



HAL
open science

Accurate initial conditions for cosmological N-body simulations: Minimizing truncation and discreteness errors

Michaël Michaux, Oliver Hahn, Cornelius Rampf, Raul E. Angulo

► To cite this version:

Michaël Michaux, Oliver Hahn, Cornelius Rampf, Raul E. Angulo. Accurate initial conditions for cosmological N-body simulations: Minimizing truncation and discreteness errors. *Monthly Notices of the Royal Astronomical Society*, 2020, 500 (1), pp.663-683. 10.1093/mnras/staa3149 . hal-02939987

HAL Id: hal-02939987

<https://hal.science/hal-02939987>

Submitted on 28 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accurate initial conditions for cosmological N -body simulations: minimizing truncation and discreteness errors

Michaël Michaux ¹★, Oliver Hahn ¹, Cornelius Rampf ¹ and Raul E. Angulo ^{2,3}

¹Observatoire de la Côte d’Azur, CNRS, Laboratoire Lagrange, Université Côte d’Azur, Boulevard de l’Observatoire, CS 34229, F-06304 Nice, France

²Donostia International Physics Center (DIPC), Paseo Manuel de Lardizabal, 4, E-20018 Donostia-San Sebastián, Spain

³IKERBASQUE, Basque Foundation for Science, E-48013 Bilbao, Spain

Accepted 2020 September 28. Received 2020 August 28; in original form 2020 May 25

ABSTRACT

Inaccuracies in the initial conditions for cosmological N -body simulations could easily be the largest source of systematic error in predicting the non-linear large-scale structure. From the theory side, initial conditions are usually provided by using low-order truncations of the displacement field from Lagrangian perturbation theory, with the first- and second-order approximations being the most common ones. Here, we investigate the improvement brought by using initial conditions based on third-order Lagrangian perturbation theory (3LPT). We show that with 3LPT, truncation errors are vastly suppressed, thereby opening the portal to initializing simulations accurately as late as $z = 12$ (for the resolution we consider). We analyse the competing effects of perturbative truncation and particle discreteness on various summary statistics. Discreteness errors are essentially decaying modes and thus get strongly amplified for earlier initialization times. We show that late starting times with 3LPT provide the most accurate configuration, which we find to coincide with the continuum fluid limit within 1 per cent for the power- and bispectrum at $z = 0$ up to the particle Nyquist wavenumber of our simulations ($k \sim 3h \text{ Mpc}^{-1}$). In conclusion, to suppress non-fluid artefacts, we recommend initializing simulations as late as possible with 3LPT. We make our 3LPT initial condition generator publicly available.

Key words: dark matter – large-scale structure of Universe – cosmology: theory.

1 INTRODUCTION

Current and upcoming space and ground-based galaxy clustering and weak gravitational lensing surveys – such as HSC (Aihara et al. 2018), the LSST (Abell et al. 2009), and the Euclid satellite (Laureijs et al. 2011), but also future instruments probing the gas distribution across cosmic history, such as the Square Kilometre Array (Weltman et al. 2020), will test the validity of, and possible deviations from, the concordance Λ CDM model with unprecedented precision.

The accuracy of future observational data is expected to be such that it will require 1 per cent accurate predictions for the spatial distribution of matter in the late-time Universe to scales of $k \sim 1 \text{ Mpc}^{-1}$ (e.g. Heitmann et al. 2010; Schneider et al. 2016). Numerical simulations are, in principle, able solve the relevant equations with such accuracy over volumes comparable to the entire visible Universe (e.g. Angulo et al. 2012; Heitmann et al. 2016; Potter, Stadel & Teyssier 2017; Heitmann et al. 2019; Cheng et al. 2020). However, numerical solutions are computationally demanding and suffer from numerical artefacts and discretization errors. On the other hand, perturbative approaches of the underlying equations are not affected by such effects and deliver computationally cheap predictions. Yet their validity is limited to certain length and time-scales, since the gravitational collapse of cosmic structure is

intimately tied to the development of extreme densities that still pose a challenge for such approaches.

The worlds of analytic and numerical approaches come together when the formation of cosmic structures is investigated via cosmological simulations. Indeed, perturbative approaches essentially solve for the structure formation at early times, while numerical solutions are best at evolving later stages. Consequently, perturbation theory is employed to generate perturbatively truncated initial conditions (ICs) for cosmological simulations (as pioneered by Klypin & Shandarin 1983; Efstathiou et al. 1985). Since both perturbation theory and numerical solutions have their strengths and weaknesses (see further below), it is important to find the optimal window where ICs can be provided while simultaneously minimizing perturbative and numerical errors. Quantifying these errors and finding this optimal cosmic time for ICs are precisely the tasks of this paper.

Observations indicate that cold dark matter (CDM) is extremely weakly (self-)interacting, implying that it can be treated to be effectively collisionless with zero temperature on cosmological scales. The evolution of a continuous and collisionless medium, such as CDM, is governed by the Vlasov–Poisson equations. Most perturbative approaches solve these equations in the single-stream limit (vanishing velocity dispersion), which appears to be well justified for sufficiently early times. There exist four exact solutions in the single-stream limit, namely for one-dimensional ICs called the Zel’dovich approximation (ZA; Zel’dovich 1970), for quasi-one-dimensional ICs (Rampf & Frisch 2017), for spherical collapse (Peebles 1967), and for quasi-spherical ICs (Rampf 2019). Except for the spherical

* E-mail: michael.michaux@oca.eu

case, those solutions are valid until, and including, the instance of the first shell-crossing, where particle trajectories overlap, leading to the generation of (effective) vorticity and velocity dispersion through multistreaming (e.g. Pichon & Bernardeau 1999; Pueblas & Scoccimarro 2009; Hahn, Angulo & Abel 2015; Buehlmann & Hahn 2019). Most recently, attempts at pushing the theory beyond shell-crossing have received increasing interest (Colombi 2015; Taruya & Colombi 2017; McDonald & Vlah 2018; Pietroni 2018; Rampf, Frisch & Hahn 2019; Valageas 2020). For a broader overview over cosmological perturbation theory (PT), we refer the reader to the review by Bernardeau et al. (2002).

In general the non-linear evolution of the post-shell-crossing regime is currently most accurately modelled with N -body simulations of the Vlasov–Poisson equations. The transition, via ‘initial conditions’, from perturbation theory to the full non-linear but discretized simulation is, however, a rather delicate matter since it involves minimizing the errors in the approaches.

Initial conditions for simulations are usually provided by particle displacements and velocities from Lagrangian perturbation theory (LPT), truncated at some low order. This truncation, however, leads to so-called transients (Scoccimarro 1998; Crocce, Pueblas & Scoccimarro 2006); effectively this is a spurious decaying mode due to missing terms relative to the true (infinite-order) solution. One alternative to ameliorate the impact of transients is to consider higher order versions of LPT. In fact, LPT has been derived at increasingly higher order over the last decades; starting from first order by Zel’dovich (1970), Buchert & Goetz (1987), over second (Bouchet et al. 1992; Buchert & Ehlers 1993) and third order (Buchert 1994; Bouchet et al. 1995) in the 1990s, to fourth order by Rampf & Buchert (2012), Tatekawa (2014), and finally all-order recursion relations by Rampf (2012), Zheligovsky & Frisch (2014), and Matsubara (2015). Another alternative could be to generate initial conditions at earlier times since the linear growth amplitude of density fluctuations in LPT, D_+ , is the small perturbative parameter.

Unfortunately, starting numerical simulations at early times might add significant sources of numerical error. As is now well known, the N -body method is prone to discreteness effects due to the self-interaction of the discrete particle lattice (cf. Joyce et al. 2005; Joyce & Marcos 2007; Garrison et al. 2016), leading to a deviation from the fluid characteristics in the continuum limit. New tessellation methods reduce particle discreteness and thus could overcome these limitations (Hahn, Abel & Kaehler 2013; Hahn & Angulo 2016; Sousbie & Colombi 2016; Stücker et al. 2020b), albeit at increased computational cost. However, this might not be possible when simulations are optimized to simulate volumes as large as possible with the least possible computational resources. Recently, Garrison et al. (2016) have argued for an explicit linear-order correction during the course of a simulation, but this is not widely used. Furthermore, popular methods such as tree-based N -body (Barnes & Hut 1986) can suffer from large force errors at early times (where the density distribution is only slightly perturbed). In addition, these errors accumulate the more time-steps are made during the early (‘linear’) stages of a simulation when the density field is still close to homogeneous.

These problems would strongly suggest starting a simulation as late as possible. However, since standard perturbative approaches break down at shell-crossing, a competition exists between (i) late starts that reduce discreteness effects and numerical errors, but require higher-order LPT, and (ii) early starts that allow lower-order LPT but are more prone to discreteness errors. This dilemma is our main focus in this paper.

The impact of the order of the LPT, up to second order, and the starting time on properties of the non-linear density field have been studied already in quite some detail in the past literature (e.g. Crocce et al. 2006; Tatekawa & Mizuno 2007; L’Huillier, Park & Kim 2014; Garrison et al. 2016). Here, we extend these results to third-order LPT, which has been studied much less (see however Buchert, Melott & Weiss 1994; Tatekawa 2014, 2019), thereby allowing us to make for the first time more definite statements about the convergence radius and thus the perturbative regime accessible by LPT for cosmological ICs. In our analysis, we pay particular attention also to the impact of aforementioned discreteness effects on all results.

This paper is organized as follows. We begin with a brief review of LPT together with explicit solutions up to third order; see Section 2. We implement these LPT solutions numerically, which allows us to perform simple convergence tests, thereby essentially pinning down until which time-value LPT solutions to any order can be trusted. Details and results on this are provided in Section 3. Then, in Section 4, we discuss the details of our numerical simulations and analysis algorithms, whereas our results are given in Section 5. We conclude and summarize our results in Section 6.

Notation: Unless otherwise stated, all functions and spatial derivatives are w.r.t. Lagrangian coordinate \mathbf{q} . We use i, j, \dots for spatial indices, and summation over repeated indices is assumed. A comma ‘ $\Phi_{,i}$ ’ denotes a spatial partial derivative w.r.t. component q_i on Φ , while an overdot denotes a Lagrangian time derivative w.r.t. the cosmic-scale factor time a , the latter governed by the usual Friedmann equations.

2 LPT, INITIAL CONDITIONS, AND ITS NUMERICAL IMPLEMENTATION

In this section, we will discuss several theoretical and practical aspects of the creation of initial conditions for cosmological simulations. Specifically, in Section 2.1, we begin with a review of LPT solutions up to third order, and describe in Section 2.2 how these solutions are employed to create a consistent particle representation at a given starting time. We discuss the technical aspects of higher-order LPT implementations in Section 2.3. The effects of particle discreteness and initial arrangement are outlined and validated in Sections 2.4 and 2.5, respectively.

2.1 LPT results to third order

Let $\mathbf{q} \mapsto \mathbf{x}(\mathbf{q}, t) = \mathbf{q} + \boldsymbol{\psi}(\mathbf{q}, t)$ be the Lagrangian map from initial position \mathbf{q} to current (Eulerian) position \mathbf{x} at time t . The Lagrangian representation of the velocity is defined with $\mathbf{v} \equiv \dot{\mathbf{x}} = \dot{\boldsymbol{\psi}}$. In LPT, the displacement field, $\boldsymbol{\psi}$, is expanded as a power series in $D_+ = D_+(t)$, the linear growth of matter fluctuations in a Λ CDM universe, i.e.

$$\boldsymbol{\psi}(\mathbf{q}, t) = \sum_{n=1}^{\infty} \boldsymbol{\psi}^{(n)}(\mathbf{q}) D_+^n. \quad (1)$$

The truncation of the series at order n is commonly called n LPT, except for the first-order truncation that is called the Zel’dovich approximation (ZA; Zel’dovich 1970). Results to third order have been first derived by Buchert (1994), Catelan (1995), and Bouchet et al. (1995). Explicitly, the 3LPT solution for the displacement is

$$\boldsymbol{\psi}_{\text{3LPT}}(\mathbf{q}, t) = \boldsymbol{\psi}^{(1)}(\mathbf{q}) D_+ + \boldsymbol{\psi}^{(2)}(\mathbf{q}) D_+^2 + \boldsymbol{\psi}^{(3)}(\mathbf{q}) D_+^3, \quad (2)$$

with

$$\boldsymbol{\psi}^{(1)} = -\nabla \Phi^{(1)}, \quad (3)$$

$$\psi^{(2)} = -\frac{3}{7}\nabla\Phi^{(2)}, \quad (4)$$

$$\psi^{(3)} = \frac{1}{3}\nabla\Phi^{(3a)} - \frac{10}{21}\nabla\Phi^{(3b)} + \frac{1}{7}\nabla \times \mathbf{A}^{(3c)}, \quad (5)$$

which are expressed in terms of the purely spatial functions

$$\Phi^{(1)} = \varphi_{\text{ini}}, \quad (6)$$

$$\Phi^{(2)} = \frac{1}{2}\nabla^{-2} \left[\Phi_{,ii}^{(1)}\Phi_{,jj}^{(1)} - \Phi_{,ij}^{(1)}\Phi_{,ij}^{(1)} \right], \quad (7)$$

$$\Phi^{(3a)} = \nabla^{-2} \left[\det \Phi_{,ij}^{(1)} \right], \quad (8)$$

$$\Phi^{(3b)} = \frac{1}{2}\nabla^{-2} \left[\Phi_{,ii}^{(2)}\Phi_{,jj}^{(1)} - \Phi_{,ij}^{(2)}\Phi_{,ij}^{(1)} \right], \quad (9)$$

$$\mathbf{A}^{(3c)} = \nabla^{-2} \left[\nabla\Phi_{,i}^{(2)} \times \nabla\Phi_{,i}^{(1)} \right]. \quad (10)$$

where φ_{ini} is the gravitational potential φ at $a \rightarrow 0$; explicit instructions how φ_{ini} can be obtained are given in Section 2.2. For convenience, in Appendix A, we express these spatial functions suitably for numerical applications.

There are three simplifications in these LPT solutions worth highlighting. First, the LPT solutions have only one degree of freedom (provided by φ_{ini}), which might be surprising considering that the underlying equations are of second order in time. Secondly, we ignore decaying-mode solutions, which is an additional independent assumption. Thirdly, these LPT solutions assume an irrotational fluid motion in Eulerian coordinates; indeed only at the third order the displacement field loses its potential character, exemplified through the appearance of the vector $\mathbf{A}^{(3c)}$, which is actually required to maintain the zero-vorticity condition in Eulerian space.

Mathematically, these three simplifications arise from the use of the so-called slaved boundary conditions at $a = 0$ on the LPT solutions (cf. Brenier et al. 2003; Rampf, Villone & Frisch 2015). These conditions impose initial homogeneity $\delta_m \rightarrow 0$ where δ_m is the density contrast, furthermore guarantee that only one initial function needs to be provided, as well as select the purely growing-mode solutions with zero vorticity. In the next subsection we will argue that these simplifications are seemingly intertwined with the standard procedure of generating ICs for simulations.

2.2 Standard initial conditions for N -body simulations

Relativistic linear Boltzmann solvers such as CAMB (Lewis, Challinor & Lasenby 2000) or CLASS (Blas, Lesgourgues & Tram 2011) are used to evolve the coupled system of relativistic and non-relativistic species, including CDM, down to the present epoch. Cosmological simulations are commonly performed within the Newtonian approximation and only solve for the matter species; therefore, their ICs must take into account this change in the underlying physical problem. Usually, this is handled by taking the present-day linear matter density from Boltzmann solvers and applying a suitable rescaling procedure (see below). This rescaling procedure provides a fictitious universe at initialization time with today's radiation content, is however in full agreement with relativistic perturbation theory (cf. Chisari & Zaldarriaga 2011; Hahn & Paranjape 2016; Fidler et al. 2017).

Upon performing numerical simulations, N -body particles are sampled on an unperturbed lattice representing the initial homogeneity (in accordance with the slaving argument from the previous section). This initial placement of particles can be thought of

happening at time zero, i.e. at $a = 0$. Growing-mode initial conditions for simulations are then established by displacing the particles from $a = 0$ until the time a_{start} when the simulation is initialized, provided we have access to the initial gravitational potential φ_{ini} at time $a = 0$.

To get that initial gravitational potential, first observe that the gravitational potential is related to the matter density contrast through the Poisson equation $\nabla^2\varphi = \delta_m/a$, which can be written as

$$\varphi = \frac{\nabla^{-2}\delta_m}{a}. \quad (11)$$

From a linear Boltzmann code such as CLASS or CAMB, we can obtain the linear matter density δ_m^{code} , which for the rescaling procedure is assumed to be in the growing mode, i.e.

$$\delta_m^{\text{code}}(\mathbf{x}, a) = C_+(\mathbf{x}) D_+(a). \quad (12)$$

Evaluating (11) in the limit $a \rightarrow 0$ and expressing the matter density through (12), we find

$$\lim_{a \rightarrow 0} \varphi = \nabla^{-2} C_+ \lim_{a \rightarrow 0} \frac{D_+(a)}{a} \equiv \varphi_{\text{ini}}, \quad (13)$$

where we note that $D_+(a)$ is analytic around $a = 0$ and can be represented as $D_+ \propto a + \mathcal{O}(a^4)$ for a Λ CDM universe. In a final step, we express the spatial constant C_+ in terms of the matter density at the present time a_0 , and thus obtain

$$\varphi_{\text{ini}} = \frac{\nabla^{-2}\delta_m^{\text{code}}(\mathbf{x}, a_0)}{D_+(a_0)} \lim_{a \rightarrow 0} \frac{D_+(a)}{a}. \quad (14)$$

Operationally, to obtain a realization of φ_{ini} , we begin with a white noise field $W(\mathbf{x})$ which we multiply with the total matter density transfer function $T[\delta_m]$, an amplitude A that gives the correct normalization in terms of σ_8 , as well as the primordial fluctuations produced during inflation with amplitude $\sim k^{n_s/2}$. Altogether, in Fourier space for all $\|\mathbf{k}\| \neq 0$, one thus has

$$\tilde{\varphi}_{\text{ini}}(\mathbf{k}) = A \tilde{W}(\mathbf{k}) k^{(n_s-4)/2} \frac{T[\delta_m](k, a_0)}{D_+(a_0)} \lim_{a \rightarrow 0} \frac{D_+(a)}{a}, \quad (15)$$

and set $\tilde{\varphi}_{\text{ini}}(\mathbf{k} = \mathbf{0}) = 0$. Note that this is a random field, determined by $W(\mathbf{x})$ that creates the so-called cosmic variance whose properties are not identical to those of the ensemble average. Note that these deviations could be suppressed by the method proposed by Angulo & Pontzen (2016). However, since we will compare simulations using the same noise field, our results are already largely insensitive to cosmic variance.

2.3 Numerical implementation of LPT

The LPT potentials (6) needed for the 3LPT displacement can be conveniently computed in Fourier space from $\tilde{\varphi}_{\text{ini}}$ given in equation (15) by using the Fast Fourier Transform (FFT). However, some extra care should be taken for 2LPT and higher-order LPT terms: those terms contain quadratic and higher-order non-linearities, and thus will suffer from aliasing if multiplied numerically in real space.

Aliasing leads to non-linear modes appearing at the wrong wave numbers. For quadratic non-linearities, aliasing can be avoided by respecting Orszag's 3/2 rule (Orszag 1971): By temporarily enlarging the computational domain in Fourier space by a factor of 3/2 per dimension while carrying out the product, the aliased modes will be all located in the padding region, which can then simply be discarded. This of course increases the memory footprint by requiring two fields of $(3/2)^3 = 3.375$ times the size of the original fields in 3D.

Note that the $\Phi^{(3a)}$ potential contains a cubic non-linearity, which we treat by two quadratic convolutions. This is strictly speaking only




Figure 1. Power spectra of $\nabla^2\Phi^{(2)}$ (top panel), $\nabla^2\Phi^{(3a)}$ (middle panel), and $\nabla^2\Phi^{(3b)}$ (bottom panel) for the $1 h^{-1}\text{Gpc}$ (darker colour) and the $250 h^{-1}\text{Mpc}$ box (lighter colour) with 1024^3 modes/particles, averaged over 50 different random realizations. We show results obtained with $3/2$ padding for the de-aliasing of quadratic and cubic non-linearities (blue lines) and without de-aliasing (orange lines). Below each panel, we show the ratio of the aliased to the de-aliased spectrum revealing substantial biases with significant box size/resolution dependence in the 3LPT potentials.

an approximation, since for data with finite lengths, the split of a cubic convolution into quadratic convolutions is not associative, implying that some terms in the cubic convolution will be lost (e.g. Roberts 2011). The exact way to deal with the problem would be to evaluate cubic convolutions without such reductions, which however would require a zero padding twice as expensive per dimension compared to the $3/2$ rule of the quadratic case, thereby significantly increasing the memory footprint. We find that the split into quadratic convolutions causes only a ~ 5 per cent error in $\Phi^{(3a)}$, which is already subdominant compared to $\Phi^{(3b)}$. Of course, there are more memory efficient and potentially faster implementations possible based on the in-place de-aliased convolution technique of Bowman & Roberts (2011), than what we consider here.

In Fig. 1, we show the effect of aliasing on the power spectra of the 2LPT and 3LPT source terms $\nabla^2\Phi^{(2)}$, $\nabla^2\Phi^{(3a)}$, and $\nabla^2\Phi^{(3b)}$. We show results for two different box sizes, using 1024^3 modes in all cases, applying de-aliasing (blue lines) and ignoring it (red lines) for the power spectra of the ‘densities’ ($\nabla^2\Phi$) of various terms in the LPT expansion. We find that while the quadratic term (2LPT) has

aliasing only at the highest wave numbers, with a relative difference of up to 30–35 per cent near the particle Nyquist wavenumber $k_{\text{Ny}} \equiv \pi N_p^{1/3}/L_{\text{box}}$, the cubic terms (3LPT) are affected at all scales. We also find a very strong aliasing effect on the largest scales for the $(3a)$ term with a difference of nearly 3 orders of magnitude. It thus appears critical to perform de-aliasing in order to achieve a correct implementation of 3LPT. We further corroborate this aspect in Appendix D, where we demonstrate that only with de-aliasing, 3LPT shows a formal third-order convergence $\mathcal{O}(D_+^3)$, which is the expected behaviour from theory grounds.

However, even after de-aliasing, we find a weak dependence of the 3LPT terms on the chosen box size, indicated by a drop of the power close to the particle Nyquist wavenumber k_{Ny} . This clearly indicates that at 3LPT we see a dependence on the UV truncation of the perturbation spectrum due to the finite resolution employed in computing the terms – computing by de-aliased discrete Fourier transform truncates at k_{Ny} . The coupling of modes with $k > k_{\text{Ny}}$ to the modes that can be numerically represented, i.e. $k \leq k_{\text{Ny}}$, is thus missing. This effect is also present at 2LPT, but is significantly smaller. Such UV sensitivities are well known in perturbation theory, especially within the context of loop integrations (Bernardeau et al. 2002), and have been discussed also in the context of the GRIDSPT method of Taruya, Nishimichi & Jeong (2018). They are not easy to circumvent in a numerical setting since it is always expensive to increase k_{Ny} .

We found however that, at least up to 3LPT, errors due to aliasing (and through UV sensitivities) are well below 1 per cent in the evolved simulations at low redshift. For more details, see the results in Appendix B. It is thus likely unproblematic to avoid de-aliasing up to 3LPT and thereby to save memory and computing time if one is only interested in the low-redshift results of a non-linear simulation. In this paper, we carry out all simulations with fully de-aliased 2LPT and 3LPT terms. A more detailed investigation of the impact of aliasing at higher resolution, and whether including the contribution of modes $k > k_{\text{Ny}}$ to the 3LPT terms could improve convergence of simulations with different mass resolution are interesting questions for future investigations.

Numerical implementations of 2LPT commonly used in the community include the 2LPTIC software package¹ introduced in Crocce et al. (2006), which is based on single resolution FFTs. For multi-resolution zoom simulations, we are aware of the implementation by Jenkins (2010), which uses a Tree-PM approach to evaluate the 2LPT Poisson equation at higher resolution in the zoom region, and the implementation by Hahn & Abel (2011) in MUSIC,² which uses a combined algebraic multigrid and FFT approach. Version 4 of the PINOCCHIO code (Munari et al. 2017) implements the longitudinal part of 3LPT to determine the large-scale clustering of haloes in the rapid mock catalogue PINOCCHIO scheme (Monaco, Theuns & Taffoni 2002). To our knowledge, there exists no publicly available implementation of 3LPT N -body initial conditions that includes transversal modes and performs a correct de-aliasing of higher-order terms, and which could therefore allow an accurate assessment of truncation errors and transients.

We make our implementation publicly available as the MUSIC2-MONOFONIC software package³ that comprises a distributed memory parallelized (MPI + threads) implementation of all algorithms discussed here. It is, however, currently restricted to single resolution

¹ Available from <https://cosmo.nyu.edu/roman/2LPT/>

² Available from <https://bitbucket.org/ohahn/music>

³ Available from <https://bitbucket.org/ohahn/monofonic>




Figure 2. Top (bottom) panel: ratio between power spectra from N -body simulations and from 3LPT at $z = 24$ (11.5). The N -body simulations were evolved from $z_{\text{start}} = 49$ (24) down to $z = 24$ (11.5) using initial conditions based on ZA, 2LPT and 3LPT (blue, green, and orange, respectively) and initialized from different perturbed lattices. Solid lines show the results of the SC lattice initial conditions including corrections from PLT (see Section 2.4, to correct for particle discreteness), and the dotted lines without. Dashed lines show the result of the particle oversampling using an FCC lattice without PLT correction. The vertical dotted line indicates the particle Nyquist wavenumber of the initial conditions. One per cent agreement is represented as a shaded area and a perfect agreement as a dotted horizontal line. With PLT the simulations agree at per cent level with 3LPT up to the particle Nyquist wavenumber when 2LPT or 3LPT is used.

(mono-grid) simulations. MUSIC2 (Hahn et al., in preparation) will be the next update to the MUSIC software package of Hahn & Abel (2011), which will at a later stage also support zoom simulations.

2.4 Initial particle placements and particle linear theory

As mentioned earlier, the creation of ICs require as input a homogeneous and isotropic particle distribution. In practice, this state is commonly realized by a simple cubic (SC) lattice, where particles are arranged on a regular grid. Note that other alternatives for homogeneous non-regular particle distributions have been proposed (e.g. White 1996; Hansen et al. 2007; Liao 2018, for ‘glass’, ‘quaquaversal’, and ‘CCVT’ particle distributions, respectively).

The case of an SC lattice is particularly convenient since one can simply have one particle per Fourier mode so that the SC lattice coincides with the FFT mesh used to compute the LPT terms. An important drawback is that Fourier modes on such lattices do not grow identically as in the continuum fluid limit owing to self-interactions of the discrete lattice, as has been pointed out in a series of papers (Joyce et al. 2005; Marcos et al. 2006; Joyce & Marcos 2007).

These discreteness effects are present in all particle distributions (but harder to quantify in non-regular lattices). This is a consequence of the gravitational softening length being smaller than the mean inter

particle separation (see also Angulo, Hahn & Abel 2013, where the effect is much more dramatic in the case of multiple particle species). The effect is, of course, strongest at early times when the lattice is still close to regular and physical density fluctuations are still small. Joyce et al. (2005) and Marcos (2008) have shown how the deviation from the fluid case can be calculated for early times in ‘particle linear theory’ (PLT). Recently, this work has been extended by Garrison et al. (2016) who use the earlier solutions to compensate the initial particle displacements and velocities by the expected discreteness of the lattice at linear order and for early times.

To quantify the impact of particle discreteness, we have implemented an *optional* PLT correction in our initial conditions. Note that our version is similar, but deviates in some respects from that of Garrison et al. (2016) (e.g. avoiding their artificial boosting of the counter-PLT modes). In Appendix C, we provide details on our implementation. None the less, we remark that the PLT correction is a decaying mode and thus vanishes over the course of the simulation. In principle, this can be avoided by artificially boosting PLT corrections, as proposed by Garrison et al. 2016, or by repeatedly applying it over time – but note that PLT is only valid at linear order and currently there is no higher-order theory available, which limits its applicability at late times.

Our preferred choice for most parts of this paper – instead of PLT – is to use a face-centred-cubic (FCC) lattice constructed by shifting four SC lattices,⁴ and imposing the displacement and velocity using a corresponding shift of the respective fields using a Fourier shift with the FFT. The FCC lattice is more isotropic than the SC lattice, with four times as many particles as our standard simulations. In fact, Marcos (2008) showed that Bravais lattices other than SC exhibit weaker deviations from the fluid limit at a fixed particle number, due to higher symmetries (cf. also Stücker et al. 2020a).

2.5 Validation

In this subsection, we present a validation of our numerical and analytic tools by comparing perturbative and numerical solutions, emphasizing the role of particle discreteness.

In Fig. 2, we show the ratios of various numerical predictions for the non-linear power spectrum to that in LPT. Specifically, we employ an N -body simulation with $N_p = 1024^3$ particles in a $L_{\text{box}} = 1h^{-1}\text{Gpc}$ box, and compute the LPT fields on a grid with 1024^3 points. Coloured lines display the case where the initial conditions were computed using third-, second-, or first-order LPT. Solid and dotted lines indicate cases where PLT corrections have been included or not, while dashed lines refer to using an FCC lattice instead of PLT. In the top panel, we start our simulations at $z = 49$ and evolve them until $z = 24$, whereas in the bottom panel we start at $z = 24$ and evolve until $z = 11.5$.

As is evident from the figure, without PLT corrections, there is a clear discrepancy between perturbative and numerical solutions, with the latter showing a significant loss of power at the particle Nyquist wavenumber of the particle lattice, $k_{\text{Ny}} \equiv \pi N_p^{1/3}/L_{\text{box}}$. The loss of power is less prominent when using an FCC lattice. Note that this power suppression is almost insensitive to the order of LPT used to generate the ICs, which could create the illusion of proper convergence in simulation results. This is, however, convergence to the discrete solution and not to the fluid solution.

⁴We have also considered body-centred-cubic lattices (half the particle load compared to FCC), leading to results between SC and FCC lattices.




Figure 3. Same as Fig. 2, but now showing the ratio between the respective equilateral bispectra at $z = 24$ (11.5) in the top (bottom) panel. Only the simulations generated with 3LPT and with PLT corrections lead to ~ 1 per cent agreement between simulations and 3LPT at all k below the particle Nyquist wavenumber, while 2LPT undershoots already at much smaller k . ZA is not even present on this figure because it is off at all scales (ratio between 0.8 and 0.5).

In contrast, there is a remarkable agreement between the N -body and perturbative solutions when PLT is enabled. Specifically, there is a 1 per cent agreement between 2LPT, 3LPT, and the numerical simulation up to k_{Ny} (indicated by vertical dotted line). Note that only with PLT corrections, one is able to see the improvement brought by higher order ICs, but even ZA performs well in this test if initialized early enough ($z_{\text{start}} = 49$).

In Fig. 3, we show the same figure as before (Fig. 2) but now for the equilateral bispectrum [i.e. with $k = k_1 = k_2 = k_3$; see equation (31) for the used bispectrum definition]. As for the power spectrum, we find a systematic power deficit in the numerical solutions when PLT is not considered, and an apparent convergence to the wrong solution. Note that also here the loss of power is less pronounced when using an FCC lattice.

With PLT corrections, a good agreement between numerical and perturbative bispectra is achieved, but only for the higher-order version of the ICs. Unlike for the power spectrum, ZA ICs perform very poorly in this test. On large scales, they underestimate the bispectrum at about 20 per cent, and up to 50 per cent on small scales, regardless of PLT corrections. Thus, it is not shown in the figure. This suggests that ZA ICs have a leading-order transient in the bispectrum and should *not* be used to set up simulations if high precision is required.

Note that for 3LPT with PLT correction, the error between our two redshift ranges changes sign, overshooting in the redshift range $z = 24 \rightarrow 11.1$ (lower panel in Fig. 2) while undershooting for $z = 49 \rightarrow 24$ (upper panel in Fig. 2). From a theoretical point of view, the next-to-leading (one-loop) correction to the bispectrum involves density correlators up to fourth order (the ‘411’ contribution)

which should be of the same magnitude as those that we implicitly determine on the grid points using 3LPT ICs. Thus, 3LPT might not be sufficiently converged for predicting the non-linear bispectrum at the few per cent level, and it is likely that 4LPT might be required for an accurate comparison with simulations.⁵ We can then speculate that this overshooting might imply that a starting redshift of 11.5 is already too late for 3LPT. We remark that 2LPT at lower redshift becomes worse than 3LPT without PLT correction or with FCC oversampling, even if the sign does not change, thus suggesting that 2LPT suffers more from the lack of some higher-order correction for the bispectrum. The impact of 4LPT ICs on the bispectrum will be investigated in future works.

From the previous plots, the impact of particle discreteness is obvious. Unfortunately, PLT corrections work for a relatively short period if not artificially boosted (as Garrison et al. 2016 propose) and/or multiply applied at various stages. Furthermore, the perturbative estimates to the discrete lattice effects have only been computed to first order. For this reason, our reference simulation will not be with the PLT corrections switched on. Instead, we use an FCC lattice as described at the end of Section 2.4.

Since the Fourier perturbation modes are initially specified on the SC lattice, we oversample the Fourier modes and phase-shift (in Fourier space) the n LPT displacements to obtain them in the FCC case. This enables us to compare how the exact same mode spectrum converges on various statistics as we change the particle number and degree of isotropy of the underlying particle lattice on to which these modes are imposed. In particular, the SC lattice samples 1:1 the Fourier modes, meaning that we have one particle per mode in an SC simulation, and perturbations at the particle Nyquist wavenumber are basically represented by two particles. By oversampling, no new perturbation modes are added, which is in contrast to usual convergence studies carried out for cosmological simulations, where with an increase in the particle number also new perturbations are added. When we compare an SC against an FCC simulation, we thus test convergence at fixed initial modes, investigating only the impact of particle sampling. Keeping the modes fixed is especially important since, as we will show in Section 3, increasing the particle Nyquist mode of sampled fluctuations increases the variance of the fluctuations realized in the box, and in turn influences the degree of non-linearity at the starting time. In contrast, allowing the ICs to change would render an accurate comparison of starting redshifts impossible. A somewhat orthogonal test of the impact of particle discreteness comparing the different lattices, as well as also ‘glass’ (White 1996) and other (e.g. Hansen et al. 2007; Liao 2018) pre-initial conditions, at fixed particle number would be a very interesting project for a future study. We found however that glass pre-initial conditions (results not shown in this paper) did not improve the convergence of power and bispectrum over those presented for the perturbed SC lattice below in Section 5. This shows that indeed the small-scale interactions between N -body particles are responsible for the discreteness errors, rather than the particular anisotropies of a given pre-initial conditions particle distribution.

In summary, the results presented above clearly indicate that, without discreteness corrections, N -body simulations deviate strongly from the fluid limit during the perturbative phase. This effect is larger, the earlier the simulation is initialized. Therefore, to minimize these errors, one should delay the initialization time to the latest possible

⁵This should be contrasted to the convergence studies related to the power spectrum (Fig. 2) where 4LPT becomes relevant only at 2-loop and beyond.

moment. In the following section, we will formally investigate what this ‘latest possible moment’ is in the context of LPT.

3 CONVERGENCE RADIUS OF LPT AND STARTING TIME FOR SIMULATIONS

In the previous section, we argued that early starting redshifts for simulations are accompanied by significant discreteness effects. These effects decrease in magnitude at lower redshifts, thus the latest possible starting redshift is desirable. On the other hand, LPT ICs are based on a single-stream fluid description implying that once particle trajectories cross for the first time (‘shell-crossing’), the single-stream fluid equations and thus LPT become invalid.

In this section, we will discuss several aspects regarding determining the point at which LPT breaks down, which assist clarifying the appropriate redshift-window for generating initial conditions. Necessary definitions related to the LPT series are provided in Section 3.1, while we outline numerical tests for estimating the convergence radius in Section 3.2. Fairly complementary to those numerical tests, there are also ways to estimate the convergence radius directly from theory; the main ideas are sketched in Section 3.3 while further technical details are provided in Appendix E.

3.1 LPT series and its convergence radius

The breakdown of LPT is intimately linked to finding the radius of convergence of the Taylor series of the displacement

$$\psi = \sum_{n=1}^{\infty} \psi^{(n)}(\mathbf{q}) D_+^n. \quad (16)$$

Indeed, as it is known from complex analysis, the radius of convergence, R_{conv} , of any Taylor series is limited by the nearest singularity in the complex domain of its argument. For the LPT series, which is a time-Taylor series with time variable D_+ , the nearest singularity could be in the real domain of D_+ but could also take complex values.⁶ For example, a real singularity could appear at shell-crossing – where the density becomes formally infinite – which in any case marks the breakdown of LPT.

Regardless of the precise nature of the singularities, a single-time push-forward displacement of particles (as done for the IC creation) is only meaningful mathematically as long as the chosen time-step is within the disc of convergence spanned by R_{conv} (cf. Rampf et al. 2015):

$$|D_+(a_{\text{max}})| < R_{\text{conv}}. \quad (17)$$

Hence, a_{max} sets formally the maximal scale factor for using LPT initial conditions for a simulation. Note that typically, N -body simulations adopt an heuristic criterion requiring the average amplitude

⁶To illustrate the argument of singularities, let us consider a toy example. Let us assume that the exact, non-perturbative displacement is $\psi(D_+) = 1/(1 + D_+^2)$, which has complex singularities at $D_+ = \pm i$. In LPT, we represent this displacement by a Taylor series, i.e. $\psi = \sum_n c_n D_+^n$, but because of the appearance of these two complex singularities, we can evolve particles only for $0 < |D_+| < 1$, i.e. within the disc of convergence. If shell-crossing has not occurred until $|D_+| = 1$, then we should seek for an analytic continuation technique, like the one of Weierstraß, that in our case amounts to evolve particles until $|D_+| = D_< < 1$, then re-expand and determine the new Taylor coefficients around $D_<$. The series around $D_<$ will generally have a new radius of convergence that allow us to continue and follow the particles for $|D_+| > D_<$, possibly involving repetitive re-expansions, until a real singularity in time and/or shell-crossing, occurs.

of fluctuations σ at the resolution scale to be small, i.e. $\sigma \ll 1$. While such empirical criteria are certainly useful, we suggest here to take also the breakdown of LPT into account.

In the following subsection, we provide two complementary methods to estimate R_{conv} : the numerical ratio test, see Section 3.2, and a fully analytical method that exploits a theoretical lower bound on R_{conv} , as outlined in Section 3.3. These estimates translate into threshold for the latest possible initialization time of the simulation (based on theory grounds), which is discussed in Section 3.4.

3.2 Numerical estimation of the radius of convergence

A particularly simple method for estimating R_{conv} is to consider the L^2 norm of the displacement (16), i.e.

$$\Psi \equiv \sum_{n=1}^{\infty} \|\psi^{(n)}(\mathbf{q})\| D_+^n, \quad (18)$$

and perform the convergence tests for that series (cf. Podvigina, Zheligovsky & Frisch 2016). For this, we use the ratio test which states that the radius of convergence R_{conv} of the series is

$$\frac{1}{R_{\text{conv}}} = \lim_{n \rightarrow \infty} \|\psi^{(n)}\| / \|\psi^{(n-1)}\| \quad (19)$$

(if that limit exists). In all generality, it is thus the large- n limit of Taylor coefficients that decides questions about convergence. Of course, in numerical implementations of the ratio test, the actual limit $n \rightarrow \infty$ can only be reached by employing numerical extrapolation methods (see next paragraph), which in principle can be performed to very high accuracy. None the less, we remark that in the present case where we have numerically implemented perturbative solutions only up to third order, the resulting numerical extrapolation is fairly crude (see the following section for a more rigorous yet more restrictive method).

Domb & Sykes (1957) introduced a numerical extrapolation method in a non-cosmological context. There, one draws the plot $\|\psi^{(n)}\|/\|\psi^{(n-1)}\|$ versus $1/n$, and takes the y-intercept (‘ $n \rightarrow \infty$ ’) as the estimate for $1/R_{\text{conv}}$. For solutions until third order, we have two such ratios, leading to the two tuples

$$\left(1/2, \|\psi^{(2)}\|/\|\psi^{(1)}\|\right) \quad \text{and} \quad \left(1/3, \|\psi^{(3)}\|/\|\psi^{(2)}\|\right), \quad (20)$$

from which we can perform a simple linear extrapolation to the y-intercept. This argument leads to the following estimate for the radius of convergence of the LPT series,

$$\frac{1}{R_{\text{conv}}(\mathbf{q})} \approx 3\|\psi^{(3)}\|/\|\psi^{(2)}\| - 2\|\psi^{(2)}\|/\|\psi^{(1)}\|. \quad (21)$$

Note the somewhat artificial dependence of \mathbf{q} in R_{conv} which we have included here by hand; the actual radius of convergence of LPT is obtained by searching for the global minimum $R_{\text{actual}} = \min_{\mathbf{q}} R_{\text{conv}}(\mathbf{q})$, which is, so to say, the worst-case scenario over the whole spatial domain (in the present case: in the simulation box). However, since we employ random initial conditions, we expect some points in the realization of this probability distribution to have extreme values. For example, if $\|\psi^{(2)}\| \simeq 0$ locally (indicating local pancake formation) but then with $\|\psi^{(3)}\| = \mathcal{O}(1)$ (non-negligible gravitational couplings from small spatial scales), then our low-order convergence test will predict a tiny radius of convergence. We stress however, that in such collapse scenarios higher-order ratios are likely to change the outcome of the convergence test significantly. We will come back to this issue in a forthcoming paper.

Therefore, we expect to observe outlier values that should be disregarded. For that reason, we choose for the ratio test a statistical




Figure 4. Maximum scale factor $a_{\max,p}$ for which various methods predict convergence of LPT, thereby pinning down the latest possible time when simulations should be initialized. Specifically, we show $a_{\max,p}$ as a function of the fluctuation scale σ for different box sizes L_{box} (colors), averaged over 10 random realizations. The solid lines are the result of the numerical ratio test using our implementation of 3LPT, the dashed and dotted lines are analytical bounds on the radius of convergence computed using (23) for $T = T_{\text{cons}} = 0.022$ (dotted) and $T = T_{\text{cross}} = 0.107$ (dashed). The values of p denote the percentage of particle displacements in the simulation that satisfy the LPT convergence condition (17). We chose here the percentiles corresponding to the 1σ , 2σ , and 3σ (light to bright) standard deviation of the normal distribution. For the theoretical bounds, we use $p = 99.9$ per cent. Our simulations (box length $1h^{-1}$ Gpc) are indicated in the figure with a square symbol for $z_{\text{start}} = 49$, a diamond for $z_{\text{start}} = 24$ and a triangle for $z_{\text{start}} = 11.5$.

approach as outlined next. See the following section for a method that is only sensitive to the deterministic radius of convergence.

Inverting $|D_+(a_{\max})| \lesssim R_{\text{conv}}$ gives the upper limit a_{\max} (or the lower limit z_{min}) on the starting time of the simulation, i.e. $a_{\text{start}} \lesssim a_{\max}$. We can compute $a_{\max}(\mathbf{q})$ for all \mathbf{q} in the simulation box and get the value of $a_{\max,p}$ at some chosen percentile p . This means that by taking $a_{\text{start}} = a_{\max,p}$, p per cent of the points will fulfil the condition $|D_+(a_{\max})| \lesssim R_{\text{conv}}$. Results of the ratio test are shown in Fig. 4 and discussed in Section 3.4.

3.3 Analytical bound on the radius of convergence

Complementary to the above numerical method, Zheligovsky & Frisch (2014) and Rampf et al. (2015) have shown that, being equipped with explicit all-order recursion relations for LPT, it is also possible to obtain an analytical bound D_+^{theory} on the radius of convergence R_{conv} , i.e.

$$|D_+^{\text{theory}}| < R_{\text{conv}}. \quad (22)$$

This bound comes actually from ‘order ∞ ’ in LPT (see Appendix E), but at the same time, since only very weak assumptions on the initial conditions are imposed, leads to a bound that is (much) smaller than R_{conv} . None the less, this analytical bound is to some extent complementary to the low-order estimate as employed in the previous section, and therefore included in our studies. The numerical

exploitation of all this and the ratio test will be done in the following section.

In Appendix E, we outline three ways how theoretical bounds on the radius of convergence can be estimated. For this work, we use the two methods that lead to the best bounds, namely case (a) where the employed bounds are conservative (‘cons’) but already slightly improved in comparison to those of Rampf et al. (2015), and (b) a new – however very optimistic – method where we ignore complex time singularities and, possibly, can increase the bound until shell-crossing (‘cross’). We find

$$D_+^{\text{theory}} = \frac{T}{\|\nabla_i \nabla_j \varphi_{\text{ini}}\|}, \quad (23)$$

where in case (a) we have $T = T_{\text{cons}} = 0.022$ and in case (b) we have $T = T_{\text{cross}} = 0.107$. In Fig. 4, we show how these theoretical bounds relate to constraints for the allowed scale factor where simulations can be initialized. As expected, these theoretical bounds are much weaker than the one coming from the ratio test. We note that for the theoretical bounds we have fixed $p = 99.9$ per cent in Fig. 4 to minimize an unwanted dependence on the box size.

3.4 Upper limits on initialization time for simulations

The analytical bound as well as the complementary numerical ratio test translate into upper limits on the initialization time $a_{\text{start}} \lesssim a_{\max}$ for the simulations. In Fig. 4, we summarize the resulting upper limits as a function of the fluctuation scale σ . Evidently, the numerical tests provide much stronger bounds on a_{\max} than the analytical method, however we remark again that the employed numerical tests only go to third order in PT – which only provides two data points for the numerical extrapolation. Thus, these stronger bounds, which could change at higher orders, should still be read with caution. The analytical method, by contrast, provides bounds that are rock-solid.

For sufficiently small σ , the results from the numerical ratio test are approximately straight lines using log–log scaling. Using this as a working assumption, we can give some simple formula that could help determine quickly whether a given starting time is inside or outside the radius of convergence, and to choose the optimal one for a simulation.

For the ratio test, the maximum starting time is well approximated by

$$a_{\max,n} \approx \frac{0.2}{n} \sigma^{-0.8}, \quad (24)$$

where σ is the standard deviation, and n is the number of standard deviations required to be inside the convergence radius (since the distribution is not Gaussian, we map this to the corresponding percentiles, i.e. $p = 68, 95$, and 99.7 for $n = 1, 2$, and 3 , respectively). We remark that the value of σ can be calculated on the random realization prior to the back-scaling step and thus, the initialization time for the simulation can be chosen at the last moment from equation (24) by specifying the number of standard deviations n .

In contrast, the theoretical bounds from the previous section are much closer to

$$a_{\max} \approx \frac{T}{12.2\sigma}, \quad (25)$$

where T is either T_{cons} or T_{cross} as defined in the previous section. The scaling $\propto \sigma^{-1}$ corresponds to what one would expect from linear theory structure growth in an Einstein–de Sitter (EdS) universe, such that these bounds basically require that the fluctuations do not exceed a fixed amplitude at the starting time.

Table 1. Force and time integration accuracy parameters of L-GADGET3 employed for the simulations.

Softening	0.016276
ErrTolForceAcc	0.002
ErrTolIntAccuracy	0.025
ErrTolTheta	0.5
MaxRMSDisplacementFac	0.25
MaxSizeTimestep	0.01
TypeOfOpeningCriterion	1

The theoretical bounds are thus much more conservative than the numerical estimates we obtained from the ratio test: the former suggest starting times $z_{\text{start}} \gtrsim 100$, while the latter suggest this could be a factor $\gtrsim 5$ lower, provided that we exclude certain outliers (due to the employed low-order numerical test and the nature of random ICs). How these numbers translate into a given precision for the summary statistics at late times is a non-trivial question due to the increased impact of discreteness errors for early starts that we have discussed before. We will attempt to disentangle the two in the next sections.

4 NON-LINEAR SIMULATIONS AND ANALYSIS

Here, we first describe our cosmological numerical simulations (Section 4.1) and then the methods we employ to quantify the properties of the respective non-linear density fields (Section 4.2).

4.1 Numerical evolution

For all numerical results, we use the L-GADGET3 code (Angulo et al. 2012) which is a heavily modified and optimized version of the GADGET-2 N -body code (Springel 2005). It implements a parallel tree-PM method with periodic boundary conditions. For all simulations, we employ a 2048^3 PM grid, with the force and time integration parameters listed in Table 1. The values of those parameters are expected to yield sub-per cent convergence in the low-redshift power spectrum on scales $k \lesssim 2h \text{ Mpc}^{-1}$ (see fig. 2 of Angulo et al. 2020).

We use the following cosmological parameters, consistent with the PLANCK2018+LSS results (Planck Collaboration 2020): $\Omega_m = 0.3111$, $\Omega_\Lambda = 0.6889$, $\Omega_b = 0.04897$, $h = 0.6766$, $\sigma_8 = 0.8102$, and $n_s = 0.9665$. The transfer function has been calculated using the CLASS code⁷ (Blas et al. 2011) at $z = 0$, and then scaled back using the linear theory growth factor $D_+(z)$ to the various starting redshifts we use (see discussion in Section 2.2).

For the simulations in this study, we use a box of linear size $L_{\text{box}} = 1h^{-1} \text{ Gpc}$, and use $N_{\text{part}} = 1024^3$ whenever we initialize particles on a simple cubic (SC) lattice. This yields an N -body particle mass of $m_{\text{part}} = 8.04 \times 10^{10} h^{-1} M_\odot$. In some cases, we oversample the white noise fields with different lattices (see discussion in Section 2.4). When an FCC lattice is used, the number of particles is 4×1024^3 and the mass $2.01 \times 10^{10} h^{-1} M_\odot$. For each of these simulations, we have carried out versions employing de-aliased (c.f. Section 2.3) ZA, 2LPT, and 3LPT initial conditions. For testing, we have also run versions with 2LPT and 3LPT ICs without dealiasing (see Appendix B). In Table 2, we list all simulations we use in our paper.

Table 2. List of the simulations we use in this paper, differing in the order of the LPT used to initialize them as well as the starting time. All simulate the same $L_{\text{box}} = 1 h^{-1} \text{ Gpc}$ box with the cosmology described in Section 4.1. The ‘SC’ runs have 1024^3 particles in an SC lattice, the ‘FCC’ runs 4×1024^3 particles in an FCC lattice. All simulations sample the same 1024^3 modes.

PT	z_{start}	Lattice	$m_p/10^{10} h^{-1} M_\odot$
ZA	49	SC	8.04
ZA	99	SC	8.04
ZA	199	SC	8.04
2LPT	11.5	SC	8.04
2LPT	24	SC	8.04
2LPT	49	SC	8.04
2LPT	99	SC	8.04
2LPT	199	SC	8.04
3LPT	11.5	SC	8.04
3LPT	11.5	FCC	2.01
3LPT	24	SC	8.04
3LPT	24	FCC	2.01
3LPT	49	SC	8.04
3LPT	49	FCC	2.01

4.2 Statistical analysis of density field data

We briefly summarize below the various quantities that we will investigate in Section 5 along with their respective definitions and the way we computed/evaluated them. Throughout, we use cloud-in-cell interpolation (CIC; cf. Hockney & Eastwood 1981) to deposit the N -body particles to a regular grid. From this density field, we compute one-, two-, and three-point statistics.

4.2.1 Density field and one-point statistics

To study the distribution function of densities, we use a grid of 1024^3 cells and consider smoothed versions of this CIC (over)density field, δ_{R_s} , by multiplying it with the Fourier transform of the spherical top-hat filter of radius R_s in Fourier space. From this smoothed density field we compute the third and fourth cumulant statistics, also called skewness S_3 and kurtosis S_4 , and defined by Bernardeau et al. (2002), e.g. as

$$S_3(R_s) = \langle \delta_{R_s}^3 \rangle_c / \langle \delta_{R_s}^2 \rangle^2, \quad S_4(R_s) = \langle \delta_{R_s}^4 \rangle_c / \langle \delta_{R_s}^2 \rangle^3, \quad (26)$$

where the angle brackets indicate volume averages, and the connected part of the third and fourth moments are defined by:

$$\langle \delta_{R_s}^3 \rangle_c = \langle \delta_{R_s}^3 \rangle - 3 \langle \delta_{R_s}^2 \rangle \langle \delta_{R_s} \rangle + 2 \langle \delta_{R_s} \rangle^3, \quad (27)$$

$$\langle \delta_{R_s}^4 \rangle_c = \langle \delta_{R_s}^4 \rangle - 4 \langle \delta_{R_s}^3 \rangle \langle \delta_{R_s} \rangle - 3 \langle \delta_{R_s}^2 \rangle^2 + 12 \langle \delta_{R_s}^2 \rangle \langle \delta_{R_s} \rangle^2 - 6 \langle \delta_{R_s} \rangle^4. \quad (28)$$

Since we only use one realization for each simulation, a bootstrapping method was used to estimate the mean and variance of these quantities. Specifically, from the density fields $\delta_{R_s, \text{num}}$ and $\delta_{R_s, \text{den}}$ appearing in the numerator and the denominator of the ratios in equation (26), we created an ensemble of new fields $\delta_{R_s, \text{num}}^i$ and $\delta_{R_s, \text{den}}^i$ with the same number of elements as the original. These elements were chosen randomly with replacement, meaning that a particular element on the original sample might not be present in the resampling, and that other elements can appear more than once. For each i , the random resampling employed the same pseudo-random generator seed for the numerator and denominator of the

⁷Available from <http://class-code.net/>

ratio. The ratios are then estimated using the means of the one point statistics on this ensemble of resamplings. We used an ensemble of 20 resamplings.

4.2.2 Power spectra and bispectra

We define the power spectrum $P(k)$ as

$$\langle \delta(\mathbf{k}) \delta(\mathbf{k}') \rangle = (2\pi)^3 \delta_D^{(3)}(\mathbf{k} + \mathbf{k}') P(k), \quad (29)$$

where $k = \|\mathbf{k}\|$ and $\delta_D^{(3)}$ is the Dirac delta, itself defined by

$$\delta_D^{(3)}(\mathbf{x}) = \int \frac{d^3k}{(2\pi)^3} e^{-i\mathbf{k}\cdot\mathbf{x}}. \quad (30)$$

The power spectra are computed on-the-fly during the simulation based on an FFT of the PM grid of 2048^3 cells with CIC assigned particles. To reduce the effect of the mass assignment, we then divide our measurements by the Fourier transform of a real-space CIC kernel (Jing 2005). We have compared our results with a measurement performed after folding the density field 64 times in each direction. This folding procedure allows to shift the range of sampled wave numbers to smaller scales (see Jenkins et al. 1998, who first proposed this method). From this, we verified that our power spectra results are not influenced by the FFT grid resolution at any scale considered here.

Analogously, the bispectrum $B(k_1, k_2, k_3)$ is defined by

$$\langle \delta(\mathbf{k}_1) \delta(\mathbf{k}_2) \delta(\mathbf{k}_3) \rangle_c = (2\pi)^3 \delta_D^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3). \quad (31)$$

We use the PYTHON package BSKIT⁸ (Foreman et al. 2020) to compute the bispectra presented in this paper. This software package is based on the NBODYKIT⁹ toolkit of Hand et al. (2018) and implements a parallel version of the ‘Scoccimarro estimator’ (cf. Scoccimarro 2000; Sefusatti et al. 2016; Tomlinson, Jeong & Kim 2019).

We recall here the basic ideas behind this bispectrum estimator. The brute-force method integrates the triple-product $\delta(\mathbf{q}_1) \delta(\mathbf{q}_2) \delta(\mathbf{q}_3)$ over spherical shells of width Δk around k_1 , k_2 , and k_3 :

$$\hat{B}(k_1, k_2, k_3) = \frac{1}{(2\pi)^3 V_\Delta} \int_{k_1} d^3q_1 \int_{k_2} d^3q_2 \int_{k_3} d^3q_3 \times \delta(\mathbf{q}_1) \delta(\mathbf{q}_2) \delta(\mathbf{q}_3) \delta_D^{(3)}(\mathbf{q}_1 + \mathbf{q}_2 + \mathbf{q}_3), \quad (32)$$

where

$$V_\Delta = \int_{k_1} d^3q_1 \int_{k_2} d^3q_2 \int_{k_3} d^3q_3 \delta_D^{(3)}(\mathbf{q}_1 + \mathbf{q}_2 + \mathbf{q}_3), \quad (33)$$

where the short-hand notation $\int_{k_i} d^3q_i$ denotes integration while applying a filter with a window function $W(\mathbf{q}_i)$ that is 1 inside the spherical shell around \mathbf{k} and 0 elsewhere. Inserting equation (30) in (32) and using Fubini’s theorem to reorder the terms, we get

$$\hat{B}(k_1, k_2, k_3) = \frac{1}{(2\pi)^3 V_\Delta} \int \frac{d^3x}{(2\pi)^3} I_{k_1}(x) I_{k_2}(x) I_{k_3}(x), \quad (34)$$

where

$$I_{k_j}(x) = \int d^3q_j W_j(\mathbf{q}_j) \delta(\mathbf{q}_j) e^{-i\mathbf{x}\cdot\mathbf{q}_j}. \quad (35)$$

Written as in equation (34), the bispectrum calculation involves three inverse Fourier transforms, their product and an integration over the real spatial domain. While the complexity is the same as a brute-force averaging over every triangle, the calculations are much more

optimizable using cache efficient and parallel computations, and exploit advantages of FFT algorithms. The volume factor (33) is calculated using the same method, with a density field of value 1, and can be precomputed and reused if multiple bispectra estimations are done with the same binning.

4.2.3 Halo finding and mass functions

We use an inlined, on-the-fly version of the SUBFIND algorithm (Springel et al. 2001) to identify gravitationally bound structures in the simulations. The underlying friend-of-friends (FoF) haloes (Davis et al. 1985), defined with a linking length of $b = 0.2$, must contain a minimum number of 32 particles. When we discuss mass function results, we present results for these FoF haloes. We did not find a significant difference in our results if instead spherical overdensity masses obtained at 200 times the critical density were used, once the FoF masses are corrected for discreteness effects (see discussion in Section 5.4.1). To compute the mass function, we use 20 mass bins with uniform logarithmic spacing in the range $[32; 5 \times 10^5] \times m_{\text{part}}^{\text{sc}}$, where $m_{\text{part}}^{\text{sc}}$ is the particle mass in the runs with initial SC lattice. The ratios and error bars are estimated using bootstrap resampling of the catalogue of halo masses, similar to the method described in Section 4.2.1.

5 RESULTS

In this section, we present our study of the impact of the starting time z_{start} and of the order of the LPT used to initialize the simulations. We consider various summary statistics that are commonly extracted from cosmological N -body simulations: one-point cumulants, power spectra, and bispectra of the matter density field, and halo mass functions. To assess the impact of discreteness errors, we compare each case to the results from a simulation that oversamples the same initial perturbation field with four times more particles.

5.1 One-point statistics

We start by considering cumulants, S_n , of the matter density field. The non-linear gravitational collapse introduces non-Gaussianity in the originally purely Gaussian density field, which can be captured in these cumulants. In addition, high-order cumulants of the matter density field are sensitive to the order of the LPT used to set up the simulation and the starting redshift (e.g. Crocce et al. 2006).

Our measurements of the skewness and kurtosis of the density distribution function of the matter density field smoothed on different scales R_s are presented in Fig. 5. Specifically, we show ratios of $S_3(R_s)$ (left column), and $S_4(R_s)$ (right column) at redshifts $z = 0, 1, \text{ and } 3$ (first to third row) initialized at different times (represented by the line styles) and with different LPT orders (represented by the line colours) with respect to the 3LPT reference run initialized with an oversampled FCC lattice at $z_{\text{start}} = 24$.

For a large enough smoothing radius, the filtered density field approaches a Gaussian field, and the agreement between simulations should be perfect. This is what we see indeed for S_3 with both 2 and 3LPT. For S_4 however, while 3LPT converges correctly, 2LPT does not agree with our reference simulation even on large scales $R_s \gtrsim 10 h^{-1} \text{Mpc}$. At the latest starting time we consider, $z_{\text{start}} = 11.5$, the disagreement for 2LPT is larger than 1 per cent at $z = 3$, reducing to 0.2 per cent at $z = 0$ for $R_s = 25 h^{-1} \text{Mpc}$. For $z_{\text{start}} = 24$, the difference stays within 1 per cent at all scales for $z \leq 3$ but the transients are still visible in S_4 with 2LPT. Munshi, Sahni &

⁸ Available from <https://github.com/sjforeman/bskit>

⁹ Available from <https://github.com/bccp/nbodykit>

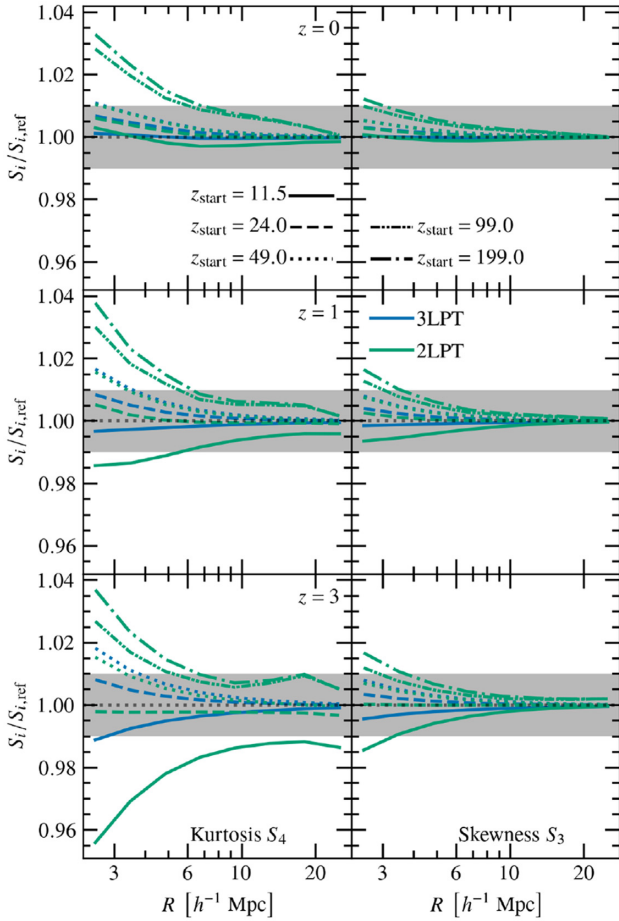


Figure 5. Dependence of one-point density field statistics on LPT order and starting redshifts. The panels show the ratios of the skewness S_3 (right-hand column) and kurtosis S_4 (left-hand column) as a function of smoothing scale R_s of simulations started at different redshifts (line styles), with different orders of LPT (line colors), with respect to the FCC run started at $z_{\text{start}} = 24$ and 3LPT initial conditions. We find these to be well converged in all cases, except the kurtosis when 2LPT is initialized too late, while 3LPT still yields about 1 per cent accuracy on S_3 and S_4 even when initialized as late as $z_{\text{start}} = 11.5$.

Starobinsky (1994) have shown in the case of spherical collapse that n LPT reproduces the one-point statistics up to S_{n+1} . We confirm here that indeed 2LPT correctly predicts S_3 but not S_4 in the large- R_s limit, while 3LPT is correct for both. Likewise, we expect all cumulants higher than 4 to be incorrect even with 3LPT. A similar result for 2LPT was presented by Crocce et al. (2006).

The deviations from Gaussianity that are produced by the gravitational collapse of overdensities affect the smallest scales first. The differences with respect to the reference simulation in Fig. 5 can be interpreted as the presence of small-scale transients due to the truncation of the LPT series. For the higher starting redshift ($z_{\text{start}} = 24$), they are all under 1 per cent even for the smallest shown scale of $2.5h^{-1}$ Mpc (since the mean particle separation is $\sim 1h^{-1}$ Mpc, we did not probe smaller scales). For the lower starting redshift $z_{\text{start}} = 11.5$, the disagreement with 2LPT for S_3 is under 2 per cent at $z = 3$ and goes under 1 per cent for $z \leq 1$. For S_4 , the relative difference is > 3 per cent at $z = 3$, 2 per cent at $z = 1$, and under 0.5 per cent at $z = 0$. This is roughly consistent with the earlier results of Tatekawa & Mizuno (2007) for 3LPT including only longitudinal

modes, who however were not studying per cent level agreements, so that a detailed comparison is not possible.

Of course, the lower starting redshift ($z_{\text{start}} = 11.5$) is extremely late by traditional wisdom of when to start a simulation, and based on our analysis of the convergence radius of LPT in Section 3, it is also dangerously close to the time when we expect LPT to break down (for the considered resolutions). It is thus no surprise that 2LPT does not fare well in this case. In contrast, however, 3LPT performs at the per cent accuracy level even for such extremely late starts. For example, in the case of S_4 , the difference is already under 2 per cent at $z = 3$, and for S_3 it is always below 1 per cent.

5.2 Two-point statistics

We next investigate the impact of the initial conditions and starting time on the late-time two-point statistics, as quantified by the matter density power spectrum. First, we quantify the magnitude of discreteness errors as a function of starting time, then the impact of LPT order and starting time on the accuracy of matter power spectra. We will consider a one per cent agreement with a reference solution to the largest possible wavenumber k as the benchmark of accuracy for all spectra.

5.2.1 Discreteness errors on power spectra

In order to investigate the impact of particle discreteness on the power spectrum, we ran the same initial perturbation spectrum with particles initially placed on an SC lattice, as well as on a four times oversampled FCC lattice (see our discussion in Section 2.4 for more details). The initial conditions were generated using 3LPT at starting redshifts $z_{\text{start}} = 49, 24$, and 11.5 . In Fig. 6, we show the ratio of the SC to the FCC runs at redshifts $z = 3, 1$, and 0 . At $z = 0$, all power spectra agree to better than one per cent up to the particle Nyquist wavenumber, irrespective of the starting time. However, at higher redshifts, the SC lattice shows an important suppression of the small scale structures near the particle Nyquist wavenumber due to discreteness errors. At $z = 3$ for $z_{\text{start}} = 49$, the underestimation of the SC compared to FCC is about 8 per cent, down to ~ 3 per cent at $z = 1$. For the late start $z_{\text{start}} = 11.5$, we find a maximum suppression of ~ 2 per cent, ~ 1 per cent and ~ 0.4 per cent at $z = 3, 1$, and 0 , respectively. As a consequence, the power spectra agree to better than one per cent for $k \gtrsim k_{\text{Ny}}/3$ at $z = 3$ and for all $k \gtrsim k_{\text{Ny}}$ at $z = 1$ and 0 . While not shown in the figure, we found that for a fixed starting redshift, the suppression does not depend on the order of LPT used. The independence from the LPT order clearly indicates an origin in discreteness errors, and the shape of the power suppression is indeed very similar to the scale-dependent PLT growth factors (cf. fig. 1 of Joyce et al. 2005) indicating that the fluid ICs relax to the discrete evolution in the quasi linear regime, and increasingly so, the earlier the starting time.

We thus find that *the higher the starting redshift is, the greater is the suppression of the high k modes, irrespective of the order of LPT*. The effect peaks at intermediate redshifts, before the high- k part of the resolved power spectrum becomes fully dominated by collapsed haloes. A similar suppression of high- k power that improves the better the non-linear scale is resolved has also been reported by Schneider et al. (2016), who also investigated the additional dependence on the gravity solver employed in the simulations. We leave an assessment of whether/how much our results depend on the choice of N -body gravity solvers for future work.




Figure 6. Impact of particle discreteness on the matter power spectrum evolution: ratios between the power spectra of the SC and oversampled FCC runs for different starting redshifts (different line styles) at $z = 0, 1,$ and 3 (top to bottom panels). All simulations shown use 3LPT ICs, but the results do not depend on the LPT order. The grey shaded area indicates per-cent-level agreement with the reference. The vertical dashed line indicates the particle Nyquist wavenumber of the initial perturbation field. Discreteness effects lead to a time-dependent suppression close to the particle Nyquist wavenumber.

5.2.2 Dependence of power spectra on LPT order and starting time

Fig. 7 shows ratios of power spectra at redshifts $z = 0, 1,$ and 3 (panels top to bottom) of simulations starting from different initial conditions with respect to a reference simulation. Specifically, we vary the order of the LPT from the ZA (orange) to 2LPT (green) to 3LPT (blue lines), and the starting time between $z_{\text{start}} = 199, 99, 49, 24,$ and 11.5 (as indicated by the different line styles). Note that we have not run all combinations of LPT order and z_{start} . We use a simulation initialized with 3LPT at $z_{\text{start}} = 24$ and using an FCC lattice (i.e. four times more particles) as our reference simulation. This reference solution should suffer from much reduced discreteness effects due to the significantly higher particle number and increased symmetry of the lattice.

The results shown in Fig. 7 indicate clearly that for a simulation with our resolution, 3LPT performs best at $z = 3$ and is accurate for $k \lesssim k_{\text{Ny}}/2$ if started at $z_{\text{start}} = 11.5$, while performing increasingly similar to 2LPT for earlier starts, with both being accurate roughly for $k \lesssim k_{\text{Ny}}/4$ and $\lesssim k_{\text{Ny}}/5$ when started at $z = 24$ and $z = 49$, respectively. Note that 2LPT, in stark contrast to 3LPT, performs very poorly for the latest start $z_{\text{start}} = 11.5$, which just reflects that this is indeed a too late start for second order LPT. The ZA runs, for which we only considered much earlier starting times $z_{\text{start}} = 49, 99,$ and 199 , are all accurate only to $\sim k_{\text{Ny}}/10$, and have the largest errors overall.

As structures collapse into haloes over time, the power spectra become significantly less sensitive to the initial conditions as evidenced by our results at $z = 1$ and $z = 0$. The 3LPT run with the latest starting time $z_{\text{start}} = 11.5$ is now accurate up to the particle Nyquist wave

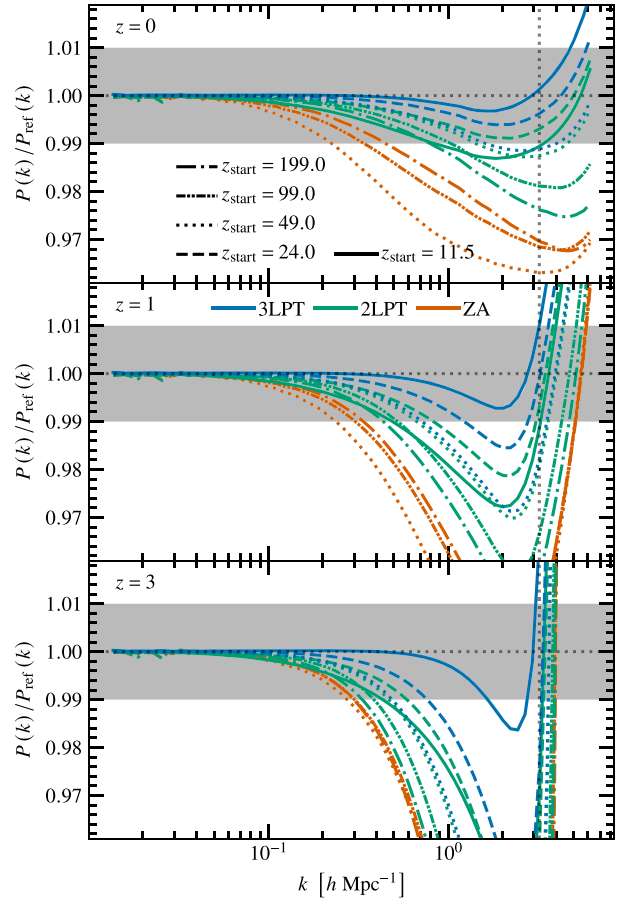


Figure 7. Impact of starting time and LPT order on the matter power spectrum evolution: ratios of power spectra for different orders of LPT (line colours) and starting redshifts (line styles) at $z = 0, 1,$ and 3 (top to bottom panels), w.r.t. the FCC 3LPT $z_{\text{start}} = 24$ reference simulation. The grey shaded area indicates per-cent-level agreement with the reference, and the dashed vertical line indicates the particle Nyquist wavenumber. Higher order LPT and later starting times improve agreement with the reference run.

numbers ($z = 1$) or even beyond ($z = 0$). This is also true for 2LPT and 3LPT at $z = 0$ if started at 24. All other runs are less accurate, to varying degrees. Most notably, and as expected from our discussion of discreteness effects, the ZA runs converge only extremely slowly with increasingly earlier starting time, and arguably converge to the discrete solution rather than the fluid solution. Even for the earliest start we considered, $z_{\text{start}} = 199$, the ZA run is only accurate for $k \lesssim k_{\text{Ny}}/8$ ($\lesssim k_{\text{Ny}}/9$) at $z = 0$ ($z = 1$). We note that for early starts, of course, ZA and 2LPT are consistent at $z = 0$ within about one per cent of each other (cf. the dash-dotted orange and green lines), which is what has also been found by Schneider et al. (2016) (see their fig. 3), and the general picture for the poor performance of ZA versus 2LPT has already been outlined by Crocce et al. (2006).

For 2LPT, as for 3LPT, a later start gives the best results with $z_{\text{start}} = 24$ among the runs we considered, but the latest possible start is of course earlier for 2LPT than for 3LPT. This is in fact perfectly consistent with the analysis for 2LPT of Nishimichi et al. (2019) who followed a different approach to find the optimal starting redshift, by finding that z_{start} for which the small-scale power spectrum has the highest amplitude, thereby finding the ‘sweet spot’ between particle discreteness and LPT truncation transients. Their $250 h^{-1} \text{Mpc}$ box

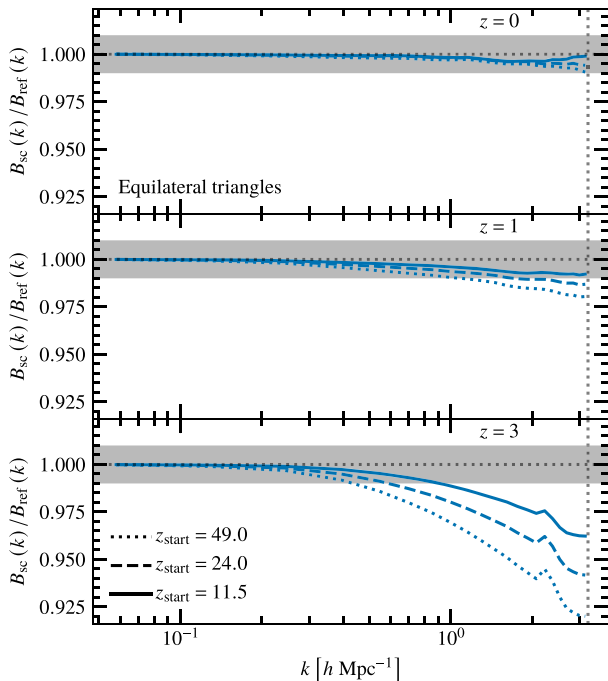


Figure 8. Same as Fig. 6 but showing the impact of particle discreteness on the equilateral bispectrum $B(k) \equiv B(k, k, k)$.

with 256^3 particles has the same Nyquist mode as our runs, and they find the optimal start to be around $z_{\text{start}} \simeq 30$. For comparison, we note that the so-called Euclid Flagship simulation (Potter et al. 2017) and those used in the EuclidEmulator (Euclid Collaboration et al. 2019) were initialized at $z = 200$ with ZA. Our results indicate that their results could be systematically biased low by about 4 per cent at $k \sim 4 h\text{Mpc}^{-1}$. In addition, the simulations of (Angulo et al. 2020), were initialized at $z = 49$ with 2LPT, which should be biased low by 1.5 per cent on the same scales.

Overall, we thus find that a picture emerges that fits nicely with our analysis of the convergence radius of LPT presented in Section 3.4. The best results are obtained with the latest start at the highest order of LPT. These results should be compared against the ones displayed in Fig. 4. The starting redshifts of 49, 24, and 11.5 are placed respectively just below the 99.7th, just above the 95th, and just above the 68th percentile of the convergence radius plot for the box size and resolution considered here (cf. symbols in Fig. 4). It thus appears that for the best accuracy on the power spectrum, pushing the starting time as low as the 68th percentile yields the best results. We did not consider even later starting times, but they might yield a small further improvement for 3LPT. One has to caution against pushing this too far however, as beyond the radius of convergence, higher orders of LPT introduce higher errors.

5.3 Three-point statistics

In this section, we repeat the analysis of the previous section for bispectra. In order to simplify the analysis, we exclusively focus on equilateral bispectra $k = k_1 = k_2 = k_3$, which are one-dimensional in that they depend only on a single scalar k , and not on the direction \mathbf{k} (due to the assumed statistical isotropy), and thus allow us to probe the scale-dependence up to the particle Nyquist wavenumber most conveniently. We analyse the bispectra for the same simulations

as those in Section 5.2. Again, we first quantify the magnitude of discreteness errors as a function of starting time, then the impact of LPT order and starting time on the accuracy of the bispectra. As before, we will consider a one per cent agreement with a reference solution in order to quantify accuracy.

5.3.1 Discreteness errors on bispectra

We again quantify the impact of particle discreteness on bispectra by taking ratios of bispectra measured from simulations that sample the same perturbations using an SC and a four times oversampled FCC lattice (cf. Section 2.4). As before, we vary the starting redshift of the 3LPT initial conditions between $z_{\text{start}} = 49, 24$, and 11.5. The results of this study are shown in Fig. 8, and are in broad agreement with our previous analysis of discreteness effects on the power spectrum. Again, we see a time and scale dependent suppression which is the strongest close to the particle Nyquist wavenumber. For the latest starting time, $z_{\text{start}} = 11.5$, agreement with the reference run is at the sub-per cent level at $z = 0$ and 1 up to k_{Ny} , with increasing errors for the earlier starts. At $z = 3$ errors are significantly larger, but still with later starts faring better. These results are in tension with the interpretation given by McCullagh, Jeong & Szalay (2015) who report a convergence of bispectra only for early enough starting times. We will discuss this more below.

5.3.2 Dependence of bispectra on LPT order and starting time

We next investigate the dependence of the accuracy of bispectra on the initial conditions varying LPT order and starting time. Fig. 9 shows the ratios of equilateral bispectra w.r.t. the reference run (3LPT, oversampled FCC lattice, $z_{\text{start}} = 24$). We vary the orders of LPT (line colour) and starting redshift (line style), and present measurements at $z = 0, 1$ and 3. One notes immediately that for bispectra, the impact of the order of LPT used to set up the simulations is much more pronounced than for the power spectra presented in the previous Section 5.2.2. At all times and all scales, and irrespective of the starting time, the bispectrum obtained from ZA initial conditions disagrees with the reference run, with differences dramatically increasing at earlier times. Changing to 2LPT ICs improves the situation dramatically, as has been reported before (e.g. McCullagh et al. 2015; Baldauf et al. 2015). However, even for 2LPT, the bispectra agree at $z = 1$ with the reference solution only for $k \lesssim k_{\text{Ny}}/15$ (if we disregard the latest starting time $z_{\text{start}} = 11.5$). The dependence on starting time is however non-trivial for 2LPT: there is clearly a sweet spot for the runs starting at $z \approx 24$ –49, with worse agreement with the reference run for both earlier and later starts. While disagreeing with the reference run, the results appear to converge for increasingly earlier starting times ($z_{\text{start}} = 99$ and 199), which are however strongly suppressed w.r.t. the reference run at $k \sim 1 h\text{Mpc}^{-1}$. Generally, the situation improves with the growth of non-linear structure on increasingly larger scales at $z = 0$. Now the later starting times $z_{\text{start}} = 49$ and 24 agree with the reference run at $k \lesssim k_{\text{Ny}}/4$, while the early starts agree perfectly with one another, but not with the reference run for $k \gtrsim k_{\text{Ny}}/8$. We are arguably witnessing the convergence to the discrete solution with the increasingly higher starting times here. The best results are obtained with 3LPT and again the latest possible starts, which agree with the reference at $z = 0$ at all scales we investigated. At $z = 1$, the agreement is still excellent for the $z_{\text{start}} = 24$ run ($k \lesssim k_{\text{Ny}}/2$). It is slightly less good for the latest start ($z_{\text{start}} = 11.5$), indicating that this is a slightly too late start even




Figure 9. Same as Fig. 7 but showing the impact of starting time and LPT order on the equilateral bispectrum. The impact of 3LPT on the bispectrum is more dramatic: At $z = 1$, we were not able to achieve per cent level agreement at this resolution with 2LPT to small scales, and ZA is so off that it should not be used for precision studies. 3LPT allows sub-per cent errors on the bispectrum at $z = 0$ up to the Nyquist mode.

for 3LPT. At $z = 3$ the 3LPT runs are accurate only for $k \lesssim k_{\text{Ny}}/15$ but have overall the smallest errors.

In contrast to the results of McCullagh et al. (2015), who suggest using very early starting times, our results appear to indicate quite the opposite. Since bispectra are impacted by particle discreteness just like power spectra, discreteness errors leading to a deviation from the fluid limit accumulate more, the higher the starting redshift and the smaller the perturbations.

In summary, we find that, as for the power spectra, the bispectrum accuracy is best when high-order LPT is combined with a very late start. N -body simulations starting from 3LPT at our relatively low resolution certainly can be accurate to predict the bispectrum at $z = 0$ to better than one per cent up to the particle Nyquist wavenumber.

Note that there have been concerns in the literature about the accuracy with which N -body simulations predict the matter bispectrum (e.g. Schmittfull, Regan & Shellard 2013; Baldauf et al. 2015; Hung, Fergusson & Shellard 2019). For instance, Baldauf et al. (2015) assumes a systematic error in their simulated results to account for lack of convergence when ZA or 2LPT is used. In contrast, our results indicate that once discreteness and truncation are taken into account, N -body simulations can provide very reliable and accurate results.

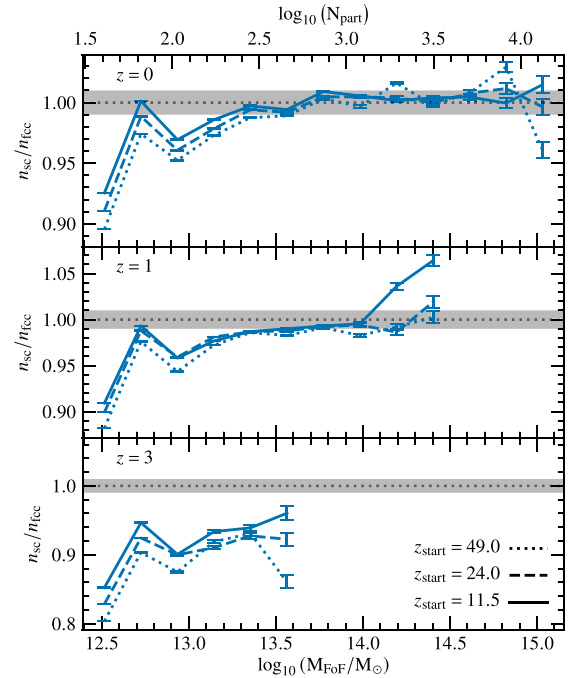


Figure 10. Same as Fig. 6 but for the halo mass function. We indicate halo mass in the bottom x -axis, and corresponding particle number in the SC run in the top x -axis. The FoF-halo masses have been corrected using the Warren et al. (2006) correction (see text for details).

5.4 Halo mass functions

Finally, we also consider the abundance of collapsed structures – as quantified by the halo mass function – as a benchmark for the dependence of simulation results on the initial conditions. To this end, we essentially repeat once again the analysis performed already in the sections before. We first quantify the impact of discreteness on the mass function, and in a second step the dependence on starting time and order of LPT on the abundance of haloes.

5.4.1 Discreteness errors in the mass function

In Fig. 10, we show the ratio of mass functions obtained from 3LPT IC runs (starting from an SC lattice) with varying starting time to our usual reference run (3LPT, oversampled FCC lattice, $z_{\text{start}} = 24$). The FCC reference has four times more particles, meaning that haloes of a given mass are much better resolved in that run, but does not introduce further small-scale modes (see discussion in Section 2.4). In this sense, our convergence test is different from the usual convergence tests of halo mass functions (e.g. Jenkins et al. 2001; Tinker et al. 2008), where simply different ranges of scales of the full Λ CDM spectrum are resolved. In our results, we show mass functions for FoF haloes. The FoF algorithm is known to be biased at low particle numbers due to percolation noise. Warren et al. (2006) have proposed to correct the mass of an FoF group due to discreteness errors at a finite number of particles N_{part} as

$$M_{\text{FoF}} = M_{\text{FoF,raw}} \left(1 - N_{\text{part}}^{-0.6} \right), \quad (36)$$

which we have applied before computing the mass functions (note that this is an ad hoc correction that might take a different form depending on the mass resolution or even starting redshift). It is also entirely possible that an SC and an FCC lattice lead to somewhat

different discreteness errors in FoF haloes, which we have not investigated but do not believe to strongly impact our conclusions.

After applying this correction, our ratio plots are very similar to those obtained when using spherical overdensity haloes (not shown) for an overdensity of 200 times the critical density, which also underpredicts the number of haloes at the low mass end compared to the higher resolution simulation. Arguably, a more optimal correction can be performed to push up the low-mass end of the FoF-mass function (cf. e.g. Nishimichi et al. 2019), but such avenues shall not be our concern here. We also found that without the correction, the FoF mass function ratio is biased high at the per cent level at all masses due to ‘Eddington bias’ (Eddington 1913), which is implicitly corrected by the ‘Warren’ correction.

Looking at our simulation results, we first notice in Fig. 10 a systematic drop below ~ 300 particles at late times $z = 0, 1$ which is more severe for the earlier start $z_{\text{start}} = 49$ than the later starts. This undershooting is of the order of $\lesssim 5$ per cent for haloes of ~ 100 particles. Finally, at the earliest time $z = 3$, we see a systematic (mass-independent) bias in the mass function that depends only weakly on the starting time of the simulation. While at this early time, the haloes in our simulations are not well resolved (all below 1000 particles), the mass function is low by 10 per cent, which is much more than at the later times. From our previous analysis of power- and bispectra, we know that simulation results are particularly sensitive to the ICs at this early time.

5.4.2 Impact of LPT order and starting time on the mass function

Finally, we show in Fig. 11 ratios of mass functions of dark matter haloes relative to the reference run (3LPT, oversampled FCC lattice, $z_{\text{start}} = 24$) for different orders of LPT and starting redshifts. In general, we find a strong dependence on both the order and starting time.

Let us first focus on the latest time, $z = 0$, when structures are more developed, and we have many well resolved haloes. We find that at the high mass end, all ICs converge to the same mass function, in per cent level agreement with the reference run. At the low-mass end, we find however a very strong dependence on the starting time of the simulation, but not the LPT order used. The earliest starts are low by more than 5 per cent at the ~ 100 particle scale. In addition, the ZA runs for the latest time we ran with ZA, $z_{\text{start}} = 49$ are low at 2–3 per cent even at high masses. This picture is consistent with earlier studies that also did not find a strong dependence of the late-time mass function on starting time (e.g. Jenkins et al. 2001; Tinker et al. 2008; Knebe et al. 2009, in the pre-precision cosmology era), and Reed et al. (2013) and Nishimichi et al. (2019) who find per cent level convergence for moderately early starts with 2LPT.

At the intermediate time, $z = 1$, a more mixed picture appears in that the ZA runs are low by at least 5 per cent at all masses, and we see a weak additional dependence on the LPT order used, with 3LPT ever so slightly better converged than 2LPT.

Finally, at early times, $z = 3$, we find a very strong dependence on the order of LPT used. ZA predicts significantly (by more than 20 per cent) fewer haloes at fixed mass than the reference run, even when started early ($z_{\text{start}} = 99$ and 199). This is improved by 2LPT to about 10–15 per cent if started at $z_{\text{start}} = 24$ or 49. The 3LPT runs are closer to the reference, and show a weaker dependence on starting time than 2LPT, but are still ~ 10 per cent low.

At all times, the very early starts, $z_{\text{start}} = 99$ and 199, show a significantly lower abundance of low-mass haloes w.r.t. the later started runs when 2LPT and 3LPT are used. At the later times $z = 0$ and 1, both 2LPT and 3LPT give results that are within about




Figure 11. Same as Fig. 7 but for the halo mass function, showing a strong dependence on starting time and LPT order, especially at high redshift.

2 per cent of one another for starting times of $z_{\text{start}} = 24$ or even 11.5. The 2LPT run with $z_{\text{start}} = 11.5$ however is clearly off at early times, $z = 3$. Again, consistent with the power- and bispectrum results, we find that discreteness effects impact most strongly the early starts, while late starts with high-order LPT give the most converged results. Only for the ZA runs did we actually observe an improvement when starting earlier.

As has been demonstrated by Reed et al. (2013) and the recent study of Ludlow, Schaye & Bower (2019), the absolute convergence of the halo mass function at the low-mass end also depends on further variables, such as force resolution and time-stepping. Our results should therefore ultimately be subjected to further studies that include also variations of time and force integration parameters to determine the ultimate errors on the mass function at the few-particle, low-mass end.

6 SUMMARY AND CONCLUSIONS

In this paper, we have presented a rigorous analysis of the interplay of discreteness effects and the order truncation in LPT used to set up initial conditions for cosmological N -body simulations. Our findings strongly suggest, contrary to common wisdom, that in order to be most economical (i.e. the most accurate to the smallest possible scales with the least computational resources), N -body simulations should be initialized with the highest possible order of LPT at the latest possible time. In that case, we are able to achieve the, admittedly, ambitious goal of sub-per cent level convergence in the matter power and bispectrum all the way to the particle Nyquist wavenumber. We shall summarize the arguments leading up to this conclusion next.

We have considered simulations based on initial conditions that use up to third order in Lagrangian perturbation theory (3LPT). The numerical implementation and proofs of correctness of the numerical implementation are somewhat more involved than for low-order LPT (see Section 2.3). In particular, high-order LPT involves the convolution of non-linear fields and thus, numerical implementations should be de-aliased. We demonstrate that for efficiency, at least up to 3LPT, this can however be disrespected in IC generation, since errors decay away quickly for simulations evaluated at low redshift.

The particle discretization used in an N -body simulation only approximates the underlying dark matter fluid, which manifests itself in a deviation from the evolution expected in the continuum limit. We have shown in Section 2.4 that this results in an amplitude suppression close to the particle Nyquist wavenumber of the initial particle grid in both power- and bispectra. When correcting for the linear discreteness error (using PLT, cf. e.g. Joyce et al. 2005), we found that the N -body simulation agrees with 3LPT to about 1–2 per cent up to the particle Nyquist wavenumber at much later times than are typically used to initialize N -body simulations. At the same time, the N -body system, initialized as a fluid, quickly relaxes to the evolution of the discrete system (cf. also Garrison et al. 2016), particularly so while perturbations are still relatively small during the quasi-linear stage of the evolution. As a consequence, it follows that the earlier the starting time, the more closely will the evolution follow the discrete and not the fluid solution. It is thus obvious that a late starting time is preferable, since it initializes the simulation using fluid perturbations that are already relatively large. Subsequent nonlinear evolution, and the transfer of power from large to small scales, will reduce the importance of initial discreteness errors, yielding progressively better convergence.

Our conclusions are in contrast with traditional wisdom (see however Nishimichi et al. 2019), that would tell us that perturbation theory is increasingly more accurate at early times, and so that cosmological N -body simulations should be initialized early. For instance, the Euclid Flagship simulation follows this guideline and starts at $z = 200$.

A late start of a cosmological N -body simulation then leads to the question of how late one could start. The answer depends, of course, on the convergence radius of LPT with Λ CDM random initial conditions, as well as on the truncation error introduced by fixed-order LPT. As regards to the former, while theoretical predictions by Zheligovsky & Frisch (2014) and Rampf et al. (2015) revealed a finite but lower bound on the convergence radius of LPT, here we have performed provided a novel quantitative analysis that demonstrates the validity of LPT at much later times; see Section 3.2. We found that the convergence radius follows relatively simple scaling laws with the RMS amplitude of density fluctuations at the resolution scale (see Fig. 4). We then assessed the impact of the LPT truncation error using standard convergence tests based on various summary statistics. In particular, we demonstrated that with 3LPT, the truncation error is so much reduced that it could be ignored even for very late starts.

In all cases, we tested convergence against a simulation with four times better mass resolution, while using the same perturbation modes. This is in contrast to previous convergence studies that usually alter the ICs by adding more perturbations when resolution is increased. In Λ CDM, adding new modes implies adding more non-linearities at a given time. This, in our reasoning, impacts the starting time and truncation errors, and thus makes any rigorous tests of LPT truncation errors and assessment of transients questionable. For production runs, the increase in particles is likely better invested in sampling also additional modes – which should however be investigated in future work.

After investigating in Section 5, the dependence on starting time and LPT truncation of one-point statistics of the density field, density power- and bispectra, as well as halo mass functions, we can summarize our main findings as follows.

(i) 3LPT initial conditions with the latest starting redshift we considered, namely $z_{\text{start}} = 11.5$, generally yielded the most accurate results in all summary statistics w.r.t. our better resolved reference simulation. Furthermore, memory-intensive de-aliasing procedures may be safely ignored as the evolved aliasing errors are of order per mille for all considered observables for $z \lesssim 3$. Note however that this optimal starting time is for the mass resolution we consider here, and increases at higher mass resolution (cf. Fig. 4 for the scaling).

(ii) In contrast, the ZA is almost never accurate enough for initial conditions for precision cosmology, especially if one is interested in higher-order statistics such as the matter bispectrum.

(iii) Lower-order LPT with much higher starting redshift, e.g. $z_{\text{start}} = 199$, leads to much worse agreement with the reference solution. Instead, it indicates a ‘spurious’ convergence to the discrete (lattice) solution, as opposed to the continuous fluid solution. It is also possible that early starts are further impacted by additional numerical force errors due to the smallness of perturbations.

(iv) The impact of the truncation error increases with the order of the statistics considered: kurtosis is more affected than skewness, and bispectra are more affected than power spectra. One can conjecture that this continues to be even more severe for higher orders.

(v) On scales smaller than the quasi-linear scales, i.e. those which are dominated by halo profiles, we find a very weak dependence on initial conditions. As a consequence, the accuracy always improves at late times, while errors are largest at high redshift.

Therefore, we deem it quite possible that the very high particle count per simulation volume size typically advocated for precision era simulations, is to some significant degree driven by discreteness errors due to early starts. Specifically, to reach one per cent accurate predictions for the power spectrum at $k \sim 1 h \text{Mpc}^{-1}$ over $z \in [0, 1]$, it is required to employ LPT at second order or higher with starting redshift $z \sim 50$ or less. At $k \sim 3 h \text{Mpc}^{-1}$ the requirement is stronger, and only 3LPT is suitable for such accuracy. These numbers are however specific to the relatively low-mass resolution that we employed in this paper, and at higher mass resolution one can expect altered requirements. We leave it for future work to establish the precise requirements for a given mass resolution.

In conclusion, our results thus suggest that late-start ICs based on higher-order LPT might allow for more economic simulations to achieve a given accuracy – a conjecture that should however be backed up with further convergence and resolution studies, and in particular also beyond 3LPT. Nevertheless, we expect that the increasing computational footprint at successively higher orders in LPT should eventually become too expensive compared to the improvement (cf. Rampf & Hahn 2020).

While our current results only concern simulations of the total matter flow in terms of a single collisionless fluid, the real Universe is composed of various distinct components, in particular baryons. Such multicomponent simulations are however prone to more theoretical uncertainties. Simulations that model baryons with Lagrangian methods (SPH or moving mesh, when coupling to gravity using N -body techniques) suffer from particle discreteness errors which appear to be even more severe than in the single-fluid case (cf. Angulo et al. 2013; Valkenburg & Villaescusa-Navarro 2017; Bird et al. 2020). While not suffering from particle discreteness errors, Eulerian hydrodynamical simulations are impacted by advection errors and intrinsic smoothing due to the order of the finite volume method that

also leave an imprint on the power spectrum (Hahn & Abel 2011, see their fig. 23). Extrapolating the logic of this paper, one way out of the situation for both Lagrangian and Eulerian hydrodynamics would be to initialize multifluid simulations also as late as possible to minimize discretization effects. This, however, requires the development of higher-order multifluid initial conditions to begin with, and will be accompanied by further hurdles that arise when incorporating baryonic physics. Such avenues will be assessed in Rampf, Uhlemann & Hahn (2020) and Hahn, Rampf & Uhlemann (2020).

ACKNOWLEDGEMENTS

We thank Bruno Marcos for useful discussions, and Romain Teyssier, Volker Springel, Simon White, and Jens Stuecker for comments on the manuscript. MM, OH, and RA acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programmes: Grant agreement No. 679145 (COSMO-SIMS) for MM and OH; and for RA, grant agreement no. 716151 (BACCO). CR is a Marie Skłodowska-Curie Fellow and acknowledges funding from the People Programme (Marie Curie Actions) of the European Union's H2020 Marie Skłodowska-Curie Actions under Grant Agreement No. 795707 (COSMO-BLOW-UP). This work was granted access to the HPC resources of TGCC/CINES under the allocation A0060410847 attributed by GENCI (Grand Equipement National de Calcul Intensif). We thank the authors of BSKIT for making their software publicly available.

DATA AVAILABILITY

The code to generate 3LPT initial conditions is freely available at <http://bitbucket.org/ohahn/monofonic>. The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Abell P. A. et al., 2009, preprint ([arXiv:0912.0201](https://arxiv.org/abs/0912.0201))
 Aihara H. et al., 2018, *PASJ*, 70, S8
 Angulo R. E., Pontzen A., 2016, *MNRAS*, 462, L1
 Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, *MNRAS*, 426, 2046
 Angulo R. E., Hahn O., Abel T., 2013, *MNRAS*, 434, 1756
 Angulo R. E., Zennaro M., Contreras S., Aricò G., Pellejero-Ibañez M., Stücker J., 2020, preprint ([arXiv:2004.06245](https://arxiv.org/abs/2004.06245))
 Baldauf T., Mercolli L., Mirbabayi M., Pajer E., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 007
 Barnes J., Hut P., 1986, *Nature*, 324, 446
 Bernardeau F., Colombi S., Gaztañaga E., Scoccimarro R., 2002, *Phys. Rep.*, 367, 1
 Bird S., Feng Y., Pedersen C., Font-Ribera A., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 002
 Blas D., Lesgourgues J., Tram T., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 034
 Bouchet F. R., Juszkiewicz R., Colombi S., Pellat R., 1992, *ApJ*, 394, L5
 Bouchet F. R., Colombi S., Hivon E., Juszkiewicz R., 1995, *A&A*, 296, 575
 Bowman J. C., Roberts M., 2011, *SIAM J. Sci. Comput.*, 33, 386
 Brenier Y., Frisch U., Henon M., Loeper G., Matarrese S., Mohayaee R., Sobolevskii A., 2003, *MNRAS*, 346, 501
 Buchert T., 1994, *MNRAS*, 267, 811
 Buchert T., Ehlers J., 1993, *MNRAS*, 264, 375
 Buchert T., Goetz G., 1987, *J. Math. Phys.*, 28, 2714
 Buchert T., Melott A. L., Weiss A. G., 1994, *A&A*, 288, 349
 Buehlmann M., Hahn O., 2019, *MNRAS*, 487, 228
 Catelan P., 1995, *MNRAS*, 276, 115

Cheng S., Yu H.-R., Inman D., Liao Q., Wu Q., Lin J., 2020, in *2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID)*. IEEE, Melbourne, p. 685
 Chisari N. E., Zaldarriaga M., 2011, *Phys. Rev. D*, 83, 123505
 Colombi S., 2015, *MNRAS*, 446, 2902
 Crocce M., Pueblas S., Scoccimarro R., 2006, *MNRAS*, 373, 369
 Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
 Domb C., Sykes M. F., 1957, *Proc. R. Soc. A*, 240, 214
 Eddington A. S., 1913, *MNRAS*, 73, 359
 Efstathiou G., Davis M., White S. D. M., Frenk C. S., 1985, *ApJS*, 57, 241
 Euclid Collaboration et al., 2019, *MNRAS*, 484, 5509
 Fidler C., Tram T., Rampf C., Crittenden R., Koyama K., Wands D., 2017, *J. Cosmol. Astropart. Phys.*, 1712, 022
 Foreman S., Coulton W., Villaescusa-Navarro F., Barreira A., 2020, *MNRAS*, 498, 2887
 Garrison L. H., Eisenstein D. J., Ferrer D., Metchnik M. V., Pinto P. A., 2016, *MNRAS*, 461, 4125
 Hahn O., Abel T., 2011, *MNRAS*, 415, 2101
 Hahn O., Angulo R. E., 2016, *MNRAS*, 455, 1115
 Hahn O., Paranjape A., 2016, *Phys. Rev. D*, 94, 083511
 Hahn O., Abel T., Kaehler R., 2013, *MNRAS*, 434, 1171
 Hahn O., Angulo R. E., Abel T., 2015, *MNRAS*, 454, 3920
 Hahn O., Rampf C., Uhlemann C., 2020, preprint ([arXiv:2008.09124](https://arxiv.org/abs/2008.09124))
 Hand N., Feng Y., Beutler F., Li Y., Modi C., Seljak U., Slepian Z., 2018, *AJ*, 156, 160
 Hansen S. H., Agertz O., Joyce M., Stadel J., Moore B., Potter D., 2007, *ApJ*, 656, 631
 Heitmann K., White M., Wagner C., Habib S., Higdon D., 2010, *ApJ*, 715, 104
 Heitmann K. et al., 2016, *ApJ*, 820, 108
 Heitmann K. et al., 2019, *ApJS*, 245, 16
 Hockney R. W., Eastwood J. W., 1981, *Computer Simulation Using Particles*. Computer Simulation Using Particles. McGraw-Hill, New York
 Hung J., Fergusson J. R., Shellard E. P. S., 2019, preprint ([arXiv:1902.01830](https://arxiv.org/abs/1902.01830))
 Jenkins A., 2010, *MNRAS*, 403, 1859
 Jenkins A. et al., 1998, *ApJ*, 499, 20
 Jenkins A., Frenk C. S., White S. D. M., Colberg J. M., Cole S., Evrard A. E., Couchman H. M. P., Yoshida N., 2001, *MNRAS*, 321, 372
 Jing Y. P., 2005, *ApJ*, 620, 559
 Joyce M., Marcos B., 2007, *Phys. Rev. D*, 75, 063516
 Joyce M., Marcos B., Gabrielli A., Baertschiger T., Sylos Labini F., 2005, *Phys. Rev. Lett.*, 95, 011304
 Klypin A. A., Shandarin S. F., 1983, *MNRAS*, 204, 891
 Knebe A., Wagner C., Knollmann S., Diekershoff T., Krause F., 2009, *ApJ*, 698, 266
 L'Huillier B., Park C., Kim J., 2014, *New Astron.*, 30, 79
 Laureijs R. et al., 2011, preprint ([arXiv:1110.3193](https://arxiv.org/abs/1110.3193))
 Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
 Liao S., 2018, *MNRAS*, 481, 3750
 Ludlow A. D., Schaye J., Bower R., 2019, *MNRAS*, 488, 3663
 McCullagh N., Jeong D., Szalay A. S., 2015, *MNRAS*, 455, 2945
 McDonald P., Vlah Z., 2018, *Phys. Rev. D*, 97, 023508
 Marcos B., 2008, *Phys. Rev. D*, 78, 043536
 Marcos B., Baertschiger T., Joyce M., Gabrielli A., Sylos Labini F., 2006, *Phys. Rev.*, D73, 103507
 Matsubara T., 2015, *Phys. Rev. D*, 92, 023534
 Monaco P., Theuns T., Taffoni G., 2002, *MNRAS*, 331, 587
 Munari E., Monaco P., Sefusatti E., Castorina E., Mohammad F. G., Anselmi S., Borgani S., 2017, *MNRAS*, 465, 4658
 Munshi D., Sahni V., Starobinsky A. A., 1994, *ApJ*, 436, 517
 Nishimichi T. et al., 2019, *ApJ*, 884, 29
 Orszag S. A., 1971, *J. Atmos. Sci.*, 28, 1074
 Peebles P. J. E., 1967, *ApJ*, 147, 859
 Pichon C., Bernardeau F., 1999, *A&A*, 343, 663
 Pietroni M., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 028
 Planck Collaboration, 2020, *A&A*, 641, A6
 Podvigina O., Zheligovsky V., Frisch U., 2016, *J. Comput. Phys.*, 306, 320
 Potter D., Stadel J., Teyssier R., 2017, *Comput. Astrophys. Cosmol.*, 4, 2
 Pueblas S., Scoccimarro R., 2009, *Phys. Rev. D*, 80, 043504

- Rampf C., 2012, *J. Cosmol. Astropart. Phys.*, 1212, 004
Rampf C., 2019, *MNRAS*, 484, 5223
Rampf C., Buchert T., 2012, *J. Cosmol. Astropart. Phys.*, 6, 021
Rampf C., Frisch U., 2017, *MNRAS*, 471, 671
Rampf C., Hahn O., 2020, preprint (arXiv:2010.12584)
Rampf C., Villone B., Frisch U., 2015, *MNRAS*, 452, 1421
Rampf C., Frisch U., Hahn O., 2019, preprint (arXiv:1912.00868)
Rampf C., Uhlemann C., Hahn O., 2020, preprint (arXiv:2008.09123)
Reed D. S., Smith R. E., Potter D., Schneider A., Stadel J., Moore B., 2013, *MNRAS*, 431, 1866
Roberts M., 2011, PhD thesis, University of Alberta
Saga S., Taruya A., Colombi S., 2018, *Phys. Rev. Lett.*, 121, 241302
Schmittfull M. M., Regan D. M., Shellard E. P. S., 2013, *Phys. Rev. D*, 88, 063512
Schneider A. et al., 2016, *J. Cosmol. Astropart. Phys.*, 2016, 047
Scoccimarro R., 1998, *MNRAS*, 299, 1097
Scoccimarro R., 2000, *ApJ*, 544, 597
Sefusatti E., Crocce M., Scoccimarro R., Couchman H. M. P., 2016, *MNRAS*, 460, 3624
Sousbie T., Colombi S., 2016, *J. Comput. Phys.*, 321, 644
Springel V., 2005, *MNRAS*, 364, 1105
Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
Stücker J., Schmidt A. S., White S. D. M., Schmidt F., Hahn O., 2020a, preprint (arXiv:2003.06427)
Stücker J., Hahn O., Angulo R. E., White S. D. M., 2020b, *MNRAS*, 495, 4943
Taruya A., Colombi S., 2017, *MNRAS*, 470, 4858
Taruya A., Nishimichi T., Jeong D., 2018, *Phys. Rev. D*, 98, 103532
Tatekawa T., 2014, *J. Cosmol. Astropart. Phys.*, 2014, 025
Tatekawa T., 2019, preprint (arXiv:1901.09210)
Tatekawa T., Mizuno S., 2007, *J. Cosmol. Astropart. Phys.*, 2007, 014
Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, *ApJ*, 688, 709
Tomlinson J., Jeong D., Kim J., 2019, *AJ*, 158, 116
Uhlemann C., Rampf C., Gosenca M., Hahn O., 2019, *Phys. Rev. D*, 99, 083524
Valageas P., 2020, *Phys. Rev. D*, 101, 123524
Valkenburg W., Villaescusa-Navarro E., 2017, *MNRAS*, 467, 4401
Warren M. S., Abazajian K., Holz D. E., Teodoro L., 2006, *ApJ*, 646, 881
Weltman A. et al., 2020, *PASA*, 37, e002
White S. D. M., 1996, in Schaeffer R., Silk J., Spiro M., Zinn-Justin J., eds, *Cosmology and Large Scale Structure*. Elsevier, Amsterdam, p. 349
Zel'dovich Ya. B., 1970, *A&A*, 500, 13
Zheligovsky V., Frisch U., 2014, *J. Fluid Mech.*, 749, 404

APPENDIX A: EXPLICIT FORMULAE FOR 2LPT AND 3LPT SCALAR AND VECTOR POTENTIALS

For convenience and reference, we give here the explicitly spelled out expressions for the n LPT potentials up to $n = 3$:

$$\begin{aligned} \nabla^2 \Phi^{(2)} = & \Phi_{,xx}^{(1)} (\Phi_{,yy}^{(1)} + \Phi_{,zz}^{(1)}) + \Phi_{,yy}^{(1)} \Phi_{,zz}^{(1)} \\ & - \Phi_{,xy}^{(1)} \Phi_{,xy}^{(1)} - \Phi_{,xz}^{(1)} \Phi_{,xz}^{(1)} - \Phi_{,yz}^{(1)} \Phi_{,yz}^{(1)}, \end{aligned} \quad (\text{A1})$$

$$\begin{aligned} \nabla^2 \Phi^{(3a)} = & \Phi_{,xx}^{(1)} \Phi_{,yy}^{(1)} \Phi_{,zz}^{(1)} + 2\Phi_{,xy}^{(1)} \Phi_{,xz}^{(1)} \Phi_{,yz}^{(1)} \\ & - (\Phi_{,yz}^{(1)})^2 \Phi_{,xx}^{(1)} - (\Phi_{,xz}^{(1)})^2 \Phi_{,yy}^{(1)} - (\Phi_{,xy}^{(1)})^2 \Phi_{,zz}^{(1)}, \end{aligned} \quad (\text{A2})$$

$$\begin{aligned} \nabla^2 \Phi^{(3b)} = & \frac{1}{2} \Phi_{,xx}^{(1)} (\Phi_{,yy}^{(2)} + \Phi_{,zz}^{(2)}) \\ & + \frac{1}{2} \Phi_{,yy}^{(1)} (\Phi_{,zz}^{(2)} + \Phi_{,xx}^{(2)}) + \frac{1}{2} \Phi_{,zz}^{(1)} (\Phi_{,xx}^{(2)} + \Phi_{,yy}^{(2)}) \\ & - \Phi_{,xy}^{(1)} \Phi_{,xy}^{(2)} - \Phi_{,xz}^{(1)} \Phi_{,xz}^{(2)} - \Phi_{,yz}^{(1)} \Phi_{,yz}^{(2)}, \end{aligned} \quad (\text{A3})$$

$$\begin{aligned} \nabla^2 A_x^{(3c)} = & \Phi_{,xy}^{(2)} \Phi_{,xz}^{(1)} - \Phi_{,xz}^{(2)} \Phi_{,xy}^{(1)} \\ & + \Phi_{,yz}^{(1)} (\Phi_{,yy}^{(2)} - \Phi_{,zz}^{(2)}) - \Phi_{,yz}^{(2)} (\Phi_{,yy}^{(1)} - \Phi_{,zz}^{(1)}) \end{aligned} \quad (\text{A4})$$

$$\begin{aligned} \nabla^2 A_y^{(3c)} = & \Phi_{,yz}^{(2)} \Phi_{,yx}^{(1)} - \Phi_{,yx}^{(2)} \Phi_{,yz}^{(1)} \\ & + \Phi_{,zx}^{(1)} (\Phi_{,zz}^{(2)} - \Phi_{,xx}^{(2)}) - \Phi_{,zx}^{(2)} (\Phi_{,zz}^{(1)} - \Phi_{,xx}^{(1)}) \end{aligned} \quad (\text{A5})$$

$$\begin{aligned} \nabla^2 A_z^{(3c)} = & \Phi_{,zx}^{(2)} \Phi_{,zy}^{(1)} - \Phi_{,zy}^{(2)} \Phi_{,zx}^{(1)} \\ & + \Phi_{,xy}^{(1)} (\Phi_{,xx}^{(2)} - \Phi_{,yy}^{(2)}) - \Phi_{,xy}^{(2)} (\Phi_{,xx}^{(1)} - \Phi_{,yy}^{(1)}). \end{aligned} \quad (\text{A6})$$

These expressions are equivalent to those in equation (6).

APPENDIX B: IMPACT OF ALIASING ERRORS ON LATE-TIME RESULTS

In Fig. B1, we show the impact of aliasing errors in the non-linear 2LPT and 3LPT terms on the power spectrum at late times. Generally and as expected, aliasing errors become increasingly more relevant with higher order non-linearities, so the errors are larger in 3LPT than in 2LPT. Since they only appear in those higher order LPT terms, but not in the ZA, and since these terms are always subdominant on larger scales, we expect the impact to be weakened by non-linear

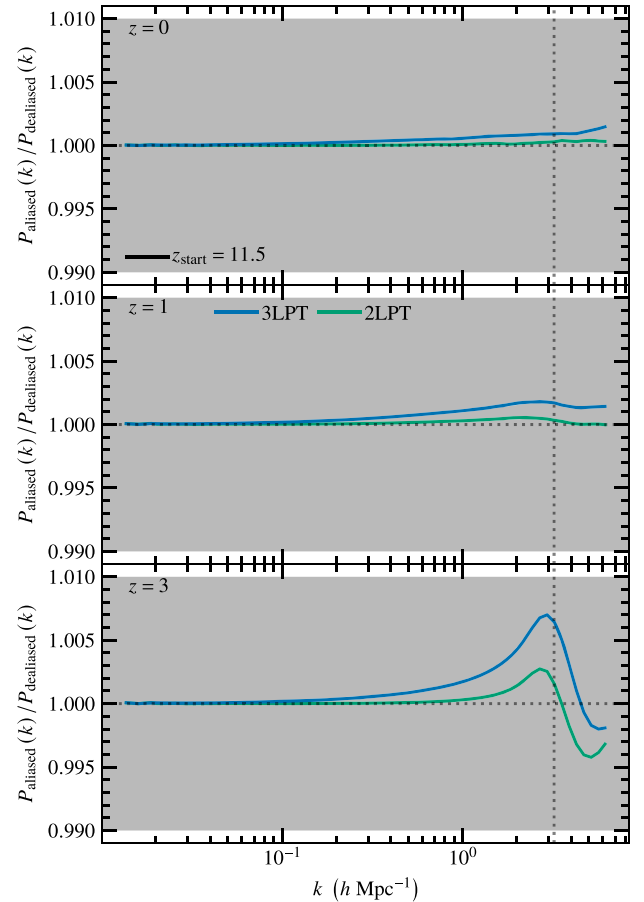


Figure B1. Comparison of low-redshift power spectra starting from initial conditions whose non-linearities have been properly de-aliased and those with aliasing errors for 2LPT and 3LPT starting at $z_{\text{start}} = 11.5$ to maximize the error. Aliasing errors appear as decaying mode transients and thus, up to 3LPT, impact the power spectrum at the sub per-cent level at redshifts below 3. As before, the shaded grey area indicates a 1 per cent error.

evolution and not amplified. This is indeed what the simulations show: We find that at $z = 0$, the error is only ~ 0.1 per cent in the power spectrum, while at $z = 3$ the error is just under 1 per cent, sharply peaking at the Nyquist mode. For the bispectrum, we arrive at fairly similar conclusions (not shown). It is thus most likely safe to work with an aliased version of 2LPT and 3LPT as long as one uses it as input for a fully non-linear simulation. Instead, if the output of nLPT was directly used, then the aliasing errors can be more significant, especially close to shell crossing.

APPENDIX C: DISCRETENESS EFFECTS AND LATTICE EVOLUTION

Following Joyce et al. (2005), Joyce & Marcos (2007), and Garrison et al. (2016), we consider the PLT of a perturbed simple cubic (SC) lattice. Since our methodology for correcting the spurious PLT excitations differs slightly from Garrison et al. (2016), we begin with a review of the PLT approach which allows us to discuss the subtleties of the numerical implementations.

For simplicity, let us assume an EdS universe and consider a simple cubic lattice of N particles, with identical mass, in a box with periodic boundary conditions. The resulting equation of motion can be written as

$$\partial_a^2 \mathbf{x}_i = -\frac{3}{2a} (\partial_a \mathbf{x}_i + \nabla_x \varphi_N), \quad (\text{C1})$$

with \mathbf{x}_i being the current comoving position of particle i with initial position \mathbf{q}_i , and φ_N is the *peculiar* gravitational potential, i.e. involving the usual subtraction of the overall background density; see e.g. section 4 in Marcos et al. (2006) for details. In PLT, we are primarily interested to investigate the linear effects that stem from the particles discretization of the Poisson source. Expanding the Poisson source to linear order one gets

$$\nabla_x \varphi_N(\mathbf{x}(\mathbf{q}_i)) = -\frac{1}{a} \sum_{\mathbf{q}_j} \mathcal{D}(\mathbf{q}_i - \mathbf{q}_j) \psi^{(1)}(\mathbf{q}_j, a) + \mathcal{O}(\psi^{(1)2}), \quad (\text{C2})$$

where \mathcal{D} is the so-called dynamical matrix known in solid state physics (cf. Marcos et al. 2006 and references therein).

In Fourier space, the evolution equation for the linear displacement can then be written as (suppressing the particle index)

$$a^2 \partial_a^2 \tilde{\psi}^{(1)}(\mathbf{k}) + (3a/2) \partial_a \tilde{\psi}^{(1)}(\mathbf{k}) - (3/2) \tilde{\mathcal{D}}(\mathbf{k}) \tilde{\psi}^{(1)}(\mathbf{k}) = 0. \quad (\text{C3})$$

To make progress, we follow Marcos et al. (2006) and solve the above equation with the *Ansatz*

$$\tilde{\psi}^{(1)}(\mathbf{k}, a) = \sum_{n=1}^3 \hat{e}_n(\mathbf{k}) \tilde{\psi}_n(\mathbf{k}, a), \quad (\text{C4})$$

leading to an eigenvalue problem as well as an ODE for $\tilde{\psi}_n$, which respectively read

$$\begin{aligned} \tilde{\mathcal{D}}(\mathbf{k}) \hat{e}_n(\mathbf{k}) &= \varepsilon_n(\mathbf{k}) \hat{e}_n(\mathbf{k}), \\ a^2 \partial_a^2 \tilde{\psi}_n + (3a/2) \partial_a \tilde{\psi}_n - (3/2) \varepsilon_n(\mathbf{k}) \tilde{\psi}_n &= 0, \end{aligned} \quad (\text{C5})$$

where the eigenvalues ε_n can be determined using standard linear algebra programmes (see further below for details). The ODE has the general solution

$$\tilde{\psi}_n = C_n^+ a^{\alpha_n^+}(\mathbf{k}) + C_n^- a^{\alpha_n^-}(\mathbf{k}), \quad (\text{C6})$$

where C_n^{\pm} are integration constants, and $\alpha_n^{\pm}(\mathbf{k}) = (1/4)[-1 \pm \sqrt{1 + 24\varepsilon_n(\mathbf{k})}]$. In the following two paragraphs, we shall enforce that the PLT solution is compatible with purely growing mode and

curl-free solutions. After that, we outline some technical details about the numerical computation of the dynamical matrix.

C1 Demanding curl-free solutions

One of the eigenmodes called \hat{e}_{\parallel} is exactly parallel to $\hat{\mathbf{k}}$ in the fluid limit, and thus, \hat{e}_{\parallel} is associated with gravity induced modes. The other two eigenmodes are orthonormal to \hat{e}_{\parallel} , and, in the fluid limit, are associated with transverse modes that are vectorial transients and fictitious in the fluid sense, due to the assumed irrotational character of CDM. Thus, to initialize N -body simulations, we need to make sure that only the lattice modes along \hat{e}_{\parallel} are excited; to achieve this, we rescale the \mathbf{k} vectors appearing in the LPT fluid solution $\tilde{\psi}_{\text{fluid}} = i\mathbf{k}\tilde{\phi} + i\mathbf{k} \times \tilde{\mathbf{A}}$ according to (cf. Garrison et al. 2016)

$$\mathbf{k} \rightarrow \hat{e}_{\parallel}/(\hat{e}_{\parallel} \cdot \hat{\mathbf{k}}), \quad (\text{C7})$$

where the division of the scalar product is included to ensure that the matter power spectrum has unchanged amplitudes. This procedure, applied to ZA, has been first outlined by Garrison et al. (2016), although we note that we apply this rescaling also beyond ZA.

C2 Demanding purely growing-mode solutions

In the fluid limit, it can be shown that $\alpha_{\parallel}^+ \rightarrow 1$ and $\alpha_{\parallel}^- \rightarrow -3/2$ and thus, the solution of the ODE that is proportional to C_- is a decaying one, and thereby not compatible with growing-mode solutions. To enforce growing-mode solutions, we first evaluate $\tilde{\psi}_{\parallel} = C_{\parallel}^+ a^{\alpha_{\parallel}^+} + C_{\parallel}^- a^{\alpha_{\parallel}^-}$ and its first time derivative at initial time a_{ini} , and express the integration constants C_{\parallel}^{\pm} in terms of $\tilde{\psi}_{\parallel}(a_{\text{ini}}) \equiv \tilde{\psi}_{\text{ini}}$ and $\dot{\tilde{\psi}}_{\parallel}(a_{\text{ini}}) \equiv \tilde{v}_{\text{ini}}$. We find that

$$C_{\parallel}^- \propto \frac{\tilde{v}_{\text{ini}}}{\alpha_{\parallel}^+} - \frac{\tilde{\psi}_{\text{ini}}}{a_{\text{ini}}}, \quad (\text{C8})$$

and thus, decaying modes in PLT are zero provided that the initial PLT velocity and displacement satisfy $\tilde{v}_{\text{ini}}/\alpha_{\parallel}^+ = \tilde{\psi}_{\text{ini}}/a_{\text{ini}}$. This relation implies that only a single initial potential (e.g. v_{ini} , ψ_{ini} , or the initial gravitational potential) needs to be specified; this is precisely expected for slaved initial conditions. Comparing this with the fluid case for which the slaved initial conditions imply

$$C_{\text{fluid}}^- \propto \tilde{v}_{\text{ini}}^{\text{fluid}} - \frac{\tilde{\psi}_{\text{ini}}^{\text{fluid}}}{a_{\text{ini}}} \rightarrow 0, \quad (\text{C9})$$

where $\mathbf{v}^{\text{fluid}} = \nabla v^{\text{fluid}}$ and $\psi^{\text{fluid}} = \nabla \psi^{\text{fluid}}$, we realize that growing-mode PLT initial solutions in the N -body simulation are established provided we rescale the fluid velocity according to

$$\tilde{\mathbf{v}}_{\text{ini}}^{\text{fluid}} \rightarrow \tilde{\mathbf{v}}_{\text{ini}}^{\text{fluid}}/\alpha_{\parallel}^+. \quad (\text{C10})$$

We found that this somewhat ad hoc modification, demonstrated to be strictly correct at leading order, works very well also in combination with higher-order LPT, as demonstrated by our results shown in Figs 2 and 3.

C3 Numerical evaluation of dynamical matrix

Following Marcos et al. (2006), we compute the dynamical matrix \mathcal{D}_{ij} using Ewald summation, i.e. we split \mathcal{D} into a short-range and a long-range contribution, $\mathcal{D}_{ij} = \mathcal{D}_{ij}^{\text{sr}} + \mathcal{D}_{ij}^{\text{lr}}$ with splitting parameter κ . Specifically for an SC lattice, we compute on a grid of coordinates

in real space, with $r := \|\mathbf{x}\|$

$$\mathcal{D}_{ij}^{\text{sr}}(\mathbf{x} \neq \mathbf{0}) = -\frac{\kappa^3}{\pi^{3/2}} \frac{x_i x_j}{r^2} e^{-\kappa^2 r^2} + \frac{1}{4\pi} \left(\frac{\delta_{ij}}{r^3} - 3 \frac{x_i x_j}{r^5} \right) \times \left[\text{erfc}(\kappa r) + \frac{2\kappa}{\sqrt{\pi}} e^{-\kappa^2 r^2} \right] \quad (\text{C11})$$

and we set $\mathcal{D}_{ij}^{\text{sr}}(\mathbf{x} = \mathbf{0}) = 0$. We then Fourier-transform this short-range response function using the FFT. The respective long-range component is given in Fourier space as

$$\tilde{\mathcal{D}}_{ij}^{\text{lr}}(\mathbf{k} \neq \mathbf{0}) = \frac{\delta_{ij}}{3} + \frac{k_i k_j}{\|\mathbf{k}\|^2} e^{-\|\mathbf{k}\|^2/(4\kappa^2)}, \quad \tilde{\mathcal{D}}_{ij}^{\text{lr}}(\mathbf{k} = \mathbf{0}) = \frac{\delta_{ij}}{3}. \quad (\text{C12})$$

We use a cut-off scale corresponding to four cells, i.e. $\kappa = (\sqrt{2}4\Delta x)^{-1}$. We sum short and long-range components to obtain the full dynamical matrix $\mathcal{D}_{ij}(\mathbf{k})$ in Fourier space, then diagonalize it to obtain the lattice growing mode eigenvectors \mathbf{e}_\parallel with associated eigenvalue ϵ_\parallel . Since the dynamical matrix is very smooth, and its computation for the actual full-resolution lattice relatively costly, we compute a relatively low resolution version of it (in fact 64^3) and then interpolate to a higher resolution lattice (as also Garrison et al. 2016 are suggesting).

APPENDIX D: INVARIANCE CONSERVATION AS CONVERGENCE AND CORRECTNESS TEST

Testing numerical implementations of LPT in a realistic cosmological set-up (i.e. with a wide spectrum of modes) is inherently difficult. Here we propose a new approach: we know that LPT violates certain symmetries of the fluid equations perturbatively beyond the ZA. Specifically, Uhlemann et al. (2019) have shown that beyond ZA, fixed-order LPT at order n excites artificial vortical modes at order $n + 1$. The relevant equations for this are the Cauchy invariants, which are the local form of Kelvin's circulation theorem. In the cosmological context the Cauchy invariants state the conservation of the zero vorticity, which is satisfied at all times for individual fluid particles. The three invariants are ($p = 1, 2, 3$)

$$C_p = \varepsilon_{pln} v_{k,l} x_{k,n} = 0. \quad (\text{D1})$$

These vectorial relations can be written in tensor form (thereby making contact with similar theorems in Hamiltonian mechanics). Contracting C_p with the Levi-Civita symbol ε_{pij} yields the antisymmetric tensor

$$\mathcal{I}_{ij} := \varepsilon_{pij} C_p = v_{k,i} x_{k,j} - v_{k,j} x_{k,i} = v_{j,i} - v_{i,j} + v_{k,i} \psi_{k,j} - v_{k,j} \psi_{k,i} = 0, \quad (\text{D2})$$

where we used that $x_{k,j} = \delta_{kj} + \psi_{k,j}$. Equation (D2) vanishes non-perturbatively. However, only in the case of the Zel'dovich approximation, \mathcal{I}_{ij} is exactly zero, when v_i and ψ_i are truncated at the perturbative order. As mentioned before, for n th order LPT with $n > 1$, the truncation error is of order $n + 1$ (cf. Uhlemann et al. 2019). Since the Cauchy invariants contain a time derivative through $v_k = \dot{x}_k$, however, the resulting truncation error grows with D_+^n (and not as D_+^{n+1}). We show the RMS amplitude of \mathcal{I}_{xy} for our implementation as a function of the starting time D_+ in equation (D1). We see that only when de-aliasing the non-linearities, we get the expected scaling of \mathcal{I}_{ij} with D_+^3 in 3LPT demonstrating the correctness of our implementation.

On a technical note, numerically, in both cases, we applied 3/2 padding when evaluating equation (D2), but compute v_i and ψ_i with and without padding for second and higher order non-linearities to test the effect of aliasing on velocities and displacements.

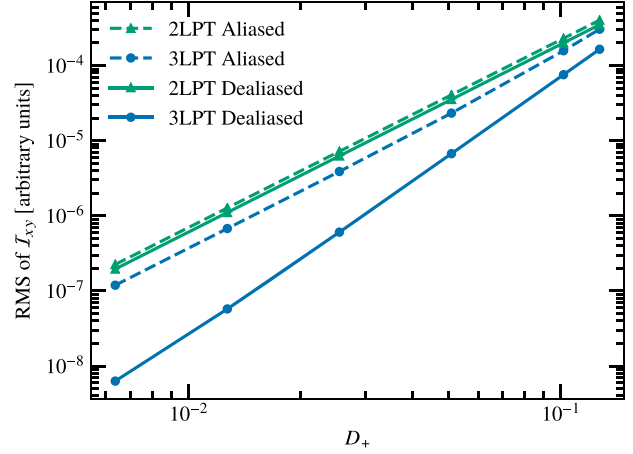


Figure D1. RMS amplitude of the violation of the Cauchy invariants (D2) as a function of starting time D_+ for 2LPT and 3LPT. We plot the standard deviation of the \mathcal{I}_{xy} field for a cosmological $1 h^{-1}$ Gpc box with 1024^3 modes and particles. The amplitudes for \mathcal{I}_{yz} and \mathcal{I}_{xz} are identical. We generated the LPT fields using both Orszag's 3/2 rule for de-aliasing (solid lines), and also not applying any de-aliasing procedure (dashed lines). Clearly only the de-aliased results scale $\mathcal{O}(D_+^3)$ for 3LPT, in all other cases we find behaviour $\mathcal{O}(D_+^2)$ that is expected for 2LPT. Note that \mathcal{I}_{ij} vanishes exactly for the Zel'dovich approximation.

Note that the lines for 2LPT and 3LPT in Fig. D1 intersect close to $a \sim 0.1$ which is further evidence that this is the point where LPT breaks down, consistent with our results in Section 3.

APPENDIX E: DETAILS TO OBTAINING THE THEORETICAL BOUND ON THE RADIUS OF CONVERGENCE

Here, we outline the derivation of the bounds on the theoretical radius of convergence of LPT, Section 3.3.

E1 Most conservative estimate

Details in this subsection are based on Zheligovsky & Frisch (2014) and Rampf et al. (2015). Beyond first order, the LPT recursion relations express the (to be determined) n th-order displacement gradients coefficients $\psi_{i,j}^{(n)}(\mathbf{q})$ in terms of quadratic and cubic combinations of (known) displacement gradients from lower orders. These quadratic and cubic combinations of displacements contain numerical coefficients that are all bounded by unity from below. Furthermore, these quadratic and cubic displacement terms contain operators of the kind $\nabla^{-2} \partial_i \partial_j$; applying suitable norms on the recursion relation of the displacement gradients, i.e. by considering the recursion relation for $\|\psi_{i,j}^{(n)}\|$, it can be shown that also such operators are bounded by unity. Exploiting those two bounds by setting the coefficients and operators effectively to unity, the recursion relation for the coefficient $\|\psi_{i,j}^{(n)}\|$ turns into a cubic *polynomial inequality*. Upon reintroducing the time variable D_+ by using the generating function $\Psi \equiv \sum_n \|\psi_{i,j}^{(n)}\| D_+^n$, this polynomial inequality turns into¹⁰

$$p_{\text{low}}(D_+, \Psi) := 6\Psi^3 + 12\Psi^2 - \Psi + D_+ \|\nabla_i \nabla_j \varphi_{\text{ini}}\| \geq 0, \quad (\text{E1})$$

¹⁰For simplicity we assume here an Einstein-de Sitter universe which is a sufficient approximation for our purpose. Rampf et al. (2015) generalized the presented results to the Λ CDM universe, however found only a weak dependence of Λ on the lower bound of convergence.

where φ_{ini} is given by equation (14). For small Ψ 's, the polynomial $p_{\text{low}}(0, \Psi)$ behaves as $-\Psi$ while asymptotically it behaves like Ψ^3 . Furthermore, for $D_+ > 0$, the cubic polynomial is shifted upwards linearly in D_+ along the y -axis, and has first three roots and later on a single root. To investigate the physical branch, let us begin at $D_+ = 0$ for which there is a root crossing the origin, in accordance with a zero displacement initially and thus this root denotes the physical branch. For times shortly after, this root moves from zero to a small positive value, and furthermore the whole branch between 0 and that root is positive, thereby satisfying the positivity condition of (E1) and thus marking the physical branch. At some critical time later, denoted D_+^{crit} , this root will merge with another one. Just shortly after D_+^{crit} , these two roots disappear (i.e. by becoming complex), implying that p loses its boundedness and so does the displacement field.¹¹ Thus, D_+^{crit} marks the last possible time for which the displacement is bounded, and thus is the maximal time for which convergence is still guaranteed, $D_+^{\text{crit}} = D_+^{\text{theory}}$. That maximal time can be determined by setting the discriminant of $p_{\text{low}}(\Psi)$ to zero. This leads to the positive critical time

$$D_+^{\text{theory}} = \frac{T}{\|\nabla_i \nabla_j \varphi_{\text{ini}}\|}, \quad (\text{E2})$$

with $T = 0.0204$ which, as promised, is a lower bound on the radius of convergence $D_+^{\text{theory}} < R$. In the following, we show how this bound can be improved.

E2 Conservative, third-order estimate

As outlined above, before obtaining the polynomial inequality (E1) from which the lower bound on the convergence is retrieved, we set all numerical coefficients in front of the n th-order displacement to unity, simply since those coefficients are bounded by unity. However, we do have explicit numerical coefficients up to 3LPT (cf. equation 3). A refinement for the theoretical bound is therefore obtained by keeping some of these coefficients as they really are, and only 'set' the higher-order coefficients to unity. Doing so leads after some straightforward calculations to the modified polynomial inequality

$$p_{3\text{LPT}}(D_+, \Psi) := 6\Psi^3 + 12\Psi^2 - \Psi + T - \frac{24}{7}T^2 - \frac{86}{7}T^3 \geq 0, \quad (\text{E3})$$

where $T := D_+ \|\nabla_i \nabla_j \varphi_{\text{ini}}\|$, which, by identification of the result from the previous section leads to the equation $T - 24T^2/7 - 86T^3/7 = 0.0204$, leading to

$$T =: T_{\text{cons}} = 0.022 \quad (\text{E4})$$

for equation (E2), a 9 per cent improvement regarding the previous section.

E3 Optimistic (shell-crossing) estimate

Observe that shell-crossing occurs when the Jacobian $J = \det[\delta_{ij} + \psi_{i,j}] = 1 + \nabla \cdot \psi + (1/2)[\psi_{i,i}\psi_{j,j} - \psi_{i,j}\psi_{i,j}] + \det[\psi_{i,j}]$ vanishes for the first time. As it was shown by Rampf et al. (2015), by the same methods as explained in Appendix E1, we can obtain a bound on the vanishing Jacobian by introducing the

¹¹The unboundedness of the polynomial p_{low} essentially implies that the higher-order displacement coefficients do not decrease at subsequent orders – which is a clear indication that the series lost its convergence property.

generating function $\Psi = \sum_n \|\psi_{i,j}^{(n)}\| D_+^n$. One finds

$$|J - 1| \leq 6\Psi^3 + 6\Psi^2 + 3\Psi \equiv Q(\Psi). \quad (\text{E5})$$

Now, from the bound as obtained in Appendix E1, one can easily determine the corresponding critical value $\Psi_{\text{crit}} \simeq 0.0404$ for which the bound (E2) holds. Plugging this value in equation (E5), one finds $|J - 1| \leq 0.132$, which is well below unity for which shell-crossing occurs. Thus, for the conservative bound D_+^{theory} , shell-crossing is not reached. Therefore, we seek here in this section to obtain stronger bounds, thereby allowing us to reach shell-crossing closely from below.

In the main text, we have mentioned that the first singularity in LPT could be in the complex time-domain; however, recently Saga, Taruya & Colombi (2018) provided solid numerical evidence that LPT converges until shell-crossing. Specifically, they estimated the radius of convergence for sine-wave initial conditions in 3D, by using LPT to 10th order. Comparison to Vlasov–Poisson simulations revealed indeed that LPT appears to converge until shell-crossing, excluding the spherical (or symmetric sine-wave) collapse that is degenerate in a universe with random initial conditions.

Thus, in our language, a complex singularity is unlikely to be decisive for determining the radius of convergence for sine-wave ICs, and the same should also be true for generic, Gaussian initial conditions.

Therefore, in this appendix we *assume* that complex singularities are ruled out, and that the convergence-determining factor is indeed shell-crossing. One reason why the methods in Appendix E1 did not allow us to come closer to the first shell-crossing, is the failure to provide better estimates for the all-order coefficients and spatial derivatives, that in particular sensitively influence some prefactors within the polynomial inequality, equation (E1). Specifically, the following altered polynomial inequality – which is not based on a theory but rather on intuition – would be in principle closer to the true inequality for the generating function,

$$p_{\text{cross}}(D_+, \Psi) := 6a\Psi^3 + 12b\Psi^2 - \Psi + T_{\text{cross}} \geq 0, \quad (\text{E6})$$

where a, b are positive but a priori unknown coefficients, whereas T_{cross} is the critical coefficient to be determined, since it sets the (possibly strongest) bound on the time of shell-crossing.

Now, we first search for the maximal value of Ψ_{max} for which the bound on the Jacobian just begins to touch the unity bound, i.e. $Q(\Psi_{\text{max}}) = 1$. We find $\Psi_{\text{max}} = 0.2178$. If the p_{cross} has its critical value Ψ_{crit} for which its discriminant vanishes at precisely Ψ_{max} , then T_{cross} would be maximized. Thus, solving for $p_{\text{cross}}(\Psi_{\text{max}}) = 0 = p'_{\text{cross}}(\Psi_{\text{max}})$ leads to updated constraints. Unfortunately, the system of equations is still under determined – in fact there is a redundant free parameter (e.g. either a or b), thereby seemingly rendering the considered task as unsolvable.

Instead of solving explicitly for the two coefficients a and b , however, we could simply scan through the allowed parameter space where shell-crossing is ruled out, as well the two constrains $a > 0$ and $b > 0$ are satisfied. Scanning through the parameter space, we arrive at the (optimistic) theoretical shell-crossing constraint

$$D_+^{\text{cross}} = \frac{T_{\text{cross}}}{\|\nabla_i \nabla_j \varphi_{\text{ini}}\|}, \quad (\text{E7})$$

with the worst-case coefficient in the allowed regime to be $T_{\text{cross}} = 0.107$.