



**HAL**  
open science

## Perfectly parallel cosmological simulations using spatial comoving Lagrangian acceleration

F. Leclercq, B. Faure, G. Lavaux, B. D. Wandelt, A.H. Jaffe, A.F. Heavens, W.J. Percival, Camille Noûs

► **To cite this version:**

F. Leclercq, B. Faure, G. Lavaux, B. D. Wandelt, A.H. Jaffe, et al.. Perfectly parallel cosmological simulations using spatial comoving Lagrangian acceleration. *Astronomy and Astrophysics - A&A*, 2020, 639, pp.A91. 10.1051/0004-6361/202037995 . hal-02536126v2

**HAL Id: hal-02536126**

**<https://hal.science/hal-02536126v2>**

Submitted on 5 May 2024

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

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

# Perfectly parallel cosmological simulations using spatial comoving Lagrangian acceleration

Florent Leclercq,<sup>1, a)</sup> Baptiste Faure,<sup>2</sup> Guilhem Lavaux,<sup>3</sup> Benjamin D. Wandelt,<sup>3, 4, 5</sup> Andrew H. Jaffe,<sup>1</sup> Alan F. Heavens,<sup>1</sup> Will J. Percival,<sup>6, 7, 8</sup> and Camille Noi s<sup>9</sup>

<sup>1)</sup>*Imperial Centre for Inference and Cosmology (ICIC) & Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, United Kingdom*

<sup>2)</sup>*AIM, CEA, CNRS, Universit  Paris-Saclay, Universit  Paris Diderot, Sorbonne Paris Cit , 91191 Gif-sur-Yvette, France*

<sup>3)</sup>*CNRS, Sorbonne Universit , UMR 7095, Institut d'Astrophysique de Paris, 98 bis bd Arago, 75014 Paris, France*

<sup>4)</sup>*Sorbonne Universit , Institut Lagrange de Paris (ILP), 98 bis bd Arago, 75014 Paris, France*

<sup>5)</sup>*Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, 10010, New York, NY, USA*

<sup>6)</sup>*Waterloo Centre for Astrophysics, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*

<sup>7)</sup>*Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*

<sup>8)</sup>*Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada*

<sup>9)</sup>*Cogitamus Laboratory*

(Dated: 19 September 2022)

Existing cosmological simulation methods lack a high degree of parallelism due to the long-range nature of the gravitational force, which limits the size of simulations that can be run at high resolution. To solve this problem, we propose a new, perfectly parallel approach to simulate cosmic structure formation, which is based on the spatial COMoving Lagrangian Acceleration (sCOLA) framework. Building upon a hybrid analytical and numerical description of particles' trajectories, our algorithm allows for an efficient tiling of a cosmological volume, where the dynamics within each tile is computed independently. As a consequence, the degree of parallelism is equal to the number of tiles. We optimised the accuracy of sCOLA through the use of a buffer region around tiles and of appropriate Dirichlet boundary conditions around sCOLA boxes. As a result, we show that cosmological simulations at the degree of accuracy required for the analysis of the next generation of surveys can be run in drastically reduced wall-clock times and with very low memory requirements. The perfect scalability of our algorithm unlocks profoundly new possibilities for computing larger cosmological simulations at high resolution, taking advantage of a variety of hardware architectures.

## I. INTRODUCTION

We live in the age of large astronomical surveys. These surveys detect and record tracers of cosmic structure across vast volumes of the Universe, using electromagnetic and gravitational waves. A non-exhaustive list includes optical and infrared imaging and spectroscopic surveys such as LSST ([LSST Science Collaboration, 2012](#)), Euclid ([Laureijs et al., 2011](#)), DESI ([DESI Collaboration, 2016](#)), and SPHEREx ([SPHEREx Science Team, 2018](#)); catalogues and intensity maps from large radio surveys such as the square kilometer array ([Square Kilometre Array Cosmology Science Working Group, 2018](#)) and its precursors; cluster catalogues from high-resolution observations of the microwave sky (Advanced ACTPol, [Simon et al., 2018](#); SPTPol, [Austermann et al., 2012](#); Simons Observatory, [Simons Observa-](#)

[tory Collaboration, 2019](#), and CMB-S4); X-ray surveys such as the eROSITA mission ([Merloni et al., 2012](#)); as well as gravitational wave sirens across cosmological volumes with successive updates of (Advanced) LIGO ([LIGO Scientific Collaboration, 2015](#)), Virgo ([The Virgo Collaboration, 2020](#)) and LISA ([Barausse et al., 2020](#)). Whilst these data sets will be prodigious sources of scientific discovery across astrophysics, their enormous volume and dense sampling of cosmic structure will make them uniquely powerful when studying some of the deepest scientific mysteries of our time: the statistical properties of the primordial perturbations, the nature of dark matter, and the physical properties of dark energy. Indeed many of these surveys were conceived to address these questions.

Accomplishing this promise requires the ability to model these surveys in sufficient detail and with sufficient accuracy. All but the most simplistic models require the production of cosmological light-cone simulations. In particular, cosmological inferences often rely on large numbers of mock catalogues, which are used to

<sup>a)</sup>Electronic mail: [florent.leclercq@polytechnique.org](mailto:florent.leclercq@polytechnique.org); <http://www.florent-leclercq.eu/>

construct unbiased estimators and study their statistical properties, such as covariance matrices. As surveys are getting deeper, these mock catalogues now need to represent a sizeable portion of the observable Universe, up to a redshift of  $\sim 2 - 3$  (e.g.  $z = 2.3$  for the Euclid Flagship simulation<sup>1</sup>). Unfortunately, cosmological simulations put a heavy load on supercomputers. Even if only dark matter is included and resolution is minimised, they can require millions of CPU hours and hundreds of terabytes of disk space to solve the gravitational evolution of billions of particles and store the corresponding data. For instance, the DEUS-FUR simulation (Alimi *et al.*, 2012), containing  $8192^3$  particles in a box of  $21 \text{ Gpc}/h$  side length, required 10 million hours of CPU time and 300 TB of storage.

While computational needs are soaring, the performance of individual compute cores attained a plateau around 2015. Traditional hardware architectures are reaching their physical limit. Therefore, cosmological simulations cannot merely rely on processors becoming faster to reduce the computational time. Current hardware development focuses on increasing power efficiency<sup>2</sup> and solving problems of heat dissipation to allow packing a larger number of cores into each CPU. As a consequence, the performance gains of the world’s top supercomputers are the result of a massive increase in the number of parallel cores, currently<sup>3</sup> to  $\mathcal{O}(10^5)$ , and soon to  $\mathcal{O}(10^{6-7})$  in systems that are currently being built.<sup>4</sup> Hybrid architectures, where CPUs work alongside GPUs and/or reconfigurable chips such as FPGAs, add to the massive parallelism. In the exa-scale world, raw compute cycles are no longer the scarce resource. The challenge is to access the available computational power when Amdahl’s law demonstrates that communication latencies kill the potential gains due to parallelisation (Amdahl, 1967).

A way to embed high-resolution simulation of objects such as galaxy clusters, or even galaxies, in a cosmological context is through the use of varying particle mass resolution and the Adaptive Mesh Refinement technique (AMR, Berger & Colella, 1989). AMR is widely employed in grid-based simulation codes such as RAMSES (Teyssier, 2002), ENZO (Bryan *et al.*, 2014), FLASH (Fryxell *et al.*, 2000), and AMIGA (Knebe & Doumler, 2010). It is also used in MUSIC (Hahn & Abel, 2011) to generate zoom-in initial conditions for simulations. The AMR technique, which uses multi-grid relaxation methods (e.g. Guillet & Teyssier, 2011), allows focusing the effort on a specific region of the computational do-

main, but requires a two-way flow of information between small and large scales. More recently, leading computational cosmology groups have been developing sophisticated schemes to leverage parallel and hybrid computing architectures (Gonnet *et al.*, 2013; Theuns *et al.*, 2015; Aubert, Deparis & Ocvirk, 2015; Ocvirk *et al.*, 2016; Potter, Stadel & Teyssier, 2017; Yu, Pen & Wang, 2018; Garrison, Eisenstein & Pinto, 2019; Cheng *et al.*, 2020).

Full simulations of large cosmological volumes, even limited to cold dark matter and at coarse resolution, involve multiple challenges. One of the main issues preventing their easy parallelisation is the long-range nature of gravitational interactions, which forestalls high-resolution, large-volume cosmological simulations. As a response, much of the classical work in numerical cosmology focused on computational algorithms (tree codes, fast multipole methods, particle-mesh methods, and hybrids such as particle-particle-particle-mesh and tree-particle-mesh) that reduced the need for  $\mathcal{O}(N^2)$  all-to-all communications between  $N$  particles across the full computational volume.

While these algorithms are and remain the backbone of computational cosmology, they fail to fully exploit the physical scale hierarchy of cosmological perturbations. This hierarchy has first been used to push the results of  $N$ -body simulations to Universe scale for cosmic velocity fields (Strauss *et al.*, 1995). At the largest scales, the dynamics of the Universe is not complicated, and in particular, is well-captured by Lagrangian Perturbation Theory (LPT; see Bouchet *et al.*, 1995). Building upon this view, Tassev *et al.* (2015) introduced spatial COmoving Lagrangian Acceleration (sCOLA). This algorithm, using a hybrid analytical and numerical treatment of particles’ trajectories, allows one to perform simulations without the need to substantially extend the simulated volume beyond the region of interest in order to capture far-field effects, such as density fluctuations due to super-box modes. The sCOLA proof-of-concept focused on one sub-box embedded into a larger simulation box.

In this paper, we extend the sCOLA algorithm and use it within a novel method for perfectly parallel cosmological simulations. To do so, we rely on a tiling of the full cosmological volume to be simulated, where each tile is evolved independently using sCOLA. The principal challenge for the accuracy of such simulations are the boundary conditions used throughout the evolution of tiles, which can introduce artefacts. In this respect, we introduce three crucial improvements with respect to Tassev *et al.* (2015): the use of a buffer region around each tile, the use of exact boundary conditions in the calculation of LPT displacements (which has the side benefit of reducing memory requirements), and the use of a Poisson solver with Dirichlet boundary conditions meant to approximate the exact gravitational potential around sCOLA boxes. The method proposed in this work shares similar goals with zoom-in simulation techniques, the main difference residing in the change

<sup>1</sup> [https://www.euclid-ec.org/?page\\_id=4133](https://www.euclid-ec.org/?page_id=4133)

<sup>2</sup> For example, Oak-Ridge National Laboratories’ (ORNL) Summit machine has a typical power consumption of about 13 MW.

<sup>3</sup> <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>

<sup>4</sup> See for example ORNL’s next supercomputer, Frontier: <https://www.olcf.ornl.gov/wp-content/uploads/2019/05/frontier-specsheet.pdf>

of frame of reference introduced in sCOLA, which accounts for the dynamics of large scales without requiring flows of information during the evolution. On the other hand, our method is independent of the  $N$ -body integrator used to calculate the numerical part of particles' trajectories within each sCOLA box, and therefore, it cannot be related to specific approaches to do so, such as force-splitting. It is slightly approximate and more CPU-expensive than the corresponding "monolithic" simulation technique (chosen in this paper as tCOLA, [Tassev, Zaldarriaga & Eisenstein, 2013](#)), but has the essential advantage of perfect scalability. This scalability comes from the removal of any kind of communication among tiles after the initialisation of the simulation. As a consequence, for its major part, the degree of parallelism of the algorithm equals the number of tiles, which means that the workload is perfectly parallel (also called embarrassingly parallel). This property can be exploited to produce cosmological simulations in very short wall-clock times on a variety of hardware architectures, as we discuss in this paper.

After reviewing Lagrangian Perturbation Theory and its use within numerical simulations in section II, we describe our algorithm for perfectly parallel cosmological simulations in section III. In section IV, we test the accuracy and speed of the algorithm with respect to reference simulations that do not use the tiling. We discuss the implications of our results for computational strategies to model cosmic structure formation, and conclude, in section V. Details regarding the implementation are provided in the appendices.

## II. COSMOLOGICAL SIMULATIONS USING LAGRANGIAN PERTURBATION THEORY

Throughout this section we denote by  $a$  the scale factor of the Universe. For simplicity, some of the equations are abridged. We reintroduce the omitted constants, temporal prefactors, and Hubble expansion in appendix A.

Particle simulators are algorithms that compute the final position  $\mathbf{x}$  and momentum  $\mathbf{p} \equiv d\mathbf{x}/da$  of a set of particles, given some initial conditions. They can also be seen as algorithms that compute a displacement field  $\Psi$ , which maps the initial (Lagrangian) position  $\mathbf{q}$  of each particle to its final (Eulerian) position  $\mathbf{x}$ , according to the classic equation (see e.g. [Bernardeau et al., 2002](#), for a review)

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

With this point of view, the outputs are  $\mathbf{x}$  and  $\mathbf{p} = \partial\Psi/\partial a$ .

### A. Lagrangian perturbation theory (LPT)

In Lagrangian perturbation theory (LPT), the displacement field is given by an analytic equation which

is used to move particles, without the need for a numerical solver. At second order in LPT, the displacement field is written

$$\Psi_{\text{LPT}}(\mathbf{q}, a) = \Psi^{(1)}(\mathbf{q}, a) + \Psi^{(2)}(\mathbf{q}, a), \quad (2)$$

where each of the terms is separable into a temporal and a spatial contribution deriving from a Lagrangian potential:

$$\Psi^{(1)}(\mathbf{q}, a) = -D_1(a) \nabla_{\mathbf{q}} \phi^{(1)}(\mathbf{q}), \quad (3)$$

$$\Psi^{(2)}(\mathbf{q}, a) = D_2(a) \nabla_{\mathbf{q}} \phi^{(2)}(\mathbf{q}). \quad (4)$$

In equations (3) and (4),  $D_1$  and  $D_2$  are the growth factor and second-order growth factor, respectively. The Lagrangian potentials obey Poisson-like equations ([Buchert, Melott & Weiß, 1994](#)):

$$\Delta_{\mathbf{q}} \phi^{(1)}(\mathbf{q}) = \delta_i(\mathbf{q}), \quad (5)$$

$$\Delta_{\mathbf{q}} \phi^{(2)}(\mathbf{q}) = \sum_{i>j} \left[ \phi_{,ii}^{(1)} \phi_{,jj}^{(1)} - \left( \phi_{,ij}^{(1)} \right)^2 \right], \quad (6)$$

where  $\delta_i(\mathbf{q})$  is the density contrast in the initial conditions, in Lagrangian coordinates, and the  $\phi_{,ij}^{(1)}$  are spatial second derivatives of  $\phi^{(1)}$ , i.e.  $\phi_{,ij}^{(1)} \equiv \partial^2 \phi^{(1)} / \partial \mathbf{q}_i \partial \mathbf{q}_j$ .

If only the first-order term is included in equation (2), the solution is known as the Zel'dovich approximation ([Zel'dovich, 1970](#)).

### B. Temporal comoving Lagrangian acceleration (tCOLA)

In contrast to the analytical equations of LPT, particle-mesh (PM) codes (see e.g. [Klypin & Holtzman, 1997](#)) provide a fully numerical solution to the problem of large-scale structure formation. The equation of motion to be solved in a PM code reads schematically

$$\partial_a^2 \Psi(\mathbf{q}, a) = -\nabla_{\mathbf{x}} \Phi(\mathbf{x}, a), \quad (7)$$

where the gravitational potential  $\Phi$  satisfies the Poisson equation,

$$\Delta_{\mathbf{x}} \Phi(\mathbf{x}, a) = \delta(\mathbf{x}, a). \quad (8)$$

Here,  $\delta(\mathbf{x}, a)$  is the density contrast at a scale factor  $a$ , which is obtained from the set of particles' positions  $\{\mathbf{x}(a)\}$  through a density assignment operator that we denote  $B$  (typically a cloud-in-cell (CiC) scheme, see [Hockney & Eastwood, 1981](#)):

$$\delta(\mathbf{x}, a) \equiv B(\{\mathbf{x}(a)\}). \quad (9)$$

We denote by  $\bar{B}$  the corresponding interpolation operator, which is needed to obtain the accelerations of particles given the acceleration field on the grid:

$$\partial_a^2 \Psi(\{\mathbf{x}(a)\}) \equiv \bar{B}(-\nabla_{\mathbf{x}} \Phi). \quad (10)$$

The temporal COMoving Lagrangian Acceleration (tCOLA) algorithm seeks to decouple large and small scales by evolving large scales using analytic LPT results, and small scales using a numerical solver. This is achieved by splitting the Lagrangian displacement field into two contributions (Tassev & Zaldarriaga, 2012):

$$\Psi(\mathbf{q}, a) \equiv \Psi_{\text{LPT}}(\mathbf{q}, a) + \Psi_{\text{res}}(\mathbf{q}, a), \quad (11)$$

where  $\Psi_{\text{LPT}}(\mathbf{q}, a)$  is the LPT displacement field discussed in section II A and  $\Psi_{\text{res}}(\mathbf{q}, a)$  is the residual displacement of each particle, as measured in a frame comoving with an ‘‘LPT observer’’, whose trajectory is given by  $\Psi_{\text{LPT}}(\mathbf{q}, a)$ . Using equation (11), it is possible to rewrite equation (7) as

$$\partial_a^2 \Psi_{\text{res}}(\mathbf{q}, a) = -\nabla_{\mathbf{x}} \Phi(\mathbf{x}, a) - \partial_a^2 \Psi_{\text{LPT}}(\mathbf{q}, a). \quad (12)$$

The term  $\partial_a^2 \Psi_{\text{LPT}}(\mathbf{q}, a)$  can be thought of as a fictitious force acting on particles, caused by our use of a non-inertial frame of reference. Importantly, it can be computed analytically given the equations of Lagrangian perturbation theory.

The equations of motions (7) and (12) are usually integrated by the use of time-stepping techniques (see appendix B). In the limit of zero time-steps used to discretise the left-hand side of equation (12),  $\Psi_{\text{res}} = 0$  and tCOLA recovers the results of LPT; therefore, tCOLA always solves the large scales with an accuracy of at least that of LPT. In contrast, PM codes require many time-steps in equation (7) just to recover the value of the linear growth factor  $D_1$ . In the limit where the number of time-steps becomes large, tCOLA reduces to a standard PM code. In the intermediate regime (for  $\mathcal{O}(10)$  time-steps), tCOLA provides a good approximation to large-scale structure formation, at the expense of not solving the details of particle trajectories in deeply non-linear halos (see Tassev, Zaldarriaga & Eisenstein, 2013; Howlett, Manera & Percival, 2015; Leclercq, Jasche & Wandelt, 2015; Koda *et al.*, 2016; Izard, Crocce & Fosalba, 2016, for further discussion). Since by construction, tCOLA always gets the large scales correct, contrary to a PM code, the trade-off between speed and accuracy only affects small scales.

### C. Spatial comoving Lagrangian acceleration (sCOLA)

During large-scale structure formation, non-linearities appear at late times and/or at small scales. tCOLA (equation (12)) decouples LPT displacements and residual non-linear contributions ‘‘in time’’, so that, for a given accuracy, fewer time-steps are required to solve large-scale structure evolution than with a PM code. Following a similar spirit, the spatial COMoving Lagrangian Acceleration (sCOLA) framework decouples LPT displacements and residual non-linear contributions ‘‘in space’’, so that numerically evolved small scales can feel far-field effects captured analytically via LPT.

More specifically, for each particle in a volume of interest (the ‘‘sCOLA box’’) embedded in a larger cosmological volume (the ‘‘full box’’), the equation of motion of particles, which reads for a traditional  $N$ -body problem

$$\partial_a^2 \Psi(\mathbf{q}, a) = \partial_a^2 \Psi_{\text{LPT}}(\mathbf{q}, a) + \partial_a^2 \Psi_{\text{res}}(\mathbf{q}, a) = \mathbf{F}(\mathbf{x}, a) \quad (13)$$

is replaced by

$$\partial_a^2 \Psi_{\text{res}}(\mathbf{q}, a) = \mathbf{F}^{\text{sCOLA}}(\mathbf{x}, a) - \partial_a^2 \Psi_{\text{LPT}}^{\text{sCOLA}}(\mathbf{q}, a). \quad (14)$$

$\partial_a^2 \Psi_{\text{res}}(\mathbf{q}, a)$  is defined by equation (11) as the residual displacement with respect to the LPT observer of the full box, whose trajectory is given by  $\Psi_{\text{LPT}}(\mathbf{q}, a)$ . In equation (14),  $\Psi_{\text{LPT}}^{\text{sCOLA}}(\mathbf{q}, a)$  is the trajectory prescribed by solving LPT equations (see section II A) in the sCOLA box. Note that  $\Psi_{\text{LPT}}^{\text{sCOLA}}(\mathbf{q}, a)$  may differ from  $\Psi_{\text{LPT}}(\mathbf{q}, a)$ , depending on the assumptions made for the boundary conditions of the sCOLA box, discussed in section III C. Denoting by  $\mathcal{S} \subseteq \llbracket 1, N \rrbracket$  the set of particles in the sCOLA box, the gravitational force, which in equation (13) reads

$$\mathbf{F}(\mathbf{x}_i, a) \equiv \sum_{\substack{j=1 \\ j \neq i}}^N \frac{\mathbf{x}_j(a) - \mathbf{x}_i(a)}{|\mathbf{x}_j(a) - \mathbf{x}_i(a)|^3}, \quad (15)$$

is replaced by

$$\mathbf{F}^{\text{sCOLA}}(\mathbf{x}_i, a) \equiv \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} \frac{\mathbf{x}_j(a) - \mathbf{x}_i(a)}{|\mathbf{x}_j(a) - \mathbf{x}_i(a)|^3}. \quad (16)$$

It is possible to evaluate  $\mathbf{F}^{\text{sCOLA}}(\mathbf{x}, a)$ , and thus to solve equation (14), like equation (13), using any numerical gravity solver, such as particle-particle-particle-mesh, tree codes, or AMR. In this paper, we choose to focus on evaluating forces via a PM scheme. In this case, the equation of motion of particles in sCOLA reads schematically (Tassev *et al.*, 2015)

$$\partial_a^2 \Psi_{\text{res}}(\mathbf{q}, a) = -\nabla_{\mathbf{x}}^{\text{sCOLA}} \Phi^{\text{sCOLA}}(\mathbf{x}, a) - \partial_a^2 \Psi_{\text{LPT}}^{\text{sCOLA}}(\mathbf{q}, a). \quad (17)$$

The gravitational potential in the sCOLA box,  $\Phi^{\text{sCOLA}}(\mathbf{x}, a)$ , obeys the near-field version of the Poisson equation,

$$\Delta_{\mathbf{x}}^{\text{sCOLA}} \Phi^{\text{sCOLA}}(\mathbf{x}, a) = \delta^{\text{sCOLA}}(\mathbf{x}, a). \quad (18)$$

The superscript ‘‘sCOLA’’ over the gradient and Laplacian operators,  $\nabla_{\mathbf{x}}^{\text{sCOLA}}$  and  $\Delta_{\mathbf{x}}^{\text{sCOLA}}$ , mean that they are restricted to the sCOLA box (contrary to that of equations (8) and (12)). Over the density contrast  $\delta^{\text{sCOLA}}(\mathbf{x}, a)$ , the superscript means that only particles in the sCOLA box  $\{\mathbf{x}(a)\}_{\text{sCOLA}} \equiv \{\mathbf{x}_i(a)\}_{i \in \mathcal{S}}$  (instead of the full box) are used within the density assignment  $\text{B}^{\text{sCOLA}}$ , i.e.

$$\delta^{\text{sCOLA}}(\mathbf{x}, a) \equiv \text{B}^{\text{sCOLA}}(\{\mathbf{x}(a)\}_{\text{sCOLA}}). \quad (19)$$

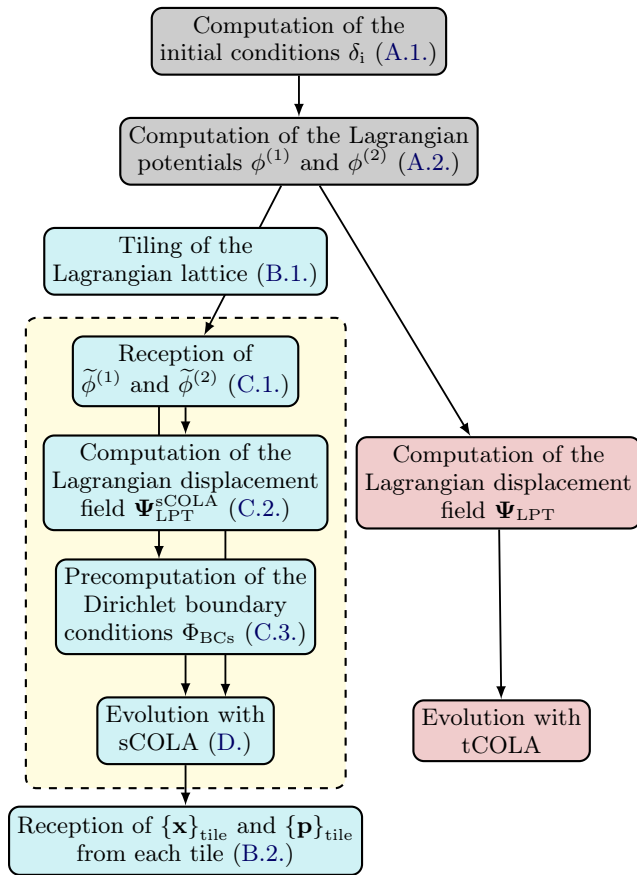


FIG. 1. Functional diagram of sCOLA (left) versus tCOLA (right). The grey boxes are common steps. sCOLA specific steps are represented in blue, and tCOLA specific steps in red. The yellow rectangle constitutes the perfectly parallel section, within which no communication is required with the master process or between processes. Arrows represent dependencies, and references to the main text are given between parentheses.

Contrary to tCOLA, which is an exact rewriting of the equations of motion of a PM code, sCOLA potentially involves approximations for the calculation of each quantity and operator with a superscript “sCOLA” instead of its full box equivalent. As a proof of concept, [Tassev et al. \(2015\)](#) showed that under certain circumstances, sCOLA provides a good approximation for the evolution of one sCOLA box embedded into a larger full box. As discussed in the introduction, we aim at generalising this result by using sCOLA within multiple sub-volumes of a full simulation box.

### III. ALGORITHM FOR PERFECTLY PARALLEL SIMULATIONS USING SCOLA

In this section, we describe an algorithm for cosmological simulations using sCOLA, for which the time evolution of independent Lagrangian sub-volumes is perfectly parallel, without any communication. A functional block

Symbol	Meaning
$N$	LPT grid size in the full box
$N_p$	Lagrangian lattice size in the full box
$N_{\text{tiles}}$	Number of tiles in each direction
$N_{p,\text{tile}}$	Number of particles per direction in each tile
$L_{\text{tile}}$	Physical size of each tile
$N_{p,\text{buffer}}$	Number of buffer particles per direction
$L_{\text{buffer}}$	Physical size of the buffer region
$N_{p,\text{sCOLA}}$	Number of particles per direction in each sCOLA box
$L_{\text{sCOLA}}$	Physical size of each sCOLA box
$N_{\text{tile}}$	LPT grid portion covering each tile
$N_{\text{sCOLA}}$	LPT grid portion covering each sCOLA box
$N_{\text{ghost}}$	Number of ghost cells depending on FDA
$N_g$	PM grid size in each sCOLA box
$r$	Over-simulation factor
$p$	Parallelisation potential factor

TABLE I. Nomenclature of symbols used in the present article.

diagram representing the main steps and their dependencies is given in figure 1. An illustration of the different grids appearing in the algorithm is presented in figure 2, and table I provides the nomenclature of some of the different variables appearing in this section.

We work in a cubic full box of side length  $L$  with periodic boundary conditions, populated by  $N_p^3$  particles initially at the nodes  $\{\mathbf{q}\}$  of a regular Lagrangian lattice. We seek to compute the set of final positions  $\{\mathbf{x}(a_f)\}$  and momenta  $\{\mathbf{p}(a_f)\}$  at final scale factor  $a_f$ . The model equations are reviewed in appendix A. The time-stepping of these equations consists of a series of “kick” and “drift” operations and is discussed in appendix B.

We approximate the Laplacians  $\Delta_{\mathbf{x}}$ ,  $\Delta_{\mathbf{q}}$  and gradient operators  $\nabla_{\mathbf{x}}$ ,  $\nabla_{\mathbf{q}}$  by finite difference approximation (FDA) at order 2, 4, or 6. The coefficients of the finite difference stencils in configuration and in Fourier space are given for example in table 1 in [Hahn & Abel \(2011\)](#). We note  $N_{\text{ghost}} = 1, 2, 3$  if FDA is taken at order 2, 4, 6, respectively.

#### A. Initial conditions and Lagrangian potentials

Before the perfectly parallel section, two initialisation steps are performed by the master process in the full box.

- A.1. The first step is to generate the initial density contrast  $\delta_i$  in the full box, on a cubic grid of  $N^3$  cells (the “LPT grid”, represented in red in the left panel of figure 2). This step can be done via the standard convolution approach (e.g. [Hockney & Eastwood, 1981](#)), given the specified initial power spectrum.
- A.2. The second step is to compute the Lagrangian potentials  $\phi^{(1)}(\mathbf{q})$  and  $\phi^{(2)}(\mathbf{q})$  on the LPT grid in the




FIG. 2. Illustration of the different grids used within sCOLA. The Lagrangian lattice is represented by dashed lines. For each tile, central particles (in black) are surrounded by buffer particles (in cyan), which are ignored at the end of the evolution. The corresponding buffer region in other grids is represented in cyan. The left panel represents the “LPT grid” on which Lagrangian potentials  $\tilde{\phi}^{(1)}$  and  $\tilde{\phi}^{(2)}$  are defined. The central region has  $N_{\text{sCOLA}}^3$  grid points (in red) and is padded by  $2N_{\text{ghost}}$  cells in each direction (pink region). The right panel shows the “PM grid” on which the density contrast  $\delta^{\text{sCOLA}}$ , the gravitational potential  $\Phi^{\text{sCOLA}}$ , and the accelerations  $-\nabla_{\mathbf{x}}^{\text{sCOLA}}\Phi^{\text{sCOLA}}$  are defined. The density contrast is defined only in the central region (which has  $N_{\text{g}}^3$  grid points, in dark green). The gravitational potential is padded by  $2N_{\text{ghost}}$  cells in each direction (light green and yellow regions), and the gridded accelerations only by  $N_{\text{ghost}}$  cells in each direction (yellow region). Solving the Poisson equation requires Dirichlet boundary conditions in six layers of  $N_{\text{ghost}}$  cells, denoted as hatched regions. For simplicity of representation, we have used here  $N_{\text{ghost}} = 1$ .

full box, which is achieved by solving equations (5) and (6).

If initial phases are generated in Fourier space, the Zel’dovich approximation (i.e. the calculation of  $\phi^{(1)}$ ) requires only one inverse fast Fourier transform (FFT) on the LPT grid. For the second-order potential, the source term on the right-hand side of equation (6) has to be computed from  $\phi^{(1)}$ ; this can either be done in Fourier space (for a cost of six inverse FFTs) or in configuration space via finite differencing (for a cost of nine one-dimensional gradient operations). In both cases, the calculation of  $\phi^{(2)}$  from its source then requires one forward and one inverse FFT.

These few FFTs in the full box are the most hardware-demanding requirement of the algorithm (particularly in terms of memory), and the only step which is not distributed and suitable for grid computing. These FFTs may however be performed on a cluster of computers with fast interconnection suitable for Message Passing Interface (Frigo & Johnson, 2005; Johnson & Frigo, 2008).

## B. Tiling and buffer region

B.1. After having computed the Lagrangian potentials, the master process splits the Lagrangian lattice (of size  $N_{\text{p}}^3$ ) into  $N_{\text{tiles}}^3$  cubic tiles (we require that  $N_{\text{p}}$  is a multiple of  $N_{\text{tiles}}$ ). Tiles are constructed to be evolved independently; therefore the main, perfectly parallel region of the algorithm starts here.

To minimise artefacts due to boundary effects (see section III D), each tile is surrounded by a “buffer region” in Lagrangian space. This buffer region consists of  $N_{\text{p,buffer}}$  particles in each direction, so that each sCOLA box contains a total of  $N_{\text{p,sCOLA}}^3$  particles, where  $N_{\text{p,sCOLA}} \equiv N_{\text{p,tile}} + 2N_{\text{p,buffer}}$  and  $N_{\text{p,tile}} \equiv N_{\text{p}}/N_{\text{tiles}}$ . Corresponding physical sizes are  $L_{\text{tile}} \equiv L N_{\text{p,tile}}/N_{\text{p}}$ ,  $L_{\text{buffer}} \equiv L N_{\text{p,buffer}}/N_{\text{p}}$ , and  $L_{\text{sCOLA}} \equiv L N_{\text{p,sCOLA}}/N_{\text{p}}$ . The fraction of the full Lagrangian lattice assigned to one child sCOLA process is represented by dotted lines in figure 2. Particles of the tile are represented in black, and particles of the buffer region are represented in cyan.

The sCOLA box is chosen to encompass the tile and its buffer region. We define the over-simulation factor  $r$  as the ratio between the total volume simulated in all sCOLA boxes and the target simulation volume, i.e.

$$\begin{aligned} r &\equiv \frac{N_{\text{tiles}}^3 N_{\text{p,sCOLA}}^3}{N_{\text{p}}^3} = \frac{N_{\text{tiles}}^3 (N_{\text{p,tile}} + 2N_{\text{p,buffer}})^3}{N_{\text{p}}^3} \\ &= \frac{N_{\text{tiles}}^3 L_{\text{sCOLA}}^3}{L^3} = \frac{N_{\text{tiles}}^3 (L_{\text{tile}} + 2L_{\text{buffer}})^3}{L^3}. \end{aligned} \quad (20)$$

Since all sCOLA boxes can be evolved independently, the degree of parallelism of the algorithm is equal to the number of sCOLA boxes,  $N_{\text{tiles}}^3$ . We call the “parallelisation potential factor” the quantity  $p \equiv N_{\text{tiles}}^3/r$ , which balances the degree of parallelism with the amount of over-simulation. It is also

$$p = \frac{N_{\text{p}}^3}{N_{\text{p,sCOLA}}^3} = \frac{L^3}{L_{\text{sCOLA}}^3}. \quad (21)$$

For each sCOLA box, the corresponding child process computes the set of final positions  $\{\mathbf{x}\}_{\text{sCOLA}}$  and momenta  $\{\mathbf{p}\}_{\text{sCOLA}}$ .

B.2. At the end of the evolution, each child process sends the set of final positions  $\{\mathbf{x}\}_{\text{tile}}$  and momenta  $\{\mathbf{p}\}_{\text{tile}}$  of particles of the tile back to the master process. Particles of the buffer region are ignored. The master process then “untiles” the simulation by gathering the results from all the tiles.

## C. Initial operations in the sCOLA boxes

A few steps are required in each sCOLA box before starting the evolution per se.

C.1. The sCOLA box receives the relevant portion of  $\phi^{(1)}(\mathbf{q})$  and  $\phi^{(2)}(\mathbf{q})$  from the master process. This is the only communication required with the master process before sending back the results at the end of the evolution.

The portion of the LPT grid received by each process from the master process corresponds to the full spatial region covered by the sCOLA box, plus an additional padding of  $2N_{\text{ghost}}$  cells in each direction. We denote by  $\tilde{\phi}^{(1)}(\mathbf{q})$  and  $\tilde{\phi}^{(2)}(\mathbf{q})$  the parts of  $\phi^{(1)}(\mathbf{q})$  and  $\phi^{(2)}(\mathbf{q})$  received from the master process (we avoid the superscript “sCOLA” since no approximation is involved at this stage). They are defined on a grid of size  $(N_{\text{sCOLA}} + 4N_{\text{ghost}})^3$ , where

$$N_{\text{tile}} \equiv \left\lceil N_{\text{p,tile}} \frac{N}{N_{\text{p}}} \right\rceil, \quad N_{\text{sCOLA}} \equiv N_{\text{tile}} + 2 \left\lceil N_{\text{p,buffer}} \frac{N}{N_{\text{p}}} \right\rceil \quad (22)$$

( $\lceil \cdot \rceil$  denotes the ceiling function). An illustration is provided in figure 2, left panel. There, the portion of the LPT grid corresponding to the sCOLA box, of size  $N_{\text{sCOLA}}$  in each direction, is represented in red and the padding region, of size  $2N_{\text{ghost}}$  in each direction, is represented in pink.

C.2. The sCOLA process locally computes the required time-independent LPT vectors  $\Psi_1^{\text{sCOLA}}$  and  $\Psi_2^{\text{sCOLA}}$  via finite differencing in configuration space and interpolation to particles’ positions.

The ghost cells included around  $\tilde{\phi}^{(1)}(\mathbf{q})$  and  $\tilde{\phi}^{(2)}(\mathbf{q})$  in the sCOLA box ensure that the proper boundary conditions are used when applying the gradient operator  $\nabla_{\mathbf{q}}^{\text{sCOLA}}$  in configuration space to get the LPT displacements on the grid. This step “consumes”  $N_{\text{ghost}}$  layers of ghost cells in each direction, so that the grid of LPT displacements has a size of  $(N_{\text{sCOLA}} + 2N_{\text{ghost}})^3$ . To use again the proper boundary conditions when going from the LPT grid to particles’ positions, another  $N_{\text{ghost}}$  layers of ghost cells is consumed by the interpolation operator  $\bar{\mathbf{B}}^{\text{sCOLA}}$ . The use of the exact boundary conditions at each of these two steps ensures that  $\nabla_{\mathbf{q}}^{\text{sCOLA}} = \nabla_{\mathbf{q}}$  and  $\bar{\mathbf{B}}^{\text{sCOLA}} = \bar{\mathbf{B}}$ . Therefore, by construction,  $\Psi_1^{\text{sCOLA}} \equiv \nabla_{\mathbf{q}}^{\text{sCOLA}} \tilde{\phi}^{(1)}(\mathbf{q})$  and  $\Psi_2^{\text{sCOLA}} \equiv \nabla_{\mathbf{q}}^{\text{sCOLA}} \tilde{\phi}^{(2)}(\mathbf{q})$  in the sCOLA box are always the same as  $\Psi_1 \equiv \nabla_{\mathbf{q}} \phi^{(1)}(\mathbf{q})$  and  $\Psi_2 \equiv \nabla_{\mathbf{q}} \phi^{(2)}(\mathbf{q})$  in the full box (as would be computed by the master process). Consequently, we do not keep track of both  $\Psi_{1,2}^{\text{sCOLA}}$  and  $\Psi_{1,2}$ , contrary to [Tassev et al. \(2015\)](#). In addition to being simpler, this scheme has the practical advantage of saving six floating-point numbers per particle in memory (three in the case of the Zel’dovich approximation).

C.3. The sCOLA process precomputes the Dirichlet boundary conditions  $\Phi_{\text{BCs}}$  that will be used at each calculation of the gravitational potential during the sCOLA evolution.

For each sCOLA box, we define a particle-mesh grid of size  $N_{\text{g}}^3$  (the “PM grid”, represented in dark green in the right panel of figure 2). The PM grid defines the force resolution; it should be equal to or finer than the LPT grid ( $N_{\text{g}} \geq N_{\text{sCOLA}}$ ). Before starting the evolution with sCOLA, each process precomputes the Dirichlet boundary conditions that will be required by the Poisson solver at each value of the scale factor  $a_{\text{K}}$ . This calculation takes as input the initial gravitational potential  $\tilde{\phi}^{(1)}(\mathbf{q})$  and outputs  $\Phi_{\text{BCs}}(\mathbf{x}, a_{\text{K}})$  for each  $a_{\text{K}}$ , defined on the PM grid with a padding of  $2N_{\text{ghost}}$  cells around the sCOLA box in each direction (light green and yellow regions in figure 2, right panel). The approximation involved in this step is further discussed in section III D 2.

## D. Evolution of sCOLA boxes

Each sCOLA box is then evolved according to the scheme reviewed in section II C and appendices A and B. Two specific approximations are needed to compute the operators and quantities with a superscript “sCOLA”; we now discuss the choices that we made.

### D.1 Density assignment ( $\mathbf{B}^{\text{sCOLA}}$ )

As mentioned in section II C, only particles of the sCOLA box should contribute to  $\delta^{\text{sCOLA}}(\mathbf{x}, a)$ . For particles that are fully in the sCOLA box, density assignment can be chosen as the same operation as would be used in a PM or tCOLA code (typically, a CiC scheme). A question is what to do with particles that have (partially) left the sCOLA box during the evolution, while keeping the requirement of no communication between boxes: this constitutes the only difference between the operators  $\mathbf{B}$  and  $\mathbf{B}^{\text{sCOLA}}$ . Possible choices include artificially periodising the sCOLA box (which is clearly erroneous) or stopping particles at its boundaries (which does not conserve momentum). Both of these choices assign the entire mass carried by the set of sCOLA particles  $\mathcal{S}$  to the PM grid, but result in artefacts in the final conditions, if the buffer region is not large enough.

An alternative choice is simply to limit the (Eulerian) PM grid volume where we compute  $\delta^{\text{sCOLA}}(\mathbf{x}, a)$  to the (Lagrangian) sCOLA box, including central and buffer regions. In practice, this means ignoring the fractional particle masses that the CiC assignment would have deposited to grid points outside the sCOLA box. We have found in our tests that this choice gives the smallest artefacts of the three choices considered.<sup>5</sup> We note that (partially) erasing some particles’ mass is an approximation

<sup>5</sup> There is a certain symmetry to this choice, since particles that would have moved into the buffer region from the outside are also neglected in the force calculation, due to the lack of communication between different sCOLA boxes.



that is only used in the  $\mathbb{B}^{\text{sCOLA}}$  operator to evaluate the source term in the Poisson equation, and therefore only affects the force calculation. The number of particles, both within each sCOLA process ( $N_{\text{p,sCOLA}}^3$ ) and in the full simulation ( $N_{\text{p}}^3$ ), is left unchanged during the evolution. Therefore, mass is always conserved both within each sCOLA process and within the full volume.

## D.2 Gravitational potential ( $\Delta_{\text{x}}^{\text{sCOLA}}$ , $\nabla_{\text{x}}^{\text{sCOLA}}$ and $\bar{\mathbb{B}}^{\text{sCOLA}}$ )

**D.2.a Poisson solver ( $\Delta_{\text{x}}^{\text{sCOLA}}$ ).** To make sure that differences between  $\Phi^{\text{sCOLA}}(\mathbf{x}, a)$  and  $\Phi(\mathbf{x}, a)$  are as small as possible, we make use of a Poisson solver with Dirichlet boundary conditions, instead of assuming periodic boundary conditions. Such a Poisson solver uses discrete sine transforms (DSTs) instead of FFTs, and requires the boundary values of  $\Phi$  in six planes (west, east, south, north, bottom, top) surrounding the PM grid (see appendix C). These planes have a thickness of  $N_{\text{ghost}}$  cells (depending on the value of the FDA used to approximate the Laplacian); they are represented by hatched regions in figure 2, right panel. At each scale factor  $a_K$  when the computation of accelerations is needed, the Dirichlet boundary conditions are extracted from the precomputed  $\Phi_{\text{BCs}}(\mathbf{x}, a_K)$  (step C.3., see section III C).

Ideally,  $\Phi_{\text{BCs}}(\mathbf{x}, a_K)$  should be the exact, non-linear gravitational potential in the full volume at  $a_K$ ,  $\Phi(\mathbf{x}, a_K)$ . However, knowing this quantity would require having previously run the monolithic simulation in the full volume, which we seek to avoid. In this paper, we rely instead on the linearly-evolving potential (LEP) approximation (Brainerd, Scherrer & Villumsen, 1993; Bagla & Padmanabhan, 1994), namely

$$\Phi_{\text{BCs}}(\mathbf{x}, a_K) \approx \Phi_{\text{LEP}}(\mathbf{x}, a_K) \equiv D_1(a_K) \tilde{\phi}^{(1)}(\mathbf{x}). \quad (23)$$

The idea behind this approximation is that the gravitational potential is dominated by long-wavelength modes, and therefore it ought to obey linear perturbation theory to a better approximation than the density field.

In equation (23), we have assumed that the linear growth factor  $D_1$  is normalised to unity at the scale factor corresponding to the initial conditions. The precomputation of  $\Phi_{\text{BCs}}$  in step C.3. is therefore an interpolation from the LPT grid to the PM grid and a simple scaling with  $D_1(a_K)$ .

The output of the Poisson solver is the gravitational potential  $\Phi^{\text{sCOLA}}(\mathbf{x}, a_K)$  on the PM grid, in the interior of the sCOLA box (dark green grid points in figure 2, right panel). Consistently with the treatment above,  $\Phi^{\text{sCOLA}}(\mathbf{x}, a_K)$  is padded using the values of  $\Phi_{\text{BCs}}(\mathbf{x}, a_K)$  in  $2N_{\text{ghost}}$  cells around the PM grid, in each direction (light green and yellow regions in figure 2, right panel).

Therefore, the only difference between  $\Delta_{\text{x}}^{\text{sCOLA}}$  and  $\Delta_{\text{x}}$  resides in using the LEP instead of the true, non-linear gravitational potential at the boundaries of the sCOLA box.

**D.2.b Accelerations ( $\nabla_{\text{x}}^{\text{sCOLA}}$  and  $\bar{\mathbb{B}}^{\text{sCOLA}}$ ).** Given the gravitational potential  $\Phi^{\text{sCOLA}}(\mathbf{x}, a_K)$ , accelerations are computed by finite differencing in configuration space and interpolation to particles' positions, similarly to step C.2. (see section III C). The application of  $\nabla_{\text{x}}^{\text{sCOLA}}$  consumes  $N_{\text{ghost}}$  cells, so that accelerations are obtained on the PM grid with a padding of  $N_{\text{ghost}}$  cells (yellow region in figure 2, right panel). Interpolation from the grid to particles' position (the  $\bar{\mathbb{B}}^{\text{sCOLA}}$  operator) further consumes  $N_{\text{ghost}}$  cells.

As for the Laplacian, the only difference between  $\nabla_{\text{x}}^{\text{sCOLA}}$  and  $\nabla_{\text{x}}$ , and  $\bar{\mathbb{B}}^{\text{sCOLA}}$  and  $\bar{\mathbb{B}}$ , resides in using the LEP in  $\Phi^{\text{sCOLA}}(\mathbf{x}, a_K)$  instead of the true, non-linear gravitational potential at the boundaries of the sCOLA box.

## IV. ACCURACY AND SPEED

We implemented the perfectly parallel sCOLA algorithm described in section III in the SIMBELMYNĚ code (Leclercq, Jasche & Wandelt, 2015), publicly available at <https://bitbucket.org/florent-leclercq/simbelmyne/> (see also Leclercq, 2015, appendix B, for technical details on the implementation of the PM and tCOLA models in SIMBELMYNĚ). This section describes some tests of the accuracy and speed of the new sCOLA algorithm. Since our implementation, relying on evaluating forces with a PM scheme, introduces some additional approximations with respect to tCOLA, we compare our results to that of corresponding monolithic tCOLA simulations. The accuracy of tCOLA with respect to more accurate gravity solvers has been characterised in the earlier literature (Tassev, Zaldarriaga & Eisenstein, 2013; Howlett, Manera & Percival, 2015; Leclercq, Jasche & Wandelt, 2015; Koda *et al.*, 2016; Izard, Crocce & Fosalba, 2016). The question of comparing the accuracy of our sCOLA algorithm to full  $N$ -body simulations would require building in a full  $N$ -body integrator for the sCOLA boxes (see equations (14) and (16)); this subject is left for future research.

Throughout the paper, we adopt the  $\Lambda$ CDM model with Planck 2015 cosmological parameters:  $h = 0.6774$ ,  $\Omega_{\Lambda} = 0.6911$ ,  $\Omega_{\text{b}} = 0.0486$ ,  $\Omega_{\text{m}} = 0.3089$ ,  $n_{\text{S}} = 0.9667$ ,  $\sigma_8 = 0.8159$  (Planck Collaboration, 2016, page 31, table 4, last column). The initial power spectrum is computed using the Eisenstein & Hu (1998, 1999) fitting function.

We base our first tests on a periodic box of comoving side length  $L = 200 \text{ Mpc}/h$  populated with  $N_{\text{p}}^3 = 512^3$  dark matter particles. For all operators, we use FDA at order 2. The LPT grid has  $N^3 = 256^3$  voxels. Particles are evolved to redshift  $z = 19$  using 2LPT. For all runs, we use 10 time-steps linearly-spaced in the scale factor to evolve particles from  $z = 19$  ( $a_i = 0.05$ ) to

$L$ [Mpc/h]	$N_p$	$N$	$N_{\text{tiles}}$	$N_{p,\text{tile}}$	$L_{\text{tile}}$ [Mpc/h]	$N_{p,\text{buffer}}$	$L_{\text{buffer}}$ [Mpc/h]	$N_g$	$r$	$p$
200	512	256	16	32	12.5	32	12.5	97	27	151.70
			8	64	25	32	12.5	129	8	64
			8	64	25	64	25	193	27	18.96
			4	128	50	32	12.5	193	3.38	18.96
			4	128	50	64	25	257	8	8
			4	128	50	128	50	385	27	2.37
			2	256	100	32	12.5	321	1.95	4.10
			2	256	100	64	25	385	3.38	2.37
1000	1024	512	16	64	62.5	14	13.7	93	2.97	1378.91
			16	64	62.5	26	25.4	117	5.95	687.90
			16	64	62.5	40	39.1	145	11.39	359.59
			16	64	62.5	64	62.5	193	27	151.70
			8	128	125	10	9.8	149	1.55	331.22
			8	128	125	20	19.5	169	2.26	226.45
			8	128	125	30	29.3	189	3.17	161.59
			8	128	125	50	48.8	229	5.65	90.59

TABLE II. Different setups used to test the accuracy and speed of our sCOLA algorithm.

$z = 0$  ( $a_f = 1$ ) (see appendix B).<sup>6</sup> For tCOLA, the PM grid, covering the full box, has  $512^3$  voxels. For sCOLA, we use eight different setups, with various parameters  $\{N_{\text{tiles}}, N_{p,\text{tile}}, L_{\text{tile}}, N_{p,\text{buffer}}, L_{\text{buffer}}, N_g, r, p\}$  given in the first part of table II.

To assess more extensively the impact of using sCOLA on large scales, we used a second ensemble of simulations with the following differences: a box with comoving side length of  $L = 1$  Gpc/h,  $N_p = 1024^3$  particles, a LPT grid with  $N^3 = 512^3$  voxels, and a PM grid of  $1024^3$  voxels for tCOLA. For sCOLA, we use eight different setups given in the second part of table II.

### A. Qualitative assessments

The redshift-zero density field is estimated by assigning all particles to the LPT grid using the CiC scheme. Results for the 200 Mpc/h box are shown in figure 3. There, the bottom right panel shows the reference tCOLA density field and other panels show the differences between sCOLA and tCOLA results, for the eight different setups. Some qualitative observations can be made: when artefacts are visible in the sCOLA results, they mainly affect over-dense regions of the cosmic web (filaments and halos), whereas under-dense regions are generally better recovered. Artefacts are of two types: the position of a structure (usually a filament) can be imprecise due to a misestimation of bulk motions (this is visible as a “dipole” in figure 3); or the density (usually of halos) can be over- or under-estimated (this is visible as a

“monopole” in figure 3). In all setups, artefacts are predominantly located close to the boundaries of tiles (represented as dashed lines) and are less visible in the centre of tiles. This can be easily understood given that the approximations made all concern the behaviour at the boundaries of sCOLA boxes. At fixed size for the buffer region, the correspondence between sCOLA and tCOLA density fields improves with increasing tile size. A minimum tile size of about 50 Mpc/h seems necessary to limit the misestimation of halo densities (“monopoles” in figure 3). At low redshift, this scale is in the mildly non-linear regime, where LPT starts to break down; therefore, the LPT frame is inaccurate for particles, and the requirement of no communication between tiles leads to mispredicted clustering. As expected, at fixed tile size, the results are improved by increasing the buffer region around tiles: in each sCOLA box, boundary approximations are pushed farther away from the central region of interest. A good compromise between reducing artefacts and increasing the size of buffer regions seems to be found for a buffer region of 25 Mpc/h, which corresponds roughly to the maximum distance travelled by a particle from its initial to its final position. In particular, the setup  $L_{\text{tile}} = 50$  Mpc/h,  $L_{\text{buffer}} = 25$  Mpc/h leads to a satisfactory approximation of the tCOLA density with a parallelisation potential factor  $p = 8$ .

In a similar fashion, the velocity field is estimated on the LPT grid from particle information, using the simplex-in-cell estimator (Hahn, Angulo & Abel, 2015; Leclercq *et al.*, 2017). Using phase-space information, this estimator accurately captures the velocity field, even in regions sparsely sampled by simulation particles. Results for the 200 Mpc/h box are shown in figure 4, where one component of the tCOLA velocity field  $v_{\text{tCOLA}}$  (in km/s) is shown in the bottom right panel. Other panels

<sup>6</sup> This means that in the case of our new sCOLA algorithm, we use COLA both “in space and time” (see Tashev *et al.*, 2015).




FIG. 3. Qualitative assessment of the redshift-zero density field from sCOLA for different tilings and buffer sizes, with respect to tCOLA. The bottom right panel shows the reference tCOLA density field in a 200 Mpc/h box with periodic boundary conditions (the quantity represented is  $\ln(2 + \delta_{\text{tCOLA}})$  where  $\delta_{\text{tCOLA}}$  is the density contrast). Other panels show the difference between sCOLA and tCOLA density fields,  $\ln(2 + \delta_{\text{sCOLA}}) - \ln(2 + \delta_{\text{tCOLA}})$ , for different sizes of tile and buffer region, as indicated above the panels. The tiling is represented by dashed lines, and the central tile's buffer region is represented by solid lines. In the third dimension, the slices represented intersect the central tile at its centre. As can be observed in this figure, artefacts are predominantly located close to the boundaries of tiles; they are reduced with increasing tile size and buffer region size.

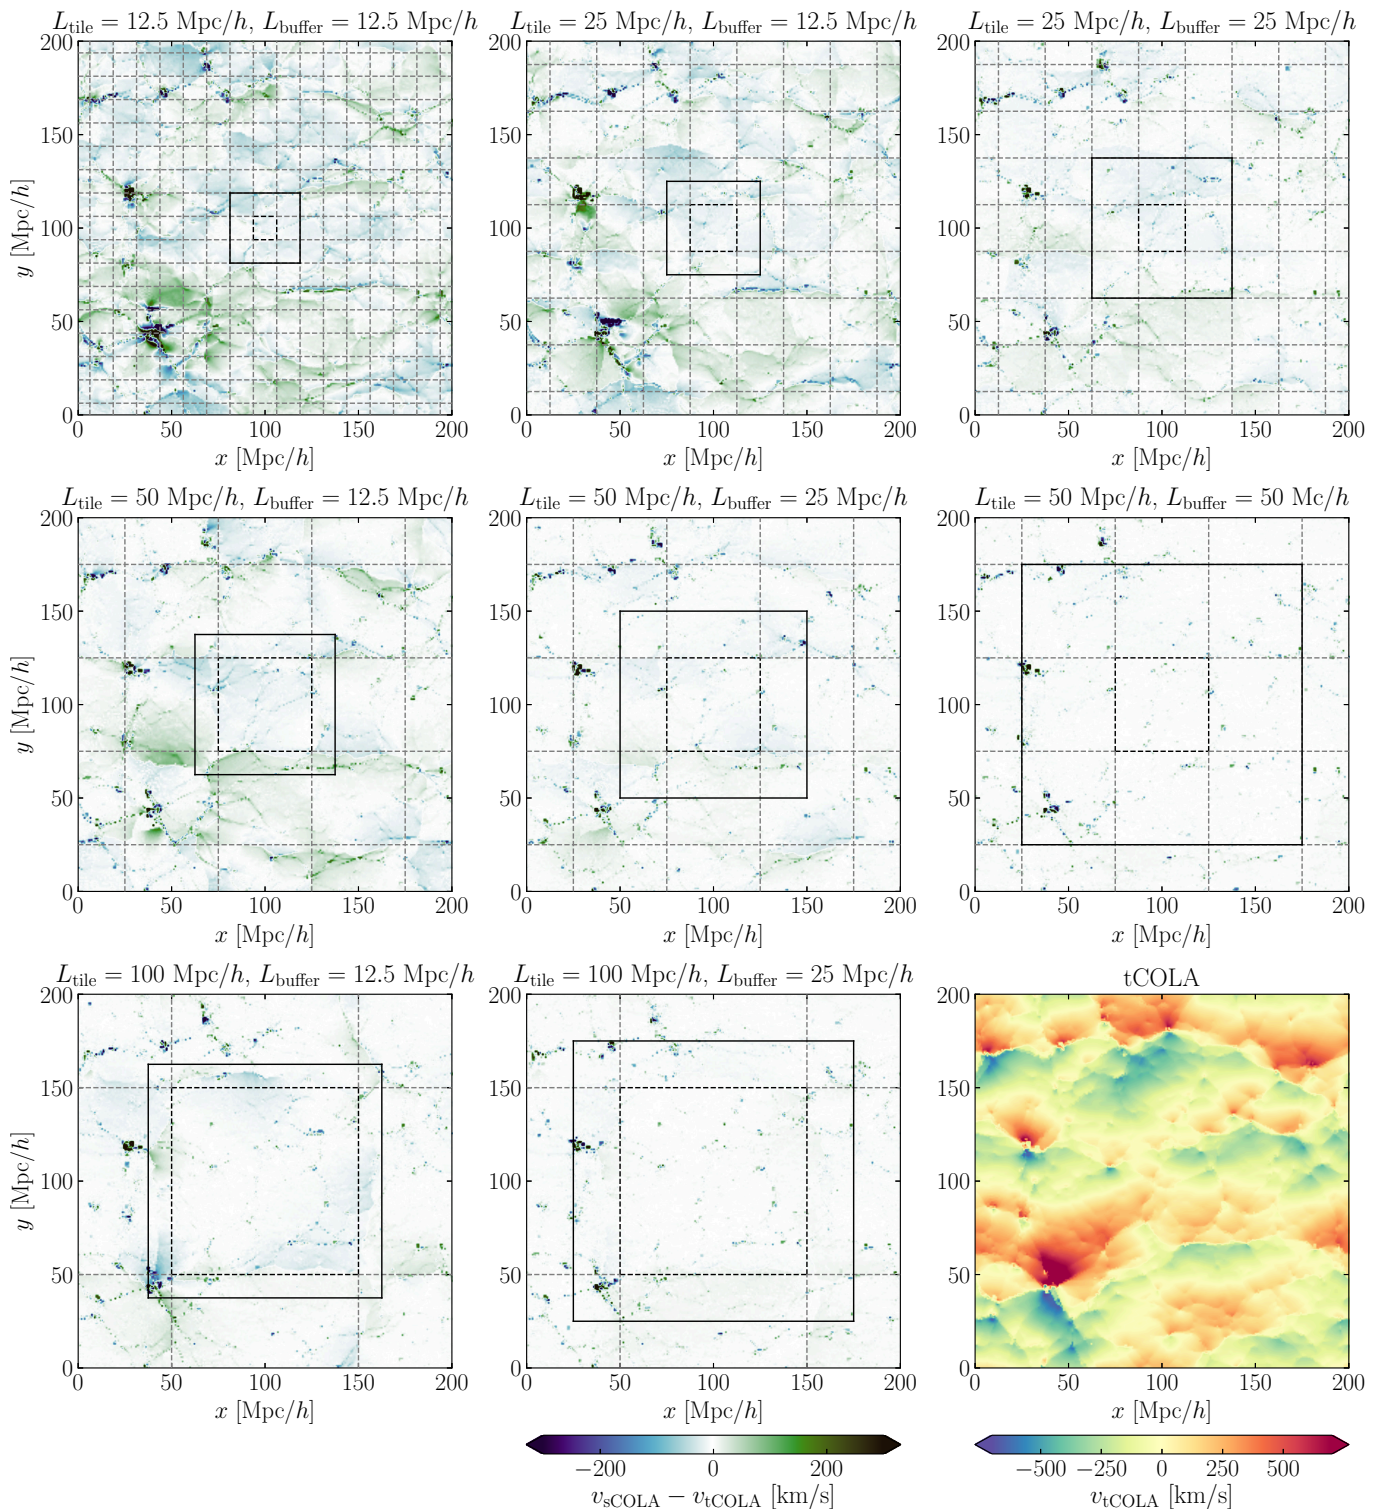


FIG. 4. Same as figure 3, but for one component of the velocity field, in km/s. Bulk flows are correctly captured if tiles and their buffer regions are large enough. Residual differences inside halos can be observed, but they are expected due to the limited number of time-steps, rendering both tCOLA and sCOLA velocities inaccurate in the deeply non-linear regime.

show the velocity error in sCOLA,  $v_{\text{sCOLA}} - v_{\text{tCOLA}}$  in km/s. Differences between tCOLA and sCOLA velocity fields are of two kinds: misestimation of bulk flows (visible as light, spatially extended regions in figure 4),

or misestimation of particle velocities inside halos (visible as dark spots in figure 4). We do not interpret the second kind of differences as errors made by our sCOLA algorithm: indeed, motions within virialised regions are




FIG. 5. Power spectrum relative to tCOLA (top panel) and cross-correlation with respect to tCOLA (bottom panel) of redshift-zero sCOLA density fields, in a 200 Mpc/h box containing  $512^3$  dark matter particles. Different sizes for the tiles (represented by different line styles) and buffer regions (represented by different colours) are used, as indicated in the legend. The vertical lines show the respective fundamental mode of different tiles, the light grey bands correspond to 3% accuracy, and the dark grey bands to 1% accuracy.

not captured accurately by any simulation using only ten time-steps, even by tCOLA in the full box. Therefore, only the first kind of differences, that is, the misestimation of coherent bulk motions is physically interpretable. In this respect, the same behaviour as for density fields can be observed: artefacts are mostly located at the boundaries of tiles, and they are reduced with increasing tile size and buffer region size, with safe minima of  $L_{\text{tile}} \gtrsim 50$  Mpc/h and  $L_{\text{buffer}} \gtrsim 25$  Mpc/h, respectively.

## B. Summary statistics

In this section, we turn to a more quantitative assessment of our results, by checking the power spectrum of final density fields and their cross-correlation to the tCOLA density field. Even if final density fields are non-Gaussian, two-point statistics (auto- and cross-spectra) are expected to be sensitive to the approximations made in our sCOLA algorithm, which involves both local and non-local operations in configuration space.

According to [Huterer & Takada \(2005\)](#) or [Audren et al. \(2013\)](#), in the best cases, observational errors for a Euclid-like survey are typically of order 3% for  $k < 10^{-2}$  (Mpc/h) $^{-1}$ . These results do not account for any of the systematic uncertainties linked to selection




FIG. 6. Same as figure 5, but in a 1 Gpc/h box containing  $1024^3$  particles.

effects or contamination of the clustering signal by foregrounds. At smaller scales, theoretical uncertainties take over, reaching 1% and above for  $k > 10^{-1}$  (Mpc/h) $^{-1}$ . In addition, the impact of baryonic physics is still largely uncertain, some models predicting an impact of at least 10% at  $k = 1$  (Mpc/h) $^{-1}$  (e.g. [van Daalen et al., 2011](#); [Chisari et al., 2018](#); [Schneider et al., 2019](#)). Any data model involving our sCOLA algorithm will be subject to these uncertainties. For this reason, we aim for no better than 3% to 1% accuracy at all scales up to  $k = 1$  (Mpc/h) $^{-1}$ , for any two-point measurement of clustering.

More precisely, we work with  $P(k)$  and  $R(k)$ , defined for two density contrast fields  $\delta$  and  $\delta' = \delta_{\text{tCOLA}}$ , with our Fourier transform convention, by

$$\begin{aligned} \delta_{\text{D}}(\mathbf{k} - \mathbf{k}')P(k) &\equiv (2\pi)^{-3}L^6 \langle \delta^*(\mathbf{k})\delta(\mathbf{k}') \rangle, \\ \delta_{\text{D}}(\mathbf{k} - \mathbf{k}')R(k) &\equiv \frac{\langle \delta^*(\mathbf{k})\delta'(\mathbf{k}') \rangle}{\sqrt{\langle \delta^*(\mathbf{k})\delta(\mathbf{k}') \rangle \langle \delta'^*(\mathbf{k})\delta'(\mathbf{k}') \rangle}}, \end{aligned} \quad (24)$$

where  $\delta_{\text{D}}$  is a Dirac delta distribution. For the estimation of  $P(k)$  and  $R(k)$ , we use 100 logarithmically-spaced  $k$ -bins from the fundamental mode of the box  $k_{\text{min}} \equiv 2\pi/L$  to  $k = 1$  (Mpc/h) $^{-1}$ .

In figures 5 and 6, we plot the power spectrum of sCOLA density fields divided by the power spectrum of the reference tCOLA density field,  $P_{\text{sCOLA}}(k)/P_{\text{tCOLA}}(k)$  (upper panels) and the cross-correlation between sCOLA and tCOLA density fields,  $R(k)$  (bottom panels), for our 200 Mpc/h (figure 5) and 1 Gpc/h box (figure 6). The grey horizontal bands represent the target accuracies of 3% and 1%, and the vertical lines mark the fundamental modes of the tiles,  $k_{\text{tile}} \equiv 2\pi/L_{\text{tile}}$ , for the different

values of  $L_{\text{tile}}$  used.

Figure 5 quantitatively confirms the considerations of section IV A. Both the amplitudes (as probed by  $P(k)/P_{\text{tCOLA}}(k)$ ) and the phase accuracy (as probed by  $R(k)$ ) of sCOLA simulations are improved with increasing tile size, for a fixed buffer region (different line styles, same colours). For a fixed tile size, results are also improved by increasing the size of the buffer region (same line styles, different colours). Remarkably, all setups yield perfect phase accuracy at large scales ( $R(k) = 1$  for  $k \leq 0.2$   $(\text{Mpc}/h)^{-1}$ ), even when the amplitude of corresponding modes deviates from the tCOLA result. Defects at small scales (lack of power and inaccurate phases) are only observed for the smallest tile sizes and are fixed by increasing the size of buffer region. This effect can be interpreted in Lagrangian coordinates: when the Lagrangian volume forming a halo is divided among different tiles that do not exchange particles, and if the buffer region is too small to contain the rest of the halo, the resulting structure is then split and under-clustered in Eulerian coordinates. In this respect, preferring a sCOLA box size ( $L_{\text{sCOLA}} \equiv L_{\text{tile}} + 2L_{\text{buffer}}$ ) of at least 100  $\text{Mpc}/h$  (and therefore  $L_{\text{tile}} \gtrsim 50$   $\text{Mpc}/h$ ,  $L_{\text{buffer}} \gtrsim 25$   $\text{Mpc}/h$ , in most situations) seems to be sensible. A more difficult issue is the amplitude of large-scale modes, for  $k < k_{\text{tile}}$ . These are sensitive to the tiling if buffer regions around tiles are too small. A safe requirement also seems to be  $L_{\text{buffer}} \gtrsim 25$   $\text{Mpc}/h$ . Putting everything together, in our 200  $\text{Mpc}/h$  box, three setups reach 3% accuracy in amplitude and phases at all scales:  $\{L_{\text{tile}} = 50$   $\text{Mpc}/h, L_{\text{buffer}} = 25$   $\text{Mpc}/h\}$  (discussed already in section IV A);  $\{L_{\text{tile}} = 100$   $\text{Mpc}/h, L_{\text{buffer}} = 25$   $\text{Mpc}/h\}$ ; and  $\{L_{\text{tile}} = 50$   $\text{Mpc}/h, L_{\text{buffer}} = 50$   $\text{Mpc}/h\}$ . The last-mentioned performs even better, reaching 1% accuracy at all scales, but at the price of over-simulating the volume by a larger factor.

Figure 6 shows the same diagnostics for a 1  $\text{Gpc}/h$  box, where the qualitative behaviour is the same as before. It confirms the requirement  $L_{\text{buffer}} \gtrsim 25$   $\text{Mpc}/h$  to get sufficient accuracy at high  $k$ . The question of the accuracy reached at the largest scales is then jointly sensitive to  $L_{\text{tile}}$  and  $L$ . In our tests, the setups  $\{L_{\text{tile}} = 62.5$   $\text{Mpc}/h, L_{\text{buffer}} = 39.1$   $\text{Mpc}/h\}$  and  $\{L_{\text{tile}} = 125$   $\text{Mpc}/h, L_{\text{buffer}} = 29.3$   $\text{Mpc}/h\}$  yield 3% accurate results at all scales, and the setups  $\{L_{\text{tile}} = 62.5$   $\text{Mpc}/h, L_{\text{buffer}} = 62.5$   $\text{Mpc}/h\}$  and  $\{L_{\text{tile}} = 125$   $\text{Mpc}/h, L_{\text{buffer}} = 48.8$   $\text{Mpc}/h\}$  almost reach 1%-level precision at all scales. We note that the two different boxes have different mass resolutions, which confirms that requirements for tile and buffer region sizes should be expressed in physical size.

### C. Tests of the approximations

As discussed in section III D, two approximations are introduced in our sCOLA algorithm with respect to a monolithic tCOLA approach. These concern density as-




FIG. 7. Tests of the approximations made in sCOLA for the density field and the gravitational potential. As in figure 5, the diagnostic tools are the power spectrum relative to tCOLA (top panel) and the cross-correlation with tCOLA (bottom panel). Our sCOLA algorithm uses the approximate interior density field  $\delta^{\text{sCOLA}}$  and the LEP approximation for the boundary gravitational potential (dash-dotted blue line). In other simulations, as indicated in the legend, we use the true density field  $\delta$  and/or the true gravitational potential  $\Phi$  at the boundaries. The approximation made for the density field dominates, especially at large scales.

signment in the interior of sCOLA boxes (approximation D.1.) and the gravitational potential at the boundaries of sCOLA boxes (approximation D.2.). In this section, we test the impact of these approximations on final results, using two-point statistics as diagnostic tools. For this test we use our sCOLA run with  $L = 200$   $\text{Mpc}/h$ ,  $N_p = 512^3$ , 64 tiles ( $N_{\text{tiles}} = 4$ ,  $N_{p,\text{tile}} = 128$ ) and  $N_{p,\text{buffer}} = 32$  (i.e.  $L_{\text{tile}} = 50$   $\text{Mpc}/h$ ,  $L_{\text{buffer}} = 12.5$   $\text{Mpc}/h$ ). We choose a small buffer size on purpose, to be sensitive to the approximations made.

Let us denote by  $\delta_{\text{int}}$  the density contrast in the interior of sCOLA boxes and by  $\Phi_{\text{BCs}}$  the gravitational potential at the boundaries of sCOLA boxes. As discussed in section III D, our algorithm involves an approximation regarding particles leaving the sCOLA box during the evolution, yielding  $\delta^{\text{sCOLA}}$ , and relies on the LEP approximation at the boundaries. It therefore uses

$$\delta_{\text{int}} = \delta^{\text{sCOLA}} \quad \text{and} \quad \Phi_{\text{BCs}} = \Phi_{\text{LEP}}. \quad (26)$$

Everything else being fixed, we ran three investigative

sCOLA simulations using respectively,

$$\delta_{\text{int}} = \delta \quad \text{and} \quad \Phi_{\text{BCs}} = \Phi_{\text{LEP}}, \quad (27)$$

$$\delta_{\text{int}} = \delta^{\text{sCOLA}} \quad \text{and} \quad \Phi_{\text{BCs}} = \Phi, \quad (28)$$

$$\delta_{\text{int}} = \delta \quad \text{and} \quad \Phi_{\text{BCs}} = \Phi, \quad (29)$$

where  $\delta$  is the “true” density contrast and  $\Phi$  is the “true” gravitational potential, extracted at each time-step from the corresponding tCOLA simulation.

Figure 7 shows the auto- and cross-spectra of resulting sCOLA density fields, with respect to the reference tCOLA result. The use of  $\delta_{\text{int}} = \delta$  yields by construction  $R(k) = 1$  at all scales, as can be checked from the bottom panel. The setup given by equation (29) is rid of the two approximations; it is therefore a consistency check: one should retrieve the tCOLA result if no bias is introduced by the tiling and different Poisson solver. As expected, figure 7 shows that our implementation recovers the tCOLA result at all scales, with only a small excess of power at  $k > 0.4 \text{ (Mpc/h)}^{-1}$  explained by the slightly higher force resolution of the sCOLA run with respect to tCOLA (the PM grid cell sizes are 0.3886 and 0.3906 Mpc/h, respectively).

The setups given by equations (27) and (28) allow disentangling the impact of approximations D.1. and D.2. In the standard run (equation (26)), averaging over tiles and timesteps,  $\sim 0.43\%$  of the  $512^3$  particles, all of which belonging to the buffer region, do not deposit all of their mass in the calculation of  $\delta^{\text{sCOLA}}$ , but  $\sim 76.5\%$  on average. This number only slightly increases with time (from  $\sim 0.35\%$  at  $a = 0.05$  to  $\sim 0.47\%$  at  $a = 1$ ); in other simulations, we have found that it has a stronger dependence on the mass resolution and on the surface of sCOLA boxes. Regarding the accuracy of the LEP approximation, the ratio of the power spectra of  $\Phi - \Phi_{\text{LEP}}$  and of  $\Phi$  goes to zero at early times and large scales, and stays below 12% for all scales with wavenumber  $k \leq 2\pi/L_{\text{sCOLA}}$  at  $a = 1$ . As can be observed in figure 7, although using the non-linear gravitational potential instead of the LEP improves both  $P(k)$  and  $R(k)$  for the final density field at all scales with wavenumber  $k > 7 \times 10^{-2} \text{ (Mpc/h)}^{-1}$ , it does not remove the  $\gtrsim 5\%$  bias in amplitude at the largest scales. On the contrary, using the true density contrast solves this problem and yields a 3% accurate result at all scales, which is remarkable given the small buffer size used in this case (the over-simulation factor is only  $r = 3.38$ ).

We conclude from these tests that the approximation made regarding the density field (D.1.) has more impact than the one regarding the gravitational potential (D.2.), especially on the largest modes. This result is consistent with the standard paradigm for structure formation, where the density contrast undergoes severe non-linearity at small scales and late times, while the gravitational potential evolves very little. It also suggests that future improvements of our algorithm should focus on finding a better approximation for  $\delta^{\text{sCOLA}}$ , rather than  $\Phi_{\text{BCs}}$ .

#### D. Computational cost

One of the main motivations for our perfectly parallel algorithm based on sCOLA is to be able to run very large volume simulations at reasonably high resolution. A detailed analysis of the speed and computational cost of our algorithm, as implemented in SIMBELMYNĚ, is therefore beyond the intent of this paper. However, in this section we discuss some performance considerations based on a sCOLA run with  $L = 1 \text{ Gpc/h}$ ,  $N_p = 1024^3$ , 512 tiles ( $N_{\text{tiles}} = 8$ ,  $N_{p,\text{tile}} = 128$ ),  $N_{p,\text{buffer}} = 30$  (i.e.  $L_{\text{tile}} = 125 \text{ Mpc/h}$ ,  $L_{\text{buffer}} = 29.3 \text{ Mpc/h}$ ),  $N_g = 199$ ; and the corresponding monolithic tCOLA simulation. In this case, the over-simulation factor is  $r \approx 3.17$  and the parallelisation potential factor is  $p \approx 161.59$ . To compare the theoretical parallelisation potential factor and the realised parallelisation efficiency, we use one process for tCOLA and 512 processes for sCOLA. Each process is run on a node with 32 cores using OpenMP parallelisation.

One of the main advantages of our sCOLA algorithm lies in its reduced memory consumption. In figure 8 (first row), we show the memory requirements for the calculation of LPT potentials in the full box (common for tCOLA and sCOLA), for the evolution of the full box with tCOLA, and for the evolution of each sCOLA box, all in single-precision floating-point format. LPT requires eight grids of size  $N^3$  (one for the initial conditions, one for the Zel’dovich potential, and six for the second-order term), occupying  $\sim 4.3 \text{ GB}$ . Evolution with tCOLA requires one integer and 12 floating-point numbers per particle (their identifier, their position  $\mathbf{x}$ , their momentum  $\mathbf{p}$ , and the vectors  $\Psi_1$  and  $\Psi_2$ ), plus a PM grid of  $1024^3$  voxels, for a total of  $\sim 60.1 \text{ GB}$ . Within each box, sCOLA requires the same memory per particle (but with  $N_{p,\text{sCOLA}}^3 \ll N_p^3$ ), a PM grid of size  $N_g^3$ , and some overhead for Dirichlet boundary conditions. The total is around 400 MB per sCOLA box with the setup considered here.

In the second row of figure 8, we show the overall cost of tCOLA versus sCOLA, both in terms of CPU time (middle left panel) and wall-clock time (middle right panel). The key feature of our algorithm is that, although the overall CPU time needed is unavoidably higher than with tCOLA, the wall-clock time spent can be drastically reduced. This owes to the degree of parallelism of our algorithm, which is equal to the number of sCOLA boxes. In particular, if as many processes as sCOLA boxes can be allocated (512 in this case), the overall wall-clock time is determined by the initial full box operations (common with tCOLA, see section III A), plus the cost of evolving only *one* sCOLA box (an average of 30.9 wall-clock seconds on 32 cores in this test). This is what is shown in the middle right panel of figure 8. The wall-clock time reduction factor is  $\approx 93$  for the evolution only ( $\approx 11$  when accounting for initialisation and writing outputs). Compared to the parallelisation potential factor  $p \approx 162$ , this number means that sCOLA-specific operations and




FIG. 8. Plots of memory requirements (first row) and of timings for two corresponding tCOLA and sCOLA simulations. Although the CPU time required is higher for sCOLA, the memory consumption and wall-clock time are significantly reduced with respect to tCOLA, due to the perfectly parallel nature of most computations (second row). In the middle left panel, the height of the white bar shows the hypothetical cost of running tCOLA for the same volume as simulated with sCOLA, when taking buffer regions into account. The relative contributions of different operations, as detailed in the legend, is shown in the third row. The main difference in computational cost in sCOLA with respect to tCOLA comes from the use of DSTs instead of FFTs, which makes the evaluation of the potential significantly more expensive.

the larger fractional parallelisation overhead in sCOLA boxes do not significantly hamper the perfectly parallel nature of the code.

The increased CPU time needed with sCOLA (see figure 8, middle left panel) is partly due to the necessity of over-simulating the volume of interest by a factor  $r > 1$  for accuracy. For comparison with the sCOLA CPU time, the height of the white bar shows the tCOLA CPU time multiplied by  $r$ . The rest of the difference in CPU time principally comes from the fact that simulations with our variant of sCOLA are intrinsically more expensive than with tCOLA for a periodic volume of the same size. This point is further discussed below.

In the third row of figure 8, we show the various relative contributions to CPU time and wall-clock time, both for full tCOLA/sCOLA runs and per tCOLA/sCOLA box. The generations of the initial conditions (brown, step A.1.) and writing of outputs to disk (grey) are common to tCOLA and sCOLA and have an overall fixed cost. LPT calculations in the full box (pink) consist of computing the Lagrangian potentials and the particle-based LPT displacements in tCOLA, but are limited to computing the Lagrangian potentials in the full box in the case of sCOLA (step A.2.). These full-box operations are only shown in the bars labelled “tCOLA” and “sCOLA”. Within each box, the different operations are evaluating the density field (yellow), solving the Poisson equation to get the gravitational potential (green), differentiating the gravitational potential to get the accelerations (blue), “kicking” particles (red), and “drifting” particles (purple). sCOLA further requires some specific operations within each box: communicating with the master process (steps B.1., B.2., and C.1.), calculating the particle-based LPT displacements (step C.2.), grouped in figure 8 and shown in orange; and pre-computing the Dirichlet boundary conditions with the LEP approximation (step C.3., cyan). sCOLA-specific operations do not contribute more than 10% of the CPU and wall-clock times per box.

A notable difference between evolving a given box with sCOLA or with tCOLA resides in the higher cost of evaluating the potential (green): in this case, 9% of CPU time and 13% of wall-clock time with sCOLA versus 6% of CPU time and 3% of wall-clock time with tCOLA. This effect is due to the use of DSTs, required by the Poisson solver with Dirichlet boundary conditions (see section III D and appendix C), instead of FFTs. Indeed, depending on the size of the PM grid, the evaluation of DSTs can be the computational bottleneck of our algorithm (up to 60% of overall CPU time is some of our runs), as opposed to the evaluation of the density field (e.g. via CiC) in traditional tCOLA or PM codes (37% of overall CPU time). For this reason, within each setup, we recommend performing experiments to find a PM grid size giving a good compromise between force accuracy and computational efficiency. In particular, it is strongly preferable that  $N_g + 1$  not contain large prime factors (this number appears in the basis functions of sine trans-



forms, see appendix C 2). Throughout this paper, we ensured that  $N_g + 1$  is always even, while keeping roughly the same force resolution as the corresponding tCOLA simulation. We note that our choice of  $N_g + 1 = 200$  in the present test, combined with the use of a power of two for the PM grid in the monolithic tCOLA run, favours tCOLA in the comparison of CPU times. The sCOLA CPU time shown in the middle left panel of figure 8 could be further optimised by making  $N_g + 1$  a power of two in sCOLA boxes.

## V. DISCUSSION AND CONCLUSION

### A. Discussion

The principal computational challenge of the gravitational  $N$ -body problem is the long-range nature of the gravitational force. Our sCOLA approach enables perfectly parallel computations and therefore opens up profoundly new possibilities for how to compute large-scale cosmological simulations. We discuss these, some consequences and possible future directions in the following.

**Gravity and physics models.** It is important to note that the sCOLA algorithm introduced in this work is general, and not limited to the gravity model used here: while we focused on a tCOLA particle-mesh implementation to evolve the sCOLA tiles, this choice was designed to facilitate the assessment of tiling artefacts against monolithic tCOLA runs. Nonetheless, any  $N$ -body method, such as particle-particle-particle-mesh, tree methods or AMR, could be used to evolve each tile. In particular, since the sCOLA approach separates quasi-linear and non-linear scales, there is no need to cut off the computation on small scales. In concert with the approaches discussed below, this fact can be exploited to perform very high-resolution, fully non-linear simulations in cosmological volumes. In this case, the spatial decoupling due to sCOLA would render computations possible that would otherwise be prohibitive.

Similar comments apply to including non-gravitational physics: since hydrodynamical or other non-gravitational forces are typically much more local than gravitational interactions, there are no algorithmic barriers to including them in each sCOLA tile.<sup>7</sup>

**Construction of light-cones and mock catalogues.** The decoupling of computational volumes achieved by our approach means that each sCOLA box can be run completely independently. Therefore, it is not necessary to define a common final redshift for all tiles. This means that to compute a cosmological light-cone, only a single tile (the one containing the observer) needs to be run

to redshift zero. Since the volume on the light-cone increases rapidly with redshift, the vast majority of tiles would only have to be run until they intersect the light-cone at high redshift. In monolithic  $N$ -body simulations, most of the computational time is spent at low redshift, since the local time-step of simulations decreases with the local dynamical time. Our approach would therefore greatly accelerate the time needed to complete light-cone simulations, by scheduling tiles in order of the redshift to which they should run (and therefore in reverse order of expected computational time), aiding load-balancing.

The construction of light-cones for surveys with large aspect ratios, such as pencil-beam surveys, can further benefit from sCOLA. Indeed, tiles that do not intersect the three-dimensional survey window do not need to be run at all for the construction of mock catalogues. In such a case, the algorithm will still capture the effects of large-scale transverse modes, even if the simulated volume is not substantially increased with respect to the survey volume.

**Low memory requirements.** sCOLA divides the computational volume into much smaller tiles and vastly reduces the memory footprint of each independent sCOLA tile computation, as shown in section IV D. As an example, simulating a  $(16 \text{ Gpc}/h)^3$  volume containing  $8192^3$  particles to achieve a mass resolution of  $10^{12.5} M_\odot$  requires  $\sim 19.8$  TB of RAM with a PM code and  $\sim 33.0$  TB of RAM with tCOLA. The setup  $\{L_{\text{tile}} = 62.5 \text{ Mpc}/h, L_{\text{buffer}} = 62.5 \text{ Mpc}/h\}$  would break down the problem into  $256^3$  tiles, each with  $(3 \times 32)^3$  particles and a memory footprint of  $\sim 53$  MB. This has important consequences, which we explore in the following.

The very modest memory requirement of our algorithm opens up multiple possibilities to accelerate the computation: even on traditional systems, the entire computation of each sCOLA tile would fit entirely into the L3 cache of a multi-core processor. This would cut out the slowest parts of the memory hierarchy, leading to a large potential performance boost and reducing code complexity. Even more promising, many such tiles could be evolved entirely independently on GPU accelerators, or even dedicated FPGAs, taking advantage of hybrid architectures of modern computational platforms while reducing the need to develop sophisticated code to manage task parallelism. At this scale, each tile computation would even fit comfortably on ubiquitous small computational platforms such as mobile phones.

**Grid computing.** The perfect scalability achieved by our approach means that large  $N$ -body simulations can even be run on very inexpensive, strongly asynchronous networks designed for large throughput computing. An extreme example would be participatory computing platforms such as Cosmology@Home,<sup>8</sup> where tens of thousands of users donate computational resources. The use

<sup>7</sup> A potential exception is long-range radiative transport of energetic (X-ray or gamma ray) photons, requiring a non-trivial extension of the approach.

<sup>8</sup> <https://www.cosmologyathome.org/>

of such platforms would be particularly suited to light-cone computations, as described above. Even if running the low-redshift part necessitates dedicated hardware, other workers could efficiently work independently to compute most of the volume, which lives at high-redshift. Only two communication steps are required for each tile: the LPT potentials are received at the beginning, and at the end of the computation each tile returns its final state at the redshift where it intersects the light-cone.

**Node Failures.** Robustness to node failure is an important consideration on all very large computational platforms. Even with extremely low failure probability for each node, since the number of nodes is high, the probability that some node fails during the course of a computation becomes high. After its initialisation steps (see section III A), our approach is entirely robust to such failure, since any individual tile can be recomputed after the fact on a modest system, for very little cost.

## B. Conclusion

In this paper, we introduced a perfectly parallel and easily applicable algorithm for cosmological simulations using sCOLA. Our approach is based on a tiling of the full simulation box, where each tile is run independently. By the use of buffer regions and appropriate Dirichlet

boundary conditions, we improved the accuracy of the algorithm with respect to Tassev *et al.* (2015). In particular, we showed that suitable setups can reach 3% to 1% accuracy at all the scales simulated, as required for data analysis of the next generation of large-scale structure surveys. In case studies, we tested the relative impact of the two approximations involved in our approach, for density assignment and the boundary gravitational potential. We considered the computational cost of our algorithm and demonstrated that even if the CPU time needed is unavoidably higher, the wall-clock time and memory footprint can be drastically reduced.

This study opens up a wide range of possible extensions, discussed in section V A. Benefiting from its perfect scalability, the approach could also allow for novel analyses of cosmological data from fully non-linear models previously too expensive to be tractable. It could straightforwardly be used for the construction of mock catalogues, but also within recently introduced likelihood-free inference techniques such as DELFI (Alsing, Wandelt & Feeney, 2018), BOLFI (Leclercq, 2018) and SELFIE (Leclercq *et al.*, 2019), which have a need for cheap simulator-based data models. We therefore anticipate that sCOLA will become an important tool in computational cosmology for the coming era.

Our perfectly parallel sCOLA algorithm has been implemented in the publicly available SIMBELMYNE code,<sup>9</sup> where it is included in version 0.4.0 and later.

## A. MODEL EQUATIONS

### 1. Model equations in the standard PM code

Denoting by  $a$  the scale factor of the Universe and  $\tau$  the conformal time, a PM code solves the equations of motion for the position  $\mathbf{x}$  and momentum  $\mathbf{p}$  of dark matter particles in comoving coordinates (the mass of particles  $m$  is absorbed in the definition of the momentum  $\mathbf{p}$ ):

$$\mathbf{p} = a \frac{d\mathbf{x}}{d\tau}, \quad (\text{A1})$$

$$\frac{d\mathbf{p}}{d\tau} = -a \nabla_{\mathbf{x}} \Phi(\mathbf{x}, \tau), \quad (\text{A2})$$

coupled to the Poisson equation for the gravitational potential, sourced by density fluctuations (equation (8)),

$$\Delta_{\mathbf{x}} \Phi(\mathbf{x}, \tau) = 4\pi G a^2 \bar{\rho}(\tau) \delta(\mathbf{x}, \tau), \quad (\text{A3})$$

where  $G$  is the gravitational constant and  $\bar{\rho}(\tau)$  is the mean matter density at conformal time  $\tau$ . The density contrast is defined from the local matter density  $\rho(\mathbf{x}, \tau)$  by

$$\delta(\mathbf{x}, \tau) \equiv \frac{\rho(\mathbf{x}, \tau)}{\bar{\rho}(\tau)} - 1. \quad (\text{A4})$$

For simplicity, from now on we note  $\nabla_{\mathbf{x}} = \nabla$ ,  $\Delta_{\mathbf{x}} = \Delta$  and  $\delta(\mathbf{x}, \tau) = \delta$ .

<sup>9</sup> <https://bitbucket.org/florent-leclercq/simbelmyne/>

It is convenient to choose the scale factor as time variable. Using  $\partial_\tau = a' \partial_a$  and the background evolution  $\bar{\rho}(\tau) = \rho^{(0)} a^{-3}$  (a prime denotes a differentiation with respect to  $\tau$  and the superscript (0) denotes quantities at the present time), the equations to solve are rewritten:

$$\frac{d\mathbf{x}}{da} = \frac{\mathbf{p}}{a'a}, \quad (\text{A5})$$

$$\frac{d\mathbf{p}}{da} = -\frac{a}{a'} \nabla \Phi, \quad (\text{A6})$$

$$\Delta \Phi = 4\pi G \rho^{(0)} a^{-1} \delta \equiv \frac{3}{2} \Omega_m^{(0)} \mathcal{H}^{(0)2} a^{-1} \delta. \quad (\text{A7})$$

We will use the equivalent formulation

$$\frac{d\mathbf{x}}{da} = \mathcal{D}(a) \mathbf{p} \quad \text{with} \quad \mathcal{D}(a) \equiv \frac{1}{a^2 \mathcal{H}(a)}, \quad (\text{A8})$$

$$\frac{d\mathbf{p}}{da} = \mathcal{K}(a) \nabla (\Delta^{-1} \delta) \quad \text{with} \quad \mathcal{K}(a) \equiv -\frac{3}{2} \frac{\Omega_m^{(0)} \mathcal{H}^{(0)2}}{a \mathcal{H}(a)}, \quad (\text{A9})$$

where we have combined equations (A6) and (A7), introduced the conformal Hubble factor  $\mathcal{H}(a) \equiv a'/a$ , and defined the ‘‘drift prefactor’’  $\mathcal{D}(a)$  and the ‘‘kick prefactor’’  $\mathcal{K}(a)$ .

## 2. Model equations with COLA

We now introduce the COLA scheme, following [Tassev, Zaldarriaga & Eisenstein \(2013\)](#) and [Tassev \*et al.\* \(2015\)](#). For each particle, we work in the frame comoving with its LPT observer, whose position is given by (see section II A)

$$\mathbf{x}_{\text{LPT}}(a) = \mathbf{q} - D_1(a) \Psi_1 + D_2(a) \Psi_2, \quad (\text{A10})$$

where we have introduced the time-independent vectors  $\Psi_1 \equiv \nabla_{\mathbf{q}} \phi^{(1)}$  and  $\Psi_2 \equiv \nabla_{\mathbf{q}} \phi^{(2)}$ . Noting  $\mathbf{x}(a) = \mathbf{x}_{\text{LPT}}(a) + \mathbf{x}_{\text{res}}(a)$  the final position of the same particle, we have

$$\frac{d\mathbf{x}}{da} = \frac{d\mathbf{x}_{\text{LPT}}}{da} + \frac{d\mathbf{x}_{\text{res}}}{da}; \quad \text{with} \quad \frac{d\mathbf{x}_{\text{LPT}}}{da} = -\frac{dD_1}{da} \Psi_1 + \frac{dD_2}{da} \Psi_2 \equiv \mathcal{D}(a) \mathbf{p}_{\text{LPT}}. \quad (\text{A11})$$

We also define  $\mathbf{p}_{\text{res}}$  such that  $d\mathbf{x}_{\text{res}}/da \equiv \mathcal{D}(a) \mathbf{p}_{\text{res}}$ . Then  $\mathbf{p} = \mathbf{p}_{\text{LPT}} + \mathbf{p}_{\text{res}}$  (see equation (A8)). Furthermore,

$$\frac{d\mathbf{p}_{\text{LPT}}}{da} = \frac{d}{da} \left( \frac{1}{\mathcal{D}(a)} \frac{d\mathbf{x}_{\text{LPT}}}{da} \right) \equiv -\mathcal{K}(a) \mathcal{V}[\mathbf{x}_{\text{LPT}}](a), \quad (\text{A12})$$

where the differential operator  $\mathcal{V}[\cdot](a)$  is defined by

$$\mathcal{V}[\cdot](a) \equiv -\frac{1}{\mathcal{K}(a)} \frac{d}{da} \left( \frac{1}{\mathcal{D}(a)} \frac{d\cdot}{da} \right). \quad (\text{A13})$$

With these notations, equation (A9) reads

$$\frac{d\mathbf{p}}{da} = \frac{d\mathbf{p}_{\text{LPT}}}{da} + \frac{d\mathbf{p}_{\text{res}}}{da} = -\mathcal{K}(a) \mathcal{V}[\mathbf{x}_{\text{LPT}}](a) + \frac{d\mathbf{p}_{\text{res}}}{da} = \mathcal{K}(a) \nabla (\Delta^{-1} \delta). \quad (\text{A14})$$

In COLA, the natural variables are therefore  $\mathbf{x}$  and  $\mathbf{p}_{\text{res}}$ .

As mentioned in section II A, the key point in COLA is that the fictitious LPT force acting on particles,  $\mathcal{V}[\mathbf{x}_{\text{LPT}}](a)$ , can be computed analytically. From equation (A10), it is straightforward to check that  $\mathcal{V}[\mathbf{x}_{\text{LPT}}](a) = -\mathcal{V}[D_1](a) \Psi_1 + \mathcal{V}[D_2](a) \Psi_2$ . The computation of  $\mathcal{V}[D_1](a)$  and  $\mathcal{V}[D_2](a)$  uses the differential equations followed by the linear and second-order growth factor, as well as the second Friedmann equation. The result is (see e.g. [Leclercq, 2015](#), equations (1.7), (1.96), (1.118) and appendix B)

$$\mathcal{V}[D_1](a) = D_1(a), \quad (\text{A15})$$

$$\mathcal{V}[D_2](a) = D_2(a) - D_1^2(a). \quad (\text{A16})$$

The equations of motion to solve are therefore, in tCOLA,

$$\frac{d\mathbf{x}}{da} = \mathcal{D}(a)\mathbf{p}_{\text{res}} - \frac{dD_1}{da}\Psi_1 + \frac{dD_2}{da}\Psi_2, \quad (\text{A17})$$

$$\frac{d\mathbf{p}_{\text{res}}}{da} = \mathcal{K}(a) [\nabla (\Delta^{-1}\delta) - D_1(a)\Psi_1 + (D_2(a) - D_1^2(a))\Psi_2]. \quad (\text{A18})$$

These are mathematically equivalent to the equations of motion of a PM code (equations (A8) and (A9)). In sCOLA, the “kick equation” (equation (A18)) is replaced for each particle of the sCOLA box by the approximation (Tassev *et al.*, 2015)

$$\frac{d\mathbf{p}_{\text{res}}}{da} = \mathcal{K}(a) \left[ \nabla^{\text{sCOLA}} ((\Delta^{\text{sCOLA}})^{-1}\delta^{\text{sCOLA}}) - D_1(a)\Psi_1^{\text{sCOLA}} + (D_2(a) - D_1^2(a))\Psi_2^{\text{sCOLA}} \right]. \quad (\text{A19})$$

with the notations introduced in section II C, as well as  $\Psi_1^{\text{sCOLA}} \equiv \nabla_{\mathbf{q}}^{\text{sCOLA}}\phi^{(1)}$  and  $\Psi_2^{\text{sCOLA}} \equiv \nabla_{\mathbf{q}}^{\text{sCOLA}}\phi^{(2)}$ . Importantly, the “drift equation” (equation (A17)) is not modified, since we are always, by definition, computing a residual displacement with respect to the LPT observer of the full box, whose position is given by equation (A10).

## B. STANDARD AND MODIFIED TIME-STEPPING

### 1. Time-stepping in the standard PM algorithm

In this paper, we adopt the second-order symplectic “kick-drift-kick” algorithm, also known as the leapfrog scheme (e.g Birdsall & Langdon, 1985) to integrate the equations of motion. This algorithm relies on integrating the equations on a small time-step and approximating the momenta ( $\mathbf{p}$  in the “drift equation” (A8)) and accelerations ( $\nabla(\Delta^{-1}\delta)$  in the “kick equation” (A9)) that appear in the integrands by their value at some time within the interval. This defines the Drift (D) and Kick (K) operators, which read using the standard discretisation (Quinn *et al.*, 1997):

$$D(t_i^D, t_f^D, t^K) : \mathbf{x}(t_i^D) \mapsto \mathbf{x}(t_f^D) = \mathbf{x}(t_i^D) + \alpha_{\mathbf{p}}(t_i^D, t_f^D, t^K)\mathbf{p}(t^K), \quad (\text{B1})$$

$$K(t_i^K, t_f^K, t^D) : \mathbf{p}(t_i^K) \mapsto \mathbf{p}(t_f^K) = \mathbf{p}(t_i^K) + \beta_{\delta}(t_i^K, t_f^K, t^D) [\nabla (\Delta^{-1}\delta)](t^D), \quad (\text{B2})$$

where

$$\alpha_{\mathbf{p}}(t_i^D, t_f^D, t^K) \equiv \int_{t_i^D}^{t_f^D} \mathcal{D}(\tilde{t}) d\tilde{t} = \int_{t_i^D}^{t_f^D} \frac{d\tilde{t}}{\tilde{t}^2 \mathcal{H}(\tilde{t})}, \quad \beta_{\delta}(t_i^K, t_f^K, t^D) \equiv \int_{t_i^K}^{t_f^K} \mathcal{K}(\tilde{t}) d\tilde{t} = -\frac{3}{2}\Omega_{\text{m}}^{(0)} \int_{t_i^K}^{t_f^K} \frac{d\tilde{t}}{\tilde{t} \mathcal{H}(\tilde{t})}, \quad (\text{B3})$$

and  $t$  is a function of the scale factor  $a$  (typically  $t(a) = a$  or  $t(a) = \exp(a)$  for time-steps linearly spaced or logarithmically spaced in the scale factor, respectively).

The time evolution between  $t_0 = t(a_i)$  and  $t_{n+1} = t(a_f)$  is then achieved by applying the following operator,  $E(t_{n+1}, t_0)$ , to the initial state ( $\mathbf{x}(t_0), \mathbf{p}(t_0)$ ):

$$K(t_{n+1/2}, t_{n+1}, t_{n+1})D(t_n, t_{n+1}, t_{n+1/2}) \left[ \prod_{i=0}^n K(t_{i+1/2}, t_{i+3/2}, t_{i+1})D(t_i, t_{i+1}, t_{i+1/2}) \right] K(t_0, t_{1/2}, t_0). \quad (\text{B4})$$

### 2. Time-stepping with COLA, standard discretisation

Using the standard discretisation (Quinn *et al.*, 1997) of equations (A17) and (A18), the Kick and Drift operators for tCOLA are defined by

$$\tilde{D}(t_i^D, t_f^D, t^K) : \mathbf{x}(t_i^D) \mapsto \mathbf{x}(t_f^D) = \mathbf{x}(t_i^D) + \alpha_{\mathbf{p}}(t_i^D, t_f^D, t^K)\mathbf{p}_{\text{res}}(t^K) - [D_1]_{t_i^D}^{t_f^D} \Psi_1 + [D_2]_{t_i^D}^{t_f^D} \Psi_2, \quad (\text{B5})$$

$$\tilde{K}(t_i^K, t_f^K, t^D) : \mathbf{p}_{\text{res}}(t_i^K) \mapsto \mathbf{p}_{\text{res}}(t_f^K) = \mathbf{p}_{\text{res}}(t_i^K) + \beta_{\delta}(t_i^K, t_f^K, t^D) \times \\ ([\nabla (\Delta^{-1}\delta)](t^D) - D_1(t^D)\Psi_1 + (D_2(t^D) - D_1^2(t^D))\Psi_2), \quad (\text{B6})$$

where the time factors  $\alpha_{\mathbf{p}}(t_i^D, t_f^D, t^K)$  and  $\beta_{\delta}(t_i^K, t_f^K, t^D)$  are the same as in the PM case (see equation (B3)). For sCOLA,  $\tilde{K}$  is given by the same expression (equation (B6)) but operates on quantities and differential operators superscripted “sCOLA” consistently with equation (A19).

In the initial conditions, generated with LPT, we have  $\mathbf{p} = \mathbf{p}_{\text{LPT}}$ ; therefore the momentum residual in the rest frame of LPT observers,  $\mathbf{p}_{\text{res}}$ , should be initialised to zero. At the end, the LPT momentum  $\mathbf{p}_{\text{LPT}}$  has to be added to  $\mathbf{p}_{\text{res}}$  to recover the full momentum of particles,  $\mathbf{p}$ . This corresponds respectively to the  $L_-$  and  $L_+$  operators (Tassev, Zaldarriaga & Eisenstein, 2013, appendix A), given by

$$L_{\pm}(t) : \quad \mathbf{p}(t) \mapsto \mathbf{p}(t) \pm \mathbf{p}_{\text{LPT}}(t) = \mathbf{p}(t) \pm \frac{1}{\mathcal{D}(t)} \left( -\frac{dD_1}{dt} \boldsymbol{\Psi}_1 + \frac{dD_2}{dt} \boldsymbol{\Psi}_2 \right). \quad (\text{B7})$$

In COLA, the time evolution between  $t_0 = t(a_i)$  and  $t_{n+1} = t(a_f)$  is therefore achieved by applying the following operator to the initial state  $(\mathbf{x}(t_0), \mathbf{p}(t_0))$ :

$$L_+(t_{n+1}) \tilde{\mathbb{E}}(t_{n+1}, t_0) L_-(t_0), \quad (\text{B8})$$

where  $\tilde{\mathbb{E}}(t_{n+1}, t_0)$  is the operator given by equation (B4), replacing  $D$  by  $\tilde{D}$  and  $K$  by  $\tilde{K}$ .

### 3. Time-stepping with COLA, modified discretisation

Another approach for the discretisation of equations (A17) and (A18) is proposed by Tassev, Zaldarriaga & Eisenstein (2013). For any arbitrary positive function  $u$  of  $t$ , we can rewrite

$$\frac{d\mathbf{x}}{dt} = \mathcal{D}(t)u(t) \left\{ \frac{1}{u(t)} \times \mathbf{p}_{\text{res}} \right\} - \frac{dD_1}{dt} \boldsymbol{\Psi}_1 + \frac{dD_2}{dt} \boldsymbol{\Psi}_2, \quad (\text{B9})$$

$$\frac{d\mathbf{p}_{\text{res}}}{dt} = \frac{du(t)}{dt} \left\{ \frac{\mathcal{K}(t)}{du(t)/dt} \times [\nabla(\Delta^{-1}\delta) - D_1(t)\boldsymbol{\Psi}_1 + (D_2(t) - D_1^2(t))\boldsymbol{\Psi}_2] \right\}. \quad (\text{B10})$$

This form is particularly relevant if  $\mathbf{p}_{\text{res}}$  has a time dependence which is entirely captured by a particular  $u(t)$ , which is universal for all particles. For each equation, considering that the part between curly brackets is constant during the time-step (instead of the momentum and accelerations, respectively), the modified  $\tilde{D}$  and  $\tilde{K}$  operators are given by equations (B5) and (B6) with the following modified time factors instead of  $\alpha_{\mathbf{p}}(t_i^D, t_f^D, t^K)$  and  $\beta_{\delta}(t_i^K, t_f^K, t^D)$ :

$$\tilde{\alpha}_{\mathbf{p}}(t_i^D, t_f^D, t^K) \equiv \frac{1}{u(t^K)} \int_{t_i^D}^{t_f^D} \mathcal{D}(\tilde{t})u(\tilde{t}) d\tilde{t}, \quad \tilde{\beta}_{\delta}(t_i^K, t_f^K, t^D) \equiv (u(t_f^K) - u(t_i^K)) \times \frac{\mathcal{K}(t^D)}{(du(t)/dt)(t^D)}, \quad (\text{B11})$$

where in  $\tilde{\beta}_{\delta}(t_i^K, t_f^K, t^D)$ , we have used the trivial integration  $\int_{t_i^K}^{t_f^K} \frac{du(\tilde{t})}{d\tilde{t}} d\tilde{t} = u(t_f^K) - u(t_i^K)$ .

Using the Ansatz suggested by Tassev, Zaldarriaga & Eisenstein (2013),  $u(a) = a^{n_{\text{LPT}}}$  when  $t(a) = a$ , we get the explicit expressions

$$\tilde{\alpha}_{\mathbf{p}}(a_i^D, a_f^D, a^K) = \frac{1}{(a^K)^{n_{\text{LPT}}}} \int_{a_i^D}^{a_f^D} \frac{\tilde{a}^{n_{\text{LPT}}-2}}{\mathcal{H}(\tilde{a})} d\tilde{a}, \quad \tilde{\beta}_{\delta}(a_i^K, a_f^K, a^D) = -\frac{3}{2} \Omega_{\text{m}}^{(0)} \frac{(a_f^K)^{n_{\text{LPT}}} - (a_i^K)^{n_{\text{LPT}}}}{n_{\text{LPT}}(a^D)^{n_{\text{LPT}}} \mathcal{H}(a^D)}. \quad (\text{B12})$$

We adopt this form and  $n_{\text{LPT}} = -2.5$  for both tCOLA and sCOLA operators, throughout this paper.

### C. POISSON SOLVER WITH DIRICHLET BOUNDARY CONDITIONS

In this appendix, we describe how to compute the interior gravitational potential  $\Phi$  with Dirichlet boundary conditions. The method is standard in computational physics and has been used at least since James (1977). Formally, we seek to solve the discrete Poisson equation,

$$\Delta\Phi = \delta, \quad (\text{C1})$$

subject to a known boundary potential  $\Phi_{\text{BCs}}$ , where  $\Delta$  is the FDA to the exact Laplacian operator, i.e.  $\Delta \equiv \partial_x^2 + \partial_y^2 + \partial_z^2$  where  $\partial_x^2$ ,  $\partial_y^2$ , and  $\partial_z^2$  are discrete one-dimensional second-order derivatives (see table 1 in Hahn & Abel, 2011, for their expressions in FDA at order 2, 4 and 6).

## 1. Modified density distribution

We define  $\Phi_{\text{BCs}}$  as having non-zero values only in a layer of  $N_{\text{ghost}}$  cells immediately outside the active domain of the PM grid. We can then write the desired potential as  $\Phi \equiv \tilde{\Phi} + \Phi_{\text{BCs}}$  where the required boundary condition for  $\tilde{\Phi}$  is  $\tilde{\Phi} = 0$ . From the definition of  $\Phi_{\text{BCs}}$ ,  $\Delta\Phi_{\text{BCs}}$  will be non-zero only in a layer of  $N_{\text{ghost}}$  active cells just inside the domain boundaries. We can thus define a modified density distribution,

$$\delta' \equiv \delta - \Delta\Phi_{\text{BCs}}, \quad (\text{C2})$$

which is the same as  $\delta$  everywhere except in the layer of  $N_{\text{ghost}}$  cells adjoining the domain boundaries. We can then employ a zero-boundary condition Poisson solver to obtain a solution of  $\Delta\tilde{\Phi} = \delta'$  (see section C 2). Within the interior, where  $\Phi_{\text{BCs}} = 0$ , this solution is the desired final solution of  $\Delta\Phi = \delta$  with the Dirichlet boundary condition  $\Phi = \Phi_{\text{BCs}}$ .

## 2. Zero-boundary condition Poisson solver

In cosmological simulations, it is conventional to use FFTs to solve the Poisson equation, since the discrete Laplacian operator is diagonal in Fourier space. FFTs assume that the input source has periodic boundary conditions. Similarly, for zero boundary conditions, we can work with three-dimensional type-I discrete sine transforms (DST-I), defined by

$$\delta^{\ell,m,n} \equiv \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{k=1}^{N_z} \delta_{i,j,k} \mathcal{X}_i^\ell \mathcal{Y}_j^m \mathcal{Z}_k^n, \quad (\text{C3})$$

where  $\delta_{i,j,k}$  is the value of the source field in the voxel indexed by  $1 \leq i \leq N_x$ ,  $1 \leq j \leq N_y$ ,  $1 \leq k \leq N_z$  ( $N_x = N_y = N_z = N_g$  in this paper). The basis functions are defined by

$$\mathcal{X}_i^\ell \equiv \sin\left(\frac{\pi\ell i}{N_x + 1}\right), \quad \mathcal{Y}_j^m \equiv \sin\left(\frac{\pi m j}{N_y + 1}\right), \quad \mathcal{Z}_k^n \equiv \sin\left(\frac{\pi n k}{N_z + 1}\right). \quad (\text{C4})$$

They ensure that the signal has zero boundary values (for  $i \in \{0, N_x + 1\}$  or  $j \in \{0, N_y + 1\}$  or  $k \in \{0, N_z + 1\}$ ). They satisfy discrete orthogonality relations, for example,

$$\frac{2}{N_x + 1} \sum_{i=1}^{N_x} \mathcal{X}_i^\ell \mathcal{X}_i^{\ell'} = \delta_{\ell\ell'} \quad \text{and} \quad \frac{2}{N_y + 1} \sum_{m=1}^{N_y} \mathcal{Y}_j^m \mathcal{Y}_j^{m'} = \delta_{jj'}, \quad (\text{C5})$$

where  $\delta_{\text{K}}$  is the Kronecker symbol. The inverse transformation is simply DST-I multiplied by  $8/[(N_x + 1)(N_y + 1)(N_z + 1)]$ , i.e.

$$\delta_{i,j,k} = \frac{8}{(N_x + 1)(N_y + 1)(N_z + 1)} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{k=1}^{N_z} \delta^{\ell,m,n} \mathcal{X}_i^\ell \mathcal{Y}_j^m \mathcal{Z}_k^n, \quad (\text{C6})$$

and similarly for the gravitational potential,

$$\Phi_{i,j,k} = \frac{8}{(N_x + 1)(N_y + 1)(N_z + 1)} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{k=1}^{N_z} \Phi^{\ell,m,n} \mathcal{X}_i^\ell \mathcal{Y}_j^m \mathcal{Z}_k^n. \quad (\text{C7})$$

It is straightforward to show that  $\mathcal{X}_i^\ell$ ,  $\mathcal{Y}_j^m$ ,  $\mathcal{Z}_k^n$  are eigenfunctions of the discrete one-dimensional second-order derivatives  $\partial_x^2$ ,  $\partial_y^2$ , and  $\partial_z^2$ , respectively. The corresponding eigenvalues  $\lambda_x^\ell$ ,  $\lambda_y^m$  and  $\lambda_z^n$  are given by

$$\lambda_x^\ell = -\frac{4}{d_x^2} \sin^2\left(\frac{k_\ell d_x}{4} \frac{N_x}{N_x + 1}\right), \quad (\text{C8})$$

$$\lambda_x^\ell = \frac{1}{3d_x^2} \left[ \sin^2\left(\frac{k_\ell d_x}{2} \frac{N_x}{N_x + 1}\right) - 16 \sin^2\left(\frac{k_\ell d_x}{4} \frac{N_x}{N_x + 1}\right) \right], \quad (\text{C9})$$

$$\lambda_x^\ell = -\frac{1}{45d_x^2} \left[ 2 \sin^2\left(\frac{3k_\ell d_x}{4} \frac{N_x}{N_x + 1}\right) - 27 \sin^2\left(\frac{k_\ell d_x}{2} \frac{N_x}{N_x + 1}\right) + 270 \sin^2\left(\frac{k_\ell d_x}{4} \frac{N_x}{N_x + 1}\right) \right], \quad (\text{C10})$$

for FDA at order 2, 4, and 6 respectively, where  $k_\ell \equiv 2\pi\ell/L_x$ ,  $d_x \equiv L_x/N_x$  and  $L_x$  is the size of the box along the  $x$ -direction ( $L_x = L_{\text{sCOLA}} \equiv L/N_{\text{tiles}}$  in this paper). Similar expressions exist for  $\lambda_y^m$  and  $\lambda_z^n$ .

Plugging equations (C6) and (C7) into (C1) and using the orthogonality relations, we obtain a simple form for the discretised Poisson equation in sine space,

$$\Phi^{\ell,m,n} = \frac{\delta^{\ell,m,n}}{\lambda_x^\ell + \lambda_y^m + \lambda_z^n}. \quad (\text{C11})$$

Therefore, the Poisson equation  $\Delta\Phi = \delta$  with zero boundary conditions can be solved by the following three steps:

1. performing a forward DST of the source ( $\delta_{i,j,k} \rightarrow \delta^{\ell,m,n}$ ), according to equation (C3) (costing  $\mathcal{O}(N_x N_y N_z \log [N_x N_y N_z])$  operations),
2. solving the Poisson equation in sine space ( $\delta^{\ell,m,n} \rightarrow \Phi^{\ell,m,n}$ ), according to equation (C11) (costing  $\mathcal{O}(N_x N_y N_z)$  operations),
3. performing an inverse DST of the gravitational potential ( $\Phi^{\ell,m,n} \rightarrow \Phi_{i,j,k}$ ), according to equation (C7) (costing  $\mathcal{O}(N_x N_y N_z \log [N_x N_y N_z])$  operations).

In practice, forward and inverse DSTs are performed using the FFTW library (Frigo & Johnson, 2005), publicly available at <http://www.fftw.org>, where the DST-I is known as FFTW\_RODFT00.

## ORCID IDS

Florent Leclercq  0000-0002-9339-1404  
 Baptiste Faure  0000-0002-6623-2344  
 Guilhem Lavaux  0000-0003-0143-8891  
 Benjamin D. Wandelt  0000-0002-5854-8269  
 Andrew H. Jaffe  0000-0003-2086-1759  
 Alan F. Heavens  0000-0003-1586-2773  
 Will J. Percival  0000-0002-0644-5727  
 Camille Noûs  0000-0002-0778-8115

## STATEMENT OF CONTRIBUTION

FL and BDW conceived the project. FL wrote the SIMBELMYNĚ code, implemented the new sCOLA algorithm, ran the simulations, performed the study, supervised BF’s internship project, and wrote the bulk of the paper. BF contributed to the first implementation of the tiling algorithm in SIMBELMYNĚ and to early tests of the method. GL suggested the investigation of an alternative Poisson solver, proposed tests of the impact of boundary effects, and contributed to writing the paper. BDW made conceptual contributions, helped designing the accuracy and speed tests, and contributed to writing the paper. AHJ prompted the use of the linearly-evolving potential approximation. AHJ and AFH contributed to the interpretation of results. WJP supported the design of the first version of the algorithm and contributed to student supervision. CN contributed to the collegial construction of the standards of science, by developing the methodological framework, the state-of-the-art, as well as post-publication procedures. All natural authors read and approved the final manuscript.

## ACKNOWLEDGEMENTS

We are grateful to Matías Zaldarriaga for stimulating discussions and useful comments throughout the realisation of this project. We thank Jens Jasche and Svetlin Tassev for discussions that triggered this project, and Oliver Hahn for constructive observations. FL and BDW acknowledge the hospitality of the Institute for Advanced Study, Princeton, where this project was initiated. FL, BF and WJP thank the Institute of Cosmology and Gravitation of the University of Portsmouth, where part of this work was prepared.

This work made use of NumPy (van der Walt, Colbert & Varoquaux, 2011), IPython (Perez & Granger, 2007), Matplotlib (Hunter, 2007), Jupyter notebooks (Kluyver *et al.*, 2016), and the colourmaps provided by the cmocool (Thyng *et al.*, 2016) and CMasher (<https://github.com/1313e/CMasher>) packages.

FL acknowledges funding from the Imperial College London Research Fellowship Scheme. GL and BDW acknowledge financial support from the ANR BIG4, under reference ANR-16-CE23-0002. The Center for Computational Astrophysics is supported by the Simons Foundation. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. Numerical computations were done on the Sciama High Performance Compute (HPC) cluster which is supported by the ICG, SEPNet and the University of Portsmouth; and on the cx1 cluster hosted by the Research Computing Service facilities at Imperial College London (doi:10.14469/hpc/2232). This work was done within the Aquila Consortium (<https://aquila-consortium.org>).

## REFERENCES

- (Alimi *et al.*, 2012) J.-M. Alimi, V. Bouillot, Y. Rasera, V. Reverdy, P.-S. Corasaniti, I. Balmes, S. Requena, X. Delaruelle, J.-N. Richet, *DEUS Full Observable  $\Lambda$ CDM Universe Simulation: the numerical challenge* (2012), arXiv:1206.2838 [astro-ph.CO].
- (Alsing, Wandelt & Feeney, 2018) J. Alsing, B. Wandelt, S. Feeney, *Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology*, Mon. Not. R. Astron. Soc. **477**, 2874 (2018), arXiv:1801.01497.
- (Amdahl, 1967) G. M. Amdahl, *Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities*, in *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference*, AFIPS '67 (Spring) (Association for Computing Machinery, New York, NY, USA, 1967) p. 483–485.
- (Aubert, Deparis & Ocvirk, 2015) D. Aubert, N. Deparis, P. Ocvirk, *EMMA: an adaptive mesh refinement cosmological simulation code with radiative transfer*, Mon. Not. R. Astron. Soc. **454**, 1012 (2015), arXiv:1508.07888 [astro-ph.IM].
- (Audren *et al.*, 2013) B. Audren, J. Lesgourgues, S. Bird, M. G. Haehnelt, M. Viel, *Neutrino masses and cosmological parameters from a Euclid-like survey: Markov Chain Monte Carlo forecasts including theoretical errors*, J. Cosmology Astropart. Phys. **2013**, 026 (2013), arXiv:1210.2194 [astro-ph.CO].
- (Austermann *et al.*, 2012) J. E. Austermann, K. A. Aird, J. A. Beall, D. Becker, A. Bender, B. A. Benson, L. E. Bleem, J. Britton, J. E. Carlstrom, C. L. Chang, *et al.*, *SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope*, in *Proceedings of the SPIE*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8452 (2012) p. 84521E.
- (Bagla & Padmanabhan, 1994) J. S. Bagla, T. Padmanabhan, *Nonlinear Evolution of Density Perturbations Using the Approximate Constancy of the Gravitational Potential*, Mon. Not. R. Astron. Soc. **266**, 227 (1994), arXiv:gr-qc/9304021 [gr-qc].
- (Barausse *et al.*, 2020) E. Barausse, E. Berti, T. Hertog, S. Hughes, P. Jetzer, P. Pani, T. Sotiriou, N. Tamanini, H. Wittek, K. Yagi, *et al.*, *Prospects for Fundamental Physics with LISA* (2020), arXiv:2001.09793 [gr-qc].
- (Berger & Colella, 1989) M. J. Berger, P. Colella, *Local Adaptive Mesh Refinement for Shock Hydrodynamics*, J. Comp. Phys. **82**, 64 (1989).
- (Bernardeau *et al.*, 2002) F. Bernardeau, S. Colombi, E. Gaztañaga, R. Scoccimarro, *Large-scale structure of the Universe and cosmological perturbation theory*, Phys. Rep. **367**, 1 (2002), astro-ph/0112551.
- (Birdsall & Langdon, 1985) C. K. Birdsall, A. B. Langdon, *Plasma Physics via Computer Simulation* (CRC Press, 1985).
- (Bouchet *et al.*, 1995) F. R. Bouchet, S. Colombi, E. Hivon, R. Juszkiewicz, *Perturbative Lagrangian approach to gravitational instability*, Astron. & Astrophys. **296**, 575 (1995), astro-ph/9406013.
- (Brainerd, Scherrer & Villumsen, 1993) T. G. Brainerd, R. J. Scherrer, J. V. Villumsen, *Linear Evolution of the Gravitational Potential: A New Approximation for the Nonlinear Evolution of Large-Scale Structure*, Astrophys. J. **418**, 570 (1993).
- (Bryan *et al.*, 2014) G. L. Bryan, M. L. Norman, B. W. O'Shea, T. Abel, J. H. Wise, M. J. Turk, D. R. Reynolds, D. C. Collins, P. Wang, S. W. Skillman, *et al.*, *ENZO: An Adaptive Mesh Refinement Code for Astrophysics*, Astrophys. J. Supp. **211**, 19 (2014), arXiv:1307.2265 [astro-ph.IM].
- (Buchert, Melott & Weiß, 1994) T. Buchert, A. L. Melott, A. G. Weiß, *Testing higher-order Lagrangian perturbation theory against numerical simulations I. Pancake models*, Astron. & Astrophys. **288**, 349 (1994), astro-ph/9309056.
- (Cheng *et al.*, 2020) S. Cheng, H.-R. Yu, D. Inman, Q. Liao, Q. Wu, J. Lin, *CUBE – Towards an Optimal Scaling of Cosmological N-body Simulations* (2020), arXiv:2003.03931 [physics.comp-ph].
- (Chisari *et al.*, 2018) N. E. Chisari, M. L. A. Richardson, J. Devriendt, Y. Dubois, A. Schneider, A. M. C. Le Brun, R. S. Beckmann, S. Peirani, A. Slyz, C. Pichon, *The impact of baryons on the matter power spectrum from the Horizon-AGN cosmological hydrodynamical simulation*, Mon. Not. R. Astron. Soc. **480**, 3962 (2018), arXiv:1801.08559 [astro-ph.CO].
- (DESI Collaboration, 2016) DESI Collaboration, *The DESI Experiment Part I: Science, Targeting, and Survey Design* (2016), arXiv:1611.00036 [astro-ph.IM].
- (Eisenstein & Hu, 1999) D. J. Eisenstein, W. Hu, *Power Spectra for Cold Dark Matter and Its Variants*, Astrophys. J. **511**, 5 (1999), astro-ph/9710252.
- (Eisenstein & Hu, 1998) D. J. Eisenstein, W. Hu, *Baryonic Features in the Matter Transfer Function*, Astrophys. J. **496**, 605 (1998), astro-ph/9709112.
- (Frigo & Johnson, 2005) M. Frigo, S. G. Johnson, *The Design and Implementation of FFTW3*, Proceedings of the IEEE **93**, 216 (2005), special issue on "Program Generation, Optimization, and Platform Adaptation".
- (Fryxell *et al.*, 2000) B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, H. Tufo, *FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes*, Astrophys. J. Supp. **131**, 273 (2000).
- (Garrison, Eisenstein & Pinto, 2019) L. H. Garrison, D. J. Eisenstein, P. A. Pinto, *A high-fidelity realization of the Euclid code comparison N-body simulation with ABACUS*, Mon. Not. R. Astron. Soc. **485**, 3370 (2019), arXiv:1810.02916 [astro-ph.CO].
- (Gonnet *et al.*, 2013) P. Gonnet, M. Schaller, T. Theuns, A. B. G. Chalk, *SWIFT: Fast algorithms for multi-resolution SPH on multi-core architectures* (2013), arXiv:1309.3783 [astro-ph.IM].
- (Guillet & Teyssier, 2011) T. Guillet, R. Teyssier, *A simple multigrid scheme for solving the Poisson equation with arbitrary domain boundaries*, J. Comp. Phys. **230**, 4756 (2011), arXiv:1104.1703 [physics.comp-ph].
- (Hahn & Abel, 2011) O. Hahn, T. Abel, *Multi-scale initial conditions for cosmological simulations*, Mon. Not. R. Astron. Soc. **415**, 2101 (2011), arXiv:1103.6031.
- (Hahn, Angulo & Abel, 2015) O. Hahn, R. E. Angulo, T. Abel, *The properties of cosmic velocity fields*, Mon. Not. R. Astron. Soc. **454**, 3920 (2015), arXiv:1404.2280.
- (Hockney & Eastwood, 1981) R. W. Hockney, J. W. Eastwood, *Computer Simulation Using Particles* (McGraw-Hill, 1981).
- (Howlett, Manera & Percival, 2015) C. Howlett, M. Manera, W. J. Percival, *L-PICOLA: A parallel code for fast dark matter simulation*, Astron. Comput. **12**, 109 (2015), arXiv:1506.03737.
- (Hunter, 2007) J. D. Hunter, *Matplotlib: A 2D Graphics Environment*, Comput. Sci. Eng. **9**, 90 (2007).
- (Huterer & Takada, 2005) D. Huterer, M. Takada, *Calibrating the nonlinear matter power spectrum: Requirements for future weak lensing surveys*, Astroparticle Phys. **23**, 369 (2005), arXiv:astro-ph/0412142 [astro-ph].
- (Izard, Crocce & Fosalba, 2016) A. Izard, M. Crocce, P. Fosalba, *ICE-COLA: towards fast and accurate synthetic galaxy catalogues optimizing a quasi-N-body method*, Mon. Not. R. Astron. Soc. **459**, 2327 (2016), arXiv:1509.04685 [astro-ph.CO].
- (James, 1977) R. A. James, *The Solution of Poisson's Equation for Isolated Source Distributions*, J. Comp. Phys. **25**, 71 (1977).
- (Johnson & Frigo, 2008) S. G. Johnson, M. Frigo, in *Fast Fourier Transforms*, edited by C. S. Burrus (Connexions, Rice University, Houston TX, 2008) Chap. 11.
- (Kluyver *et al.*, 2016) T. Kluyver, B. Ragan-Kelley, F. Pérez, B. E. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. B. Hamrick, J. Grout, S. Corlay, *et al.*, *Jupyter Notebooks - a publishing format for reproducible computational workflows*, in *ELPUB* (2016).
- (Klypin & Holtzman, 1997) A. Klypin, J. Holtzman, *Particle-Mesh code for cosmological simulations* (1997), arXiv:astro-



- ph/9712217 [astro-ph].
- (Knebe & Doumler, 2010) A. Knebe, T. Doumler, *AMIGA: Adaptive Mesh Investigations of Galaxy Assembly* (2010), ascl:1007.006.
- (Koda *et al.*, 2016) J. Koda, C. Blake, F. Beutler, E. Kazin, F. Marin, *Fast and accurate mock catalogue generation for low-mass galaxies*, *Mon. Not. R. Astron. Soc.* **459**, 2118 (2016), arXiv:1507.05329 [astro-ph.CO].
- (Laureijs *et al.*, 2011) R. Laureijs, J. Amiaux, S. Arduini, J. L. Augères, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvel, A. Ealet, *et al.*, *Euclid Definition Study Report* (2011), arXiv:1110.3193 [astro-ph.CO].
- (Leclercq, 2015) F. Leclercq, *Bayesian large-scale structure inference and cosmic web analysis*, Ph.D. thesis, Institut d'Astrophysique de Paris (2015), arXiv:1512.04985 [astro-ph.CO].
- (Leclercq, 2018) F. Leclercq, *Bayesian optimization for likelihood-free cosmological inference*, *Phys. Rev. D* **98**, 063511 (2018), arXiv:1805.07152.
- (Leclercq, Jasche & Wandelt, 2015) F. Leclercq, J. Jasche, B. Wandelt, *Bayesian analysis of the dynamic cosmic web in the SDSS galaxy survey*, *J. Cosmology Astropart. Phys.* **6**, 015 (2015), arXiv:1502.02690 [astro-ph.CO].
- (Leclercq *et al.*, 2019) F. Leclercq, W. Enzi, J. Jasche, A. Heavens, *Primordial power spectrum and cosmology from black-box galaxy surveys*, *Mon. Not. R. Astron. Soc.* **490**, 4237 (2019), arXiv:1902.10149 [astro-ph.CO].
- (Leclercq *et al.*, 2017) F. Leclercq, J. Jasche, G. Lavaux, B. Wandelt, W. Percival, *The phase-space structure of nearby dark matter as constrained by the SDSS*, *J. Cosmology Astropart. Phys.* **6**, 049 (2017), arXiv:1601.00093 [astro-ph.CO].
- (LIGO Scientific Collaboration, 2015) LIGO Scientific Collaboration, *Advanced LIGO*, *Class. Quantum Gravity* **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].
- (LSST Science Collaboration, 2012) LSST Science Collaboration, *Large Synoptic Survey Telescope: Dark Energy Science Collaboration* (2012), arXiv:1211.0310 [astro-ph.CO].
- (Merloni *et al.*, 2012) A. Merloni, P. Predehl, W. Becker, H. Böhringer, T. Boller, H. Brunner, M. Brusa, K. Dennerl, M. Freyberg, P. Friedrich, *et al.*, *eROSITA Science Book: Mapping the Structure of the Energetic Universe* (2012), arXiv:1209.3114 [astro-ph.HE].
- (Ocvirk *et al.*, 2016) P. Ocvirk, N. Gillet, P. R. Shapiro, D. Aubert, I. T. Iliev, R. Teyssier, G. Yepes, J.-H. Choi, D. Sullivan, A. Knebe, *et al.*, *Cosmic Dawn (CoDa): the First Radiation-Hydrodynamics Simulation of Reionization and Galaxy Formation in the Local Universe*, *Mon. Not. R. Astron. Soc.* **463**, 1462 (2016), arXiv:1511.00011 [astro-ph.GA].
- (Perez & Granger, 2007) F. Perez, B. E. Granger, *IPython: A System for Interactive Scientific Computing*, *Comput. Sci. Eng.* **9**, 21 (2007).
- (Planck Collaboration, 2016) Planck Collaboration, *Planck 2015 results. XIII. Cosmological parameters*, *Astron. & Astrophys.* **594**, A13 (2016), arXiv:1502.01589.
- (Potter, Stadel & Teyssier, 2017) D. Potter, J. Stadel, R. Teyssier, *PKDGRAV3: beyond trillion particle cosmological simulations for the next era of galaxy surveys*, *Comput. Astrophys. Cosmology* **4**, 2 (2017), arXiv:1609.08621 [astro-ph.IM].
- (Quinn *et al.*, 1997) T. Quinn, N. Katz, J. Stadel, G. Lake, *Time stepping N-body simulations* (1997), arXiv:astro-ph/9710043 [astro-ph].
- (Schneider *et al.*, 2019) A. Schneider, R. Teyssier, J. Stadel, N. E. Chisari, A. M. C. Le Brun, A. Amara, A. Refregier, *Quantifying baryon effects on the matter power spectrum and the weak lensing shear correlation*, *J. Cosmology Astropart. Phys.* **3**, 020 (2019), arXiv:1810.08629 [astro-ph.CO].
- (Simon *et al.*, 2018) S. M. Simon, J. A. Beall, N. F. Cothard, S. M. Duff, P. A. Gallardo, S. P. Ho, J. Hubmayr, B. J. Koopman, J. J. McMahon, F. Nati, *et al.*, *The Advanced ACTPol 27/39 GHz Array*, *J. Low Temp. Phys.* **193**, 1041 (2018).
- (Simons Observatory Collaboration, 2019) Simons Observatory Collaboration, *The Simons Observatory: science goals and forecasts*, *J. Cosmology Astropart. Phys.* **2**, 056 (2019), arXiv:1808.07445 [astro-ph.CO].
- (SPHEREx Science Team, 2018) SPHEREx Science Team, *The SPHEREx All-Sky Spectral Survey, in American Astronomical Society Meeting Abstracts #231*, American Astronomical Society Meeting Abstracts, Vol. 231 (2018) p. 354.21.
- (Square Kilometre Array Cosmology Science Working Group, 2018) Square Kilometre Array Cosmology Science Working Group, *Cosmology with Phase 1 of the Square Kilometre Array; Red Book 2018: Technical specifications and performance forecasts* (2018), arXiv:1811.02743 [astro-ph.CO].
- (Strauss *et al.*, 1995) M. A. Strauss, R. Cen, J. P. Ostriker, T. R. Lauer, M. Postman, *Can Standard Cosmological Models Explain the Observed Abell Cluster Bulk Flow?*, *Astrophys. J.* **444**, 507 (1995), arXiv:astro-ph/9406038 [astro-ph].
- (Tassev & Zaldarriaga, 2012) S. Tassev, M. Zaldarriaga, *Estimating CDM particle trajectories in the mildly non-linear regime of structure formation. Implications for the density field in real and redshift space*, *J. Cosmology Astropart. Phys.* **12**, 011 (2012), arXiv:1203.5785 [astro-ph.CO].
- (Tassev, Zaldarriaga & Eisenstein, 2013) S. Tassev, M. Zaldarriaga, D. J. Eisenstein, *Solving large scale structure in ten easy steps with COLA*, *J. Cosmology Astropart. Phys.* **6**, 036 (2013), arXiv:1301.0322 [astro-ph.CO].
- (Tassev *et al.*, 2015) S. Tassev, D. J. Eisenstein, B. D. Wandelt, M. Zaldarriaga, *sCOLA: The N-body COLA Method Extended to the Spatial Domain* (2015), arXiv:1502.07751 [astro-ph.CO].
- (Teyssier, 2002) R. Teyssier, *Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES*, *Astron. & Astrophys.* **385**, 337 (2002), arXiv:astro-ph/0111367 [astro-ph].
- (The Virgo Collaboration, 2020) The Virgo Collaboration, *Advanced Virgo Status*, in *Journal of Physics Conference Series*, *Journal of Physics Conference Series*, Vol. 1342 (2020) p. 012010.
- (Theuns *et al.*, 2015) T. Theuns, A. Chalk, M. Schaller, P. Gonnet, *SWIFT: task-based hydrodynamics and gravity for cosmological simulations* (2015), arXiv:1508.00115 [astro-ph.IM].
- (Thyng *et al.*, 2016) K. M. Thyng, C. A. Greene, R. D. Hetland, H. M. Zimmerle, S. F. DiMarco, *True Colors of Oceanography: Guidelines for Effective and Accurate Colormap Selection*, *Oceanography* **29** (2016).
- (van Daalen *et al.*, 2011) M. P. van Daalen, J. Schaye, C. M. Booth, C. Dalla Vecchia, *The effects of galaxy formation on the matter power spectrum: a challenge for precision cosmology*, *Mon. Not. R. Astron. Soc.* **415**, 3649 (2011), arXiv:1104.1174 [astro-ph.CO].
- (van der Walt, Colbert & Varoquaux, 2011) S. van der Walt, S. C. Colbert, G. Varoquaux, *The NumPy Array: A Structure for Efficient Numerical Computation*, *Comput. Sci. Eng.* **13**, 22 (2011), arXiv:1102.1523 [cs.MS].
- (Yu, Pen & Wang, 2018) H.-R. Yu, U.-L. Pen, X. Wang, *CUBE: An Information-optimized Parallel Cosmological N-body Algorithm*, *Astrophys. J. Supp.* **237**, 24 (2018), arXiv:1712.06121 [astro-ph.CO].
- (Zel'dovich, 1970) Y. B. Zel'dovich, *Gravitational instability: An approximate theory for large density perturbations.*, *Astron. & Astrophys.* **5**, 84 (1970).