DOCTORAL THESIS

Producing simulated catalogues for next generation galaxy surveys

Author:
ALBERT IZARD ALBERICH

Supervisors:
Dr. PABLO FOSALBA VELA
Dr. HÉCTOR MARTÍN CROCCE

Tutor:
Dr. EDUARD MASSÓ SOLER

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy

in the

UNIVERSITAT AUTÒNOMA DE BARCELONA
DEPARTAMENT DE FÍSICA

ICE INSTITUT DE CIÈNCIES DE L'ESPAI
CSIC CONSEJO SUPERIOR DE INVESTIGACIONES CIENTÍFICAS
IEEC

July 7, 2016

Abstract

Current and future galaxy surveys will be able to map the large-scale structure of the Universe with unprecedented detail and measure cosmological parameters with exquisite precision. In order to develop the science cases and the analysis pipelines, it is necessary an accurate modelling of the non-linear gravitational evolution. This thesis presents a methodology for producing accurate mock catalogues, much faster than conventional methods (2-3 orders of magnitude) and incorporating past light cone effects.

First, we present the optimization of a quasi N-body method in the compromise between accuracy and computational cost. We studied how variations in the code parameter space have and impact on the accuracy of observables such as the halo abundance and distribution and matter clustering. We propose optimal parameter configurations for achieving high accuracy as compared to exact N-body simulations and we explore different calibration techniques to match even better two-point halo clustering statistics.

The next step is mimicking the geometry of real astrophysical observations, in which distant objects are seen in the past light cone. We introduce `ICE-COLA`, a simulation code developed for this thesis that implements the production of light cone catalogues on-the-fly. The user can generate three different kinds of data types. The first contains all the information of the phase-space matter distribution while the others store high-level data catalogues ready to use to model galaxy surveys. This enables large compression factors of ~ 2 orders of magnitude in the data volume to be stored. In particular, the code can generate halo catalogues in the light cone and pixelated two-dimensional projected matter density maps in spherical concentric shells around the observer.

Using `ICE-COLA` we produce large light cone simulations and perform an extensive validation of the catalogues. We introduce a novel methodology to model weak gravitational lensing with an approximate method and we show that we can resolve most of the scales probed by current weak lensing experiments. Finally we extend the results to halo mock catalogues with weak lensing quantities, which represents a key step forward modelling galaxy clustering and weak lensing observables consistently in a quasi N-body approach.

Acknowledgements

En primer lugar, mis agradecimientos van a Pablo y a Martín. Pablo, vas ser qui em va obrir la primera porta al món de la cosmologia. La meva trajectòria hauria sigut diferent si no fos per tu. Martín, me has brindado muchas otras posibilidades durante la tesis y el haber podido trabajar en temas que han acabado siendo centrales. Gracias por vuestro esfuerzo, tiempo, conocimientos y oportunidades que me habéis dado.

Al grup de cosmologia de l'ICE, especialment a Francisco i a Enrique perquè la meva feina també ha estat possible gràcies a vosaltres. And to many people that have been part of the group: Jorge, Santi, Ricard, Christopher, Anne, Antonino, Stéphanie, Noélia, Raoul, Anna, Àlex, Kwan, Martin... També al Pau del PIC. Sam and Darren, it would have been a pleasure to collaborate with you during a longer period of time. Kai, you are probably waiting for me to finish the thesis before having a beer. It will be soon... Linda, having the words deadline, debugging and simulations in the same sentence is not a good idea. Thanks for the great effort you did.

My acknowledgement goes also to Ariel and Pierluigi, my work is also possible thanks to you. També a l'Eduard i al Marc. I want to mention Jun Koda and Yu Feng for making codes publicly available. I hope to follow the same line soon. And Alina and Jason, for this really nice opportunity.

Vull agraïr a totes les persones que he pogut conèixer i que han fet que les hores a l'ICE, els dinars i les pauses, o fora d'allà, fossin genials. Arnau, no sé si recordar els llençols compartits, el mapamundi o haver estat instigador d'un concurs que... no cal dir res més. Laura, seguiré celebrando el aniversario de nuestro cumpleaños. Cadascú em fa pensar en algo però aniré al gra: Jacobo, Jonatan, Núria, Carmen, Antonia, Padu, Carles, Marina, Hilda, Daniel, Daniele, Fran, Ferran, Nataly, Josep, Manuel, Jian. Heu fet que aquesta època hagi valgut molt la pena!

Als meus pares. Més que per l'encert amb aquell llibre de la infància, sobretot per tot allò que va més enllà d'un gest, frase o fet puntual. Pels valors i llibertat en què ens heu fet créixer. Oriol, Marta, Berta, Ricard, Arnau, Berta, Aniol i Martí, pels molts somriures.

Andrea, Arnau, Miquel, Roser, Sebi i Víctor: un trosset de la tesi va per vosaltres.

Al CSIC, per decidir que al meu any encara es mantenia el programa JAE. Al Barcelona Supercomputing Center, pel temps concedit al MareNostrum (AECT-2014-2-0013, AECT-2014-3-0012, AECT-2015-1-0013, AECT-2015-2-0011). Ah, i al tècnic de l'aire acondicionat que ha activat el sistema una setmana abans de dipositar la tesi.

Certified GF Gluten-Free

CONTENTS

INTRODUCTION

Seeing is in some respect an art, which must be learnt.

Sir William Hershchel

Motivation

Present and planned galaxy surveys like the Dark Energy Survey[1] (DES, The Dark Energy Survey Collaboration, 2005), the Kilo Degree Survey (KiDS; de Jong et al., 2013) the Large Synoptic Survey Telescope[2] (LSST, LSST Science Collaboration et al., 2009), Euclid[3] (Laureijs et al., 2011) , the Wide-Field Infrared Survey Telescope – Astrophysics Focused Telescope Asset[4] (WFIRST–AFTA; Spergel et al., 2013), the Extended Baryon Oscillation Spectroscopic Survey[5] (eBOSS, Dawson et al., 2016), the Dark Energy Spectroscopic Instrument[6] (DESI, Levi et al., 2013), will generate a wealth of high-quality data that will allow to test the nature of dark matter and dark energy and constrain possible deviations from the standard cosmological model based on General Relativity (for reviews or textbooks see e.g. Peebles, 1980; Dodelson, 2003; Albrecht et al., 2006; Weinberg et al., 2013; Joyce et al., 2016).

Galaxy clustering encompasses a wide class of observational probes that aim at extracting cosmological information from the distribution of galaxies in the sky. Galaxy clusters are the final result of the amplified primordial fluctuations by gravity, and analysing their correlations it is possible to test the theories of gravity, the expansion history and the matter content of the Universe (Allen et al., 2011; Silk et al., 2013). In particular, Baryon Acoustic Oscillations (BAO) leave an imprint on the matter distribution that can be measured as an excess of clustering of galaxies at a certain scale (for a review, see Bassett et al., 2010). Therefore, it provides a standard ruler that can be measured with two-point statistics. The first detection of this signature was reported in 2005 by Eisenstein et al., 2005 using data from the Sloan Digital Sky Survey (SDSS, York et al., 2000).

[1] http://www.darkenergysurvey.org/
[2] http://www.lsst.org
[3] http://www.euclid-ec.org/
[4] http://wfirst.gsfc.nasa.gov/
[5] http://www.sdss.org/surveys/eboss/
[6] http://desi.lbl.gov/

In the present, the most precise BAO measurements derive from the Baryonic Oscillation Spectroscopic Survey (BOSS) DR12 galaxy sample (Cuesta et al., 2016; Gil-Marin et al., 2016), which provided a measurement of the BAO scale at the mean redshift $z = 0.57$ with a precision almost at the percent level. See the left panel in Fig. 1 for the cosmological constraints from BOSS.

Another major observable of galaxy surveys is weak gravitational lensing. This effect is the consequence of light rays being bended by massive objects as they travel through the Universe, as predicted by Albert Einstein's general theory of relativity. It is a phenomenon that carries information about the matter distribution along the line-of-sight and therefore it is an invaluable tool for probing cosmological models, such as the theory of gravity, the properties of dark matter and dark energy or the physics at inter-cluster scales (for comprehensive reviews, see Bartelmann et al., 2001; Huterer, 2010; Kilbinger, 2015). There are many applications of weak gravitational lensing. Cosmic shear is the analysis of tidal deformation of images, causing galaxy shapes to have coherent distortions. By averaging over many galaxies, their intrinsic orientations (if random), are expected to be cancelled, and the lensing signal can then be measured. The first detections of cosmic shear were done in 2000 by several independent groups (Bacon et al., 2000; Kaiser et al., 2000; Van Waerbeke et al., 2000; Wittman et al., 2000; Miyazaki et al., 2002), in fields that ranged from 0.5 to $2\,\mathrm{deg}^2$. Next milestones were set by the Canada-France Hawaii Legacy Survey (CFHTLS; Hoekstra et al., 2006) and the CFHT Lens Survey (CFHTLenS; Kilbinger et al., 2013), which by measuring the two-dimensional shear correlation function in an observed area of more than $100\,\mathrm{deg}^2$ were able to constrain cosmological parameters (see right panel in Fig 1).

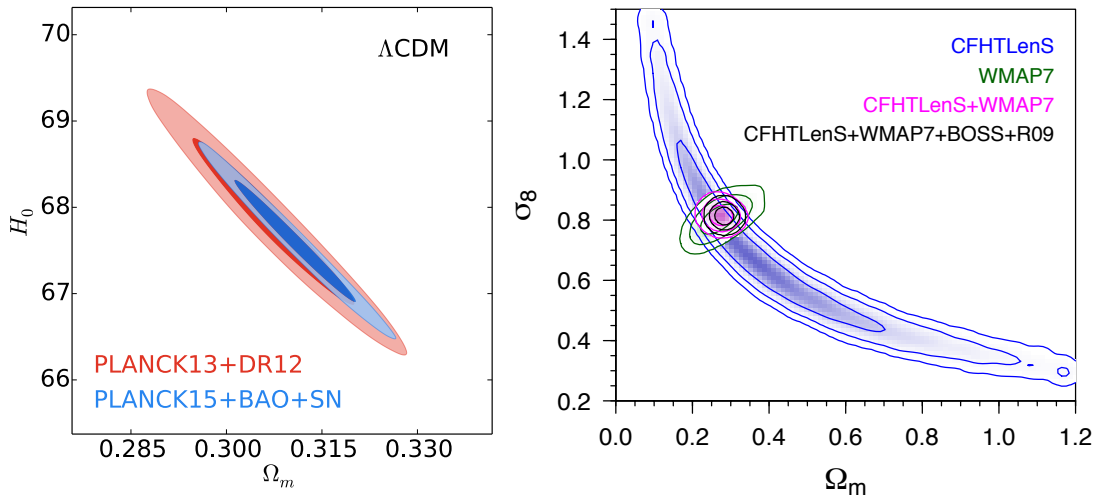

FIGURE 1: Left: Constraints on the expansion rate and the matter content from BOSS in combination with other probes. Right: contours on the $\sigma_8 - \Omega_m$ plane from the CFHTLLens survey and other data sets. (From Cuesta et al., 2016 and Kilbinger et al., 2013 for BOSS and CFHTLenS respectively).

Building upon these first measurements, ongoing shear surveys rely on observing

large samples of galaxies in multiple photometric bands and estimating radial distances thanks to photometric redshifts. However, the covered area is now being increased by an order of magnitude. This is the case of KiDS and DES, which will observe 1500 and 5000 deg^2 respectively once they are completed. Nonetheless, next generation projects will cover most of the extragalactic sky, such as the LSST, Euclid and WFIRST–AFTA. Weak lensing is currently entering a golden era in which statistical uncertainties are being reduced drastically. In the near future it may become one of the most powerful cosmological probes if systematic effects are controlled in the modelling and analysis so that become sub-dominant in the overall error budget. Likelihood analyses use theoretical predictions of covariance matrices, which are hard to model for weak lensing since it is an observable that probes highly non-linear scales at the same time that samples very large volumes. Standard numerical simulation methods of large-scale structure formation are able to accurately model cosmological observables but become prohibitively expensive if a large number of realizations are required ljto estimate covariances by Monte-Carlo methods.

An optimal extraction of cosmological parameters from those very large and complex datasets will ultimately rely on our ability to model cosmological observables and their covariances with high accuracy. This entails the development of synthetic observations based on mock catalogues produced from numerical simulations, that allow to optimize the pipelines, model and mitigate systematic effects, calibrate algorithms, test new techniques, etc. The requirement of sampling large cosmological volumes while still resolving small scales is a big challenge to current N-body simulation codes (some state-of-the-art simulations are Teyssier et al., 2009; Kim et al., 2011; Angulo et al., 2012; Alimi et al., 2012; Watson et al., 2013; Klypin et al., 2013; Skillman et al., 2014; Heitmann et al., 2014; Fosalba et al., 2015a; for a review see Kuhlen et al., 2012). Moreover, hundreds or thousands of realizations are needed for robustly estimating covariance matrices (crucial for cosmological parameter estimation, see Taylor et al., 2013; Blot et al., 2015) or for propagating errors in complex and non-linear analysis (e.g. Baryon Acoustic Oscillations reconstruction, see Takahashi et al., 2009; Manera et al., 2013; Kazin et al., 2014; Ross et al., 2015). Yet, producing massive ensembles of N-body mocks is computationally prohibitive and alternative routes need to be devised in order to face the enormous challenge. This is what approximate methods aim to solve, which overcome the problem of explicitly solving the non-linear evolution by incorporating a simpler yet accurate modelling of the small-scale physics.

Most of the experiments mentioned in this introduction probe both weak lensing and galaxy clustering simultaneously. Combining both in a single analysis is definitely the best way to tighten cosmological constraints, since degeneracies affecting individual probes can be broken, mitigating also the impact of their corresponding systematic effects. In the roadmap for designing and optimally exploiting next generation of galaxy surveys it is therefore necessary to model accurately both observables in a consistent way and produce thousands of mock catalogues to analyse the data. This thesis presents a methodology to achieve that goal. A fast method is optimized in the compromise between accuracy and computational cost. Efficient light cone simulations are developed

and they are used to model galaxy clustering and weak lensing coherently. The method and its accuracy are validated against full N-body simulations, which typically require between 2 and 3 orders of magnitude more computing time.

Contents of this thesis

This thesis is organized as follows. Chapter 1 gives a general introduction to observational cosmology, with special emphasis on large-scale structure concepts and the formalism of weak gravitational lensing. Chapter 2 reviews the methods to produce mock catalogues, both from N-body simulations and approximate methods. It explains the N-body simulation used in this thesis as a benchmark, the Marenostrum Institut de Ciències de l'Espai-Grand Challenge (MICE-GC), as well as the fast method extensively used, tested and expanded in this work: `COLA`. In the last section we detail all the simulations developed for this thesis. Chapter 3 is where the first results are discussed, presenting an optimization of the simulation method employed in a compromise between accuracy and computational cost. It studies the impact of variations of the internal code parameters on observable quantities and determines the optimal set-up. From Chapter 4 onwards the work is based on new light cone simulations developed with `ICE-COLA`. First is explained the algorithm for generating such catalogues on-the-fly, able to deliver high-level data products tailored for modelling galaxy surveys, and the performance of the method is discussed at length. After, Chapter 5 presents a basic validation of such simulations and a description of the pipeline developed to model weak lensing quantities. A further validation of the maps is discussed in terms of the two-point statistics. Finally, Chapter 6 shows how halo catalogues are combined together with weak lensing maps to model consistently both weak lensing and galaxy clustering from the same simulation. Lastly, some prospects of ongoing and future projects are outlined in Chapter 7 and a summary with conclusions is given in Chapter 8. In addition, Appendix A explains how transient effects were considered when comparing `COLA` simulations with MICE-GC, whereas Appendix B compares the gain of the `COLA` method with respect to Particle-Mesh-only runs.

1

OBSERVATIONAL COSMOLOGY

During the recent history, our conception of the Universe has radically changed by rejecting those theories that do not surpass the exam of observational evidences. The Ptolemaic model placed the Earth in the centre of the cosmos. The Copernican Revolution, starting in 1543, displaced ourselves for the first time from the centre by using an heliocentric model. After important contributions by Tycho Brahe, Johannes Kepler and Galileo Galilei, the revolution concluded in 1687 when Newton published in his *Principia* the law of universal gravitation, which allowed to describe mathematically the dynamics of the orbits of the planets. The Copernican Revolution implies the renounce of a privileged position for the Earth and a greater generalization of it is the so called Cosmological Principle:

> *The Universe is homogeneous and isotropic at large scales*

Or in other words, the same physical laws are valid everywhere, i.e., there are no privileged observers. Modern cosmology assumes this axiom, which actually has been proved to a high degree of accuracy (see e.g., Planck Collaboration et al., 2015a for the isotropy of the Cosmic Microwave Background, CMB, as measured by the European Space Agency Planck satellite mission). In this framework, the evolution of the Universe can be decomposed into two pieces: the background evolution, that describes the metric and properties of the homogeneous Universe, and the structure formation at small scales, that studies the departures from homogeneity. This Chapter gives an introduction to both aspects in § 1.2 and § 1.3 respectively. Before, § 1.1 describes the consensus cosmological model briefly. Finally, § 1.4 gives the basic principles of weak gravitational lensing, which is a powerful observational probe.

1.1 The ΛCDM model

The ΛCDM model is the current consensus cosmological model (Hamilton, 2013) that explains most of the observations, such as the abundances of light elements, the properties of the CMB radiation (Planck Collaboration et al., 2015b), supernova type Ia data (Suzuki et al., 2012) and the large-scale structure clustering (Anderson et al., 2012). The model is built on the cosmological principle and adopts the concept of Big Bang, in which the Universe has been expanding from an initial singularity. Assumes the General Relativity theory as a model for gravity but needs to postulate the existence of some unknown fluids: dark matter and dark energy.

In 1933, Zwicky studied the dynamical state of the Coma cluster and found the first evidences of the problem of the missing matter (Zwicky, 1933). This was later confirmed by other observational techniques, such as velocity rotational curves of galaxies, gravitational lensing and the large-scale structure formation. The ΛCDM model postulates the existence of a form of Cold Dark Matter (CDM), that has the properties of only interacting through gravity and being non-relativistic.

Dark energy was introduced after 1998, when two different teams (High-Z Supernova Research Team and Supernova Cosmology Project) were searching for type Ia supernovae for using them as standard candles. Their findings showed an unexpected accelerated expansion rate (Riess et al., 1998; Perlmutter et al., 1999). The simplest for of the standard cosmological model includes a dark energy fluid with a constant equation of state parameter $w = {}^{P}/_{\rho} = -1$, corresponding to a cosmological constant, that causes the repulsive force.

Current tightest constraints on cosmological parameters suggest that the energy budget of the Universe consists on: 5% of baryonic matter, 27% of dark matter and 68% of dark energy, being consistent with a flat geometry (Planck Collaboration et al., 2015b). Our understanding of the Universe has progressed much in the last decades, but there are still some missing pieces that have to be addressed. This includes topics such as the mass of neutrinos, possible deviations from general relativity on the largest scales, early massive black hole formation, the mechanism responsible of inflation..., as well as the nature of the elusive dark matter and dark energy.

1.2 Dynamics of the expansion

In general relativity, the metric has a central role in the description of the Universe as a whole. The Friedmann-Lemaître-Robertson-Walker (FLRW) metric provides the geometry of an homogeneous and isotropic Universe

$$ds^2 = c^2 dt^2 - a(t) \left[dr^2 + S_k^2(r)(d\theta^2 + \sin^2\theta \, d\phi^2) \right], \tag{1.1}$$

where r is a radial comoving distance, θ and ϕ are angular coordinates, a is the scale factor and S_k depends on the curvature k of the Universe and can be

$$S_k(r) = \begin{cases} \sinh(r) & k = -1 \\ r & k = 0 \\ \sin(r) & k = +1. \end{cases} \tag{1.2}$$

By introducing the FLRW metric into the General Relativity field equations, the Friedmann equations can be derived

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G\rho}{3} - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3}, \tag{1.3}$$

where H is the Hubble parameter, ρ includes the contributions to the density from the matter and the radiation and Λ is the cosmological constant. This is the first of the Friedmann's equations, that relates the geometry and the density. For instance, the critical density can be defined as the value that corresponds to a flat Universe ($k = 0$)

$$\rho_c = \frac{3H^2}{8\pi G}. \tag{1.4}$$

It is useful to use the density parameters, which re-scale the density contributions relative to the critical density

$$\Omega_i \equiv \frac{\rho_i}{\rho_c} \tag{1.5}$$

$$\Omega_m = \frac{8\pi G\rho_0}{3H_0^2} \tag{1.6}$$

$$\Omega_r = \frac{8\pi G\rho_{r,0}}{3H_0^2} \tag{1.7}$$

$$\Omega_\Lambda = \frac{\Lambda c^2}{3H_0^2} \tag{1.8}$$

$$\Omega_k = \frac{-kc^2}{(a_0 H_0)^2}, \tag{1.9}$$

where the subscript 0 indicates that a quantity is evaluated at the present. These parameters have been defined to fulfil the relation $1 = \Omega_m + \Omega_r + \Omega_\Lambda + \Omega_k$. Thereby, the Friedmann equation read

$$H^2(z) = H_0^2 E^2(z) = H_0^2 \left(\Omega_r(1+z)^4 + \Omega_m(1+z)^3 + \Omega_k(1+z)^2 + \Omega_\Lambda\right), \tag{1.10}$$

where it was introduced the function $E(z)$ and the present day Hubble constant: $H_0 \equiv 100h\,\mathrm{km/s/Mpc}$. Applying the FLRW metric (equation 1.1) to a photon, $ds^2 = 0$, and integrating over the radial coordinate, the relation between the comoving distance and the

redsfhit can be written as

$$\chi(z) = \frac{c}{H_0} \int_0^z \frac{dz'}{E(z')}. \tag{1.11}$$

1.3 Large-scale structure formation

The large-scale structure is the gravitationally amplified result of the initial quantum fluctuations that arose in the early and hot Universe. Inhomogeneities grow driven by the gravitational attraction, resulting in bound objects, such as stars, galaxies and galaxy clusters, that were born the firsts few hundreds millions of years after the Big Bang. The underlying physics of the processes leading to structure formation have been extensively explained in many books (see e.g. Peacock, 1999; Peebles, 1980; Liddle et al., 2000). This section focuses on the gravitational evolution of density perturbations at large scales and on their statistical description. In what follows, the density refers to the contribution of both the dark matter and the baryon components.

The dimensionless density contrast is defined as the relative deviation from the mean matter density $\bar{\rho}$

$$\delta(\boldsymbol{x}) = \frac{\rho(\boldsymbol{x})}{\bar{\rho}} - \bar{\rho}. \tag{1.12}$$

The two-point correlation function is the expected excess of probability with respect to a homogeneous distribution of having two over-densities separated by a distance r

$$\xi(r) =< \delta(\boldsymbol{x})\delta(\boldsymbol{x}+\boldsymbol{r}) >, \tag{1.13}$$

which thanks to isotropy it does not depend on the angular part. The power spectrum is the equivalent correlator in Fourier space

$$P(k) = \int \xi(r) e^{i\boldsymbol{k}\boldsymbol{x}}\, d^3x\,. \tag{1.14}$$

Fourier transforming the density field gives

$$\delta(\boldsymbol{k}) = \int \delta(\boldsymbol{x}) e^{i\boldsymbol{k}\boldsymbol{x}}\, d^3x \tag{1.15}$$

and for Gaussian fields, different modes are uncorrelated

$$< \delta(\boldsymbol{k}_i)\delta(\boldsymbol{k}_j) >= (2\pi)^2 \delta_D(\boldsymbol{k}_i - \boldsymbol{k}_j) P(k_i), \tag{1.16}$$

where δ_D is the Dirac delta function and k is the modulus of the wavenumber $\boldsymbol{k}$.

A Gaussian field is completely characterized by its 2-point statistics: any non-vanishing higher order moment can be expressed as a function of the former. Or in other words, all reduced higher order moments vanish.

1.3.1 Linear growth of structure

Inhomogeneities in the initial density field $\delta(\boldsymbol{x})$ grow due to the gravitational forces. In a matter-dominated epoch, the Poisson equation can be written as

$$\nabla^2\phi = \frac{3H_0^2}{2a}\Omega_0\delta. \tag{1.17}$$

In the linear regime, perturbations are small (i.e. $\delta \ll 1$) and follow a second order partial differential equation

$$\ddot{\delta} + 2H\dot{\delta} - \frac{3}{2a^3}\Omega_0 H_0^2\delta = 0. \tag{1.18}$$

The fact that there are only temporal derivatives implies that the temporal evolution is the same at all positions, and the solution can be expressed by the linear growth rate function (see e.g., Dodelson, 2003)

$$D_1(z) = \frac{5\Omega_0 E(z)}{2}\int_z^\infty \frac{1+z'}{E(z')}\,dz', \tag{1.19}$$

which in the Einstein-de Sitter case, it is simply $D_1(a) = a$.

Therefore, the evolution of perturbations in the linear regime is given by the linear growth factor: $\delta(\boldsymbol{x}, z) = \delta(\boldsymbol{x}, z_0)D_1(z)/D_1(z_0)$. And the matter power spectrum grows as $P(k, z) = [D_1(z)/D_1(z_0)]^2 P(k, z_0)$.

The evolution of differential mass elements (or particles, which are the elements that sample the phase space in an N-body simulation, see § 2.1) can be expressed in terms of the growth factor as well in a perturbative approach. The displacement field $\boldsymbol{s}(\boldsymbol{q}, a)$ is defined as the difference between the Eulerian position $\boldsymbol{x}(\boldsymbol{q}, a)$ and the Lagrangian (initial) position $\boldsymbol{q}$:

$$\boldsymbol{s}(\boldsymbol{q}, a) = \boldsymbol{x}(\boldsymbol{q}, a) - \boldsymbol{q}. \tag{1.20}$$

In Lagrangian Perturbation Theory (LPT, for reviews see Bernardeau et al., 2002; Bouchet et al., 1995), the displacement field can be expanded to second order as

$$\boldsymbol{s}(\boldsymbol{q}, a) = D_1(a)\boldsymbol{s}_1(\boldsymbol{q}, a) + D_2(a)\boldsymbol{s}_2(\boldsymbol{q}, a), \tag{1.21}$$

where sub-indices now distinguish between the linear (or Zel'dovich, see Zel'dovich, 1970) and the second-order perturbative terms. The linear displacement field depends on the gradient of the potential at lagrangian positions, while the second order component includes tidal distortions.

1.3.2 Non-linear evolution

At late times, the condition $\delta \ll 1$ is no longer satisfied for all scales. Small-scale fluctuations grow beyond the collapse threshold and start a gravitational collapse (see e.g. Press et al., 1974). As a result, bounded objects of dark matter are formed and sustained by

virialization. These are called haloes and galaxies, the baryonic component, are believed to form inside them.

Initially the density field can be well described as a Gaussian field. Different modes are uncorrelated (see equation 1.16) and the covariance of the power spectrum is diagonal, with the variance given by

$$\sigma^2(P(k)) = \frac{2P^2(k)}{N_{\text{modes}}}, \tag{1.22}$$

where N_{modes} is the number of k-modes contributing to the measurement at the wavenumber k (i.e. $N_{\text{modes}} \approx k^2 \Delta k V/(2\pi)^2$ with V being the volume and Δk the width of the bin). The non-linear evolution couples different modes, there is a transfer of power from large to small scales and the amplitude of small-scale fluctuations deviates from the linear growth rate. This is visible in the matter power spectrum as the so-called non-linear bump for modes $k \gg 0.1\,h^{-1}\,\text{Mpc}$. Fig. 1.1 displays the non-linear matter power spectrum measured in a COLA simulation, where it is visible the deviations from linear theory at small scales and the smoothing of the Baryon Acoustic Oscillations (BAO) feature due to mode coupling (see Beutler et al., 2015; Fosalba et al., 2015a; Ross et al., 2015 for how non-linear effects in the BAO are modelled by simulations). The non-linear covariance matrix is shown in Fig. 1.2, where it becomes evident a large correlation between different modes for scales $k > 0.2\,h\,\text{Mpc}^{-1}$.

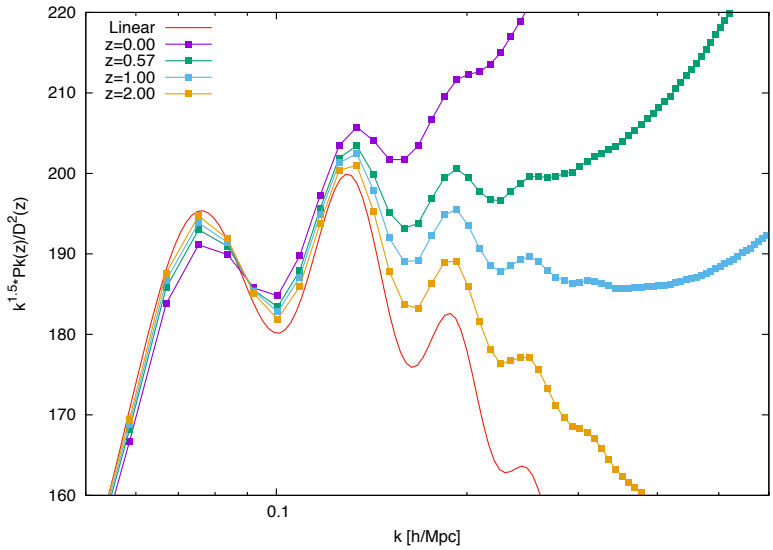

FIGURE 1.1: Matter power spectra at four different redshifts from a COLA simulation compared to the linear power spectrum. All power spectra have been scaled by the linear growth factor (see equation 1.19) to redshift 0. The non-linear evolution is clearly visible at scales $k > 0.15\,h\,\text{Mpc}^{-1}$, where the power deviates progressively from linear theory with time. Note as well how Baryon Acoustic Oscillations, that extend to large scales, are damped.

The non-linear regime becomes much more difficult to determine analytically as the perturbative scheme breaks down and for an accurate modelling it is necessary to resolve

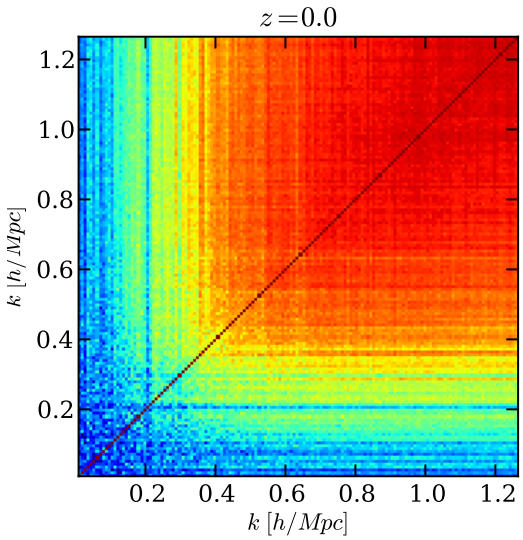

FIGURE 1.2: Correlation matrix of the matter power spectrum measured from a set of 100 COLA simulations at redshift 0. The prediction for a Gaussian field is the identity matrix, but departures from that are evident in this plot for $k > 0.2\,h\,\mathrm{Mpc}^{-1}$, where off-diagonal elements display strong correlations between different modes.

the full dynamics exactly using N-body simulations.

1.3.3 Halo and galaxy Bias

Haloes are formed in high density regions of the matter field. Thus their distribution is correlated with the peaks of the matter distribution. However, and according to the peak-background split theory (Bardeen et al., 1986; Cole et al., 1989; Sheth et al., 1999), it is harder for density peaks in low density environments to reach the collapse threshold and less haloes will be able form in these regions. Therefore, the halo formation is enhanced in rich environments and suppressed in voids, producing relative variations in the abundance larger than those of the density. The consequence is that haloes are biased tracers of the matter field, where biasing means an enhanced clustering. The concept applies as well to galaxies, since these inhabit in haloes, and are referred as halo or galaxy bias. At first order, the bias can be expressed as (Fry et al., 1993)

$$\delta_h(k,z) = b_h(k,z)\delta(k,z). \tag{1.23}$$

In this case, the halo power spectrum would be given by

$$P_h(k,z) = b_h^2(k,z)P(k,z). \tag{1.24}$$

Bias might depend on the scale, the epoch and the properties of the tracer object (in the case of haloes for instance, their mass). The expansion can also be carried to higher orders, yielding a full hierarchy of bias parameters.

1.4 Principles of weak gravitational lensing

One of the predictions of Albert Einstein's general theory of relativity is that light paths are bent by the matter inhomogeneities of the Universe via the distortions that are produced in the metric tensor of space-time. In fact, Newtonian physics can also explain this phenomenon, but it predicts an effect a factor of two weaker. Images of distant objects are distorted by the intervening matter distribution , inducing changes in shapes, apparent positions and fluxes of the background population. In most situations the effect is small, since relative variations are at the percent level, and then it is referred as weak gravitational lensing. It is a phenomena that carries information about the matter distribution along the line-of-sight and therefore it is an invaluable tool for probing cosmological models, such as the theory of gravity, the properties of dark matter and dark energy or the physics at inter-cluster scales. This section gives an overview of the mathematical description of this phenomena in order to present expressions that are used in later chapters. For extensive reviews see Narayan et al., 1996; Bartelmann et al., 2001; Van Waerbeke et al., 2003; Munshi et al., 2008; Huterer, 2010; Kilbinger, 2015.

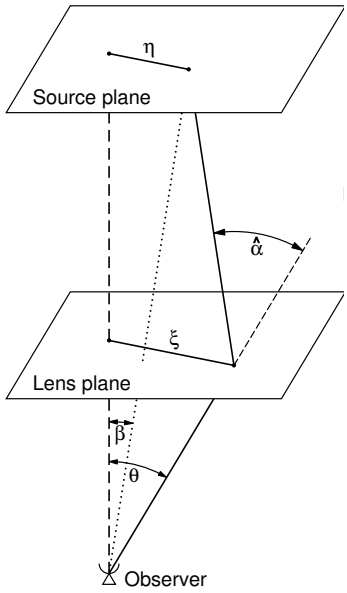

FIGURE 1.3: Sketch of the geometry of the gravitational lensing effect (from Bartelmann et al., 2001).

1.4.1 Lens equation

The geometry of the gravitational lensing effect is mainly a three-object system, formed by a source, a lens and an observer, as sketched in Fig. 1.3. The source object at the background emits a ray that traverses the surroundings of a massive body, where it is deflected and is received by the observer at an angle $\boldsymbol{\theta}$, different than the angle $\boldsymbol{\beta}$ subtended

by the straight line that connects the observer with the source. For small deflection angles, the phenomenon can be approximated to the problem in optics of a single thin lens, governed by the lens equation

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \frac{\chi_s - \chi_l}{\chi_s}\hat{\boldsymbol{\alpha}} \equiv \boldsymbol{\theta} - \boldsymbol{\alpha}, \tag{1.25}$$

where χ_s and χ_l are the comoving distances to the source and the lens respectively and $\hat{\boldsymbol{\alpha}}$ is the deflection angle caused by the three-dimensional gravitational potential Φ of the lens

$$\hat{\boldsymbol{\alpha}} = -\frac{2}{c^2}\int \boldsymbol{\nabla}_\perp \Phi(\chi')\, d\chi' \,. \tag{1.26}$$

The integral is taken along the light path and the gradient perpendicular. In the last step of eq. 1.25 we defined the scaled deflection angle $\boldsymbol{\alpha}$. The expression 1.26 is valid not only for a single point source but for any gravitational potential arising from a matter distribution at the lens plane.

1.4.2 Effective lensing potential

It is useful to define the effective lensing potential as a projection of the three-dimensional Newtonian potential of the lenses

$$\Psi(\boldsymbol{\theta}, \chi) = \frac{2}{c^2}\int_0^\chi \frac{\chi - \chi'}{\chi\chi'}\Phi(\boldsymbol{\theta}, \chi')\, d\chi', \tag{1.27}$$

since then the scaled deflection is equal to its gradient in the plane of the sky

$$\boldsymbol{\alpha}(\chi) = \boldsymbol{\nabla}_\theta \Psi(\boldsymbol{\theta}, \chi). \tag{1.28}$$

Equation 1.27 is valid for an extended three-dimensional distribution of matter, that is, it accounts for the distortions induced by all lenses at any distance between the source and the observer. However, it assumes the Born approximation, in which integrals over the line-of-sight are computed along the unperturbed path $\boldsymbol{\theta}$ in Φ. This is accurate in most cosmological situations, since the deflection angle is typically small and we shall assume it hereafter.

1.4.3 Surface mass density

Next we define the convergence for a single lens plane (having a comoving width of $\Delta\chi$) as a dimensionless surface mass density

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{\Sigma(\boldsymbol{\theta}, \chi)}{\Sigma_{cr}(\chi)}, \tag{1.29}$$

where

$$\Sigma(\boldsymbol{\theta}, \chi) = \rho(\boldsymbol{\theta}, \chi)\Delta\chi \qquad \Sigma_{cr}(\chi) = \frac{c^2}{4\pi G}\frac{\chi}{\chi_l(\chi - \chi_l)}. \tag{1.30}$$

It is more convenient to convert the density in the last equation to the density contrast $\delta \equiv \rho/\bar{\rho} - 1$, where $\bar{\rho} = \Omega_o \rho_c = \Omega_0 {}^{3H_0^2}/_{8\pi G}$. Integrating all the lenses along the line-of-sight in the Born approximation, the convergence is

$$\kappa(\boldsymbol{\theta}, \chi) = \frac{3H_0^2\Omega_m}{2c^2} \int_0^{\chi} \delta(\boldsymbol{\theta}, \chi') \frac{(\chi - \chi')\chi'}{a\chi} \, d\chi' \,. \tag{1.31}$$

1.4.4 Linearized lensing quantities

At linear order, the mapping from image to source coordinates is given by the Jacobian matrix of the transformation

$$A_{ij} = \frac{\partial \beta_i}{\partial \theta_j} = \delta_{ij} - \frac{\partial \alpha_i}{\partial \theta_j} = \delta_{ij} - \partial_i \partial_j \Psi = \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix} \tag{1.32}$$

where the above equalities use the lens equation 1.25 and equation 1.28. The last equality introduces the convergence and shear as second derivatives of the potential

$$\kappa = \frac{1}{2} \Delta \Psi \tag{1.33}$$

$$\gamma_1 = \frac{1}{2}(\partial_1 \partial_1 - \partial_2 \partial_2)\Psi \qquad \gamma_2 = \partial_1 \partial_2 \Psi. \tag{1.34}$$

Note that the first expression is the equivalent of the Poison equation in two dimensions. Convergence is a scalar that causes isotropic distortions, while the shear is a spin-2 tensor that deforms circles into ellipses. The latter can be expressed as a complex number: $\gamma = \gamma_1 + i\gamma_2 = |\gamma| e^{2i\phi}$. In the case of galaxy surveys, the effect of the convergence is a variation of the sizes of objects, while the shear distorts the intrinsic ellipticity e of a galaxy into a measured one of $e + \gamma$. Furthermore, since surface brightness is conserved in weak lensing, there is a magnification effect that modifies the observed brightness by a factor $\mu = [\det A]^{-1} = ((1-\kappa^2) - |\gamma|^2)^{-1}$, that is due to the change in apparent area of the lensed object.

2

GALAXY AND HALO MOCK CATALOGUES

Our conception of the Universe has changed considerably in the last decades thanks to observational cosmology. The ΛCDM model (see §1.1) successfully explains most of the handful of observations available, but there are still unresolved puzzles such as the components of the dark sector. Ongoing surveys sample volumes one order of magnitude larger than what used to be one decade ago, and next generation of experiments will increase one order of magnitude further. It is said that we are entering the era of precision cosmology, in which new experiments will provide measurements with error-bars within the percent level.

Such tremendous narrowing of the statistical errors has to go along with an equivalent reduction of the systematics, otherwise the final error-bars on cosmological parameter will not scale down accordingly. A good theoretical modelling is essential to predict the signal that will be observed. Furthermore, in order to place constraints on cosmological parameters it is required to have accurate predictions for covariance matrices. Analytical prescriptions are unable to model the non-linear evolution (see §1.3.2) and numerical simulations become an invaluable tool, which allow to follow the growth of structure formation down to very small scales. Besides, mock galaxy catalogues can be produced from them and used to develop the science cases of future surveys. Mock catalogues are essential for the analysis of complex and huge astronomical datasets because they allow to optimize the pipelines, model and mitigate systematic effects, calibrate algorithms, test new techniques, etc.

However, conventional N-body codes are challenged by the requirement of sampling increasingly larger cosmological volumes while still resolving with very fine detail small

scales. This translates into very large simulations, including billions of particles, what in turn demands millions of CPU-hours to develop in current high performance computing platforms. Besides, an accurate estimation of covariance matrices demands the development of massive ensembles of simulations. For these many reasons, nowadays there is a growing interest in the so-called approximate methods, that overcome the problem of explicitly solving the non-linear evolution by incorporating a smart modelling of the small-scale physics, at a much lower computational cost.

This chapter begins in §2.1 with an overview of large-scale structure simulations by means of the common N-body method. Through the thesis, a particular simulation is taken as the fiducial value to benchmark results. It is the Marenostrum Institut de Ciències de l'Espai-Grand Challenge simulation (hereafter MICE-GC) a state-of-the-art cosmological simulation (Fosalba et al., 2015a; Crocce et al., 2015; Fosalba et al., 2015b). Next, §2.2 summarizes the existing approximate methods to produce mock catalogues. A very promising method among them, `COLA`, in which this thesis has been based, is introduced in §2.3. Lastly, §2.4 summarizes all the simulations that have been developed for this thesis.

2.1 N-body simulations of the large-scale structure

Cosmological simulations solve the evolution of a self-gravitating system within an expanding volume (for reviews, see Binney et al., 1987; Bertschinger, 1998; Dehnen et al., 2011). The expansion can be scaled out of the dynamics of the system and becomes irrelevant in the numerics. The problem is how forces are computed and integrated in time. Hydrodynamical simulations also account for friction forces due to collision terms in the right hand side of the Boltzmann equation 2.2 from the baryon component, which modifies the dynamics at cluster scales. But those demand much more computational resources and baryons can be neglected on large enough scales (of few Megaparsecs). Therefore, it is common to assume a dark-matter only system that is governed just by gravitational interactions.

A gravitational system that is affected by long-range interactions in which close encounters are inefficient for re-distributing kinetic energy is called collissionless. In that situation relaxation is not important, where the time needed to achieve relaxation in a system can be estimated as

$$t_{\rm relax} \simeq \frac{N}{8 \ln \Lambda} t_{\rm cross}, \tag{2.1}$$

where N is the number of particles, $t_{\rm cross}$ is the crossing time (the time required to cross the radius of the object at the typical velocity), and $\ln \Lambda = \ln(b_{\rm max}/b_{\rm min})$ is called the Coulomb Logarithm, which depends on the maximum and minimum impact parameter of the system[1]. for large N (as in N-body systems), the relaxation time becomes arbitrarily large.

[1]For a system of point particles, the minimum impact parameter is $b_{\rm min} = Gm/ < v^2 >$, where m is the mass of a single particle (S. D. M. White, 1976a).

The process of large-scale structure formation follows a collisionless evolution and, in fact, most of the haloes have a low relaxation level. The evolution of the density field in phase-space, $f(\boldsymbol{x}, \dot{\boldsymbol{x}}, t)$, is given by the collisionless Boltzmann equation

$$\frac{df}{dt} = \frac{\partial f}{\partial t} + \dot{\boldsymbol{x}} \cdot \frac{\partial f}{\partial \boldsymbol{x}} + \left(\frac{\partial \phi}{\partial \boldsymbol{x}}\right) \cdot \frac{\partial f}{\partial \dot{\boldsymbol{x}}} = 0 \tag{2.2}$$

together with the Poisson equation 1.17. The N-body method consists on sampling the phase space discretely with point-like particles that move within the simulated volume. This introduces Poisson noise due to the finite sampling, but is less computational intensive that solving the system in the six-dimensional phase space (although attempts have been made in that direction, e.g. Hahn et al., 2016). The evolution is solved by computing the force field and integrating discretely the equations of motion for particles in repeated time steps.

There are many N-body algorithms and differences are found mainly in how forces are computed. The gravitational potential of a collection of particle has the following expression

$$\phi(\boldsymbol{r}) = -G \sum_{i=1}^{N} \frac{m_i}{[(\boldsymbol{r} - \boldsymbol{r}_i)^2 + \epsilon^2]^{1/2}}, \tag{2.3}$$

where ϵ is the softening length, that it is necessary to suppress close pair encounters, where the potential would diverge. A direct summation approach (Peebles, 1970; S. D. M. White, 1976b; Efstathiou et al., 1981; Aarseth et al., 1979) means that this expression (or the analogue for the forces) is evaluated at each particle position in order to compute their forces, which results in a problem that scales as $\mathcal{O}(N^2)$, with N being the number of particles. This becomes prohibitive to compute as soon as there are few thousands particles in the simulation. For that reason, many numerical techniques have emerged that scale roughly as $\mathcal{O}(N \log(N))$. There are a collection of algorithms that are based on interpolating the particle distribution (or rather its density) to a grid, such as the Particle-Mesh (PM see § 2.3.1), the Particle-Particle Particle-Mesh (P^3M, Hockney et al., 1988; Efstathiou et al., 1985; Couchman, 1991). They use Fourier transforms in a cartesian grid to compute long-range interactions and in some cases are complemented with a direct summation at short distances. More sophisticated are multi-grid techniques such as the Adaptive Mesh Refinment (AMR), that adjust the resolution at each region of space depending on its dynamical state (Teyssier, 2002). On the other hand, the Tree algorithm computes forces by grouping particles that are close (Appel, 1985; Stadel, 2001). A successful model has been combining both the PM and the tree algorithms in a hybrid scheme to compute long- and short-range forces, as is the case of the Gadget code (Springel, 2005), which was used to develop the MICE-GC simulation (see § 2.1.1). Thanks to the improved software algorithms and the usage of the state-of-the-art hardware infrastructures, over the past years many large simulations have been developed with more than 10^{10} particles (Kim et al., 2011; Angulo et al., 2012; Alimi et al., 2012; Habib et al., 2012; Skillman et al., 2014; Heitmann et al., 2014; Fosalba et al., 2015a; for a review see Kuhlen et al., 2012).

2.1.1 The MICE-GC simulation

The Marenostrum Institut de Ciències de l'Espai-Grand Challenge (MICE-GC)[2] is a state-of-the-art N-body simulation. This thesis uses it to benchmark results obtained from the approximate method COLA and next is described briefly.

MICE-GC evolved 4096^3 particles in a volume of $(3072\,h^{-1}\,\mathrm{Mpc})^3$ using the GADGET-2 code (Springel, 2005) assuming a flat ΛCDM cosmology with $\Omega_m = 0.25$, $\Omega_\Lambda = 0.75$, $\Omega_b = 0.044$, $n_s = 0.95$, $\sigma_8 = 0.8$ and $h = 0.7$, that is consistent with the best fit to WMAP 5-year data (Dunkley et al., 2009). This results in a particle mass of $2.93 \times 10^{10}\,h^{-1}\,\mathrm{M}_\odot$. The initial conditions were generated at $z_i = 100$ using the Zel'dovich approximation (Zel'dovich, 1970) and a linear power spectrum generated with CAMB[3].

This simulation and its products have been extensively validated. The dark-matter and halo outputs are described in Fosalba et al., 2015a and Crocce et al., 2015. In addition, lensing maps are described in Fosalba et al., 2015b while Carretero et al., 2015 and Hoffmann et al., 2015 detail the HOD implementation used to produce galaxies mocks and the higher-order clustering, respectively.

The work in this thesis uses dark matter and halo catalogues of comoving outputs at $z = 0, 0.5, 1.0$ and 1.5. Haloes were identified using a Friends-of-Friends (FoF) algorithm (Davis et al., 1985) with a linking length of 0.2 (in units of the mean particle separation).

It is known that long-lived transients from the initial conditions affect the abundance of massive haloes and the clustering towards small scales, even for high starting redshifts if the Zel'dovich approximation is used (Scoccimarro, 1998). To investigate and correct such effects, an additional set of dedicated GADGET-2 N-body simulations were run. This is discussed in Appendix A for comoving catalogues. Results in Chapter 3 take into account these corrections, which are $\lesssim 2\%$ for 2-point matter clustering and $2-5\%$ for halo abundance on the regime and redshifts of interest. The results showed that transient effects are below the 1 per cent level for halo clustering, so the correction was safely neglected for these measurements.

2.2 Approximate methods for producing mock catalogues

Approximate methods are becoming a key tool in the observational cosmology community, as outlined in the introduction of this chapter (for a recent review of the theoretical foundations, history and a comparison of all the recent methods see Monaco, 2016). Most of them have in common that small-scale physics are unresolved (which represent the most expensive computations in an ordinary N-body simulation, thereby the speed-up), but this lack of information is supplied by introducing a motivated theoretical modelling. This is possible because non-linearities are restricted basically within haloes and, once these have collapsed, they display a more or less uniform profile. Their spatial distribution is connected with the large-scale density field (see § 1.3.3), and a simple biasing prescription might be able to predict one distribution from the other.

[2] More information is available at http://www.ice.cat/mice.
[3] http://camb.info

Instead, COLA is a quasi N-body that solves approximately the dynamics at small scales and incorporates a modelling that assures recovering the correct evolution at large and linear scales. The internal halo sub-structure is not accurately resolved, although it is capable of reliably predicting halo mean properties.

A convenient approach to reproduce the observed galaxy distribution on large scales is to populate dark matter haloes with galaxies using either Semi-Analytic Models or techniques relying on the Halo model, such as the Halo Occupation distribution or the Sub-halo Abundance Matching (see Knebe et al., 2015, and references therein for a comparison of models). Some of these techniques do not need to fully capture internal halo substructure since they only use reliable positions and mass estimates for haloes, which can be accurately predicted by approximate methods. A simplified and cheaper evolution of the density field is enough to approximately predict the abundance and clustering of collapsed regions to build halo mock catalogues. Available methods that are based on this idea include: PINOCCHIO (PINpointing Orbit-Crossing Collapsed HIerarchical Objects, Monaco et al., 2002; Monaco et al., 2013; Munari et al., 2016), PThaloes (Scoccimarro et al., 2002; Manera et al., 2015) and recently COLA (COmoving Lagrangian Acceleration, Tassev et al., 2013; Tassev et al., 2015).

An alternative approach is to use prescriptions to assign haloes in a density field produced by a simple gravity solver. This is the case of the log-normal model (Coles et al., 1991), QPM (Quick Particle Mesh, M. White et al., 2014), PATCHY (PerturbAtion Theory Catalog generator of Halo and galaxY distributions, Kitaura et al., 2015), EZmocks (effective Zel'dovich approximation mock catalogues, Chuang et al., 2015), and HALOGEN (Avila et al., 2015). Thus these methods constitute a more direct modelling for the galaxy distribution. The drawback is that they contain many internal parameters describing properties such as the galaxy clustering and abundance that have to be fit in order to correctly reproduce observations. COLA (see § 2.3) can be categorized as a quasi-N-body method and therefore has higher computational requirements than other methodologies. But it has no internal parameters to tweak beyond a generic optimization of code parameters to trade speed versus accuracy. It is more predictive and yields accurate high order clustering statistics. This has motivated the development of ICE-COLA, to extend its capabilities to light cone catalogues and to the modelling of weak lensing.

2.3 The COLA method

COmoving Lagrangian Acceleration (COLA, Tassev et al., 2013; Tassev et al., 2015), is a novel method for producing fast and approximate N-body simulations for large-scale structure formation of the Universe. A perturbative approach is used to obtain a first approximate solution to the dynamics, which serve as the basis for solving easily a more accurate solution though a simplified numerical integration.

A common feature in most of the approximate methods is that they evolve mass particles using Lagrangian Perturbation Theory (LPT). COLA is unique because on top of the analytical trajectory it adds a residual displacement computed by an N-body solver.

Equations of motion are solved in a frame comoving with LPT observers which, at a given perturbative order, encodes more non-linear growth information than the corresponding Eulerian approach. This guarantees an accurate description of the dynamics on large scales where the evolution is quasi-linear. The numerical evolution is simplified with respect to full N-body codes, making use of a fine Particle-Mesh (PM) algorithm (see § 2.3.1), and evaluating forces for a few (i.e, order of ten) time steps. Haloes can then be identified running a group finder in the evolved dark matter particle distribution, in the same way that it is done for an N-body.

The method decomposes the displacement field $\boldsymbol{x}$ into two terms: $\boldsymbol{x}_{\rm LPT}$ describes the LPT trajectory (see equation 1.21) and $\boldsymbol{x}_{\rm res}$ is the residual displacement with respect to the LPT path

$$\boldsymbol{x}_{\rm res}(t) \equiv \boldsymbol{x}(t) - \boldsymbol{x}_{\rm LPT}(t). \tag{2.4}$$

The equation of motion in a pure gravitational simulation relates the acceleration to the Newtonian potential Φ: $\partial_t^2 \boldsymbol{x}(t) = -\nabla\Phi(t)$. Therefore, equation 2.4 can be rewritten as

$$\partial_t^2 \boldsymbol{x}_{\rm res}(t) = -\nabla\Phi(t) - \partial_t^2 \boldsymbol{x}_{\rm LPT}(t), \tag{2.5}$$

where Φ is evaluated at the Eulerian position $\boldsymbol{x} = \boldsymbol{x} + \boldsymbol{q}$ (see § 1.3.1) by the PM algorithm, that solves the Poisson equation 1.17. COLA discretizes the time derivatives only on the left-hand side (see § 2.3.2), while uses the LPT expression at the right-hand side (see § 2.3.1). More recently, Tassev et al., 2015 extended this reformulation to the Eulerian domain, which allow the simulation of a sub-volume embedded in a larger effective simulation box.

COLA takes the advantage that the LPT displacement field is already computed at the beginning of a simulation when setting the initial conditions. Instead of dropping it, it is kept during the whole run time and simply re-scaled by the growth factors in order to integrate the equations of motion (see § 2.3.2).

2.3.1 The Particle-Mesh method

In the PM method, the density is estimated in a cartesian grid from the particle positions with a mass assignment function (see Hockney et al., 1988). Then, the Poisson equation 1.17 is solved in Fourier space, where derivatives become simple multiplications

$$\boldsymbol{k}^2 \tilde{\phi}(\boldsymbol{k}) = 4\pi G \tilde{\rho}(\boldsymbol{k}). \tag{2.6}$$

Forces are then computed from the potential and are interpolated from the grid nodes to the position of the particles.

The method uses Fast Fourier Transform (FFT) in two steps: just before and after solving the Poisson equation. Indeed, this represents the most expensive part of the method. Another consequence of this is that it results in assuming periodic boundary

conditions. Simulations can take advantage of this and build box replicas in adjacent positions (see § 4.2.2).

The cell size of the mesh determines the spatial resolution of the simulation. Tassev et al. (2013) suggest to use a mesh in COLA that is finer than the mean inter-particle distance. We shall refer to this scaled distance as the PM_{grid} factor, which these authors suggest to set to ~ 3 in order to adequately resolve small-mass haloes. Hence, in total there are 3^3 more mesh cells than particles (the total number of cells is thus the $\mathrm{number\,particles} \times \mathrm{PM}^3_{grid}$).

2.3.2 Temporal Drift and Kick operators

Once forces have been computed with the PM method (see § 2.3.1), the leapfrog integrator is used to solve the equation of motion of particles. It is a second order method that alternates the so-called drift (D) and kick (K) operators, that update positions and velocities respectively. In its standard form, these operators are:

$$D(a_i, a_{i+1}): \quad \boldsymbol{s}(a_i) \mapsto \boldsymbol{s}(a_{i+1}) = \boldsymbol{s}(a_i) + \boldsymbol{v}(a_{i+1/2})\Delta t \tag{2.7}$$

$$K(a_{i+1/2}, a_{i+3/2}): \quad \boldsymbol{v}(a_{i+1/2}) \mapsto \boldsymbol{v}(a_{i+3/2}) = \boldsymbol{v}(a_{i+1/2}) + \frac{1}{2}(-\nabla\phi(a_{i+1}))\Delta t. \tag{2.8}$$

In COLA, the operators are re-written, separating the contribution coming from LPT and the numerical integration (the residual component). Velocities contain only the residual component and are updated with the residual acceleration, while positions are found integrating these velocities and adding the LPT displacement:

$$\begin{aligned} D(a_i, a_{i+1}): \quad & \boldsymbol{s}(a_i) \mapsto \boldsymbol{s}(a_{i+1}) = \boldsymbol{s}(a_i) + \boldsymbol{v}(a_{i+1/2})\Delta t \\ & + [D_1(a_{i+1}) - D_1(a_i)]\boldsymbol{s}_1 + [D_2(a_{i+1}) - D_2(a_i)]\boldsymbol{s}_2 \end{aligned} \tag{2.9}$$

$$\begin{aligned} K(a_{i+1/2}, a_{i+3/2}): \quad & \boldsymbol{v}(a_{i+1/2}) \mapsto \boldsymbol{v}(a_{i+3/2}) = \boldsymbol{v}(a_{i+1/2}) + \Delta t \times \\ & \left(-\frac{1}{2}\nabla\phi(a_{i+1}) - \partial_t^2\boldsymbol{s}_1(t) - \partial_t^2\boldsymbol{s}_2(t)\right). \end{aligned} \tag{2.10}$$

Then, the evolution of n time steps can be written as

$$L_+(a)\left(\prod_{i=0}^{n} K(a_{i+\frac{1}{2}}, a_{i+\frac{3}{2}})D(a_i, a_{i+1})\right)K(a_i, a_{i+\frac{1}{2}})L_-(a). \tag{2.11}$$

In this notation, the rightmost operator applies first. Note that at the beginning it is necessary to perform a kick of half time step. Before this, it is necessary to transform the initial conditions to the rest frame of LPT observers (where the residual velocity is zero at the start) and add it back at the end of the simulation to re-write the solution for a comoving Euler-space observer with the operators

$$L_\pm(a): \quad \boldsymbol{v}(a) \mapsto \boldsymbol{v}(a) = \boldsymbol{v}(a) \pm (\partial_t\boldsymbol{s}_1(t) + \partial_t\boldsymbol{s}_2(t)). \tag{2.12}$$

The procedure is fully described in the appendix A of Tassev et al., 2013, where it is also explained how COLA uses a modified discretization for integrating the temporal intervals, which is based on assuming an ansatz for the time dependence of the residual velocities.

2.3.3 The Parallel COLA code

Koda et al., 2015 developed a parallel version of COLA, suitable for the massive production of mock catalogues, which was used for modelling the BAO signal in the WiggleZ survey (Kazin et al., 2014). The code is named Parallel COLA and it is publicly available[4]. FFT's are computed by the Fastest Fourier Transform in the West (FFTW) package[5] in its Message Passing Interface (MPI)[6] version for distributed memory parallelization. Interpolations in the PM algorithm (see § 2.3.1) use the Cloud-in-Cell linear interpolation (Hockney et al., 1988). It includes both the generation of random Gaussian initial conditions using 2LPT (Crocce et al., 2006) and the Friends-of-Friends (FoF, see also § 4.3.3) halo finder (Davis et al., 1985) running on-the-fly (for comoving outputs).

ICE-COLA (Institut de Ciències de l'Espai-COmoving Lagrangian Acceleration), the simulation code presented in this thesis, is based on Parallel COLA. They differ in the new features that have been added to the former such as using arbitrary time stepping distributions, the generation of matter power spectra on-the-fly, the storage of halo catalogues in binary format and the parallel storage of the comoving matter density field interpolated onto large mesh grids using the Cloud-in-Cell assignment (it has been used with grids up to 1024^3 cells). But the main novelty is the implementation of the algorithm to generate many flavours of light cone catalogues on-the-fly, as it is described in Chapter 4 and validated in Chapters 5 and 6. These include two compressed data formats: the two-dimensional projected matter field in several concentric pixelated maps (see § 4.3.2) and halo catalogues (see § 4.3.3).

Recently, other groups have produced other parallel implementations. L-PICOLA[7] (see Howlett et al., 2015a) is similar to Parallel COLA in its technical details and has been used to generate mock galaxy catalogues for the SDSS-II Main Galaxy Sample (Ross et al., 2015; Howlett et al., 2015b). This group advocate for low values for PM_{grid}. It is also capable of producing light cone data, but as far as we know they seem to be restricted to the three-dimensional particle distribution, that demand huge disk storage and an intensive post-processing. FASTPM[8] (see Feng et al., 2016) is a parallel PM code that includes the COLA method as an option. It uses a distinct library for performing FFT's: the Parallel FFT (PFFT)[9], that has the advantage of using a two-dimensional domain decomposition. This allows a good scalability of the code to thousands of cores. Finally,

[4] https://github.com/junkoda/cola_halo
[5] http://www.fftw.org/
[6] https://www.mpi-forum.org/
[7] https://cullanhowlett.github.io/l-picola/
[8] https://github.com/rainwoodman/fastpm
[9] https://github.com/mpip/pfft

there is a multi-threaded code written in Python/Cython[10] dubbed `pyCOLA` written by the original authors of the `COLA` method that implements both the temporal and spatial domain decoupling of the dynamics (as explained in Tassev et al., 2013; Tassev et al., 2015 respectively).

2.4 Simulation developed for this thesis

A large suite of simulations has been produced for this thesis. Although the method is faster in terms of wall-clock time when compared to full resolution N-body simulations, it still requires large dedicated memory. Subsections below describe the simulation sets that were used in different chapters of this thesis. Most of them have some common characteristics that are explained next and are assumed, unless stated differently.

The simulation code used was `ICE-COLA`. The cosmological model, the input linear power spectrum and the particle mass are identical as MICE-GC (see §2.1.1, with some exceptions in the latter for exploring mass resolution effects). The fiducial set of runs contain 2048^3 particles in a box size of $1536\,h^{-1}\,\mathrm{Mpc}$, that is, a factor 8 smaller volume than MICE-GC. It provides nonetheless a very large cosmological volume. The default configuration assumes $\mathrm{PM}_{grid} = 3$, 40 time steps and a distribution of these linear in the scale factor starting at redshift $z_i = 19$. A typical run without light cone outputs (the performance of these is discussed in §4.4) uses 1024 cores, with maximum memory consumptions of 2.6 TB for $\mathrm{PM}_{grid} = 3$ and takes around 40 minutes wall-clock time for 40 time steps. This means that the CPU-time consumed is less than 1 khour, to be compared with the 3 Mhours that needed MICE-GC having 8 times more particles, which gives a speed-up factor between two and three orders of magnitude with respect to a full N-body simulation with the same number of particles.

2.4.1 Simulations for Chapter 3

The results presented in Chapter 3 were derived from comoving outputs of a set of `ICE-COLA` simulations whose code parameters are listed in Table 2.1. Parameters that were varied are: the PM_{grid} factor, the number of time steps, the time sampling and the initial redshift. Additionally, it was explored the effect of decreasing the mass resolution by reducing the number of particles while keeping the box size fixed. Or in the opposite direction, the box size was reduced while keeping constant the particle load for a better mass resolution. For some particular runs of interest additional realizations were produced in order to reduce sampling variance.

All runs use the same seed number for the generation of the initial conditions (except for those which add more realizations to the same parameter configuration), what cancels out cosmic variance between different simulations (but not with respect to MICE-GC, which uses a different box size). The data products from those simulations are the following comoving ouputs: FoF halo catalogue, matter field interpolated in a 1024^3 mesh,

[10] http://cython.org/

$N_{\text{realizations}}$	Particle number	L_{box}	PM_{grid}	N_{steps}	z_{i}	Time steps distribution
1	$\underline{1024}^3$	1536	3	10	9	$\propto a$
1	$\underline{1536}^3$	1536	3	10	9	$\propto a$
1	2048^3	1536	$\underline{2}$	10	9	$\propto a$
2	2048^3	1536	3	10	9	$\propto a$
1	2048^3	1536	3	$\underline{20}$	9	$\propto a$
1	2048^3	1536	3	$\underline{40}$	9	$\propto a$
1	2048^3	1536	$\underline{2}$	40	19	$\propto a$
48	2048^3	1536	3	40	19	$\propto a$
1	2048^3	$\underline{768}$	3	40	19	$\propto a$
8	2048^3	1536	3	40	$\underline{39}$	$\propto a$
1	2048^3	1536	3	40	$\underline{39}$	$\underline{\propto \log a}$
1	2048^3	1536	3	40	$\underline{39}$	$\underline{\propto a^{0.8}}$
1	2048^3	1536	3	40	$\underline{100}$	$\propto a$

TABLE 2.1: ICE-COLA code parameters used in the simulations for Chapter 3. The box sizes are in units of $h^{-1}\,\text{Mpc}$. Underlined are parameter values that are distinctive for a run and the highlighted row corresponds to the optimal set-up, which provides the best accuracy. For the latter, each realization needed 1024 cores during 40 minutes and 2.6Tb of memory. How to extrapolate those computational requirements to other parameter configurations is explained in § 3.1.

and for some runs the matter power spectrum. The redshifts outputted were 0, 0.5, 1 and in some cases also 1.5. This results in a typical data volume of 16 GB per run.

2.4.2 Light cone simulations

All the results from light cone simulations presented in this thesis (see Chapters 4, 5 and 6) derive from a single run using that new feature of `ICE-COLA`. This has already been augmented to 68 realizations in total, that are being analysed for ongoing projects on weak lensing covariance matrix estimation (as outlined in §7.2), and it is foreseen that more realizations will be produced in a future. The code parameters used are the fiducial ones explained in §2.4. An all-sky light cone was built starting at redshift $z_{LC} = 1.4$, which corresponds to a comoving distance two times larger than the box size (see equation 4.6). This generates 64 box replicas (see equation 4.7), i.e., two in each cartesian direction. The outputs that were generated consist on FoF halo catalogues in the light cone (for haloes with 20 or more particles, see §4.3.3), pixelated two-dimensional projected matter density maps (see §4.3.2) and the matter power spectrum at some selected redshifts (0,0.5, 1, 1.5 and 2). The data volume generated for each realization is roughly 100 GB. For more technical details on the numerical performance of these simulations with light cones, see §4.4.

Besides, ten additional realizations were generated with only particle light cone catalogues, with the same parameters as described before except that the volume was restricted to one octant of the sky. These are being used to generate projected matter maps at post-processing with different angular resolutions in order to asses its impact on weak lensing observables.

2.4.3 Other simulations

As described on §7.1, an on-going project consists on comparing covariance matrices estimated from the approximate methods `COLA` and `PINOCCHIO` with respect to a set of 100 N-body simulations, dubbed Minerva. The suites from approximate methods consist of 1000 realisations on a box of $1500\,h^{-1}\,\mathrm{Mpc}$ with a ΛCDM cosmology that matches the best-fitting analysis of WMAP9 and the correlation function from BOSS DR9 (Sánchez et al., 2013). `COLA` evolved 1536^3 particles while `PINOCCHIO` and Minerva 1000^3. The outputs produced include FoF halo catalogues at redshifts 0, 0.57, 1, 2, and the matter power spectrum at same redshifts. The total CPU-time needed was is 250 khrs. In addition, the set is completed with 100 extra realizations that use the same initial conditions and number of particles than the Minerva simulations. These are useful to cancel out cosmic variance. The latter required ten more khrs in total. These simulations produce 388 MB (92 MB) of output data per run for the ensembles of 1000 (100) realizations.

Besides, Appendix A and B use additional simulations to estimate transient effects and the performance of PM-only simulations respectively. The details are described there.

3

Optimization of a quasi N-body method for clustering

COLA is able to accurately determine the evolution of the matter field on scales above the typical size of dark matter haloes. However, reproducing the birth and growth of haloes is more challenging because they display high density contrasts and non-linear dynamics sustained by virialisation. Halo formation is very sensitive to the degree of approximation in the dynamics at small scales and a minimum accuracy is indispensable to generate reliable halo mock catalogues. Therefore, it is essential to assess the performance of the method under different values of the internal code parameters that describe the spatial and temporal discretization of the gravitational evolution.

This chapter uses a suite of large ICE-COLA simulations with comoving outputs, described in §2.4.1, and the N-body simulation MICE-GC as a reference (see §2.1.1, for an extensive validation see Fosalba et al., 2015a; Crocce et al., 2015; Fosalba et al., 2015b). We summarize the scaling of computational requirements of Parallel COLA in §3.1 and we explain the details for measuring observables in §3.2. In §3.3 we discuss the capabilities and limitations of COLA when run with as few as 10 time steps and §3.4 explores the impact of varying internal code parameters on basic observables such as the matter real-space power spectrum and the halo mass function. In particular the parameters explored are: the size of the force evaluation mesh, the number and distribution of time steps, and the initial redshift, in combination with mass resolution. The optimal code set-up is described in §3.4.5, showing that it yields the best accuracy in the power spectrum and the mass function (for a wide range of halo masses, comoving scales and

redshifts) without a large increase in computational cost. Next in § 3.5 we discuss the performance of the optimal parameter choice on halo clustering in real and redshift space. This procedure yields what can be regarded as the optimal accuracy of the code out of the box. To improve it further one needs to rely on simple corrections to halo masses using an external simulation, as it is shown in § 3.5.1 by matching the halo abundance and in § 3.5.2 by matching the halo clustering amplitude on large scales. In that way, deviations in halo clustering can be reduced to within the percent level in most situations. Besides, Appendix A shows how transient effects have been corrected in the simulations and Appendix B demonstrates the advantage of COLA over a PM-only simulation when there are as many as 40 time steps.

Finally, § 3.6 gives a comparison of the performance of COLA against other approximate methods in the literature.

The work of this chapter resulted to the following publication (Izard et al., 2016):

Towards fast and accurate synthetic galaxy catalogues optimizing a quasi N -body method, **Izard A.**, Crocce M., and Fosalba P., 2016, MNRAS, 459, 2327–2341, arXiv: 1509.04685.

3.1 Computational requirements and their scaling

This section explains how the run-time and memory consumption of a single COLA realization scale with code parameters, so that combined with the information provided in § 2.4.1 for the suite of simulations used one can extrapolate the numbers to other configurations.

The choice of the initial redshift and the time sampling distribution have no effects at all on the computational requirements. The run-time is largely dominated by the computations of FFTs during force evaluation at each step. For the default configuration and 40 time steps, the code spends only 10 per cent of the time in computations not related to the PM algorithm, such as the initial set up and I/O. For this reason, the run-time increases roughly linearly with the size of the FFT and the number of time steps. Such transforms scale as $\propto \mathcal{O}(n \log n)$, where n is the number of grid points, and since this is proportional to PM_{grid}^3, the run-time scales roughly as $\propto \mathrm{PM}_{grid}^3$ (for large numbers the $\log n$ factor can be neglected).

On the other hand, given a constant number of particles, the memory consumption depends only on the size of the force mesh. The allocation of memory from the PM part represents around 60 per cent for $\mathrm{PM}_{grid} = 3$ and it scales as PM_{grid}^3.

3.2 Power spectrum and mass function measurements

The matter and halo power spectra are determined interpolating the particles into a cubic grid of 1024^3 cells via a Cloud-in-Cell (CiC) assignment. We then obtain the density in k-space by doing a FFT and estimate the band power by averaging the square over the range of modes corresponding to a k-bin: $P(k) = \langle |\delta_{\boldsymbol{k}}|^2 \rangle$. The mass assignment into

Sample	Mass range $[\log(M/\,h^{-1}\,\mathrm{M}_\odot)]$
M1	$12.5 - 13.0$
M2	$13.0 - 13.5$
M3	≥ 13.5

TABLE 3.1: Halo mass samples used throughout this chapter, defined by halo mass, at $z = 0, 0.5$ and 1.

a finite size grid introduces a filtering artifact that we compensate by deconvolving the CiC window function, which in Fourier space is simply a division. We correct for aliasing effects due to the finite sampling as in Jing (2005). Lastly, monopole measurements of halo power spectrum are corrected for shot-noise assuming a Poisson distributed noise.

For haloes, we restrict the analysis to those having more than 100 particles, corresponding to $M \geq 10^{12.5}\,h^{-1}\,\mathrm{M}_\odot$ for most of the runs. Halo masses are defined using the Warren correction (Warren et al., 2006): $M = m_p N(1 - N^{-0.6})$, where m_p is the particle mass and N the number of particles. This is irrelevant for most of the runs that use the same mass resolution as MICE-GC, but provides better agreement for lower mass resolution runs. We build three halo samples according to the mass cuts listed in Table 3.1. Mass function measurements contain error bars estimated by Jack-knife re-sampling using 64 different cubic sub-volumes and only mass-bins whose relative error is less than 5% are shown (see e.g. Fig. 3.1).

3.3 Limitations of 10 time steps

The COLA method is designed to use very few time-steps, so that a high speed-up of more than two orders of magnitude with respect to a full N-body is achieved (Tassev et al., 2015). In this section we explicitly test the accuracy, as a function of scale and redshifts, of the original COLA configuration of 10 time steps with scale factor linearly distributed between redshift 9 and 0 and a PM_{grid} factor of 3.

With only 10 time steps, the matter power spectrum is accurate at the 5% level to $k \sim 0.5\,h\,\mathrm{Mpc}^{-1}$ (see § 3.4.1). This allows the exploration of non-linear scales, but only to some limited extent. At the halo level, however, the situation is more complex. Figure 3.1 shows that the mass function is severely underestimated at $z = 1$, especially at large masses. The problem is not so visible at $z = 0.5$, where the disagreement is at most at the 15% level and at $z = 0$ it is within the 5% for all masses. Likewise, we have checked that the halo bias is overestimated by as much as 20% at $z = 1$ and agrees within few percent level at lower redshifts. Both effects (underestimation of the mass function and overestimation of the halo bias) can be explained by a halo mass underestimation at high redshift, when the evolution has been computed with very few time steps.

We suggest that these problems could come from a higher difficulty of achieving relaxation inside haloes when the time discretization is too coarse. Particles evolve according to the mean gravitational potential that arises from the smooth distribution, but are also affected by individual encounters. The relaxation time is related to the moment when

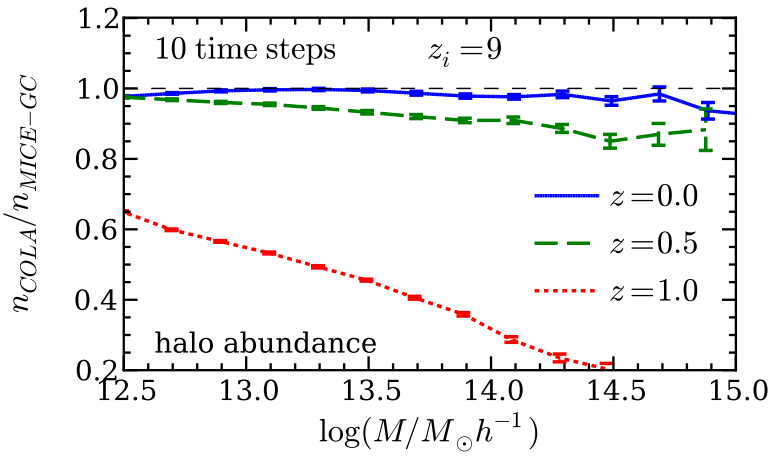

FIGURE 3.1: Mass function when using COLA with 10 time steps starting at $z_i = 9$. Solid, dashed and dotted lines correspond to $z = 0$, 0.5 and 1 respectively. At high redshift there is a large discrepancy that is partially solved at $z = 0.5$.

the latter start to significantly contribute to the dynamics, boosting the re-distribution of kinetic energy and achieving a dynamical equilibrium in the system (Dehnen et al., 2011).

Each time step is a chance for particles to interact with each other, but if we reduce them drastically the re-distribution of energy is unphysical suppressed. This is critical for those haloes that have not relaxed yet. Since the relaxation time is proportional to the number of particles of a halo (Binney et al., 1987), the effect is larger for high mass haloes. The halo formation time increases and merging processes are poorly captured, producing halo masses largely underestimated for $z > 0$, before 10 time steps have been completed. Note that full N-body codes with adaptive time steps schemes trigger finer time samplings at high density regions and halo formation is properly tracked.

To visually confirm this idea, we show in Fig. 3.2 the halo distribution of two runs with 10 and 40 time steps in red open and blue filled circles, respectively. The initial conditions and the rest of parameters are kept the same, so that differences are due only to the number of time steps. In the left panel, we confirm that at $z = 1.0$ massive haloes are in general not properly tracked with 10 time steps and they seem to appear fragmented as smaller mass haloes. And not all of the low mass haloes are identified. Nevertheless, in the right panel at $z = 0$, when 10 time steps have been completed, the agreement is much better on halo masses, positions and abundance at all masses. This visual inspection suggests that one needs $\gtrsim 10$ time steps before halo properties converge at the redshift of interest.

We find that relaxation effects when only 10 time steps are used can be reduced using a higher particle mass (e.g. above $10^{11}\,h^{-1}\,\mathrm{M}_\odot$). The number of particles of the haloes decreases and, therefore, their relaxation time as well. We checked that disagreements in the mass function and halo clustering are indeed lower, but this apparent improvement is however lost for other statistics where mass resolution is important.

Evolving particles with just ten time steps before the redshift of interest, therefore,

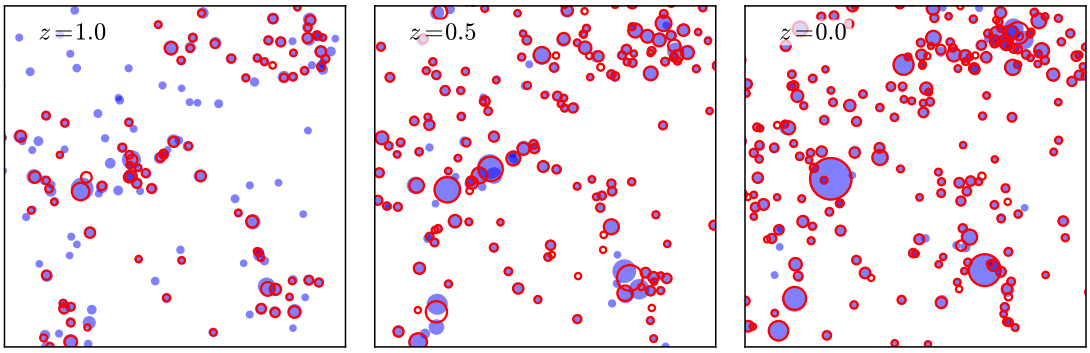

FIGURE 3.2: Spatial distribution of haloes in two different COLA runs that differ only in the number of time steps: 10 are represented by open circles and 40 by filled ones. Different panels show redshifts 1.0, 0.5 and 0 from left to right. The slices have a width of $50\,h^{-1}\,\mathrm{Mpc}$ and are $25\,h^{-1}\,\mathrm{Mpc}$ thick. The radii of circles are proportional to $M^{1/3}$ and match the r_{200} values, so that they reflect the physical size of haloes. Only those with more than 200 particles are shown (which corresponds to $6 \times 10^{12}\,h^{-1}\,\mathrm{M}_\odot$). The largest halo at $z = 0$ has a mass of $1.87 \times 10^{15}\,h^{-1}\,\mathrm{M}_\odot$ for a run with 40 time steps and a 4.5% less for 10 time steps. At $z = 1$ the matching between the runs is poor, with the abundance under-estimated by 50% or more with 10 time-steps. The agreement improves as one approaches $z = 0$.

provides accurate results only at large scales ($k \leq 0.3\,h\,\mathrm{Mpc}^{-1}$) and low redshifts. We might have stronger requirements that clearly push to go beyond 10 time steps to surpass these limitations.

3.4 Optimization

A gravity solver algorithm discretizes both temporal and spatial dimensions (and the mass as well) in order to numerically solve for the dynamical evolution. The idea behind COLA is to reduce numerical computations as much as possible while still capturing the growth of structure on large scales (see § 2.3). This can be controlled with few internal code parameters: the number of time steps, the time sampling distribution, the initial redshift and the size of the force mesh grid, in combination with the mass resolution and/or the box size. Note however that COLA is not fundamentally different from a full N-body in the sense that as one increases the requirements on such parameters the numerical integration of particle trajectories becomes more accurate and COLA would eventually converge to a full N-body.

In this section we explore the code parameter space in order to understand their impact on observables and determine regions that provide optimal results in terms of accuracy versus running time (or memory usage). We assess the performance by comparing the COLA dark matter power spectra up to $k \sim 1\,h\,\mathrm{Mpc}^{-1}$, and halo mass functions for $M \gtrsim 10^{12.5}\,h^{-1}\,\mathrm{M}_\odot$, against those in the reference N-body run. A key difference from previous works (Kazin et al., 2014; Howlett et al., 2015b; Leclercq et al., 2015) is that we aim at reproducing those observables simultaneously across a large redshift range

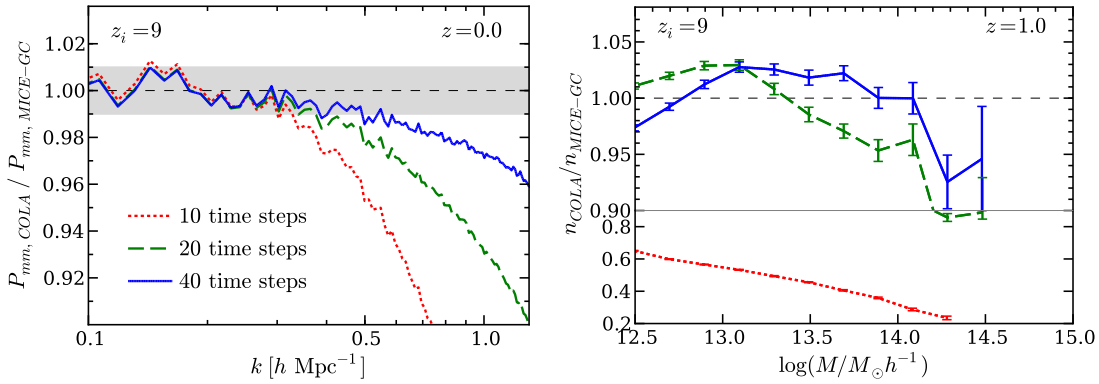

FIGURE 3.3: Matter power spectrum in real space at $z = 0$ (left panel) and mass function at $z = 1.0$ (right panel) for three different choices for the number of time steps: 10, 20 and 40 in dotted, dashed and solid lines, respectively. The initial redshift is fixed at $z_i = 9$. An increase in the number of time steps directly translates into a better accuracy at small scales. As for haloes, a high number of time steps is necessary to correctly predict the mass function at high redshift.

($0 \leq z \leq 1$). As we will see next, the requirements in this case are more stringent that those needed for a single redshift output or halo mass bin.

3.4.1 Number of time steps

The first parameter we vary is the number of time steps, with the initial redshift fixed to 9. The upper panel of Fig. 3.3 displays the $z = 0$ matter power spectrum for COLA runs with increasing number of time steps divided by the one measured in MICE-GC. The characteristic scale at which the ratio deviates from unity is progressively shifted towards higher wave-numbers as more time steps are included: there is a 2 per cent agreement up to scales of $k \sim 0.4$, 0.6 and 0.8 $h\,\mathrm{Mpc}^{-1}$ for 10, 20 and 40 time steps respectively. This can still be improved adjusting the initial redshift according to the number of time steps (see § 3.4.3). In particular, we check that doubling the number of time steps almost doubles the characteristic wavenumber where the power spectrum is significantly underestimated. This in turn means that there is room for higher accuracies with more than 40 time steps, although presumably the force mesh resolution would then soon become a limiting factor. We also find that these results for the matter power spectrum are to a good extent independent of the redshift analyzed.

However, for the mass function there is a higher sensitivity to the number of time steps at high redshift. As shown in the right panel of Fig. 3.3, the large underestimation of the mass function at $z = 1$ is solved by doubling the total number of time steps. With 20 time steps, 10 of them are computed before the redshift of interest. The abundance at high masses further increases by ~ 5 per cent when moving from 20 to 40 time steps and the mass function is consistent with the one from MICE-GC.

At redshift 0 and 0.5 the mass function also increases with the number of time steps at masses above $\sim 10^{13.5}\,h^{-1}\,\mathrm{M}_\odot$, although more moderately. Moving from 10 to 20 time

steps, the mass function augments by $5-10\%$ at the high mass range and from 20 to 40 by $\lesssim 5\%$. At low masses, differences remain within 1 per cent for 20 and 40 time steps. We conclude, therefore, that the low mass regime of the mass function converges for 20 time steps but that 40 are necessary for the most massive haloes. In Appendix B we show that, even with as many steps as 20 or 40, the 2LPT contribution in the COLA method is still key to achieve accurate results, as compared to PM only simulations.

3.4.2 Time sampling distribution

The scale factor a is the variable used to discretize the temporal variable in regular time steps. For that, we can choose a time sampling function $f(a)$ and distribute n steps in intervals of constant width $\Delta f(a)$,

$$\Delta f(a) = \frac{f(a_f) - f(a_i)}{n-1}, \tag{3.1}$$

where a_i is the initial scale factor and $a_f \equiv 1$ the final. If the resulting $\Delta a << 1$ then $\Delta a \approx [f'(a)]^{-1}\Delta f(a)$. For the linear case, we simply have $f(a) = a$ and the step width Δa is constant. We can define the step density as the inverse of the step width: $\rho \equiv 1/\Delta a$. Since $\Delta f(a)$ is constant, then $\rho \propto f'(a)$.

In §3.4.1 we showed for the linear case how increasing the step density (which in that case is only set by the number of time-steps) improves the accuracy of the simulation, but at the expense of a higher computational cost. In this section we explore which function f, and the corresponding step density $\rho \propto f'(a)$, produce a step distribution that is balanced in terms of accumulated errors over time. In Fig. 3.4 we show the step density for four different choices of the time sampling function. We distribute 40 time steps between $z_i = 39$ and $z = 0$ using a linear (circles), logarithmic (squares) or power law function $f(a) = a^p$ (diamonds and triangles for $p = 0.5$ and $p = 0.8$ respectively).

A logarithmic time sampling is useful in full N-body codes for global time steps that affect all particles. But these algorithms in general refine the global time stepping as the simulation evolves and also implement adaptive time stepping schemes for individual particles that sample more accurately the time evolution when non-linearities start to grow. In the implementation of COLA we are using there is no such refinement at late times and we see in Fig. 3.4 that the logarithmic choice oversample the early evolution of the particle distribution at the price of a low step density at later times. On the other hand, the linear case has large relative variations on the scale factor during the first steps, which might presumably lead to larger inaccuracies. For that reason we have considered the power law function to sample intermediate situations if $0 < p < 1$.

In Fig. 3.5 we compare the matter power spectrum for three time sampling functions: linear, $a^{0.8}$ and logarithmic in solid, dashed and dotted lines respectively. Upper and lower panels correspond to redshifts 0 and 1 respectively. We see that any gain we might have at high redshift by concentrating there more time steps is lost as soon as the step density decreases later. This is evident for the logarithmic case, which has the lowest

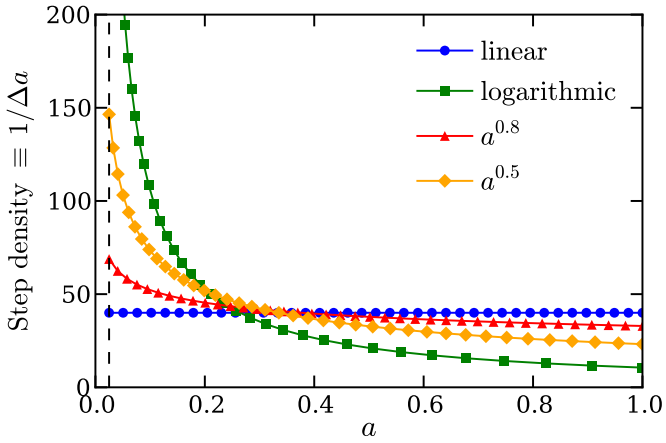

FIGURE 3.4: Step density (or the inverse of the step width) as a function of the scale factor for different schemes of time sampling: linear (circles), logarithmic (squares) and using a power of the scale factor, where for the latter we show two cases with exponents 0.5 (diamonds) and 0.8 (triangles). In all cases we distribute 40 steps between $z_i = 39$ and $z = 0$ and the markers are located at the position of each step.

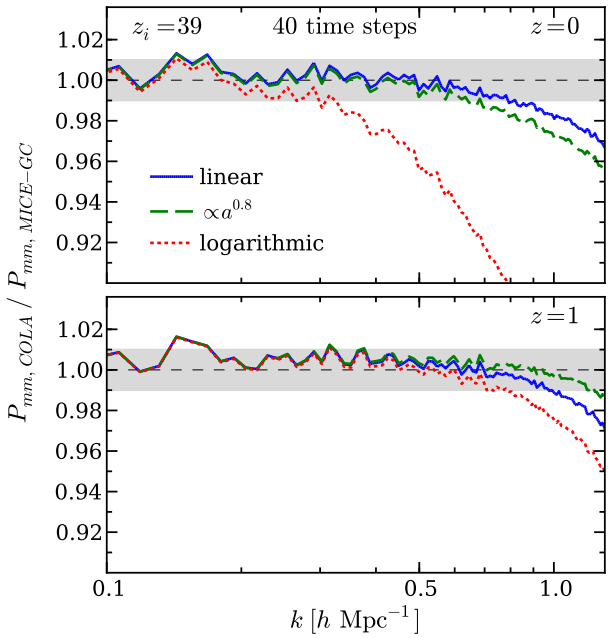

FIGURE 3.5: Matter power spectrum for the following time sampling functions: linear (solid line), $a^{0.8}$ (dashed) and logarithmic (dotted). All runs contain 40 time steps starting at $z_i = 39$. The upper and lower panels show redshifts 0 and 1, respectively. The case that performs better at all redshifts is the linear one.

step density for $z < 2$. In particular, at $z = 0$ the power spectrum is close to that for the case of 10 time steps linearly distributed (see dotted line in the upper panel of Fig. 3.3) and indeed they have a similar step density. The distribution of time steps using the function $a^{0.8}$ provides better results at high redshift, but falls below the linear distribution for lower redshifts where it has a lower step density.

Therefore, an optimal distribution should have a step density without strong variations and we conclude from the measurements that the linear case offers the best global performance at all redshifts. Although large relative variations on the scale factor during the first steps could lead to inaccuracies, this is balanced by the fact that at early times the dynamics is close to linear and can be well approximated by the 2LPT evolution. Hence a better time sampling at the beginning is not as critical as for low-z. Since all these arguments are built based on relative variations of the step density across time, *the conclusions are independent of the absolute number of time steps and the initial redshift*. Therefore in the rest of this chapter (and the thesis as well) we shall adopt a linear time-stepping distribution.

Lastly, we can use the concept of step density to frame the results from § 3.4.1 for the abundance of haloes. We find that whenever the step density is low ($\rho \lesssim 20$), the mass function suffers an underestimation for masses above $10^{13.5}\, h^{-1}\,\mathrm{M}_\odot$.

3.4.3 Initial redshift

The optimal selection of the initial redshift is coupled with the number of time steps (but it is independent of the time sampling distribution, as we have already shown).

A first guess we can do is to set the initial scale factor equal to the step width, which for 10 time steps gives $z_i = 9$ (Tassev et al., 2013). Starting later would introduce more transient effects whereas doing it earlier would produce large relative variations in the scale factor in the first time step, which seems not optimal. So there is not much room for optimization using few time steps. Instead we now focus on the situation in which we have 40 time steps.

Using the same rule as before we can estimate a good guess as $z_i = 39$. In Fig. 3.6 we show the resulting matter power spectrum at $z = 0$ when the initial redshift is varied from 9 to 100. A low value of $z_i = 9$ yields transient effects at all wave-numbers with an amplitude up to one per cent. The rest of cases are almost indistinguishable, only for $z_i = 100$ there is slightly less power at $k \sim 1\, h\,\mathrm{Mpc}^{-1}$. On the other hand, we detect that the mass function is underestimated at low masses if the initial redshift is too high. One possible explanation is that for high starting redshift the density contrast in the initial conditions are too smooth and then, due to the coarse time sampling in `COLA`, the smallest density peaks that seed small mass haloes are blurred. This pushes towards using an initial redshift not as high so fluctuations are larger, even if they are given by 2nd order in LPT.

Nonetheless the dependence we observe on this parameter once $z_i \sim 20 - 40$ is weak, since in practice only affects the position of the first(s) time step. But given the slightly better performance on the mass function of $z_i = 19$ when using 40 time steps, we adopt this value as our fiducial choice in what follows.

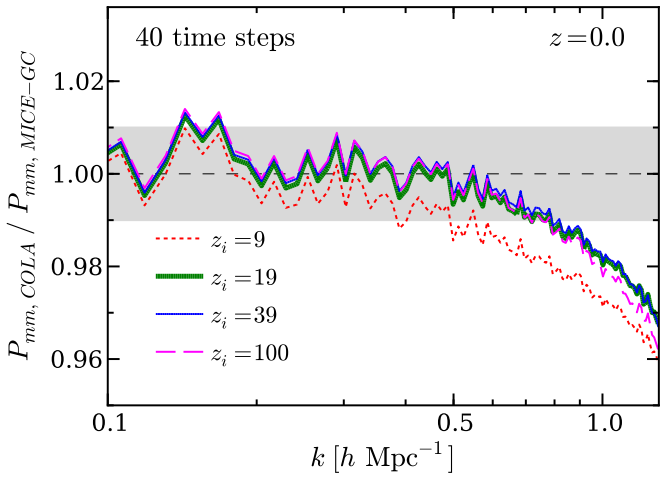

FIGURE 3.6: Matter power spectrum at $z = 0$ as a function of the initial redshift. All runs distribute 40 time steps linearly along the scale factor. A too low starting redshift produces transients at non-linear scales. The cases with $z_i \gtrsim 19$ are almost indistinguishable from each other, except for a slightly less power for $z_i = 100$ at $k \sim 1\,h\,\mathrm{Mpc}^{-1}$ (see text for details).

3.4.4 Force mesh grid size

Previous subsections were devoted to parameters that define the temporal discretization of the simulation. We use now the most accurate configuration (that is, 40 time steps linearly distributed from $z_i = 19$) to study the effects of the spatial discretization in the force computations. In particular, we compare runs with PM_{grid} factors of 2 and 3. We note that this parameter is of particular relevance as it has a large impact in the computational cost of the runs. For example, $\mathrm{PM}_{grid} = 2$ allows a saving of 70% of the computing time and 30% of the memory consumption with respect to $\mathrm{PM}_{grid} = 3$.

We find that changing from $\mathrm{PM}_{grid} = 3$ to 2 only changes the matter power spectrum by $\sim 1\%$ for $k \lesssim 1\,h\,\mathrm{Mpc}^{-1}$. However there is a more important effect on the halo mass function. For $\mathrm{PM}_{grid} = 3$, what corresponds to a comoving cell size of $0.25\,h^{-1}\,\mathrm{Mpc}$ given our box-size and particle load, we recover the mass function within 5% of the one measured in MICE-GC down to $10^{12.5}\,h^{-1}\,\mathrm{M}_\odot$. This is shown in Fig. 3.7 with a solid blue line. The mass scale where the mass function deviates at low masses is related to the spatial resolution of the mesh. If the latter is degraded by a factor of 1.5 when PM_{grid} is changed from 3 to 2, we expect a shift of that mass scale by a factor of $1.5^3 \simeq 3$. For that reasoning, we assume that the mass of a halo scales as the third power of the radius, or equivalently, that the density inside a halo is constant, in agreement to the halo model (Scoccimarro et al., 2001). Indeed, for $\mathrm{PM}_{grid} = 2$, the mass function is within $\sim 5\%$ for haloes of $M \sim 10^{13}\,h^{-1}\,\mathrm{M}_\odot \simeq 3 \times 10^{12.5}\,h^{-1}\,\mathrm{M}_\odot$, in agreement with the expectations.

A discrepancy in the mass function might have two sources: a genuine difference on the abundance or that halo mass estimates are systematically biased. The first case, and assuming that the difference is spatially homogeneous, does not produce differences in

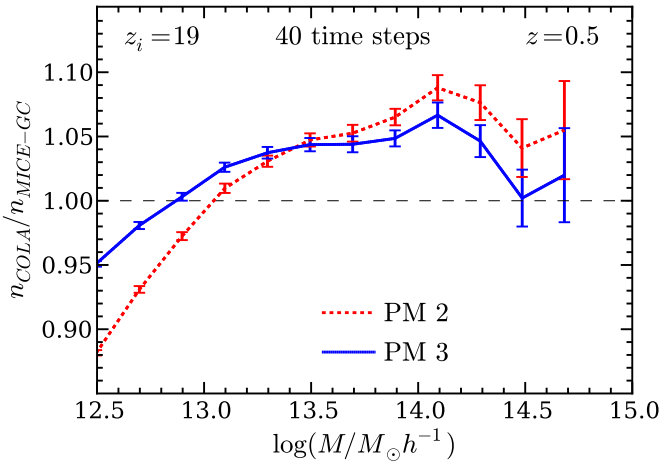

FIGURE 3.7: Mass function at $z = 0.5$ using a $\mathrm{PM}_{grid} = 2$ and 3 (dotted and solid lines, respectively). Both runs contain 40 time steps and have $z_i = 19$. Decreasing the size of the PM grid produces a larger overestimation of masses for large haloes and increases the incompleteness for small haloes.

clustering for samples selected by mass cuts (see Table 3.1), while the second does. Since we do not detect any significant difference in halo clustering at the low mass range for different PM_{grid} factors, we infer that there is a completeness problem at those masses due to the size of the force mesh. Not all haloes that should form are detected in the simulation, and in a mass-dependent way.

At high masses, on the contrary, we observe a lower clustering amplitude ($\sim$ 1 per cent at linear scales) for the run that produces a higher overestimation on the mass function. Both facts can be explained by a halo mass overestimation. One possible interpretation is that the puffier the haloes due to the force resolution, more easily the FoF algorithm bridges neighboring particles or small groups to a halo that really do not belong to it, hence systematically biasing high the mass estimate, as we observe.

3.4.5 Optimal setup

So far, we have given an exploration of the main code parameters in `COLA` and their impact on the dark matter clustering and on halo abundance. To achieve percent accuracy on both quantities, at the very least 10 time steps have to be done before the redshift of interest[1], which means that in total we might need 20 or more until $z = 0$. The more we do, the higher is the wavenumber where the dark matter power spectrum starts to miss power. This is true to at least $\sim$ 40 time steps, above that one should probably set $\mathrm{PM}_{grid} > 3$ (for the reference mass resolution we use, $2.9 \times 10^{10}\, h^{-1}\,\mathrm{M}_\odot$), so that the force resolution does not limit the accurate sampling of power up to $k \sim 1\, h\,\mathrm{Mpc}^{-1}$. In Appendix B we show that still with these configurations, the `COLA` method yields better results than a PM only evolution. The optimal time sampling distribution is given by

[1]Using a particle mass of $2.9 \times 10^{10}\, h^{-1}\,\mathrm{M}_\odot$.

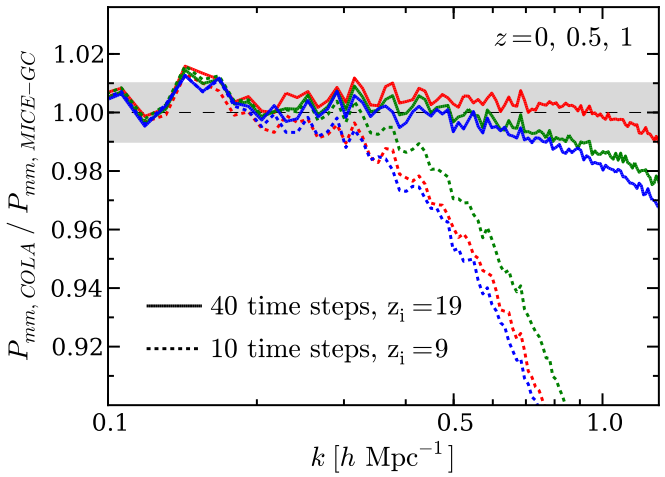

FIGURE 3.8: Comparison of the matter power spectrum in real space for the run using optimal parameters (40 time steps and $z_i = 19$, solid lines) and the default configuration (10 time steps and $z_i = 9$, dotted lines). We show redshifts 0, 0.5 and 1. This optimal setup delivers a ~ 1 per cent accuracy at scales $k \sim 1\,h\,\mathrm{Mpc}^{-1}$.

spacing steps linearly in the scale factor, regardless of the rest of parameters, and for a large number of time steps best results are achieved with an initial redshift of 19. A high PM_{grid} factor is required for percent accuracy in halo abundance and matter clustering and thus we set $\mathrm{PM}_{grid} = 3$ despite its relatively higher computational cost. The prize of further increasing it is not well justified in terms of additional accuracy.

A good choice of code parameters depends on which accuracy requirements need to be accomplished by the final mock catalogues. In this work, our target is to achieve per cent level accuracy on the matter power spectrum and halo abundance, in the wide ranges of $10^{12.5} - 10^{15.0}\,h^{-1}\,\mathrm{M}_\odot$ in mass, scales up to $k \sim 1\,h\,\mathrm{Mpc}^{-1}$ and redshifts comprised between 0 and 1.5.

Given these requirements we find that the best setup is set by 40 time steps linearly distributed along the scale factor, starting at $z_i = 19$ and with a PM_{grid} factor of 3. Fig. 3.8 shows the matter power spectrum using this configuration (solid lines) compared with the case of 10 time steps starting at $z_i = 9$ (dotted lines) at redshifts 0, 0.5 and 1. At $z = 0$ (1), there is a 1 per cent accuracy up to $k = 0.8$ (1.3) $h\,\mathrm{Mpc}^{-1}$. Regarding the mass function, the solid line in Fig. 3.7 depicts results at $z = 0.5$ for the optimal setup. For other redshifts the conclusions are quite robust, i.e., a 5 per cent underestimation at $M = 10^{12.5}\,h^{-1}\,\mathrm{M}_\odot$ and a ~ 5 per cent excess for $M \gtrsim 10^{13.5}\,h^{-1}\,\mathrm{M}_\odot$.

3.5 Halo clustering

In §3.4 we found an optimal configuration set-up for COLA by benchmarking the matter clustering and the halo abundance as a function of redshift against those measured in MICE-GC. We now study what that configuration implies for the clustering of haloes.

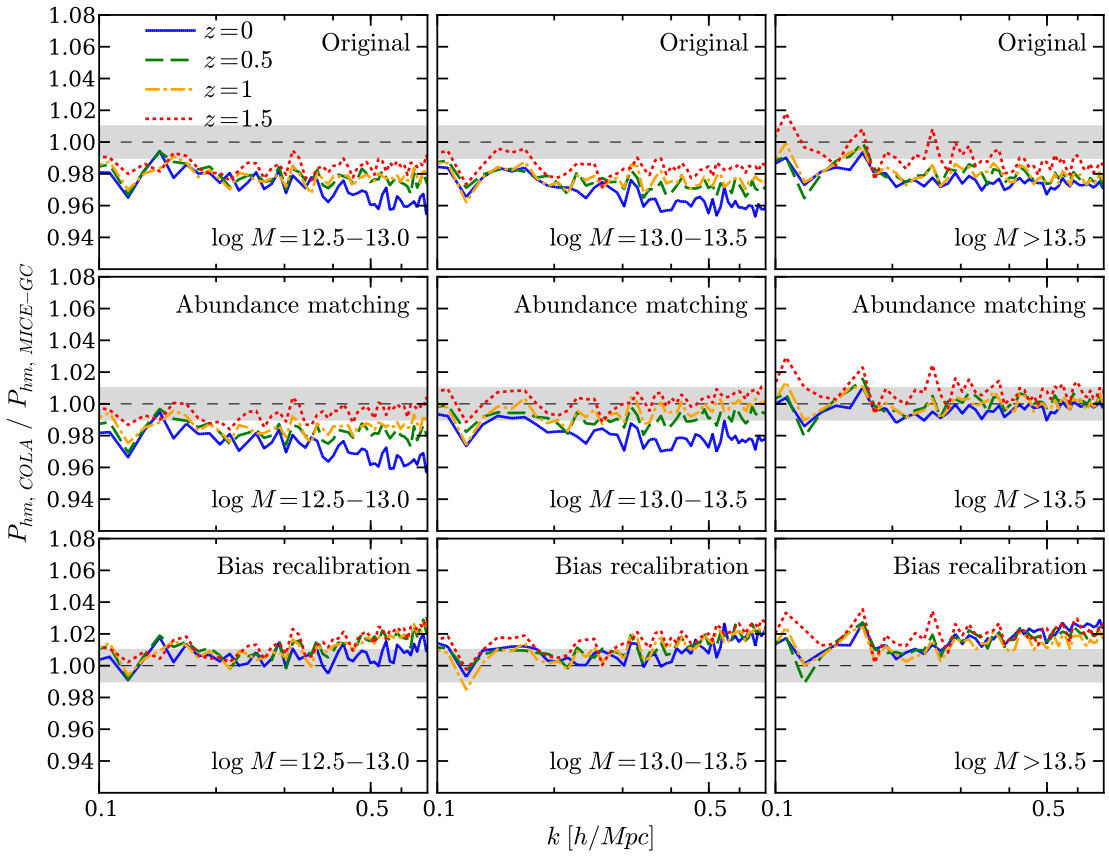

FIGURE 3.9: Halo-matter cross power spectrum. Rows correspond to different halo mass corrections: original mass, after abundance matching (see §3.5.1) and after bias re-calibration (see §3.5.2) from top to bottom. Columns separate halo samples M1, M2 and M3 from left to right. The original catalogue has a bias underestimation of $\sim$ 2 per cent in all cases. The abundance matching performs well at high masses while the bias re-calibration is able to achieve a 1 % agreement on large scales (i.e, $k < 0.3 - 0.4\,h\,\mathrm{Mpc}^{-1}$) and 2 % at intermediate scales ($0.4 < k < 0.7\,h\,\mathrm{Mpc}^{-1}$).

The first row in Fig. 3.9 shows the halo-matter cross power spectrum in real space and without applying any correction to the catalogues. Different columns separate mass samples M1, M2 and M3 (see Table 3.1) from left to right. Solid, dashed, dot-dashed and dotted lines display redshifts 0, 0.5, 1 and 1.5 respectively. The other two rows are explained in the following subsections §3.5.1 and §3.5.2. We notice that there is a general ~ 2 per cent under-estimation of the clustering amplitude at all mass bins and redshifts. This constitutes a remarkable result: it is possible to predict the halo linear bias with an accuracy of ~ 2 per cent without doing any correction nor the necessity of calibrating against a reference full N-body simulation.

Note, however, that we have also found evidence that halo masses are biased. Hence, we now explore two different corrections on the mass with the aim of reducing further the deviations in the halo bias to the $\lesssim 1\%$ level. One is based on fitting abundances (abundance matching, §3.5.1) and the other on fitting clustering (bias re-calibration, §3.5.2 §3.5.2).

3.5.1 Abundance matching

The cumulative halo mass function gives a monotonic relationship between the mass and the abundance of haloes. Biased halo mass estimates makes this function in COLA to have deviations with respect to a reference N-body simulation. If we have an external fiducial mass function (coming from a full N-body simulation for instance), we can re-assign the halo masses in the catalogue so that the reference abundance is fitted. When the incompleteness is negligible, we expect this calibration to greatly reduce disagreements among both catalogues if the ranking of halo masses has the correct ordering. If the incompleteness is present, there are missing entries in the catalogue and trying to match abundances will not produce the desired effect but a mixing of haloes with different clustering properties.

The second row in Fig. 3.9 shows the halo bias after correcting halo masses by abundance matching, using the measured mass function in MICE-GC as reference. The small disagreements in the top panels are greatly corrected in mass samples M2 at $z > 0$ and M3 at all redshifts, but not in M1. This is consistent with the impact of incompleteness in the mass sample described above: abundance matching works well as long as the incompleteness is not present, i.e. $M \gtrsim 10^{13}\, h^{-1}\, \mathrm{M}_{\odot}$ (see the solid line in Fig. 3.7).

We have tested as well the capabilities of the abundance matching for runs using only 10 time steps, in which the "uncorrected" mass function is highly under-estimated at $z = 1$ and the halo bias deviates by ~ 20 per cent (see Fig. 3.1). After abundance matching, the bias is recovered at the 3 per cent level for all mass bins and redshifts, but only for $k < 0.5\, h\, \mathrm{Mpc}^{-1}$, what illustrates that mass calibration performs worse when non-optimal parameters are used in COLA.

3.5.2 Halo bias re-calibration

We now explore an alternative mass re-calibration that is targeted to fit the halo bias. Note in the first row of Fig. 3.9 the COLA run (with optimal set-up) yields always a residual bias mismatch of 2 per cent, regardless of the mass sample and redshift. We can use this fact to build an alternative correction independent of any parent simulation (assuming the 2% factor is roughly independent of cosmology). In the framework of the halo model (Cooray et al., 2002) halo bias and halo mass are related through a function that only depends on cosmology $b = b(M)$. Thus, to first order, a fractional reduction in the bias of $\delta \ln b$ can be recovered with a shift in halo mass $\ln M \to \ln M - \delta \ln M$ given by[2],

$$\delta \ln M = \left(\frac{\partial b}{\partial \ln M} \right)^{-1} b \, \delta \ln b. \tag{3.2}$$

In what follows we set $\delta \ln b = 0.02$ (the bias calibration value we found) and evaluate the derivative in Eq. 3.2 at the corresponding mass and redshift using the bias prediction from Sheth et al. (1999) but we have checked that other fitting functions provide similar results.

The recovered halo bias values after doing such mass re-calibration are shown in the third row in Fig. 3.9. Now the agreement with MICE-GC is within 1% up to scales $k \lesssim 0.5\, h\,\mathrm{Mpc}^{-1}$ for all redshifts and masses. However, the correction is not working perfectly for the mass sample M3, where it is sub-percent up to $k \lesssim 0.3\, h\,\mathrm{Mpc}^{-1}$ but yields an over-estimate of $\sim 2\%$ beyond. We believe this could be due to the limited accuracy of the bias predictions coming from the theory of the peak background split (Manera et al., 2010), used to evaluate Eq. (3.2). Provided with a better bias prescription (or maybe the bias-mass relation measured from a reference N-body itself) one would expect the bias re-calibration to give very good results by construction. Nonetheless the accuracy remains within 1% for most cases and deviations from that are small.

As was the case for the abundance, this correction solves disagreements due to biased halo masses but not due to incompleteness. The over-abundance at high masses is removed but the underestimation at low masses persists and the true incompleteness, that might be larger, becomes visible. Despite that, haloes have the right clustering amplitude.

3.5.3 Redshift space

We now turn into discussing the performance of COLA for reproducing observables in redshift space, as this is what is actually observed in large scale structure galaxy surveys. Redshift space positions $\boldsymbol{s}$ are obtained by,

$$\boldsymbol{s} = \boldsymbol{r} + \frac{\boldsymbol{v}_r}{aH(a)}, \tag{3.3}$$

[2]We choose logarithmic quantities as we are interested on relative variations.

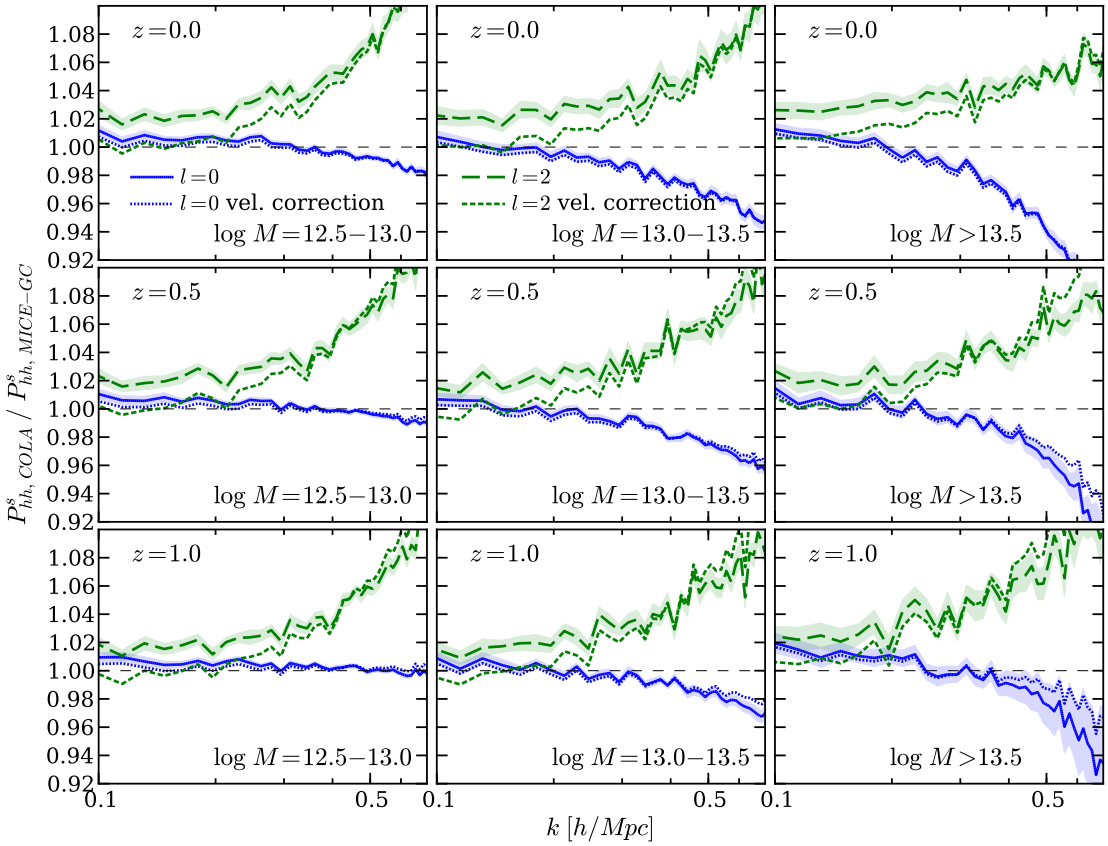

FIGURE 3.10: Monopole ($l = 0$) and quadrupole ($l = 2$) of the halo power spectrum in redshift space (solid and dashed lines respectively) in COLA vs. the MICE-GC N-body simulation. Different rows correspond to redshifts 0, 0.5 and 1 from top to bottom and columns separate mass samples from left to right (halo masses have been corrected by the bias calibration method). Monopoles have been corrected for shot-noise. Measurements in COLA correspond to the mean over 48 runs using the optimal setup. At large scales ($k < 0.3\,h\,\mathrm{Mpc}^{-1}$) the agreement is within 1 per cent for the monopole and $2\% - 3\%$ for the quadrupole. Dotted and short-dashed lines are the monopole and quadrupole after reducing halo velocities by 2 per cent, what brings the latter to better agreement while leaving the former unchanged.

where $\boldsymbol{r}$ is the position in real space, $\boldsymbol{v}_r$ is the peculiar velocity along the line of sight direction, a is the scale factor and $H = a^{-1}da/dt$ the Hubble expansion rate. For concreteness we will focus in halo power spectrum multipoles and assume the plane-parallel approximation, that is, fixing the line of sight to one of the three Cartesian axes.

In order to reduce the statistical errors in higher order multipoles we have produced 48 COLA runs using the optimal setup described in § 3.4.5. We split the halo catalogue in the 3 mass bins as in Table 3.1, with halo masses re-calibrated using the bias method with $\delta \ln b = 0.02$ (see § 3.5.2).

Figure 3.10 shows the mean of the monopole (l=0) and the quadrupole (l=2) over the suite of COLA runs divided by the corresponding quantities measured in MICE-GC. Rows separate redshifts $0, 0.5$ and 1 from top to bottom and columns mass bins M1, M2 and M3

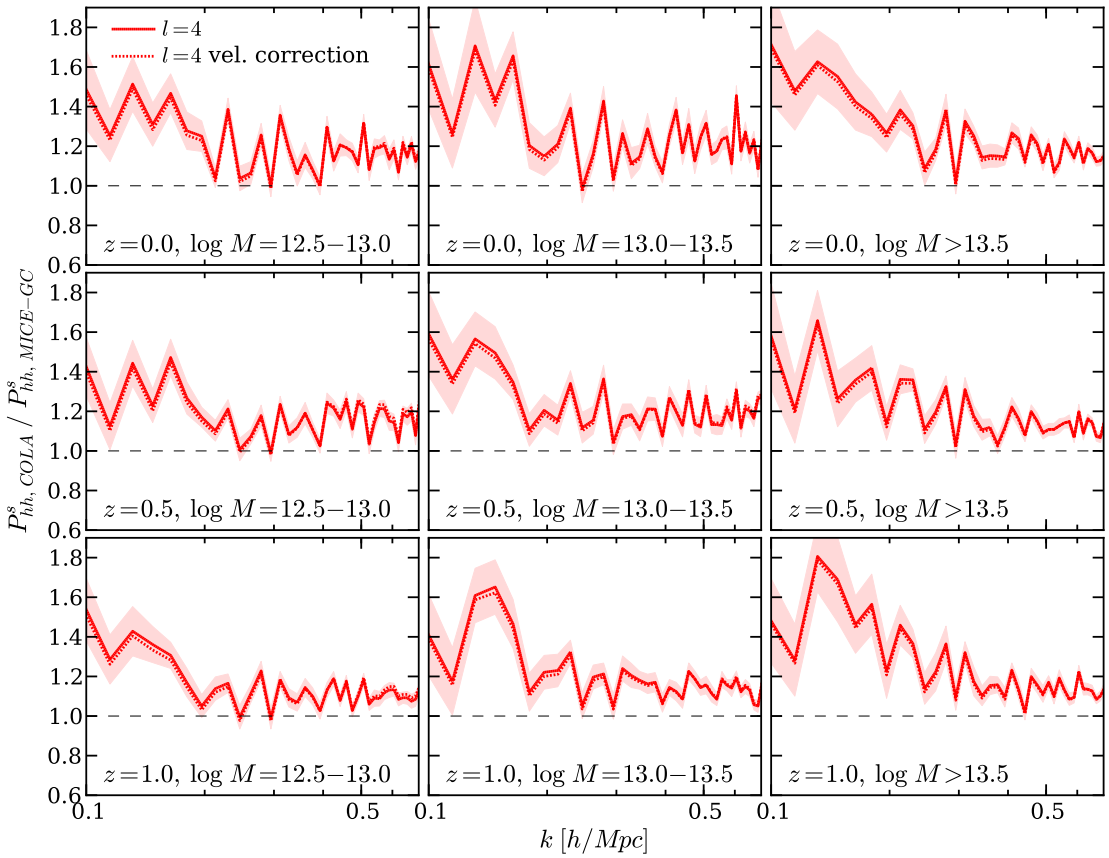

FIGURE 3.11: Mean halo power spectrum hexadecapole in COLA (after the correction in the halo masses by the bias re-calibration method with $\delta b/b = 0.02$) compared to the one in MICE-GC. Different rows correspond to redshifts 0, 0.5 and 1 from top to bottom and columns separate mass samples M1, M2 and M3 from left to right. For $k > 0.2\,h\,\mathrm{Mpc}^{-1}$ the agreement with the N-body is at the $\lesssim 20\%$ level across redshifts and mass bins. The dotted line includes as well a reduction in halo velocities of 2 per cent (see text for details).

from left to right.

For reference we recall the large-scale limit expressions for these quantities (Kaiser, 1987) assuming a simple linear bias model,

$$\begin{aligned} P^{\rm s}_{l=0,hh}(k) &= \left(b^2 + \tfrac{2}{3}bf + \tfrac{1}{5}f^2\right) P^{\rm r}_{mm}(k), \\ P^{\rm s}_{l=2,hh}(k) &= \left(\tfrac{4}{3}bf + \tfrac{4}{7}f^2\right) P^{\rm r}_{mm}(k), \end{aligned} \tag{3.4}$$

where b is the bias and $f \equiv \frac{{\rm d}\ln D}{{\rm d}\ln a}$ the linear growth rate.

At large scales ($k < 0.3\,h\,{\rm Mpc}^{-1}$), the agreement with MICE-GC is within 1 per cent for the monopole. Recall that we are using bias-recalibrated masses what ensures that the halo clustering is well reproduced in real space. And this contribution is the leading order for the monopole in redshift space, on large scales (i.e. the $b^2 P^r_{mm}$ term in Eq. 3.4). Had we used the actual COLA halo masses instead we would have obtained biases off by 2 per cent and the monopole underestimated by at least $\sim 4\%$. In turn, the quadrupole in Fig. 3.10 is systematically overestimated by ~ 2 per cent across all mass bins and redshifts ($k < 0.3\,h\,{\rm Mpc}^{-1}$). On large scales, the leading order contribution to the quadrupole is the cross-correlation between halo densities and halo velocities, i.e. the term $bf P^r_{mm}$ in Eq. 3.4. This means that any inaccuracies in reproducing the velocity field by COLA will have a direct impact in the quadrupole. For instance we have checked that the differences on large-scales can be corrected by reducing by 2 per cent each halo velocity (what would amount to reduce the overall bulk flow). We over plot (without error bars) the monopole and the quadrupole with dotted and short dashed lines respectively after applying such velocity correction. As expected, the quadrupole is now perfectly in agreement at large scales and the monopole remains almost unaltered.

At smaller scales ($k > 0.3\,h\,{\rm Mpc}^{-1}$) we observe larger discrepancies. The monopole is underestimated, specially at high masses, and the quadrupole is overestimated. We believe this is due to the details of the full velocity PDF[3] but we do not attempt to tune the results further to those of MICE-GC as i) the results on these scales will eventually depend on the galaxy sample under consideration and ii) these are scales that start to be smaller than those used in standard large-scale structure probes such as BAO. For instance, small-scale corrections can be postponed to a later stage when haloes are populated with galaxies using an HOD prescription, in which the velocity dispersion can be fitted to have agreement with observations. For reference, we have checked that adding a dispersion component to the halo velocities drawn from a Gaussian distribution with a width of $\sim 35\,{\rm km\,s}^{-1}$ and zero mean reduces the quadrupole for $k > 0.3\,h\,{\rm Mpc}^{-1}$ and is then in agreement within 2 per cent for most scales, mass samples and redshifts, whereas the monopole is not substantially affected.

Figure 3.11 shows the equivalent of Fig. 3.10 but for the hexadecapole ($l = 4$). We find that our optimal configuration for COLA yields an excess of ~ 50 per cent at large

[3]For example, we have measured the halo 1-dimensional velocity distribution and found that the fraction of haloes with center of mass velocities larger than 500 km/s is slightly underestimated by few percent in COLA (the exact number varies for mass samples and redshift), although the halo velocity rms agrees within 1 per cent with MICE-GC.

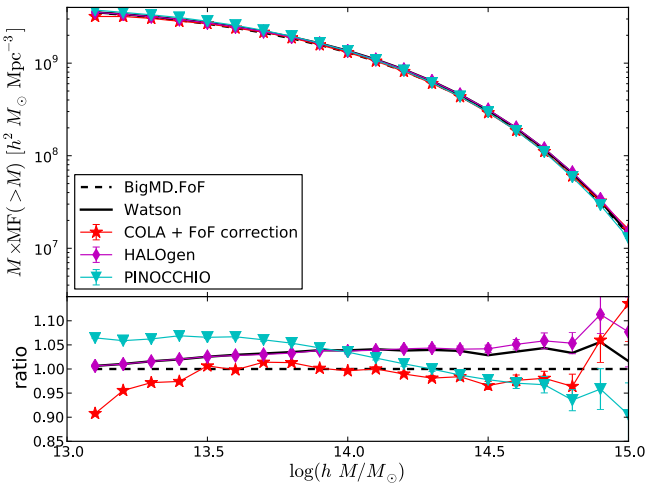

FIGURE 3.12: Comparison of the mass function for some of the methodologies. Only COLA and PINOCCHIO have predictive power on that observable, although the second one has been tuned to reproduce it.

scales, while for $k > 0.2\,h\,\mathrm{Mpc}^{-1}$ the agreement is significantly improved, down to ~ 20 per cent. These differences are not significantly changed when the velocity correction of 2 per cent is applied. If we further add an ad-hoc velocity dispersion term, as discussed above for lower multipoles, we achieve an agreement within 10 per cent at small scales.

3.6 Fast methods comparison project

This last section is devoted to a comparison project among most of the existing approximate methods, that is: COLA, EZmocks, HALOgen, PATCHY, PINNOCHIO and PTHalos (see §2.2 for references). For the comparison, all the methodologies were run using the same cosmology and box size as the reference N-body simulation BigMultiDark (Klypin et al., 2016). Then a halo sample with the threshold $M \gtrsim 10^{13}\,h^{-1}\,\mathrm{M}_\odot$ was drawn and the clustering measured.

The comparison of the mass function is shown in Fig. 3.12. The agreement is within 5% for the different methods. In particular, COLA was run with only 10 time steps and a particle mass ~ 20 times larger than our fiducial set-up (see §2.4), and taking that into account the findings are compatible with previous results (see Fig. 3.1). PINOCCHIO also accurately predicts the abundance of haloes and HALOgen is included in the plot for a consistency check, since it should reproduce the input mass function of Watson et al., 2013, as it does. The rest of methodologies fit reference halo abundances and therefore there is no interest in showing them in the plot.

Fig. 3.13 shows the cuadrupole of the power spectrum. These confirm that COLA provides the most accurate results, that are within 5% with respect to BigMultiDark, even that for this particular project the standard 10 time steps configuration was used. PINOCCHIO reproduces the cuadrupole within 10%. Note however that the code version used

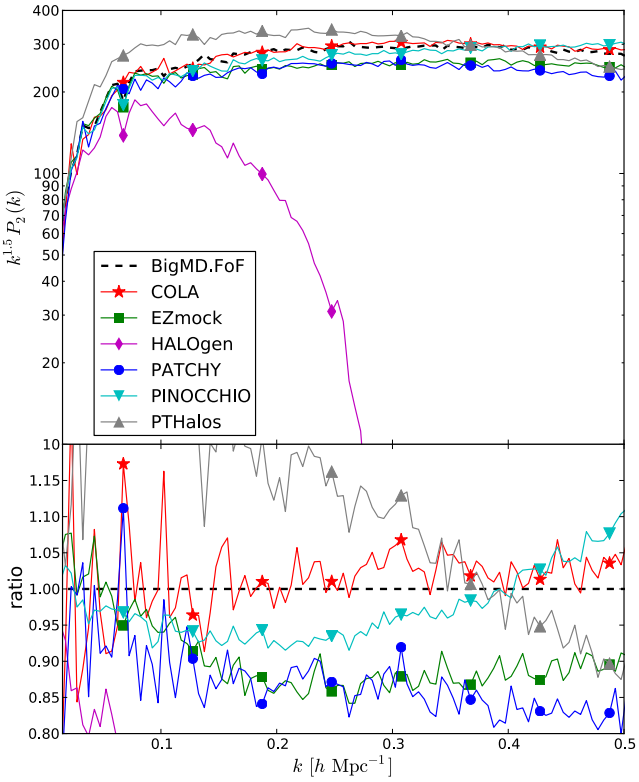

FIGURE 3.13: Comparison of the quadrupole of the power spectrum. Only COLA gives predictions within 5% at all scales and PINOCCHIO within 10%.

is old and that the most recent one (Munari et al., 2016) expands the dynamics to higher perturbative orders for a better accuracy. EZmocks and PATCHY have an accuracy at the 15% level. Nonetheless, they demand much less computational resources.

Finally, Fig. 3.14 shows the three-point correlation function in real space, which is defined similarly as the two-point correlation function (see equation 1.13 in §1.3) as $\zeta(r_1, r_2, r_3) =< \delta(\boldsymbol{r}_1)\delta(\boldsymbol{r}_2)\delta(\boldsymbol{r}_3) >$. The measurements are for triangles with two fixed lengths of 10 and $20\,h^{-1}\,\mathrm{Mpc}$ while the third one varies and is represented in the x-axis. Results are more noisy but the conclusions are similar than the ones for the cuadrupole of the power spectrum.

Comparing COLA against the rest of techniques, it is clear that COLA is the one that has higher computational requirements, specially on the memory and the CPU-time. However, it is still 2-3 orders of magnitude faster than a conventional N-body simulation and it is the most accurate method. Furthermore, it is a stand-alone technique that does not rely on any assumption nor calibration with an external simulation.

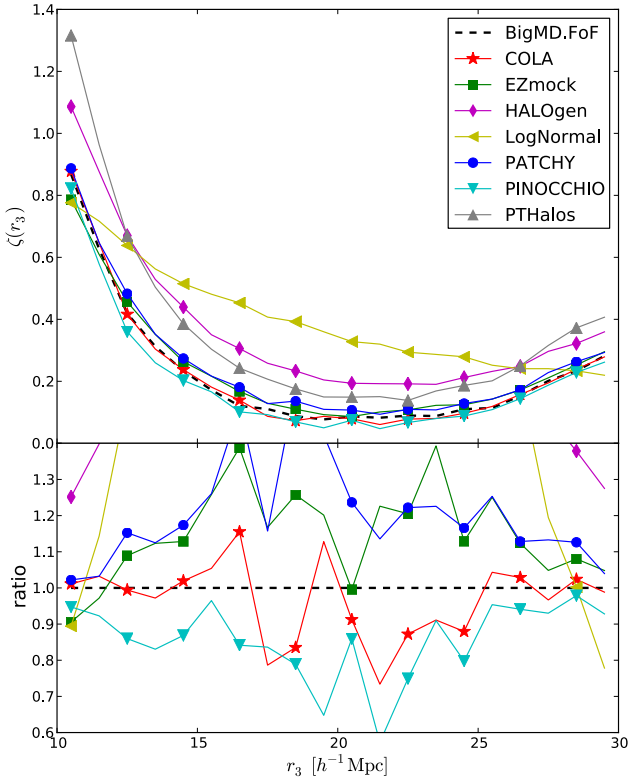

FIGURE 3.14: Results for the three-point correlation function in real space.

4

SIMULATIONS IN THE LIGHT CONE

According to special relativity, light travels at a finite speed. As a consequence, all physical information in the Universe is obtained in the so-called past light cone, in which light from more distant objects that is received by observers today was emitted at an earlier stage of the cosmic evolution. Hence, realistic simulations of the galaxy distribution should incorporate such structure in the output data, where objects are less evolved as their distance to the observer increases.

This chapter is devoted to how `ICE-COLA` implements this feature and produces catalogues on-the-fly. We start in § 4.1 by exposing the problem of generating light cone data in cosmological simulations. How this is solved in `ICE-COLA` is presented in § 4.2, where is sketched the algorithm for identifying the particles that enter into the light cone. This data can be used to produce different kind of catalogues, as it is explained in § 4.3. Finally, in § 4.4 is discussed the performance of the code in various computational aspects.

4.1 Mimicking real observations

Astronomy is a scientific discipline in which distances are tremendously large and, in particular, cosmology stands at the highest limit of the distance ladder (see Rowan-Robinson, 1985 for an historical review or e.g. Freedman et al., 2012; Efstathiou, 2014; Riess et al., 2011; Beaton et al., 2016; Riess et al., 2016 for the state-of-the-art in direct measurements of the Hubble constant). Observations are one of the main sources of knowledge in astrophysics and consist in collecting photons emitted by celestial objects. The light we receive in the present was emitted in the past and therefore the properties we can measure are

relative to that time. In terms of the special relativity, we observe our past light cone, where more distant objects are seen at an earlier time. Within the context of an expanding universe, the *lookback time* t_L to an object is the increment of the age of the Universe between the emission time of the photons and now (Hogg, 1999)

$$t_L(z) = t_H \int_0^z \frac{dz'}{(1+z')\,E(z')}, \tag{4.1}$$

where $t_H \equiv 1/H_0$ is the Hubble time nowadays and $E(z)$ is defined in equation 1.10. Observational data consists thus in a mixture of properties of objects at different evolutionary stages according to their redshift. Therefore, mock catalogues oriented towards understanding cosmological surveys will be more realistic if they are built in the form of a past light cone, where the relative distance of an object from an hypothetical observer determines the time when its properties are to be computed. There are many galaxy surveys planned for the next years that will sample the cosmic evolution also well before dark energy dominates, i.e., $z \gtrsim 2$, such as Euclid[1], the Wide-Field Infrared Survey Telescope (WFIRST)[2], the Large Synoptic Survey Telescope (LSST)[3], Dark Energy Spectroscopic Instrument (DESI)[4] and the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS)[5]. They will probe large variations in the cosmic growth from matter to dark energy dominated regimes and this evolution needs to be modelled in mock catalogues.

Numerical N-body simulations evolve the matter distribution inside a box by computing its gravitational interactions (and hydrodynamical as well if fluid properties of gas particles need to be considered). The simplest way to store the data is creating snapshots at certain steps of the simulation, in which all the outputted information correspond to the same epoch. Then, having multiple snapshots allows to sample discretely the cosmic evolution and these can be stitched together to build a light cone (Blaizot et al., 2005; Kiessling et al., 2011a; Merson et al., 2013). If snapshots are spaced in the temporal axis by Δz, each one covers the interval $(z_i - \Delta z/2, z_i + \Delta z/2)$, which corresponds to a comoving distance interval (χ_{i-}, χ_{i+}). An observer is placed inside the box and the comoving sub-volume delimited by the distance interval is copied from snapshot i to the light cone. This can be done either in the plane-parallel approximation (in which distances with respect to the observer are large and line-of-sight vectors are assumed to be parallel), or in a spherical geometry (where sub-volumes form multiple concentric shells). The former is suited for small patches of the sky but the latter is necessary for large area mock catalogues that will be needed by upcoming surveys that will observe large fractions of the sky.

Another limitation of appending sub-volumes from different snapshots is that discontinuities are generated at the boundaries between them. These can be minimized if

[1] http://www.euclid-ec.org/
[2] http://wfirst.gsfc.nasa.gov/
[3] http://www.lsst.org
[4] http://desi.lbl.gov/
[5] http://pan-starrs.ifa.hawaii.edu/public/

a finer time sampling is used by producing a large number of snapshots, but this would lead to huge data volumes. Alternatively, one can generate a light cone on-the-fly, that is, at running time of the simulation, what overcomes these difficulties by skipping the production of snapshots as well as the post-processing step of stitching them together. This enables a dramatic reduction of the data volume to be stored and the generation of all-sky light cones with a smooth evolution across all the volume. However, it requires a more elaborate code to identify during the simulation those particles that are entering into the light cone and compute their coordinates at crossing time.

Fast methods aim at reducing computational requirements, and these includes as well data storage. Even more, since they are fast, they enable the production of massive ensembles of simulations. Therefore, it becomes essential to incorporate techniques to minimize catalogue sizes and therefore simulated data need to be produced in the light cone format at run time. This chapter is devoted to the modifications that have been implemented in `ICE-COLA` in order to produce light cone catalogues on-the-fly. First, § 4.2 describes the algorithm for determining the particles that enter into the light cone in multiple box replicas. Then in § 4.3 it is explained how different kind of catalogues can be produced. Lastly, § 4.4 assesses the numerical performance of the `ICE-COLA` in three different aspects: the overhead in the running time, the memory usage and the data volume of the catalogues produced.

4.2 Light cone construction

4.2.1 Crossing time

The core part of the algorithm is to determine for each particle the exact time when it enters into the light cone, which is called the crossing time t_c, or a_c if it is expressed as a scale factor. Typically, collissionless N-body simulations evolve particle coordinates by the drift and kick temporal operators (see § 2.3.2), which update particle positions and velocities respectively under the sole force of gravity for a dark matter-only simulation. When the evolution has reached redshift z, the concentric shell around an observer defined by $(z - \Delta z/2, z + \Delta z/2)$ can be constructed, where Δz defines the width and separation of shells. Conventional N-body simulations have high temporal resolution and calling the routines for the light cone construction at each time-step would not be efficient. For this reason it should be assigned a value to Δz considerably larger than the interval between consecutive time-steps. In `COLA`, however, time-steps are much broader and it is more adequate to adjust the light cone steps to those of the simulation itself. Therefore, **the light cone routine is called after each temporal operator**.

At any moment of the simulation, particle positions and velocities are given by the last drift and kick operators respectively, computed at scale factors a_D and a_K. Coordinates at intermediate time values can be interpolated using an extra pair of drift and kick operators to move the particle to a desired scale factor a (while it obeys $\min(a_D, a_K) < a < \max(a_D, a_K)$). By converting a_D and a_K to comoving distances (see equation 1.11) we find the radii that define a spherical shell in the light cone, as it is shown in Fig. 4.1.

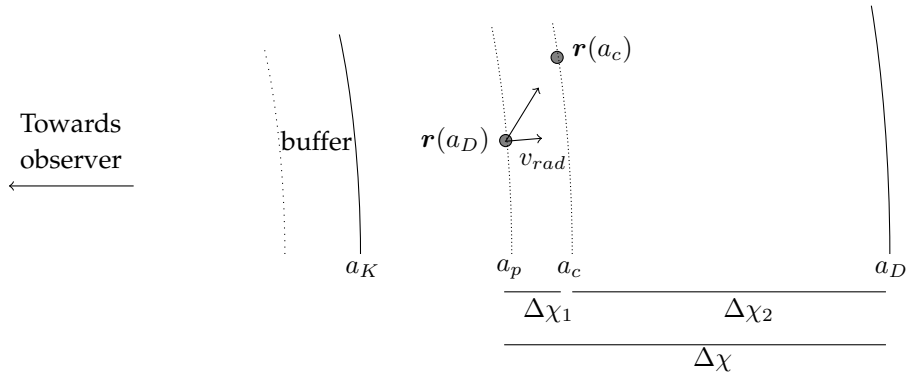

FIGURE 4.1: Diagram of a particle at its crossing time. Initially, at drift time, the light cone is at position a_D and the particle at $\boldsymbol{r}(a_D)$. Knowing these two quantities, together with the radial velocity, is enough to evaluate equation 4.5 that gives the crossing time a_c. Then, coordinates for the particles can be interpolated at that time. Note that the velocity buffer is only necessary around a_K.

Particles inside that volume are going to be included in the light cone in the current call, but before we have to compute their coordinates at crossing time.

A particle will move from position $\boldsymbol{r}(a_D)$ at drift time to $\boldsymbol{r}(a_c)$ at crossing time t_c that can be computed from its radial velocity, having a variation on the comoving distance along the line-of-sight of

$$\Delta\chi_1 \equiv |\boldsymbol{r}(a_c)| - |\boldsymbol{r}(a_D)| = \frac{v_{rad}(t_c - t_D)}{a}, \tag{4.2}$$

where $v_{rad} = \boldsymbol{v} \cdot \hat{\boldsymbol{r}}$ is the physical radial velocity and we assume that variations on the scale factor are negligible. During the same time, the light cone boundary shifts by a comoving distance

$$\Delta\chi_2 = \frac{c(t_c - t_D)}{a}. \tag{4.3}$$

Similarly, the expression for the total distance is

$$\Delta\chi = \frac{c(t_p - t_D)}{a}, \tag{4.4}$$

where t_p denotes the time when the light cone is at the comoving distance $|\boldsymbol{r}(a_D)|$. Then, by simply equating $\Delta\chi = \Delta\chi_1 + \Delta\chi_2$ the crossing time can be expressed as

$$t_c = \frac{c}{c + v_{rad}}(t_p - t_D) + t_D. \tag{4.5}$$

To evaluate it one needs the velocity of the particle, the time of the last drift operator and the initial position of the particle, that fixes t_p. For slowly moving particles, $v_{rad} \ll c$ and $t_c \approx t_p$: the crossing time is nearly equal to the moment when the light cone reaches the original position of the particle, as expected.

After determining the crossing time, particles are displaced to that time by computing $D(a_D \to a_c)$ and $K(a_K \to a_c)$. Note that particles might enter inside the volume of the shell defined by a_D and a_K when the displacement is applied and therefore an extra buffer zone is needed at the edge associated with a_K, as shown in Fig. 4.1. The crossing time is computed as well for particles in that region but only those which are inside the shell after the displacement are kept in the light cone output. Note that the buffer is only necessary around the boundary associated with the kick operator, since close to the other boundary we have that $t_p \to t_D$, from equation 4.5 it follows that $t_c \to t_D$ and therefore $D(a_D \to a_c) \to D(a_D \to a_D) = 0$.

Equation 4.5 has to be evaluated for every particle that is a candidate to be inside the current shell of the light cone. This involves conversions between the scale factor, the age of the Universe and comoving distances. For a fast performance of the algorithm, values for those functions are computed once, at the beginning, and later they are just interpolated at the desired value. Since they are smoothly-varying quantities, with as few as 10 interpolating points sampling the range of values needed at a given call to the light cone routine and using double precision is typically enough to ensure sufficient accuracy. The temporal operators that displace particles to the crossing time need to evaluate as well time-dependent variables and are also pre-computed and interpolated later on.

The accuracy on the crossing time determined by evaluating the approximate eq. 4.5 with interpolated quantities, can be tested calculating the difference between the radial distance to the particle at crossing time, $|\boldsymbol{r}(a_c)|$, and the light cone position at the same time, $\chi(a_c)$. We have checked that deviations are only $\sim 1\,h^{-1}\,\mathrm{Kpc}$ on average and $\sim 50\,h^{-1}\,\mathrm{Kpc}$ the largest value of all particles[6]. This is orders of magnitude below the typical force resolution length employed in `COLA` so we can hence safely neglect this error.

A pseudo-code that sketches how the light cone algorithm works is shown in Fig. 4.2. It also includes the generation of box replicas, explained next in Sec. § 4.2.2.

4.2.2 Box replicas

The position of the hypothetical observer, that determines the centre of the light cone, is arbitrary for a simulation with periodic boundary conditions. For simplicity, it is usually chosen to be at coordinates $(0, 0, 0)$, at one corner of the box, so that the full simulated volume extends up to $(L_{box}, L_{box}, L_{box})$. If the outputs of the simulation are snapshots (or data derived from them), in which all the information correspond to a single epoch, the volume $(0 - L_{box}, 0 - L_{box}, 0 - L_{box},)$ contains all the possible information that can be extracted from the simulation at a fixed time. Nonetheless, in the production of light cones this is no longer the case since the time of writing information in a region of space depends on the relative distance between the particles and the observer. This fact, together with the periodic boundary conditions, allows to replicate the box to adjacent positions and fill in a larger volume around the observer. As a consequence, the total volume of the catalogues will extend to higher redshifts and project a larger area on the sky (for

[6] These numbers are taken from a simulation with 2048^3 particles.

$\chi_1 = \chi(max(a_D, a_K))$;
$\chi_2 = \chi(min(a_D, a_K))$;
$\chi_{1,buffer}$, $\chi_{2,buffer} = f(\chi_1, \chi_2,$ `velocity_buffer_size`$)$;
Compute quantities evaluated at a_D or a_K;
Setup spline interpolation for (a_D, a_K);
for *(`i = 0 → n_replicas - 1`)* **do**
  Compute origin of the replica `i`: $\boldsymbol{r}_{0,i}$;
  if *(not inside replica `i`)* **then**
    Skip replica ;
  end
  for *(`j = 0 → n_particles - 1`)* **do**
    if *($\chi_{1,buffer} < |\boldsymbol{r}_j(a_D) + \boldsymbol{r}_{0,i}| < \chi_{2,buffer}$)* **then**
      Compute crossing time a_c;
      Compute Drift$(a_D \to a_c)$;
      if *($\chi_1 < |\boldsymbol{r}_j(a_c) + \boldsymbol{r}_{0,i}| < \chi_2$)* **then**
        Compute Kick$(a_K \to a_c)$;
        Save particle j coordinates;
      end
    end
  end
end

FIGURE 4.2: Algorithm of the light cone construction. Quantities such as the crossing time or the temporal operators are computed just for those particles which is needed. Here it is shown a simplified version, the full code has additional features and some of them are given in Fig. 4.7 for the case of running a halo finder in the light cone.

instance, full sky). As a disadvantage, modes are sampled multiple times and repeated structures could potentially appear in the light cone, although separated by a distance comparable to the box size. However, since a given particle in two different replicas will be seen by the observer at different times, their evolutionary stage will differ and thus no obvious repetitions of structures will be observed in the light cone volume.

The algorithm produces box replicas by simply shifting positions by the vector $\boldsymbol{n} \equiv (n_x, n_y, n_z)\ L_{box}$, where $\boldsymbol{n}$ is expressed in units of the box size and therefore its components have integer values. `ICE-COLA` builds the necessary number of replicas according to the values of two parameters specified in the parameter file:

Starting redshift, z_{LC}: converting it to a comoving distance and dividing by L_{box} gives the necessary number of replicas in one direction:

$$n_{rep,1D} = \text{ceiling}\left(\frac{\chi(z_{LC})}{L_{box}}\right) \tag{4.6}$$

Number of dimensions replicated, n_{dim}: It determines in which axis i it is allowed to have replicas along the negative semi-axis: $n_i < 0$. If set to 0, the light cone is built only in the octant where $x, y, z > 0$. If set to 1, n_x iterates over positive and negative values; with 2 also n_y; and with 3 n_z as well, which produces an all-sky light cone.

Then, the total number of box replicas obeys the formula

$$n_{rep,total} = n_{rep,1D}^3 \times 2^{n_{dim}}. \tag{4.7}$$

It scales with the third power of the number of replicas in one direction, so a high value for the starting redshift may produce a very large number of box replicas, making the light cone computations to become a significant fraction of the simulation time. The configuration that has been used for the light cone simulations of this thesis uses $n_{rep,1D} = 2$ and $n_{dim} = 3$, which results in 64 box replicas (as was the case for the reference N-body simulation, MICE-GC).

When the code iterates over replicas, it is very straightforward to check if the the sub-volume is intersected by the current shell being constructed. The replicas where this condition is not true are skipped, reducing the number of computations, as it is sketched in Fig. 4.2.

4.3 Catalogues on-the-fly

Section § 4.2 (see also Fig. 4.2) gives an overview of the core of the algorithm to build the light cone on-the-fly in `ICE-COLA`. Now we turn in how this information can be used to generate different kinds of catalogues. Some of them are complementary or represent different levels of data compression.

After running the algorithm that builds the light cone, the code can access a subsample of particles with their coordinates (computed at the crossing time). The simplest catalogue that can be produced is just a dark-matter particle distribution, as described in § 4.3.1. It represents a low-level output that contains all the information. Therefore this turns out into huge data volumes, that are not suited for producing massive ensembles of realizations. For instance, an all sky light cone starting at redshift 1.4 in a simulation with a particle mass of $3 \times 10^{10}\, h^{-1}\, \mathrm{M}_\odot$ generates 6 TB of data (matching the MICE-GC light cone specifications), which is hard to write, store, transfer and analyse. For these reasons it is essential to produce derived and compressed data catalogues that, although will not contain all the information, allow to model the necessary observables for galaxy surveys. With that motivation, `ICE-COLA` implements two higher level catalogues: § 4.3.2 describes the projected matter density field in pixelated concentric spherical shells (that enable modelling weak lensing observables) and § 4.3.3 is devoted to the implementation of the Friends-of-Friends algorithm (Davis et al., 1985) in the light cone (that produce halo catalogues). Thanks to these data formats, large compression factors can be achieved.

4.3.1 Dark matter particle distribution

As discussed above, a particle catalogue may become extremely large, specially if the volume sampled by the light cone covers many box replicas, and thus it requires massive I/O[7] operations. Using binary (unformatted) files is mandatory for such outputs, in

[7] Input/Output.

order to reduce their size and for a better performance of the I/O routines. A quite standard format in cosmological simulations is the Gadget one, very popular due to the well known code with the same name (Springel, 2005)[8]. `ICE-COLA` also adopts this format, in which data is plit into several files that contain a header of 256 bytes. The number of particles contained in a file has to be known in advance at the moment of starting to write it, since it is contained in the header. But during the light cone construction, only the particles in the current shell are accessible in memory after being determined, and therefore the Gadget file can only include those particles. Therefore, the code would have to produce a distinct file for every shell, process and replica, which for large runs might turn into several thousands. This is often regarded as a bad practice, since most file systems may be inefficient managing those numbers.

One simple solution would be to re-write files to append new data after each call to the light cone algorithm, which would keep the header updated. This might still be acceptable for conventional N-body simulations, but not for `ICE-COLA` since it is necessary to have optimized I/O operations that do not spoil the speed-up of the method. Instead, the solution adopted is to write distinct files for each replica and shell, but gather processes in groups that write in parallel into the same file. The number of files is controlled by the parameter `n_files` and the number of processes that are grouped for each file is determined by trying to have a similar number of particles `np_target` among the files:

$$\texttt{np\_target} = \frac{\texttt{np\_remaining}}{\texttt{n\_files\_remaining}}, \tag{4.8}$$

where `np_remaining` is the sum of the number of particles of processes still not assigned to any file and `n_files_remaining` is the number of files remaining to have processes assigned. Initially, processes are added to the group that will write to the first file until their joint number of particles is similar to `np_target`. Then the values for `np_remaining` and `n_files_remaining` are updated and the new target number of particles is used for determining the second group of processes that will write to the second file. The process is repeated until all processes have been assigned a file, as it is depicted in more detail in Fig. 4.3.

However, the algorithm described in Fig. 4.3 is a simplified version and there are some exceptions that the code also takes into account. The reason for most of them is that some processes may not have particles inside the shell of the light cone. We refer to these as inactive processes. If the number of active processes is lower than `n_files`, then the number of files is set to the number of active processes. In the situation that all processes are inactive, the code still writes a file, that contains just the header and zero particles. The reason is to avoid gaps in the file nomenclature, that would complicate the input routine in a posterior analysis of the data. Finally, the code also checks at each iteration of the algorithm whether the number of active processes pending to be assigned to a file is equal to `n_files_remaining`. In such case, these active processes are only grouped with inactive ones (so that at least each file has one active process).

[8]The user guide of the Gadget-2, where also the Gadget format is explained, is available at http://wwwmpa.mpa-garching.mpg.de/~volker/gadget/users-guide.pdf

```
rank_new = 0;
rank_old = 0;
np_remaining = sum(np[:]);
n_files_remaining = n_files;
for ( ifile = 0 → n_files - 1) do
    n_target = np_remaining / n_files_remaining;
    while True do
        if ( sum(np[rank_old:rank_new+1]) ≥ np_target) then
            dif1 = np_target - sum(np[rank_old:rank_new]);
            dif2 = sum(np[rank_old:rank_new+1]) - np_target;
            if (dif2<dif1) then
                rank_new++;
            end
            break;
        end
        rank_new++;
    end
    rank_files[ifile] = rank_new;
    np_remaining -= sum(np[rank_old:rank_new]);
    rank_new++;
    rank_old = rank_new;
    n_files_remaining--;
end
```

FIGURE 4.3: Algorithm of the dark matter particle light cone for aggreging processes into `n_files` groups that will write files in parallel. The variable `np` must be an array containing the number of particles dumped to the light cone by each process. `rank_files` will contain the maximum rank, i.e., process number, that will write to each file. The actual code is bit more elaborated than this pseudo-code, see the text for more details.

Once finished the grouping of processes, a new MPI communicator is built for each file. The header and the two 4-byte integers that give the length of each data block are written by the process with rank 0 in the file communicator, and the actual data is written in parallel by collective MPI-I/O writing routines[9]. This is done by setting the correct file displacement for each process, that is, the absolute byte position relative to the beginning of the file. It is computed by communicating how many particles will be written by the processes in the left. Then, each section of the file is written in parallel by the processes that belong to its communicator.

4.3.2 Two-dimensional projected matter density fields

A way to compress the data volume of a simulation is to build a grid and store for each cell the occupation number of particles. The product is a density field with the resolution of the cell size. And the compression factor comes from the fact that the number of cells will be typically much smaller than the number of particles. This technique is employed

[9]For a complete description of the MPI-I/O routines see chapter 13 of the MPI-3.1 standard, available at http://www.mpi-forum.org/docs/mpi-3.1/mpi31-report.pdf

frequently in simulations by choosing a regular grid on the Cartesian axes and enables studies even in the non-linear regime if a good resolution is used.

However, for the catalogues described in this section, ICE-COLA uses a different approach by using spherical coordinates, which are the natural choice for the light cone. Thus, the matter field is projected onto several thin concentric shells around the observer, which can be later used to model weak lensing observables (see § 5.2). This idea was originally proposed and implemented in N-body simulations by Fosalba et al., 2008, using the Healpix discretization of the sky area into pixels (Górski et al., 2005) and later implemented in other light cone simulations (Das et al., 2008; Teyssier et al., 2009). For this thesis, the idea has been adapted to the ICE-COLA algorithm and peculiarities.

In the first place, the discretization of the volume is done using spherical coordinates centred at the observer. Placing several hundred bins along the radial direction suffices to recover accurately the redshift evolution. In particular, there are 265 shells until redshift $z = 1.4$, having a width of $\sim 7(18)\, h^{-1}\,\mathrm{Mpc}$ at low (high) redshift, corresponding to steps of $\sim 35\,\mathrm{Myears}$ in look-back time. The binning was chosen to be the same than the used for MICE-GC for an ease comparison to that reference simulation. It is left for future work the study of how many radial bins are enough to recover accurately the matter field within the accuracy of COLA.

Then, each shell is discretized into equal-area pixels using the Healpix format. The simulations developed for this thesis use $N_{side} = 2048$, that produces $12N_{side}^2 \approx 50$ million pixels with a size of $\theta \simeq 1.7\,\mathrm{arcmin}$. Maps consists on number counts of dark matter particles, represented by 4-bytes integers. A single map has 192MB of data, so all the maps need 50 GB of disk space. This represents a compression factor of two orders of magnitude with respect to the size of the dark matter three-dimensional particle light cone (see § 4.3.1).

Note that the light cone routine is called twice at each time step: after the drift and the kick operators. So a simulation with 40 time steps will generate less than 80 calls (since at the beginning the light cone has still not started) but will output 265 Healpix maps. Which means that many of them have to be constructed simultaneously. The procedure is the following. Before the algorithm starts the loop over replicas (see Fig. 4.2), all maps are initialized to zero. At the point where particle coordinates can be saved, the code instead adds a count to the pixel of the map that correspond to its angular and radial positions. The maps are finished after iterating over all the replicas and the particles inside them.

However, the wide shell built at each call to the light cone routine is used to produce several thinner maps, as discussed above. These maps that are in the edges of the wide shell are incomplete, since their volume extend to the shell computed during another call to the routine. In these cases, the second call writes a separate file and both partial maps have to be joined at post-processing by simply adding them.

We can see in Fig. 4.4 the matter distribution at full sky and two zoomed regions of a stack of various shells at redshift 0.5. The homogeneity of the Universe is evident at large scales, while filaments and voids display large fluctuations at more moderate scales.

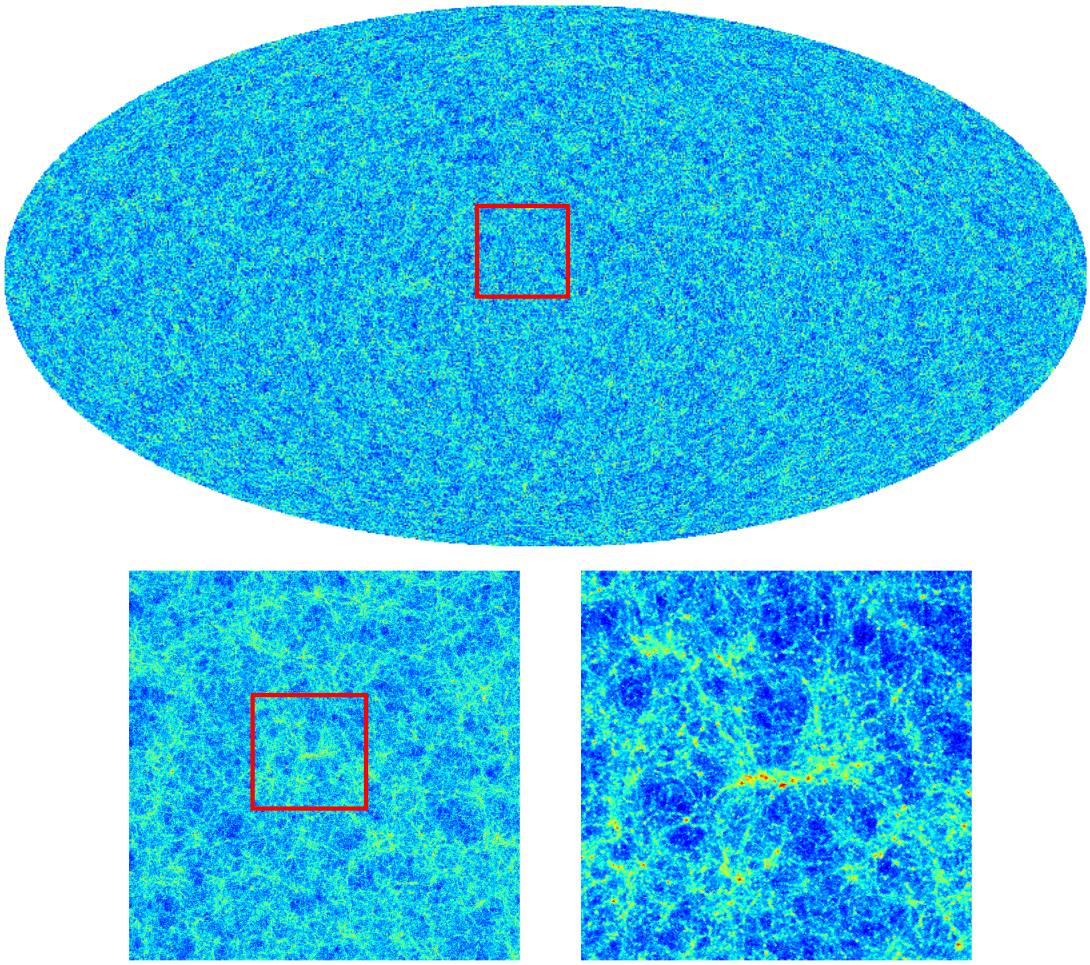

FIGURE 4.4: Visual representation of the two-dimensional projected matter field. The upper graphic shows all the sky, while the lower left and right zoom in patches of $30 \times 30\text{deg}^2$ and $8 \times 8\text{deg}^2$ respectively. 20 Healpix maps with an angular resolution of $1.7'$ were stacked in the redshift interval $z = 0.45 - 0.55$.

The maps thus obtained serve to model weak lensing observables on the full-sky light cone, incorporating naturally the spherical geometry. Furthermore, the Healpix software is specially designed to efficiently transform maps in the spherical harmonic decomposition, which allows measuring angular power spectra or solving the equations of weak lensing efficiently with built.in tools. §5.2 explains the pipeline developed with that aim and the accuracy of the whole method with `ICE-COLA`, which is studied further in Chapter 6 where they are applied to halo catalogues.

4.3.3 Halo catalogues

Halo catalogues are one of the main outcomes of large-scale structure simulations. `Parallel COLA` includes a parallel Friends-of-Friends (FoF) algorithm (Davis et al.,

1985) that runs on-the-fly on particle snapshots, and it is based on the publicly available serial code from the N-body shop of the University of Washington[10]. This thesis explains how in ICE-COLA it has been adapted to produce FoF catalogues in the light cone. Although the modification of the algorithm may seem rather simple from a theoretical point of view, there are many technical underlying complications that are briefly explained in this section.

To start with, it is necessary to understand how this halo finder works. The FoF algorithm links together particles that are separated by a distance smaller than the so called *linking length*, b. Groups are formed by adding the neighbours that are within the linking length of particles members of the group. The linking length is expressed in units of the mean inter-particle distance and for this work we use the widely used standard value $b = 0.2$, which corresponds approximately to an overdensity of $\Delta = 180$ in the context of the spherical collapse model (Cole et al., 1996). Once groups have been found, the halo coordinates are simply given by the centre of mass of the member particles and their mean velocity.

The FoF does basically local operations, that consist first in searching for neighbouring particles. To do so, the serial code version from the N-body shop uses a k-dimensional tree (or k-d tree for short), in which the space is recursively divided into halves until each leaf nodes contains few particles. In ICE-COLA, when the FoF is called from the light cone, some processes may not have particles in the current shell. This is something that the code does not expect and triggers a bug while building the k-d tree. It can be avoided by simply adding an isolated fake particle to these processes, that will not form any group.

Buffer zones and ghost particles

In Parallel COLA, groups are built within the sub-volume of each process. Afterwards, a communication step is required to merge the structure that lies near the edges of each slab. The code assumes periodic boundary conditions and edges corresponding to the box limits are mapped to the opposite side of the box. This has no confusion when the particle distribution comes from a snapshot. In the light cone, however, it needs a special care since the crossing time varies with the position. The environment of a particle close to the box edge has to be completed with the distribution corresponding to the neighbouring replica. But since only one replica is computed at a time, it is necessary to add buffer zones around box edges and provide in that way the required environment to the algorithm. For the same reason, buffer zones are also needed around the shell limits and are added to the velocity buffer (for the latter, see Sec. §4.2.1). The particles inside either the shell or the boundary regions are passed to the FoF code, but at the end only those haloes whose centre of mass position is inside the volume of the current shell and the local replica are written to the catalogue.

A visual representation of this geometry is depicted in Fig. 4.5, that shows a box replica intersected by a shell of the light cone, coloured in blue. The two red strips are

[10] http://www-hpcc.astro.washington.edu/tools/fof.html

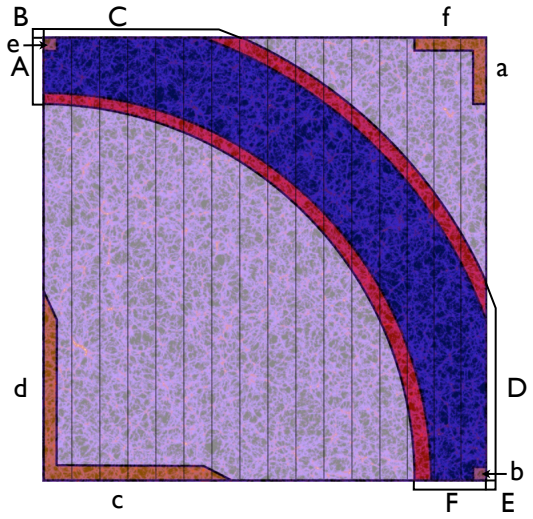

FIGURE 4.5: Sketch of the buffer zones needed for running the FoF algorithm in the light cone. The blue band is the shell under construction, which is surrounded by a buffer zone. Buffer zones within the local replica are coloured in red, while those that correspond to neighbouring replicas are in orange. Particles belonging to the latter are called ghost. Letters identify the different regions with ghost particles, using capital letters for their original physical location and the lower case for their associated region mapped inside the local replica. Vertical lines depict the volume decomposition of the parallelization (in that case it would correspond to 16 processes).

buffer zones belonging to the local replica, while the orange ones correspond to neighbouring ones. We refer to the particles within the latter as *ghost particles*, since their physical information really belongs to the white regions outside the box. For these, their crossing time is computed as if they were placed in the neighbouring replica, but they coordinates are trimmed to the range $0 - L_{box}$, as shown in the figure. The reason is that the FoF algorithm expects values inside this range, and thanks to periodic boundary conditions the final result will be the same. However, an additional step will be necessary at the end to shift by $\pm L_{box}$ the position of some of the haloes, as we explain next.

The width of buffer zones has to be larger than the maximum separation between any particle and its halo center of mass position. Due to the approximated dynamics in COLA, haloes are puffier and more extended (Tassev et al., 2013) and therefore it is difficult to predict theoretically a correct size for the buffer zones. The pragmatic approach used is computing few statistics during a real run (see §2.4.2 for the characteristics of the light cone simulations developed for this thesis). For simplicity of its technical implementation, it is registered the distance between the first particle that starts forming a group and the centre of mass of the final object. We shall assume that this particle has the same properties as if any random member of the group was picked up. The rms of the distance thus defined is $\sim 0.3(0.4)\, h^{-1}\,\mathrm{Mpc}$ for all FoF with more than 20 (50) particles. This separation is only representative for haloes close to the minimum mass, since they are the more abundant.

To encompass all possible haloes it is necessary to look at the maximum separations,

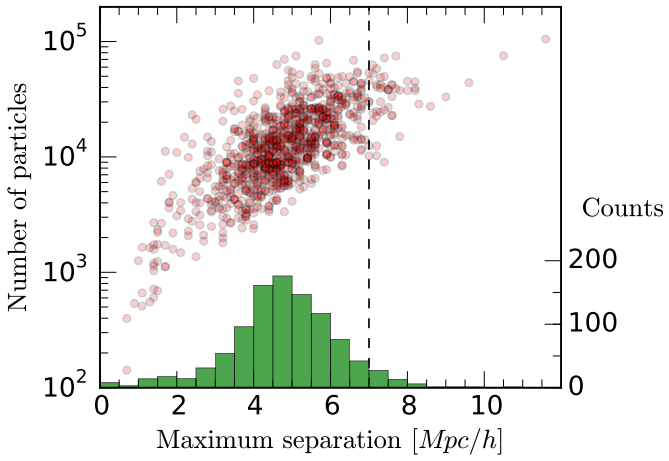

FIGURE 4.6: Distribution of the maximum distance between a particle and its halo center of mass. Each event is a call to the FoF routine. The y-axis of the two dimensional plane is the number of particles of the halo: mass is correlated with the maximum separation. The distribution is projected into a histogram below, peaking at $5\,h^{-1}\,\mathrm{Mpc}$. The vertical line displays the width of buffer zones, see the text why this value is an optimal choice.

as displays Fig. 4.6. This scatter plot has one thousand events, corresponding to all the calls to the FoF routine in the whole simulation. Represented in a two-dimensional plane is the maximum separation of all groups found in a single call and the number of particles of that object. The distribution is projected into a histogram in the lower part, showing that it peaks at $5\,h^{-1}\,\mathrm{Mpc}$ and declines quickly thereafter.

With Fig. 4.6 in hand, the most conservative choice would be to set the buffer zone to $\sim 10\,h^{-1}\,\mathrm{Mpc}$, which encompasses almost the whole range of values. However, for the sake of performace, values as low as possible are desired, in order to reduce the volume where computations are duplicated. In particular, it was set to $7\,h^{-1}\,\mathrm{Mpc}$. Maximum separations higher than that were only found in 50 calls to the light cone. For these, there may be the risk of providing a too small environment to the FoF algorithm for haloes for which: **i**) They are very massive and have particles further than $7\,h^{-1}\,\mathrm{Mpc}$, **ii**) their center of mass position is very close to the shell edges, and **iii**) their more distant particles happen to be in the direction of the buffer zone. The probability of **ii** can be estimated to be $\sim 10\%$ by considering the relative width of the region with incomplete environment ($12 - 7 = 5\,h^{-1}\,\mathrm{Mpc}$ at most) with respect to the width of the shell of the light cone ($40\,h^{-1}\,\mathrm{Mpc}$ or $100\,h^{-1}\,\mathrm{Mpc}$ at low and high redshift respectively)[11]. Therefore, out of the 50 calls to the FoF that have maximum separations larger than $7\,h^{-1}\,\mathrm{Mpc}$, very few of them may cut out some particles of the largest haloes. Taking into account that the total volume of the light cone is $108\,h^{-3}\,\mathrm{Gpc}^3$, these very few events represent a negligible fraction of the haloes. In fact, this is confirmed in §5.1.3, where it is demonstrated that

[11]The volume of boundary regions around the box edges has been neglected in this estimation, since it is much smaller.

the abundance of haloes found in the light cone is compatible with measurements from snapshots.

Tracking shifts of particles

`Parallel COLA` assumes periodic boundary conditions and that all the particles assigned to a process have positions that are inside the sub-volume corresponding to that process, according to the slab volume decomposition. Since particles are drifted during the light cone construction, it is necessary to transfer to another process those that have moved to the sub-volume assigned to another process before running the FoF. Positions are hence trimmed to the range $0 - L_{box}$. Particles that suffer a shift of $\pm Lbox$ in one axis for that reason have to be tracked in advance by the light cone code, so that the displacement is undone before writing haloes to the catalogue and the physical position in the light cone is recovered.

At this point, there are two reasons for which some particles may need to be shifted after the FoF algorithm: because their positions were trimmed or because they are ghosts (and belong to a neighbouring replica). That can be foreseen during the light cone construction and by saving this information, positions can be finally restored to the correct physical location. To do so, a flag variable is used to encode the information. Two bites are used for each spatial coordinate, with the meaning:

00 No displacement for this coordinate
01 A shift of $-L_{box}$ will be necessary
10 A shift of $+L_{box}$ will be necessary

So six bits store the information of the shift in the three coordinates: $zzyyxx$. For instance, a value 001000 encodes a shift of $(0, +L_{box}, 0)$.

To avoid the allocation of additional memory, the field of the particles ID's, that is not needed any more, is used to store the flags. The FoF code, however, adds particles to groups and discards individual particle information (including the ID). It stores the coordinates of the first member, that serves as a reference frame, and averages relative vectors to that. Then, to keep track of the displacements, the 6-bit flag of the first particle[12] is encoded into the 32-bit integer that saves the number of particle members of the halo. This is possible because haloes will always have less than $2^{24} = 16777216$ particles: the 24 rightmost bites suffice to store that integer and the remaining 8 bits in the left are unused. The 6-bit flag is stored there. Therefore, a 32-bits integer `n` can contain simultaneously the number of particles and the 6-bit flag, which can be extracted using some bitwise operations in the C language: the 24 rightmost bits (i.e., the number of particles) are extracted with the AND operator `n & 0x00ffffff`, and the 8 leftmost bits (i.e., the 6-bits flag) are recovered by shifting the variable 24 bits to the right as `n >> 24`.

[12] Only the displacement of the first particle of the group is relevant, since the halo center of mass will use that as a reference frame. If the first particle has a shift associated, it will be applied to the whole halo at the end of the FoF algorithm. The shifts of other particles of a halo do not affect the FoF algorithm since it assumes periodic boundary conditions and opposite box edges are linked.

```
Add buffer zones to χ1 and χ2;
for (ireplica = 0 → n_replicas - 1) do
    for ( i = 0 → n_particles - 1) do
        if (distance to any box edge {x,y,z} < fof_buffer) then
            Ghost particle candidate: start a loop over neighbouring replicas
            including the local;
            if (particle within the shell including vel.&fof buffer) then
                Compute crossing time;
                Drift particle;
                if (particle within the shell including fof buffer) then
                    if (Not a ghost particle in multiple replicas) then
                        Save particle position wrt. the origin of the local replica;
                        Compute 6-bit flag and store it in the ID;
                    else
                        Set 7th bit flag to 1 and save it in the ID;
                        Skip other neighbouring replicas and jump to next particle;
                    end
                end
            end
            Exit loop over neighbouring replicas;
        end
    end
    Find FoF haloes (using periodic boundary conditions);
    Write FoF haloes (shifts of particles are restored);
end
```

FIGURE 4.7: Algorithm for producing catalogues of FoF in the light cone. Only the variations with respect to the basic algorithm (see Fig. 4.2) are shown here, depicting how buffer zones are constructed before calling the FoF routines.

Note, however, that a ghost particle might correspond simultaneously to multiple neighbouring replicas. This is the case of the four corners in Fig. 4.5. Then, instead of adding 6 extra bits for each occurrence, a flag is stored in the 7th bit. For those particles, the code will check at the end all the possible $3^3 = 27$ replicas and select only those for which the position falls inside the shell.

Fig. 4.7 shows the extended algorithm to build FoF in the light cone. The parameter `fof_buffer` controls the width of the buffer. Note that each replica is built independently and once its particle distribution has been computed it calls the FoF code. First to find the haloes and right after to write them in disk. In the latter routine is where the necessary shifts are performed to the halo positions, just before writing.

That algorithm to apply the necessary shifts to the haloes before being written to the catalogues is given in Fig. 4.8. First, a shift is applied if the 6-bite flag is different than zero. Haloes not inside the local replica nor inside the shell are discarded. If the reference particle of the halo is present in multiple replicas, there is an additional loop that iterates over all the neighbours. When the correct replica is found, the loop is exited (because the halo can belong only to a single replica).

```
for ( i = 0 → n_haloes - 1) do
    if (First particle of the halo is not a ghost in multiple replicas) then
        if (shift flag ≠ 0) then
            Add ±L_box
        end
        if (not inside box) then
            Skip halo;
        end
        Add displacement related to the position of the replica's origin;
        if (not inside shell) then
            Skip halo;
        end
        Copy halo coordinates;
    else
        if (shift flag ≠ 0) then
            ±L_box
        end
        for ( j = 0 → n_neigh_replicas - 1) do
            if (not inside box) then
                Skip halo;
            end
            Add displacement related to the position of replica's origin;
            if (not inside shell) then
                Skip halo;
            end
            Copy halo coordinates;
            Exit loop over neighbouring replicas;
        end
    end
    Write this halo;
end
```

FIGURE 4.8: Algorithm to apply the necessary shifts to the haloes before are added to the catalogue. The code checks wether haloes have to be shifted by $\pm L_{box}$ in any of the dimensions according to flag variables and saves those which are inside the volume being built.

4.4 Numerical performance

When producing numerical simulations, there are usually three main limitations coming from the hardware resources that threaten the scalability to larger problem sizes. First is the availability of computing time in the hosting machine or, in other words, the number of dedicated CPU-computing hours. Another constraint is the memory available per node. In a parallel code relying on MPI communications that is run on a distributed memory machine, each process can access a memory size given by the memory per node divided by the number of active cores in it. That limit sets the maximum allocatable memory per process, which means that the larger the problem size, the higher the division into smaller pieces. Of course, this requires implementing efficient parallelization strategies that scale to the desired configuration. And lastly, the volume of the data products generated is another limiting factor, first due to the time dedicated to I/O during the run and also once the simulation has finished. To begin with, data has to be transferred from the supercomputer to local hard drives or a data storage centre. For illustration, to copy 1 TB per day implies at least a sustained bandwidth of 10 MB/s. And also, because any further post-processing or the posterior analysis are data-intensive operations that take a comparable amount of time and resources than the production of the simulation itself.

This section reviews how `ICE-COLA` performs in the three aspects aforementioned and which strategies were implemented in order to use resources efficiently.

4.4.1 Running time

During the build up of the light cone there are large particle load imbalances between processes. The volume intersected by a shell and the slab assigned to a given process can vary substantially, as can be seen in Fig. 4.5, where slabs in the right have 2-3 times more particles in the light cone than the leftmost. The situation would even worsen with volume decompositions that produce more compact regions. For that reason, some processes may be idle while others are still working on the current box replica.

In the light cone simulations developed for this thesis, the process with rank 0 was idle during 50% of the total time spent in the light cone. There is no simple solution for that problem and optimizing that imbalance is left for future versions of the code. The effort so far has been focused on trying to minimize the number of operations carried out for each particle. With that spirit, § 4.2.1 explained how time varying quantities are computed only once and are interpolated later on. Distance checks are done for all the particles at each time step, except for replicas that do not intersect the shell and are skipped. The crossing time and the drift operator are computed just for the relevant particles (see Fig. 4.2). Due to the buffer zones, this might happen multiple times at the end of the simulation for particles in the buffer zones.

The code can build as many replicas as desired, since they are independent. But for large number of replicas the light cone construction will dominate the time budget of the simulation. The exact numbers depend on the particular code parameters. For instance,

the simulations developed for this thesis, that generated two-dimensional maps and halo catalogues in 64 replicas, dedicated 44% of the time for building the light cone. Inside that routine, the time was spent in the following way: 13% to find haloes, 12% to write catalogues, 5% to communications between processes inside the FoF routines (that could be added to the previous 13%) and only 1.6% to the actual algorithm of the light cone, that is, the selection of particles, computation of crossing times and evaluation of the temporal operators. The remaining fraction was spent in smaller parts or was just idle time.

Regarding the particle light cone catalogues, the time consumption is mainly driven by the I/O performance, which may vary considerably in different architectures and file systems. The results presented next correspond to a simulation that used 1024 cores on MarenostrumIII, which uses the IBM General Parallel File System. The parameter of the code `n_files` was set to 24 and the code was compiled with gcc[13] using the OpenMPI[14] libraries[15]. There were 129 calls to the light cone routine that produced data, which consisted in one octant and 8 box replicas in total. Each call generated a volume of data in the range between $1 - 15GB$ in most of the cases. The I/O bandwidth achieved varied considerably for each call but typically it was in the range between $5 - 10GB/s$ when the total data size was $\gtrsim 10GB$. For smaller data volumes the bandwidth was lower, but in all the 129 calls the wall-clock time dedicated purely to I/O was always less than 3 seconds, to be compared to the 40 minutes that took to complete the simulation.

4.4.2 Memory usage

The peak of memory consumption in `COLA` happens during the evaluation of forces in the PM algorithm. This is specially true if the mesh is finer than the mean particle distance. Other parts of the code can take advantage of that and use memory that has been freed. In particular, `Parallel COLA` allocates two large shared memory blocks that are re-used by different routines separated in time:

- `mem1`: used by the initial conditions, the density grid of the PM, the k-d tree of the FoF, and the particle buffer for moving particles between processes.
- `mem2`: used by the density grid in Fourier space of the PM and the snapshot data structure.

The size of each block is set to the maximum memory required for the routines of each group. Shared memories are allocated at the beginning of the code, are used elsewhere and are freed at the end. The light cone routine in `ICE-COLA` does not allocate more memory, all the data is stored in the shared memory blocks. Depending on the kind of catalogues that are requested, the memory needs vary:

[13] https://gcc.gnu.org/

[14] https://www.open-mpi.org/

[15] The support of the machine recommended the usage of the following MPI-I/O hints, that supply extra information to the MPI implementation for a better performance: `striping_unit= 4194304` and `romio_cb_write=enable`.

- **Dark matter particle distribution**. Particle positions interpolated at the crossing time are stored in the already existing snapshot data structure.
- **Projected density maps**. It is necessary to store multiple Healpix maps. Particle positions are not kept, i.e., the snapshot data structure is not necessary.
- **Halo catalogues**. A snapshot is build similarly as for the dark matter particle outputs, which now includes buffer regions as well. In addition, space is needed for the k-d tree of the FoF.

The memory set-up in `Parallel COLA` is already suitable to construct the snapshot and the FoF with particles in the light cone, which use `mem2` and `mem1` respectively. However, new room is needed for the Healpix maps. If only projected density maps are produced, then they can be saved in the shared memory blocks. However, this is not possible if multiple outputs are requested since `mem1` and `mem2` are already used. Some memory is still left at the end of the shared memory blocks, but are non-contiguous.

To avoid that, the memory layout has been slightly changed in `ICE-COLA`. A single shared memory block `mem` is allocated at the beginning, and it is explicitly broken into the two pieces `mem1` and `mem2` for compatibility with the rest of the code. The light cone routine uses `mem`, inside which memory blocks are located one after the other in a compact configuration. For example, the simulations of Chapter 5 requested 1728 and 864 MB/process for `mem1` and `mem2` respectively and hence `mem` had 2592 MB. From that, the FoF used 708 MB, the snapshot 320 MB and the Healpix maps[16] 1346 MB, leaving just 218 MB unused. In this way, all the three kind of outputs can be generated simultaneously without increasing the memory usage and running the light cone routine just a single time at each `COLA` time step.

4.4.3 Disk storage requirements

Catalogues of dark matter particles in the light cone turn out to be extremely massive, as has been already pointed out. Given the volume inside the light cone, the mean particle density and that 24 bytes are written per particle, the data volume can be computed as

$$\begin{aligned} \mathrm{MEM_{DM}} &= \frac{4\pi}{3}\left(\chi(z_{max})\right)^3 \frac{2^{n_{dim}}}{8} \times \frac{N_{part}}{L_{box}^3} \times 24\,\mathrm{bytes} \\ &\simeq 5.9 \times \left(\frac{\chi(z_{max})}{3000\,h^{-1}\,\mathrm{Mpc}}\right)^3 \times \frac{2^{n_{dim}}}{8} \times \left(\frac{m_{part}}{2.9\times 10^{10}\,h^{-1}\,\mathrm{M}_\odot}\right)^{-3} \mathrm{TB}. \end{aligned} \tag{4.9}$$

So an all-sky simulation having the reference scaling values (corresponding to a maximum redshift of 1.4), $\sim 6\,\mathrm{TB}$ of data are generated. This might be affordable only for a single realization, or a few at most, in which detailed information of the matter field is necessary at full resolution. But since `COLA` is an approximate method, very accurate

[16] Using the value $n_{\mathrm{side}} = 2048$, a single map has ~ 50 million 4-byte pixels, or 192 MB. Up to 7 shells in the light cone might be placed between consecutive `COLA` time steps, which adds up to 1346 MB in total.

results are out of the scope and presumably one will run ensembles of many realizations. Then, it is necessary to generate only the two derived and compressed data formats.

The disk space used by the projected matter maps depends on the number of 4-byte angular pixels and the radial bins (i.e., on the size of the spherical cell of the light cone):

$$\begin{aligned} \mathrm{MEM_{HP}} &= 12 n_{side}^2 \times n_{z,bins} \times 4\mathrm{bytes} \\ &\simeq 50 \times \left(\frac{n_{side}}{2048}\right)^2 \times \frac{n_{z,bins}}{265}\,\mathrm{GB}. \end{aligned} \tag{4.10}$$

That means that it would be possible to have projected matter density fields for 1000 realizations in a disk space of 50 TB. Furthermore, it will be studied in the future if a broader binning in redshift, i.e., the width of the light cone shells projected onto pixelated maps, does not have an impact on the accuracy of the clustering statistics modelled. In that way, a reduction of a factor of 2 or even more in the storage would be possible.

Finally, let's turn into the halo catalogues. Their size depends strongly on the minimum number of particles to write a halo. Haloes with few tens of particles are not well resolved, but if the simulation is populated by galaxies in a later stage using for instance a Halo Occupation Distribution method, undersampling issues (e.g., the Warren correction) might be corrected and therefore make sense to push the catalogue towards low mass haloes (Howlett et al., 2015a). In the current code version, halo catalogues are written in ascii format containing the following fields: number of particles, centre of mass position, and halo bulk velocity[17]. `ICE-COLA` is already able to produce halo catalogues in binary format for comoving outputs, that allows a compression factor of 3 of the file sizes. In the light cone, however, there is the additional complication that data is appended to existing files. This has still not been resolved and will be implemented in a future version. Following with the same reference light cone simulations, cutting at haloes with 20 or more particles translates into 460 million haloes and 37 GB for an all sky light cone, which is comparable or somewhat smaller than the data of the projected matter maps. In order to extrapolate to other mass cuts, the size scales roughly inversely proportional to the minimum number of particles[18]. Therefore, the rule-of-thumb for the data volume for a halo catalogue that starts at redshift 1.4 is

$$\mathrm{MEM_{FoF}} \simeq 37 \times \left(\frac{N_{part}}{2048}\right)^3 \times \left(\frac{N_{min}}{20}\right)^{-1} \times \frac{2^{n_{dim}}}{8}\,\mathrm{GB}, \tag{4.11}$$

which assumes that there is no dependence with the particle mass. This will be reduced by a factor of three when halo catalogues are written in binary.

[17] The current version of the code does not keep the full halo detailed information that is necessary to compute other halo statistics such as the velocity dispersion. In a future version this issue shall be adressed.

[18] Scaling that holds for masses below $10^{13}\,h^{-1}\,\mathrm{M}_\odot$, where the abundance of haloes decreases approximately as the inverse of the mass.

5

ANALYSIS OF LIGHT CONE CATALOGUES

In this chapter we present light cone catalogues produced with the new ICE-COLA features described in Chapter 4 and discuss a first validation of them in terms of projected dark matter and weak lensing observables. First, § 5.1 summarizes some basic tests on the three kind of light cone catalogues produced by the method in order to provide a first validation. A more elaborated analysis is presented later in § 5.2, where it is described a full pipeline developed to model weak lensing and it is shown the accuracy with which observables are reproduced. Chapter 6 extends this analysis in combination with halo catalogues.

5.1 Basic validation of the light cone catalogues

5.1.1 Particle light cone

The particle light cone (see § 4.3.1) is the lower level and more massive light cone output implemented in ICE-COLA. Given the great advantages of using the higher-level data formats developed later in time (see § 4.3.2 and § 4.3.3) in terms of data volume and post-processing computing requirements, the particle light cone had a limited use. However and for completeness, some tests are presented next that demonstrate that this feature of the code is validated.

Ten runs were developed with particle light cone. All the parameters used are the same as the the other light cone simulations, except that the sky area covered was one octant instead of the full sky (see § 2.4). Fig. 5.1 shows a slice of $1500\,h^{-1}\,\mathrm{Mpc}$ wide and

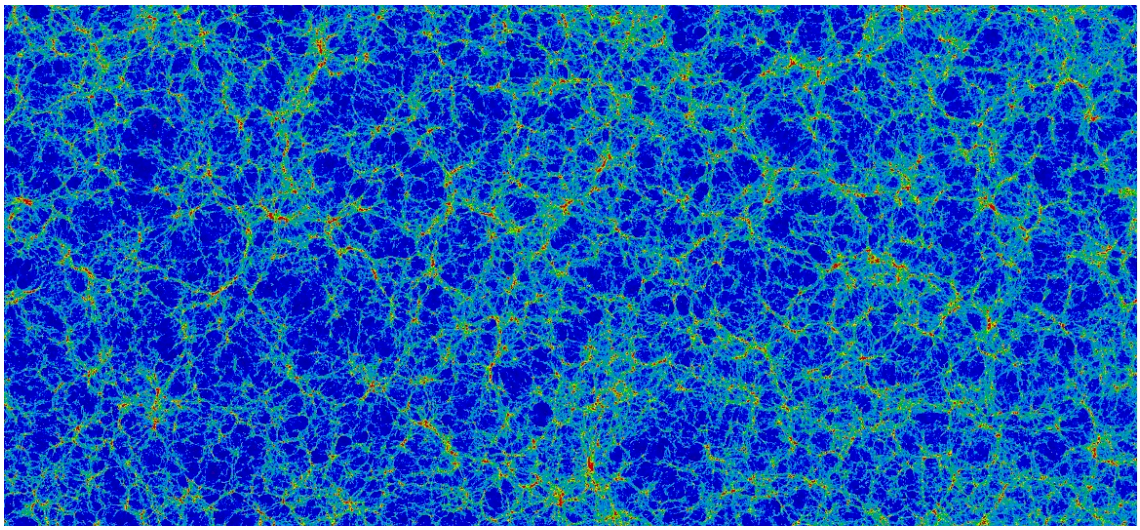

FIGURE 5.1: A slice of $1500 \times 700 \times 25\, h^{-1}\,\mathrm{Mpc}$ of the particle light cone, displaying the cosmic web. See the text for more details. A very high resolution version of the image (of 37MB size) can be downloaded from https://github.com/albertizard/thesis_public_material/blob/master/dm_particle_lc_slice.pdf

a thickness of $25\, h^{-1}\,\mathrm{Mpc}$, showing the expected cosmic web resulting from non-linear gravitational growth. Particles were interpolated onto a grid of 3000^2 cells and convolved with a 3×3 pixel gaussian kernel with a beam equal to the pixel size. The number count at each pixel is then transformed into a colour value using the transfer function

$$\mathrm{counts} \to \mathrm{minimum}(\mathrm{counts}^{0.6} + 3, 110), \tag{5.1}$$

which enhances colour variations at low density regions and samples evenly the wide dynamic range. Also, a saturation value is set in order to limit colour variations at very high density regions.

Aside from this visual inspection, projected matter maps were also built from these catalogs at post-processing and compared to the ones obtained on-the-fly from simulations with identical parameters and initial conditions. The results matched perfectly, only with small differences arising from accumulated errors due to numerical precision. In particular, number counts were identical in 99.87% of the pixels, while in the remaining differences were at most 2 counts (take into account that the mean number of counts was 50 and the maximum 3000).

5.1.2 Projected density maps

In §4.3.2 we introduced the two-dimensional projected matter density maps using the Healpix discretization. The main scientific motivation for constructing them is to model weak lensing observables, as it is done in §5.2. As a first step towards modelling galaxy angular clustering, this sub-section shows two-point clustering statistics measured from the dark matter maps directly obtained from the simulation. In particular, we choose

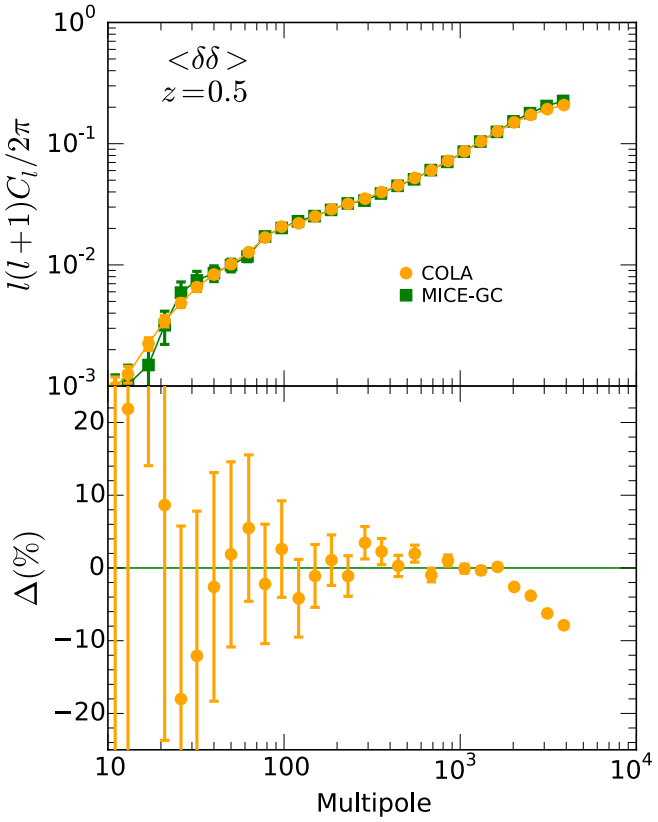

FIGURE 5.2: Matter angular power spectrum for a shell at mean redshift of $z = 0.5$ and a width of $\Delta z = 0.1$. The lower panel shows the relative differences between COLA and MICE-GC, that are within the error bars up to $l \sim 2000$ and within 10% up to $l = 4000$.

a fiducial case where maps are stacked into a slice centred at redshift 0.5 and having a width of 0.1. Fig. 5.2 shows the matter angular power spectrum measured for that Healpix map. Circles and squares represent COLA and MICE-GC respectively. The shot noise contribution has been subtracted according to the formula

$$C_l^{noise} = f_{sky}\frac{4\pi}{N}, \tag{5.2}$$

where N is the number of objects in the sample and f_{sky} is the fractional area of the sky of the catalogue (1 if it is all-sky). At small scales it represents a 1% correction for COLA and a factor 8 larger for MICE-GC, since the sample used in the latter is restricted to one octant of the sky. Also, the errors in the measurements of angular power spectra can be computed for the case of a Gaussian density field by considering the number of modes sampled at each l-bin

$$\sigma(C_l) = \sqrt{\frac{2}{f_{sky}\Delta l(2l+1)}}(C_l + C_l^{noise}), \tag{5.3}$$

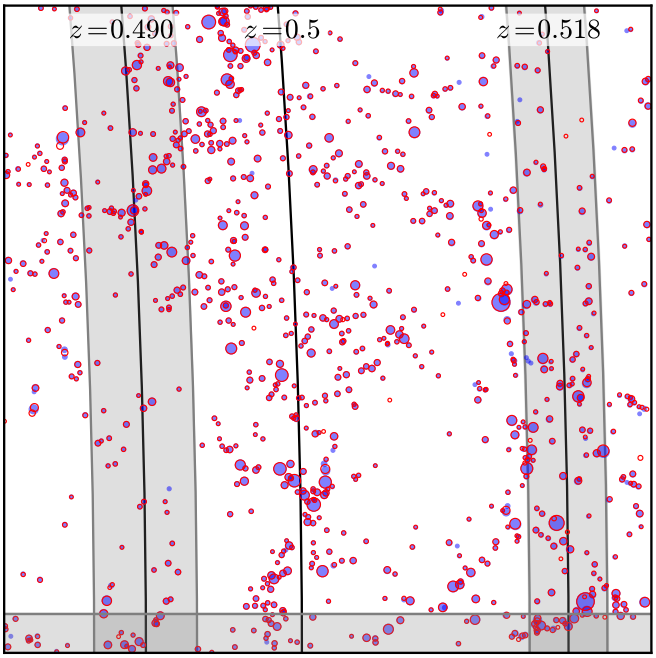

FIGURE 5.3: Slice $100\,h^{-1}\,\mathrm{Mpc}$ wide and $40\,h^{-1}\,\mathrm{Mpc}$ thick comparing FoF haloes from a catalogues in the light cone (red open circles) and a comoving output at redshift 0.5 (blue filled circles). The volume is centred at a radius of $1350\,h^{-1}\,\mathrm{Mpc}$ from the origin, where the time of both catalogues coincides. Grey areas denote light cone buffer zones, either around shell (vertical shaded areas) or box (horizontal shaded) edges.

where Δl is the bin width. For the ratio displayed in the lower panel, the error bars are computed adding in quadrature the errors coming from each simulation.

The agreement is very good up to $l = 2000$, as can be seen better in the lower panel displaying the relative differences, that are within 10% up to $l = 4000$. This is to be expected as dynamics within haloes are not fully resolved in COLA, which contribute to the power at these scales (see §4 and fig. 7 in Fosalba et al., 2015a for a discussion on the mass-resolution effects).

5.1.3 Halo catalogues in the light cone

Catalogues in the light cone can be compared to catalogues coming from snapshots in the small range of spatial and temporal coordinates where they almost coincide (i.e., assuming that an observer is located somewhere within the snapshot volume). For instance, the part of the light cone around redshift 0.5 matches with comoving outputs taken at the same redshift for objects at a radial distance $\chi(z = 0.5) \sim 1350\,h^{-1}\,\mathrm{Mpc}$. Fig. 5.3 compares the distribution of haloes in the light cone (red open circles) with haloes from a comoving catalogue at $z = 0.5$ (blue filled circles). At object by object the agreement is excellent in both the identification of haloes, their position and mass (which is related to the radius of the circles drawn). There are very few cases of haloes without correspondence in the other catalogue, or groups that appear as two distinct objects in one case but

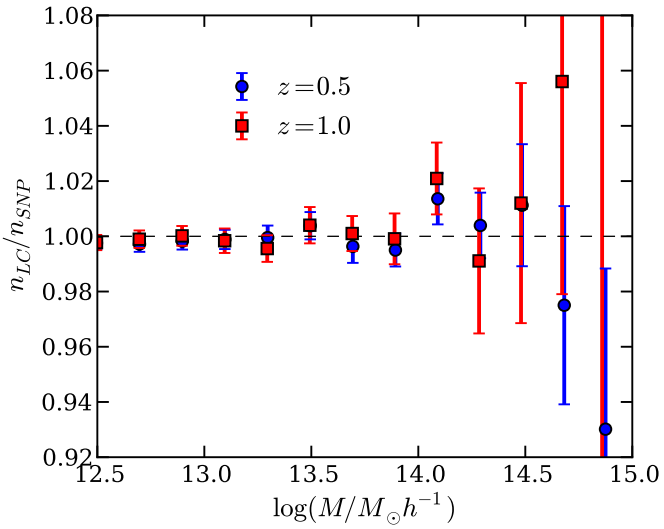

FIGURE 5.4: Comparison of the mass function measured from light cone or comoving catalogues. Error bars display Jack-Knife estimates from the measurements of the comoving catalogue. Differences are very small or within the error bars, what validates the algorithm for generating FoF in the light cone.

merged into a single halo in the other. Actually, some of these differences may be real, since the time difference between both catalogues (comoving and light cone) is as much as $\lesssim 200 M$years: the only position where the time exactly matches the light cone solution is given by the line at $z = 0.5$. Also, buffer zones around shell and box edges (see 4.5) are indicated as grey shaded areas, and by visual inspection there are no significant differences between both catalogues in the limit where they intersect.

A more quantitative test is to compare the abundance of haloes between a comoving catalogue and a shell of the light cone that has the same mean redshift . Fig. 5.4 displays the ratio of the corresponding mass functions, at redshifts 0.5 and 1.0[1]. The width of the shell in the light cone is chosen so that the volume is the same as the comoving box. The agreement is excellent again, with differences within the 1% for most of the mass range and within the error bars estimated by jack-knife re-sampling.

The angular power spectrum of haloes is discussed in Fig. 6.1 in Chapter 6, where additional measurements of the clustering and cross-correlation of halo samples with weak lensing are shown.

5.2 Modelling weak lensing in cosmological simulations

The shear and convergence fields constitute basic weak lensing quantities that characterize the Jacobian matrix (see equation 1.32 in Chapter 1) from which other lensing fields can be derived. In fact, they are both related to the effective lensing potential as is shown

[1]Redshift 0 is impossible because the light cone would have no volume.

in §1.4, which is a central quantity for modelling weak lensing distributions. Cosmological weak lensing simulations aim at producing these observables, which in turn can be used to model multiple observational tools that probe weak lensing (e.g., shear mapping, magnification, mass mapping, aperture mass, peak statistics, lensing of the cosmic microwave background... with two/three-dimensional statistics or in a tomographic approach).

Weak lensing can be implemented in simulations via ray-tracing techniques, which follow the ray propagation from the source to the observer along the perturbed path, or rather the back propagation from the observer to tge source plane in practise (see e.g. Blandford et al., 1986; Jain et al., 2000; Das et al., 2008; Teyssier et al., 2009; Li et al., 2011). This involves intensive computations because the deflection angle needs to be constantly updated as the ray travels in order to determine the geometry at each encounter of the multiple-lens system.

However, if the Born approximation is assumed (see §1.4.2), integrations along the straight line-of-sight are much faster as they rely on the unperturbed path from previous lens planes along the line-of-sight (i.e., lens planes are uncorrelated) and have been successfully implemented in simulations (in which different discretization schemes can be used, see e.g. M. White et al., 2000; Fosalba et al., 2008; Kiessling et al., 2011a). This removes higher-order contributions and the coupling of lenses at different distances (Krause et al., 2010), but these represent sub-percent corrections for most of the relevant scales and can be neglected, specially in the case considered here in which an approximate gravitational simulation method is being used. In the Born approximation, convergence can be estimated by employing a discretized version of equation 1.31, that is, the projected matter density field at unperturbed photon positions weighted by the lensing kernel. Other weak lensing quantities, such as the shear, can be derived from the convergence (Kaiser et al., 1993). This is specially simple to implement in harmonic space for all-sky light cone simulations (Hu, 2000). Next it is explained the methodology employed in this thesis to convert projected density maps generated with `ICE-COLA` to lensing quantities. For this we follow Fosalba et al., 2008 but adapted to the peculiarities of `COLA`. Thanks to using all-sky light cones, there are no problems associated with edges when finite areas are sampled.

5.2.1 Production of convergence maps

The density field outputted by the simulation has been discretized into multiple angular pixels and radial bins. Therefore, the integral in equation 1.31 is converted into a sum running over the shells

$$\kappa(\boldsymbol{\theta}_i,\chi) = \frac{3H_0^2\Omega_m}{2c^2}\sum_j^{\chi_j<\chi}\delta(\boldsymbol{\theta}_i,\chi_j)\frac{(\chi-\chi_j)\chi_j}{a_j\chi}\,\Delta\chi_j\,x, \tag{5.4}$$

where $\Delta\chi_j$ is the width of the radial bins and the convergence map obtained inherits the same angular pixelization as the density maps. Fig. 5.5 shows a patch of $4\times4\,\mathrm{deg}^2$ of the

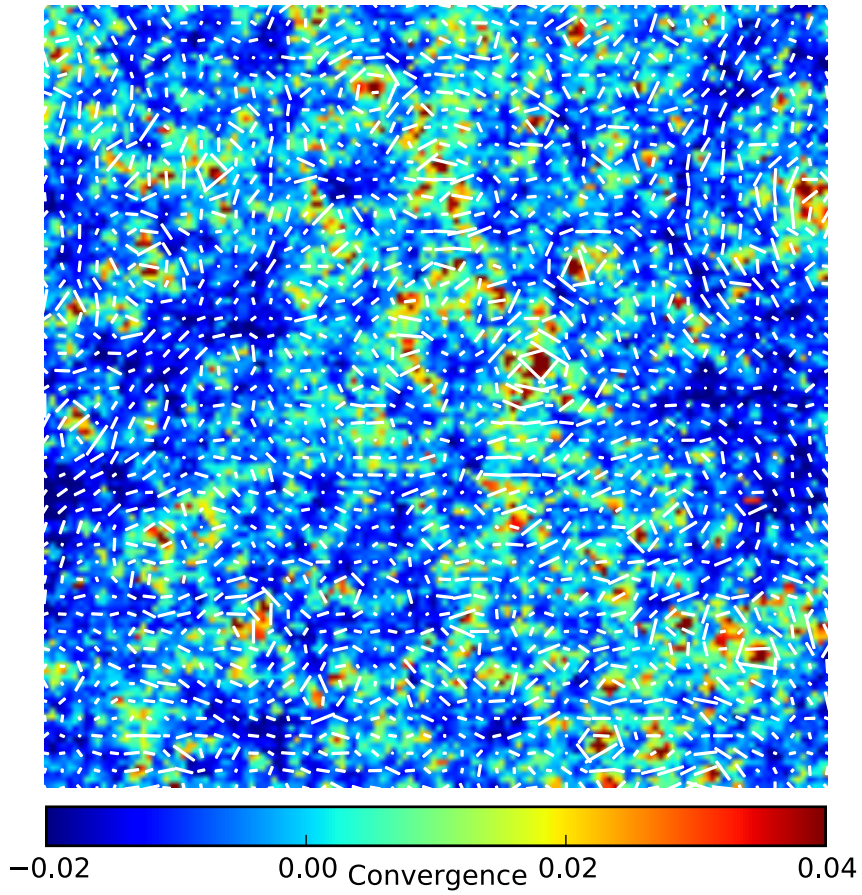

FIGURE 5.5: Color map of the convergence field for sources at redshift 1 in a patch of $4 \times 4\,\mathrm{deg}^2$. The distribution features large voids with slightly negative values and concentrated peaks due to massive haloes or the superposition of several of them at intermediate distances. White ticks indicate the orientation and amplitude of the shear field (see § 5.2.2), and as expected for a field without B-modes they are preferentially oriented tangential to the peaks of the convergence field and point towards underdense regions.

maps thus obtained, which basically shows a weak lensing-weighted version of the dark matter distribution (i.e., the cosmic web). Note that typical values of the convergence fluctuation field are at the percent level, what illustrates that weak lensing statistics have smaller amplitude than the underlying dark matter ones. Besides, white ticks show the shear field, which will be explained in § 5.2.2.

Auto-power spectrum

The convergence angular power spectrum can be measured by transforming the map to harmonic space via spherical harmonics

$$\hat{\kappa}_{lm} = \frac{4\pi}{N_{pix}} \sum_{i=1}^{N_{pix}} \kappa(\boldsymbol{\theta}_i) Y^*_{lm}(\boldsymbol{\theta}_i) \tag{5.5}$$

and averaging the square of the coefficients over the m-modes:

$$C_l^\kappa = \frac{1}{2l+1} \sum_{m=-l}^{l} |\hat{\kappa}_{lm}|^2. \tag{5.6}$$

The actual measurements can be compared to a theoretical prediction for the convergence power spectrum that can be derived starting from eq. 1.31 and using the Limber approximation (Limber, 1953) to compute the two-point statistics, which results in the expression (Kaiser, 1992; Kaiser, 1998)

$$C_l^\kappa(\chi) = \frac{9H_0^4\Omega_m^2}{4c^4} \int P(k = l/\chi', z) \frac{\chi - \chi'}{\chi^2 a^2} \, d\chi', \tag{5.7}$$

where $P(k, z)$ is the measured three-dimensional matter power spectrum. [2]

The shot noise contribution to the convergence power spectrum, $C_l^{\kappa,\,noise}$, that is corrected for, is given by replacing the power spectrum in equation 5.7 by the inverse of the mean number density of particles: $P(k, z) = 1/\bar{n}$.

Fig. 5.6 compares the convergence power spectra for a source redshift of $z_s = 1$ that are obtained both from the prediction 5.7 that uses the measured matter power spectrum (dotted curve) and from direct measurements of the angular spectrum (circles), including as well the fiducial values from MICE-GC (squares). First, note that C_l^κ measurements have more power than the prediction at small scales. This can be explained because the integral at equation 5.7, in the limit of low distances and high multipoles, needs to evaluate the matter power spectrum at larger wavenumbers than those for which it has been measured. Therefore it lacks some of the contributions at small scales and the power of the convergence is hence suppressed at high multipoles. Another deviation of the theoretical prediction is present at linear scales, where it has a slight excess of power with respect to the the linear prediction. However, they agree within sample variance errors.

In summary, the agreement between the actual measurements and the prediction from equation 5.7 is quite remarkable, which constitutes a first successful validation test of the convergence maps. Comparing now with MICE-GC, measurements are compatible up to $l \sim 10^3$, beyond which COLA recovers less power. The difference is $\sim 10(20)\%$ at $l = 2 \times 10^3 (4 \times 10^3)$. Note that theses are highly non-linear scales, roughly one decade above the transition scale to the non-linear regime (i.e., where the non-linear growth starts being significant). slightly more than one order of magnitude above the point where the linear prediction breaks down.

[2] In ICE-COLA this is measured on-the-fly at each time step of the simulation thanks to a modification of the PM routines that sums in radial k-bins the density field already transformed to Fourier space (the three-dimensional equivalent of equation 5.6).

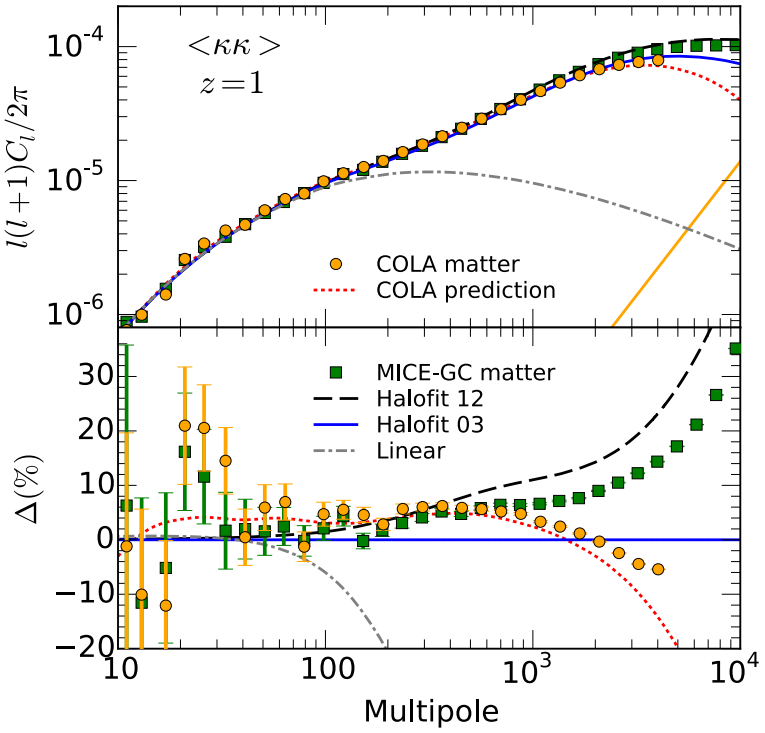

FIGURE 5.6: Angular power spectrum for the convergence map for $z_s = 1$. Different symbols represent: real measurements (circles), MICE-GC measurements (squares), theoretical prediction from equation 5.7 (dotted line), two non-linear predictions from Halofit (Smith et al., 2003 solid line) and revised Halofit (Takahashi et al., 2012 dashed line) and the linear case (dot-dashed). Shot noise error is denoted by the solid line in the lower right corner. COLA is in very good agreement with MICE-GC up to $l \sim 1000$ and deviates thereafter to 10% (20%) at $l = 2 \times 10^3$ (4×10^3). The theoretical prediction for COLA reproduces reasonably well the measurements (see the text why at small scales there is a small deviation).

Convergence-matter cross-power spectrum

The convergence field is largely correlated with the intervining mass distribution due precisely to the lensing effect. This is shown in Fig. 5.7 for the cross-correlation between the convergence at $z_s = 1$ and a lens sample at $z_l = 0.5$ (having a width of $\Delta z = 0.1$). The matter catalogues reproduce well the signal up to $l \sim 2000$ and there is a lack of power at smaller scales, although deviations are within 10% at all scales. The agreement with MICE-GC is better in the cross-correlation than the auto-power spectrum of the convergence (see Fig. 5.6). In the cross power, only the contribution to the convergence field coming from the distribution at the lens redshift is relevant. Therefore, the accuracy in that case depends on how well scales at this lens redshift can be resolved. However, the convergence has contributions as well from the low-redshift structure, which span larger angular sizes: higher multipoles associate with larger wavenumbers in the matter power spectrum, hence more non-linear and hard to model in COLA. For that reason, the convergence auto-power spectrum shows larger deviations than the cross component at

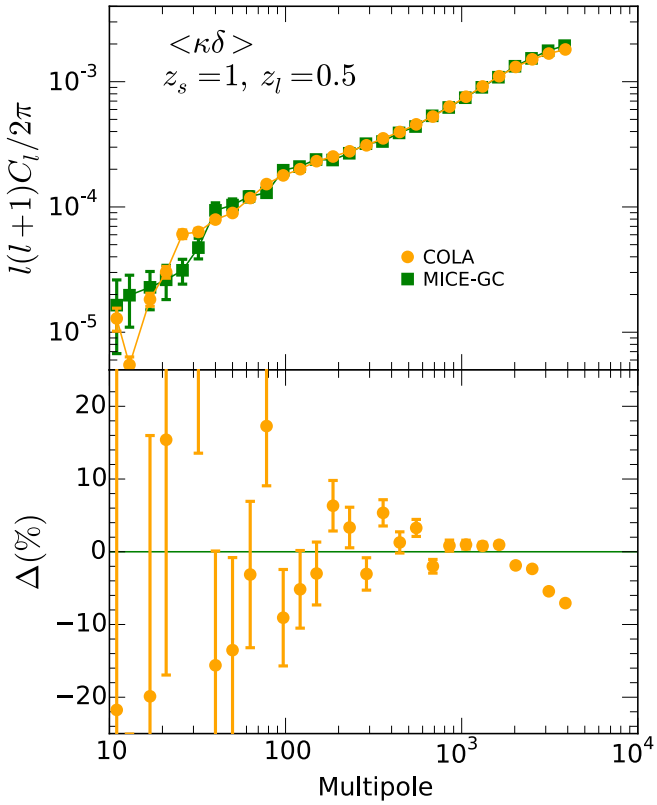

FIGURE 5.7: Cross-power spectrum between the convergence at $z_s = 1$ and the density field at $z_l = 0.5$ (with a width for the latter of $\Delta z = 0.1$). Deviations are found only at scales above $l = 2000$, as in the case of the matter power spectrum (see Fig. 5.2).

a given non-linear scale.

5.2.2 Production of shear maps

The shear field is a spin-2 tensor that in the complex notation can be expressed as two polarization Stokes parameters: $\gamma = \gamma_1 + i\gamma_2$. In the all-sky formalism, the decomposition in harmonic space is written in terms of the spin-2 spherical harmonics ${}_{\pm2}Y_{lm}$ (Hu, 2000)

$$ {}_{\pm}\hat{\gamma}_{lm} = \hat{\epsilon}_{lm} \pm i\hat{\beta}_{lm} = \sum_{lm} \gamma(\boldsymbol{\theta})\, {}_{\pm2}Y^*_{lm}(\boldsymbol{\theta}), \tag{5.8} $$

where in the first equality we have introduced the so-called E- and B-modes as $\hat{\epsilon}_{lm}$ and $\hat{\beta}_{lm}$ respectively. In the Born approximation, the shear and convergence fields are originated by a single scalar potential (see § 1.4.4) and cannot have a curl component due to parity symmetry of the weak lensing fields. These can be decomposed into scalar and vector modes, or E-modes and B-modes respectively in analogy with electromagnetism. Therefore, in this approximation we have that the B-mode shear is zero: $\beta_{lm} = 0$.

The shear and convergence are expressed as derivatives of the lensing potential in configuration space (see equations 1.33 and 1.34). By transforming to harmonic space,

these become multiplications by multiple wavenumbers and one can efficiently transform from one weak lensing quantity to another. For instance, the shear and convergence coefficients are related as

$$\hat{\epsilon}_{lm} = -\sqrt{\frac{(l+2)(l-1)}{l(l+1)}}\hat{\kappa}_{lm}, \tag{5.9}$$

where we assume that B-modes are null. Note that in the small angle limit ($l \gg 1$) the coefficient is just -1 and the power spectra of both quantities coincide (see Hu, 2000).

Finally, the Stokes parameters of the shear field can be expressed as the inverse of the transformation 5.8, which in the convention adopted by Healpix (and in the absence of B-modes) is

$$\gamma_1 = -\sum_{lm} \hat{\epsilon}_{lm} X_{1,lm} \tag{5.10}$$

$$\gamma_2 = \sum_{lm} i\hat{\epsilon}_{lm} X_{2,lm}, \tag{5.11}$$

with $X_{\frac{1}{2},lm}$ defined as the combination: $X_{\frac{1}{2},lm} = ({}_2Y_{lm} \pm {}_{-2}Y_{lm})/2$.

In summary, the steps that are taken in order to produce a shear map given a convergence map are the following. First, the convergence is transformed to harmonic space according to equation 5.5. Then the E-mode shear is obtained using 5.9 and finally the two components in angular space are given by the inverse harmonic transforms 5.10 and 5.11. The transforms are done using Healpix (in particular Healpy[3], its Python version) and the full post-processing pipeline can be run on a modest personal computer, since the memory and processing requirements are moderate. Using the value $n_{side} = 2048$ for the Healpix maps, that produces ~ 50 million pixels, a map of float values needs 192 MB (or twice if it is stores complex numbers), which is an order of magnitude less than the memory that a nowadays personal computer has. An harmonic transform takes 9 minutes in a laptop equipped with a 2.8 GHz Intel Core i7 single processor. Considering that there are 265 maps in total and that two transforms are needed to produce a shear map, the total time ascends to 3 days. This is proportionally reduced if many maps are read simultaneously and processed each one by a different core, enabling doing the complete post-processing in a single day. A visual impression of the final shear maps is displayed in Fig. 5.5 by the white ticks. They are oriented perpendicular to the convergence peaks, as it is expected for a field without B-mode (i.e., in the weak lensing limit), which would generate instead swirl patterns. At empty regions, the field is radially pointing towards under-densities. A more extensive validation of the shear is shown in Chapter 6, where we present two-point clustering statistics for halo mock catalogues with weak lensing quantities.

[3] https://healpy.readthedocs.io/en/latest/

6

HALO CATALOGUES FOR WEAK LENSING AND CLUSTERING

We have described the capabilities of ICE-COLA to generate halo catalogues in the light cone (see §4.3.3) and two-dimensional projected matter maps (see §4.3.2), both at run time. The latter are converted to weak lensing quantities following the method described in §5.2. The last step of the pipeline is to combine both in a catalogue of haloes with weak lensing properties. To do so, each halo is assigned the convergence and shear values of the pixel where its centre of mass position lays. It is left for future work using a better scheme by interpolating the pixelated maps to the particular position of each halo. Furthermore, the same approach could be used for galaxies in an HOD pipeline for the assignment of weak lensing properties.

In this chapter, two-point statistics of the halo catalogues are presented in order to validate the full simulation method producing light cone catalogues and the post-processing pipeline, as well as to test the accuracy in predicting weak lensing observables. For that purpose, two halo samples are built: the source sample at redshift 1 and the lens sample at 0.5. Both are all-sky, have a radial width of $\Delta z = 0.1$ and include those objects with 100 or more particles (i.e., a mass of $3 \times 10^{12}\, h^{-1}\, \mathrm{M}_{\odot}$). This results in 6M and 11M objects for the lens and source samples respectively. The same criteria are used to draw the reference N-body samples from MICE-GC, except that those are for one octant of the sky.

6.1 Halo angular power spectrum

The halo angular power spectrum of the lens sample is shown in Fig. 6.1. The lower panel displays the halo bias as the ratio $\sqrt{<\delta_{\mathrm{h}}\delta_{\mathrm{h}}>/<\delta_{\mathrm{m}}\delta_{\mathrm{m}}>}$, where the sub-indices h and m

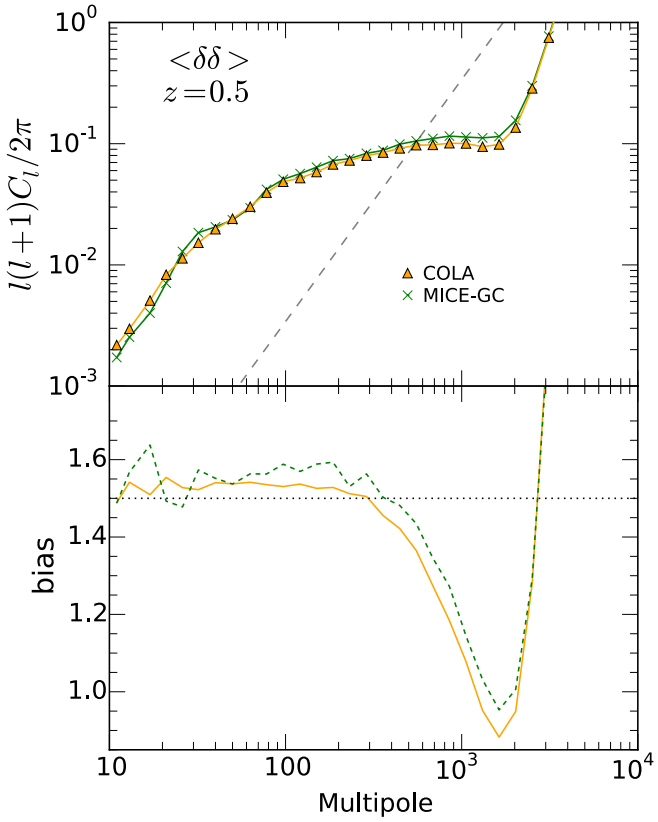

FIGURE 6.1: Angular power spectra for the lens halo sample. The lower panel shows the halo bias as measured with $\sqrt{<\delta_h\delta_h>/<\delta_m\delta_m>}$, which is close to 1.5. The shot-noise term is shown as a dashed line in the upper panel and dominates for $l > 500$, where both COLA and MICE-GC display a scale-dependent bias with a similar behaviour.

refer to haloes and matter repectively. At large scales it is recovered a value for the linear halo bias of ~ 1.5, indicated by a dotted horizontal line. However, it is $\sim 3\%$ lower for COLA, in concordance with results discussed in § 3.5.

The shot noise term has been corrected as equation 5.2[1]. In fact, it dominates the measurements for $l > 500$, where both COLA and MICE-GC display a scale-dependent bias with a similar behaviour.

6.2 Halo-convergence cross-power spectrum

Fig. 6.2 shows the halo-convergence cross-power spectrum. The halo bias, estimated as the ratio $<\kappa_h\delta_h>/<\kappa_m\delta_m>$, is consistent with a value of 1.5 in both samples and with the value from the halo power spectrum (see Fig. 6.1). At small scales (i.e., high l's) the measurements feature a scale-dependent bias, with a similar trend in both COLA and MICE-GC. For instance, $l = 2000$ is associated with a scale of $2\,h^{-1}\,\mathrm{Mpc}$ at $z = 0.5$, only a factor ~ 4 larger than the mean size of the haloes in the sample.

[1]Note that the shot-noise correction can be improved by modelling the halo exclusion effect (see e.g., Baldauf et al., 2013).

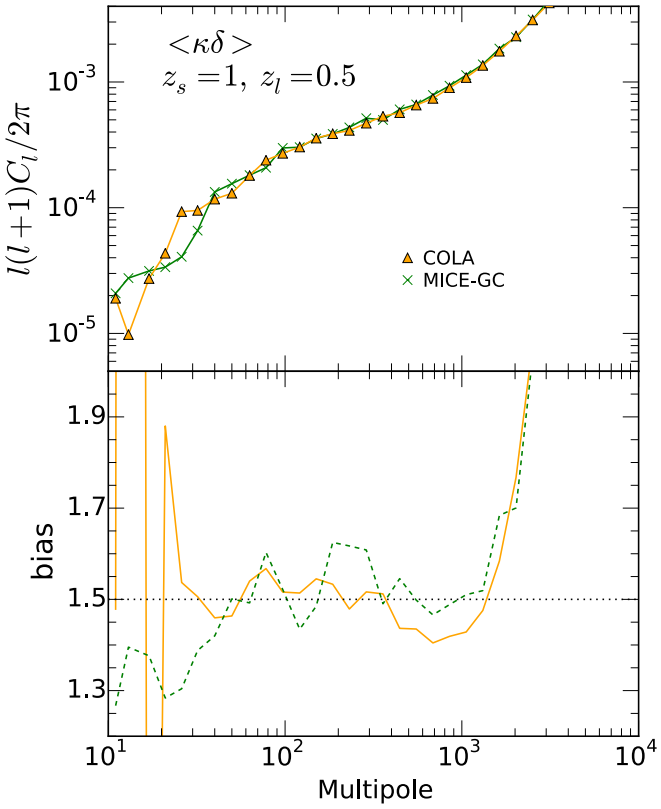

FIGURE 6.2: Halo-convergence cross-power spectrum for the lens and source samples. The halo bias, displayed in the lower panel, coincides to 1.5 in both samples and with the value from the halo power spectrum (see Fig. 6.1). Scale-dependent bias shows up in a similar fashion for both samples at $l > 1000$, in the highly non-linear regime.

6.3 Convergence auto-correlation

In § 5.2.1 we shown the clustering in harmonic space for the convergence maps. Next we compute lensing second-order correlation functions in angular space using the public code Athena[2] based on a two-dimensional tree. Fig. 6.3 shows the convergence field correlation function for the source sample at $z_s = 1.0$. COLA reproduces the reference MICE-GC measurements at large scales and down to separations of only 3 arcmin. Despite the fact that the pixel size employed is 1.7 arcmin (4 times smaller in MICE-GC), the power is still within 25% at 1 arcmin. Hence, it is not clear whether deviation are due only to the limited capability of COLA to resolve small scales or to the angular resolution of the Healpix maps. It is also remarkable that the signal is well reproduced to separations that are ~ 5 times smaller than the scale where non-linearities arise. These results are in agreement with Heitmann et al., 2010 (see their Fig. 1 and also §3.3 in Fosalba et al., 2015b), where it is investigated how a systematic effect in the matter power spectrum translates into a suppression of the shear correlation function at small scales. In this work, COLA recovers 50% of the power in the matter power spectrum at $k \sim 5\,h\,\mathrm{Mpc}^{-1}$,

[2] http://www.cosmostat.org/software/athena/

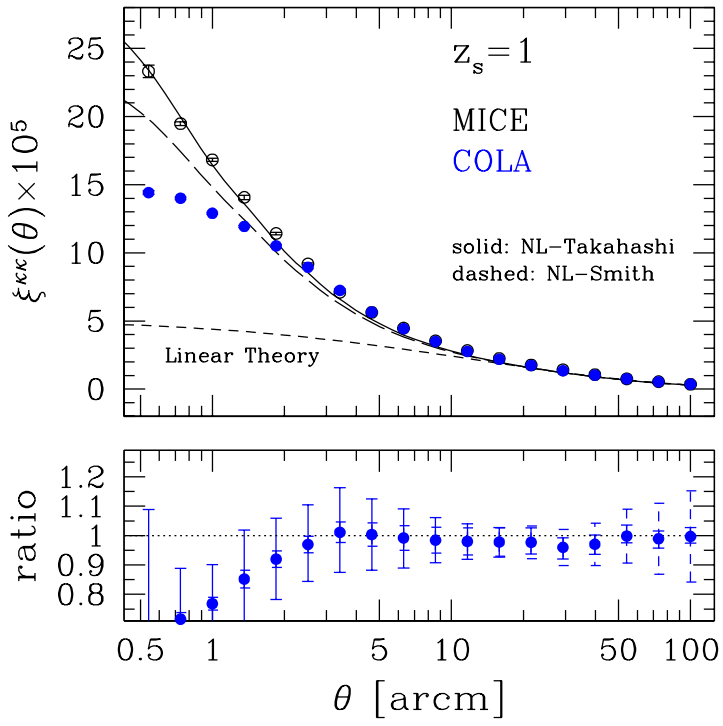

FIGURE 6.3: Convergence correlation function for the source sample. The lower panel displays the ratio of COLA with respect to MICE-GC. COLA recovers accurately the signal down to 3 arcmin. Error bars with solid lines are estimated with Jack-Knife, while dashed error bars show the contribution from sample variance alone. The non-linear predictions are computed using Halofit (Smith et al., 2003 dashed line) and revised Halofit (Takahashi et al., 2012 solid line).

which according to Heitmann et al., 2010 should translate into a scale of ~ 3 arcmin in the convergence correlation function, as we find.

Another conclusion that can be drawn is that the convergence two-point statistics is better reproduced in configuration space than in harmonic space. The scale of 3 arcmin is associated with the multipole of ~ 3600, where the convergence power spectrum largely deviates from the full N-body (see Fig. 5.6). For reference, this scale projects onto a comoving separation of $1\,h^{-1}\,\mathrm{Mpc}$ at $z = 0.5$. A plausible explanation is that COLA does not accurately trace the evolution of dark matter particles within haloes, introducing a scatter in the position of the particles with a certain characteristic scale related to the halo mean size. Correlations in configuration space are correct at those scales that are above. But when transforming to harmonic space, small and large scales are mixed and errors are propagated to a wider range of scales (see e.g., Monaco et al., 2013).

6.4 Shear correlations

When measuring cosmic shear around galaxy clusters, it is convenient to decompose the shear signal in two components, the tangential γ_t and the cross $\gamma_\times$ components. The

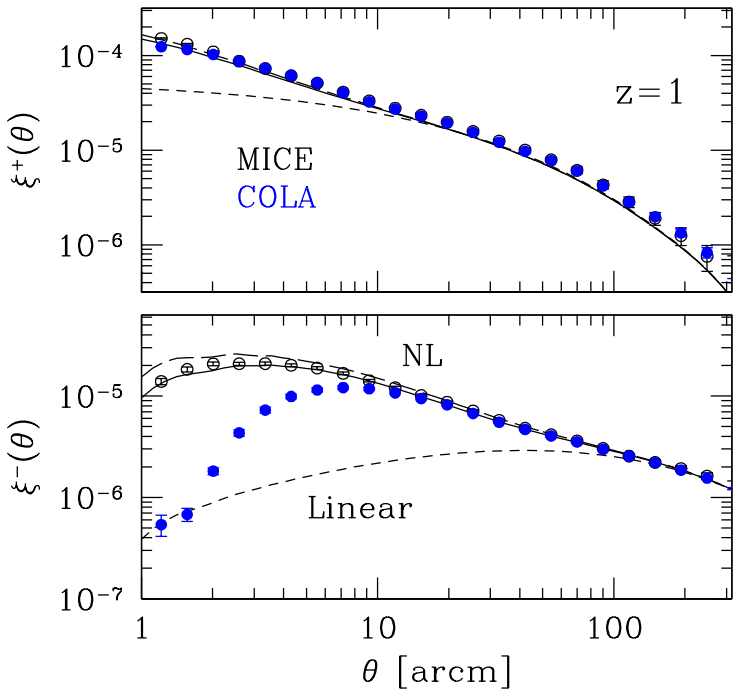

FIGURE 6.4: Shear correlation functions for the source sample. The plus component (upper panel) is accurately modelled down to 2 arcmin, while the minus (lower panel) down to ~ 15 arcmin. In both cases, it corresponds to an order magnitude beyond the non-linear scale.

polar angle ϕ of the separation vector between two points is used to project the real and imaginary part:

$$\gamma_t = -\Re(\gamma e^{-2i\phi}) \tag{6.1}$$

$$\gamma_\times = -\Im(\gamma e^{-2i\phi}), \tag{6.2}$$

where the minus sign is by convention in order to have positive values in correlation functions, since an over-density at the lens plane generates a tangential alignment of the background sources. The mixed correlation function $< \gamma_t\gamma_\times >$ vanishes due to the parity symmetry of these components. The non-zero shear correlation functions are usually expressed in the two following combinations of the tangential and cross components (see e.g. Bartelmann et al., 2001)

$$\xi_\pm(\theta) =< \gamma_t\gamma_t > (\theta) \pm < \gamma_\times\gamma_\times > (\theta). \tag{6.3}$$

Fig. 6.4 shows the shear correlation functions for the sources. The signal is well reproduced down to 2 and ~ 15 arcmin for the plus and minus components respectively.

Such different values come from the fact that the minus component is much more non-linear, but in both cases COLA is able to accurately resolve scales an order of magnitude smaller than the transition scale to the non-linear regime (i.e., where it deviates from the linear prediction).

6.5 Tangential shear

Finally, the last observable studied is the correlator of the tangential shear $< \gamma_t \gamma_t >$ between the source and lens samples. It is just the configuration space counterpart of the cross-power spectrum $C_l^{\kappa\delta_h}$ between the convergence of sources and the density of lenses (see Fig. 6.2). In the flat-sky limit, which is accurate enough for separations angles below few degrees[3], it is given by

$$\gamma_t(\theta) = \frac{1}{2\pi} \int J_2(l\theta) C_l^{\kappa\delta_h} l \, dl, \tag{6.4}$$

where J_2 is a Bessel function of the first kind. Fig. 6.5 shows that the measurements of the tangential shear correlation function in COLA are in close agreement with MICE-GC for scales larger than 5 arcmin. The theoretical curves have been multiplied by a bias factor of 1.5 in order to match the amplitude of the matter and halo clustering. We observe disagreement with the non-linear prediction for separations below 5 arcmin, that can be due to non-linear bias or the accuracy in which COLA captures the non-linear evolution within the halo structure. That scale is associated with a multipole of $l \simeq \pi/\theta = 2000$, and indeed it is consistent with the deviations found in harmonic space (see Fig. 6.2).

[3]The exact expression is given by (de Putter et al., 2010).

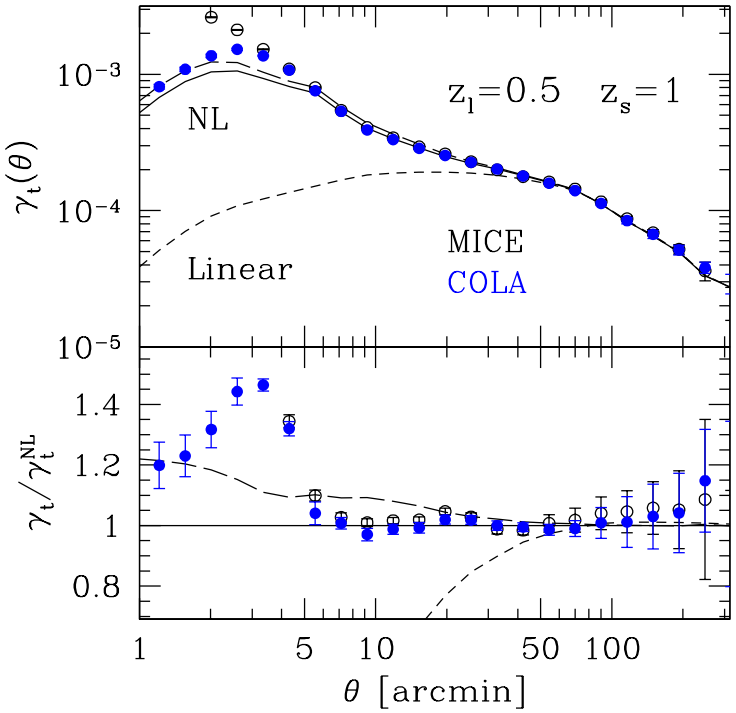

FIGURE 6.5: Tangential shear for the source and lens samples. The lower panel displays the ratio with respect to the non-linear prediction from Halofit (Smith et al., 2003, solid line), and the dashed line corresponds to the revised Halofit (Takahashi et al., 2012). Both COLA and MICE-GC are consistent with theoretical expectations down to 5 arcmin. At smaller scales, deviations arise due to non-linearities and exclusion effects.

7

ONGOING PROJECTS AND FUTURE WORK

This thesis presents the accuracy of the ICE-COLA method in modelling observables in comoving (Chapter 3) and light cone catalogues (Chapters 5 and 6). The work presented focuses on the mean value of observables. The final purpose of fast methods is to produce massive ensembles of mocks, either to estimate covariance matrices or to test analysis pipelines, and this is precisely the aim of some ongoing work. This chapter describes two ongoing projects for estimating covariance matrices, one from comoving data and the other in light cone catalogues. Lastly, there are given future directions which can be taken, many of them covering possible new features that can be implemented in ICE-COLA.

7.1 Covariance estimates comparison project

The Euclid mission[1], with a planned launch date on December 2020, as well as other future surveys that are scheduled for the next decade such as WFIRST[2] or LSST[3] will pose the most stringent requirements on ensembles of cosmological simulations, needed both for estimating covariance matrices and to develop data reduction and analysis pipelines. To be representative of the volume surveyed by the mission, they should cover large

[1] http://www.euclid-ec.org
[2] http://wfirst.gsfc.nasa.gov/
[3] http://www.lsst.org

volumes (many cubic Giga-parsecs) while retaining sufficient resolving power at small-scales. If the requirements demand a 1% accuracy on the covariance matrices, this translates into $\sim \mathcal{O}(10^4)$ realizations using just brute-force (Taylor et al., 2013; Blot et al., 2016). This big challenge seems unattainable by current approximate methods, unless they are combined with smart techniques that allow to reduce the number of realizations, such as a theoretical modelling (Grieb et al., 2016; O'Connell et al., 2015; Pearson et al., 2016), the tapering method (Paz et al., 2015), the shrinkage technique (Pope et al., 2008), data compression (Tegmark et al., 1997) or re-sampling techniques (Escoffier et al., 2016).

It is therefore essential to answer the question of how many mock catalogues are necessary for the Euclid project and other future missions to define the future roadmap. This is the motivation of a project for comparing covariance estimates produced by different approximate methods. One thousand realizations have been produced with `COLA` and `PINOCCHIO` that are being compared with a set of 100 full N-body simulations (for more details go to § 2.4.3). Other methods (e.g., HALOGEN, PATCHY, EZmocks) are expected to be incorporated in the near future. Preliminary results for the variance of the halo auto-power spectrum in redshift space are encouraging and indicate that both methods reproduce accurately the signal of the covariance on scales $k < 0.5\,h\,\mathrm{Mpc}^{-1}$ (see Fig. 7.1).

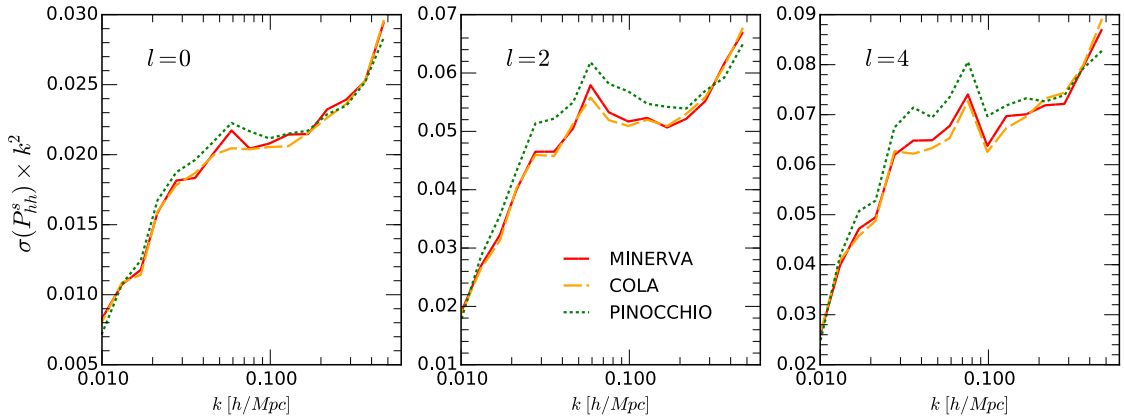

FIGURE 7.1: Preliminary results of the variance of the halo auto-power spectrum in redshift space for the monopole ($l = 0$), quadrupole ($l = 2$) and hexadecapole ($l = 4$) from left to right. All methods use the same initial conditions and indeed, the results show that the noise around the variance is well reproduced by the approximate methods. The sample is for haloes with $M > 10^{13}\,h^{-1}\,\mathrm{M}_\odot$ at redshift 1.

7.2 Covariances for weak lensing observables

Modelling covariance matrices of weak lensing observables is one of the main unresolved problems in computational cosmology. This is due to the complexity of modelling non-linear scales accurately, what typically demands large computational resources even for a single simulation. This problem becomes prohibitive when a massive number of realizations are needed. The `ICE-COLA` code presented in thesis represents a step forward

in such challenge. The next step is to analyse an ensemble of 68 realizations that we have developed so far (which hopefully will be enlarged in the future) to have accurate estimates of covariances for this particular probe. Furthermore, the simulations implement halo catalogues as well and this enables a joint analysis of galaxy clustering and weak lensing consistently, which has never done before with approximate methods. This in turn is a key tool for many other future projects, that are outlined below.

This project shall optimize the binning in redshift for weak lensing observables in ICE-COLA, that currently matches the values used for MICE-GC. Another aspect being studied is the impact of the resolution of the Healpix maps on small scales, as well as the impact of the interpolation to assign weak lensing properties at halo positions.

7.3 Future work

Given that the basic tools needed have been already developed and validated, on-going projects are expected to be completed in a short time scale. Besides, below we describe new ideas or projects that still need a more concrete definition or work before achieving substantial progress.

Mid term

Partial maps

The light cone routine stores all the pixels of each Healpix maps in the current version of ICE-COLA. However, a single process in general only modifies values of a sub-sample of pixels that are in a compact region of the sky. This fact can be used to generate partial Healpix maps that cover the need area, which translates into an efficient usage of memory and probably a slightly faster code (due to the reduction of cache misses).

Impact of weak lensing covariance on cosmological inference

Regarding weak lensing covariances, the final interest of generating those is modelling the impact of their accuracy on cosmological parameters estimation. Non-gaussianities and non-linear mode coupling are hard to model and have a strong impact on parameter inference, as has been shown in Kiessling et al., 2011b. With the complete set of light cone simulations planned this issue can be addressed. Note that the simulations are all-sky: in a study modelling a survey with 1000 deg^2 each simulation would contain 40 of such footprints, effectively multiplying the number of survey realizations.

COLA HOD galaxy mocks

Another interesting project is populating the halo catalogues of the existing simulations with galaxies using a hybrid Halo Occupation Distribution and Halo Abundance Matching technique, as has been done for MICE-GC (Carretero et al., 2015; Crocce et al., 2015). This is the last necessary step to model galaxy clustering in the current pipeline and realistic observational effects could be implemented, such as photometric-redshifts errors,

masking effects, survey varying conditions... (Leistedt et al., 2015) In turn, the galaxy mocks thus produced could be used by the collaborations of surveys such as DES or Euclid, distributing the catalogues to the community through a dedicated web portal, CosmoHUB[4].

Long term

Code optimization

There are many directions in which `ICE-COLA` can be further developed. For instance, a first idea would be implementing the spatial-COLA idea presented in Tassev et al., 2015 to simulate a sub-volume embedded within a larger environment. This could enable evolving independently sub-volumes of a very large simulation that fit in the computational resources available. In that way, larger and higher-resolution simulations could be produced. Alternatively, a different approach would be implementing a two-dimensional volume decomposition as is the case of `FastPM`, although it would involve a deep modification of the code. Other ideas is implementing the `COLA` method in a more sophisticated N-body solver such as an adaptive mesh refinement algorithm. In that case, a study would be necessary to gauge that the LPT contribution is not becoming superfluous.

Modelling of weak lensing systematics

Aside from modifications to the code, other future lines could be implementing intrinsic alignments into the analysis pipeline (Joachimi et al., 2015; Kiessling et al., 2015; Kirk et al., 2015). This effect is known to introduce systematic errors in cosmological analysis of weak lensing observations and a better modelling is needed in order to understand them and mitigate their effects. Using prescriptions of the phenomena, galaxy mock catalogues could include a value for the orientations consistent with observations (see e.g., Singh et al., 2016).

[4] http://cosmohub.pic.es/

8

CONCLUSIONS

Next generation galaxy surveys will produce a wealth of data that will allow to test the nature of the dark sector and test the cosmological model, narrowing down error-bars within the percent level. However, these achievements are only possible if systematic errors are controlled and reduced accordingly to that precision. Mock catalogues are essential in that aspect, but producing many of them using N-body simulations becomes very inefficient due to their huge computational cost. Instead, approximate methods provide the adequate tools to tackle the challenge, offering an optimal balance between accuracy and speed.

This thesis provides a series of developments in a competitive approximate method, `COLA`, that can arguably represent a step forward in the optimal exploitation of future galaxy surveys. We used a semi-N-body method and determined for the first time optimal internal code parameters that provide the highest accuracy in reproducing observables while maintaining low computational requirements. We discuss its advantages and limitations. Thanks to that optimization of parameters it is possible to resolve the necessary scales of the matter field to model weak gravitational lensing. Indeed, we implemented and validated the production of light cone simulations in a new version of the code, named `ICE-COLA`, that brings the method closer to the analysis of real data. In particular, this allowed us to accurately model weak gravitational lensing observables.

Optimization of a quasi N-body method for clustering

In order to optimize the code internal parameters, in Chapter 3 we developed a suite of 2048^3 particles `COLA` simulations that sampled the internal code parameter space and we compared observables to a reference state-of-the-art N-body simulation, the MICE-GC.

We determined the optimal configuration using as observables the two-point clustering of matter and the abundance of haloes across a wide range in halo mass and cosmic time (see § 3.4). We showed that too few time steps (of order 10) produce largely underestimated halo masses, but that a good temporal resolution reproduces the mass function within 5% and yields a matter power spectrum with one percent accuracy up to highly non-linear scales, $\sim 1\,h\,\mathrm{Mpc}^{-1}$. Therefore, for an accurate prediction of matter clustering and halo abundance across redshift, one needs to increase the default 10 time-steps by a factor of a few. Above 40, however, further gains might be limited by the force resolution. We show that, at least up to the fiducial 40 time steps, the COLA method is still preferred over a PM-only simulation (see Appendix B). We also explored the time sampling distribution and found that a linear distribution along the scale factor gives the best performance regardless of the number of time steps and the initial redshift. In addition, transients effects from the initial redshift were found to be negligible once $z_i \gtrsim 19$ within the COLA accuracy. Starting earlier had no clear benefits. Finally, We explicitly showed that a good force resolution is indispensable to mitigate an intrinsic systematic effect in COLA that generates an incompleteness in the abundance of low-mass haloes.

For the optimal set-up (see § 3.4.5) we then studied the halo clustering in real and redshift space (see § 3.5). Without applying any re-calibration against an external N-body, we found the halo clustering to be accurate within 2%. Incorporating the distortions in the radial direction that the peculiar velocities add in redshift space, we found that the signal of the monopole and the quadrupole was reproduced within 4%. These conclusions hold for wide ranges of redshifts and halo masses. Note that other approximate methods depend to a certain degree on full N-body simulations for calibrations. We further improved the accuracy of COLA by investigating two particular recalibration schemes: matching the abundance or the clustering of haloes. In this manner, we were able to achieve percent level agreement in halo bias at all mass bins and redshifts and a better agreement in redshift space down to mildly non-linear scales.

In summary, we have shown how an optimal choice of COLA code parameters, plus a minimal halo mass recalibration, can yield clustering results to per cent agreement with respect to full N-body for all mass bins, scales and redshifts of interest for the new generation of galaxy surveys. These results have been published in the following paper:

Towards fast and accurate synthetic galaxy catalogues optimizing a quasi N -body method, **Izard A.**, Crocce M., and Fosalba P., 2016, MNRAS, 459, 2327–2341, arXiv: 1509.04685.

Light cone simulations

Real observations come in the format of a light cone, in which distant objects are seen as they were in the past[1]. Hence, simulations that produce catalogues in this format enable a more realistic modelling of surveys. We implemented light cone simulations in ICE-COLA, producing catalogues on-the-fly (see Chapter 4). It is possible to generate

[1] At a look-back time corresponding to the time it takes light to travel from the source to the observer.

directly two-dimensional projected matter density maps in concentric spherical shells around the observer. Moreover, the Friend-of-Friends algorithm has been adapted to run on-the-fly in the light cone, providing halo catalogues in the light cone. In both cases, the catalogues are generated as the simulation evolves, skipping the massive production and post-processing of the particle three-dimensional light cone output. Thanks to that, compression factors of two orders of magnitude in the data volume is possible. The light cone is constructed in many box replicas, allowing for instance all-sky catalogues, and has been carefully designed to consume the minimum computational resources. The data products obtained with the new code are validated (see §5.1), showing that the outputs are reliable.

Modelling of weak gravitational lensing

We described the pipeline that we developed to model weak gravitational lensing observables from light cone catalogues (see §5.2). This probe holds the promise of providing some of the most stringent cosmological constraints in the near future if systematic effects are kept under control (see the introduction). But its theoretical modelling is challenging because it probes highly non-linear scales at the same time that samples very large volumes, which is an unsolved problem in numerical cosmology. The methodology of this thesis aims at meeting this important challenge in computational cosmology. In particular, we produced all-sky convergence and shear maps with adequate resolution in redshift and angular scale on the sky to model weak lensing observables from the largest (linear) to small (non-linear) scales. The pipeline uses efficiently the compressed data catalogues and can be run on a personal laptop. The maps were compared against MICE-GC (see §5.2), finding that the angular auto-power spectrum is accurately reproduced up to multipoles of 2000 within 10% accuracy. More interestingly, we find that the convergence-matter cross-power spectrum resolves scales of $l \sim 4000$ with the same accuracy, since in this case non-linear contributions at low redshifts (that are seen at larger angles) have a smaller contribution to the signal.

Finally, we show how to combine halo catalogues and weak lensing maps to study in a consistent way galaxy clustering and weak lensing from a single simulation (see Chapter 6). Comparing also to MICE-GC, we found the promising results that we can reproduce shear two-point statistics to angular separations of few arc-minutes. The particular scale depends on each observable, but in general we find that the method predicts correctly the signal on scales an order of magnitude beyond the linear regime. Note that other approximate methods are not able to directly predict such regimes, since they explicitly avoid resolving small-scale structure (within haloes).

These results constitute a further and thorough validation of the light cone simulations presented before.

All the findings related to the light cone simulations and the modelling of weak lensing observable are included in a forthcoming publication.

BIBLIOGRAPHY

Aarseth S. J., Turner E. L., and Gott J. R., 1979, *N-body simulations of galaxy clustering. I - Initial conditions and galaxy collapse times, Astrophys. J.*, **228,** 664–683.

Albrecht A., Bernstein G., Cahn R., Freedman W. L., Hewitt J., Hu W., Huth J., Kamionkowski M., Kolb E. W., and Knox L., 2006, *Report of the Dark Energy Task Force,* ArXiv Astrophysics e-prints, eprint: astro-ph/0609591.

Alimi J.-M., Bouillot V., Rasera Y., Reverdy V., Corasaniti P.-S., Balmes I., Requena S., Delaruelle X., and Richet J.-N., 2012, *DEUS Full Observable {\Lambda}CDM Universe Simulation: the numerical challenge,* ArXiv e-prints, arXiv: 1206.2838 [astro-ph.CO].

Allen S. W., Evrard A. E., and Mantz A. B., 2011, *Cosmological Parameters from Observations of Galaxy Clusters, Annu. Rev. Astron. Astrophys.*, **49,** 409–470. arXiv: 1103.4829 [astro-ph.CO].

Anderson L., Aubourg E., Bailey S., Bizyaev D., Blanton M., Bolton A. S., Brinkmann J., Brownstein J. R., Burden A., and Cuesta A. J., 2012, *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Release 9 spectroscopic galaxy sample, MNRAS*, **427,** 3435–3467. arXiv: 1203.6594.

Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., and Frenk C. S., 2012, *Scaling relations for galaxy clusters in the Millennium-XXL simulation, MNRAS*, **426,** 2046–2062. arXiv: 1203.3216 [astro-ph.CO].

Appel A. W., 1985, *An Efficient Program for Many-Body Simulation,* SIAM Journal on Scientific and Statistical Computing, vol. 6, no. 1, January 1985, p. 85-103., **6,** 85–103.

Avila S., Murray S. G., Knebe A., Power C., Robotham A. S. G., and Garcia-Bellido J., 2015, *HALOGEN: a tool for fast generation of mock halo catalogues, MNRAS*, **450,** 1856–1867. arXiv: 1412.5228.

Bacon D. J., Refregier A. R., and Ellis R. S., 2000, *Detection of weak gravitational lensing by large-scale structure, MNRAS*, **318,** 625–640. eprint: astro-ph/0003008.

Baldauf T., Seljak U., Smith R. E., Hamaus N., and Desjacques V., 2013, *Halo stochasticity from exclusion and nonlinear clustering, Phys. Rev. D*, **88,**(8), 083507, 083507. arXiv: 1305.2917 [astro-ph.CO].

Bardeen J. M., Bond J. R., Kaiser N., and Szalay A. S., 1986, *The statistics of peaks of Gaussian random fields, Astrophys. J.*, **304,** 15–61.

Bartelmann M., and Schneider P., 2001, *Weak gravitational lensing, Phys. Rep.*, **340,** 291–472. eprint: astro-ph/9912508.

Bassett B., and Hlozek R., 2010, *Baryon acoustic oscillations, Dark Energy: Observational and Theoretical Approaches*. Ed. by P. Ruiz-Lapuente, p. 246.

Beaton R. L., Freedman W. L., Madore B. F., Bono G., Carlson E. K., Clementini G., Durbin M. J., Garofalo A., Hatt D., and Jang I. S., 2016, *The Carnegie-Chicago Hubble Program. I. A New Approach to the Distance Ladder Using Only Distance Indicators of Population II,* ArXiv e-prints, arXiv: 1604.01788.

Bernardeau F., Colombi S., Gaztañaga E., and Scoccimarro R., 2002, *Large-scale structure of the Universe and cosmological perturbation theory, Phys. Rep.*, **367,** 1–248. eprint: astro-ph/0112551.

Bertschinger E., 1998, *Simulations of Structure Formation in the Universe, Annu. Rev. Astron. Astrophys.*, **36,** 599–654.

Beutler F., Blake C., Koda J., Marin F., Seo H.-J., Cuesta A. J., and Schneider D. P., 2015, *The BOSS-WiggleZ overlap region I: Baryon Acoustic Oscillations,* ArXiv e-prints, arXiv: 1506.03900.

Binney J., and Tremaine S., 1987, *Galactic dynamics*.

Blaizot J., Wadadekar Y., Guiderdoni B., Colombi S. T., Bertin E., Bouchet F. R., Devriendt J. E. G., and Hatton S., 2005, *MoMaF: the Mock Map Facility, MNRAS*, **360,** 159–175. eprint: astro-ph/0309305.

Blandford R., and Narayan R., 1986, *Fermat's principle, caustics, and the classification of gravitational lens images, Astrophys. J.*, **310,** 568–582.

Blot L., Corasaniti P. S., Alimi J.-M., Reverdy V., and Rasera Y., 2015, *Matter power spectrum covariance matrix from the DEUS-PUR ΛCDM simulations: mass resolution and non-Gaussian errors, MNRAS*, **446,** 1756–1764. arXiv: 1406.2713.

Blot L., Corasaniti P. S., Amendola L., and Kitching T. D., 2016, *Non-linear matter power spectrum covariance matrix errors and cosmological parameter uncertainties, MNRAS*, **458,** 4462–4470. arXiv: 1512.05383.

Bouchet F. R., Colombi S., Hivon E., and Juszkiewicz R., 1995, *Perturbative Lagrangian approach to gravitational instability., A&A*, **296,** 575. eprint: astro-ph/9406013.

Carretero J., Castander F. J., Gaztañaga E., Crocce M., and Fosalba P., 2015, *An algorithm to build mock galaxy catalogues using MICE simulations, MNRAS*, **447,** 646–670. arXiv: 1411.3286.

Chuang C.-H., Kitaura F.-S., Prada F., Zhao C., and Yepes G., 2015, *EZmocks: extending the Zel'dovich approximation to generate mock galaxy catalogues with accurate clustering statistics, MNRAS*, **446,** 2621–2628. arXiv: 1409.1124.

Cole S., and Kaiser N., 1989, *Biased clustering in the cold dark matter cosmogony, MNRAS*, **237,** 1127–1146.

Cole S., and Lacey C., 1996, *The structure of dark matter haloes in hierarchical clustering models, MNRAS*, **281,** 716. eprint: astro-ph/9510147.

Coles P., and Jones B., 1991, *A lognormal model for the cosmological mass distribution, MNRAS*, **248,** 1–13.

Cooray A., and Sheth R., 2002, *Halo models of large scale structure, Phys. Rep.*, **372,** 1–129. eprint: astro-ph/0206508.

Couchman H. M. P., 1991, *Mesh-refined P3M - A fast adaptive N-body algorithm, Astrophys. J. Lett.*, **368,** L23–L26.

Crocce M., Castander F. J., Gaztañaga E., Fosalba P., and Carretero J., 2015, *The MICE Grand Challenge lightcone simulation - II. Halo and galaxy catalogues, MNRAS*, **453,** 1513–1530. arXiv: 1312.2013.

Crocce M., Fosalba P., Castander F. J., and Gaztañaga E., 2010, *Simulating the Universe with MICE: the abundance of massive clusters, MNRAS*, **403,** 1353–1367. arXiv: 0907.0019 [astro-ph.CO].

Crocce M., Pueblas S., and Scoccimarro R., 2006, *Transients from initial conditions in cosmological simulations, MNRAS*, **373,** 369–381. eprint: astro-ph/0606505.

Cuesta A. J., Vargas-Magaña M., Beutler F., Bolton A. S., Brownstein J. R., Eisenstein D. J., Gil-Marin H., Ho S., McBride C. K., and Maraston C., 2016, *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the correlation function of LOWZ and CMASS galaxies in Data Release 12, MNRAS*, **457,** 1770–1785. arXiv: 1509.06371.

Das S., and Bode P., 2008, *A Large Sky Simulation of the Gravitational Lensing of the Cosmic Microwave Background, Astrophys. J.*, **682,** 1-13, 1–13. arXiv: 0711.3793.

Davis M., Efstathiou G., Frenk C. S., and White S. D. M., 1985, *The evolution of large-scale structure in a universe dominated by cold dark matter, Astrophys. J.*, **292,** 371–394.

Dawson K. S., Kneib J.-P., Percival W. J., Alam S., Albareti F. D., Anderson S. F., Armengaud E., Aubourg É., Bailey S., and Bautista J. E., 2016, *The SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data, Astr. J.*, **151,** 44, 44. arXiv: 1508.04473.

de Jong J. T. A., Kuijken K., Applegate D., Begeman K., Belikov A., Blake C., Bout J., Boxhoorn D., Buddelmeijer H., and Buddendiek A., 2013, *The Kilo-Degree Survey,* The Messenger, **154,** 44–46.

de Putter R., and Takada M., 2010, *Halo-galaxy lensing: A full sky approach, Phys. Rev. D*, **82,**(10), 103522, 103522. arXiv: 1007.4809 [astro-ph.CO].

Dehnen W., and Read J. I., 2011, *N-body simulations of gravitational dynamics,* European Physical Journal Plus, **126,** 55, 55. arXiv: 1105.1082 [astro-ph.IM].

Dodelson S., 2003, *Modern cosmology*.

Dunkley J., Komatsu E., Nolta M. R., Spergel D. N., Larson D., Hinshaw G., Page L., Bennett C. L., Gold B., and Jarosik N., 2009, *Five-Year Wilkinson Microwave Anisotropy Probe Observations: Likelihoods and Parameters from the WMAP Data, Astrophys. J. Suppl.*, **180,** 306–329. arXiv: 0803.0586.

Efstathiou G., 2014, H_0 *revisited, MNRAS*, **440,** 1138–1152. arXiv: 1311.3461.

Efstathiou G., Davis M., White S. D. M., and Frenk C. S., 1985, *Numerical techniques for large cosmological N-body simulations, Astrophys. J. Suppl.*, **57,** 241–260.

Efstathiou G., and Eastwood J. W., 1981, *On the clustering of particles in an expanding universe, MNRAS*, **194,** 503–525.

Eisenstein D. J., Zehavi I., Hogg D. W., Scoccimarro R., Blanton M. R., Nichol R. C., Scranton R., Seo H.-J., Tegmark M., and Zheng Z., 2005, *Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies, Astrophys. J.*, **633,** 560–574. eprint: astro-ph/0501171.

Escoffier S., Cousinou M.-C., Tilquin A., Pisani A., Aguichine A., de la Torre S., Ealet A., Gillard W., and Jullo E., 2016, *Jackknife resampling technique on mocks: an alternative method for covariance matrix estimation,* ArXiv e-prints, arXiv: 1606.00233.

Feng Y., Chu M.-Y., and Seljak U., 2016, *FastPM: a new scheme for fast simulations of dark matter and halos,* ArXiv e-prints, arXiv: 1603.00476.

Fosalba P., Crocce M., Gaztañaga E., and Castander F. J., 2015a, *The MICE grand challenge lightcone simulation - I. Dark matter clustering, MNRAS*, **448,** 2987–3000.

Fosalba P., Gaztañaga E., Castander F. J., and Crocce M., 2015b, *The MICE Grand Challenge light-cone simulation - III. Galaxy lensing mocks from all-sky lensing maps, MNRAS*, **447,** 1319–1332. arXiv: 1312.2947 [astro-ph.CO].

Fosalba P., Gaztañaga E., Castander F. J., and Manera M., 2008, *The onion universe: all sky lightcone simulations in spherical shells, MNRAS*, **391,** 435–446. arXiv: 0711.1540.

Freedman W. L., Madore B. F., Scowcroft V., Burns C., Monson A., Persson S. E., Seibert M., and Rigby J., 2012, *Carnegie Hubble Program: A Mid-infrared Calibration of the Hubble Constant, Astrophys. J.*, **758,** 24, 24. arXiv: 1208.3281.

Fry J. N., and Gaztanaga E., 1993, *Biasing and hierarchical statistics in large-scale structure, Astrophys. J.*, **413,** 447–452. eprint: astro-ph/9302009.

Gil-Marin H., Percival W. J., Cuesta A. J., Brownstein J. R., Chuang C.-H., Ho S., Shu Kitaura F.-., Maraston C., Prada F., and Rodriguez-Torres S., 2016, *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: BAO measurement from the LOS-dependent power spectrum of DR12 BOSS galaxies, MNRAS*, arXiv: 1509.06373.

Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., and Bartelmann M., 2005, *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere, Astrophys. J.*, **622,** 759–771. eprint: astro-ph/0409513.

Grieb J. N., Sánchez A. G., Salazar-Albornoz S., and Dalla Vecchia C., 2016, *Gaussian covariance matrices for anisotropic galaxy clustering measurements, MNRAS*, **457,** 1577–1592. arXiv: 1509.04293.

Habib S., Morozov V., Finkel H., Pope A., Heitmann K., Kumaran K., Peterka T., Insley J., Daniel D., and Fasel P., 2012, *The Universe at Extreme Scale: Multi-Petaflop Sky Simulation on the BG/Q,* ArXiv e-prints, arXiv: 1211.4864 [cs.DC].

Hahn O., and Angulo R. E., 2016, *An adaptively refined phase-space element method for cosmological simulations and collisionless dynamics, MNRAS*, **455,** 1115–1133. arXiv: 1501.01959.

Hamilton J.-C., 2013, *What have we learned from observational cosmology ?,* ArXiv e-prints, arXiv: 1304.4446 [astro-ph.CO].

Heitmann K., Lawrence E., Kwan J., Habib S., and Higdon D., 2014, *The Coyote Universe Extended: Precision Emulation of the Matter Power Spectrum, Astrophys. J.*, **780,** 111, 111. arXiv: 1304.7849.

Heitmann K., White M., Wagner C., Habib S., and Higdon D., 2010, *The Coyote Universe. I. Precision Determination of the Nonlinear Matter Power Spectrum, Astrophys. J.*, **715,** 104–121. arXiv: 0812.1052.

Hockney R. W., and Eastwood J. W., 1988, *Computer simulation using particles*.

Hoekstra H., Mellier Y., van Waerbeke L., Semboloni E., Fu L., Hudson M. J., Parker L. C., Tereno I., and Benabed K., 2006, *First Cosmic Shear Results from the Canada-France-Hawaii Telescope Wide Synoptic Legacy Survey, Astrophys. J.*, **647,** 116–127. eprint: astro-ph/0511089.

Hoffmann K., Bel J., and Gaztañaga E., 2015, *Comparing halo bias from abundance and clustering, MNRAS*, **450,** 1674–1692. arXiv: 1503.00313.

Hogg D. W., 1999, *Distance measures in cosmology,* ArXiv Astrophysics e-prints, eprint: astro-ph/9905116.

Howlett C., Manera M., and Percival W. J., 2015a, *L-PICOLA: A parallel code for fast dark matter simulation,* Astronomy and Computing, **12,** 109–126. arXiv: 1506.03737.

Howlett C., Ross A. J., Samushia L., Percival W. J., and Manera M., 2015b, *The clustering of the SDSS main galaxy sample - II. Mock galaxy catalogues and a measurement of the growth of structure from redshift space distortions at z = 0.15, MNRAS*, **449,** 848–866. arXiv: 1409.3238.

Hu W., 2000, *Weak lensing of the CMB: A harmonic approach, Phys. Rev. D*, **62,**(4), 043007, 043007. eprint: astro-ph/0001303.

Huterer D., 2010, *Weak lensing, dark matter and dark energy,* General Relativity and Gravitation, **42,** 2177–2195. arXiv: 1001.1758 [astro-ph.CO].

Izard A., Crocce M., and Fosalba P., 2016, *ICE-COLA: towards fast and accurate synthetic galaxy catalogues optimizing a quasi-N-body method, MNRAS*, **459,** 2327–2341. arXiv: 1509.04685.

Jain B., Seljak U., and White S., 2000, *Ray-tracing Simulations of Weak Lensing by Large-Scale Structure, Astrophys. J.*, **530,** 547–577. eprint: astro-ph/9901191.

Jing Y. P., 2005, *Correcting for the Alias Effect When Measuring the Power Spectrum Using a Fast Fourier Transform, Astrophys. J.*, **620,** 559–563. eprint: astro-ph/0409240.

Joachimi B., Cacciato M., Kitching T. D., Leonard A., Mandelbaum R., Schäfer B. M., Sifon C., Hoekstra H., Kiessling A., and Kirk D., 2015, *Galaxy Alignments: An Overview, Space Sci. Rev.*, **193,** 1–65. arXiv: 1504.05456.

Joyce A., Lombriser L., and Schmidt F., 2016, *Dark Energy vs. Modified Gravity,* ArXiv e-prints, arXiv: 1601.06133.

Kaiser N., 1987, *Clustering in real space and in redshift space, MNRAS*, **227,** 1–21.

— 1992, *Weak gravitational lensing of distant galaxies, Astrophys. J.*, **388,** 272–286.

— 1998, *Weak Lensing and Cosmology, Astrophys. J.*, **498,** 26–42. eprint: astro-ph/9610120.

Kaiser N., and Squires G., 1993, *Mapping the dark matter with weak gravitational lensing, Astrophys. J.*, **404,** 441–450.

Kaiser N., Wilson G., and Luppino G. A., 2000, *Large-Scale Cosmic Shear Measurements*, ArXiv Astrophysics e-prints, eprint: astro-ph/0003338.

Kazin E. A., Koda J., Blake C., Padmanabhan N., and *et al.* 2014, *The WiggleZ Dark Energy Survey: improved distance measurements to z = 1 with reconstruction of the baryonic acoustic feature*, *MNRAS*, **441,** 3524–3542. arXiv: 1401.0358.

Kiessling A., Cacciato M., Joachimi B., Kirk D., Kitching T. D., Leonard A., Mandelbaum R., Schafer B. M., Sifon C., and Brown M. L., 2015, *Galaxy Alignments: Theory, Modelling & Simulations*, *Space Sci. Rev.*, **193,** 67–136. arXiv: 1504.05546.

Kiessling A., Heavens A. F., Taylor A. N., and Joachimi B., 2011a, *SUNGLASS: a new weak-lensing simulation pipeline*, *MNRAS*, **414,** 2235–2245. arXiv: 1011.1476.

Kiessling A., Taylor A. N., and Heavens A. F., 2011b, *Simulating the effect of non-linear mode coupling in cosmological parameter estimation*, *MNRAS*, **416,** 1045–1055. arXiv: 1103.3245.

Kilbinger M., 2015, *Cosmology with cosmic shear observations: a review*, Reports on Progress in Physics, **78,**(8), 086901, 086901. arXiv: 1411.0115.

Kilbinger M., Fu L., Heymans C., Simpson F., Benjamin J., Erben T., Harnois-Déraps J., Hoekstra H., Hildebrandt H., and Kitching T. D., 2013, *CFHTLenS: combined probe cosmological model comparison using 2D weak gravitational lensing*, *MNRAS*, **430,** 2200–2220. arXiv: 1212.3338.

Kim J., Park C., Rossi G., Lee S. M., and Gott J. R., 2011, *The New Horizon Run Cosmological N-Body Simulations*, Journal of Korean Astronomical Society, **44,** 217–234. arXiv: 1112.1754 [astro-ph.CO].

Kirk D., Brown M. L., Hoekstra H., Joachimi B., Kitching T. D., Mandelbaum R., Sifon C., Cacciato M., Choi A., and Kiessling A., 2015, *Galaxy Alignments: Observations and Impact on Cosmology*, *Space Sci. Rev.*, **193,** 139–211. arXiv: 1504.05465.

Kitaura F.-S., Gil-Marin H., Scoccola C. G., Chuang C.-H., Muller V., Yepes G., and Prada F., 2015, *Constraining the halo bispectrum in real and redshift space from perturbation theory and non-linear stochastic bias*, *MNRAS*, **450,** 1836–1845. arXiv: 1407.1236.

Klypin A., Prada F., Yepes G., Hess S., and Gottlober S., 2013, *Halo Abundance Matching: accuracy and conditions for numerical convergence*, ArXiv e-prints, arXiv: 1310.3740 [astro-ph.CO].

Klypin A., Yepes G., Gottlöber S., Prada F., and Heß S., 2016, *MultiDark simulations: the story of dark matter halo concentrations and density profiles*, *MNRAS*, **457,** 4340–4359. arXiv: 1411.4001.

Knebe A., Pearce F. R., Thomas P. A., Benson A., Blaizot J., Bower R., Carretero J., Castander F. J., Cattaneo A., and Cora S. A., 2015, *nIFTy cosmology: comparison of galaxy formation models*, *MNRAS*, **451,** 4029–4059. arXiv: 1505.04607.

Koda J., Blake C., Beutler F., Kazin E., and Marin F., 2015, *Fast and accurate mock catalogue generation for low-mass galaxies*, ArXiv e-prints, arXiv: 1507.05329.

Krause E., and Hirata C. M., 2010, *Weak lensing power spectra for precision cosmology. Multiple-deflection, reduced shear, and lensing bias corrections*, *A&A*, **523,** A28, A28. arXiv: 0910.3786.

Kuhlen M., Vogelsberger M., and Angulo R., 2012, *Numerical simulations of the dark universe: State of the art and the next decade,* Physics of the Dark Universe, **1,** 50–93. arXiv: 1209.5745 [astro-ph.CO].

Laureijs R., Amiaux J., Arduini S., Augueres J. .-.-., Brinchmann J., Cole R., Cropper M., Dabin C., Duvet L., and Ealet A., 2011, *Euclid Definition Study Report,* ArXiv e-prints, arXiv: 1110.3193 [astro-ph.CO].

Leclercq F., Jasche J., and Wandelt B., 2015, *Bayesian analysis of the dynamic cosmic web in the SDSS galaxy survey, J. Cosm. Astr. Phys.*, **6,** 015, 15. arXiv: 1502.02690.

Leistedt B., Peiris H. V., Elsner F., Benoit-Lévy A., Amara A., Bauer A. H., Becker M. R., Bonnett C., Bruderer C., and Busha M. T., 2015, *Mapping and simulating systematics due to spatially-varying observing conditions in DES Science Verification data,* ArXiv e-prints, arXiv: 1507.05647.

Levi M., Bebek C., Beers T., Blum R., Cahn R., Eisenstein D., Flaugher B., Honscheid K., Kron R., and Lahav O., 2013, *The DESI Experiment, a whitepaper for Snowmass 2013,* ArXiv e-prints, arXiv: 1308.0847 [astro-ph.CO].

Li B., King L. J., Zhao G.-B., and Zhao H., 2011, *An analytic ray-tracing algorithm for weak lensing, MNRAS*, **415,** 881–892. arXiv: 1012.1625.

Liddle A. R., and Lyth D. H., 2000, *Cosmological Inflation and Large-Scale Structure*, 414.

Limber D. N., 1953, *The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field., Astrophys. J.*, **117,** 134.

LSST Science Collaboration, Abell P. A., Allison J., Anderson S. F., Andrew J. R., Angel J. R. P., Armus L., Arnett D., Asztalos S. J., and Axelrod T. S., 2009, *LSST Science Book, Version 2.0,* ArXiv e-prints, arXiv: 0912.0201 [astro-ph.IM].

Manera M., Samushia L., Tojeiro R., Howlett C., Ross A. J., Percival W. J., Gil-Marin H., Brownstein J. R., Burden A., and Montesano F., 2015, *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: mock galaxy catalogues for the low-redshift sample, MNRAS*, **447,** 437–445. arXiv: 1401.4171.

Manera M., Scoccimarro R., Percival W. J., Samushia L., McBride C. K., Ross A. J., Sheth R. K., White M., Reid B. A., and Sánchez A. G., 2013, *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: a large sample of mock galaxy catalogues, MNRAS*, **428,** 1036–1054. arXiv: 1203.6609.

Manera M., Sheth R. K., and Scoccimarro R., 2010, *Large-scale bias and the inaccuracy of the peak-background split, MNRAS*, **402,** 589–602. arXiv: 0906.1314.

Merson A. I., Baugh C. M., Helly J. C., Gonzalez-Perez V., Cole S., Bielby R., Norberg P., Frenk C. S., Benson A. J., and Bower R. G., 2013, *Lightcone mock catalogues from semi-analytic models of galaxy formation - I. Construction and application to the BzK colour selection, MNRAS*, **429,** 556–578. arXiv: 1206.4049.

Miyazaki S., Hamana T., Shimasaku K., Furusawa H., Doi M., Hamabe M., Imi K., Kimura M., Komiyama Y., and Nakata F., 2002, *Searching for Dark Matter Halos in the Suprime-Cam 2 Square Degree Field, Astrophys. J. Lett.*, **580,** L97–L100. eprint: astro-ph/0210441.

Monaco P., 2016, *Approximated methods for the generation of dark matter halo catalogs in the age of precision cosmology*, ArXiv e-prints, arXiv: 1605.07752.

Monaco P., Sefusatti E., Borgani S., Crocce M., Fosalba P., Sheth R. K., and Theuns T., 2013, *An accurate tool for the fast generation of dark matter halo catalogues, MNRAS*, **433,** 2389–2402. arXiv: 1305.1505 [astro-ph.CO].

Monaco P., Theuns T., and Taffoni G., 2002, *The pinocchio algorithm: pinpointing orbit-crossing collapsed hierarchical objects in a linear density field, MNRAS*, **331,** 587–608. eprint: astro-ph/0109323.

Munari E., Monaco P., Sefusatti E., Castorina E., Mohammad F. G., Anselmi S., and Borgani S., 2016, *Improving fast generation of halo catalogs with higher-order Lagrangian perturbation theory,* ArXiv e-prints, arXiv: 1605.04788.

Munshi D., Valageas P., van Waerbeke L., and Heavens A., 2008, *Cosmology with weak lensing surveys, Phys. Rep.*, **462,** 67–121. eprint: astro-ph/0612667.

Narayan R., and Bartelmann M., 1996, *Lectures on Gravitational Lensing*, ArXiv Astrophysics e-prints, eprint: astro-ph/9606001.

O'Connell R., Eisenstein D., Vargas M., Ho S., and Padmanabhan N., 2015, *Large Covariance Matrices: Smooth Models from the 2-Point Correlation Function,* ArXiv e-prints, arXiv: 1510.01740.

Paz D. J., and Sánchez A. G., 2015, *Improving the precision matrix for precision cosmology, MNRAS*, **454,** 4326–4334. arXiv: 1508.03162.

Peacock J. A., 1999, *Cosmological Physics*, 704.

Pearson D. W., and Samushia L., 2016, *Estimating the power spectrum covariance matrix with fewer mock samples, MNRAS*, **457,** 993–999. arXiv: 1509.00064.

Peebles P. J. E., 1970, *Structure of the Coma Cluster of Galaxies, Astr. J.*, **75,** 13.

— 1980, *The large-scale structure of the universe*.

Perlmutter S., Aldering G., Goldhaber G., Knop R. A., Nugent P., Castro P. G., Deustua S., Fabbro S., Goobar A., and Groom D. E., 1999, *Measurements of Ω and Λ from 42 High-Redshift Supernovae, Astrophys. J.*, **517,** 565–586. eprint: astro-ph/9812133.

Planck Collaboration, Ade P. A. R., Aghanim N., Akrami Y., Aluri P. K., Arnaud M., Ashdown M., Aumont J., Baccigalupi C., and Banday A. J., 2015a, *Planck 2015 results. XVI. Isotropy and statistics of the CMB,* ArXiv e-prints, arXiv: 1506.07135.

Planck Collaboration, Ade P. A. R., Aghanim N., Arnaud M., Ashdown M., Aumont J., Baccigalupi C., Banday A. J., Barreiro R. B., and Bartlett J. G., 2015b, *Planck 2015 results. XIII. Cosmological parameters,* ArXiv e-prints, arXiv: 1502.01589.

Pope A. C., and Szapudi I., 2008, *Shrinkage estimation of the power spectrum covariance matrix, MNRAS*, **389,** 766–774. arXiv: 0711.2509.

Press W. H., and Schechter P., 1974, *Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation, Astrophys. J.*, **187,** 425–438.

Riess A. G., Filippenko A. V., Challis P., Clocchiatti A., Diercks A., Garnavich P. M., Gilliland R. L., Hogan C. J., Jha S., and Kirshner R. P., 1998, *Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant, Astr. J.*, **116,** 1009–1038. eprint: astro-ph/9805201.

Riess A. G., Macri L. M., Hoffmann S. L., Scolnic D., Casertano S., Filippenko A. V., Tucker B. E., Reid M. J., Jones D. O., and Silverman J. M., 2016, *A 2.4% Determination of the Local Value of the Hubble Constant*, ArXiv e-prints, arXiv: 1604.01424.

Riess A. G., Macri L., Casertano S., Lampeitl H., Ferguson H. C., Filippenko A. V., Jha S. W., Li W., and Chornock R., 2011, *A 3% Solution: Determination of the Hubble Constant with the Hubble Space Telescope and Wide Field Camera 3*, *Astrophys. J.*, **730,** 119, 119. arXiv: 1103.2976.

Ross A. J., Samushia L., Howlett C., Percival W. J., Burden A., and Manera M., 2015, *The clustering of the SDSS DR7 main Galaxy sample - I. A 4 per cent distance measure at z = 0.15*, *MNRAS*, **449,** 835–847. arXiv: 1409.3242.

Rowan-Robinson M., 1985, *The cosmological distance ladder: Distance and time in the universe*.

Sánchez A. G., Kazin E. A., Beutler F., Chuang C.-H., Cuesta A. J., Eisenstein D. J., Manera M., Montesano F., Nichol R. C., and Padmanabhan N., 2013, *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological constraints from the full shape of the clustering wedges*, *MNRAS*, **433,** 1202–1222. arXiv: 1303.4396.

Schneider A., Teyssier R., Potter D., Stadel J., Onions J., Reed D. S., Smith R. E., Springel V., and Pearce F. R., 2015, *Matter power spectrum and the challenge of percent accuracy*, ArXiv e-prints, arXiv: 1503.05920.

Scoccimarro R., 1998, *Transients from initial conditions: a perturbative analysis*, *MNRAS*, **299,** 1097–1118. eprint: astro-ph/9711187.

Scoccimarro R., and Sheth R. K., 2002, *PTHALOS: a fast method for generating mock galaxy distributions*, *MNRAS*, **329,** 629–640. eprint: astro-ph/0106120.

Scoccimarro R., Sheth R. K., Hui L., and Jain B., 2001, *How Many Galaxies Fit in a Halo? Constraints on Galaxy Formation Efficiency from Spatial Clustering*, *Astrophys. J.*, **546,** 20–34. eprint: astro-ph/0006319.

Sheth R. K., and Tormen G., 1999, *Large-scale bias and the peak background split*, *MNRAS*, **308,** 119–126. eprint: astro-ph/9901122.

Silk J., Di Cintio A., and Dvorkin I., 2013, *Galaxy formation*, ArXiv e-prints, arXiv: 1312.0107 [astro-ph.CO].

Singh S., and Mandelbaum R., 2016, *Intrinsic alignments of BOSS LOWZ galaxies - II. Impact of shape measurement methods*, *MNRAS*, **457,** 2301–2317. arXiv: 1510.06752.

Skillman S. W., Warren M. S., Turk M. J., Wechsler R. H., Holz D. E., and Sutter P. M., 2014, *Dark Sky Simulations: Early Data Release*, ArXiv e-prints, arXiv: 1407.2600.

Smith R. E., Peacock J. A., Jenkins A., White S. D. M., Frenk C. S., Pearce F. R., Thomas P. A., Efstathiou G., and Couchman H. M. P., 2003, *Stable clustering, the halo model and non-linear cosmological power spectra*, *MNRAS*, **341,** 1311–1332. eprint: astro-ph/0207664.

Spergel D., Gehrels N., Breckinridge J., Donahue M., Dressler A., Gaudi B. S., Greene T., Guyon O., Hirata C., and Kalirai J., 2013, *Wide-Field InfraRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA Final Report*, ArXiv e-prints, arXiv: 1305.5422 [astro-ph.IM].

Springel V., 2005, *The cosmological simulation code GADGET-2, MNRAS*, **364,** 1105–1134. eprint: astro-ph/0505010.

Stadel J. G., 2001, *Cosmological N-body simulations and their analysis,* PhD thesis. UNIVERSITY OF WASHINGTON.

Suzuki N., Rubin D., Lidman C., Aldering G., Amanullah R., Barbary K., Barrientos L. F., Botyanszki J., Brodwin M., and Connolly N., 2012, *The Hubble Space Telescope Cluster Supernova Survey. V. Improving the Dark-energy Constraints above z > 1 and Building an Early-type-hosted Supernova Sample, Astrophys. J.*, **746,** 85, 85. arXiv: 1105.3470 [astro-ph.CO].

Takahashi R., Sato M., Nishimichi T., Taruya A., and Oguri M., 2012, *Revising the Halofit Model for the Nonlinear Matter Power Spectrum, Astrophys. J.*, **761,** 152, 152. arXiv: 1208.2701.

Takahashi R., Yoshida N., Takada M., Matsubara T., Sugiyama N., Kayo I., Nishizawa A. J., Nishimichi T., Saito S., and Taruya A., 2009, *Simulations of Baryon Acoustic Oscillations. II. Covariance Matrix of the Matter Power Spectrum, Astrophys. J.*, **700,** 479–490. arXiv: 0902.0371 [astro-ph.CO].

Tassev S., Eisenstein D. J., Wandelt B. D., and Zaldarriaga M., 2015, *sCOLA: The N-body COLA Method Extended to the Spatial Domain,* ArXiv e-prints, arXiv: 1502.07751.

Tassev S., Zaldarriaga M., and Eisenstein D. J., 2013, *Solving large scale structure in ten easy steps with COLA, J. Cosm. Astr. Phys.*, **6,** 036, 36. arXiv: 1301.0322 [astro-ph.CO].

Taylor A., Joachimi B., and Kitching T., 2013, *Putting the precision in precision cosmology: How accurate should your data covariance matrix be?, MNRAS*, **432,** 1928–1946. arXiv: 1212.4359 [astro-ph.CO].

Tegmark M., Taylor A. N., and Heavens A. F., 1997, *Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets?, Astrophys. J.*, **480,** 22–35. eprint: astro-ph/9603021.

Teyssier R., 2002, *Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES, A&A*, **385,** 337–364. eprint: astro-ph/0111367.

Teyssier R., Pires S., Prunet S., Aubert D., Pichon C., Amara A., Benabed K., Colombi S., Refregier A., and Starck J.-L., 2009, *Full-sky weak-lensing simulation with 70 billion particles, A&A*, **497,** 335–341. arXiv: 0807.3651.

The Dark Energy Survey Collaboration, 2005, *The Dark Energy Survey,* ArXiv Astrophysics e-prints, eprint: astro-ph/0510346.

Van Waerbeke L., and Mellier Y., 2003, *Gravitational Lensing by Large Scale Structures: A Review,* ArXiv Astrophysics e-prints, eprint: astro-ph/0305089.

Van Waerbeke L., Mellier Y., Erben T., Cuillandre J. C., Bernardeau F., Maoli R., Bertin E., McCracken H. J., Le Fèvre O., and Fort B., 2000, *Detection of correlated galaxy ellipticities from CFHT data: first evidence for gravitational lensing by large-scale structures, A&A*, **358,** 30–44. eprint: astro-ph/0002500.

Warren M. S., Abazajian K., Holz D. E., and Teodoro L., 2006, *Precision Determination of the Mass Function of Dark Matter Halos, Astrophys. J.*, **646,** 881–885. eprint: astro-ph/0506395.

Watson W. A., Iliev I. T., D'Aloisio A., Knebe A., Shapiro P. R., and Yepes G., 2013, *The halo mass function through the cosmic ages, MNRAS*, **433,** 1230–1245. arXiv: 1212.0095.

Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., and Rozo E., 2013, *Observational probes of cosmic acceleration, Phys. Rep.*, **530,** 87–255. arXiv: 1201.2434 [astro-ph.CO].

White M., and Hu W., 2000, *A New Algorithm for Computing Statistics of Weak Lensing by Large-Scale Structure, Astrophys. J.*, **537,** 1–11. eprint: astro-ph/9909165.

White M., Tinker J. L., and McBride C. K., 2014, *Mock galaxy catalogues using the quick particle mesh method, MNRAS*, **437,** 2594–2606. arXiv: 1309.5532.

White S. D. M., 1976a, *Dynamical friction in spherical clusters, MNRAS*, **174,** 19–28.

— 1976. b, *The dynamics of rich clusters of galaxies, MNRAS*, **177,** 717–733.

Wittman D. M., Tyson J. A., Kirkman D., Dell'Antonio I., and Bernstein G., 2000, *Detection of weak gravitational lensing distortions of distant galaxies by cosmic dark matter at large scales, Nature*, **405,** 143–148. eprint: astro-ph/0003014.

York D. G., Adelman J., Anderson J. E., Anderson S. F., Annis J., Bahcall N. A., Bakken J. A., Barkhouser R., Bastian S., and Berman E., 2000, *The Sloan Digital Sky Survey: Technical Summary, Astr. J.*, **120,** 1579–1587. eprint: astro-ph/0006396.

Zel'dovich Y. B., 1970, *Gravitational instability: An approximate theory for large density perturbations., A&A*, **5,** 84–89.

Zwicky F., 1933, *Die Rotverschiebung von extragalaktischen Nebeln*, Helvetica Physica Acta, **6,** 110–127.

A

TRANSIENT EFFECTS CORRECTION

The MICE-GC simulation used first order LPT initial conditions (Zeldovich approximation) at $z_i = 100$. This approximation is known to yield transient effects in the distribution of matter and haloes mainly depending on the actual initial redshift used (Scoccimarro, 1998; Crocce et al., 2006). In this appendix we quantify these effect on the matter power spectrum and the halo mass function by using first or second order LPT for the concrete configuration of MICE-GC. These differences are then corrected in MICE-GC whenever in Chapter 3 we compare any measurement against those in COLA, because otherwise it would be a source of a systematic error. We estimate them using additional GADGET-2 simulations using the same cosmology and starting redshift as in MICE-GC (see §2.1.1), and evolving 1024^3 particles using either first or second order LPT. The box size is $768\,h^{-1}\,\mathrm{Mpc}$ in order to keep the same mass resolution as in the reference MICE-GC N-body simulation. For the mass function we use as well a larger box size of $3072\,h^{-1}\,\mathrm{Mpc}$ in order to have good statistics at the high mass end.

The left panel of Fig. A.1 shows the transient effects in the matter power spectrum in real space for redshifts 0, 0.5 and 1 in solid, dashed and dotted lines respectively. The correction is always below 2 per cent in the scales studied and remarkably similar to the results that found by (Schneider et al., 2015) using another simulation code (see their fig. 2). The right panel displays the mass function and uses the same line styles. The vertical line at $M = 10^{13.5}\,h^{-1}\,\mathrm{M}_\odot$ marks the matching mass-scale for the two runs used (the smaller one for smaller masses and the other way around). Remind that halo masses are defined using the Warren correction and this enables a good overlapping of measurements at that matching mass-scale (in agreement with other tests for such correction, e.g.

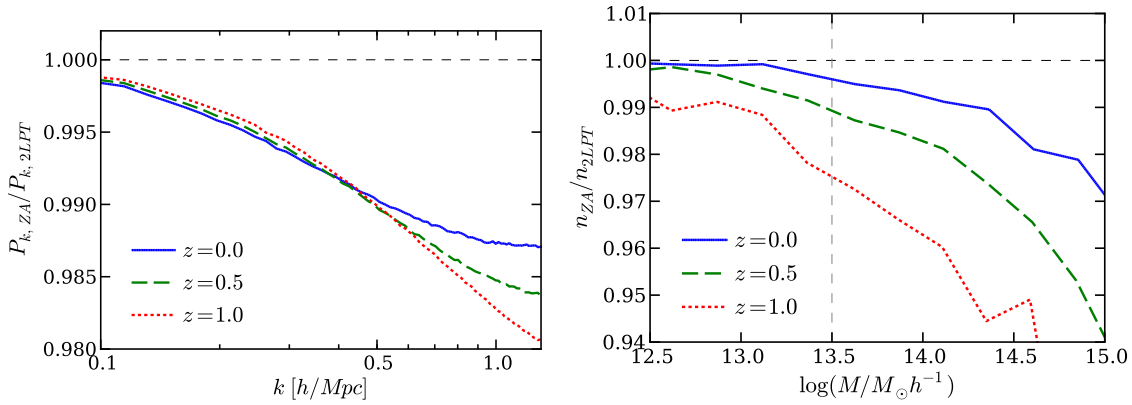

FIGURE A.1: Transient effects on the matter power spectrum (left panel) and the mass function (right panel). Solid, dashed and dotted lines correspond to redshifts 0, 0.5 and 1 respectively. We show the ratio of the observables measured on a pair of identical full N-body runs differing only in the set up of the initial conditions: first order versus second order LPT. The simulations used the same particle mass as MICE-GC in a box of $768\,h^{-1}\,\mathrm{Mpc}$, except for the mass function plot that for $M > 10^{13.5}\,h^{-1}\,\mathrm{M}_{\odot}$ (marked by a vertical dashed line) used a larger box of $3072\,h^{-1}\,\mathrm{Mpc}$.

Crocce et al., 2010). In the mass function, differences are more important at high masses and redshifts, going up to 5 per cent at $z = 1$ for the mass range of interest and they are within 3% at $z = 0$. We measured as well the correction for the halo-matter cross power spectrum, but we found it to be always within the 1 per cent so that the effect is negligible. Thus, whenever we show in Chapter 3 a ratio of either mass functions or matter power spectra with respect to MICE-GC, we have multiplied it by the corresponding ratio shown in Fig. A.1. In the case of halo clustering observables we find that transient effects are below the 1 per cent level, so we consider that we can neglect the correction for those measurements.

B

PERFORMANCE OF COLA WITH RESPECT TO PARTICLE-MESH ONLY RUNS

Tassev et al. (2013) showed that COLA simulations with as few as ten time steps recover the matter density field much better than just doing either a particle-mesh (PM) only simulation with the same number of time steps or a 2LPT evolution. In this thesis we advocate the use of more time steps in order to produce mock catalogues that are accurate in a large span of redshifts. After increasing the number of time-steps one might think that the 2LPT part of the COLA method has a very little contribution to the dynamics and much of the information comes from the PM integration. In this appendix we show the relative impact of the 2LPT information when many time steps are used.

For this exercise we use the FastPM[1] parallel implementation of COLA (Feng et al., 2016). We run several PM only and COLA runs using the same set-up: 768^3particles in a box of $576\,h^{-1}\,\mathrm{Mpc}$ by side, and we vary the number of time steps for the PM only runs (the initial redshift is fixed at $z_i = 19$). The green line in the left panel in Fig. B.1 shows that the PM only method recovers less power in the matter power spectrum than COLA for the same number of time steps. The deficit is larger at small scales and at high redshift (dashed and dotted lines correspond to redshifts 0.0 and 1.5). The plausible explanation is that the PM only has more difficulties to accurately integrate the equations of motions at high redshifts, when few time steps sample each e-fold of the growth of structures[2], and differences persist until $z = 0$. The PM method slowly converges to COLA at large

[1] https://github.com/rainwoodman/fastPM.

[2] Note that in those runs, time steps are linearly distributed with the scale factor.

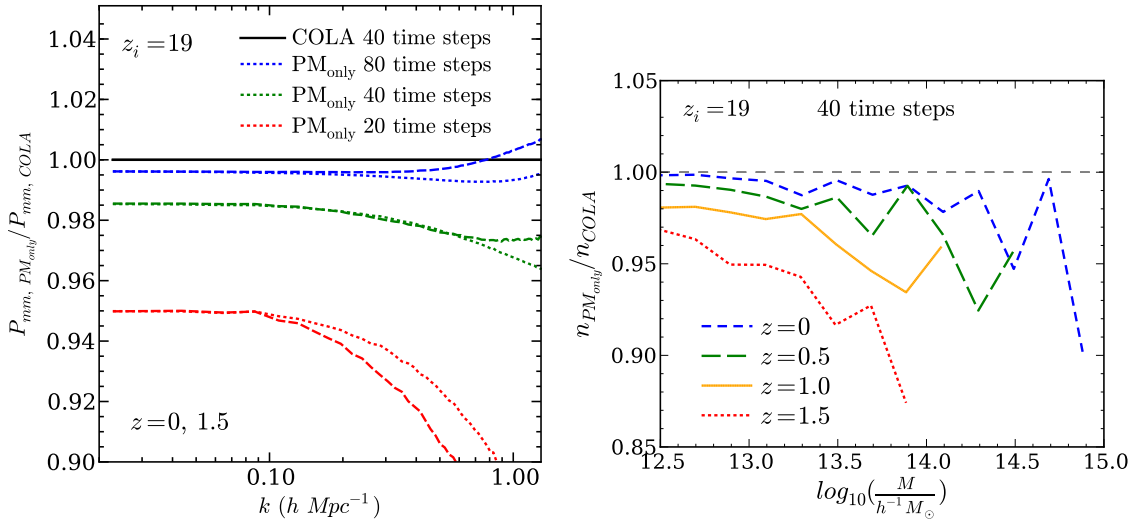

FIGURE B.1: *Left panel*: matter power spectra for three PM only simulations with 20, 40 and 80 time steps from bottom to top. The reference simulation is a COLA run with 40 time steps. Dotted and dashed lines correspond to redshifts 1.5 and 0 respectively. *right panel*: mass function of a PM only run with respect to a COLA one, both with 40 time steps. Plain PM simulations introduce additional systematic effects, even with as much as 40 time steps.

scales by increasing the number of time steps. In turn, we recall that COLA reproduces the linear growth rate accurately regardless of the number of time steps (see Fig. 3.3 in § 3.4.1). The right panel in Fig. B.1 displays the ratio of mass functions between the PM only and COLA runs with 40 time steps, for various redshifts.

There is a clear underestimation that reaches $5-10\%$ for both high redshifts and high masses. We have also studied the halo linear bias and found a corresponding excess exhibiting the same trends. Both differences can be explained by a systematic underestimation of the halo masses for plain PM simulations that is mass and redshift dependent. Both panels of Fig. B.1 show that, for a similar number of time-steps, differences between PM only and COLA decrease towards lower redshifts. As shown by the matter power spectrum, discrepancies originate at high redshift, and late non-linear evolution masks them (in a similar fashion than transient effects showed in Appendix A). We do not show in Fig. B.1 full N-body values since we want to focus only on the relative effect of both methods. Also, it is clear that plain PM simulations converge to COLA in the limit of a large number of time steps, as both tend asymptotically to a full N-body.

One might argue that this effect on the mass function could apparently solve the overestimation studied in § 3.4.4. But this seems just a cancellation of errors that might introduce even more undesired systematic effects. We conclude that it is worth using the COLA method even with as many as 40 time steps. Less time-steps (e.g., 20 or fewer) still produce accurate results for COLA, while plain PM simulations show non-negligible biases.