and numerical methods and to clearly identify their range of validity. Our findings in this Letter lay the groundwork for further comparison studies at the intersection between analytical and numerical methods.

We thank Raul Angulo and Jens Stücker for insightful discussions. O.H. thanks Tom Abel for many past discussions on discreteness and the sheet. A software package implementing the discussed algorithms will be released in the near future.

*florian.list@univie.ac.at

- [1] P. Peebles, *The Large-scale Structure of the Universe*, Princeton Series in Physics (Princeton University Press, Princeton, 1980).
- [2] F. Bernardeau, S. Colombi, E. Gaztanaga, and R. Scoccimarro, *Phys. Rep.* **367**, 1 (2002).
- [3] C. Rampf, *Rev. Mod. Plasma Phys.* **5**, 10 (2021).
- [4] R. E. Angulo and O. Hahn, *Living Rev. Comput. Astrophys.* **8**, 1 (2022).
- [5] Ya. B. Zel’dovich, *Astron. Astrophys.* **5**, 84 (1970).
- [6] T. Buchert and J. Ehlers, *Mon. Not. R. Astron. Soc.* **264**, 375 (1993).
- [7] F. R. Bouchet, S. Colombi, E. Hivon, and R. Juszkiewicz, *Astron. Astrophys.* **296**, 575 (1995).
- [8] J. Ehlers and T. Buchert, *Gen. Relativ. Gravit.* **29**, 733 (1997).
- [9] C. Rampf, *J. Cosmol. Astropart. Phys.* **12** (2012) 004.

- [10] V. Zheligovsky and U. Frisch, *J. Fluid Mech.* **749**, 404 (2014).
- [11] T. Matsubara, *Phys. Rev. D* **92**, 023534 (2015).
- [12] S. Saga, A. Taruya, and S. Colombi, *Phys. Rev. Lett.* **121**, 241302 (2018).
- [13] S. Colombi, *Mon. Not. R. Astron. Soc.* **446**, 2902 (2015).
- [14] A. Taruya and S. Colombi, *Mon. Not. R. Astron. Soc.* **470**, 4858 (2017).
- [15] C. Rampf and O. Hahn, *Mon. Not. R. Astron. Soc.* **501**, L71 (2021).
- [16] S. Saga, A. Taruya, and S. Colombi, *Astron. Astrophys.* **664**, A3 (2022).
- [17] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **07** (2012) 051.
- [18] J. J. M. Carrasco, M. P. Hertzberg, and L. Senatore, *J. High Energy Phys.* **09** (2012) 082.
- [19] G. Cabass, M. M. Ivanov, M. Lewandowski, M. Mirbabayi, and M. Simonović, *Phys. Dark Universe* **40**, 101193 (2023).
- [20] M. Joyce, B. Marcos, A. Gabrielli, T. Baertschiger, and F. Sylos Labini, *Phys. Rev. Lett.* **95**, 011304 (2005).
- [21] B. Marcos, T. Baertschiger, M. Joyce, A. Gabrielli, and F. Sylos Labini, *Phys. Rev. D* **73**, 103507 (2006).
- [22] L. H. Garrison, D. J. Eisenstein, D. Ferrer, M. V. Metchnik, and P. A. Pinto, *Mon. Not. R. Astron. Soc.* **461**, 4125 (2016).
- [23] F. R. Bouchet, R. Juszkiewicz, S. Colombi, and R. Pellat, *Astrophys. J. Lett.* **394**, L5 (1992).
- [24] M. Michaux, O. Hahn, C. Rampf, and R. E. Angulo, *Mon. Not. R. Astron. Soc.* **500**, 663 (2021).
- [25] T. Buchert, *Mon. Not. R. Astron. Soc.* **267**, 811 (1994).
- [26] P. Catelan, *Mon. Not. R. Astron. Soc.* **276**, 115 (1995).
- [27] F. List and O. Hahn, [arXiv:2301.09655](https://arxiv.org/abs/2301.09655).
- [28] R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles* (1st ed.) (CRC Press, Boca Raton, 1988).
- [29] C. Rampf, S. O. Schobesberger, and O. Hahn, *Mon. Not. R. Astron. Soc.* **516**, 2840 (2022).
- [30] M. Fasiello, T. Fujita, and Z. Vlah, *Phys. Rev. D* **106**, 123504 (2022).
- [31] Y. Feng, M.-Y. Chu, U. Seljak, and P. McDonald, *Mon. Not. R. Astron. Soc.* **463**, 2273 (2016).
- [32] Y. Brenier, U. Frisch, M. Hénon, G. Loeper, S. Matarrese, R. Mohayaee, and A. Sobolevskii, *Mon. Not. R. Astron. Soc.* **346**, 501 (2003).
- [33] U. Frisch, S. Matarrese, R. Mohayaee, and A. Sobolevski, *Nature (London)* **417**, 260 (2002).
- [34] C. Rampf, *Mon. Not. R. Astron. Soc.* **484**, 5223 (2019).
- [35] V. Arnold, *Mathematical Methods of Classical Mechanics* (Springer, New York, 1989), Vol. 60.
- [36] A. Bravetti, H. Cruz, and D. Tapias, *Ann. Phys. (Amsterdam)* **376**, 17 (2017).
- [37] C. Rampf, C. Uhlemann, and O. Hahn, *Mon. Not. R. Astron. Soc.* **503**, 406 (2021).
- [38] S. Colombi, *Astron. Astrophys.* **647**, A66 (2021).
- [39] M. Feistl and P. Pickl, [arXiv:2307.06146](https://arxiv.org/abs/2307.06146).
- [40] J. S. Bagla, *J. Astrophys. Astron.* **23**, 185 (2002).
- [41] P. Bode and J. P. Ostriker, *Astrophys. J. Suppl. Ser.* **145**, 1 (2003).
- [42] O. Hahn and R. E. Angulo, *Mon. Not. R. Astron. Soc.* **455**, 1115 (2016).
- [43] J. Stücker, P. Busch, and S. D. White, *Mon. Not. R. Astron. Soc.* **477**, 3230 (2018).
- [44] A. Chaniotis and D. Poulidakos, *J. Comput. Phys.* **197**, 253 (2004).
- [45] J. Stücker, O. Hahn, R. E. Angulo, and S. D. White, *Mon. Not. R. Astron. Soc.* **495**, 4943 (2020).
- [46] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.132.131003> for mathematical details and additional results, which includes Refs. [47–65].
- [47] Y. Brenier, U. Frisch, M. Hénon, G. Loeper, S. Matarrese, R. Mohayaee, and A. Sobolevskii, *Mon. Not. R. Astron. Soc.* **346**, 501 (2003).
- [48] M. Fasiello, T. Fujita, and Z. Vlah, *Phys. Rev. D* **106**, 123504 (2022).
- [49] S. Nadkarni-Ghosh and D. F. Chernoff, *Mon. Not. R. Astron. Soc.* **410**, 1454 (2011).
- [50] C. Rampf and O. Hahn, *Phys. Rev. D* **107**, 023515 (2023).
- [51] C. Rampf, S. Saga, A. Taruya, and S. Colombi, *Phys. Rev. D* **108**, 103513 (2023).
- [52] M. Crocce, S. Pueblas, and R. Scoccimarro, *Mon. Not. R. Astron. Soc.* **373**, 369 (2006).
- [53] T. Abel, O. Hahn, and R. Kaehler, *Mon. Not. R. Astron. Soc.* **427**, 61 (2012).
- [54] S. Shandarin, S. Habib, and K. Heitmann, *Phys. Rev. D* **85**, 083005 (2012).
- [55] O. Hahn, T. Abel, and R. Kaehler, *Mon. Not. R. Astron. Soc.* **434**, 1171 (2013).
- [56] T. Sousbie and S. Colombi, *J. Comput. Phys.* **321**, 644 (2016).
- [57] J. Wang and S. D. M. White, *Mon. Not. R. Astron. Soc.* **380**, 93 (2007).
- [58] A. B. Romeo, O. Agertz, B. Moore, and J. Stadel, *Astrophys. J.* **686**, 1 (2008).
- [59] V. Springel, *Mon. Not. R. Astron. Soc.* **364**, 1105 (2005).
- [60] O. Hahn and T. Abel, *Mon. Not. R. Astron. Soc.* **415**, 2101 (2011).
- [61] E. Sefusatti, M. Crocce, R. Scoccimarro, and H. M. Couchman, *Mon. Not. R. Astron. Soc.* **460**, 3624 (2016).
- [62] L. Chen, A. Bruce Langdon, and C. K. Birdsall, *J. Comput. Phys.* **14**, 200 (1974).
- [63] A. W. Appel, *SIAM J. Sci. Comput.* **6**, 85 (1985).
- [64] J. Barnes and P. Hut, *Nature (London)* **324**, 446 (1986).
- [65] C. Uhlemann, C. Rampf, M. Gosenca, and O. Hahn, *Phys. Rev. D* **99**, 083524 (2019).
- [66] jax-finufft: JAX bindings to the Flatiron Institute Non-uniform Fast Fourier Transform (FINUFFT) library (2021), <https://github.com/flatironinstitute/jax-finufft>.
- [67] A. H. Barnett, J. Magland, and L. af Klinteberg, *SIAM J. Sci. Comput.* **41**, C479 (2019).
- [68] A. H. Barnett, *Appl. Comput. Harmon. Anal.* **51**, 1 (2021).
- [69] Y.-h. Shih, G. Wright, J. Andén, J. Blaschke, and A. H. Barnett, in *Proceedings of the 2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (IEEE, New York, 2021), pp. 688–697; [arXiv:2102.08463](https://arxiv.org/abs/2102.08463).
- [70] V. Springel, R. Pakmor, O. Zier, and M. Reinecke, *Mon. Not. R. Astron. Soc.* **506**, 2871 (2021).