

# PySCo: A fast particle-mesh $N$ -body code for modified gravity simulations in Python

Michel-Andrès Breton<sup>1,2,3,\*</sup>

<sup>1</sup> Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191, Gif-sur-Yvette, France

<sup>2</sup> Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain

<sup>3</sup> Laboratoire Univers et Théorie, Observatoire de Paris, Université PSL, Université Paris Cité, CNRS, 92190 Meudon, France

Received 27 October 2024 / Accepted 18 February 2025

## ABSTRACT

We present PySCo, a fast and user-friendly Python library designed to run cosmological  $N$ -body simulations across various cosmological models, such as  $\Lambda$ CDM ( $\Lambda$  with cold dark matter) and  $w_0w_a$ CDM, and alternative theories of gravity, including  $f(R)$ , MOND (modified newtonian dynamics) and time-dependent gravitational constant parameterisations. PySCo employs particle-mesh solvers, using multigrid or fast Fourier transform (FFT) methods in their different variations. Additionally, PySCo can be easily integrated as an external library, providing utilities for particle and mesh computations. The library offers key features, including an initial condition generator based on up to third-order Lagrangian perturbation theory (LPT), power spectrum estimation, and computes the background and growth of density perturbations. In this paper, we detail PySCo's architecture and algorithms and conduct extensive comparisons with other codes and numerical methods. Our analysis shows that, with sufficient small-scale resolution, the power spectrum at redshift  $z = 0$  remains independent of the initial redshift at the 0.1% level for  $z_{\text{ini}} \geq 125, 30$ , and 10 when using first, second, and third-order LPT, respectively. Moreover, we demonstrate that acceleration (or force) calculations should employ a configuration-space finite-difference stencil for central derivatives with at least five points, as three-point derivatives result in significant power suppression at small scales. Although the seven-point Laplacian method used in multigrid also leads to power suppression on small scales, this effect can largely be mitigated when computing ratios. In terms of performance, PySCo only requires approximately one CPU hour to complete a Newtonian simulation with  $512^3$  particles (and an equal number of cells) on a laptop. Due to its speed and ease of use, PySCo is ideal for rapidly generating vast ensemble of simulations and exploring parameter spaces, allowing variations in gravity theories, dark energy models, and numerical approaches. This versatility makes PySCo a valuable tool for producing emulators, covariance matrices, or training datasets for machine learning.

**Key words.** methods: numerical – cosmology: miscellaneous – large-scale structure of Universe

## 1. Introduction

On the largest scales, the  $\Lambda$ CDM ( $\Lambda$  with cold dark matter) model provides a robust description of our Universe, where dark energy is represented as a cosmological constant ( $\Lambda$ ) and dark matter as non-relativistic cold dark matter (Planck Collaboration VI 2020; Alam et al. 2021; Riess et al. 2022). The  $\Lambda$ CDM model also assumes General Relativity (GR), which offers an accurate representation of the Universe by incorporating small perturbations on a Friedman–Lemaître–Robertson–Walker (FLRW) background (Green & Wald 2014). The formation of the large-scale structure is driven by these small perturbations, which grow over time through gravitational interactions, giving rise to the cosmic web observed today (Peebles 1980).

While linear theory can analytically describe the growth of perturbations on large scales above  $\sim 100$  Mpc (Yoo et al. 2009; Bonvin & Durrer 2011; Challinor & Lewis 2011), the non-linear nature of structure formation on smaller scales necessitates more advanced approaches. Higher-order perturbation theory, both in its Eulerian (Peebles 1980; Bernardeau et al. 2002) and Lagrangian (Zel'dovich 1970; Buchert & Ehlers 1993) formulations, allows for a more accurate understanding down to scales of  $\sim 20$  Mpc. Below this scale, structure formation becomes highly non-linear, and the evolution of gravitationally interacting

particles can only be accurately modeled through  $N$ -body simulations (Efstathiou et al. 1985).

Over time, numerous cosmological  $N$ -body codes have been developed (Couchman 1991; Kravtsov et al. 1997; Knebe et al. 2001; Teyssier 2002; Ishiyama et al. 2009; Potter et al. 2017; Springel et al. 2021; Garrison et al. 2021), with recent simulations including trillions of particles (Ishiyama et al. 2021; Euclid Collaboration: Castander et al. 2025). While significant effort has been invested in developing efficient Newtonian simulations for cosmology (see Angulo & Hahn 2022 for a review), many  $N$ -body codes for alternative gravity theories are derived from Newtonian codes. For instance, ECOSMOG (Li et al. 2012), ISIS (Llinares et al. 2014), RayMOND (Candlish et al. 2015), and Phantom of RAMSES (Lüghausen et al. 2015) are all based on the RAMSES code (Teyssier 2002). The former two implement the  $f(R)$  model (Hu & Sawicki 2007) and the nDGP model (Dvali et al. 2000), while the latter two implement MOND (modified newtonian dynamics, Milgrom 1983). Although it is possible to integrate modified gravity theories into TreePM codes like Gadget (Springel et al. 2021), as demonstrated by MG-Gadget (Puchwein et al. 2013), RAMSES has emerged as the preferred code for such implementations. This is due to the fact that alternative theories of gravity generally introduce additional fields governed by non-linear partial differential equations, which cannot be efficiently solved using standard

\* Corresponding author; michel-andres.breton@cea.fr

tree-based methods. As a result, RAMSES, a particle-mesh (PM) code with adaptive-mesh refinement (AMR) and a multigrid solver, becomes an ideal choice for implementing these features. However, PM-AMR, TreePM, or Fast Multipole Method (FMM) codes typically require significant computational resources, even when highly optimised. This challenge is exacerbated in simulations of alternative gravity theories; for instance, [Euclid Collaboration: Adamek et al. \(2025\)](#) showed that  $f(R)$  or nDGP simulations can run up to ten times slower than their Newtonian counterparts.

To address these computational demands, researchers have explored PM codes ([Knebe et al. 2001](#); [Merz et al. 2005](#); [Feng et al. 2016](#); [Adamek et al. 2016](#); [Klypin & Prada 2018](#)) to reduce the cost at the expense of small-scale accuracy. These faster, albeit less precise, simulations are suitable for specific applications, such as producing large numbers of realisations for covariance matrix estimation or calculating the ratio (or boost) of specific statistical quantities relative to a reference case exploring large parameter spaces. Furthermore, exact PM codes can accurately reproduce structure formation at small scales given a higher resolution of the uniform mesh. Consequently, such codes have been ideal for developing solvers for modified gravity theories (see [Llinares 2018](#) for a review), and have been used to study their impact relative to Newtonian simulations ([Valogiannis & Bean 2017](#); [Winther et al. 2017](#); [Hassani & Lombriser 2020](#); [Ruan et al. 2022](#); [Hernández-Aguayo et al. 2022](#)).

Given the need for speed, most  $N$ -body simulations have traditionally been written in compiled languages such as Fortran, C, or C++. In contrast, Python has become the most popular programming language in data science, owing to its straightforward syntax, rapid development speed, and extensive community-driven libraries, especially in astronomy. This popularity has created a gap between simulators and the broader scientific field. Despite its advantages, Python is often viewed as a slow language when used natively. To address this, significant effort has gone into developing efficient Python libraries, either as wrappers for C-based codes (such as NumPy, [Harris et al. 2020](#)) or through compiling Python code to machine language, as seen in Numba ([Lam et al. 2015](#)) and Cython ([Behnel et al. 2011](#)). Recently, the latter approach was used to develop the  $P^3M$  code CONCEPT ([Dakin et al. 2022](#)), demonstrating the viability of Python for high-performance applications.

In this paper, we present PySCo<sup>1</sup> (Python Simulations for Cosmology), a cosmological PM  $N$ -body code written in Python and utilising the Numba library for increased performance and multithreading. The paper is organised as follows: Section 2 introduces the different models implemented in PySCo, including the modified gravity theories  $f(R)$  from [Hu & Sawicki \(2007\)](#), MOND ([Milgrom 1983](#)), parameterised gravity ([Amendola et al. 2008](#)), and dynamic dark energy ([Chevallier & Polarski 2001](#); [Linder 2003](#)). Section 3 details the structure and algorithms implemented in the code, covering initial condition generation and various  $N$ -body solvers. In Section 4, we validate PySCo against other codes and analyse the impact of different numerical methods on the matter power spectrum. Finally, we conclude in Section 5.

## 2. Theory

### 2.1. Newtonian gravity

Let us consider only scalar perturbations on FLRW metric, in the Newtonian gauge ([Ma & Bertschinger 1995](#))

$$ds^2 = a^2(\eta) \left[ -(1 + 2\psi) d\eta^2 + (1 - 2\phi) dx^2 \right], \quad (1)$$

with  $a$  the scale factor,  $\eta$  the conformal time and  $\psi$  and  $\phi$  the Bardeen potentials ([Bardeen 1980](#)). In GR, we have  $\psi = \phi$ , and the Einstein equation gives

$$\nabla^2 \phi - 3\mathcal{H}(\phi' + \mathcal{H}\phi) = 4\pi G a^2 \delta \rho, \quad (2)$$

with  $\mathcal{H} = aH$  the conformal Hubble parameter,  $\phi'$  the time derivative of the potential and  $\rho$  the Universe's components density. It is common to apply the Newtonian and quasi-static (neglecting time derivatives) approximations, which involve neglecting the second term on the left-hand side of Eq. (2), as it is only significant at horizon scales. In the context of Newtonian cosmology, the Einstein equation takes the same form as the classical Poisson equation, with an additional dependence on the scale factor

$$\nabla^2 \phi = 4\pi G a^2 \delta \rho_m, \quad (3)$$

with  $\rho_m$  the matter density. There are, however, cosmological General Relativity (GR) simulations ([Adamek et al. 2016](#); [Barrera-Hinojosa & Li 2020](#)) that solve the full Einstein equations in the weak-field limit, taking into account gauge issues ([Fidler et al. 2015, 2016](#)). Additionally, there are methods to interpret Newtonian simulations within a relativistic framework ([Chisari & Zaldarriaga 2011](#); [Adamek & Fidler 2019](#)). In this work, however, we focus on smaller scales where relativistic effects are negligible, and the Newtonian approximation remains well justified.

### 2.2. Dynamical dark energy

The  $w_0 w_a$ CDM model provides a useful phenomenological extension of the standard  $\Lambda$ CDM framework by offering a dynamic, time-dependent description of dark energy. In this model, the cosmological constant is replaced by a variable dark energy component, which affects the formation of cosmic structures by modifying the universe's expansion history. For a flat geometry ( $\Omega_k = 0$ ), the Hubble parameter  $H(z)$  is given by

$$H(z) = H_0 \sqrt{\Omega_{m,0}(1+z)^3 + \Omega_{r,0}(1+z)^4 + \Omega_\Lambda(z)}, \quad (4)$$

where a subscript zero indicates a present-day evaluation, and  $\Omega_m$ ,  $\Omega_r$ , and  $\Omega_\Lambda$  represent the density fractions of matter, radiation, and dark energy, respectively. In this model, the dark-energy density evolves with redshift, following the relation

$$\Omega_\Lambda(z) = \Omega_{\Lambda,0} \exp \left\{ \int_0^z \frac{3[1+w(z')] dz'}{1+z'} \right\}. \quad (5)$$

Using the widely adopted CPL parametrisation ([Chevallier & Polarski 2001](#); [Linder 2003](#)), the dark-energy equation of state is expressed as

$$w(z) = w_0 + w_a \frac{z}{1+z}. \quad (6)$$

This simple modification does not alter the Einstein field equations nor the equations of motion, and it recovers the standard  $\Lambda$ CDM model when  $w_0 = -1$  and  $w_a = 0$ .

<sup>1</sup> <https://github.com/mianbreton/pysco>

### 2.3. MOND gravity

The Modified Newtonian Dynamics (MOND) theory, introduced by [Milgrom \(1983\)](#), was proposed as a potential solution to the dark matter problem by suggesting a deviation from Newtonian gravity in a Universe where the matter content is entirely in the form of baryons (for a detailed review, see [Famaey & McGaugh 2012](#)). In MOND, Newton's second law is modified as follows

$$\mu \left( \frac{|\mathbf{g}|}{g_0} \right) \mathbf{g} = \mathbf{g}_N, \quad (7)$$

where  $\mathbf{g}$  and  $\mathbf{g}_N$  represent the MOND and Newtonian accelerations, respectively, and  $g_0$  is a characteristic acceleration scale, approximately  $g_0 \approx cH_0/[2\pi] \approx 10^{-10} \text{ m.s}^{-2}$ . The function  $\mu(x)$  is an interpolating function that governs the transition between the Newtonian regime (where gravitational force scales as  $r^{-2}$ ) and the MOND regime (where the force scales as  $r^{-1}$ ), with  $r$  being the separation between masses. The interpolating function has the following limits

$$\mu(x) = \begin{cases} 1, & x \gg 1, \\ x, & x \ll 1. \end{cases} \quad (8)$$

Similarly, the inverse interpolating function  $\nu(y)$  can be defined as

$$\mathbf{g} = \nu \left( \frac{|\mathbf{g}_N|}{g_0} \right) \mathbf{g}_N, \quad (9)$$

where  $\nu(y)$  follows the limits

$$\nu(y) = \begin{cases} 1, & y \gg 1, \\ y^{-1/2}, & y \ll 1. \end{cases} \quad (10)$$

In the MOND framework, the classical Poisson equation is modified as follows ([Bekenstein & Milgrom 1984](#))

$$\nabla \left[ \mu \left( \frac{|\nabla \phi|}{g_0} \right) \nabla \phi \right] = 4\pi G \delta \rho, \quad (11)$$

a formulation known as AQUAL, derived from a quadratic Lagrangian. In this paper, however, we consider the QUMOND (quasi-linear MOND, [Milgrom 2010](#)) formulation, where the non-linearity in the Poisson equation is re-expressed in terms of an additional effective dark matter fluid in the source term. The modified Poisson equations are

$$\nabla^2 \phi^N = 4\pi G \rho, \quad (12)$$

$$\nabla^2 \phi = \nabla \left[ \nu \left( \frac{|\nabla \phi^N|}{g_0} \right) \nabla \phi^N \right], \quad (13)$$

where  $\phi^N$  and  $\phi$  are the Newtonian and MOND potentials, respectively. This system of equations is more convenient to solve, as it involves only two linear Poisson equations. Moreover, numerical simulations have demonstrated that the AQUAL and QUMOND formulations yield very similar results ([Candlish et al. 2015](#)).

The missing component in the MOND framework is the specific form of the  $\nu(y)$  function. Several families of interpolating functions have been proposed, each with different characteristics

– Simple function:

$$\nu(y) = \frac{1}{2} + \frac{\sqrt{1+4/y}}{2}, \quad (14)$$

which corresponds to the simple function proposed by [Famaey & Binney \(2005\)](#), equivalent to  $\mu(x) = x/(1+x)$ .

– The  $n$ -family:

$$\nu(y) = \left[ \frac{1}{2} + \frac{\sqrt{1+4/y^n}}{2} \right]^{1/n}, \quad (15)$$

a commonly used parametrisation for the interpolating function ([Milgrom & Sanders 2008](#)). The case  $n = 2$  is particularly well studied and is known as the standard interpolating function ([Begeman et al. 1991](#)). Additionally, [Milgrom & Sanders \(2008\)](#) introduced other functional forms

– The  $\beta$ -family:

$$\nu(y) = (1 - e^{-y})^{-1/2} + \beta e^{-y}, \quad (16)$$

– The  $\gamma$ -family:

$$\nu(y) = (1 - e^{-y^{\gamma/2}})^{-1/\gamma} + (1 - \gamma^{-1}) e^{-y^{\gamma/2}}, \quad (17)$$

– The  $\delta$ -family:

$$\nu(y) = (1 - e^{-y^{\delta/2}})^{-1/\delta}, \quad (18)$$

which is a subset of the  $\gamma$ -family. While we focus on non-relativistic formulations of MOND for simplicity, it is important to acknowledge that various relativistic frameworks have been developed ([Bekenstein 2006; Milgrom 2009; Skordis & Złoński 2021](#)), along with a recent generalisation of QUMOND ([Milgrom 2023](#)). These topics, however, are beyond the scope of this paper.

### 2.4. Parametrised gravity

A straightforward and effective phenomenological approach to modifying the theory of gravity is through the  $\mu - \Sigma$  parametrisation of the Einstein equations ([Amendola et al. 2008](#)). This is particularly useful when considering unequal Bardeen potentials  $\phi \neq \psi$ . Under the Newtonian and quasi-static approximations, the Einstein equations can be expressed as

$$\nabla^2 \psi = 4\pi G \mu(a) a^2 \delta \rho_m, \quad (19)$$

$$\nabla^2 (\psi + \phi) = 8\pi G \Sigma(a) a^2 \delta \rho_m, \quad (20)$$

where  $\mu(a)$  represents the time-dependent ‘effective gravitational coupling’, which can be interpreted as a modification of the gravitational constant, and  $\Sigma(a)$  is the ‘light deflection parameter’. Since our focus is on the evolution of dark-matter particles, we only need to implement Eq. (19), which involves  $\mu(a)$ . In practice, the gravitational coupling  $\mu(a)$  could be a function of both time and scale in Fourier space,  $\mu(a, k)$ , as in the ‘effective-field theory of dark energy’ ([Frusciante & Perenon 2020](#)). However, for simplicity, we prefer methods that can be solved numerically in both Fourier and configuration space. The inclusion of scale-dependent corrections will be considered in future work.

For the functional form of  $\mu(a)$ , we use the parametrisation from [Simpson et al. \(2013\); Planck Collaboration XIV \(2016\); Planck Collaboration VI \(2020\); Abbott et al. \(2019\)](#), which allows for deviations from GR during a dark-energy dominated era

$$\mu(a) = 1 + \mu_0 \frac{\Omega_\Lambda(a)}{\Omega_{\Lambda,0}}, \quad (21)$$

where  $\mu_0$  is the only free parameter, representing the gravitational coupling today.

### 2.5. $f(R)$ gravity

In  $f(R)$  gravity, the Lagrangian extends the Einstein–Hilbert action (GR) by including an arbitrary function of the Ricci scalar curvature  $R$ . The total action is given as (Buchdahl 1970; Sotiriou & Faraoni 2010)

$$S = \int d^4x \sqrt{-g} \left[ \frac{R + f(R)}{16\pi G} + \mathcal{L}_m \right], \quad (22)$$

where  $\mathcal{L}_m$  represents the matter Lagrangian,  $G$  is the gravitational constant,  $g$  is the determinant of the metric, and  $f(R)$  is the additional function of the curvature, which reduces to  $-2\Lambda$  in the standard  $\Lambda$ CDM model. A commonly used parametrisation of  $f(R)$  gravity is provided by Hu & Sawicki (2007), with the following functional form

$$f(R) = -m^2 \frac{c_1 (R/m^2)^n}{c_2 (R/m^2)^n + 1}, \quad (23)$$

$$\approx -\frac{c_1}{c_2} m^2 + \frac{c_1}{c_2^2} m^2 \left( \frac{m^2}{R} \right)^n, \quad R \gg m^2, \quad (24)$$

where  $n, c_1$  and  $c_2$  are the model parameters, and  $m$  represents the curvature scale, given by  $m^2 = \Omega_m H_0^2/c^2 = 8\pi G \bar{\rho}_{m,0}/(3c^2)$ , where  $\bar{\rho}_{m,0}$  is the current mean matter density. This model incorporates a Chameleon screening mechanism (Khoury & Weltman 2004; Burrage & Sakstein 2018) to suppress the fifth force caused by the scalar field  $f_R$  (also known as the ‘scalarmon’). The scalaron is given by

$$f_R = \frac{df(R)}{dR} \approx -n \frac{c_1}{c_2^2} \left( \frac{m^2}{R} \right)^{n+1}, \quad (25)$$

which allows the theory to recover GR in high-density environments, ensuring consistency with Solar System tests. Observational evidence for dark energy as a cosmological constant imposes the constraint

$$\frac{c_1}{c_2} = 6 \frac{\Omega_{\Lambda,0}}{\Omega_{m,0}}, \quad (26)$$

and we can also express

$$\frac{c_1}{c_2^2} = -\frac{3}{n} \left[ 1 + 4 \frac{\Omega_{\Lambda,0}}{\Omega_{m,0}} \right]^{n+1} f_{R0}, \quad (27)$$

where  $f_{R0}$  is the present-day value of the scalaron, with current observational constraints from galaxy clusters indicating  $\log_{10} f_{R0} < -5.32$  (Vogt et al. 2025). In this framework, the Poisson equation is modified compared to its Newtonian counterpart. There is an additional term that depends on the scalaron field

$$\nabla^2 \phi = \frac{16\pi G a^2}{3} \delta \rho - \frac{1}{6} \delta R, \quad (28)$$

$$\nabla^2 f_R = -\frac{8\pi G a^2}{3c^2} \delta \rho + \frac{1}{3} \delta R, \quad (29)$$

where the difference in curvature is given by

$$\delta R = R - \bar{R} = \bar{R} \left[ \left( \frac{\bar{f}_R}{f_R} \right)^{1/(n+1)} - 1 \right]. \quad (30)$$

Here,  $\bar{R}$  represents the background curvature, expressed as

$$\bar{R} = 3m^2 \left( a^{-3} + 4 \frac{\Omega_\Lambda}{\Omega_m} \right), \quad (31)$$

and  $\bar{f}_R$  is the background value of  $f_R$

$$\bar{f}_R = \left( \frac{\bar{R}_0}{\bar{R}} \right)^{n+1} f_{R,0}. \quad (32)$$

The primary observational distinction between  $f(R)$  gravity and GR lies in the enhanced clustering on small scales, with the amplitude and shape of these deviations being dependent on  $f_{R0}$ , despite the two models sharing the same overall expansion history.

### 3. Methods

This section reviews the numerical methods used in PySCo, from generating initial conditions to evolving dark-matter particles in  $N$ -body simulations across various theories of gravity.

PySCo is entirely written in Python and uses the open-source library Numba (Lam et al. 2015), which compiles Python code into machine code using the LLVM compiler. This setup combines Python’s high development speed and rich ecosystem with the performance of C/Fortran. To optimise performance, PySCo relies on writing native Python code with ‘for’ loops, similar to how it would be done in C or Fortran.

Numba integrates seamlessly with NumPy (Harris et al. 2020), a widely used package for numerical operations in Python. Parallelisation in PySCo is simplified: by replacing ‘range’ with ‘prange’ in loops, the code takes advantage of multi-core processing. Numba functions are typically compiled just-in-time (JIT), meaning they are compiled the first time the function is called. Numba infers input and output types dynamically, supporting function overloading for different types.

In PySCo, however, most functions are compiled ahead-of-time (AOT), meaning they are compiled as soon as the code is executed or imported. This is because the simulation uses 32-bit floating point precision for all fields, allowing for AOT compilation. Since the simulation operates on a uniform grid, unlike AMR simulations, there is no need for fine-grained grids. Therefore, using 32-bit precision is sufficient and does not result in any loss of accuracy. Additionally, 32-bit floats improve performance by enabling SIMD (Single Instruction, Multiple Data) instructions, which the compiler implicitly optimises for.

#### 3.1. Units and conventions

We adopt the same strategy as RAMSES and use supercomoving units (Martel & Shapiro 1998), where the Poisson equation takes the same form as in classical Newtonian dynamics but includes a multiplicative scale factor

$$\nabla^2 \tilde{\phi} = \frac{3}{2} a \Omega_m (\tilde{\rho} - 1), \quad (33)$$

where a tilde denotes a quantity in supercomoving units,  $a$  is the scale factor,  $\phi$  the gravitational potential, and  $\rho$  the matter density. We also define conversion units from comoving coordinates and super-conformal time to physical SI units

$$\tilde{x} = \frac{x}{x_*}, \quad d\tilde{t} = \frac{dt}{t_*}, \quad \tilde{v} = v \frac{t_*}{x_*}, \quad (34)$$

$$\tilde{\phi} = \phi \frac{t_*^2}{x_*^2}, \quad \tilde{c} = c \frac{t_*}{ax_*}, \quad (35)$$

where  $\tilde{x}$ ,  $d\tilde{t}$ ,  $\tilde{v}$ ,  $\tilde{\phi}$  and  $\tilde{c}$  represent the particle position, time and velocity, gravitational potential and speed of light in simulation units. The conversion factors are defined as follows

$$x_* = 100aL_{\text{box}}/H_0, \quad (36)$$

$$t_* = a^2/H_0, \quad (37)$$

$$\rho_* = \Omega_m \rho_c / a^3, \quad (38)$$

as the length, time and density conversion units to  $\text{km}^2$ , seconds and  $\text{kg}/\text{m}^3$  respectively, where  $L_{\text{box}}$  is the box length in comoving coordinates and  $H_0$  is the Hubble parameter today (in seconds). The particle mass is given by

$$m_{\text{part}} = \rho_* x_*^3 / N_{\text{part}}, \quad (39)$$

where  $N_{\text{part}}$  is the total number of particles in the simulation.

### 3.2. Data structure

In this section, we discuss how PySCo handles the storage of particles and meshes using C-contiguous NumPy arrays. This approach was chosen for its simplicity and readability, allowing functions in PySCo to be easily reused in different contexts.

For particles, the position and velocity arrays are stored with the shape  $[N_{\text{part}}, 3]$ , where the elements for each particle are contiguous in memory. This format is more efficient than using a shape of  $[3, N_{\text{part}}]$ , particularly for mass assignment, where operations are performed particle by particle. To further enhance performance, particles are ordered using Morton indices (also known as z-curve indices) rather than linear or random ordering. Morton ordering improves cache usage and, thus, increases performance. This ordering is applied every  $N_{\text{reorder}}$  steps, as defined by the user, to maintain good data locality and avoid performance losses (see also Appendix E). We did not consider space-filling curves with better data locality properties (such as the Hilbert curve), because the associated encoding and decoding algorithms are much more computationally expensive, and Morton curves already provide excellent data locality. While more complex data structures are available (such as linked lists, fully-threaded trees, octrees or kdtrees), preliminary tests showed that using Morton-ordered NumPy arrays strikes a good balance between simplicity and performance, without the overhead of creating complex structures.

For scalar fields on the grid, arrays are stored with a shape  $[N_{\text{cells}}^{1/3}, N_{\text{cells}}^{1/3}, N_{\text{cells}}^{1/3}]$  using linear indexing. Although linear indexing does not offer optimal data locality, it is lightweight and does not require additional arrays to store indices. Moreover, it works well with predictable (optimizable) memory-access patterns, such as those used in stencil operators. For vector fields, such as acceleration, the arrays have a shape  $[N_{\text{cells}}^{1/3}, N_{\text{cells}}^{1/3}, N_{\text{cells}}^{1/3}, 3]$ , similar to the format used for particle arrays, to maintain consistency and performance.

### 3.3. Initial conditions

In this section, we describe how PySCo handles the generation of initial conditions for simulations, although it can also read from pre-existing snapshots in other formats. PySCo can

load data directly from RAMSES/pFoF format used in the Ray-Gal simulations (Breton et al. 2019; Rasera et al. 2022) or from Gadget format using the Pylians library (Villaescusa-Navarro 2018). PySCo computes the time evolution of the scale factor, growth factors, and Hubble parameters, with the Astropy library (Astropy Collaboration 2022) and internal routines (see Appendix A).

To generate the initial conditions, the code requires a linear power spectrum  $P(k, z = 0)$ , which is rescaled by the growth factor at the initial redshift  $z_{\text{ini}}$ . Additionally, PySCo generates a realisation of Gaussian white noise  $W$ , which is used to apply initial displacements to particles.

Some methods generate Gaussian white noise in configuration space, which is particularly useful for zoom simulations (Pen 1997; Sirko 2005; Bertschinger 2001; Prunet et al. 2008; Hahn & Abel 2013). However, in PySCo, the white noise is computed as

$$W(\mathbf{k}) = A(\mathbf{k}) e^{i\theta(\mathbf{k})}, \quad (40)$$

with  $A(\mathbf{k})$  an amplitude drawn from a Rayleigh distribution given by  $\mathcal{R}_{\text{dist}} = \sqrt{-\ln \mathcal{U}[0, 1]}$  with  $\mathcal{U}[0, 1]$  a uniform random sampling between 0 and 1, and  $\theta(\mathbf{k}) = \mathcal{U}[0, 1]$ . We also fix  $W(\mathbf{k}) = \bar{W}(-\mathbf{k})$  where a bar denotes a complex conjugate, to ensure that the configuration-space field is real valued. In our case, since we consider a regular grid with periodic boundary conditions, we can generate the white noise directly in Fourier space. An initial realisation of a density field is computed using

$$\delta_{\text{ini}}(\mathbf{k}) = (2\pi)^{3/2} \sqrt{P(k)} W(\mathbf{k}), \quad (41)$$

with  $k = |\mathbf{k}|$ . This ensures that we recover the Gaussian properties

$$\langle \delta_{\text{ini}}(\mathbf{k}) \rangle = 0, \quad (42)$$

$$\langle \delta_{\text{ini}}(\mathbf{k}) \delta_{\text{ini}}(\mathbf{k}') \rangle = (2\pi)^3 P(k) \delta_D(\mathbf{k} + \mathbf{k}'), \quad (43)$$

where  $\delta_D$  is a Dirac delta and  $\delta_{\text{ini}}(\mathbf{0}) = 0$ . We have also implemented the option to use ‘paired and fixed’ initial conditions (Angulo & Pontzen 2016). This method greatly reduces cosmic variance by running paired simulations with opposite phases, at the cost of introducing some non-Gaussian features. The concept here is that instead of averaging the product of modes to match the power spectrum, the individual modes are set directly to  $(2\pi)^{3/2} \sqrt{P(k)}$ . In practice, the density field is used to compute the initial particle displacement from a homogeneous distribution, rather than directly sampling  $\delta_{\text{ini}}$ . We use Lagrangian perturbation theory (LPT), with options for first-order 1LPT (also called ‘Zel’dovich approximation’, Zel’dovich 1970), second-order 2LPT (Scoccimarro 1998; Crocce et al. 2006), or third-order 3LPT (Catelan 1995; RAMPF & BUCHERT 2012). The displacement field at the initial redshift  $z_{\text{ini}}$  up to third order is expressed as

$$\Psi(z_{\text{ini}}) = D_+^{(1)} \Psi^{(1)} + D_+^{(2)} \Psi^{(2)} + D_+^{(3a)} \Psi^{(3a)} + D_+^{(3b)} \Psi^{(3b)} + D_+^{(3c)} \Psi^{(3c)}, \quad (44)$$

with  $D_+ \equiv D_+(z_{\text{ini}})^{(1)}$  is the linear (first order) growth factor at the initial redshift and  $\Psi^{(n)}$  are the different orders of the displacement field at  $z = 0$ , which can be written as (Michaux et al. 2021)

$$\Psi^{(1)} = -\nabla \phi^{(1)}, \quad (45)$$

<sup>2</sup> While RAMSES converts to cm.

$$\Psi^{(2)} = \nabla \phi^{(2)}, \quad (46)$$

$$\Psi^{(3a)} = \nabla \phi^{(3a)}, \quad (47)$$

$$\Psi^{(3b)} = \nabla \phi^{(3b)}, \quad (48)$$

$$\Psi^{(3c)} = \nabla \times \mathbf{A}^{(3c)}, \quad (49)$$

with

$$\phi^{(1)} = \nabla^{-2} \delta_{\text{ini}},$$

$$\phi^{(2)} = \frac{1}{2} \nabla^{-2} [\phi_{,ii}^{(1)} \phi_{,jj}^{(1)} - \phi_{,ij}^{(1)} \phi_{,ij}^{(1)}], \quad (50)$$

$$\phi^{(3a)} = \nabla^{-2} [\det \phi_{,ij}^{(1)}], \quad (51)$$

$$\phi^{(3b)} = \frac{1}{2} \nabla^{-2} [\phi_{,ii}^{(2)} \phi_{,jj}^{(1)} - \phi_{,ij}^{(2)} \phi_{,ij}^{(1)}], \quad (52)$$

$$\mathbf{A}^{(3c)} = \nabla^{-2} [\nabla \phi_{,i}^{(2)} \times \nabla \phi_{,i}^{(1)}], \quad (53)$$

and

$$\nabla^{(2)} \phi^{(2)} = \phi_{,xx}^{(1)} (\phi_{,yy}^{(1)} + \phi_{,zz}^{(1)}) + \phi_{,yy}^{(1)} \phi_{,zz}^{(1)} - \phi_{,xy}^{(1)} \phi_{,xy}^{(1)} - \phi_{,xz}^{(1)} \phi_{,xz}^{(1)} - \phi_{,yz}^{(1)} \phi_{,yz}^{(1)}, \quad (54)$$

$$\nabla^2 \phi^{(3a)} = \phi_{,xx}^{(1)} \phi_{,yy}^{(1)} \phi_{,zz}^{(1)} + 2 \phi_{,xy}^{(1)} \phi_{,xz}^{(1)} \phi_{,yz}^{(1)} - \phi_{,yz}^{(1)} \phi_{,yz}^{(1)} \phi_{,xx}^{(1)} - \phi_{,xz}^{(1)} \phi_{,xz}^{(1)} \phi_{,yy}^{(1)} - \phi_{,xy}^{(1)} \phi_{,xy}^{(1)} \phi_{,zz}^{(1)}, \quad (55)$$

$$\begin{aligned} \nabla^2 \phi^{(3b)} = & \frac{1}{2} \phi_{,xx}^{(1)} (\phi_{,yy}^{(2)} + \phi_{,zz}^{(2)}) + \frac{1}{2} \phi_{,yy}^{(1)} (\phi_{,xx}^{(2)} + \phi_{,zz}^{(2)}) \\ & + \frac{1}{2} \phi_{,zz}^{(1)} (\phi_{,xx}^{(2)} + \phi_{,yy}^{(2)}) - \phi_{,xy}^{(1)} \phi_{,xy}^{(2)} - \phi_{,xz}^{(1)} \phi_{,xz}^{(2)} - \phi_{,yz}^{(1)} \phi_{,yz}^{(2)}, \end{aligned} \quad (56)$$

$$\begin{aligned} \nabla^2 \mathbf{A}_x^{(3c)} = & \phi_{,xy}^{(2)} \phi_{,xz}^{(1)} - \phi_{,xz}^{(2)} \phi_{,xy}^{(1)} \\ & + \phi_{,yz}^{(1)} (\phi_{,yy}^{(2)} - \phi_{,zz}^{(2)}) - \phi_{,yz}^{(2)} (\phi_{,yy}^{(1)} - \phi_{,zz}^{(1)}), \end{aligned} \quad (57)$$

$$\begin{aligned} \nabla^2 \mathbf{A}_y^{(3c)} = & \phi_{,yz}^{(2)} \phi_{,yx}^{(1)} - \phi_{,yx}^{(2)} \phi_{,yz}^{(1)} \\ & + \phi_{,xz}^{(1)} (\phi_{,zz}^{(2)} - \phi_{,xx}^{(2)}) - \phi_{,xz}^{(2)} (\phi_{,zz}^{(1)} - \phi_{,xx}^{(1)}), \end{aligned} \quad (58)$$

$$\begin{aligned} \nabla^2 \mathbf{A}_z^{(3c)} = & \phi_{,xz}^{(2)} \phi_{,yz}^{(1)} - \phi_{,yz}^{(2)} \phi_{,xz}^{(1)} \\ & + \phi_{,xy}^{(1)} (\phi_{,xx}^{(2)} - \phi_{,yy}^{(2)}) - \phi_{,xy}^{(2)} (\phi_{,xx}^{(1)} - \phi_{,yy}^{(1)}). \end{aligned} \quad (59)$$

In practice, we compute all derivatives directly in Fourier space, since configuration-space finite-difference gradients would smooth the small scales and thus create inaccuracies in the initial power spectrum. The second- and third-order contributions in the initial conditions can be prone to aliasing effects due to the quadratic and cubic non-linearities involved. To mitigate this, we apply Orszag's 3/2 rule (Orszag 1971), as suggested in Michaux et al. (2021). The impact of this correction is minimal, by around 0.1% on the power spectrum at small scales, when using 3LPT initial conditions and for a relatively late start with  $z_{\text{ini}} \approx 10$ . The initial position and velocity are then (up to third order)

$$\mathbf{x}_{\text{ini}} = \mathbf{q} + \Psi^{(1)} D_{+}^{(1)} + \Psi^{(2)} D_{+}^{(2)} \quad (61)$$

$$+ \Psi^{(3a)} D_{+}^{(3a)} + \Psi^{(3b)} D_{+}^{(3b)} + \Psi^{(3c)} D_{+}^{(3c)}, \quad (62)$$

$$\mathbf{v}_{\text{ini}} = \Psi^{(1)} H f^{(1)} D_{+}^{(1)} + \Psi^{(2)} H f^{(2)} D_{+}^{(2)} \quad (63)$$

$$+ \Psi^{(3a)} H f^{(3a)} D_{+}^{(3a)} + \Psi^{(3b)} H f^{(3b)} D_{+}^{(3b)} + \Psi^{(3c)} H f^{(3c)} D_{+}^{(3c)}, \quad (64)$$

with  $H$  the Hubble parameter at  $z_{\text{ini}}$ ,  $f_n$  the growth rate contribution to the  $n$ -th order, and  $\mathbf{x}_{\text{uniform}}$  the position of cell centres (or cell edges, see also Appendix B.2) when  $N_{\text{part}} = N_{\text{cells}}$ . We internally compute the growth factor and growth rate contributions as described in Appendix A.

### 3.4. Integrator

In our simulations, we employ the second-order symplectic Leapfrog scheme, often referred to as Kick-Drift-Kick, to integrate the equations of motion for the particles. The steps in the scheme are as follows

$$\begin{aligned} \mathbf{v}_{i+1/2} &= \mathbf{v}_i + \mathbf{a}_i \Delta t / 2, & \text{Kick}, \\ \mathbf{x}_{i+1} &= \mathbf{x}_i + \mathbf{v}_{i+1/2} \Delta t, & \text{Drift}, \\ \mathbf{v}_{i+1} &= \mathbf{v}_{i+1/2} + \mathbf{a}_{i+1} \Delta t / 2, & \text{Kick}, \end{aligned} \quad (65)$$

where the subscript  $i$  indicates the integration step, while  $\mathbf{x}$ ,  $\mathbf{v}$  and  $\mathbf{a}$  are the particle positions, velocities and accelerations respectively. There are several ways to set the time step  $\Delta t$ . Some authors use linear or logarithmic spacing, depending on the user input. In our case, we followed a similar strategy as RAMSES (Teyssier 2002) and use several time stepping criteria. The first criterion is based on a cosmological time step that guarantees the scale factor does not change by more than a specified amount (by default 2%, see also Appendix B.1), which is particularly effective at high redshift. The second criterion is based on the minimum free-fall time given by

$$\Delta t_{\text{ff}} = \sqrt{\frac{h}{\max(|\mathbf{a}_i|)}}, \quad (66)$$

with  $h$  the cell size. We select the smallest value between this criterion and the cosmological time step. Additionally, we implemented a third criterion based on particle velocities  $\Delta t_{\text{vel}} = h / \max(|\mathbf{v}_i|)$ , though we found this value often exceeds  $\Delta t_{\text{ff}}$  in practice. This approach ensures that the time step dynamically adapts to the structuration of dark matter in the simulation. To further refine the time step, we multiply it by a user-defined Courant-like factor.

An interesting prospect for the future is the use of integrators based on Lagrangian perturbation theory (LPT) that could potentially reduce the number of time steps required while maintaining high accuracy (as suggested by Rampf et al. 2025). However, such methods couple the integration scheme with specific theories of gravity through growth factors. Since these factors may not always be accurately computed (for example, in MOND), we prefer to maintain the generality of the standard leapfrog integration scheme for now. We may explore the implementation of such LPT-based integrators in the future for specific theories of gravitation.

### 3.5. Iterative solvers

To displace the particles, we first need to compute the force (or acceleration). There are various algorithms available for this purpose, either computing the force directly from a particle distribution, or determining the gravitational potential from which the force can subsequently be derived. When using the gravitational potential approach, the force can be recovered by applying a finite-difference gradient operator  $\mathbf{g} = -\nabla \phi$ . Given that our results are sensitive to the order of the operator used, we have implemented several options for central difference methods, each characterised by specific coefficients. These coefficients are detailed in Table 1. We aim to solve the following problem

$$\mathcal{L}u = f \quad (67)$$

Where  $u$  is unknown,  $f$  is known and  $\mathcal{L}$  is an operator. For the classical Poisson equation,  $u \equiv \phi$ ,  $f \equiv 4\pi G\rho$  and  $\mathcal{L} \equiv \nabla^2$  with the seven-point Laplacian stencil

$$\nabla^2 u_{i,j,k} = \frac{1}{h^2} (L_{i,j,k}(u) - 6u_{i,j,k}), \quad (68)$$

**Table 1.** Stencil operator coefficients for central derivatives.

Points	Accuracy	-3	-2	-1	0	1	2	3
3	$O(h^2)$	0	0	$-\frac{1}{2}$	0	$\frac{1}{2}$	0	0
5	$O(h^4)$	0	$\frac{1}{12}$	$-\frac{2}{3}$	0	$\frac{2}{3}$	$-\frac{1}{12}$	0
7	$O(h^6)$	$-\frac{1}{60}$	$\frac{3}{20}$	$-\frac{3}{4}$	0	$\frac{3}{4}$	$-\frac{3}{20}$	$\frac{1}{60}$

**Notes.** The columns refer to the stencil operator, accuracy and coefficients  $c_n$  of the indices  $n$  such as the gradient operator at the index  $i$  can be written as  $\nabla = \frac{1}{h} \sum_{n=-3}^{n=3} c_n u_{i+n}$ , with  $h$  the grid size.

with the subscripts  $i, j, k$  the cell indices and  $L_{i,j,k}(u) = u_{i+1,j,k} + u_{i,j+1,k} + u_{i,j,k+1} + u_{i-1,j,k} + u_{i,j-1,k} + u_{i,j,k-1}$ .

Lastly,  $f$  is the density (source) term of Eq. (67), which is directly estimated from the position of dark-matter particles. In code units, the sum of the density over the full grid must be

$$\sum_{i,j,k} \tilde{\rho}_{i,j,k} = N_{\text{part}}, \quad (69)$$

where the density in a given cell is computed using the ‘nearest-grid point’ (NGP), ‘cloud-in-cell’ (CIC) or ‘triangular-shaped cloud’ (TSC) mass-assignment schemes

$$W_{\text{NGP}}(x_i) = \begin{cases} 1 & \text{if } |x_i| < 0.5, \\ 0 & \text{otherwise,} \end{cases} \quad (70)$$

$$W_{\text{CIC}}(x_i) = \begin{cases} 1 - |x_i| & \text{if } |x_i| < 1.0, \\ 0 & \text{otherwise,} \end{cases} \quad (71)$$

$$W_{\text{TSC}}(x_i) = \begin{cases} 0.75 - x_i^2 & \text{if } |x_i| < 0.5, \\ (1.5 - |x_i|)^2 / 2 & \text{else if } 0.5 < |x_i| < 1.5, \\ 0 & \text{otherwise.} \end{cases} \quad (72)$$

Here,  $x_i$  is the normalised separation between a particle and cell positions, scaled by the cell size. In a three-dimensional space, this implies that a dark-matter particle contributes to the density of one, eight, or twenty-seven cells depending on whether the NGP, CIC, or TSC scheme is employed, respectively. For consistency, we use the same (inverse) scheme to interpolate the acceleration from cells to particles.

### 3.6. The Jacobi and Gauss-Seidel methods

Let us consider  $\mathcal{L} = \begin{bmatrix} \ell_{11} & \ell_{12} & \ell_{13} & \dots & \ell_{1n} \\ \ell_{21} & \ell_{22} & \ell_{23} & \dots & \ell_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \ell_{n1} & \ell_{n2} & \ell_{n3} & \dots & \ell_{nn} \end{bmatrix}$ ,  $u = \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_n \end{bmatrix}$  and  $f = \begin{bmatrix} f_1 \\ f_2 \\ \vdots \\ f_n \end{bmatrix}$ . The Jacobi method is a naive iterative solver which, for Eq. (67), takes the form

$$\ell_{11}u_1^{\text{new}} + \ell_{12}u_2^{\text{old}} + \ell_{13}u_3^{\text{old}} + \dots + \ell_{1n}u_n^{\text{old}} = f_1, \quad (73)$$

$$\ell_{21}u_1^{\text{old}} + \ell_{22}u_2^{\text{new}} + \ell_{23}u_3^{\text{old}} + \dots + \ell_{2n}u_n^{\text{old}} = f_2, \quad (74)$$

$$\ell_{31}u_1^{\text{old}} + \ell_{32}u_2^{\text{old}} + \ell_{33}u_3^{\text{new}} + \dots + \ell_{3n}u_n^{\text{old}} = f_3, \quad (75)$$

⋮

$$\ell_{n1}u_1^{\text{old}} + \ell_{n2}u_2^{\text{old}} + \ell_{n3}u_3^{\text{old}} + \dots + \ell_{nn}u_n^{\text{new}} = f_n, \quad (76)$$

where the superscripts ‘new’ and ‘old’ refer to the iteration. This means that for one Jacobi sweep, the new value is directly given by that of the previous iteration.

In practice, we use the Gauss-Seidel method instead, which has a better convergence rate and lower memory usage. This method is akin to Jacobi’s but enhances convergence by incorporating the most recently updated values in the computation of the remainder. It follows

$$\ell_{11}u_1^{\text{new}} + \ell_{12}u_2^{\text{old}} + \ell_{13}u_3^{\text{old}} + \dots + \ell_{1n}u_n^{\text{old}} = f_1, \quad (77)$$

$$\ell_{21}u_1^{\text{new}} + \ell_{22}u_2^{\text{new}} + \ell_{23}u_3^{\text{old}} + \dots + \ell_{2n}u_n^{\text{old}} = f_2, \quad (78)$$

$$\ell_{31}u_1^{\text{new}} + \ell_{32}u_2^{\text{new}} + \ell_{33}u_3^{\text{new}} + \dots + \ell_{3n}u_n^{\text{old}} = f_3, \quad (79)$$

⋮

$$\ell_{n1}u_1^{\text{new}} + \ell_{n2}u_2^{\text{new}} + \ell_{n3}u_3^{\text{new}} + \dots + \ell_{nn}u_n^{\text{new}} = f_n, \quad (80)$$

which seems impossible to parallelise because each subsequent equation relies on the results of the previous one. However, we implement a strategy known as ‘red-black’ ordering in the Gauss-Seidel method. This technique involves colouring the cells like a chessboard, where each cell is designated as either red or black. When updating the red cells, we only utilise information from the black cells, and vice versa. This approach is equivalent to performing two successive Jacobi sweeps—one for the red cells and another for the black cells. For the Laplacian operator as expressed in Eq. (68), the Jacobi sweep for a given cell can be formulated as follows

$$u_{i,j,k}^{\text{new}} = \frac{1}{6} \left( L_{i,j,k}(u^{\text{old}}) - h^2 f_{i,j,k} \right). \quad (81)$$

In scenarios involving a non-linear operator, finding an exact solution may not be feasible. In such cases, we linearise the operator using the Newton-Raphson method, expressed as

$$u^{\text{new}} = u^{\text{old}} - \frac{\mathcal{L}(u^{\text{old}})}{\partial \mathcal{L} / \partial u^{\text{old}}}, \quad (82)$$

This approach is commonly referred to as the ‘Newton Gauss-Seidel method’, allowing for more effective convergence in non-linear contexts.

### 3.7. Multigrid

In practice, both the Jacobi and Gauss-Seidel methods are known for their slow convergence rates, often requiring hundreds to thousands of iterations to reach a solution and typically unable to achieve high accuracy. To overcome these limitations, a popular and efficient iterative method known as ‘multigrid’ is employed. This algorithm significantly accelerates convergence by solving the equation iteratively on coarser meshes, effectively addressing large-scale modes. The multigrid algorithm follows this procedure (Press et al. 1992): we first discretise the problem on a regular mesh with grid size  $h$  as

$$\mathcal{L}_h u_h = f_h, \quad (83)$$

which can be solved using the Gauss-Seidel method, and where  $\mathcal{L}_h$  is the numerical operator on the mesh with grid size  $h$  which

approximates  $\mathcal{L}$ .  $\tilde{u}_h$  denotes our approximate solution and  $v_h$  is the ‘error’ on the true solution

$$v_h = u_h - \tilde{u}_h, \quad (84)$$

and the ‘residual’ is given by

$$d_h = f_h - \mathcal{L}\tilde{u}_h. \quad (85)$$

Depending on whether the operator  $\mathcal{L}$  is linear or non-linear, different multigrid schemes will be employed to solve the equations effectively.

### 3.7.1. Linear multigrid

Considering the case where  $\mathcal{L}$  is linear, meaning that  $\mathcal{L}_h(u_h - \tilde{u}_h) = \mathcal{L}_h u_h - \mathcal{L}_h \tilde{u}_h$ , we have the relation

$$\mathcal{L}_h v_h = d_h. \quad (86)$$

From there, the goal is to estimate  $v_h$  to find  $u_h$ . We use numerical methods to find the approximate solution  $\tilde{v}_h$  by solving

$$\mathcal{L}_h \tilde{v}_h = d_h, \quad (87)$$

using Gauss-Seidel. The updated approximation for the field is

$$\tilde{u}_h^{\text{new}} = \tilde{u}_h + \tilde{v}_h. \quad (88)$$

The issue is that the approximate operator  $\mathcal{L}_h$  is usually local and finite-difference based, for which long-range perturbations are slowly propagating and therefore very inefficient computationally. The multigrid solution to this issue is to solve the error on coarser meshes to speed up the propagation of long-range modes. First, we use the ‘restriction’ operator  $\mathcal{R}$  which interpolates from fine to coarse grid

$$d_H = \mathcal{R}d_h, \quad (89)$$

where  $H = 2h$  is the grid size of the coarse mesh. We then solve Eq. (87) on the coarser grid to infer  $\tilde{v}_H$  as

$$\mathcal{L}_H \tilde{v}_H = d_H. \quad (90)$$

We then use the prolongation operator  $\mathcal{P}$ , which interpolates from coarse to fine grid

$$\tilde{v}_h = \mathcal{P}\tilde{v}_H, \quad (91)$$

and we finally update our approximation on the solution

$$\tilde{u}_h^{\text{new}} = \tilde{u}_h + \tilde{v}_h. \quad (92)$$

We have provided a brief overview of the multigrid algorithm using two grids, but in practice, to solve for  $\tilde{v}_H$ , we can extend the scheme to even coarser meshes. This results in a recursive algorithm where the coarser level in PySCo contains  $16^3$  cells. There are several strategies for navigating the different mesh levels, commonly referred to as V, F, or W-cycles. The V cycle is the simplest and quickest to execute, although it converges at a slower rate than the F and W cycles (see Appendix C).

For the restriction and prolongation operators, the lowest-level schemes employ averaging (where the field value on the parent cell is the mean of its children’s values) and straight injection (where child cells inherit the value of their parent), respectively. However, as noted by Guillet & Teyssier (2011), to

minimise inaccuracies in the estimation of the final solution, we use a higher-order prolongation scheme defined as

$$\mathcal{P} = \begin{cases} 27/64, & x_{0,1,2} < 0.5, \\ 9/64, & x_{0,1} < 0.5 < x_2 < 1.5, \\ 3/64, & x_0 < 0.5 < x_{1,2} < 1.5, \\ 1/64, & 0.5 < x_{0,1,2} < 1.5, \\ 0, & \text{otherwise,} \end{cases} \quad (93)$$

where  $\mathbf{x} = (x_0, x_1, x_2)$  with  $x_0 \leq x_1 \leq x_2$ , is the separation (normalised by the fine grid size) between the centre of the fine and coarse cells.

We consider our multigrid scheme to have converged when the residual is significantly lower than the ‘truncation error’, defined as:

$$\tau \equiv \mathcal{L}_h u - f_h, \quad (94)$$

and which can be estimated as (Press et al. 1992; Li et al. 2012)

$$\tilde{\tau}_h \equiv \mathcal{L}_H(\mathcal{R}\tilde{u}_h) - \mathcal{R}\mathcal{L}_h\tilde{u}_h. \quad (95)$$

We consider that we reached convergence when

$$|d_h| < \alpha \tilde{\tau}_h, \quad (96)$$

where  $|d_h|$  is the square root of the quadratic sum over the residual in each cell of the mesh,  $\alpha$  is the stopping criterion. It is noteworthy that (Knebe et al. 2001) proposed using an alternative approach for estimating the truncation error

$$\tau_{h,K01} = \mathcal{P}(\mathcal{L}_H(\mathcal{R}\tilde{u}_h)) - \mathcal{L}_h\tilde{u}_h, \quad (97)$$

which we can approximate by

$$\tau_{h,K01} \approx \mathcal{P}(\mathcal{R}f_h) - f_h, \quad (98)$$

to reduce computational time, when the source term is non-zero. The relation between these two approaches is  $\tau_h \approx 0.1\tau_{h,K03}$ . In practice, we use the first estimation  $\tau_h$ . Since the Jacobi and Gauss-Seidel methods are iterative, an initial guess is still required. While the multigrid algorithm is generally insensitive to the initial guess in most cases, some small optimisations can be made. For instance, if we initialise the full grid with zeros, one Jacobi step directly provides

$$u_{i,j,k}^{\text{ini}} = -\frac{h^2}{6} f_{i,j,k}, \quad (99)$$

which is used to initialise  $u_h$  for the first simulation step, and  $v_H$  the error on coarser meshes.

In addition to the Gauss-Seidel method, RAMSES also incorporates ‘successive over-relaxation’ (SOR), which allows the updated field to include a contribution from the previous iteration, expressed as  $\phi^{n+1} = \omega\phi^{n+1} + (1 - \omega)\phi^n$ , where  $n$  denotes the iteration step, and  $\omega$  is the relaxation factor. Typically, we perform two Gauss-Seidel sweeps before the restriction (termed pre-smoothing) and one sweep after the prolongation (post-smoothing), controlled by the parameters  $\text{Npre}$  and  $\text{Npost}$ . We set  $\omega = 1.25$  by default (see also Appendix C), similarly to Kravtsov et al. (1997).

A pseudo code for the V cycle is depicted in Fig. 1, where the components that require modification based on the gravitational theory are the smoothing and residual functions. As noted by Press et al. (1992), omitting SOR (setting  $\omega = 1$ ) for the first and last iterations can enhance performance since this

```

1 def V_cycle(u_h, f_h, nlevel = 0):
2
3     h = 0.5** (ncoarse - nlevel) # Cell size
4
5     smoothing(u_h, f_h, h, Npre) # Gauss-Seidel
6     d_h = residual(u_h, f_h, h)
7     d_H = restriction(d_h)
8
9     v_H = np.zeros_like(d_H)
10    if nlevel >= levelmax:
11        smoothing(v_H, d_H, 2 * h, Npre)
12    else:
13        V_cycle(v_H, d_H, nlevel + 1)
14
15    u_h += prolongation(v_H)
16    smoothing(u_h, f_h, h, Npost)

```

**Fig. 1.** Pseudo code in Python for the multigrid V-cycle algorithm. We highlight in cyan the lines where the theory of gravitation enters.

allows the prolongation operator (line 15) to operate on only half the mesh, since the other half will be directly updated during the first Gauss-Seidel sweep. Similarly, the restriction operator can also be applied to only half of the mesh because, following one Gauss-Seidel sweep, the residual will be zero on the other half by design. It is thus feasible to combine lines 6 and 7 (the residual and restriction operators) to compute the coarser residual from half of the finer mesh. In this example, we use a zero initial guess for the error at the coarser level, but in principle, we could initialise it with a Jacobi step (as shown in Eq. (99) for Newtonian gravity).

This optimisation, however, should be approached cautiously, as the residual is zero up to floating-point precision. In certain scenarios, this can significantly impact the final solution. For example, in PySCo, we primarily utilise 32-bit floats, rendering this optimisation less accurate for the largest modes. Consequently, we ultimately decided to retain only the optimisation concerning the prolongation on half the mesh when employing the non-linear multigrid algorithm, as we do not utilise SOR in this context (see Appendix C).

### 3.7.2. Non-linear multigrid

If  $\mathcal{L}$  is a non-linear operator instead, then we need to solve for

$$\mathcal{L}_h(\tilde{u}_h + v_h) - \mathcal{L}_h\tilde{u}_h = d_h. \quad (100)$$

Contrarily to the linear case in Eq. (90) where we only need to solve for the error at the coarser level, we now use

$$\mathcal{L}_H\tilde{u}_H = \mathcal{L}_H(\mathcal{R}\tilde{u}_h) + \mathcal{R}d_h, \quad (101)$$

where we need to store the full approximation of the solution  $\tilde{u}$  at every level, hence the name ‘Full Approximation Storage’. Finally, we update the solution as

$$\tilde{u}_h^{\text{new}} = \tilde{u}_h + \mathcal{P}(\tilde{u}_H - \mathcal{R}\tilde{u}_h). \quad (102)$$

### 3.8. Fast Fourier transforms

We implemented three different fast Fourier transforms (FFT) procedures to compute the force field, most of which were already implemented in FastPM (Feng et al. 2016).

– FFT: The Laplacian operator is computed through the Green’s function kernel

$$\nabla^{-2} = -k^{-2}W_{\text{MAS}}^{-2}(k), \quad (103)$$

where  $W_{\text{MAS}}(k)$  is the mass-assignment scheme filter, given by

$$W_{\text{MAS}}(k) = \left[ \prod_{d=x,y,z} \text{sinc}\left(\frac{\omega_d}{2}\right) \right]^p, \quad (104)$$

with  $\omega_d = k_d h$  between  $[-\pi, \pi]$  and where  $p = 1, 2, 3$  for NGP, CIC and TSC respectively (Hockney & Eastwood 1981; Jing 2005).

– FFT\_7pt: instead of using the exact kernel for the Laplacian operator, we use the Fourier-space equivalent of the seven-point stencil Laplacian (Eq. (68)), which reads

$$\nabla^{-2} = - \left[ \sum_{d=x,y,z} \left( h\omega_d \text{sinc} \frac{\omega_d}{2} \right)^2 \right]^{-1}, \quad (105)$$

where no mass-assignment compensation is used. For both FFT and FFT\_7pt methods, the force is estimated through the finite-difference stencil as shown in Table 1.

– FULL\_FFT: The force is directly estimated in Fourier space through the differentiation kernel

$$\nabla\nabla^{-2} = -ikk^{-2}W_{\text{MAS}}^{-2}(k). \quad (106)$$

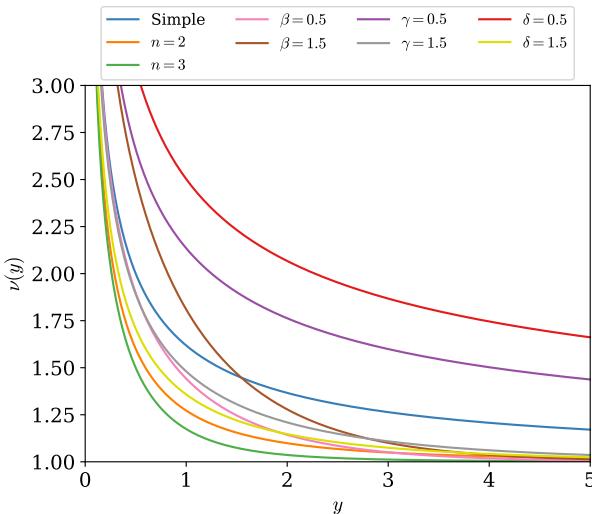
As we will see in Section 4.2, this naive operator can become very inaccurate when the simulation has more cells than particles.

### 3.9. Newtonian and parametrised simulations

In Eq. (99), we showed how to provide a generic initial guess for the multigrid algorithm. However, leveraging our understanding of the underlying physics can enable us to formulate an even more accurate first guess, thus reducing the number of multigrid cycles required for convergence. In the context of  $N$ -body simulations, we anticipate that the potential field will closely resemble that of the preceding step, especially when using sufficiently small time steps. This allows us to adopt the potential from the previous step as our initial guess for the multigrid algorithm. While this approach facilitates faster convergence to the true solution, it necessitates storing one additional grid in memory. Moreover, it is important to note that, in the linear regime in Newtonian cosmology, the density contrast evolves as a function of redshift with a scale-independent growth factor. Consequently, we can optimise our first guess by rescaling the potential field from the previous step according to the following equation

$$\tilde{\phi}(z_1) = \frac{(1+z_1)D_+(z_1)}{(1+z_0)D_+(z_0)}\tilde{\phi}(z_0), \quad (107)$$

where  $z_0$  and  $z_1$  denote the initial and subsequent redshifts, respectively. This rescaling is performed for every time step of the simulation, except for the initial time step, ensuring that we maintain an efficient and accurate estimate for our initial guess in the multigrid process. More details can be found in Appendix C.



**Fig. 2.** MOND interpolating functions for the different families and parameters shown in Sect. 2.3.

### 3.10. MOND simulations

The classical formulations of MOND (such as AQUAL and QUMOND) have already been implemented in several codes (Nusser 2002; Llinares et al. 2008; Angus et al. 2012; Candlish et al. 2015; Lüghausen et al. 2015; Visser et al. 2024). In PySCo, we further allow, if specified by the user, for a time-dependent acceleration scale  $g_0 \rightarrow a^N g_0$ , which basically delays (or accelerates) the entry of perturbations in the MOND regime, and which is set to  $N = 0$  by default. In Fig. 2, we show the behaviour of the interpolating functions  $v(y)$  described in Sect. 2.3. We see that, as expected from Eq. (13),  $v(y) \rightarrow 1$  when  $y \gg 1$ , meaning that we recover Newtonian gravity in a regime of strong acceleration.

To implement the right-hand side of Eq. (13), we employ a method analogous to that described by Lüghausen et al. (2015), and using the same notation

$$\begin{aligned} \nabla^2 \phi^N = & \frac{1}{h} \left[ \nu_{B_x} \nabla (\phi^N)_{B_x, x} - \nu_{A_x} \nabla (\phi^N)_{A_x, x} \right. \\ & + \nu_{B_y} \nabla (\phi^N)_{B_y, y} - \nu_{A_y} \nabla (\phi^N)_{A_y, y} \\ & \left. + \nu_{B_z} \nabla (\phi^N)_{B_z, z} - \nu_{A_z} \nabla (\phi^N)_{A_z, z} \right], \end{aligned} \quad (108)$$

where  $B_i$  and  $A_i$  are the points at  $+0.5h \mathbf{e}_i$  and  $-0.5h \mathbf{e}_i$  respectively, with  $\mathbf{e}_i \in \{\mathbf{e}_x, \mathbf{e}_y, \mathbf{e}_z\}$ , the unit vectors of the simulation. We have also defined

$$\nu_{B_x} = \nu \left( \frac{\sqrt{[\nabla (\phi^N)_{B_x, x}]^2 + [\nabla (\phi^N)_{B_x, y}]^2 + [\nabla (\phi^N)_{B_x, z}]^2}}{g_0} \right), \quad (109)$$

with  $\nabla (\phi^N)_{B_x, i}$  the  $i$ -th component of the force (with a minus sign) at the position  $B_x$ . The force components are estimated as

$$\nabla (\phi)_{B_x, x} = \frac{\phi_{1,0,0}^N - \phi_{0,0,0}^N}{h}, \quad (110)$$

$$\nabla (\phi)_{B_x, y} = \frac{(\phi_{1,1,0}^N - \phi_{1,-1,0}^N) + (\phi_{0,1,0}^N - \phi_{0,-1,0}^N)}{4h} \quad (111)$$

$$\nabla (\phi)_{B_x, z} = \frac{(\phi_{1,0,1}^N - \phi_{1,0,-1}^N) + (\phi_{0,0,1}^N - \phi_{0,0,-1}^N)}{4h}, \quad (112)$$

and similarly for other points. We note that for  $B_x, x$  (as well as for  $B_y, y$  and  $B_z, z$ ), we perform a three-point central derivative with half the mesh size, differing from the approach taken by (Lüghausen et al. 2015), who implemented a non-uniform five-point stencil. This decision was made to ensure that in the Newtonian case (that is,  $\nu = 1$ ), we exactly recover the seven-point Laplacian, maintaining consistency with the Laplacian operator employed in our multigrid scheme. Consequently, we opted to retain three-point derivatives for the other components as well (although we can still use a different stencil order when computing the acceleration from the MOND potential).

Additionally, we exclusively solve Eq. (108) using either the multigrid or FFT\_7pt solvers. Using the FFT solver presents challenges because, although we deconvolve  $\phi^N$  by the mass-assignment scheme kernel, the uncorrected mesh discreteness in the force computation introduces inaccuracies that can significantly affect the matter power spectrum.

### 3.11. $f(R)$ simulations

In supercomoving units, the  $f(R)$  field equations from Eqs. (28)–(29) are given by (Li et al. 2012)

$$\nabla^2 \tilde{\phi} = \frac{3}{2} \Omega_m a (\tilde{\rho} - 1) - \frac{\tilde{c}^2}{2} \nabla^2 \tilde{f}_R, \quad (113)$$

$$\nabla^2 \tilde{f}_R = -\frac{1}{\tilde{c}^2} \Omega_m a (\tilde{\rho} - 1) + \frac{1}{3\tilde{c}^2} \tilde{R} a^4 \left[ \left( \frac{\tilde{f}_R}{\tilde{f}_R} \right)^{1/(n+1)} - 1 \right], \quad (114)$$

where  $\tilde{f}_R = a^2 f_R \tilde{c}^2 / c^2$ . While Eq. (113) is linear and can be solved by standard techniques, Eq. (114) is not and needs some special attention. To this end, Oyaizu (2008) used a nonlinear multigrid algorithm (see Section 3.7.2) with the Newton-Raphson method (as shown in Eq. (82)), making the change of variable  $u \equiv \ln(\tilde{f}_R / \tilde{f}_R)$  to avoid unphysical zero-crossing of  $f_R$ . Bose et al. (2017) proposed instead that for this specific model, one could perform a more appropriate change of variable, that we generalise here to  $u \equiv (\tilde{f}_R / \tilde{f}_R)^{1/(n+1)}$ . We can then recast Eq. (114) as

$$u^{n+1} + pu + q = 0, \quad (115)$$

where

$$p = \frac{h^2}{6\tilde{c}^2 \tilde{f}_R} \left[ \Omega_m a (1 - \tilde{\rho}) - \frac{a^4 \tilde{R}}{3} \right] - \frac{1}{6} L_{i,j,k} (u^{n+1}), \quad (116)$$

$$q = \frac{a^4 h^2 \tilde{R}}{18\tilde{c}^2 \tilde{f}_R}. \quad (117)$$

We note that  $q$  is necessarily negative (because  $\tilde{f}_R < 0$ ), which is useful to determine the branch of the solution for Eq. (115).

– Case  $n = 1$ : As noticed in Bose et al. (2017), when making the change of variable  $u = \sqrt{-f_R}$ , the field equation could be recast as a depressed cubic equation,

$$u^3 + pu + q = 0, \quad (118)$$

which possesses the analytical solutions (Ruan et al. 2022)

$$u = \begin{cases} (-q)^{1/3}, & p = 0, \\ -[C + \Delta_0/C]/3, & p > 0, \\ -[C + \Delta_0/C]/3, & p < 0 \text{ and } \Delta_1^2 - 4\Delta_0^3 > 0, \\ -2\sqrt{\Delta_0} \cos(\theta/3 + 2\pi/3)/3, & \text{else,} \end{cases} \quad (119)$$

with  $\Delta_0 = -3p$ ,  $\Delta_1 = 27q$ ,  $C = \left[ \frac{1}{2} \left( \Delta_1 + \sqrt{\Delta_1^2 - 4\Delta_0^3} \right) \right]^{1/3}$  and  $\cos \theta = \Delta_1 / (2\Delta_0^{3/2})$ .

– Case  $n = 2$ : The field equation can be rewritten as a quartic equation (Ruan et al. 2022)

$$u^4 + pu + q = 0, \quad (120)$$

with the roots

$$u = \begin{cases} -S + \frac{1}{2} \sqrt{-4S^2 + p/S}, & p > 0, \\ (-q)^{1/4}, & p = 0, \\ S + \frac{1}{2} \sqrt{-4S^2 - p/S}, & p < 0, \end{cases} \quad (121)$$

where  $S = \frac{1}{2} \sqrt{\frac{1}{3} (Q + \Delta_0/Q)}$ ,  $Q = \left( \frac{1}{2} \left[ \Delta_1 + \sqrt{\Delta_1^2 - 4\Delta_0^3} \right] \right)^{1/3}$ ,  $\Delta_1 = 27p^2$  and  $\Delta_0 = 12q$ . Due to non-zero residuals in our multigrid scheme, the  $q$  term can become positive. This situation results in inequalities such as  $\Delta_1^2 - 4\Delta_0^3 < 0$ , which lack an analytical solution, or  $Q + \Delta_0/Q < 0$ . In both cases, we enforce  $u = (-q)^{1/4}$ .

While the Gauss-Seidel smoothing procedure remains necessary (as  $p$  depends on the values of the field  $u$  in adjacent cells), this method eliminates the requirement for the Newton-Raphson step and the computationally expensive exponential and logarithmic operations employed in the Oyaizu (2008) method, resulting in significant performance enhancements. Given that the operations needed to determine the branch of solutions for cubic and quartic equations are highly sensitive to machine precision, we conduct all calculations using 64-bit floating-point precision.

We must use the non-linear multigrid algorithm outlined above to solve the scalaron field, with its initial guess provided directly by the solution from the previous step, without any rescaling since we are solving for  $u$  rather than  $f_R$ . Because the tolerance threshold is heavily dependent on redshift, we cannot apply the same criterion used for the linear Poisson equation; by default, we consider convergence to be achieved after one F cycle. In fact, we do not solve Eq. (113) directly; instead, we incorporate the  $f(R)$  contribution during force computation as follows

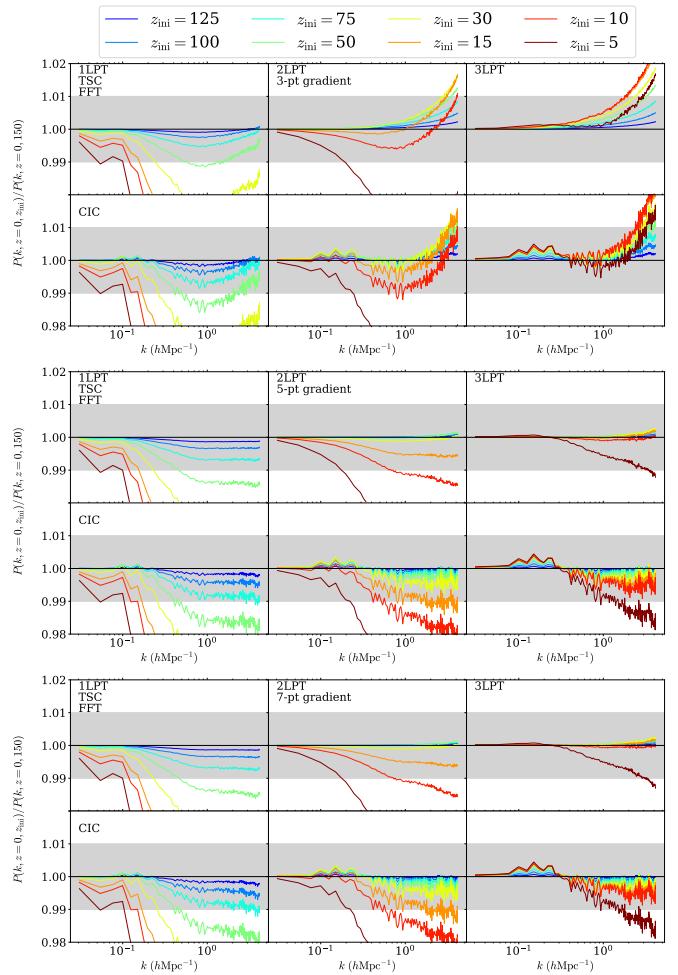
$$\mathbf{F} = \mathbf{F}_{\text{Newton}} + \frac{\tilde{c}^2}{2} \nabla \tilde{f}_R. \quad (122)$$

This choice was made because replacing  $\nabla^2 \tilde{f}_R$  in Eq. (113) with Eq. (114) could lead to a right-hand side that has a non-zero mean due to numerical inaccuracies, resulting in artificially large residuals that cannot be reduced below the error threshold.

We direct interested readers to Winther et al. (2015); Euclid Collaboration: Adamek et al. (2025), along with references therein, for comparisons of numerical methods used to solve the modified Poisson equation in Hu & Sawicki (2007)  $f(R)$  gravity.

## 4. Results

We consider a  $\Lambda$ CDM linear power spectrum computed by CAMB (Lewis et al. 2000) with parameters  $h = 0.7$ ,  $\Omega_m = 0.3$ ,  $\Omega_b = 0.05$ ,  $\Omega_r = 8.5 \cdot 10^{-5}$ ,  $n_s = 0.96$  and  $\sigma_8 = 0.8$ . We run simulations with  $512^3$  particles and as many cells (unless specifically stated) within a box of  $256 h^{-1}$  Mpc. It leads to a spatial resolution of  $0.5 h^{-1}$  Mpc (which is then reduced to  $0.25 h^{-1}$  Mpc and  $0.125 h^{-1}$  Mpc when using  $1024^3$  and  $2048^3$  cells, respectively). All the power spectrum results are shown for the snapshot at  $z = 0$ , and computed with a simple (without dealiasing) estimator implemented within PySCo. Point-mass tests are also shown in Appendix D.



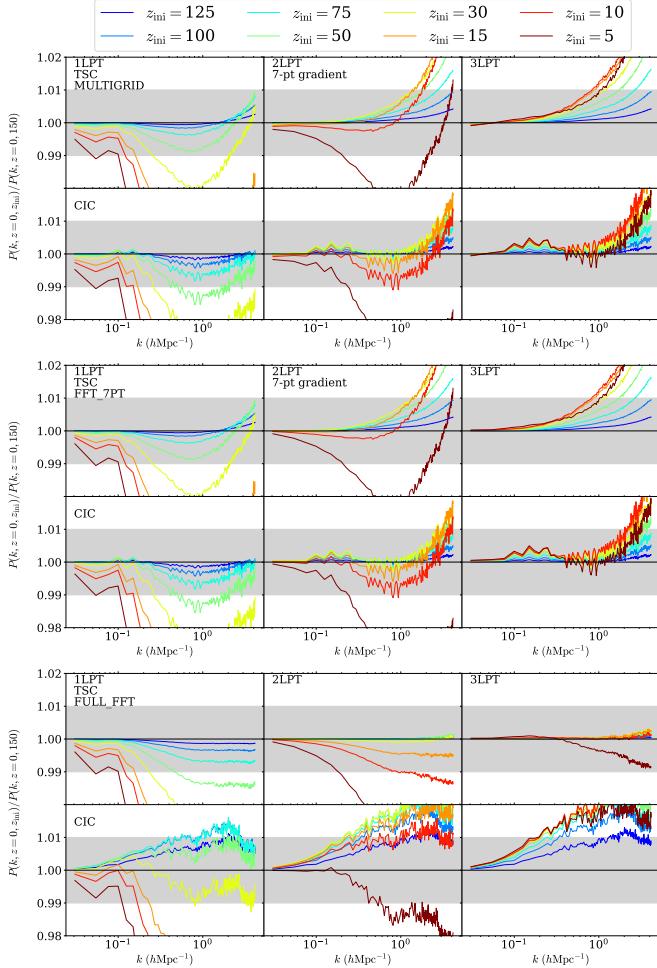
**Fig. 3.** Ratio of the matter power spectrum at  $z = 0$  for different starting redshift  $z_{\text{ini}}$  (in coloured lines), with respect to a reference simulation with  $z_{\text{ini}} = 150$ . We use 1LPT, 2LPT or 3LPT initial conditions from left to right and three-, five- and seven-point gradients from top to bottom panels. We use the FFT solver in any case (see Section 3.8). In each subplot the top and bottom panels show simulations which use TSC and CIC mass-assignments schemes respectively. Grey shaded area denotes a 1% discrepancy w.r.t. the reference power spectrum.

### 4.1. Initial conditions

Ideally, the statistical properties of  $N$ -body simulations at late times should be independent of initial conditions, but studies have shown this is not the case. For example, Crocce et al. (2006) suggested that using 2LPT instead of 1LPT could allow simulations to begin at a later time, reducing computational effort. Similar approaches were extended to 3LPT by Michaux et al. (2021) and to fourth-order (4LPT) by List et al. (2024), which detailed improvements in particle resampling and aliasing mitigation.

This section systematically examines how the perturbative order, starting redshift, Poisson solver, and gradient order for force computations affect the results. Particles are initialised at cell centres, with results for initialisation at cell edges shown in Appendix B.2.

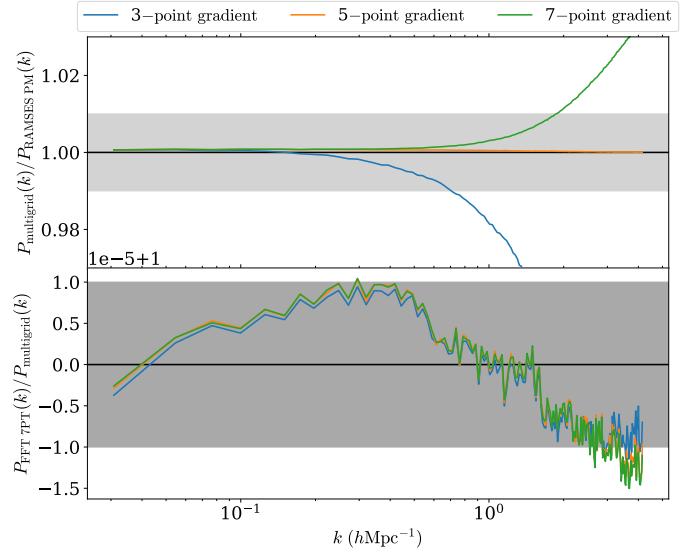
In Fig. 3, the ratio of the matter power spectrum at  $z = 0$  for various starting redshifts is compared to a reference simulation starting at  $z_{\text{ini}} = 150$ . It indicates that using five- and seven-point gradient methods produces nearly identical results, while the three-point gradient shows significant deviations for both



**Fig. 4.** Same as Fig. 3, but varying the  $N$ -body solver instead of gradient order, which is kept to seven-point (except for the FULL\_FFT solver which does not make use of finite gradients).

2LPT and 3LPT. Thus, at least a five-point gradient is necessary to achieve convergence regarding the influence of initial conditions on late-time clustering statistics. Additionally, the CIC interpolation method results in more scattered data compared to the TSC method, which produces smoother density fields and is less prone to large variations in potential and force calculations. The results show a remarkable agreement within 0.1% at  $z_{\text{ini}} \geq 125, 30$  and 10 for 1LPT, 2LPT and 3LPT respectively. This contrasts with Michaux et al. (2021), where power suppression at intermediate scales was observed before increasing around the Nyquist frequency. In the present study, a maximum wavenumber  $k_{\text{max}} = 2k_{\text{Nyq}}/3$  was used to avoid aliasing, yielding excellent agreement even at  $z_{\text{ini}} = 10$  with 3LPT. However, small-scale agreement can break down if these scales are not well-resolved, for instance, if the simulation box size increases but the number of particles and cells is kept constant.

Fig. 4 explores variations in the  $N$ -body solver. Multigrid and FFT\_7pt solvers produce nearly identical results, as both assume a seven-point Lagrangian operator, though Multigrid computes it in configuration space, while FFT\_7pt does so in Fourier space. Simulations initialised at later times exhibit excess small-scale power compared to the FFT solver in Fig. 3, likely due to their poorer resolution of small scales compared to FFT (similarly to the effect of the three-point gradient stencil, as we will see in Section 4.2). The FULL\_FFT solver shows similar results



**Fig. 5.** Top panel: ratio of the power spectrum for three-, five- and seven-point gradient in blue, orange and green lines respectively using multigrid, with respect to RAMSES (with AMR disabled and also using multigrid). Because RAMSES does not stop exactly at  $z = 0$  we rescaled its power spectrum the linear growth factor. The grey shaded area show the  $\pm 1\%$  limits. Bottom panel: ratio of the power spectrum using the PySCo FFT\_7pt solver w.r.t PySCo multigrid. The dark grey shaded area show the  $10^{-5}$  limits. In any case we use a TSC algorithm.

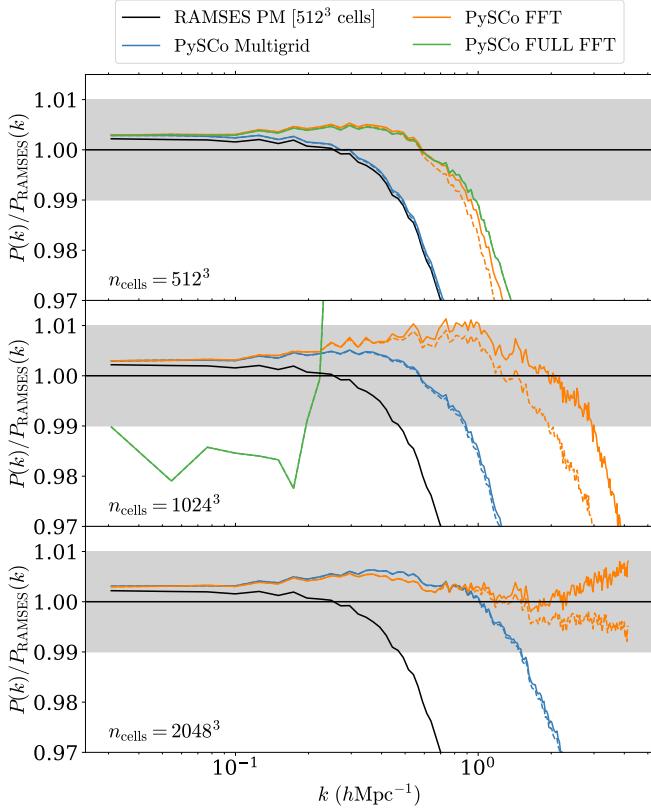
to FFT when a TSC scheme is applied. However, using CIC with FULL\_FFT fails to achieve convergence, likely due to the solver's sensitivity to the smoothness of the density field, as discussed in Section 4.2.

These findings suggest that simulations must use at least a five-point gradient to yield accurate results, with the ideal starting redshift depending on the LPT order. Additionally, employing TSC provides smoother results, and the FULL\_FFT solver should be avoided with the CIC scheme. Finally, to achieve good convergence at all scales, simulations need to adequately resolve the small scales. This thus validates our implementation of LPT initial conditions in PySCo.

#### 4.2. Comparison to RAMSES

This section compares PySCo with RAMSES by running a PM-only RAMSES simulation (disabling AMR) starting at  $z_{\text{ini}} = 50$  with 2LPT initial conditions generated by MPGRAFIC (Prunet et al. 2008). The same initial conditions are used for all comparisons between PySCo and RAMSES.

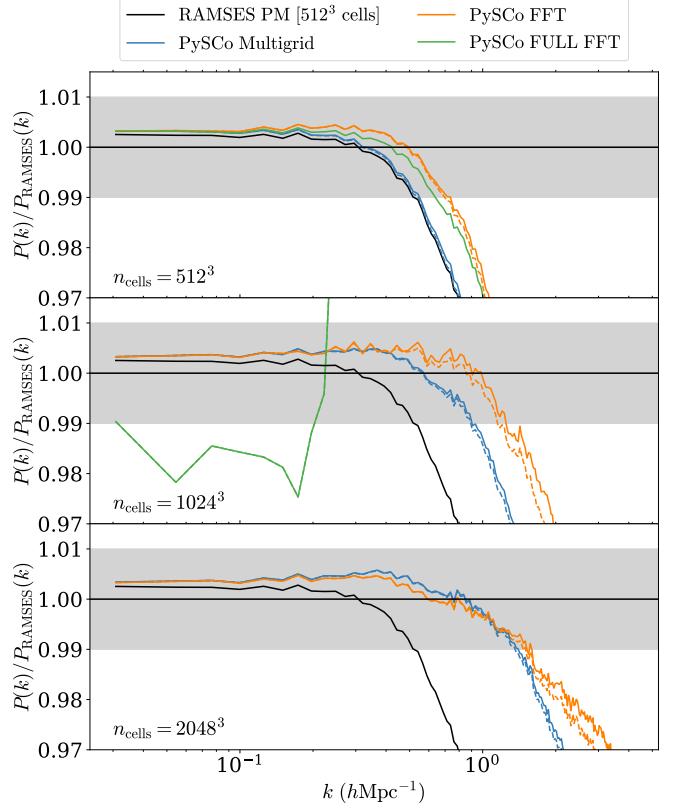
In Fig. 5, we observe remarkable agreement between PySCo with a five-point gradient and RAMSES (which also uses a five-point stencil with multigrid), with differences at only the 0.01% level. This validates the multigrid implementation in PySCo. Using a three-point gradient leads to a significant damping of the power spectrum at small scales, while the seven-point gradient shows an increase at even smaller scales. Based on these results and those from Section 4.1, it is clear that a three-point stencil is suboptimal, as the small runtime gain is outweighed by the power loss at small scales. Also shown in Fig. 5 is the power spectrum ratio between the multigrid and FFT\_7pt solvers, with both solvers agreeing at the  $10^{-5}$  level independently of the gradient order. This confirms the agreement already seen in Fig. 4. Small fluctuations around unity could be due to the convergence threshold of the multigrid algorithm (set at  $\alpha = 0.005$ ). Given



**Fig. 6.** Ratio of the power spectrum with respect to RAMSES (with AMR). In black lines we show results for RAMSES PM (no AMR), while in blue, orange and green we show results for PySCo using multigrid, FFT and FULL\_FFT solvers respectively. In top, middle and bottom panels we use  $512^3$ ,  $1024^3$  and  $2048^3$  cells respectively for PySCo. In dashed and solid lines we plot results using five- and seven-point gradient operator. The grey shaded area show the  $\pm 1\%$  limits. We use TSC in any case, and we do not plot FULL\_FFT for  $n_{\text{cells}} = 2048^3$ .

this close agreement, FFT\_7pt results are not shown further in this paper, except for performance analysis in Section 4.6.

In Fig. 6, the comparison between RAMSES (AMR) and PySCo reveals that using a seven-point Laplacian operator in a PM-only code results in a significant suppression of small-scale power compared to AMR (as seen in RAMSES PM, as well as PySCo multigrid and FFT\_7pt). However, using FFT or FULL\_FFT solvers improves small-scale resolution by a factor of two in wavenumbers before resolution effects become significant. With  $n_{\text{cells}} = 1024^3$ , the FULL\_FFT solver fails entirely due to large scatter in the density grid, which contains eight times more cells than particles. Therefore, FULL\_FFT can only be used with a smooth field, where  $n_{\text{part}} \geq n_{\text{cells}}$  and when a TSC scheme is employed. For other cases, a more sophisticated approach would be required, such as computing the mass-assignment kernel in configuration space for force computation and then Fourier-transforming it (Hockney & Eastwood 1981), which would be computationally expensive. For  $n_{\text{cells}} = 1024^3$  and  $2048^3$ , PySCo with multigrid gains factors of 2 and 4 in wavenumbers, respectively, as expected. The approximated seven-point Laplacian operator already smooths the field significantly, so there is little difference between using five- or seven-point gradients. Using the FFT solver instead leads to more accurate small scales. For  $n_{\text{cells}} = 1024^3$ , using a seven-point gradient achieves higher wavenumbers than a five-point gradient. With  $n_{\text{cells}} = 2048^3$ , FFT agrees with RAMSES at



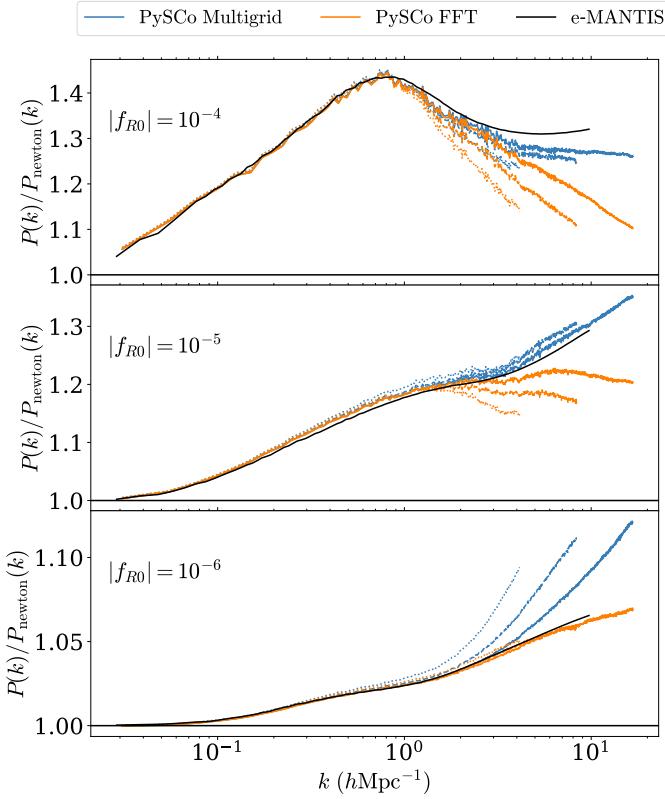
**Fig. 7.** Same as Fig. 6 but using CIC.

the percent level across the full range, although the plots are restricted to  $k_{\text{max}} = 2k_{\text{Nyq}}/3$  (where  $k_{\text{Nyq}}$  assumes  $n_{\text{cells}} = 512^3$ ). In Fig. 7, similar results are presented using CIC instead of TSC (also for the reference RAMSES simulation). The conclusions for multigrid and FULL\_FFT remain unchanged. However, the wavenumber gain with FFT is smaller compared to TSC, particularly for  $n_{\text{cells}} = 2048^3$ , where multigrid and FFT exhibit similar behaviour and deviate from RAMSES at the same scale. This indicates that for FFT, a smooth field is critical, though to a lesser extent than for FULL\_FFT.

To achieve the most accurate results compared to RAMSES, it is necessary to use the FFT solver with TSC and a seven-point stencil for the gradient operator. Otherwise, a five-point gradient can be used without a loss of accuracy. We also remark a slight shift on large scales between the reference (AMR case) and PM runs. Because the initial conditions are the same and the cosmological tables very similar, we expect this small difference (roughly 0.1%) to come from the fact that for RAMSES we do not stop exactly at  $z = 0$  (although we correct analytically for that), and that the time stepping is also slightly different, as for the AMR case we enter in a regime where the free-fall time step dominates over the cosmological time step at earlier times (see also Appendix B.1).

#### 4.3. $f(R)$ gravity

In this section, we validate the implementation of the  $f(R)$  gravity model from Hu & Sawicki (2007) described in Section 2.5. To assess this, we run  $f(R)$  simulations with varying  $f_{R0}$  values, along with a reference Newtonian simulation. Furthermore, because  $f(R)$  corrections are irrelevant at very high redshifts, we use the same Newtonian initial conditions in both cases. In Fig. 8, we compare the power spectrum boost from PySCo with



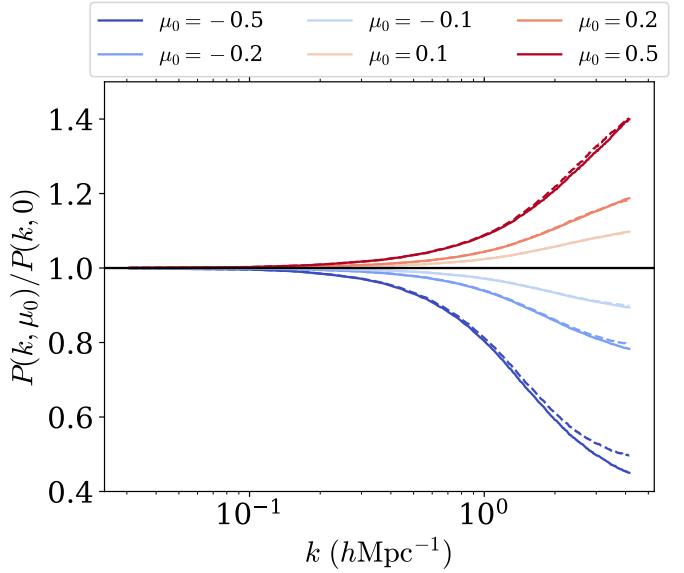
**Fig. 8.** Power spectrum boost from  $f(R)$  gravity w.r.t. to the Newtonian case. Blue and orange lines refer respectively to multigrid and FFT Poisson solvers with PySCo (in any case, the scalaron field equation is solved with non-linear multigrid), while black lines show the e-MANTIS emulator. In dotted, dashed and solid lines we use  $n_{\text{cells}} = 512^3, 1024^3$  and  $2048^3$ . Top, middle and bottom panels have the values  $|f_{R0}| = 10^{-4}, 10^{-5}$  and  $10^{-6}$ . In any case we use a seven-point gradient operator.

the e-MANTIS emulator (Sáez-Casares et al. 2024), which is based on the ECOSMOG code (Li et al. 2012), itself a modified version of RAMSES. The results show excellent agreement between PySCo and e-MANTIS up to  $k \sim 1 \text{ hMpc}^{-1}$ . The agreement improves further for higher  $n_{\text{cells}}$ , except in the case of  $|f_{R0}| = 10^{-6}$ , where the curves for PySCo overlap when using FFT. This indicates that PySCo converges well towards the e-MANTIS predictions. Notably, the best agreement is found when using the multigrid solver for  $|f_{R0}| = 10^{-4}$  and  $10^{-5}$ , while FFT performs better for  $|f_{R0}| = 10^{-6}$ . This last behaviour could be explained by the multigrid PM solver struggling to accurately compute small-scale features of the scalaron field, as lower values of  $|f_{R0}|$  result in sharper transitions between Newtonian and modified gravity regimes. No noticeable impact from the gradient order is observed in any of the cases.

In summary, PySCo demonstrates excellent agreement with e-MANTIS, aligning with prior validation efforts against other codes for similar tests, such as those conducted by Euclid Collaboration: Adamek et al. (2025). For this setup, using the multigrid solver for the Newtonian part seems advantageous for consistency, given that its non-linear version is already employed for the additional  $f(R)$  field without approximations.

#### 4.4. Parametrised gravity

The focus here shifts to simulations with parametrised gravity, where deviations from Newtonian gravity are governed by



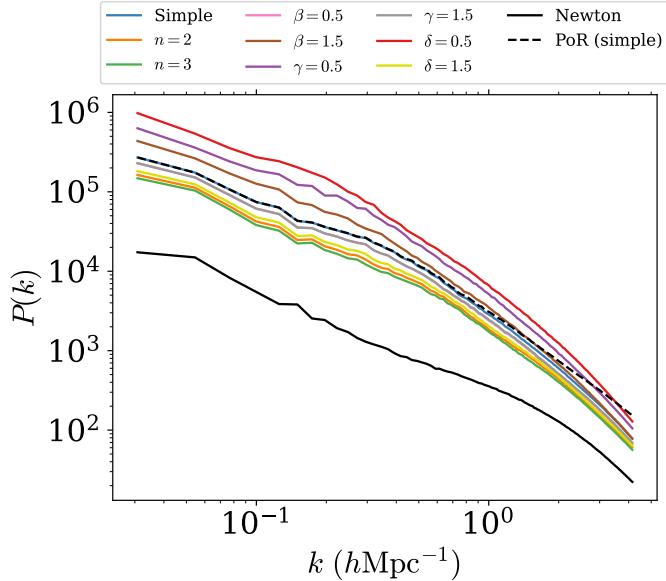
**Fig. 9.** Power spectrum boost from parametrised gravity w.r.t to the Newtonian case ( $\mu_0 = 0$ ). Coloured lines refer to different values of  $\mu_0$ , while solid and dashed lines indicates the use of FFT and multigrid solvers respectively.

a single parameter,  $\mu_0$ , representing the gravitational coupling today (as discussed in Section 2.4). In Fig. 9, the power spectrum boost is shown for various values of  $\mu_0$  compared to a Newtonian simulation. On large scales, the power spectrum ratio approaches unity, which aligns with expectations since the power spectrum is rescaled at the initial redshift according to  $\mu_0$  (detailed in Appendix A). On smaller scales, the behaviour changes: negative values of  $\mu_0$  result in a suppression of power, while positive values lead to an excess of power. The magnitude of these deviations increases with larger  $|\mu_0|$ , and the asymmetry between positive and negative values becomes evident. For instance, the departure from Newtonian behaviour is around 60% for  $\mu_0 = -0.5$  and around 40% for  $\mu_0 = 0.5$ .

There is a slight discrepancy between the results of the FFT and multigrid solvers for larger values of  $|\mu_0|$ , although no significant impact from the gradient stencil order or the number of cells is observed at the same scales.

#### 4.5. MOND

This section discusses the testing of the MOND implementation within PySCo. MOND was originally proposed as an alternative explanation for dark matter, modifying Newton's gravitational law in low-acceleration regimes. However, for validation purposes, the same cosmological parameters and initial conditions from Section 4.2 are used, rather than the typical MOND universe with  $\Omega_m = \Omega_b$  as the goal is to test the MOND gravity solver. The results are illustrated in Fig. 10. The MOND power spectra are noticeably higher than the Newtonian reference. This result aligns with known characteristics of MOND, which accelerates structure formation (Sanders 2001; Nusser 2002), explaining why MOND simulations are usually initialised with lower values of  $A_s$  (or  $\sigma_8$ ) (Knebe & Gibson 2004; Llinares et al. 2008). To validate the MOND implementation, it is compared to the PoR (Phantom of RAMSES) code (Lüghausen et al. 2015), a MOND patch for RAMSES that uses the simple interpolating function from Equation (14). The agreement between PoR and PySCo is excellent for scales  $k \lesssim 1 \text{ hMpc}^{-1}$ .



**Fig. 10.** Power spectra for several MOND interpolating functions and parameters (in coloured lines). In back solid line we show the Newtonian counterpart, while the black dashed line refers to a MOND run with the simple parametrization using the code `Phantom` of `RAMSES`. For MOND simulations we use  $g_0 = 10^{-10} \text{ m s}^{-2}$  and  $\mathcal{N} = 0$ . In PySCo, we use the multigrid solver in any case.

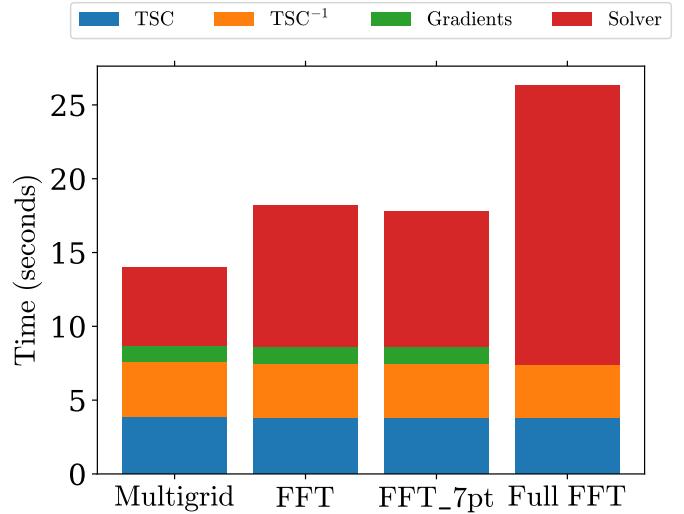
The discrepancies observed at small scales stem from differences between the PM and PM-AMR solvers, a pattern also seen in Fig. 5. Furthermore, the impact of the interpolating function on the power spectrum follows the same trend as observed in Fig. 2. For consistency, we also verified that MOND power spectra converge towards the Newtonian case when  $g_0 \ll 10^{-10} \text{ m s}^{-2}$ .

#### 4.6. Performances

Finally, we present performance metrics for PySCo on the Adastrà supercomputer at CINES, using AMD Genoa EPYC 9654 processors. For FFT-based methods, the PyFFTW package<sup>3</sup>, a Python wrapper for the FFTW library (Frigo & Johnson 2005), was used. All performance tests were run five times, and the median timings were taken to avoid outliers from potential node-related issues.

The first benchmark focuses on the time required to compute a single time step, as shown in Fig. 11. In a simulation with  $512^3$  particles and cells, the computation takes between 15 to 25 seconds on a single CPU for the various solvers, with the multigrid solver being the fastest and the FULL\_FFT solver being the slowest. The force computation using a seven-point gradient from the gravitational potential grid contributes minimally to the overall runtime. Since the multigrid solver outperforms FFT\_7pt with similar accuracy, the latter will not be used in the remainder of the paper.

Fig. 12 illustrates the strong scaling efficiency of PySCo's FFT and multigrid solvers. The scaling improves as the number of cells increases, suggesting that the workload per CPU becomes more efficient with larger grids. For smaller grids ( $n_{\text{cells}} = 128^3$ ), the multithreading is less effective, whereas for larger grids ( $n_{\text{cells}} = 1024^3$ ), multigrid reaches an efficiency



**Fig. 11.** Time to compute one time step of a Newtonian simulation for different solvers. In blue, orange, green and red we highlight the time needed to compute a density grid using TSC, an inverse TSC to interpolate the force from grid to particles, a seven-point gradient and the solver respectively. The FULL\_FFT method directly computes the force with the kernel described in Eq. (106), hereby removing the need for a finite-derivative gradient part. All timings were done using a single CPU.

of roughly 90% with 64 CPUs. Overall, multigrid consistently exhibits better efficiency than FFTW.

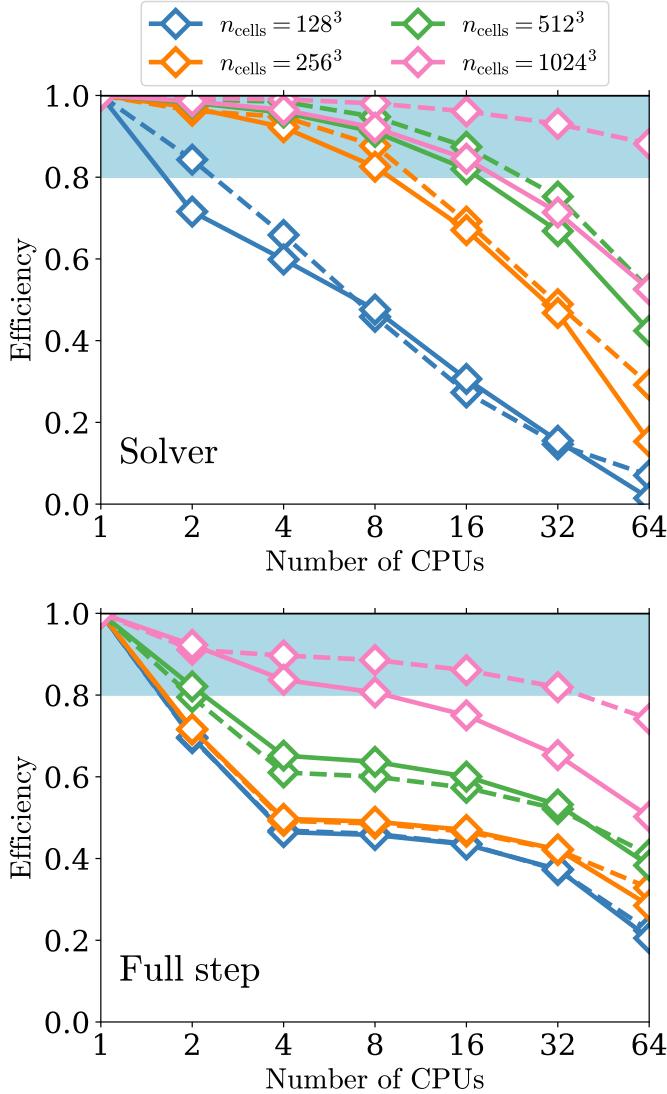
When analysing the total time per time step, a slightly different picture emerges. Efficiency still improves with larger grids but is generally lower than solver-only performance. For smaller grids ( $n_{\text{cells}} = 128^3$  and  $256^3$ ), there is little difference between multigrid and FFT, indicating that particle-grid interactions are the primary factor influencing performance, as  $512^3$  particles were used in all cases. A significant drop in efficiency occurs when using fewer than four CPUs, which is followed by a more gradual decline. This is caused by race conditions in the TSC algorithm, where multiple threads attempt to write to the same element in the density grid. To address this, atomic operations were implemented to ensure thread-safe modifications, but these operations slow the mass-assignment process by a factor of four<sup>4</sup>. Hence, when  $N_{\text{CPU}} < 4$  we use the sequential TSC version (which thus does not scale at all by definition), and the parallel-safe version for  $N_{\text{CPU}} \geq 4$ , thus giving a better scaling afterwards.

For grids with  $n_{\text{cells}} = 512^3$ , FFT achieves better efficiency than multigrid because, while multigrid is more efficient as a solver, it constitutes a smaller portion of the overall runtime. However, with larger grids  $n_{\text{cells}} = 1024^3$ , the efficiency aligns more closely with the solver-only case, as the solver dominates the runtime. Efficiency reaches approximately 50% for FFT and 75% for multigrid. To ensure optimal efficiency, the number of grid cells should thus be at least eight times the number of particles  $n_{\text{cells}} \geq 8n_{\text{part}}$ .

For comparison, a full simulation for PySCo with  $512^3$  particles and as many cells takes roughly 0.8 CPU hours for

<sup>4</sup> Unfortunately, Numba does not currently contain a way to deal with race conditions on CPUs natively. We thus had to rely on a modified version of atomic operations developed in <https://github.com/KatanaGraph/katana>. We will move on the official implementation of Numba atomics once they are developed.

<sup>3</sup> <https://github.com/pyFFTW/pyFFTW>



**Fig. 12.** Strong scaling efficiency for the solver only (top panel) and full time step (bottom panel). The solid and dashed lines refer to FFT and multigrid solvers respectively. The efficiency is computed as the time ratio for one CPUs w.r.t  $N$  CPUs, divided by  $N$ . An efficiency equal to unity denotes perfect parallel scaling.

~200 time steps, while a RAMSES run with the same setup and AMR enabled takes roughly 3000 CPU hours for ~1000 time steps.

## 5. Conclusions

In this paper, we presented PySCo, a fast particle-mesh  $N$ -body code designed for Newtonian and modified-gravity cosmological simulations. PySCo currently supports the  $f(R)$  gravity model (Hu & Sawicki 2007), the quasi-linear formulation of MOND (Milgrom 2010), and a time-dependent modification of the gravitational constant via a parameterised model. The code includes multiple solvers, such as multigrid and several FFT methods, each with exact or approximated kernels. We validated PySCo against RAMSES (using multigrid and no AMR), the e-MANTIS emulator for  $f(R)$  gravity and Phantom of RAMSES for MOND,

and found overall good agreement. The main conclusions of our study are the following:

- The TSC mass assignment should be preferred over CIC. CIC leads to larger scatters in the power spectrum and inaccuracies, particularly with FFT solvers when the number of cells exceeds the number of particles;
- Three-point gradient operators result in a loss of power at small scales. Using at least a five-point gradient is recommended for better accuracy;
- Seven-point Laplacian operators, whether used in multi-grid (configuration space) or FFT\_7pt (Fourier space), cause power suppression at small scales. Multigrid is favoured because it provides undistinguishable results from FFT\_7pt, while being faster and more efficient in parallel computing;
- The FFT solver (using the exact Laplacian kernel) is a solid choice overall, especially with TSC instead of CIC when the number of cells exceeds the number of particles;
- The FULL\_FFT solver is too simplistic to be trusted, except in configurations where the number of cells matches the number of particles ( $n_{\text{cells}} = n_{\text{part}}$ ) and TSC is used;
- Resolving small scales sufficiently is critical for obtaining a power spectrum at  $z = 0$  that is relatively insensitive to the initial redshift. Otherwise, there will be bias at small scales. Solutions include increasing the resolution (for example, reducing box size), using a higher-order gradient operator, or applying an exact Laplacian operator;
- When small scales are well-resolved, the power spectrum at  $z = 0$  becomes insensitive to the starting redshift at around 0.1% for initial redshifts greater than 125, 30, and 10 for 1LPT, 2LPT, and 3LPT, respectively;
- We recommend running PySCo with  $n_{\text{cells}} = 8n_{\text{part}}$  for optimal efficiency. Increasing to  $n_{\text{cells}} = 64n_{\text{part}}$  would significantly raise the runtime, contradicting the main advantage of PM-only simulations, which is precisely to be faster than AMR or tree codes;
- For standalone simulations, the FFT solver is ideal, as it aligns best with PM-AMR results. For calculating boost factors, such as  $f(R)$  simulations against Newtonian ones, multigrid is a faster option, with power losses at small scales cancelling out.

In this paper, we focused on power spectrum, but it would be also possible to compute accurate boost factors for halo mass functions (HMF), as in [Euclid Collaboration: Adamek et al. \(2025\)](#), although the HMF itself (not the boost) is very sensitive to small scales.

PySCo is usable both through command line and as a Python package. It also includes a wide range of built-in utilities, such as initial condition generators and power spectrum estimators. To minimise third-party dependencies, PySCo integrates several in-house modules for particle and mesh algorithms.

In the future, we will port PySCo to GPUs using Numba for CUDA kernels. However, there is currently no plan to implement MPI parallelisation, as the typical PM use case involves small, fast simulations on a single node. PySCo is particularly interesting because it comes with the fast development speed of Python and C/Fortran-like performance. Compared to PM-AMR codes like RAMSES, PySCo offers a dramatic reduction in runtime—potentially simulating in roughly 1/1000th of the time, albeit with less accuracy on small scales. These features make it useful for producing emulators, generating covariance matrices, and training neural networks for upcoming cosmological surveys without requiring massive computational resources. Our aim is to expand PySCo to include a wider range of gravity theories and

non-standard dark sector components, creating a versatile framework for exploring cosmological phenomena and aiding in the development of observational constraints.

**Acknowledgements.** I would like to thank Yann Rasera, Iñigo Sáez-Casares and Himanshu Ganjoo for comments on the draft and testing a preliminary version of PySCo. I am also grateful to Stéphane Colombi and Oliver Hahn for insightful discussions regarding initial conditions, as well as Hans Winther on modified gravity simulations and Claudio Llinares for interesting prospects with MOND. I thank Jorge Carretero and Pau Tallada, whose work on the SciPic package convinced me of the viability of Python for performant scientific computing. This project was provided with computer and storage resources by GENCI at CINES thanks to the grant 2023-A0150402287 on the supercomputer Adastrà's GENOA partition.

## References

Abbott, T. M. C., Abdalla, F. B., Avila, S., et al. 2019, *Phys. Rev. D*, **99**, 123505

Adamek, J., & Fidler, C. 2019, *J. Cosmology Astropart. Phys.*, **2019**, 026

Adamek, J., Daverio, D., Durrer, R., & Kunz, M. 2016, *JCAP*, **7**, 053

Alam, S., Aubert, M., Avila, S., et al. 2021, *Phys. Rev. D*, **103**, 083533

Amendola, L., Kunz, M., & Sapone, D. 2008, *J. Cosmology Astropart. Phys.*, **2008**, 013

Angulo, R. E., & Hahn, O. 2022, *Liv. Rev. Computat. Astrophys.*, **8**, 1

Angulo, R. E., & Pontzen, A. 2016, *MNRAS*, **462**, L1

Angus, G. W., van der Heyden, K. J., Famaey, B., et al. 2012, *MNRAS*, **421**, 2598

Astropy Collaboration (Price-Whelan, A. M., et al.) 2022, *ApJ*, **935**, 167

Bardeen, J. M. 1980, *PrD*, **22**, 1882

Barrera-Hinojosa, C., & Li, B. 2020, *J. Cosmology Astropart. Phys.*, **2020**, 007

Begeman, K. G., Broeils, A. H., & Sanders, R. H. 1991, *MNRAS*, **249**, 523

Behnel, S., Bradshaw, R., Citro, C., et al. 2011, *Comput. Sci. Eng.*, **13**, 31

Bekenstein, J. 2006, *Contemp. Phys.*, **47**, 387

Bekenstein, J., & Milgrom, M. 1984, *ApJ*, **286**, 7

Bernardeau, F., Colombi, S., Gaztañaga, E., & Scoccimarro, R. 2002, *PhysRep.*, **367**, 1

Bertschinger, E. 2001, *APJS*, **137**, 1

Bonvin, C., & Durrer, R. 2011, *PrD*, **84**, 063505

Bose, S., Li, B., Barreira, A., et al. 2017, *J. Cosmology Astropart. Phys.*, **2017**, 050

Bouchet, F. R., Colombi, S., Hivon, E., & Juszkiewicz, R. 1995, *A&A*, **296**, 575

Breton, M.-A., Rasera, Y., Taruya, A., Lacombe, O., & Saga, S. 2019, *MNRAS*, **483**, 2671

Buchdahl, H. A. 1970, *MNRAS*, **150**, 1

Buchert, T., & Ehlers, J. 1993, *MNRAS*, **264**, 375

Burridge, C., & Sakstein, J. 2018, *Liv. Rev. Relativ.*, **21**, 1

Candlish, G. N., Smith, R., & Fellhauer, M. 2015, *MNRAS*, **446**, 1060

Catelan, P. 1995, *MNRAS*, **276**, 115

Challinor, A., & Lewis, A. 2011, *PrD*, **84**, 043516

Chevallier, M., & Polarski, D. 2001, *Int. J. Mod. Phys. D*, **10**, 213

Chisari, N. E., & Zaldarriaga, M. 2011, *PrD*, **83**, 123505

Couchman, H. M. P. 1991, *ApJ*, **368**, L23

Crocce, M., Pueblas, S., & Scoccimarro, R. 2006, *MNRAS*, **373**, 369

Dakin, J., Hannestad, S., & Tram, T. 2022, *MNRAS*, **513**, 991

Dvali, G., Gabadadze, G., & Porrati, M. 2000, *Phys. Lett. B*, **485**, 208

Efstathiou, G., Davis, M., White, S. D. M., & Frenk, C. S. 1985, *ApJS*, **57**, 241

Euclid Collaboration (Adamek, J., et al.) 2025, *A&A*, in press, <https://doi.org/10.1051/0004-6361/202452180>

Euclid Collaboration (Castander, F. J., et al.) 2025, *A&A*, in press, <https://doi.org/10.1051/0004-6361/202450853>

Famaey, B., & Binney, J. 2005, *MNRAS*, **363**, 603

Famaey, B., & McGaugh, S. S. 2012, *Liv. Rev. Relativ.*, **15**, 10

Feng, Y., Chu, M.-Y., Seljak, U., & McDonald, P. 2016, *MNRAS*, **463**, 2273

Fidler, C., Rampf, C., Tram, T., et al. 2015, *PrD*, **92**, 123517

Fidler, C., Tram, T., Rampf, C., et al. 2016, *JCAP*, **9**, 031

Frigo, M., & Johnson, S. G. 2005, *Proc. IEEE*, **93**, 216

Frusciante, N., & Perenon, L. 2020, *Phys. Rep.*, **857**, 1

Garrison, L. H., Eisenstein, D. J., Ferrer, D., Maksimova, N. A., & Pinto, P. A. 2021, *MNRAS*, **508**, 575

Green, S. R., & Wald, R. M. 2014, *Class. Quant. Grav.*, **31**, 234003

Guillet, T., & Teyssier, R. 2011, *J. Computat. Phys.*, **230**, 4756

Hahn, O., & Abel, T. 2013, *Astrophysics Source Code Library* [[record ascl:1311.011](https://ascl.net/1311.011)]

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357

Hassani, F., & Lombriser, L. 2020, *MNRAS*, **497**, 1885

Hernquist, L., Bouchet, F. R., & Suto, Y. 1991, *ApJS*, **75**, 231

Hernández-Aguayo, C., Ruan, C.-Z., Li, B., et al. 2022, *J. Cosmol. Astropart. Phys.*, **2022**, 048

Hockney, R. W., & Eastwood, J. W. 1981, *Computer Simulation Using Particles* (New York: McGraw-Hill)

Hu, W., & Sawicki, I. 2007, *Phys. Rev. D*, **76**, 064004

Ishiyama, T., Fukushige, T., & Makino, J. 2009, *PASJ*, **61**, 1319

Ishiyama, T., Prada, F., Klypin, A. A., et al. 2021, *MNRAS*, **506**, 4210

Jeong, D. 2010, PhD thesis, University of Texas, Austin, USA

Jing, Y. P. 2005, *ApJ*, **620**, 559

Khoury, J., & Weltman, A. 2004, *Phys. Rev. Lett.*, **93**, 171104

Klypin, A., & Prada, F. 2018, *MNRAS*, **478**, 4602

Knebe, A., & Gibson, B. K. 2004, *MNRAS*, **347**, 1055

Knebe, A., Green, A., & Binney, J. 2001, *MNRAS*, **325**, 845

Kravtsov, A. V., Klypin, A. A., & Khokhlov, A. M. 1997, *APJS*, **111**, 73

Lam, S. K., Pitrou, A., & Seibert, S. 2015, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, 1

Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, **538**, 473

Li, B., Zhao, G.-B., Teyssier, R., & Koyama, K. 2012, *JCAP*, **1**, 051

Linder, E. V. 2003, *Phys. Rev. Lett.*, **90**, 091301

List, F., Hahn, O., & Rampf, C. 2024, *Phys. Rev. Lett.*, **132**, 131003

Llinares, C. 2018, *Int. J. Mod. Phys. D*, **27**, 1848003

Llinares, C., Knebe, A., & Zhao, H. 2008, *MNRAS*, **391**, 1778

Llinares, C., Mota, D. F., & Winther, H. A. 2014, *A&A*, **562**, A78

Lüghausen, F., Famaey, B., & Kroupa, P. 2015, *Can. J. Phys.*, **93**, 232

Ma, C.-P., & Bertschinger, E. 1995, *ApJ*, **455**, 7

Martel, H., & Shapiro, P. R. 1998, *MNRAS*, **297**, 467

Merz, H., Pen, U.-L., & Trac, H. 2005, *New Astron.*, **10**, 393

Michaux, M., Hahn, O., Rampf, C., & Angulo, R. E. 2021, *MNRAS*, **500**, 663

Milgrom, M. 1983, *ApJ*, **270**, 365

Milgrom, M. 2009, *Phys. Rev. D*, **80**, 123536

Milgrom, M. 2010, *MNRAS*, **403**, 886

Milgrom, M. 2023, *Phys. Rev. D*, **108**, 084005

Milgrom, M., & Sanders, R. H. 2008, *ApJ*, **678**, 131

Nusser, A. 2002, *MNRAS*, **331**, 909

Orszag, S. A. 1971, *J. Atmos. Sci.*, **28**, 1074

Oyaizu, H. 2008, *Phys. Rev. D*, **78**, 123523

Peebles, P. J. E. 1980, *The large-scale structure of the universe* (Princeton University Press)

Pen, U.-L. 1997, *ApJ*, **490**, L127

Planck Collaboration XIV. 2016, *A&A*, **594**, A14

Planck Collaboration VI. 2020, *A&A*, **641**, A6

Potter, D., Stadel, J., & Teyssier, R. 2017, *Computational Astrophys. Cosmol.*, **4**, 2

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in C*, 2nd edn. (Cambridge, USA: Cambridge University Press)

Prunet, S., Pichon, C., Aubert, D., et al. 2008, *ApJS*, **178**, 179

Puchwein, E., Baldi, M., & Springel, V. 2013, *MNRAS*, **436**, 348

Rácz, G., Szapudi, I., Csabai, I., & Dobos, L. 2021, *MNRAS*, **503**, 5638

Rampf, C., & Buchert, T. 2012, *J. Cosmology Astropart. Phys.*, **2012**, 021

Rampf, C., List, F., & Hahn, O. 2025, *J. Cosmology Astropart. Phys.*, **2025**, 020

Rasera, Y., Breton, M. A., Corasaniti, P. S., et al. 2022, *A&A*, **661**, A90

Riess, A. G., Yuan, W., Macri, L. M., et al. 2022, *ApJ*, **934**, L7

Ruan, C.-Z., Hernández-Aguayo, C., Li, B., et al. 2022, *J. Cosmology Astropart. Phys.*, **2022**, 018

Sáez-Casares, I., Rasera, Y., & Li, B. 2024, *MNRAS*, **527**, 7242

Sanders, R. H. 2001, *ApJ*, **560**, 1

Scoccimarro, R. 1998, *MNRAS*, **299**, 1097

Simpson, F., Heymans, C., Parkinson, D., et al. 2013, *MNRAS*, **429**, 2249

Sirk, E. 2005, *ApJ*, **634**, 728

Skordis, C., & Złotnicki, T. 2021, *Phys. Rev. Lett.*, **127**, 161302

Sotiriou, T. P., & Faraoni, V. 2010, *Rev. Mod. Phys.*, **82**, 451

Springel, V., Pakmor, R., Zier, O., & Reinecke, M. 2021, *MNRAS*, **506**, 2871

Teyssier, R. 2002, *A&A*, **385**, 337

Valogiannis, G., & Bean, R. 2017, *Phys. Rev. D*, **95**, 103515

Villaescusa-Navarro, F. 2018, *Pylions: Python libraries for the analysis of numerical simulations*, Astrophysics Source Code Library [[record ascl:1811.008](https://ascl.net/1811.008)]

Visser, P. M., Eijt, S. W. H., & de Nijs, J. V. 2024, *A&A*, **681**, A90

Vogt, S. M. L., Bocquet, S., Davies, C. T., et al. 2025, *Phys. Rev. D*, **111**, 043519

Winther, H. A., Schmidt, F., Barreira, A., et al. 2015, *MNRAS*, **454**, 4208

Winther, H. A., Koyama, K., Manera, M., Wright, B. S., & Zhao, G.-B. 2017, *J. Cosmol. Astropart. Phys.*, **2017**, 006

Yoo, J., Fitzpatrick, A. L., & Zaldarriaga, M. 2009, *PrD*, **80**, 083514

Zel'dovich, Y. B. 1970, *A&A*, **5**, 84

## Appendix A: Growth factors

Under the Lagrangian formulation of perturbation theory, the position of tracers at some given time and position is given by

$$\mathbf{x}(\eta) = \mathbf{q} + \Psi(\mathbf{q}, \eta), \quad (\text{A.1})$$

with  $\mathbf{q}$  an initial position and  $\Psi(\mathbf{q}, \eta)$  a displacement field. The Euler equation is then

$$\frac{d^2\mathbf{x}}{d\eta^2} + \mathcal{H} \frac{d\mathbf{x}}{d\eta} = -\nabla\phi, \quad (\text{A.2})$$

which, when taking the gradient, yields

$$\nabla_{\mathbf{x}} \left[ \frac{d^2\mathbf{x}}{d\eta^2} + \mathcal{H} \frac{d\mathbf{x}}{d\eta} \right] = -\frac{3}{2} \mathcal{H}^2 \Omega_m(\eta) \delta(\mathbf{x}, \eta). \quad (\text{A.3})$$

Mass conservation imposes  $\bar{\rho}(\eta) d^3\mathbf{q} = \bar{\rho}(\eta) [1 + \delta(\mathbf{x}, \eta)] d^3\mathbf{x}$ , meaning that  $\nabla_{\mathbf{x}} = J^{-1} \nabla_{\mathbf{q}}$ , with  $J(\mathbf{q}, \eta) = \det[\delta_{ij} + \Psi_{ij}(\mathbf{q}, \eta)]$  where  $\delta_{ij}$  is a Kronecker Delta. We compute the time evolution of growth factors through the general equation (Jeong 2010)

$$J \nabla_{\mathbf{x}} \left[ \frac{d^2\mathbf{x}}{d\eta^2} + \mathcal{H} \frac{d\mathbf{x}}{d\eta} \right] = \frac{3}{2} \mathcal{H}^2 \Omega_m(\eta) [J - 1]. \quad (\text{A.4})$$

A perturbative expansion at third order gives (Rampf & Buchert 2012)

$$\Psi(\mathbf{q}, \eta) = \varepsilon D(\eta) \Psi^{(1)}(\mathbf{q}) + \varepsilon^2 E(\eta) \Psi^{(2)}(\mathbf{q}) + \varepsilon^3 F(\eta) \Psi^{(3)}(\mathbf{q}) + O(\varepsilon^4), \quad (\text{A.5})$$

$$\begin{aligned} J(\mathbf{q}, \eta) = 1 + \varepsilon D(\eta) \mu_1^{(1)}(\mathbf{q}) + \varepsilon^2 [E(\eta) \mu_1^{(2)}(\mathbf{q}) + D^2(\eta) \mu_2^{(1)}(\mathbf{q})] \\ + \varepsilon^3 [F(\eta) \mu_1^{(3)}(\mathbf{q}) + 2D(\eta) E(\eta) \mu_2^{(1,2)}(\mathbf{q}) + D^3(\eta) \mu_3^{(1)}(\mathbf{q})] \\ + O(\varepsilon^4), \end{aligned} \quad (\text{A.6})$$

with  $D \equiv D^{(1)}$ ,  $E \equiv D^{(2)}$  and  $F \equiv D^{(3)}$  are the first, second and third-order growth factors respectively,  $\varepsilon \ll 1$  and

$$\mu_1^{(n)}(\mathbf{q}) = \Psi_{i,i}^{(n)}(\mathbf{q}), \quad (\text{A.7})$$

$$\mu_2^{(n,m)}(\mathbf{q}) = \frac{1}{2} [\Psi_{i,i}^{(n)}(\mathbf{q}) \Psi_{j,j}^{(m)}(\mathbf{q}) - \Psi_{i,j}^{(n)}(\mathbf{q}) \Psi_{j,i}^{(m)}(\mathbf{q})], \quad (\text{A.8})$$

$$\mu_3^{(n)}(\mathbf{q}) = \det[\Psi_{i,j}^{(n)}(\mathbf{q})]. \quad (\text{A.9})$$

Eq. (A.4) then becomes

$$J [\delta_{ij} - \Psi_{i,j}] \left[ \frac{d^2\Psi_{i,j}}{d\eta^2} + \mathcal{H} \frac{d\Psi_{i,j}}{d\eta} \right] = \frac{3}{2} \mathcal{H}^2 \Omega_m(\eta) [J - 1], \quad (\text{A.10})$$

with  $\partial/\partial x_i = [\delta_{ij} + \Psi_{i,j}]^{-1} \partial/\partial q_j$  and  $[\delta_{ij} + \Psi_{i,j}]^{-1} \approx [\delta_{ij} - \Psi_{i,j}]$ .

- The first-order solution (Zel'dovich 1970) is

$$\ddot{D} + \mathcal{H}\dot{D} - \beta D = 0, \quad (\text{A.11})$$

with  $\beta = \frac{3}{2} \mathcal{H}^2 \Omega_m(\eta)$ . where  $D \equiv D_+$  the first-order growth factor, and  $\dot{D} = dD/d\eta$ .

- The second-order solution is then

$$\ddot{E} + \mathcal{H}\dot{E} - \beta [E - D^2] = 0. \quad (\text{A.12})$$

- The third-order solutions are given by

$$[\ddot{F} + \mathcal{H}\dot{F} - \beta F] \mu_1^{(3)}$$

$$\begin{aligned} &+ 2 [\ddot{D} \mathcal{H} + \mathcal{H} \dot{D} + \ddot{D} E + \mathcal{H} E \dot{D} - \beta D E] \mu_2^{(1,2)} \\ &+ [3 \ddot{D} D^2 + 3 \mathcal{H} \dot{D} D^2 - \beta D^3] \mu_3^{(1)} = 0, \end{aligned} \quad (\text{A.13})$$

which give rise to two tangential and one transversal modes. Their ordinary differential equations (ODEs) are given by

$$\ddot{F}_a + \mathcal{H} \dot{F}_a - \beta [F_a - 2D^3] = 0, \quad (\text{A.14})$$

$$\ddot{F}_b + \mathcal{H} \dot{F}_b - \beta [F_b - D(E - D^2)] = 0, \quad (\text{A.15})$$

$$\ddot{F}_c + \mathcal{H} [\dot{E} D - E \dot{D}] + \beta D^3 = 0, \quad (\text{A.16})$$

where the last equation can be derived using the equation for irrotational fluids (Catelan 1995). In practice, we implement these ODEs as functions of  $\ln(a)$ , leading to the following set of equations

$$D'' = -\gamma D' + \beta D, \quad (\text{A.17})$$

$$E'' = -\gamma E' + \beta [E - D^2], \quad (\text{A.18})$$

$$F_a'' = -\gamma F_a' + \beta [F_a - 2D^3], \quad (\text{A.19})$$

$$F_b'' = -\gamma F_b' + \beta [F_b - 2D(E - D^2)], \quad (\text{A.20})$$

$$F_c'' = (1 - \gamma) F_c' + E D' - D E' - \beta D^3, \quad (\text{A.21})$$

with  $D' = dD/d\ln a$  and

$$\beta = \frac{3}{2} \Omega_m(a), \quad (\text{A.22})$$

$$\gamma = \frac{1}{2} \{1 - 3\Omega_\Lambda(a) [w_0 + w_a(1 - a)] - \Omega_r(a)\}. \quad (\text{A.23})$$

The  $n$ th-order growth factors can be estimated through  $f^{(n)} = D^{(n)}/D^{(n)}$ . It is straightforward to implement a scale-independent parametrised form of modified gravity (as in Section 2.4) by making the small change  $\beta = 3\mu(a)\Omega_m(a)/2$ .

Approximated analytical solutions for the growing mode in a matter-dominated era (and neglecting radiation) are (Catelan 1995; Rampf & Buchert 2012)

$$D_+^{(2)} = -\frac{3}{7} D_+^2, \quad (\text{A.24})$$

$$D_+^{(3a)} = -\frac{1}{3} D_+^3, \quad (\text{A.25})$$

$$D_+^{(3b)} = \frac{10}{21} D_+^3, \quad (\text{A.26})$$

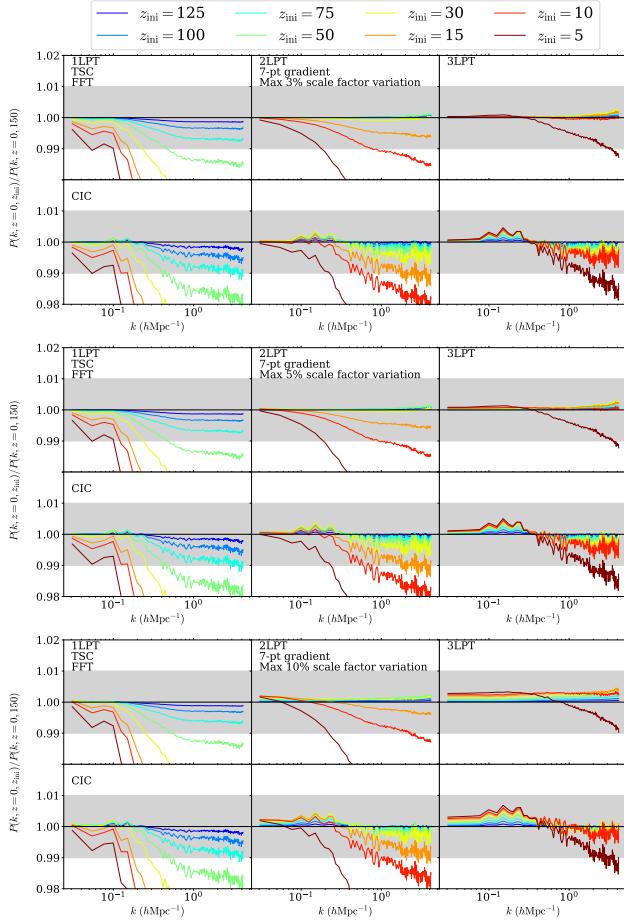
$$D_+^{(3c)} = -\frac{1}{7} D_+^3. \quad (\text{A.27})$$

Approximated fits with dependence on  $\Omega_m$  for the growth factors and growth rates are also available in Bouchet et al. (1995).

## Appendix B: Initial conditions

### Appendix B.1: Time stepping

As described in Section 3.4, we use two kinds of time stepping criteria, one based on cosmological time (in RAMSES, the scale factor cannot change by more than 10% in a single time step), and another based on the free-fall time (acceleration) or maximum velocity compared to the mesh size. For the latter two we also use a Courant-like factor, which multiplies the final time step and is equal to 0.8 by default.

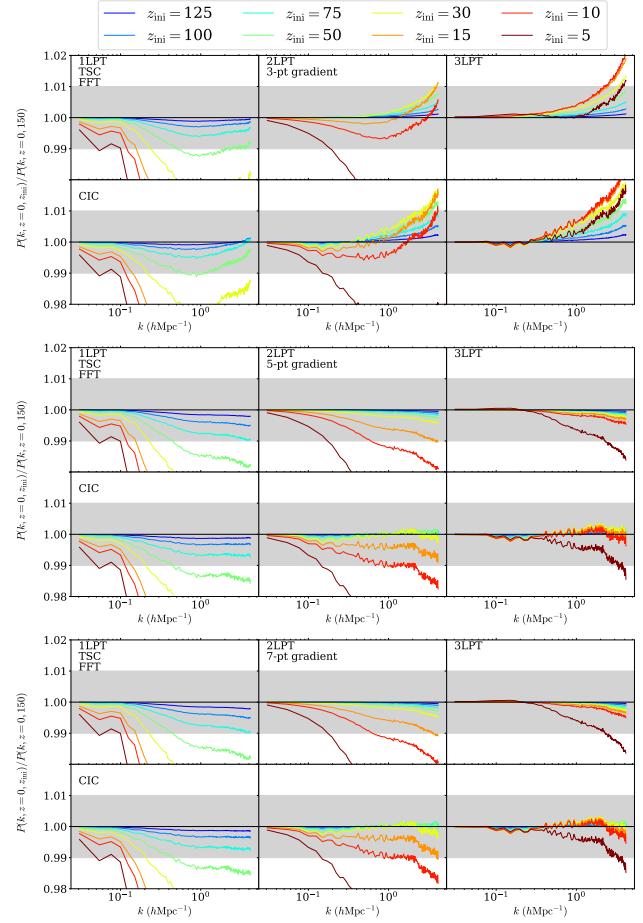


**Fig. B.1.** Same as Fig. 3, but using a seven-point gradient and varying the maximal scale factor variation in the cosmological time criterion, from top to bottom.

While small changes in the Courant factor do not significantly affect the results, using a time step that is too large can introduce a scale-independent bias in the power spectrum at redshift  $z = 0$ , as shown in Fig. B.1. We observe a bias of approximately 0.3% in the power spectrum at  $k \sim 0.1 h^{-1} \text{ Mpc}$  when the time step allows for a 10% maximum variation in the scale factor. This bias decreases to about 0.1% with a 5% variation and seems negligible when using variations below 3%. To ensure accuracy and avoid such biases, we used a more conservative 2% variation in scale factor by default. Additionally, we note that while 1LPT appears to work well on large scales, this is primarily due to the limited box size of the simulation, as we directly see the moment when the power spectrum is damped compared to the reference at initial redshift  $z = 150$ .

#### Appendix B.2: Initial positions

In PySCo, two methods for initialising particle positions were implemented: at cell centres or at cell edges. Given the periodic boundary conditions, these methods should ideally produce no differences due to translational symmetry. However, discrepancies can arise depending on the mass-assignment scheme used. For example, using the nearest grid point (NGP) method, where the density is calculated directly from the number of particles within each grid cell, will produce a uniform density grid if the particles are positioned at cell centres. This is because, at small displacements, each cell would contain exactly one particle. In

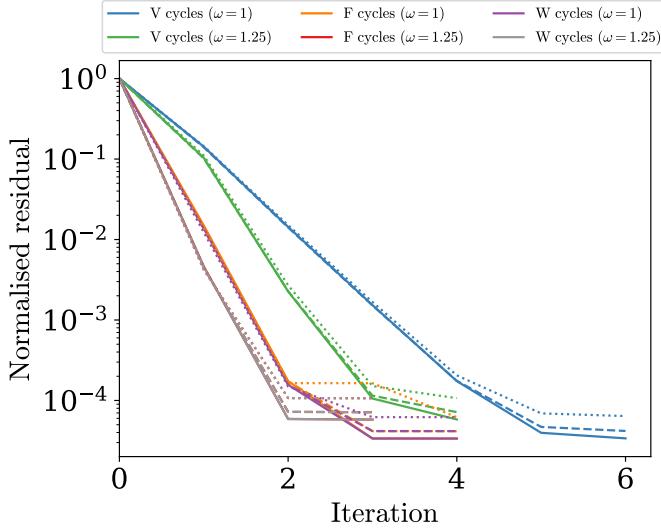


**Fig. B.2.** Same as Fig. 3, but initialising the particles at cell edges.

contrast, if the particles are initialised at the cell edges, the density field can become inhomogeneous, as it depends more strongly on the displacement field. This difference is demonstrated in Fig. B.2, which shows the ratio of the power spectrum for various initial redshifts compared to a reference case where the initial redshift is  $z_{\text{ini}} = 150$  and particles are initialised at cell edges. The results are quite similar to those from Fig. 3, where particles were initialised at cell centres. However, for five- and seven-point gradient operators with 3LPT, a slight bias appears at small scales, with a loss of power even for  $z_{\text{ini}} > 10$  compared to the reference case. As a result of these findings, we decided to initialise particles at cell centres in the main text.

#### Appendix C: Multigrid convergence

This section presents convergence tests for the multigrid algorithm (detailed in Section 3.7). Fig. C.1 shows the convergence rate of different multigrid cycles (V, F, and W). First, we remark that the convergence rate is nearly independent of redshift. The only redshift-dependent variation is the maximum residual suppression due to numerical roundings, which becomes evident after many iterations. This discrepancy arises from the normalisation by the first-guess residual, which is less accurate at lower redshifts due to the Universe being more structured compared to higher redshifts. Additionally, F and W cycles show very similar behaviour but take approximately 1.9 and 2.1 times longer than V cycles, respectively. Without overrelaxation ( $\omega = 1$ ), F and W cycles reduce the residual by a factor of 100 per iteration, while



**Fig. C.1.** Convergence rate of multigrid cycles as function of the number of cycles (or iterations). Blue, orange and purple lines refer to V, F and W cycles respectively, while in green, red and grey we use SOR with relaxation parameter  $\omega = 1.25$ . In solid, dashed and dotted lines we show results at  $z = 0, 1$  and  $3$  respectively. For the first guess we use one Jacobi step, thereby neglecting the information from previous step.

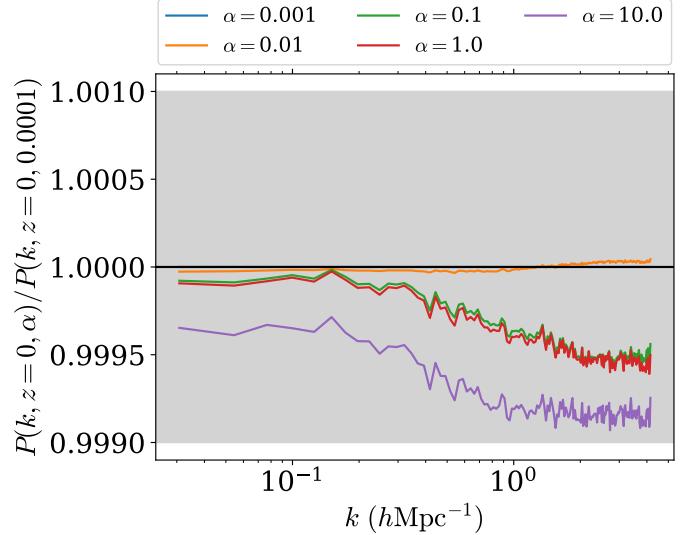
**Table C.1.** Ratio of residuals for different first guesses.

Redshift	No first guess	Last step	Last step + rescale
0	9.09	1.01	1
1	5.99	1.02	1
3	7.81	1.11	1

**Notes.** ‘No first guess’ means that we initialise the potential with one Jacobi step (see Section 3.6), ‘last step’ means that we use the potential field computed at the last step has first guess, and ‘last step + rescale’ has an additional rescaling (as described in Section 3.9). The latter is used as reference.

V cycles achieve a factor of 10 reduction. Overrelaxation with  $\omega = 1.25$  significantly improves the convergence rate, although it may not immediately help in the first V cycle due to the initial guess still being inaccurate. In practice,  $\omega = 1.25$  is applied for linear Poisson equations, whereas for non-linear cases,  $\omega = 1.0$  is used since SOR proves less effective. In Fig. C.1, one Jacobi step was used as the first guess (see Section 3.6), but Table C.1 suggests that using the potential field from the last step can significantly reduce the initial residuals. Further improvements through rescaling this last-step solution using the linear growth factor (discussed in Section 3.9) only offer marginal reductions in the residual. Rescaling becomes more useful at higher redshifts, where time steps are larger. Overall, using the last step as the first guess is a good strategy, sometimes saving one V cycle, though it requires keeping the gravitational potential grid in memory. Rescaling, however, seems unnecessary unless the time steps are particularly large.

The final test concerns the threshold criterion for convergence of the multigrid algorithm. Fig. C.2 shows the power spectrum ratio at  $z = 0$  for different values of  $\alpha$ . Even with a very large threshold value, like  $\alpha = 10$ , the bias in the power spectrum is only around 0.1%. The results for  $\alpha = 0.1$  and  $1$  are almost identical, likely because these thresholds are passed within the same V cycle, which typically reduces the residual by a factor



**Fig. C.2.** Power spectrum ratio with varying multigrid threshold parameter  $\alpha$  (see Eq. 96) w.r.t to the reference where  $\alpha = 10^{-4}$ . In blue, orange, green, red and purple we have  $\alpha = 0.001, 0.01, 0.1, 1$  and  $10$ . The grey shaded area indicates the  $\pm 0.1\%$  limits. In any case we use V cycles.

of 20–30 per iteration when using SOR. With  $\alpha = 0.01$ , the bias drops to less than 0.01%, and convergence is essentially achieved by  $\alpha = 0.001$ . These results are consistent regardless of starting redshift, LPT order, or mass-assignment scheme. In practice, an  $\alpha$  value of  $0.01$  is recommended, as it produces highly accurate results without introducing any significant bias, whereas increasing  $\alpha$  further results in a small but noticeable decrease in the power spectrum.

## Appendix D: Point-mass test

In this appendix, we evaluate the performance of PySCo in the simplified scenario of a single point mass while maintaining periodic boundary conditions. The normalised gravitational field for a homogeneous and isotropic sphere of radius  $R$  is given by

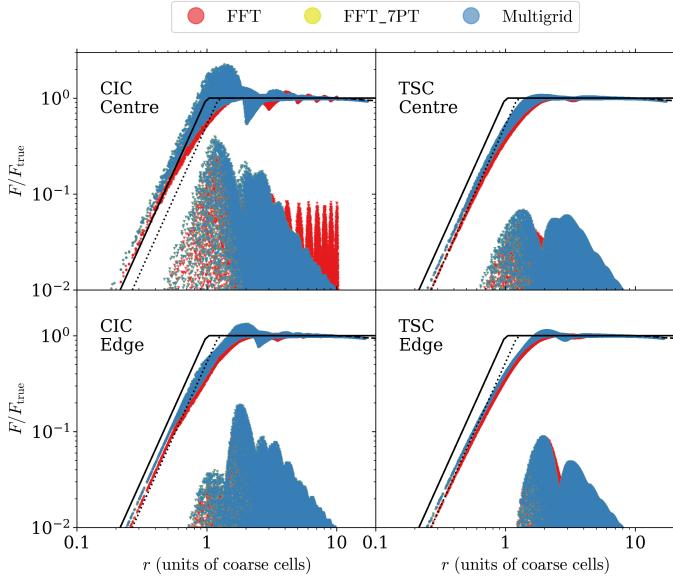
$$\mathbf{F}(\mathbf{r}, R) = \begin{cases} -\mathbf{r}/R^3, & r < R, \\ -\mathbf{r}/r^3, & r > R, \end{cases} \quad (\text{D.1})$$

where  $r = |\mathbf{r}|$ . To incorporate periodicity, we apply the Ewald summation method (Hernquist et al. 1991). For simplicity, we only account for the contribution from replicas, as in Rácz et al. (2021)

$$\tilde{\mathbf{F}}(\mathbf{r}, R) = \sum_{\mathbf{n}} \mathbf{F}(\mathbf{r} + \mathbf{n}, R), \quad (\text{D.2})$$

where a unit simulation box is assumed,  $\mathbf{n}$  represents integer triplets, and  $|(\mathbf{r} + \mathbf{n})| < 2.6$  ensures that the force decays to zero between replicas.

Fig. D.1 illustrates the acceleration of  $128^3$  massless particles surrounding a point mass in a simulation with  $64^3$  cells. The radial accelerations closely follow the predictions for the inner sphere at scales below the coarse cell size  $h$ . At larger scales, the acceleration transitions to the standard  $1/r^2$  behaviour, albeit with damping due to periodicity. This damping is well captured by Eq. (D.2). Notably, the TSC interpolation scheme demonstrates significantly improved isotropy compared to CIC, as evidenced by lower tangential acceleration values. For the CIC

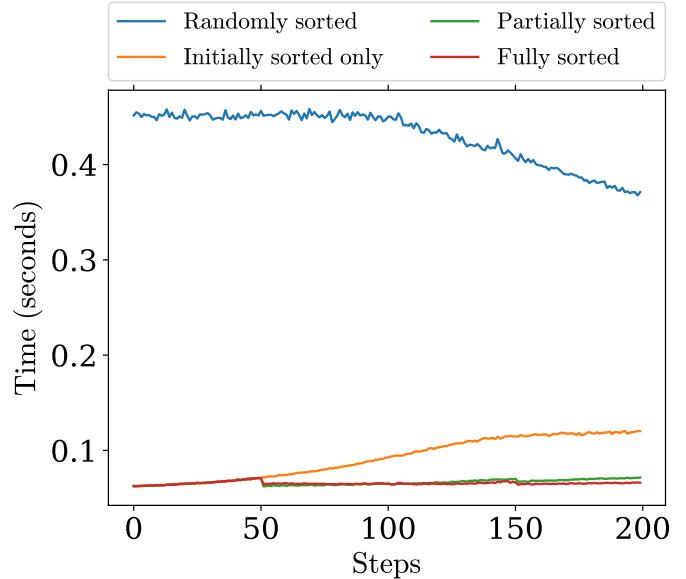


**Fig. D.1.** Acceleration of massless particles randomly sampled around a point mass, as a function of distance. The left and right panels illustrate results using the CIC and TSC mass interpolation (and inverse interpolation) schemes, respectively. The top panels show cases where the point mass is positioned at the centre of a cell, while the bottom panels depict cases where it is placed at the edge of a cell. The red, yellow, and blue data points represent accelerations computed using the FFT solver (uncorrected for the mass-assignment scheme), the FFT\_7pt solver, and the multigrid solver, respectively. Circles denote radial acceleration, while triangles indicate tangential acceleration. The black solid line represents the theoretical prediction  $F(r, h)$  from Eq. (D.1). For comparison, dotted lines correspond to  $F(r, r_s)$ , where  $r_s = (6/\pi)^{1/3} h$  is the radius of a sphere with the same volume as a cell of size  $2h$ . Dashed lines illustrate the periodic case  $\tilde{F}(r, h)$ .

scheme, when the point mass is positioned at a cell centre, only one cell exhibits a non-zero density with a value of  $\rho_0$ . However, for a point mass located at a cell edge, eight cells have non-zero densities, each with a value of  $\rho_0/8$  for CIC and TSC. The sole distinction in this configuration arises from the inverse interpolation from the grid back to the particles. In contrast, for the TSC scheme with the point mass placed at the cell centre, 27 cells exhibit non-zero densities. In terms of solver performance, the results from the FFT\_7pt and multigrid solvers are consistent and identical, as expected. While the accelerations align at larger distances, the multigrid and FFT\_7pt solvers exhibit a slight overestimation of acceleration at small scales. Put differently, these solvers generate a deeper potential compared to the FFT solver. This increased acceleration results in higher particle velocities, which in turn suppresses clustering at small scales. This behaviour is consistent with the observed damping of the matter power spectrum, as shown in Fig. 6.

## Appendix E: Data locality

Finally, we highlight the importance of data locality in particle-mesh interactions, focusing on the impact it has on runtime performance. Fig. E.1 illustrates the runtime for the inverse TSC algorithm under different particle array orderings. When the particle array is randomly sorted, the inverse TSC takes approximately 0.45 seconds for the first 100 steps, with a gradual decrease down to  $\sim 0.37$  seconds. In a randomly sorted array, consecutive particles are unlikely to be spatially close, leading to



**Fig. E.1.** Runtime of the inverse TSC algorithm (from mesh to particles) for a force field. In coloured lines we show the results with different ordering of the particle arrays. ‘Randomly sorted’ means that in the initial conditions we randomly shuffle the particle positions. ‘Initially sorted only’ means that in the initial conditions the particles are linearly ordered, with the same contiguous properties as the cells they are in. ‘Partially’ and ‘Fully sorted’ means than every  $N$  steps (here 50), we sort the particle array according to Morton indexing based on their positions (see Section 3.2). The simulation has  $512^3$  particles and as many cells, and was run with 64 CPUs.

a higher number of cache misses because grid elements are also distant in memory. However, as particles cluster at later times, even a randomly sorted array experiences improved performance due to the higher likelihood that two neighbouring particles are closer spatially, reducing cache misses.

In contrast, if the particle array is initially sorted, the runtime for inverse TSC starts much lower at  $\sim 0.04$  seconds and increases up to  $\sim 0.11$  seconds over time. This improvement comes from neighbouring particles in the array being spatially closer at the beginning, allowing the cache to be used more efficiently. The ten-fold improvement in performance, achieved solely by optimising data locality, is a significant gain without requiring complex programming techniques. Over time, as particles move, the initial sorting deteriorates, leading to less efficient cache usage, even though clustering increases.

To maintain performance, PySCo offers the ability to re-sort the particles every  $N$  steps (default  $N = 50$ , adjustable by the user). With periodic sorting, the ‘fully sorted’ runtime remains flat as a function of time. However, the sorting operation can become a bottleneck, as Numba does not support parallel sorting algorithms, and resorting to sequential sorting is prohibitive for large particle counts. To address this, PySCo implements a parallelised, partial sorting method. The particle array is divided into chunks, with  $N_{\text{chunks}} = N_{\text{CPUs}}$ , where each chunk is sorted independently by each CPU. The results of this partial sorting are comparable to those of fully sorted arrays, with a slight runtime increase at late times due to the partial nature of the sorting. This increase is marginal, even when using up to 64 CPUs. Therefore, PySCo uses this partial sorting method by default as a balance between performance and computational overhead, especially for large simulations.