



HAL
open science

A fast semidiscrete optimal transport algorithm for a unique reconstruction of the early Universe

Bruno Lévy, Roya Mohayaee, Sebastian von Hausegger

► To cite this version:

Bruno Lévy, Roya Mohayaee, Sebastian von Hausegger. A fast semidiscrete optimal transport algorithm for a unique reconstruction of the early Universe. *Monthly Notices of the Royal Astronomical Society*, 2021, 506 (1), pp.1165-1185. <10.1093/mnras/stab1676>. <hal-03115841>

HAL Id: hal-03115841

<https://hal.science/hal-03115841v1>

Submitted on 6 Jul 2023

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

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



HAL Authorization

A fast semidiscrete optimal transport algorithm for a unique reconstruction of the early Universe

Bruno Levy¹ ,¹★ Roya Mohayaee²★ and Sebastian von Hausegger^{1,2,3}★ 

¹Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

²CNRS, Institut d'Astrophysique de Paris, Sorbonne Université, 98bis Bld Arago, E-75014 Paris, France

³Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Parks Road, Oxford OX1 3PU, UK

Accepted 2021 May 23. Received 2021 April 15; in original form 2020 December 16

ABSTRACT

We leverage powerful mathematical tools stemming from optimal transport theory and transform them into an efficient algorithm to reconstruct the fluctuations of the primordial density field, built on solving the Monge-Ampère-Kantorovich equation. Our algorithm computes the optimal transport between an initial uniform continuous density field, partitioned into Laguerre cells, and a final input set of discrete point masses, linking the early to the late Universe. While existing early universe reconstruction algorithms based on fully discrete combinatorial methods are limited to a few hundred thousand points, our algorithm scales up well beyond this limit, since it takes the form of a well-posed smooth convex optimization problem, solved using a Newton method. We run our algorithm on cosmological N -body simulations, from the AbacusCosmos suite, and reconstruct the initial positions of $\mathcal{O}(10^7)$ particles within a few hours with an off-the-shelf personal computer. We show that our method allows a unique, fast, and precise recovery of subtle features of the initial power spectrum, such as the baryonic acoustic oscillations.

Key words: software: data analysis – software: development – early Universe – large scale structure of Universe .

1 INTRODUCTION

Optimal Transport theory has found spectacular applications in diverse areas of science, from economics to biology, physics, data science, and machine learning to name but a few. The emerging applications of the two-centuries-old theory is mainly due to advances in mathematics and the developments of fast new algorithms. It is indeed these fundamental advances that have paved the way for major breakthroughs in artificial intelligence, since they made it possible to compute a natural, *Wasserstein*, distance between entities of various nature, essential for object recognition and classification.

The reason behind the success of optimal transport theory in physics might be that it describes a universal foundation of Nature, where most processes seem to be governed by the optimization of an underlying, sometimes unknown, quantity. Light follows a path that, roughly, minimizes time by Fermat's principle, and freely moving test particles follow time-like geodesics in general relativity. As perhaps best phrased by Euler more than two centuries ago: 'nothing arises in the universe in which one cannot see the sense of some maximum or minimum.' The variational principle founded by Euler himself and at the basis of classical mechanics and quantum field theory has now become a subset of the vast field of variational calculus.

The Euler–Lagrange action optimization, often referred to as the least action principle, found its application in cosmology, thanks to the pioneering work of Peebles back in 1989. Peebles aimed at reconstructing the past history of the Local Group by retracing the trajectories of around 10 member-galaxies back in time (Peebles 1989).

He showed that the recovery of the initial conditions (IC) without prior knowledge of present velocities is possible by considering that, at early times, the peculiar velocities of matter is negligible and their spatial distribution uniform. He thus solved a mixed boundary-value problem instead of the usual Cauchy problem (Peebles 1989). The method not only provided valuable information on the orbits of the members of the Local Group but also put constraints on cosmological parameters. Indeed a low value of the matter density parameter was favoured by Peebles' action minimization at a time when the Cold Dark Matter (CDM) model was the standard paradigm. Peebles' algorithm was made more efficient for application to larger data sets but lacked uniqueness: existence of multiple minima, maxima, and saddle points in the landscape of solutions lead to multiple trajectories, all physically viable (Peebles 1994; Shaya, Peebles & Tully 1995; Nusser & Branchini 2000; Branchini, Eldar & Nusser 2002). Most recently, the method's excess complexity has been reduced in new numerical schemes and also by smoothing out strongly non-linear scales (Sarpa et al. 2019, 2020).

Most action optimizations involve finding minimum energy trajectories for fixed end points. The added complexity of the cosmological setting in Peebles' formulation is that in addition to the trajectories, the initial positions of galaxies are also unknown, which renders the problem highly under-determined. Can we add suitable constraints on the trajectories to achieve a unique solution? On large scales¹ – roughly scales of tens of megaparsecs – the velocities of galaxies, as tracers of the underlying dark matter fluid, remain a potential flow, as

* E-mail: [Bruno.Levy@inria.fr](mailto: Bruno.Levy@inria.fr) (BL); [mohayaee@iap.fr](mailto: mohayaee@iap.fr) (RM); [sebastian.vonhausegger@physics.ox.ac.uk](mailto: sebastian.vonhausegger@physics.ox.ac.uk) (SVH)

¹Unlike fluid mechanics, cosmology lacks a proper control parameter, e.g. a Reynolds number, and we can only separate single from multistream regimes in an empirical manner.

has been shown numerically by N -body simulations and theoretically at least up to the third-order in the Lagrangian perturbation theory (e.g. see Catelan (1995)). In previous works, we showed that where the trajectories of the fluid elements have not crossed or when their velocity is the gradient of a convex potential, the cosmological reconstruction problem is a well-posed instance of what is called *the optimal mass transportation problem* and has a unique solution (Frisch et al. 2002; Brenier et al. 2003).

The mass transportation problem² was initially formulated by Monge during the French revolution. Monge aimed at finding how to transport soil from N number of excavated holes to the same number of rubbles while minimizing the total product of the transported mass and the travelled distance (Monge 1784). This problem has a rich mathematical structure that was revealed later by a continuous stream of advances both in fundamental mathematics (see review by Villani (2009)) and in applied mathematics (see review by Peyré & Cuturi (2017)). The most prominent ‘quantum leap’ was made during WWII by Kantorovich, who invented the ‘mathematical toolbox’ to study the existence and uniqueness (Kantorovich 1942). Kantorovich studied a relaxation of the problem,³ and introduced the dual formulation, with Lagrange multipliers. The relaxed problem was subsequently referred to as the Monge–Kantorovich mass transportation problem. From our cosmological perspective, these Lagrange multipliers correspond to the initial and final gravitational potentials, related to each other by the Legendre–Fenchel transform.

In the 90’s, another ‘quantum leap’ was made by Brenier with his celebrated *polar factorization theorem*, which states that the optimal-transport map corresponds to the gradient of a convex potential (Brenier 1991). By injecting the gradient of the potential into the mass conservation constraint of the Monge–Kantorovich problem, one obtains a non-linear partial-differential equation (PDE) known as the *Monge–Ampère equation*, which can be solved to find the potential.⁴ After the polar factorization theorem was discovered, Benamou & Brenier (2000) revealed that the minimized quantity, the *Wasserstein distance*, corresponds to the *action integral* in an incompressible Euler fluid.

In our cosmological setting, it also corresponds to the action integral in a simple model of self-gravitating matter – that is the integrated kinetic energy. These theorems give us a way of computing the potential, and reconstructing the trajectories. For example, in the Zel’dovich regime, each fluid element follows a rectilinear motion with a constant speed that corresponds to the gradient of the reconstructed primordial gravitational potential. In the language of optimal transport, the gravitational potential can be deduced from the Lagrange multiplier of the mass conservation constraint, through the Legendre–Fenchel transform. This characterization implies that the potential is a convex function. In terms of physics, it implies that there is no multistreaming in the reconstructed dynamics.

To our knowledge, in the cosmological context, one of the first methods of this kind was by Nusser et al., who studied and exploited the same type of relation between the initial density and velocity field (Nusser et al. 1991).⁵ Later, the Monge–Ampère–Kantorovich (MAK) cosmological reconstruction method was developed based on

advances in optimal transport theory and subsequently solved using a fully discrete combinatoric algorithm (Frisch et al. 2002; Brenier et al. 2003). The algorithm was tested on simulations (Mohayaee et al. 2003, 2006; Lavaux et al. 2008; Mohayaee & Sobolevskii 2008) and also applied to galaxy redshift surveys (Mohayaee & Tully 2005; Lavaux et al. 2010) and found applications in condensed-matter physics (e.g. Aurell et al. (2012)). However, its cubic algorithmic complexity rendered it impractical for applications to challenging and forefront cosmological problems such as the reconstruction of the sound horizon at decoupling, i.e. the scale of the baryon acoustic oscillations (BAO; Eisenstein et al. 2005, 2007b).

In this article, we design a highly efficient new algorithm, which yields a *unique* solution by construction and can be applied to computationally demanding problems in cosmology. We demonstrate the efficiency and accuracy of our algorithm through the example of BAO reconstruction. In our case, the IC is a uniform density field (continuous), and the final one corresponds to a set of galaxies, represented by a (discrete) set of points, hence a *semidiscrete* optimal transport problem. The theory of optimal transport is written in a mathematical language (theory of probability measures) that is general enough to encompass such irregular settings (with mass concentrated on points). Not only this mathematical language is exactly what we need to model our cosmological problem, but also it can be directly translated into a computational algorithm that can be efficiently implemented on a computer: solving the underlying semidiscrete Monge–Ampère equation is equivalent to minimizing a *smooth* and *convex* function that depends on the final gravitational potential at each discrete point. It is much faster than solving a *discrete* Monge–Ampère equation, as we did before, which required exploring a huge combinatorial space (Frisch et al. 2002; Brenier et al. 2003). The convergence between the three aspects of the problem (cosmology, mathematics, and computer science) results in a new algorithm that solves the assignment problem with an empirical complexity of $M \log N$ (as compared to N^3), and that can be applied to reconstruction problems of unprecedented sizes: $O(10^6)$ particles in *minutes*, to be compared with *months*, and $O(10^8)$ particles in hours. Using our semidiscrete algorithm on a set of N -body simulations taken from the AbacusCosmos suite (Garrison et al. 2018; Garrison, Eisenstein & Pinto 2019), we examine the complexity of our algorithm and compare reconstructed and simulated initial density fields, their power spectra and correlation functions, starting from two different redshifts. We show that these quantities can be accurately reconstructed above scales of a few Mpc with competitive numerical speed. In particular, we show that BAO can be reconstructed with high accuracy and speed both in the power spectrum and the correlation function.

Indeed, the BAO scale has much decisional power, by providing a rather robust standard ruler of cosmology. As a powerful distance indicator, BAO measurement can probe the acceleration phase in the expansion history of the Universe and distinguish between the theory of general relativity and those of modified gravity. It is also a pleasantly obvious feature of the two-point correlation function of the galaxy field. Non-linear gravitational evolution at late times slightly softens this feature and therewith the statistical certainty with which the BAO scale can be determined. Beyond general (but not less important) questions about, e.g., statistical properties of the initial density fluctuations, undoing the non-linear evolution to obtain the linear density field from low-redshift measurements of the large-scale structure give strong motivation for reconstruction techniques (e.g. Eisenstein et al. (2005), Eisenstein et al. (2007b), Seo et al. (2010), Padmanabhan et al. (2012)). Since the pioneering variational method of Peebles, numerous reconstruction techniques, for different

²Which historically precedes the variational calculus of Euler and Lagrange mentioned at the beginning of this introduction.

³Where the unknown is the ‘graph’ of the function in the product space, referred to as the ‘transport plan’. More on this in Section 3.2.

⁴Interestingly, this very class of PDE was also first studied by Monge during the French revolution and later generalized by Ampère at the beginning of the 1800’s.

⁵even if they do not mention the Monge–Ampère equation explicitly.

tasks and not just the BAO retrieval, have been proposed. Here, we can only mention a few. Many of these methods take a probabilistic approach (e.g. Weinberg (1992), Kitaura & Enßlin (2008), Enßlin, Frommert & Kitaura (2009), Neyrinck, Szapudi & Szalay (2011), Jasche & Wandelt (2013), Cautun et al. (2014), Bos, Kitaura & van de Weygaert (2019)), a few others are perturbative (e.g. Nusser & Dekel (1992), Gramann (1993), Kashlinsky (1998), Eisenstein et al. (2007b), Schmittfull, Baldauf & Zaldarriaga (2017)), and many are variational (e.g. Croft & Gaztanaga (1997), Narayanan & Croft (1999), Monaco & Efstathiou (1999), Wang et al. (2017), Shi, Cautun & Li (2018)). We refer the reader to section 7 of Brenier et al. (2003) for a detailed discussion of these categories of reconstruction methods. Here, we present our algorithm that, by contrast to many of the other methods, yields a unique solution by construction, is deterministic, model-independent, and computationally efficient.

This article is structured as follows: we first provide a review of the Lagrangian dynamics in an expanding Universe and show the generality and limitations of our two hypothesis of gradient flow and convexity (Section 2). Then we explain old methods and our new algorithm to solve the underlying assignment problem (Section 3), before giving the details of our numerical solution mechanism (Section 4). Finally, we test the algorithm against the ABACUSCOSMOS simulations in Section 5. In Section 6, we conclude and comment on how galaxy biasing and redshift space distortion can be incorporated.

2 MONGE-AMPÈRE-KANTOROVICH (MAK) RECONSTRUCTION

2.1 Problem setting

We consider the problem of reconstructing the fluctuations in the IC of self-gravitating matter governed by the following equations, in Eulerian form (see Peebles (1980), or Appendix A of Brenier et al. (2003)):

$$\partial_\tau \mathbf{v} + (\mathbf{v} \cdot \nabla_x) \mathbf{v} = -\frac{3}{2\tau} (\nabla_x \phi + \mathbf{v}) \quad (1)$$

$$\partial_\tau \rho + \nabla_x \cdot (\rho \mathbf{v}) = 0 \quad (2)$$

$$\nabla_{xx}^2 \phi = \frac{\rho-1}{\tau} \quad (3)$$

where ρ is the density field defined over $V \subset \mathbb{R}^3$, ϕ the gravitational potential, τ is the growth rate of structures, used as a time variable, and normalized in such a way that $\tau = \tau_I = 0$ corresponds to the IC, and $\tau = \tau_F = 1$ corresponds to present time, and \mathbf{x} denotes the co-moving coordinates. The peculiar velocity field \mathbf{v} is also expressed as a function of the co-moving coordinates \mathbf{x} .

Equation (1) is the momentum (Euler) equation. Its right-hand side has two terms that have opposite effects: the first term ($-\frac{3}{2\tau} \nabla_x \phi$) corresponds to the effect of gravity, which tends to collapse structures, and the second term ($-\frac{3}{2\tau} \mathbf{v}$), called the Hubble–Lemaître drag, corresponds to the effect of expansion, which slows down the collapsing effect of gravity. Equation (2) enforces mass conservation (continuity equation) and equation (3) is the Poisson equation that governs the gravitational potential.

Given the density field ρ_F at time $\tau_F = 1$ that corresponds to the present distribution of mass, our goal is to reconstruct the initial fluctuations of $\rho(\cdot, \tau_I + \epsilon)$ for a small ϵ . One can also consider the problem of reconstructing the full dynamics of the system $\rho(\tau, \mathbf{x})$ and $\mathbf{v}(\tau, \mathbf{x})$ for $\tau \in [\tau_I, \tau_F]$ and $\mathbf{x} \in V$. Clearly, the problem is under-determined but, as we shall show later, under some reasonable

simplifying assumptions, it can be replaced by a well-posed convex optimization problem.

2.2 Lagrangian perturbation theory

One can observe that at the IC $\tau_I = 0$, for the right-hand side of the Poisson equation (3) to be defined, density needs to be uniform $\rho_I(\cdot) = \rho(\cdot, \tau_I) = 1$. The same consideration for the right-hand side of the momentum equation (1) implies that at the IC, the velocity coincides with the (negated) gradient of the potential $\mathbf{v}_I(\cdot) = \mathbf{v}(\cdot, \tau_I) = -\nabla_x \phi(\cdot, \tau_I)$. This condition, which also equivalently arises as a solution to the linearized set of equations (1–3), is sometimes referred to as *slaving*⁶. Consider now the Lagrangian point of view, and focus on the mass element that is at position \mathbf{q} at time τI . Denote its trajectory by $\mathbf{x}(\mathbf{q}, \tau)$. Its initial speed at time τI is given by $\mathbf{v}_I(\mathbf{q}) = -\nabla_q \phi(\mathbf{q}, \tau_I) = -\nabla_q \phi I(\mathbf{q})$. In 1D, one can prove that the speed of the mass element remains constant at any time τ (see e.g. Brenier et al. (2003) for the proof). In other words, integrating (1), (2), and (3) in 1D results in a uniform rectilinear motion for all mass elements:

$$\mathcal{D}_\tau \mathbf{x}(\mathbf{q}, \tau) = (1/\tau_F)(\mathbf{x}_F(\mathbf{q}) - \mathbf{q}) = -\nabla_q \phi I(\mathbf{q}) \quad \forall \tau \quad (4)$$

$$\mathbf{x}(\mathbf{q}, \tau) = (1 - \tau/\tau_F)\mathbf{q} + (\tau/\tau_F)\mathbf{x}_F(\mathbf{q}) \quad (5)$$

where $\mathbf{x}_F(\mathbf{q}) = \mathbf{x}(\mathbf{q}, \tau_F)$ and where \mathcal{D}_τ denotes the Lagrangian derivative w.r.t. time τ . It also means that in 1D, to determine the entire motion, one only needs to know the initial potential $\phi I(\cdot) = \phi(\cdot, \tau_I)$ at time $\tau_I = 0$.

In 3D, for a small time τ , the speed of a mass element is still given by (4), but it is no longer strictly the case at any time. However it is considered to be a reasonable approximation (Zel’dovich 1970). This approximation means that the r.h.s. of the momentum equation (1) vanishes. Physically, it means that the Hubble–Lemaître drag exactly counter-balances the effect of gravity, implying that each mass element has a uniform rectilinear motion (5). In this setting, to reconstruct the full dynamics, one just needs to determine the $V \rightarrow V$ map $\mathbf{q} \mapsto \mathbf{x}_F(\mathbf{q})$. This map is in turn completely determined by the potential $\phi I(\cdot) = \phi(\cdot, 0)$, using the relation $\mathbf{x}_F(\mathbf{q}) = \mathbf{q} + \tau_F \mathbf{v}_I(\mathbf{q}) = \mathbf{q} - \tau_F \nabla_q \phi I(\mathbf{q})$.

At this point, given the potential ϕI at the IC (we will see how to compute it in Section 4), we can reconstruct the Zel’dovich approximation. This gives us for the mass particle that was located at \mathbf{q} at time τI its position $\mathbf{x}_F(\mathbf{q}) = \mathbf{q} - \tau_F \nabla_q \phi I(\mathbf{q})$ at the present time $\tau_F = 1$. In other words, this gives us the *assignment* between the IC and the present distribution of mass from which we can obtain the particle positions at arbitrary times (τ), up to the first-order Lagrangian perturbation theory, as

$$\mathbf{x}(\tau, \mathbf{q}) = \mathbf{q} + \frac{\tau}{\tau_F}(\mathbf{x}_F(\mathbf{q}) - \mathbf{q}). \quad (6)$$

Although the assignment between \mathbf{q} and \mathbf{x}_F is valid for as long as the convexity holds, we limit ourselves to the first order only for obtaining the particle positions at intermediate times. The main reason is that here we test our method with the goodness of reconstruction of

⁶Here, the IC is considered from a mathematical point of view. From a physical point of view, clearly, there cannot be a non-uniform potential associated with a uniform density. In fact, it is at $\tau_I + \epsilon$ that the potential is non-uniform, but one can make this ϵ arbitrarily small. In a certain sense, one can ‘push’ the non-uniform potential from $\tau_I + \epsilon$ towards τ_I : the right-hand side of the Poisson equation for the potential (equation 3) with τ in the denominator results in a significant potential yielded by tiny fluctuations of the density at the IC.

BAO. It happens that often one adds additional, broad-band, fitting terms to the power spectrum which takes care of the mode-coupling as well as other effects such as the shot noise. The implementation of the second- and higher order Lagrangian perturbation theory into our algorithm shall be reported in the forthcoming works.

2.3 Least action principle and optimal transport

One can also obtain the momentum equation (1) by extremizing the following action integral (Brenier et al. (2003); appendix D):

$$I = \frac{1}{2} \int_{\tau I}^{\tau F} \int_V (\rho |\mathbf{v}|^2 + \frac{3}{2} |\nabla_x \phi|^2) \tau^{3/2} d^3 \mathbf{x} d\tau \quad (7)$$

subject to mass conservation (2), to the Poisson equation for the potential (3), and to the boundary conditions:

$$\rho(\cdot, \tau I) = \rho I = 1 \quad ; \quad \rho(\cdot, \tau F) = \rho_F(\cdot) \quad (8)$$

where ρI denotes the (uniform) density at the IC and ρ_F denotes the density map at present time $\tau = \tau F$. Using the method of Lagrange multipliers and varying ρ , one obtains the momentum equation (1).

Consider now an approximation, where the second term of the integrand and the $\tau^{3/2}$ factor are removed (Giavalisco et al. 1993), which may be thought of as replacing the 3/2 coefficient by $3\alpha/2$ and making α tend to zero. The action integral (7) then becomes:

$$I = \frac{1}{2} \int_{\tau I}^{\tau F} \int_V \rho |\mathbf{v}|^2 d^3 \mathbf{x} d\tau. \quad (9)$$

Note that the integrand has only the kinetic energy, and no longer any potential energy. This again corresponds to the Zel'dovich approximation. Given the boundary conditions (8), now we want to find the motion that minimizes the kinetic energy. If we have a single mass particle, it is easy to see that minimizing the action results in a uniform rectilinear motion (Landau & Lifshitz 1975). It can be proved (Benamou & Brenier 2000) that this is still the case for any number of particles, or even for a continuous density field ρ : extremizing the action (9) means that all mass particles (or all elementary mass elements for a continuous ρ) move in a uniform rectilinear manner. In other words, finding the motion $\mathbf{x}(\mathbf{q}, \tau)$ that minimizes the action I on $V \times [\tau I, \tau F]$ is equivalent to finding the map $\mathbf{x}_F: V \rightarrow V$ that gives the position at present time $\mathbf{x}_F(\mathbf{q}) = \mathbf{x}(\mathbf{q}, \tau F)$ of the mass element that was at position \mathbf{q} at the IC $\tau I = 0$. The map \mathbf{x}_F minimizes the following functional:

$$\inf_{\mathbf{x}_F} \int_V \rho(\mathbf{q}) |\mathbf{q} - \mathbf{x}_F(\mathbf{q})|^2 d^3 \mathbf{q} \quad (10)$$

subject to mass conservation (2) and to the boundary conditions (8). Now, it may be more natural to write mass conservation in Lagrangian coordinates. Using $\rho(\mathbf{x}(\mathbf{q}, \tau), \tau) = \rho(\mathbf{q}) / (\det \nabla_{\mathbf{q}} \mathbf{x})$, the mass conservation constraint writes:

$$\rho_F(\mathbf{x}_F(\mathbf{q})) \det(\nabla_{\mathbf{q}} \mathbf{x}_F(\mathbf{q})) = \rho I \quad \forall \mathbf{q} \quad (11)$$

The minimization of expression (10) subject to (11) is referred to as *Monge's optimal transport problem* (Monge 1784).

2.4 The convex Kantorovich potential and the Monge–Ampère equation

Consider Monge's optimal transport problem (10). Introducing the Lagrange multiplier $\Psi: V \rightarrow V$ associated with the constraint (11), and using the identity $|\mathbf{x}_F(\mathbf{q}) - \mathbf{q}|^2 = \mathbf{x}_F(\mathbf{q})^2 - 2\mathbf{q} \cdot \mathbf{x}_F(\mathbf{q}) + \mathbf{q}^2$, the optimal transport problem can be written as the following saddle

point problem:

$$\sup_{\mathbf{x}_F} \inf_{\Psi} \left[J(\mathbf{x}_F) = \int_V \rho I \mathbf{x}_F(\mathbf{q}) \cdot \mathbf{q} d^3 \mathbf{q} - \int_V \Psi(\mathbf{x}_F(\mathbf{q})) \rho(\mathbf{q}) d^3 \mathbf{q} + \int_V \Psi(\mathbf{x}) \rho_F(\mathbf{x}) d^3 \mathbf{x} \right]. \quad (12)$$

Formally, the first-order optimality condition w.r.t. \mathbf{x}_F writes:

$$\begin{aligned} \frac{\partial J}{\partial \mathbf{x}_F} = 0 &\Rightarrow \rho I \mathbf{q} = \rho I \nabla \Psi(\mathbf{x}_F(\mathbf{q})) \\ &\Rightarrow \mathbf{q} = \nabla_x \Psi(\mathbf{x}_F(\mathbf{q})) \end{aligned} \quad (13)$$

The second-order optimality condition writes:

$$\begin{aligned} \frac{\partial^2 J}{\partial \mathbf{x}_F^2} \geq 0 &\Rightarrow D^2 \Psi \geq 0 \\ &\Rightarrow \Psi \text{ is a convex function} \end{aligned} \quad (14)$$

where $D^2 \Psi$ denotes the determinant of the Hessian matrix of Ψ (the derivations above (13), (14) just give an intuition, the reader is referred to (Brenier 1991) for a rigorous proof).

From the first-order and second-order optimality conditions, we learn that the function $\mathbf{x} \mapsto \mathbf{q}$ that maps a mass element \mathbf{x} at present time τF back to its initial position \mathbf{q} at time τI is the gradient of a convex potential Ψ (called the *Kantorovich potential*).

Next, we study the forward map $\mathbf{q} \mapsto \mathbf{x}_F$. It can be observed that the variables \mathbf{q} and \mathbf{x}_F play a symmetric role in the optimal transport problem. By exchanging the roles of $\mathbf{x}_F \leftrightarrow \mathbf{q}$, one can find a similar relation for the forward map $\mathbf{q} \mapsto \mathbf{x}_F(\mathbf{q})$:

$$\mathbf{x}_F(\mathbf{q}) = \nabla_{\mathbf{q}} \Phi(\mathbf{q}) \quad ; \quad \Phi \text{ is a convex function.} \quad (15)$$

It can be shown that this symmetry implies a relation between Ψ and Φ : they are the Legendre–Fenchel transform⁷ of each other, given by:

$$\forall \mathbf{q}, \Phi(\mathbf{q}) = \Psi^*(\mathbf{q}) \quad \text{where } \Psi^*(\mathbf{q}) = \sup_{\mathbf{x}} [\mathbf{x} \cdot \mathbf{q} - \Psi(\mathbf{x})]. \quad (16)$$

Next, we recall that the map $\mathbf{x}_F(\cdot)$ is determined by the gravitational potential $\phi I(\cdot)$ at the IC by $\mathbf{x}_F(\mathbf{q}) = \mathbf{q} - \nabla_{\mathbf{q}} \phi I(\mathbf{q})$. This gives us the relation between the gravitational potential ϕI at time τI and the convex Kantorovich potential Φ associated with the map $\mathbf{q} \mapsto \mathbf{x}_F(\mathbf{q})$:

$$\begin{aligned} \mathbf{x}_F(\mathbf{q}) &= \mathbf{q} - \nabla_{\mathbf{q}} \phi I(\mathbf{q}) \\ &= \nabla_{\mathbf{q}} \Phi(\mathbf{q}) \\ &\Rightarrow \Phi(\mathbf{q}) = 1/2 \mathbf{q}^2 - \phi I(\mathbf{q}) \end{aligned} \quad (17)$$

The insertion of $\mathbf{x}_F(\mathbf{q}) = \nabla_{\mathbf{q}} \Phi(\mathbf{q})$ into the (Lagrangian) mass conservation constraint (11) yields

$$\rho_F(\nabla_x \Phi(\mathbf{q})) D^2 \Phi(\mathbf{q}) = \rho I, \quad (18)$$

where $D^2 \Phi(\cdot)$ denotes the determinant of the Hessian matrix of Φ . equation (18) is referred to as the *Monge–Ampère equation* (or MA equation for short).

In our context, the convexity of the Kantorovich potential has an interesting consequence: it implies that there is no multistreaming in the reconstructed dynamics. It can be proven by contraction: Consider two distinct points $\mathbf{q}_1, \mathbf{q}_2$ and their images $\mathbf{x}_F(\mathbf{q}_1), \mathbf{x}_F(\mathbf{q}_2)$ through the $\mathbf{x}_F(\cdot)$ map. They move along the following trajectories:

$$\begin{aligned} \mathbf{x}_1(\tau) &= (1 - \tau) \mathbf{q}_1 + \tau \mathbf{x}_F(\mathbf{q}_1) \\ \mathbf{x}_2(\tau) &= (1 - \tau) \mathbf{q}_2 + \tau \mathbf{x}_F(\mathbf{q}_2). \end{aligned} \quad (19)$$

⁷The Legendre–Fenchel transform plays an important role in mechanics and thermodynamics, since it corresponds to the relation between Hamilton and Lagrange equations, see for instance (Landau & Lifshitz 1975), chapter 7.

If there was multistreaming, then there would exist a time τ such that $\mathbf{x}_1(\tau) = \mathbf{x}_2(\tau)$, or:

$$\begin{aligned} (1 - \tau)\mathbf{q}_1 + \tau\mathbf{x}_F(\mathbf{q}_1) &= (1 - \tau)\mathbf{q}_2 + \tau\mathbf{x}_F(\mathbf{q}_2) \\ (1 - \tau)\mathbf{q}_1 + \tau\nabla_q\Phi(\mathbf{q}_1) &= (1 - \tau)\mathbf{q}_2 + \tau\nabla_q\Phi(\mathbf{q}_2) \\ (1 - \tau)(\mathbf{q}_2 - \mathbf{q}_1) + \tau(\nabla_q\Phi(\mathbf{q}_2) - \nabla_q\Phi(\mathbf{q}_1)) &= 0 \\ (1 - \tau)(\mathbf{q}_2 - \mathbf{q}_1) + \tau(\nabla_q\Phi(\mathbf{q}_2) - \nabla_q\Phi(\mathbf{q}_1)) &= 0. \end{aligned} \quad (20)$$

The last line in (20) contradicts the convexity of Φ , that implies that $(1 - \lambda)(\mathbf{q}_2 - \mathbf{q}_1) + \lambda\nabla_q\Phi(\mathbf{q}_2 - \mathbf{q}_1)$ is strictly greater than zero for all $\mathbf{q}_1 \neq \mathbf{q}_2 \in V$ and $\lambda \in [0, 1]$ ■

2.5 An overall account of this section

To summarize, given the density at present time ρ_F and the density at the IC $\rho_I = 1$, our goal is to find the assignment map $\mathbf{q} \mapsto \mathbf{x}_F(\mathbf{q})$ that determines the assignment between the points \mathbf{q} at the IC and the points \mathbf{x} at the current time. It has the following properties:

- (i) $\mathbf{x}_F(\cdot)$ is the minimizer of

$$\inf_{\mathbf{x}_F} \left[\int_V |\mathbf{q} - \mathbf{x}_F(\mathbf{q})|^2 \rho_I d^2\mathbf{q} \right]$$

subject to mass conservation (11);

- (ii) $\mathbf{x}_F(\cdot)$ is also the gradient of the (convex) Kantorovich potential Φ ;

- (iii) Φ is the solution of the Monge–Ampère equation (18);

(iv) the convexity of Φ implies that there is no multistreaming in the reconstructed dynamics;

- (v) Φ is related to the gravitational potential at the IC ϕ_I by:

$$\Phi(\mathbf{q}) = 1/2\mathbf{q}^2 - \phi_I(\mathbf{q}).$$

From the assignment map $\mathbf{x}_F(\cdot)$, it is (optionally) possible to reconstruct higher order dynamics using Lagrangian perturbation theory (Buchert 1993; Catelan 1995) and at first order using the expression we have given in (6).

3 SOLVING THE ASSIGNMENT PROBLEM

In this section, we describe numerical solution mechanisms to compute the assignment map $\mathbf{x}_F(\cdot)$. We first review the existing methods that are based on a discretization of the density ρ_I at the IC and a discretization of the density ρ_F at current time (Sections 3.1 and 3.2). Then we present our method (based on semidiscrete optimal transport), which uses a continuous representation of the initial density ρ_I and a discrete representation of the density ρ_F , hence a *semidiscrete* method (Section 3.3).

3.1 Discrete-discrete MAK reconstruction

We consider (for now) that the density at the IC ρ_I and the density at present time ρ_F are both represented in discrete form, by a set of N particles. We consider that each particle i has a mass $\mu_i = 1/N$:

- (i) At the IC τI , the mass distribution ρ_I is represented by a set of N points \mathbf{q}_i , $i = 1 \dots N$. Since the initial distribution of mass is uniform at τI , the points \mathbf{q}_i are organized on a regular grid;

- (ii) at present time τ_F , the distribution of mass ρ_F is represented by a set of (the same number N) points \mathbf{x}_j , $j = 1 \dots N$.

In this setting, the initial problem of finding the (continuous) map $\mathbf{x}_F(\cdot): V \rightarrow V$ is replaced with finding which point \mathbf{x}_j corresponds to

each point \mathbf{q}_i . The discrete version of Monge’s problem (10) writes:

$$\inf_{\pi} \sum_j |\mathbf{q}_j - \mathbf{x}_{\pi(j)}|^2 \quad (21)$$

where $\pi: [1 \dots N] \rightarrow [1 \dots N]$ is a permutation of the indices.

Note that the discrete Monge problem (21) is purely combinatorial. Conceptually, one can imagine solving it by systematically testing the $N!$ possible permutations π . Clearly, it is not feasible in practice. However, there exists more efficient algorithms which guarantee that the optimal assignment is found. Faster assignment algorithms have been developed with polynomial complexity (Hénon 2002; Bertsekas & Castanon 1989). The latest algorithm developed by M. Hénon and used in our previous works (Brenier et al. 2003), which is a cosmological adaptation of the auction algorithm of Bertsekas, scales approximately as N^3 (for relevant details see, e.g. (Bertsekas 1992)). Later improvements of the auction algorithm allowed to make it faster (see Mérigot & Thibert (2020), Section 3). However, even with these improvements, these combinatorial algorithms remain slow for all practical purposes and this has been a major obstacle for the progress of MAK reconstruction in the past few years and since its first application to cosmology. Such algorithms have been proved too slow for the cosmological analyses of large data sets, or those that require repeated reconstructions. A notable example of such an instance is the reconstruction of detailed features of the primordial density fluctuation field or the primordial power-spectrum and in particular the reconstruction of the baryonic acoustic oscillations. For a proper reconstruction of acoustic peaks, one needs to treat extremely large data sets and/or carry reconstruction on a very large number of simulations for a proper handling of errors.

3.2 MAK duality

We now exhibit more structure in the discrete Monge problem (21), and its relation with the gravitational potential ϕ_I , which we will use to design a more efficient algorithm.

Instead of searching for the (combinatorial) assignment $i \rightarrow j = \pi(i)$, we consider now the following optimization problem, introduced by Kantorovich (1942):

$$\inf_{\gamma} \sum_i \sum_j \gamma_{ij} |\mathbf{x}_i - \mathbf{q}_j|^2 \quad \text{subject to} \quad (22)$$

$$\sum_i \gamma_{ij} = \mu_i \quad (23)$$

$$\sum_j \gamma_{ij} = \mu_j \quad (24)$$

$$\gamma_{ij} \geq 0 \quad \forall i, j. \quad (25)$$

The objective function (22) depends on an $[1, N]^2 \rightarrow \mathbb{R}^+$ array of coefficients γ_{ij} . Intuitively, each coefficient γ_{ij} indicates how much matter goes from \mathbf{q}_j to \mathbf{x}_i . In this setting, matter can split and merge between different particles, for instance, a particle \mathbf{q}_j can send half of its matter to particle \mathbf{x}_k and the other half to particle \mathbf{x}_i (using $\gamma_{jk} = \gamma_{ji} = 0.5$). Clearly, the mass of all the matter that gathers at a particle \mathbf{x}_i should sum as the mass μ_i of the particle (constraint (23)), and the mass of all the matter originated from a particle \mathbf{q}_j should sum as the mass μ_j of the particle (constraint (24)). Since no matter disappears, all coefficients γ_{ij} should be positive (constraint (25)). An array of coefficients γ_{ij} that satisfies the three constraints is called a *transport plan* (and an *optimal* transport plan if it minimizes (22)).

At first sight, it may seem to be a rather convoluted re-formulation of Monge’s problem, in particular, we now need to find N^2 unknowns, to be compared with the $N \rightarrow N$ permutation we had to find initially.

However, it can be observed that (22) is a *linear* optimization problem with *linear* constraints. Introduce $\psi \in \mathbb{R}^N$ and $\phi \in \mathbb{R}^N$ the Lagrange multipliers associated with constraints (23) and (24), respectively (note that we use the same notation ψ for the Lagrange multiplier of the constraint (24) and the gravitational potential, we elaborate on that in the next subsection). The dual of the optimization problem (22) writes (see the tutorials in Villani (2009), Santambrogio (2015), Lévy & Schwindt (2018), Mérigot & Thibert (2020), and the references herein):

$$\begin{aligned} & \sup_{\psi, \phi} \left[\sum_i \psi_i \mu_i + \sum_j \phi_j \mu_j \right] \\ & \text{subject to } \psi_i + \phi_j \leq 1/2 |\mathbf{x}_i - \mathbf{q}_j|^2 \quad \forall i, j. \end{aligned} \quad (26)$$

In addition, given a pair ψ, ϕ that satisfies the constraint (26), it is easy to check that replacing ϕ with ψ^c still satisfies the constraints while always increasing the objective function (26), where ψ^c is defined by:

$$\psi_j^c = \inf_i [1/2 |\mathbf{x}_i - \mathbf{q}_j|^2 - \psi_i]. \quad (27)$$

There exists several methods that exploit the structure of the problem (22) and its dual (26), we refer the reader to Peyré & Cuturi (2017) and Santambrogio (2015) for a survey. Among these methods, we mention the *entropic regularized* method, which solves:

$$\begin{aligned} & \inf_{\gamma} \left[\sum_i \sum_j \gamma_{ij} |\mathbf{x}_i - \mathbf{q}_j|^2 + \epsilon \sum_i \sum_j \gamma_{ij} \log(\gamma_{ij}) \right] \\ & \text{subject to (23) (24) (25)} \end{aligned} \quad (28)$$

where ϵ is a (small) regularization parameter. If both the \mathbf{q}_j 's and \mathbf{x}_i 's are supported by regular grids, it is possible to exploit the structure of (28) to design a fast and efficient algorithm (Cuturi 2013). The advantage of this algorithm is its speed and simplicity. The drawbacks are the need for re-sampling everything on regular grids and the difficulty of tuning the parameter ϵ (a too large value of ϵ results in a blurry, imprecise transport plan, and a too small value of ϵ makes the algorithm slow to converge). Various ways of leveraging its speed while overcoming its drawbacks, and in particular those based on entropic regularization (Benamou et al. 2015), are currently being studied (Benamou 2018).

In the next subsection, we describe a different method and although the algorithm that we eventually obtain is more complicated than those based on entropic regularized schemes, it does not depend a regularization parameter ϵ , and does not require ρ_F to be re-sampled on a regular grid.

3.3 Semidiscrete MAK reconstruction

Consider the discrete assignment problem expressed by (21). The density ρ_I is represented by a set of N particles (\mathbf{q}_j) $_{j=1}^N$ located on a regular grid (in blue in Fig. 1A). The distribution of mass ρ_F at current time is a set of N particles (\mathbf{x}_i) $_{i=1}^N$ (in red in the figure). Suppose one wishes to increase the precision by using a finer grid for the \mathbf{q}_j 's. For instance, in Fig. 1(B), four points \mathbf{q}_j are coupled to each point \mathbf{x}_i . Up to now, we supposed that there was the same number of points on both sides, but one may imagine that each \mathbf{x}_i is replaced by four points located at the same position with 1/4 the mass allocated to each of them. We can further refine the grid, as shown in Fig. 1(C), where 16 points \mathbf{q}_j are coupled to each point \mathbf{x}_i . Clearly, doing so will significantly increase computation time, but consider now that the number of \mathbf{q}_j particles tends to infinity, while




Figure 1. From discrete-discrete to semidiscrete transport. (A): coupling between the \mathbf{x}_i points (in red) and the \mathbf{q}_j points (in blue). (B): each red \mathbf{x}_i point is coupled with 4 \mathbf{q}_j points, each of them with 1/4 the mass. (C): each red \mathbf{x}_i point is coupled with 16 \mathbf{q}_j points. (D): at the limit, when the number of \mathbf{q}_j points tends to infinity, it can be proved that the optimal assignment couples each \mathbf{x}_i point with a polygonal area (called a Laguerre cell).

each particle's mass tends to zero accordingly. At the limit, ρ_I tends to the uniform density, while ρ_F is still supported by the set of points (\mathbf{x}_i).

The above setting corresponds exactly to our early universe reconstruction problem, where the initial density ρ_I is uniform, and the density at present time ρ_F corresponds to a set of galaxies, each of them represented by a single point \mathbf{x}_i . In this setting, as detailed below, it can be proved that each point \mathbf{x}_i is coupled to a polyhedral region (or polygonal in 2D, see Fig. 1D) that can be computed explicitly by an algorithm. Not only the so-computed result is more precise, but also the algorithm is much faster than the combinatorial one used in Fig. 1(A–C).

Next, we give more details about these polyhedral regions by first considering the dual problem in expression (26). The relation between the primal and dual problem that we wrote in the discrete-discrete setting remains valid in our continuous-discrete case (with integrals instead of discrete sums), and the dual problem writes:

$$\begin{aligned} & \sup_{\phi, \psi} \left[\int_V \phi(\mathbf{q}) \rho_I d^3 \mathbf{q} + \int_V \psi(\mathbf{x}) \rho_F(\mathbf{x}) d^3 \mathbf{x} \right] \\ & \text{subject to } \phi(\mathbf{q}) + \psi(\mathbf{x}) \leq 1/2 |\mathbf{x} - \mathbf{q}|^2 \quad \forall \mathbf{x}, \mathbf{q}. \end{aligned} \quad (29)$$

We recall that ρ_F is discrete, supported by the \mathbf{x}_i 's. Then the Lagrange multiplier ψ is determined by the vector $\psi_i = \psi(\mathbf{x}_i)$, and the dual problem becomes:

$$\begin{aligned} & \sup_{\phi, \psi} \left[\int_V \phi(\mathbf{q}) \rho_I d^3 \mathbf{q} + \sum_i \psi_i \mu_i \right] \\ & \text{subject to } \phi(\mathbf{q}) + \psi_i \leq 1/2 |\mathbf{x}_i - \mathbf{q}|^2 \quad \forall \mathbf{x}_i, \mathbf{q} \end{aligned} \quad (30)$$

The optimization problem depends on a vector of N values (ψ_i) and a function $\phi : V \rightarrow \mathbb{R}$. As we mentioned, replacement of ϕ by ψ^c defined in (27) always increases the objective function, thus we can now deduce $\phi = \psi^c$ from ψ and consider the following optimization problem:

$$\begin{aligned} & \sup_{\psi} \left[\int_V \inf_i [1/2|\mathbf{x}_i - \mathbf{q}|^2 - \psi_i] \rho I d^3 \mathbf{q} + \sum_i \psi_i \mu_i \right] \\ & \text{subject to } \psi^{cc} = \psi \end{aligned} \quad (31)$$

that solely depends on a vector ψ of N components. The constraint in (31) ensures that ψ is such that there can exist at least one pair ϕ, ψ that satisfies (26). It also means that the associated Kantorovich potentials Φ, Ψ (Section 2.4) are convex.

Consider now the integrand of (31). It is possible to partition V into the set of N regions (V_i^ψ), defined according to the index i that realizes the infimum in the integrand of (31). This lets rearrange the integral as follows:

$$\begin{aligned} & \sup_{\psi} \left[\sum_i \int_{V_i^\psi} 1/2 [|\mathbf{x}_i - \mathbf{q}|^2 - \psi_i] \rho I d^3 \mathbf{q} + \sum_i \psi_i \mu_i \right] \\ & \text{subject to } V_i^\psi \neq \emptyset \quad \forall i \end{aligned} \quad (32)$$

where

$$V_i^\psi = \{ \mathbf{q} \mid 1/2|\mathbf{x}_i - \mathbf{q}|^2 - \psi_i < 1/2|\mathbf{x}_j - \mathbf{q}|^2 - \psi_j \quad \forall j \neq i \}. \quad (33)$$

The so-defined partition of V into the regions V_i^ψ is called a *Laguerre diagram* (or power diagram in our specific case). An individual region is called a *Laguerre cell*. Laguerre diagram is a generalization of *Voronoi diagram*, parametrized by the vector $\psi \in \mathbb{R}^N$. If $\psi_i = 0$ for all i , then the Laguerre diagram is a Voronoi diagram. The use of Voronoi diagrams in cosmology was pioneered by van de Weygaert in the 90s to analyse the distribution of matter in the Universe (van de Weygaert 1994). Another famous example of Voronoi diagrams used in cosmology is the ALE⁸ code developed by White & Springel (1999).

To summarize the derivations above, the Laguerre diagram appears as the combinatorial structure arising from extremizing the simplified action. This is the main reason for using this geometrical structure for our algorithm. Besides that, one may ask why Voronoi diagrams do not suffice in our case. We give some additional intuition here: a Voronoi diagram can be composed of cells with arbitrary volumes, depending on the position of the points. For instance, in zones where points are clustered, Voronoi cells are smaller. In our setting, we want to *constrain* the volumes of the cells: they should all be the same. Cell volumes can be (indirectly) controlled through the additional vector ψ , which corresponds to the gravitational potential and that is obtained by solving the semidiscrete Monge–Ampère equation. Besides the need of controlling the volumes, we need to overcome another constraint Voronoi diagrams have in this context: a Voronoi cell always contains the point it is associated with, whereas this constraint does not exist with Laguerre diagrams. Physically, it means that the matter that ended up at a given point may have come from an arbitrary location in space that may be further away than the immediate (Voronoi) neighbourhood of the point.

Examples of Laguerre diagrams in 2D and 3D are shown in Fig. 2. The Laguerre diagram is completely defined by the points $(\mathbf{x}_i)_{i=1}^N$ and the vector $\psi \in \mathbb{R}^N$ of coefficients. Note that in contrast to a Voronoi




Figure 2. (A): the Laguerre cell V_i^ψ associated with the point \mathbf{x}_i does not necessarily contain \mathbf{x}_i . (B): a 2D Laguerre diagram with 1500 cells. (C and D): a 3D Laguerre diagram with 1500 cells and a cross-section (cells are slightly shrunk to ease visualization).

cell, depending on the vector ψ , the Laguerre cell V_i^ψ associated with a point \mathbf{x}_i does not necessarily contain \mathbf{x}_i (see Fig. 2A)). It is even possible for a cell V_i^ψ to be empty. The examples shown in Fig. 2 are solutions of the optimal transport problem, thus all cells have the same area/volume. They have different shapes though, this is because the MA equation is non-linear, with potentially a highly anisotropic solution. This anisotropy influences the shapes of the Laguerre cells.

In terms of the Laguerre diagram, the convexity constraint (31) means that no Laguerre cell is empty (32). The notion of Laguerre diagram is well known and well studied in computational geometry (Aurenhammer 1987), this equips us with efficient computational tools to solve the optimization problem (32), as explained in the next section.

The solution to the optimization problem (32) results in a vector ψ of N coefficients, from which one can subsequently deduce the gravitational potential ϕI using the relation $\phi I = \psi^c$, or:

$$\begin{aligned} \phi I(\mathbf{q}) &= \inf_i [1/2|\mathbf{q} - \mathbf{x}_i|^2 - \psi_i] \\ &= 1/2|\mathbf{q} - \mathbf{x}_{i(\mathbf{q})}|^2 - \psi_{i(\mathbf{q})}, \end{aligned} \quad (34)$$

where $i(\mathbf{q})$ is the index of the Laguerre cell V_i that contains \mathbf{q} . Note that all the \mathbf{q} 's located in the Laguerre cell V_i^ψ are mapped to \mathbf{x}_i through $\mathbf{x}_F(\cdot)$:

$$\begin{aligned} \mathbf{x}_F(\mathbf{q}) &= \mathbf{q} - \nabla_{\mathbf{q}} \phi I(\mathbf{q}) \\ &= \mathbf{q} - \nabla_{\mathbf{q}} (1/2|\mathbf{q} - \mathbf{x}_{i(\mathbf{q})}|^2 - \psi_{i(\mathbf{q})}) \\ &= \mathbf{x}_{i(\mathbf{q})}. \end{aligned} \quad (35)$$

To summarize, the solution of the optimization problem (32) gives a vector (ψ_i) of N coefficients. These coefficients define a partition of V into N Laguerre cells (V_i^ψ). Each Laguerre cell V_i^ψ corresponds

⁸Arbitrary Lagrangian Eulerian.

to the (continuous) set of points \mathbf{q} at the IC that collapses into a given point \mathbf{x}_i at current time. In other words, the Laguerre cell V_i^ψ corresponds to the pre-image of \mathbf{x}_i through \mathbf{x}_F .

4 NUMERICAL SOLUTION MECHANISM

Let us denote by $K(\psi) : \mathbb{R}^N \rightarrow \mathbb{R}$ the objective function of the optimization problem (32):

$$K(\psi) = \sum_i \int_{V_i^\psi} [1/2|\mathbf{x}_i - \mathbf{q}|^2 - \psi_i \rho I] d^3 \mathbf{q} + \sum_i \psi_i \mu_i \quad (36)$$

It can be shown that $K(\cdot)$ is a concave C^2 function, which suggests that it can be efficiently maximized by a Newton algorithm (Aurenhammer, Hoffmann & Aronov 1992; Lévy 2015; Kitagawa, Mérigot & Thibert 2016). The Newton algorithm iteratively maximizes second-order approximations of $K(\cdot)$ until it converges to the (unique) maximum. A 2D example of the Laguerre diagrams corresponding to each iteration is shown in Fig. 3. The algorithm starts with $\psi = 0$, then updates ψ by solving a series of linear system. In the end, the algorithm finds the unique solution, and all the Laguerre cells have the prescribed volumes.

The algorithm has two main components, a *numerical* component that computes the second-order estimates of K (Section 4.1), and a *geometrical* component that computes the associated series of Laguerre diagram (Section 4.2). The output of the algorithm is a Lagrangian representation. To ease inter-operability with classical tools, such as NBODYKIT, we describe a method to convert the output of our algorithm into Eulerian grids (Section 4.3).

4.1 Numerical aspects

Newton's algorithm for semidiscrete optimal transport can be summarized as follows:

- Input:** – the set of N points $\mathbf{x}_i \in V = [0, 1]^3$
– the masses μ_i such that $\sum_i \mu_i = 1$
(for instance, $\mu_i = 1/N$)
- Output:** – the (unique) vector $\psi \in \mathbb{R}^N$ that maximizes $K(\cdot)$
– the Laguerre cells $(V_i^\psi)_{i=1}^N$ that give for each \mathbf{x}_i
the (continuous) set of \mathbf{q} points mapped to \mathbf{x}_i .

- (1): $\psi \leftarrow [0 \dots 0]$
- (2): Loop
- (3): Compute the Laguerre diagram $(V_i^\psi)_{i=1}^N$
- (4): Compute the gradient $\nabla K(\psi)$
- (5): If $\|\nabla K(\psi)\|_\infty < \epsilon$ then Exit loop
- (6): Compute the Hessian matrix $\nabla^2 K(\psi)$
- (7): Solve for $\mathbf{p} \in \mathbb{R}^N$ in $\nabla^2 K(\psi)\mathbf{p} = -\nabla K(\psi)$
- (8): Find the descent parameter α
- (9): $\psi \leftarrow \psi + \alpha \mathbf{p}$
- (10): End loop

Fig. 3 shows some iterations of the Newton algorithm: the vector ψ is optimized until each Laguerre cell V_i^ψ has the prescribed area $\mu_i = 1/N$.

The algorithm above needs to compute multiple evaluations of the gradient and Hessian matrix of $K(\cdot)$. The coefficients of the gradient and Hessian matrix can be deduced from the Laguerre diagram $(V_i^\psi)_{i=1}^N$ that is computed at step (3). The associated algorithm is detailed later in the next subsection on the geometric aspects. Once the Laguerre diagram is computed, the coefficients $\partial K / \partial \psi_i$ of the gradient $\nabla K(\cdot)$ are given by the following expression (Kitagawa et al.

(2016), Lévy & Schwindt (2018)):

$$\frac{\partial K}{\partial \psi_i} = \mu_i - \int_{V_i^\psi} \rho I d^3 \mathbf{q}. \quad (37)$$

In other words, this corresponds to the mass μ_i associated with a point \mathbf{x}_i at present time, minus the mass of the matter that was transported there through the assignment map $\mathbf{x}_F(\cdot)$ (remember that the region transported to \mathbf{x}_i corresponds to V_i^ψ). For the vector ψ that maximizes $K(\cdot)$, all components of the gradient vanish, which means that each Laguerre cell V_i^ψ has exactly the prescribed mass μ_i . Since $\rho I(\cdot)$ is uniform, the integrated density over V_i^ψ simply corresponds to the volume of V_i^ψ (but the formula above is valid for an arbitrary $\rho I(\cdot)$ density).

This expression of the gradient leads also to a natural stopping criterion (line 5), the largest component of the gradient corresponds to the maximum error of transported mass. We stop the algorithm as soon as it is smaller than a prescribed ϵ (typically one per cent of μ_i that is, $\epsilon = 0.01/N$).

We now consider the Hessian matrix computed at step (6). Still following (Kitagawa et al. 2016; Lévy & Schwindt 2018), its coefficients are given by:

$$\begin{aligned} \frac{\partial^2 K}{\partial \psi_i \partial \psi_j} &= \frac{1}{2\|\mathbf{x}_j - \mathbf{x}_i\|} \int_{V_{ij}^\psi} \rho I d^3 \mathbf{q} \quad \text{if } i \neq j \\ \frac{\partial^2 K}{\partial \psi_i^2} &= - \sum_{j \neq i} \frac{\partial^2 K}{\partial \psi_i \partial \psi_j} \end{aligned} \quad (38)$$

where V_{ij}^ψ denotes the polygonal facet that is common to the Laguerre cell V_i and V_j . Note that the Hessian matrix is sparse, and has a non-zero entry at coefficient (i, j) if and only if the Laguerre cells V_i^ψ and V_j^ψ touch each other along a common facet.

Remark: the Hessian matrix coincides with the Finite Element \mathbb{P}_1 Laplacian. It can be explained as follows: the MA equation can be considered as a non-linear generalization of the Poisson equation, and its second-order expansion in the Newton algorithm naturally corresponds to the Laplacian.

Step (7) of the algorithm computes the Newton step vector \mathbf{p} , by solving a linear system. This linear system is typical of a Poisson equation discretized with Finite Elements, and can be solved using classical methods: we use the Conjugate Gradient algorithm (Hestenes & Stiefel 1952) pre-conditioned by Jacobi. The (sparse) Hessian matrix is stored using the Compressed Row Storage format, constructed and assembled using a specialized and highly optimized algorithm, interfaced with the Laguerre diagram. To tune the stopping criterion of the Conjugate Gradient algorithm, we used as a ‘ground truth’ a direct solver (SuperLU) on small point sets (thousands of points). For solving the linear system $H\mathbf{p} = -\mathbf{g}$, we found that the stopping criterion $|H\mathbf{p} - \mathbf{g}|/|\mathbf{g}| \leq 10^{-3}$ results in nearly the same step vector \mathbf{p} as with the direct solver.

We implemented two versions of the linear solver, a multithreaded CPU version and a GPU version. On a high-end GPU (Nvidia V100), the algorithm is typically 45 times faster (90 GFlops) than the multithreaded CPU version (2 GFlops).

Once the step vector \mathbf{p} is computed, we need to find a good descent parameter α . The KMT algorithm (Kitagawa et al. 2016), provably




Figure 3. Newton’s algorithm applied to a 2D problem with 10 000 points. Iterations 0, 4, 8, and 12. At iteration 12, all cells have the same area (largest cell area error is smaller than 1 per cent). The diagram is the unique optimal transport solution.

convergent, works as follows:

- (1) : $\alpha \leftarrow 1$
- (2) : Loop
- (3) : If $\inf_i |V_i^{\psi+\alpha\mathbf{p}}| > a_0$
- (4) : and $|\nabla K(\psi + \alpha\mathbf{p})| \leq (1 - \alpha/2)|\nabla K(\psi)|$
- (5) : then Exit loop
- (6) : $\alpha \leftarrow \alpha/2$
- (7) : Compute Laguerre diagram $(V_i^{\psi+\alpha\mathbf{p}})_{i=1}^N$
- (8) : End loop

where $a_0 = \frac{1}{2} \min(\inf_i |V_i^0|, \inf_i(\mu_i))$ and where $|V_i^\psi|$ denotes the volume of a Laguerre cell.

The KMT algorithm iteratively halves the descent parameter α until two criteria are met: the volume of the smallest Laguerre cell needs to be larger than a threshold a_0 (line 3), and the norm of the gradient needs to decrease sufficiently (line 4). The threshold a_0 for the minimum Laguerre cell volume corresponds to (half) the minimum Laguerre cell volume for $\psi = 0$ (also called Voronoi diagram) and minimum prescribed area μ_i (in our case $1/N$).

Equipped with the KMT algorithm above, we can now compute the descent parameter α , by plugging the algorithm above into line (8) of the Newton algorithm at the beginning of this section.

The only thing we need to explain now is how to compute a Laguerre diagram.

4.2 Geometrical aspects

To compute the Laguerre diagram, we use the classical algorithm developed simultaneously by Bowyer and Watson (Bowyer 1981;




Figure 4. Internally, the Laguerre diagram is represented by its dual structure, called the regular triangulation (in grey).

Watson 1981; initially for Voronoi diagrams). We do not completely detail this algorithm, but give the general idea below. Then we mention some specificities of our implementation.

4.2.1 Bowyer-Watson and the dual triangulation

While it would be possible to directly represent the polygons/polyhedra of the Laguerre diagram, it would be costly, because each polygon/polyhedron can have a different number of vertices. The Bowyer–Watson algorithm uses the dual triangulation instead, displayed in grey in Fig. 4: each Laguerre vertex is shared by three Laguerre cells. Thus, what is represented in the computer is the set of (i, j, k) indices triplets such that the Laguerre cells V_i^ψ , V_j^ψ , and V_k^ψ have a common vertex. This forms a triangulation of the point set




Figure 5. (A): the Hilbert curve. (B): a 2D Hilbert-ordered point set. (C): a 3D Hilbert-ordered point set. (D): adaptive Hilbert ordering using the median, resulting in a balanced Hilbert curve, adapted to a heterogeneous point set.

$(\mathbf{x}_i)_{i=1}^N$, known as the *regular triangulation* (Aurenhammer 1987). In 3D, each Laguerre vertex is shared by four Laguerre cells, and the dual structure is made of (i, j, k, l) tetrahedra (instead of triangles in 2D).

This triangulation is constructed by inserting the points \mathbf{x}_i one by one. Each time a point \mathbf{x}_i is inserted, the triangles/tetrahedra that correspond to the Laguerre vertices that fall inside the cell V_i^ψ of \mathbf{x}_i are discarded, and the triangles that correspond to the vertices of V_i^ψ are created. The Boywer–Watson algorithm uses the fact that the set of triangles to be discarded is connected, and comprises the triangle that contains \mathbf{x}_i . This remark makes it possible to speed-up the algorithm: starting from the triangle that contains \mathbf{x}_i , found by navigating the triangulation, a greedy algorithm traverses the set of triangles to be discarded. This dramatically reduces the number of triangles to be tested.

4.2.2 Spatial sorting

At this point, the execution time is dominated by finding the triangle/tetrahedron that contains each point \mathbf{x}_i . In 3D, starting from a random tetrahedron, the algorithm needs to traverse an average of $\sqrt[3]{N}$ tetrahedra to find the one that contains \mathbf{x}_i .

The algorithm is made significantly faster by sorting the vertices spatially (Amenta, Choi & Rote 2003; Alauzet & Loseille 2009), along the Hilbert curve (see Fig. 5A). Sorting the vertices this way ensures that two points near to each other in 3D are mapped to close indices. Hilbert sorting is classical in high performance large-scale cosmological simulation, for instance, it is a key component of the code used in the DEUS project (Reverdy et al. 2015). Fig. 5B and C shows what the computed order looks like for a homogeneous point distribution. In our case, the distribution of points can be highly heterogeneous, with a large number of points clustered in some

zones. As suggested in Delage & Devillers (2004), to make the ordering well adapted to the point distribution, we use the median of the points coordinates to hierarchically subdivide the domain, see Fig. 5(D) for a 2D example.

Each time a new point is inserted, the tetrahedron that contains it is searched by navigating the triangulation starting from a tetrahedron incident to the previously inserted point. Since points with consecutive indices are near to each other in 3D, this considerably reduces the number of traversed tetrahedra (from $\sqrt[3]{N}$ to typically 10–20).

Spatial sorting not only accelerates the computation of the Laguerre diagram, but also it speeds-up the iterative conjugate gradient solver: since it maps neighbouring points to as-contiguous-as-possible locations in memory, it significantly improves cache locality. Without spatial sorting, on the GPU, we obtain 70 GFlops without it and 90 GFlops with it. On the CPU, we obtain 1.5 GFlops without it and 2 GFlops with it.

4.2.3 Numerical precision and geometric predicates

To determine which tetrahedron needs to be created or discarded, the algorithm needs to take combinatorial decisions based on the relative locations of some geometric elements. So the algorithm depends on a limited number of functions, called *geometric predicates*. A predicate is a function that takes as arguments a set of points $\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k, \dots$ and $\psi_i, \psi_j, \psi_k, \dots$ coefficients, and that returns a discrete value $-1, 0$, or $+1$. For computing a Laguerre diagram, we need two geometric predicates:

(i) *orient*, which indicates whether the three vectors $(\mathbf{x}_j - \mathbf{x}_i, \mathbf{x}_k - \mathbf{x}_i, \text{ and } \mathbf{x}_l - \mathbf{x}_i)$ form a direct ($+1$), degenerate (0), or indirect (-1) basis. It is used to navigate the triangulation and find the tetrahedron that contains \mathbf{x}_i :

$$\text{orient}(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k, \mathbf{x}_l) = \text{sign} \begin{vmatrix} x_j - x_i & y_j - y_i & z_j - z_i \\ x_k - x_i & y_k - y_i & z_k - z_i \\ x_l - x_i & y_l - y_i & z_l - z_i \end{vmatrix};$$

(ii) *conflict*, which indicates whether the Laguerre vertex that corresponds to the tetrahedron $(j, k, l, \text{ and } m)$ falls inside the Laguerre cell V_i^ψ of \mathbf{x}_i ($+1$), on its boundary (0) or outside (-1). It is used to determine which tetrahedra need to be discarded when inserting \mathbf{x}_i in the diagram:

$$\text{conflict}(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k, \mathbf{x}_l, \mathbf{x}_m, \psi_i, \psi_j, \psi_k, \psi_l, \psi_m) = \text{sign} \begin{vmatrix} x_j - x_i & y_j - y_i & z_j - z_i & h_j - h_i \\ x_k - x_i & y_k - y_i & z_k - z_i & h_k - h_i \\ x_l - x_i & y_l - y_i & z_l - z_i & h_l - h_i \\ x_m - x_i & y_m - y_i & z_m - z_i & h_m - h_i \end{vmatrix},$$

$$\text{where } h_i = x_i^2 + y_i^2 + z_i^2 - \psi_i \text{ (resp. } j, k, l, m).$$

The two predicates *orient* and *conflict* correspond to the sign of polynomials of the points coordinates and coefficients of ψ . It is of crucial importance that these signs are coherent: for instance, if at one moment the algorithm considers that point \mathbf{x}_i is strictly above point \mathbf{x}_j , the algorithm should not consider later that \mathbf{x}_j is strictly above \mathbf{x}_i , else it will create a triangulation that is not coherent. Since floating point numbers have a limited precision, avoiding this type of inconsistencies requires special care. It is especially true in our case, since we are computing a large number of Laguerre diagrams (typically tenths) with a huge number of points (typically tenths millions). In the 2000s, it was a major obstacle to the early development of cosmological codes with the semidiscrete setting in 3D. To ensure that the combinatorial decisions taken by the

algorithm are coherent, we developed the Predicate Construction Kit programming language (Lévy 2016), which transforms the formula of a predicate into a function that evaluates the exact sign. We used it to implement `orient`, `conflict`, and other specialized predicates (Yan et al. 2010) involved in the periodic boundary condition (next paragraph). Internally, we use exact expansion-based arithmetic (Shewchuk 1997) that represent each number by an array of double-precision floating point numbers. To speed-up computations, we also use arithmetic filters (Meyer & Pion 2008) that quickly determine the signs in the easy cases and avoid costly expansion-arithmetic computations in most cases. Finally, we use symbolic perturbation (Edelsbrunner & Mücke 1990) to ensure that the decisions remain coherent even in degenerate configurations.

4.2.4 Periodic boundary conditions

Remember that our computational domain V is the unit cube $[0, 1]^3$ with periodic boundary conditions. With earlier discrete-discrete methods, like the ‘auctions’ algorithm, it is easy to take this into account: each time the squared distance $|\mathbf{x}_i - \mathbf{x}_j|^2$ needs to be computed, it is replaced with:

$$\min_{k=1}^{27} |\mathbf{x}_i - T_k(\mathbf{x}_j)|^2,$$

where T_k denotes one of the 27 possible translations obtained using $-1, 0, 1$ coordinates.

In our case, the situation is more complicated, because we need to compute the (continuous) Laguerre diagram. A trivial solution consists in copying the points 27 times (and keep only the centre part with zero translation). Clearly this would dramatically increase the computation time. What we do instead is first computing the Laguerre diagram of the points, then determining which cell intersects the boundary of the $[0, 1]^3$ cube, and copy these points with the right translations, depending on which face, edge, or vertex of the cube was intersected (Yan et al. 2011). Note that the (back-translated) neighbours of these translated points need to be inserted as well. This typically concerns 5–10 per cent of the points (1.1 times, to be compared with 27 times). Fig. 6(A) shows an example of a 2D periodic Laguerre diagram. The diagram paves the 2D space (Fig. 6B). Fig. 6 shows a 3D example (that paves the 3D space).

For the sake of completeness, we also mention the alternative approach in (Caroli & Teillaud 2009) for computing periodic Laguerre diagrams that is more elegant theoretically. It iteratively inserts the vertices, starting with 27 copies, then it switches to a periodic triangulation where each vertex and tetrahedron is only represented once, with combinatorics that represent the periodic space, as soon as a criterion on the points location and ψ_i 's is respected. However, while their approach works well in practice for Voronoi diagrams (with $\psi = 0$), their criterion is never met in our case, resulting in 27 copies, because the ψ_i coefficients vary too much. Moreover, even with a single copy, the memory consumption of the data structures they use (CGAL library, pointer-based) makes it not practical for large-scale cosmological simulations.

4.2.5 Multicore

The Bowyer–Watson algorithm is not well suited to multicore parallelization, because it inserts the points one by one in a way that globally updates the diagram under construction. However, still using spatial sorting, it is possible to split the point set into batches that are geometrically well separated and unlikely to interact. Each batch is assigned to a different thread. Then all the threads insert




Figure 6. Laguerre diagrams with periodic boundary conditions in 2D and in 3D.

their batches of points into the diagram simultaneously. We use light weight synchronization primitives (spinlocks) to detect whenever two threads try to modify the same tetrahedron. Such conflicts are resolved by rolling back the involved modifications of the diagram, and redoing them in sequential mode. In our computations, for a 80 million points diagram, typically a few tenth of conflicts are encountered (performance penalty is negligible).

4.2.6 Performance of Laguerre diagram computation

We measured the time used by our algorithm to compute a single Laguerre diagram. Statistics for point sets of different sizes are reported in the table below (as computed on an Intel Xeon 5122 3.6 GHz CPU):

| nb points | 1M | 2M | 4M | 8M | 16M | 85M |
|-----------|-------|-------|------|------|------|-------|
| time | 2.8 s | 5.4 s | 10 s | 20 s | 44 s | 451 s |

4.2.7 Performance of SDMAK reconstruction

As can be seen in Fig. 7, the computation time indeed scales as $\mathcal{O}(N \log N)$, massively outperforming previous approaches. To provide a realistic setting within which we aim to use this algorithm, we employ snapshots from the cosmological N -body simulation suite ABACUSCOSMOS (Garrison et al. 2018; see next section). The convergence time of the reconstruction among others slightly depends on the degree to which non-linear clustering has occurred in the samples. We therefore run the complexity analysis on snapshots of different redshifts, where per sample size N , the particles are kept




Figure 7. Empirical computational complexity of the semidiscrete algorithm, running on an Intel Xeon 5122 3.6 GHz with 128 Gb RAM and an Nvidia V100 GPU with 16 Gb RAM. Tests were done on ABACUSCOSMOS simulations. The reconstructions were performed on particle samples of different sizes at redshifts of $z_s = 0.3, 0.7,$ and 1.5 , resulting in different computation times (grey dots). Due to increase of non-linearities towards lower redshifts, the reconstructions generally perform faster at higher redshifts. This is shown by the data points, and emphasized by analytical fits to those points (red lines).

across redshifts, and find that the computation time increases with decreasing redshift.

A timing breakdown for a run on a 2M point data set is reported in the table below:

| | <i>Total</i> | <i>Lag.</i> | <i>LinSolve</i> | <i>grad.</i> | <i>Hess.</i> | <i>Other</i> |
|----------|--------------|---------------|-----------------|--------------|--------------|--------------|
| time | 453 s | 147 s | 255 s | 22 s | 27 s | 2 s |
| per cent | 100 per cent | 32.5 per cent | 56 per cent | 4.8 per cent | 5.9 per cent | 0.8 per cent |

In this example, the algorithm converged in 16 iterations. Timing is dominated by the linear solve phase (56 per cent of total time). The other costly phase is computing Laguerre diagrams (32.5 per cent). The rest of the timing is mostly spent in computing the gradient and assembling the Hessian of the objective function (less than 10 per cent for both).

We now report timings for the same numerical experiment, but this time running the linear solver on the GPU (Quadro T2000):

| | <i>Total</i> | <i>Lag.</i> | <i>LinSolve</i> | <i>grad.</i> | <i>Hess.</i> | <i>Other</i> |
|----------|--------------|-------------|-----------------|--------------|--------------|--------------|
| time | 261 s | 140 s | 73 s | 21 s | 25 s | 2 s |
| per cent | 100 per cent | 53 per cent | 28 per cent | 8.2 per cent | 10 per cent | 0.8 per cent |

Using our GPU version of the linear solver significantly reduces execution time. It will be possible to further reduce execution time by also computing the Laguerre diagram on the GPU (Ray et al. 2018). However, note that efficiently using the GPU is only possible for problems that fit in GPU memory (typically up to a few tens of millions of points).

4.3 Lagrangian to Eulerian conversion

The result of our algorithm is a Lagrangian representation, that is, Laguerre cells linearly interpolated to the points \mathbf{x}_i by equation (6). From this representation, it is possible to generate an Eulerian grid. Intuitively, this Lagrangian-to-Eulerian conversion means ‘painting’ the Laguerre cells into the regular grid (see Fig. 8), which can be done at any redshift z_s by moving the vertices of the Laguerre cell using equation (39) page 30. All the cells of the Laguerre diagram are subsequently ‘painted’ on to the grid in parallel, by determining the regular grid cells contained by each Laguerre cell. Along the boundary of the shrunken Laguerre cell, we compute the intersection volume between the Laguerre cell and the grid cells. Since both objects (Laguerre cell and grid cell) are convex, their intersection can

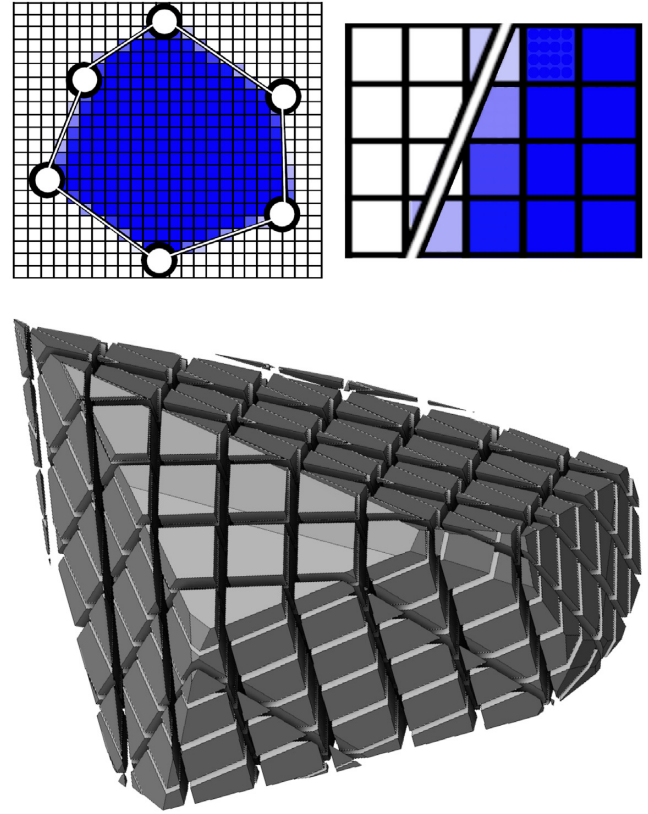


Figure 8. Converting from Lagrangian (Laguerre cell) to Eulerian (density grid) by measuring the area of the intersection between the Laguerre cell and the grid cells, in 2D (top) and in 3D (bottom).

be easily computed using a dual representation, see Aurenhammer (1987) and references herein for more details.

This Lagrangian-to-Eulerian conversion typically takes 30 minutes for converting a 16 million cells Laguerre diagram into a 512^3 Eulerian density grid.

5 TESTS OF THE SEMIDISCRETE ALGORITHM WITH COSMOLOGICAL SIMULATIONS

Finally, we put our algorithm to the proof. We employ a set of cosmological N -body simulations, in specific, 10 simulations from the ABACUSCOSMOS suite (Garrison et al. 2018, 2019).⁹ We present results of both qualitative and quantitative measures, which capture the basic capability of our code, and further highlight its special features.

The simulations: The main field of application for reconstruction algorithms is the recovery of the linearly perturbed density field, e.g. for improving the precision with which BAO can be measured. At the same time, we aim to test our algorithm on smaller scales, where high resolution simulations are required. We select ten ‘AbacusCosmos_1100box_planck’ simulations from the ABACUSCOSMOS suite, which are highly resolved large-scale simulations for Λ CDM cosmologies with parameters fixed to those of Ade et al. (2016). Each of the ten simulations has a box size of $(1100 h^{-1} \text{Mpc})^3$ and, for most of what follows, we sample 256^3 of the 1440^3 particles with

⁹<https://garrison.github.io/AbacusCosmos/>

which each simulation was run. For later visualization and for the computation of power spectra, we paint the particle positions on to a mesh of size 512^3 by use of the PYTHON package NBODYKIT¹⁰ (Hand et al. 2018b).

While snapshots of the simulations are provided at a range of low redshifts, namely $z = 0.3, 0.5, 0.7, 1.0,$ and 1.5 , we additionally generate the density field corresponding to their ICs at $z = 49.0$ via ZELDOVICHPLT¹¹ (Garrison et al. 2016). Also the density field is cast on to a mesh of size 512^3 .

The reconstructions: We begin by computing the Laguerre diagram of a simulation’s snapshot at a given redshift, z_s , as described in Section 4. This assigns to each point \mathbf{x}_i the Laguerre cell V_i^ψ . Recall that each Laguerre cell V_i^ψ represents the set of mass elements at initial positions \mathbf{q} in Lagrangian coordinates arriving at a given point \mathbf{x}_i . The first-order Lagrangian approximation for the motion of each mass element as a function of the redshift z_f , i.e. the Zel’dovich approximation, is given by equation (5), which we rephrase here for convenience:

$$\mathbf{x}(z_f) = \mathbf{q} + \frac{D(z_f)}{D(z_s)}(\mathbf{x}(z_s) - \mathbf{q}), \quad (39)$$

where $D(z)$ are the linear growth factors at redshifts z , which in our case are taken from the ABACUSCOSMOS simulations themselves.¹²

To this effect, each cell V_i^ψ at the IC is shrunk towards a single point \mathbf{x}_i at the current time by interpolation, such that each mass element \mathbf{q} undergoes the motion governed by equation (39). In order to analyse the reconstructed density field (e.g. by computing its correlation function or power spectrum), we need to compute the Fourier transform of the corresponding density field. While it is possible to compute the Fourier transform of a Laguerre diagram (Wuttke 2017), it is computationally very expensive. Hence, we simply convert our Lagrangian representation (Laguerre cells) into an Eulerian one (regular grid), as explained in Appendix 4.3. Once the density is represented in Euler form, we can use standard tools based on the Fast Fourier Transform (FFT).

In the following subsections we mainly focus on the computation of the initial positions of 256^3 particles from two of the available redshifts – the lowest available redshift, $z_s = 0.3$, and the highest available, $z_s = 1.5$ – that can be seen as representative for low- and high-redshift samples of present galaxy surveys. We estimate the algorithm’s accuracy by reconstructing the density at $z_f = 49$ in order to compare with the linear density corresponding to each of the simulations. We present a non-exhaustive range of tests concerning the accuracy of the reconstruction algorithm as well as the ability to extract cosmological information from the obtained reconstructions.

5.1 Qualitative diagnostics

5.1.1 Density slices

First, we present a purely visual comparison between the reconstructed and true density at $z_f = 49$ by the example of a single simulation. Fig. 9 shows the density contrast as computed on a 512^3 -cell mesh averaged over a slice roughly of dimensions $(500 \times 500 \times 10)h^{-1}\text{Mpc}$, as well as a $(150 \times 150 \times 10)h^{-1}\text{Mpc}$ zoom-in to

highlight details on comparably smaller scales. Each slice was further smoothed with a $S = 2h^{-1}\text{Mpc}$ Gaussian filter¹³ to remove noise on the smallest scales. The left-hand panels show the slice in the original, $z_s = 0.3$, snapshot that exhibits strongly pronounced over-dense regions. The centre panels show the same slices in the $z_f = 49$ IC and the reconstructed density contrast, respectively, and are nearly indistinguishable. The right-hand panels highlight differences between the initial and the reconstructed density contrast. All slices have been scaled with the linear growth factor to match the amplitudes at $z_f = 0.3$. This rather qualitative comparison already anticipates corresponding agreement between more quantitative measures that are presented in the following paragraphs.

5.1.2 1-point distributions

As a second qualitative diagnostic, we show the one-point distribution function of the smoothed densities in all three, original, initial, and reconstruction, see Fig. 10. Note that again the density contrasts were scaled linearly to match the growth at $z_f = 0.3$. The left-hand panel, $z_s = 0.3$, highlights the known skewed distribution of the density contrast in the non-linearly evolved universe with highly collapsed overdensities among a mostly underdense matter distribution. This stands in contrast to the (by construction) Gaussian distribution seen at $z_f = 49$. As is intuitive, our reconstruction leads to an almost Gaussian density contrast, which exhibits only weak skewness due to residuals of the very non-linear structures at the input re-shift. The idea to morph the skewed density contrast to become more Gaussian has indeed inspired some of the first variants on reconstruction methods, so-called Gaussianization techniques (Weinberg 1992). Comparisons of one-point distributions of true and reconstructed density contrasts can also be found in e.g. Schmittfull et al. (2017).

5.2 Cosmological quantities

We here characterize the power spectra and correlation functions of our reconstructions by comparing them with their linear expectations, both individually and after averaging. We account for various noise contributions, such as sample variances, shot noise, and conventional broad-band power contributions, and finally demonstrate the reconstructions’ excellent agreement with their expectation.

5.2.1 Power spectra

Even though the ultimate goal is to obtain a good estimate of the true power spectrum, which underlies any particle sample, ‘cosmic sample’ variance – arising from selecting a particular box of the simulated universe – is an unyielding obstacle inherent to any such endeavour. However, for the sake of inspecting only the power induced (or deduced) by our reconstruction algorithm, we here compare the power spectra of simulation and reconstruction for each simulation, respectively, thereby artificially circumventing cosmological sample variance. This approach might prove helpful in future work, when biases potentially present in our reconstruction, should be removed from reconstructions of real data. A second source of sample variance – in the following referred-to as ‘subsample’ variance – appears when sampling the density field with a finite number of particles. To avoid the false impression of not recovering the expected power on large scales, we here create a set of five

¹³The smoothing filter used here is defined as $W_S(k) = \exp[-k^2 S^2]$.

¹⁰<https://github.com/bccp/nbodykit>

¹¹<https://github.com/lgarrison/zeldovich-PLT>

¹²In practice, when the exact growth factors are not known, it suffices to use approximate values for a fiducial cosmology, for e.g. Lukic et al. (2007), and to relate the resulting amplitudes of the reconstructed linear densities to the correct amplitudes by a constant bias, cf. Section 5.2.




Figure 9. Slices of snapshots and reconstruction of an ABACUSCOSMOS simulation. The panels show the same $(500 \times 500 \times 10) h^{-3} \text{Mpc}^3$ slice of the density contrast (top panels) as well as a $(150 \times 150 \times 10) h^{-3} \text{Mpc}^3$ zoom-in (bottom panels, marked by white dashed lines), smoothed with a $2 h^{-1} \text{Mpc}$ Gaussian kernel to remove small-scale noise in the reconstruction. All slices have been scaled to match the linear growth amplitude at $z = 0.3$.




Figure 10. Histograms of the density contrast δ corresponding to the left three images depicted in Fig. 9.

subsamples per simulation, each with the same number of particles, yet selected with different random seeds. For each simulation i and subsample j , we compute the power spectrum $P_{ij}^{\text{rec}}(k)$ of the sample reconstructed to redshift $z_f = 49$ and compare it with the initial power spectrum $P_i(k)$ at the same redshift as follows.

$$\delta P_{ij}^{\text{rec}}(k) := \frac{P_{ij}^{\text{rec}}(k)}{P_i(k)} - 1 \quad (40)$$

This quantity is then averaged over all i and all j . We show this average and its standard deviation in the bottom panels of Fig. 11 below the average power spectra for both ICs and reconstructions from $z_s = 0.3$ (left-hand panel) and $z_f = 1.5$ (right-hand panel).¹⁴ We further attempted to correct the reconstructed power spectra for shot noise. Even though, as also will be seen below, the influence of shot noise is usually removed along with subtracting functions

¹⁴To provide better understanding, we isolate the effect of subsample variance, by the example of one simulation only, in Appendix A2.

describing anomalous broad-band power (Seo et al. 2008) we chose to isolate the computation of shot noise first, in order to disentangle noise that scales with particle number from other effects intrinsic to our method, see appendix A1. However, on the scales of our interest, we observe an additional offset that we correct for by allowing a constant, A_0 , to remove any difference between the reconstructed and the initial power spectra,

$$P_i(k) = A_0 + P_{ij}^{\text{rec}}(k). \quad (41)$$

We chose to fit A_0 in the range of $0.01 h \text{Mpc}^{-1} < k < 0.1 h \text{Mpc}^{-1}$, well in the linear regime, where we expect our method to deliver satisfactory results.

Already at this stage, the averages and the standard errors of the reconstructed power spectra agree with the linear expectation at the ~ 5 per cent-level at wavenumbers $k \lesssim 0.15 h \text{Mpc}^{-1}$ for both starting redshifts. Furthermore, relative deviations of reconstructed and initial power spectra seen in the bottom panels exhibit a smooth dependency on k that indeed is known to arise from mode-coupling




Figure 11. Average power spectra of ten ABACUSCOSMOS simulations and their reconstructions at redshift $z = 49$. Reconstructions were performed beginning with samples at redshift $z_s = 0.3$ (left) and $z_s = 1.5$ (right). *Upper panels:* Average power spectra and standard deviations for simulations and reconstructions, respectively. *Lower panels:* Averaged relative differences of reconstructed and true power spectra, cf. equation (40). The reconstructed power spectra were corrected for shot noise as described in the text. Shaded bands show 1σ deviation.

in the non-linear regime (Meiksin & White 1999; Scoccimarro, Zaldarriaga & Hui 1999; Crocce & Scoccimarro 2008; Seo et al. 2008; Xu et al. 2012). Hence, and in line with other approaches for initial density reconstruction (e.g. Xu et al. 2012), we introduce a broad-band term to our reconstructed power spectra, that, after fitting to the corresponding initial power spectrum, $P_i(k)$, compensates for this deviation:

$$P_i(k) = B(k) \cdot P_{ij}^{\text{rec}}(k), \quad (42)$$

where we found $B(k) = B_0 + B_1 \cdot k + B_2 \cdot k^2$ to be sufficient in accounting for the observed discrepancies when performing the fit in the range shown in the figures. The data show no further evidence for introducing additive corrections instead or in addition, and also gives no support for higher-order terms in $B(k)$. This supports our understanding of shotnoise contributions as well as their sufficient subtraction as above. While the linear growth factors in equation (39) were chosen to match exactly those of the simulations, we expect B_0 to be close to 1. Indeed, we find this to be the case, and for completeness list the fitted values in the table below.

| | $z_i = 0.3$ | $z_i = 1.5$ |
|-------|-------------|-------------|
| B_0 | 1.01 | 1.01 |
| B_1 | -0.86 | -0.43 |
| B_2 | 4.55 | 0.53 |

It should be noted that in practice the fitting is done by considering templates of linear power spectra that are each generated with a set of cosmological parameters, and are subsequently modified to include effects of non-linear growth. Since the intention of this paper is simply to show the accuracy of recovering the expected power of linear fluctuations, instead of recovering cosmological parameters, we chose to fit directly to the power spectrum $P_i(k)$ of the $z_f = 49$ density.

While the accounting of broad-band power is also necessitated by a combination of effects, such as redshift-space distortions and surveying effects, we must attribute the power discrepancy to the

reconstruction itself. In practice, however, such nuisance terms would absorb any such unexpected power, regardless of its source.

Fig. 12 shows the relative differences after having included the broad-band terms. The range of k over which the reconstructed power matches that of the initial density is striking, reaching $k \approx 0.4 \text{ hMpc}^{-1}$ (0.5 hMpc^{-1}) for $z_s = 0.3$ ($z_s = 1.5$) before deviating past the 5 per cent-level.

5.2.2 Correlation functions

Visually more revealing of the BAO signal is the two-point correlation function (the Hankel transform of the power spectrum). We repeat previous exercise for the correlation functions computed from simulations and reconstructions. Instead of the relative difference, as above, we here compute the absolute difference for each pair ij of simulation and reconstruction.

$$\Delta \xi_{ij}(r) := \xi_{ij}^{\text{rec}}(r) - \xi_i(r) \quad (43)$$

Correlation functions, their spread in simulations and reconstructions, and the absolute difference are shown in Fig. 13. The BAO peak and its shape are recovered well and with uncertainties comparable to the sample variance of the ICs themselves. As before, a discrepancy growing towards low separations r is removed by fitting for broad-band influences in the shown range. As opposed to broad-band noise in the power spectrum, we here find no evidence for any scale-dependencies and simply allow for a constant bias, again only performing the fit in the shown range,

$$\xi_i(r) = B \cdot \xi_{ij}^{\text{rec}}(r). \quad (44)$$

| | $z_i = 0.3$ | $z_i = 1.5$ |
|-----|-------------|-------------|
| B | 1.025 | 1.011 |

The corresponding average correlation functions are shown in Fig. 14. A slight ($\sim 1\sigma$) deviation around the BAO feature reveals residual dampening of the peak for the $z_s = 0.3$ reconstructions,




Figure 12. Same as Fig. 11 but after accounting for broad-band noise, cf. equation (42).




Figure 13. Average correlation functions of ten ABACUSCOSMOS simulations and their reconstructions at redshift $z = 49$. Reconstructions were performed beginning with samples at redshift $z_s = 0.3$ (left) and $z_s = 1.5$ (right). *Upper panels:* Average power spectra and standard deviations for simulations and reconstructions, respectively. *Lower panels:* Averaged differences of reconstructed and true correlation functions, cf. equation (43). Shaded bands show 1σ deviation.

which, however, is not significant for those reconstructions beginning at the higher redshift, $z_s = 1.5$.

To demonstrate the ability of our reconstruction algorithm on individual cases of correlation functions, we show in Fig. 15 three different examples, in which we compare the correlation functions of the reconstructed densities with those of the initial densities, as well as those of the samples with which the reconstruction algorithm was fed. Due to cosmic sample variance in the simulated volume, the BAO feature visible at $z_s = 0.3$ or $z_s = 1.5$ (dotted black) may be more or less pronounced (left-hand to right-hand panels). However, in all cases, the reconstruction is able to recover well the precise shape of the BAO bump at the redshift of reconstruction, $z_f = 49$, as shown by the close agreement between simulation (solid black) and reconstruction (red). In all cases, one can observe the well-

known sharpening of the peak (Eisenstein, Seo & White 2007a) that reconstruction methods generally aim for. Especially the third example of each panel exhibits this effect, wherein the peak is hardly recognizable at $z_s = 0.3$.

5.2.3 Accuracy of acoustic scale recovery

Finally, we provide a tentative quantification of the accuracy with which the sound horizon at decoupling, i.e. the position of the BAO peak, is recovered using the same simulations as above. To this effect, we localize the peak of the correlation function in the vicinity of its input value, $r_s(z_{\text{drag}}) = 99.09 h^{-1} \text{Mpc}$, after having interpolated the function with a cubic spline to smooth out discreteness effects. This




Figure 14. Same as Fig. 13, but after accounting for broad-band noise, cf. equation (44).




Figure 15. Examples of correlation functions in three different simulations and their reconstructions. The correlation functions as measured in snapshots at $z_s = 0.3$ (dotted black) and $z_f = 49$ (solid black) are compared with the reconstruction at $z_f = 49$ (red). All curves have been scaled to match the linear growth amplitude at $z = 49$ for visual comparison.

is done for both the initial correlation functions $\xi_i(r)$, leading to r_s^{IC} , as well as the reconstructed correlation functions $\xi_{ij}^{\text{rec}}(r)$ from both input redshifts, $z_s = 0.3$ and $z_s = 0.5$, resulting in r_s^{rec} . This results in 50 peak positions each. For both redshifts, we then compute the mean and standard deviation of the fractional deviation

$$\delta r_s = \frac{r_{s,ij}^{\text{rec}} - r_{s,i}^{\text{IC}}}{r_{s,i}^{\text{IC}}}, \quad (45)$$

$$\delta r_s \pm \sigma \text{ [per cent]} \Big|_{z_i = 0.3}^{z_i = 1.5} \Big| 0.02 \pm 0.37 \quad 0.08 \pm 0.52$$

Even with this rough attempt, we are able to recover the BAO scale with sub- per cent scatter. We will return to this point in future, dedicated works containing more comprehensive methods for the localization of the BAO peak that take into account the full shape of the peak, cf. Schmittfull et al. (2017), and including comparisons with existing reconstruction algorithms.

6 CONCLUSIONS

To extract information on the early Universe from the present Universe requires the undoing of the non-linear growth of struc-

tures. This complex inverse problem, sometimes referred to as the cosmological reconstruction problem, is often tackled in a forward iterative manner: essentially an initial model with a set of parameters is assumed and simulated and then compared to the data at the present epoch within an iterative loop till the best match between the model and the data is achieved *statistically*. In previous works, we have shown that the cosmological reconstruction problem is a subset of the general class of mass transportation problems that are solved through optimal transport theory and as such can be tackled deterministically to yield a unique solution. The problem is well-formulated by the Monge–Ampère equation. Our previous solution to the Monge–Ampère equation was obtained through a *combinatorial* fully discrete algorithm, which explored a huge solution space to find the optimal assignment. These algorithms have a complexity of N^3 , for a data-set with N points, which renders them impractical for applications to big data in cosmology. Although there have been reports of faster variants of combinatoric algorithms with a complexity of at best $N^2 \log(N)$ (see review in Mériqot & Thibert (2020)), their performance however remains too slow for large-scale cosmological problems.

In this article, we have presented a new *semidiscrete* algorithm which makes a direct use of the *variational* nature of the cosmological reconstruction problem. It finds a quick path to the solution by fully exploiting the first- and second-order derivatives of the objective function (that is both *smooth* (C^2) and *convex*). This is made possible by a fortuitous yet elegant convergence between the physical, mathematical, and computational aspects of the problem: the specific cosmological setting that we considered (continuous mass transported to a point set) has nice mathematical properties – Monge–Ampère equation translated into a smooth and convex optimization problem, with an underlying geometric structure (Laguerre diagram) that can be *exactly* computed by our algorithm. Our semidiscrete algorithm has a complexity of $M \log(N)$ which makes it significantly more efficient than any combinatoric algorithm.

As a concrete example of a practical consequence, our previous combinatoric code could reconstruct a typical 128^3 Λ CDM dark-matter-only simulation in a month on a desktop computer station. In contrast, with our new semidiscrete algorithm described in this article, the same reconstruction can be done on a portable personal computer in less than five minutes.

Subtle information in the matter power spectrum, such as BAO, can only be retrieved through the analyses carried out on horizon scales and on hundreds of millions of particles: a task completely unfeasible through combinatorial algorithms. Our new semidiscrete algorithm processes such massive datasets within hours, which makes it a powerful tool to reconstruct BAO and consequently test the theory of general relativity at cosmological scales. Here, we have chosen the BAO measurement as a challenge to test the power of our algorithm, however it is needless to say that other decisive features of the primordial power spectrum, *e.g.* primordial non-Gaussianity can also be detected by our algorithm and consequently we can also provide constraints on inflationary models (Mohayaee et al. 2006).

Aside from the BAO and primordial non-Gaussianity, our deterministic algorithm can also recover the velocity field, including the relevant phase information. Hence it can provide priors for Bayesian predictions of large-scale structure. The density and velocity field information in the reconstructed ICs can be used to develop Bayesian priors to probe CMB polarization and temperature anisotropy maps for evidence of any phase anomalies. This potentially provides a powerful new insight into the validity of the standard cosmological model.

Our reconstruction method holds for as long as the convexity holds which implies that the reconstructed map between initial and final distribution remains valid into the non-linear regime and at least up to the third order in the Lagrangian perturbation theory. However, in extracting the density field from these maps, we have used the Zel’dovich approximation for convenience as it gives us a simple analytic expression (39). The reason why higher order terms have not been implemented here is that for BAO reconstruction, a broadband fitting function is often used to account for mode-coupling as well as other effects. We have shown that our BAO reconstruction enjoys a high sub- per cent accuracy with the least number of fitting polynomial parameters. In forthcoming work, we shall implement higher orders which could reduce the need to fitting parameters and render the BAO reconstruction model-independent.

In this work, our code ran on particle samples where all points have the same mass and reside in real space within periodic boundary conditions. In the forthcoming works, we shall generalize this computational setting to other geometrical configurations, more relevant for observational surveys. We shall account for non-linear halo bias, adapt our code to redshift space and to more general boundary conditions, and survey geometries and make it available

for applications to data expected from future telescopes, as outlined below:

(i) In the case of haloes, one can use points of different masses. Our algorithm can be generalized in such a setting, with points of different masses (rather than $1/N$). The same generalized algorithm with points of different masses can be used to take into account galaxy biasing.

(ii) In the case of a very diffuse matter distribution, one can discretize a continuous density field using points of different masses. A well adapted (optimal in the sense of approximation theory) sampling can be generated using different techniques, see for instance Liu et al. (2009) and the references therein;

(iii) General boundary conditions and survey geometries can be taken into account by computing the intersection of the Laguerre cells with the survey geometry Yan et al. (2013);

(iv) Redshift distortions can be taken into account by recasting the redshift-space reconstruction as a physical-space reconstruction, on which our method can be directly applied. Such a technique is described in Brenier et al. (2003), equations (44) to (47);

(v) It is possible to reconstruct more precise dynamics using higher order Lagrangian perturbation theory, by introducing in the trajectories reconstructed by our algorithm some higher order polynomial terms, and finding the associated coefficients by minimizing the action (7), similarly to what is done in Branchini et al. (2002) and Sarpa et al. (2020).

ACKNOWLEDGEMENTS

The authors thank Jean-David Benamou, Enzo Branchini, Yann Brenier, Quentin M erigot, Joe Silk for discussions, and Alain Filbois who developed an early prototype of the parallel triangulation code. This work was supported by an Inria internal grant (Action exploratoire AeX EXPLORAGRAM). SvH thanks Lehman Garrison for assistance with the AbacusCosmos simulation suite. For analysing the results, this work made use of the NUMPY, MATPLOTLIB, NBODYKIT (Hand et al. 2018a), and HMF PYTHON packages.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author. The data sets were derived from the ABACUSCOSMOS suite (Garrison et al. 2018, 2019) (<https://lgarrison.github.io/AbacusCosmos/>).

REFERENCES

- Ade P. et al., 2016, *A&A*, 594, A13
 Alauzet F., Loseille A., 2009, Proceedings of the 18th International Meshing Roundtable, IMR. Springer, Salt Lake City, UT, USA, p. 337
 Amenta N., Choi S., Rote G., 2003, Proceedings of the Nineteenth Annual Symposium on Computational Geometry. SCG ’03. ACM, New York, NY, USA, p. 211
 Aurell E., Gaw edzki K., Mej a-Monasterio C., Mohayaee R., Muratore-Ginanneschi P., 2012, *J. Stat. Phys.*, 147, 487
 Aurenhammer F., 1987, *SIAM J. Comput.*, 16, 78
 Aurenhammer F., Hoffmann F., Aronov B., 1992, Proceedings of the Eighth Annual Symposium on Computational Geometry, ACM, Berlin, Germany, p. 350
 Benamou J., Brenier Y., 2000, *Numer. Math.*, 84, 375
 Benamou J.-D., 2018, Private communication
 Benamou J.-D., Carlier G., Cuturi M., Nenna L., Peyr  G., 2015, *SIAM J. Sci. Comput.*, 2, A1111
 Bertsekas D. P., 1992, *Comput. Optim. Appl.*, 1, 7

- Bertsekas D. P., Castanon D. A., 1989, *Ann. Ope. Res.*, 20, 67
- Bos E. G. P., Kitaura F.-S., van de Weygaert R., 2019, *MNRAS*, 488, 2573
- Bowyer A., 1981, *Comput. J.*, 24, 162
- Branchini E., Eldar A., Nusser A., 2002, *MNRAS*, 335, 53
- Brenier Y., 1991, *Commun. Pure Appl. Math.*, 44, 375
- Brenier Y., Frisch U., Hénon M., Loeper G., Matarrese S., Mohayaee R., Sobolevskii A., 2003, *MNRAS*, 346, 501
- Buchert T., 1993, *A&A*, 267, L51
- Caroli M., Teillaud M., 2009, *Proceedings of Algorithms - ESA 2009*, 17th Annual European Symposium. Springer, Copenhagen, Denmark, p. 59
- Catelan P., 1995, *MNRAS*, 276, 115
- Cautun M., van de Weygaert R., Jones B. J. T., Frenk C. S., 2014, *MNRAS*, 441, 2923
- Crocce M., Scoccimarro R., 2008, *Phys. Rev. D*, 77, 023533
- Croft R. A. C., Gaztanaga E., 1997, *MNRAS*, 285, 793
- Cuturi M., 2013, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013. ACM, Lake Tahoe, Nevada, United States*, p. 2292
- Delage C., Devillers O., 2004, *Spatial Sorting*. Available at: http://doc.cgal.org/latest/Spatial_sorting/index.html
- Edelsbrunner H., Mücke E. P., 1990, *ACM Trans. Graph.*, 9, 66
- Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
- Eisenstein D. J., Seo H.-J., White M. J., 2007a, *AJ*, 664, 660
- Eisenstein D. J., Seo H.-J., Sirko E., Spergel D. N., 2007b, *ApJ*, 664, 675
- Enßlin T. A., Frommert M., Kitaura F. S., 2009, *Phys. Rev. D*, 80, 105005
- Frisch U., Matarrese S., Mohayaee R., Sobolevskii A., 2002, *Nature*, 260, 417
- Garrison L. H., Eisenstein D. J., Ferrer D., Metchnik M. V., Pinto P. A., 2016, *MNRAS*, 461, 4125
- Garrison L. H., Eisenstein D. J., Ferrer D., Tinker J. L., Pinto P. A., Weinberg D. H., 2018, *ApJS*, 236, 43
- Garrison L. H., Eisenstein D. J., Pinto P. A., 2019, *MNRAS*, 485, 3370
- Giavalisco M., Mancinelli B., Mancinelli P.-J., Yahil A., 1993, *ApJ*, 411, 9
- Gramann M., 1993, *ApJ*, 405, 449
- Hand N., Feng Y., Adematti Li Y., Modi C., 2018a, *AJ*, preprint ([arXiv:1712.05834](https://arxiv.org/abs/1712.05834))
- Hand N., Feng Y., Beutler F., Li Y., Modi C., Seljak U., Slepian Z., 2018b, *AJ*, 156, 160
- Hénon M., 2002, *A mechanical model for the transportation problem*. preprint ([arXiv:math/0209047](https://arxiv.org/abs/math/0209047))
- Hestenes M. R., Stiefel E., 1952, *J. Res. NIST*, 49, 409
- Jasche J., Wandelt B. D., 2013, *MNRAS*, 432, 894
- Kantorovich L.-V., 1942, *CR URSS Acad. Sci.*, 37, 199
- Kashlinsky A., 1998, *ApJ*, 492, 1
- Kitagawa J., Mérigot Q., Thibert B., 2019, *Journal of the European Mathematical Society*, 21, 2603
- Kitaura F. S., Enßlin T. A., 2008, *MNRAS*, 389, 497
- Landau L., Lifshitz E., 1975, *Course of Theoretical Physics, Volume 1: Mechanics*. Butterworth-Heinemann, Oxford, UK
- Lavaux G., Mohayaee R., Colombi S., Tully R. B., Bernardeau F., Silk J., 2008, *MNRAS*, 383, 1292
- Lavaux G., Tully R. B., Mohayaee R., Colombi S., 2010, *ApJ*, 709, 483
- Lévy B., 2015, *ESAIM: M2AN (special issue - optimal transport)*, Vol. 49, EDP Sciences, Les Ulis, France, p. 1693
- Lévy B., 2016, *Comput. Aided Des.*, 72, 3
- Lévy B., Schwindt E. L., 2018, *Comput. Graph.*, 72, 135
- Liu Y., Wang W., Lévy B., Sun F., Yan D., Lu L., Yang C., 2009, *ACM Trans. Graph.*, 28, 101
- Lukic Z., Heitmann K., Habib S., Bashinsky S., Ricker P. M., 2007, *ApJ*, 671, 1160
- Meiksin A., White M., 1999, *MNRAS*, 308, 1179
- Mérigot Q., Thibert B., 2021, *Handbook of Numerical Analysis*, Vol. 22, Elsevier, p. 133
- Meyer A., Pion S., 2008, *Real Numbers and Computers*. Santiago de Compostela, Espagne, p. 47
- Mohayaee R., Sobolevskii A., 2008, *Physica D Nonlinear Phenom.*, 237, 2145
- Mohayaee R., Tully R. B., 2005, *ApJ*, 635, L113
- Mohayaee R., Frisch U., Matarrese S., Sobolevskii A., 2003, *A&A*, 406, 393
- Mohayaee R., Mathis H., Colombi S., Silk J., 2006, *MNRAS*, 365, 939
- Monaco P., Efstathiou G., 1999, *MNRAS*, 308, 763
- Monge G., 1781, *Edition de l'imprimerie royale*, p. 666
- Narayanan V. K., Croft R. A. C., 1999, *ApJ*, 515, 471
- Neyrinck M. C., Szapudi I., Szalay A. S., 2011, *ApJ*, 731, 116
- Nusser A., Branchini E., 2000, *MNRAS*, 313, 587
- Nusser A., Dekel A., 1992, *ApJ*, 391, 443
- Nusser A., Dekel A., Bertschinger E., Blumenthal G. R., 1991, *ApJ*, 379, 6
- Padmanabhan N., Xu X., Eisenstein D. J., Scalzo R., Cuesta A. J., Mehta K. T., Kazin E., 2012, *MNRAS*, 427, 2132
- Peebles P., 1980, *The Large-Scale Structure of the Universe*. Princeton Univ. Press, Princeton, NJ
- Peebles P. J. E., 1989, *ApJ*, 344, L53
- Peebles P. J. E., 1994, *ApJ*, 429, 43
- Peyré G., Cuturi M., 2019, *Foundations and Trends in Machine Learning*, Vol. 11, Mike Casey, Delft, The Netherlands, p. 355
- Ray N., Sokolov D., Lefebvre S., Lévy B., 2018, *ACM Trans. Graph.*, 37, 265
- Reverdy V. et al., 2015, *Int. J. High Perform. Comput. Appl.*, 29, 249
- Santambrogio F., 2015, *Optimal Transport for Applied Mathematicians. Progress in Nonlinear Differential Equations and their Applications Vol. 87*. Birkhäuser/Springer, Cham
- Sarpa E., Schimd C., Branchini E., Matarrese S., 2019, *MNRAS*, 484, 3818
- Sarpa E., Veropalumbo A., Schimd C., Branchini E., Matarrese S., 2019, *Monthly Notices of the Royal Academy Society*, 484, 3818
- Schmittfull M., Baldauf T., Zaldarriaga M., 2017, *Phys. Rev. D*, 96, 023505
- Scoccimarro R., Zaldarriaga M., Hui L., 1999, *ApJ*, 527, 1
- Seo H.-J. et al., 2010, *ApJ*, 720, 1650
- Seo H.-J., Siegel E. R., Eisenstein D. J., White M., 2008, *ApJ*, 686, 13
- Shaya E. J., Peebles P. J. E., Tully R. B., 1995, *ApJ*, 454, 15
- Shewchuk J. R., 1997, *Discrete Comput. Geom.*, 18, 305
- Shi Y., Cautun M., Li B., 2018, *Phys. Rev. D*, 97, 023505
- van de Weygaert R., 1994, *A&A*, 283, 361
- Villani C., 2009, *Optimal Transport : Old and New*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin
- Wang X., Yu H.-R., Zhu H.-M., Yu Y., Pan Q., Pen U.-L., 2017, *ApJ*, 841, L29
- Watson D., 1981, *Comput. J.*, 24, 167
- Weinberg D. H., 1992, *MNRAS*, 254, 315
- White S. D. M., Springel V., 1999, *Comput. Sci. Eng.*, 1, 36
- Wuttke J., 2017, *Form factor (Fourier shape transform) of polygon and polyhedron*. preprint ([arXiv:math/1703.00255](https://arxiv.org/abs/math/1703.00255))
- Xu X., Padmanabhan N., Eisenstein D. J., Mehta K. T., Cuesta A. J., 2012, *MNRAS*, 427, 2146
- Yan D., Wang W., Lévy B., Liu Y., 2010, *Proceedings of Advances in Geometric Modeling and Processing, 6th International Conference. GMP, Castro Urdiales, Spain*, p. 269
- Yan D., Wang K., Lévy B., Alonso L., 2011, *Eighth International Symposium on Voronoi Diagrams in Science and Engineering. ISVD, Qingdao, China*, p. 177
- Yan D., Wang W., Lévy B., Liu Y., 2013, *Comput. Aided Des.*, 45, 843
- Zel'dovich Y.-B., 1970, *A&A*, 5, 84

APPENDIX A: DETAILS OF POWER SPECTRUM PREPARATION

A1 Shot noise

Shot noise generally refers to Poisson noise of a discretized continuous density field and in most practical examples scales inversely proportional to the number of discretizing sources N . The power spectrum corresponding to the particle distribution in question is




Figure A1. Upper panels: power spectra of reconstructions using particle samples of different size N , highlighting the N -dependent shot noise. Reconstructions were done with an input redshift of $z = 0.3$ (left) and $z = 1.5$ (right). Lower panels: differences of each of the power spectra with that of the largest sample size, $N = 4 \times 256^3 = 67, 108, 864$.




Figure A2. Illustration of subsample variance given our choice of $N = 256^3$ particles. Shown are mean and standard deviation of the relative difference $\delta P_1(k)$, equation (40), for simulation $i = 1$ averaged over all five subsamples j , at input redshifts of $z = 0.3$ (left) and $z = 1.5$ (right).

thus increased by the k -independent noise contribution as $P(k, N) = P(k) + P_{sn}(N)$. The shot noise scales as

$$P_{sn}(N) = A \cdot N^{-1}, \quad (\text{A1})$$

where $A = l^3$ is the volume of the box. However, the power spectra shown in Fig. 11 were not computed from a discrete set of point particles. Instead, the continuous density field was effectively discretized by extended objects, the shrunken Laguerre cells, e.g. Fig. 6. It is non-trivial to predict the resulting shot noise for these structures. Empirically, however, equation (A1) seems to describe the shot noise contribution of this arrangement rather well, simply after promoting A to be a free parameter. The following paragraph describes how the shot noise is computed and finally subtracted from the corresponding power spectra in Fig. 11.

By definition of what should here be referred to as shot noise, the scaling with particle number N allows an estimation of its amplitude A from comparing reconstructed power spectra with one another, without the knowledge of the desired, ‘true’ power spectrum that provided the IC. This is especially promising for future reconstructions using real data, where the true solution is obviously unknown. We perform reconstructions of the initial density at $z = 49$ as in the main body, using particle samples of increasing size. Their uncorrected power spectra are shown in the upper panels of Fig. A1 along with the IC power spectrum, clearly demonstrating the

presence of shot noise. Without reference to the IC power spectrum, the shot noise amplitude of equation (A1) is fit by comparing power spectra of different particle number with that of one selected particle number N_* :

$$P(k, N) - P(k, N_*) = A \cdot \frac{N_* - N}{NN_*} \quad (\text{A2})$$

These quantities are shown in the bottom panels of Fig. A1 with $N_* = 2 \times 256^3$, the largest particle sample used here, and clearly show the corresponding N -dependent contributions to the reconstructed power spectra. It should be stated here that any discrepancies that remain between the reconstructed power spectra after proper shot noise subtraction and the IC power spectrum are therewith shown *not* to be dependent on the size of the particle sample, and must be accounted for by other means. In each of the same panels, indicated by dashed lines, the relation A2 was fitted collectively to all curves within the range $0.25 \text{ hMpc}^{-1} < k < 0.5 \text{ hMpc}^{-1}$. For the reconstructions from $z = 0.3$ and 1.5 (left-hand and right-hand panels) we find:

$$\frac{A \text{ [h}^{-3} \text{Mpc}^3]}{z_i} \Big|_{z_i = 0.3} \quad \Big|_{z_i = 1.5} \quad \frac{A \text{ [h}^{-3} \text{Mpc}^3]}{z_i} \Big|_{z_i = 0.3} \quad \Big|_{z_i = 1.5}$$

Finally, this allows for calculating the shot noise for a reconstruction from a given particle number via equation (A1). For the particle sample of size $N = 256^3$, as employed in the main body of this paper,

we find

$$\frac{P_{sn}(N = 256^3)}{P_{sn}(N = 256^3)} \Big|_{z_i = 0.3}^{z_i = 1.5} \begin{matrix} 0.0393 \\ 0.1387 \end{matrix}$$

Note that, both values are smaller than the shot noise expected for a discrete particle sample, even after accounting for a scaling by the linear growth factors, $D^2(z = 49)/D^2(z_i) \cdot \bar{l}^3/N$.

A2 Subsample variance

Subsample variance decreases with increasing particle number. Sampling as many as 256^3 particles therefore suppresses the influence of these terms considerably. Nevertheless, we show in Fig. A2 the relative difference of the power spectra of each of one simulations' five subsamples and their average. The variance on average makes up about 75 per cent of the total variance shown in Figs 11–12 in the range shown.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.