

Lawrence Berkeley National Laboratory

LBL Publications

Title

encore : an $O(N^2)$ estimator for galaxy N -point correlation functions

Permalink

<https://escholarship.org/uc/item/71t6g5vt>

Journal

Monthly Notices of the Royal Astronomical Society, 509(2)

ISSN

0035-8711

Authors

Philcox, Oliver HE
Slepian, Zachary
Hou, Jiamin
et al.

Publication Date

2021-11-23

DOI

10.1093/mnras/stab3025

Peer reviewed

ENCORE: An $O(N_g^2)$ Estimator for Galaxy N -Point Correlation Functions

Oliver H. E. Philcox^{1,2*}, Zachary Slepian^{3,4}, Jiamin Hou³, Craig Warner³,
Robert N. Cahn⁴ and Daniel J. Eisenstein⁵

¹*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08540, USA*

²*School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA*

³*Department of Astronomy, University of Florida, 211 Bryant Space Science Center, Gainesville, FL 32611, USA*

⁴*Physics Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94709, USA*

⁵*Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA*

14 October 2021

ABSTRACT

We present a new algorithm for efficiently computing the N -point correlation functions (NPCFs) of a 3D density field for arbitrary N . This can be applied both to a discrete spectroscopic galaxy survey and a continuous field. By expanding the statistics in a separable basis of isotropic functions built from spherical harmonics, the NPCFs can be estimated by counting pairs of particles in space, leading to an algorithm with complexity $O(N_g^2)$ for N_g particles, or $O(N_{\text{FFT}} \log N_{\text{FFT}})$ when using a Fast Fourier Transform with N_{FFT} grid-points. In practice, the rate-limiting step for $N > 3$ will often be the summation of the histogrammed spherical harmonic coefficients, particularly if the number of radial and angular bins is large. In this case, the algorithm scales linearly with N_g . The approach is implemented in the ENCORE code, which can compute the 3PCF, 4PCF, 5PCF, and 6PCF of a BOSS-like galaxy survey in ~ 100 CPU-hours, including the corrections necessary for non-uniform survey geometries. We discuss the implementation in depth, along with its GPU acceleration, and provide practical demonstration on realistic galaxy catalogs. Our approach can be straightforwardly applied to current and future datasets to unlock the potential of constraining cosmology from the higher-point functions.

Key words: methods: statistical, numerical – Cosmology: large-scale structure of Universe, theory – galaxies: statistics

1 INTRODUCTION

Amongst the most powerful tools in the survey analyst’s workshop are the N -point correlation functions (NPCFs), or their Fourier-space counterparts, the polyspectra. These encode the statistical properties of the galaxy overdensity field at sets of N positions, and may be compared to data to give constraints on properties such as the Universe’s expansion rate and composition. Most inflationary theories predict density fluctuations in the early Universe to follow Gaussian statistics; in this case all the information is contained within the two-point correlation function (2PCF), or, equivalently, the power spectrum. For a homogeneous and isotropic Universe, both are simple functions of one variable, and have been the subject of almost all galaxy survey analyses to date (e.g., Alam et al. 2017; Beutler et al. 2017b; Alam et al. 2021; Ivanov et al. 2020; Philcox et al. 2020; d’Amico et al. 2020; Colas et al. 2020).

The late-time Universe is far from Gaussian. Statistics beyond the 2PCF have been long discussed in the literature (e.g., Peebles & Groth 1975; Peebles 1978, 2001; Scoccimarro 2000; Bernardeau et al. 2002; Coles 2005), and, as shown in Schmittfull et al. (2015), are of importance since the bulk motion of matter in structure formation causes a cascade of information from the 2PCF to higher-order statistics, such as the three- and four-point functions (3PCF and 4PCF). Recent work by Gualdi et al. (2021) demonstrated the power of such correlators; for the matter field, forecasted constraints on key cosmological parameters were shown to tighten significantly when the bispectrum and trispectrum (the Fourier-space counterparts to the 3PCF and 4PCF) were included in the analysis. A number of works have demonstrated similar effects for the galaxy density field, including improved constraints on Λ CDM parameters (Gagrani & Samushia 2017; Agarwal et al. 2021), neutrino masses (Chudaykin & Ivanov 2019; Kamalinejad & Slepian 2020; Hahn & Villaescusa-Navarro 2021), modified gravity (Alam et al. 2020), and primordial non-Gaussianity (Karagiannis et al. 2018; Moradinezhad Dizgah et al. 2021). Samushia et al. (2021) provides another notable example, demonstrating that, when higher-order correlators are included, parameters such as the Alcock-Paczynski dilation statistic α can be constrained more strongly from late-time measurements than from the initial matter power spectrum. This occurs since such parameters (which are a key source of information on dark energy) appear only in the forward model rather than in the initial conditions. It remains to be seen whether the purported gains hold in the more realistic scenarios of shot-noise limited galaxy surveys.

* E-mail: ohep2@cantab.ac.uk (OP)

The above discussion suggests that large-scale structure analyses should certainly be including statistics beyond the 2PCF. This is far from trivial however, since the higher-order NPCFs (or polyspectra) are (a) expensive to estimate, (b) difficult to model and (c) of high dimension. Naïve NPCF estimation requires counting all N -tuples of galaxies (and random particles) and assigning them to bins; whilst possible for $N = 2$, this is prohibitively slow for the higher-point statistics, and scales as N_g^N for N_g particles. For the simplest beyond-Gaussian statistic, the isotropic 3PCF, a number of algorithms of varying efficiency have been proposed to circumvent this problem, using techniques such as kd -trees, Fourier Transforms, and basis decompositions (e.g., Szapudi & Szalay 1998; Moore et al. 2001; Gray et al. 2004; Szapudi 2004; Zhang & Pen 2005; Slepian & Eisenstein 2015; March et al. 2012; Slepian & Eisenstein 2016; Friesen et al. 2017; Portillo et al. 2018; Philcox & Eisenstein 2020; Philcox 2021b as well as Sabiu et al. 2019; Tomlinson et al. 2019 for the (possibly compressed) 4PCF.)

In terms of theory models, whilst a number exist for the galaxy 3PCF and bispectrum (e.g., Scoccimarro 2000; Takada & Jain 2003; Wang et al. 2004; Gaztañaga & Scoccimarro 2005; Marín et al. 2008; Slepian & Eisenstein 2017; Simonović et al. 2018; Umeh 2021), many of which are semi-analytic, much work remains to be done. For the matter field, perturbative approaches have been computed up to one-loop order (Baldauf et al. 2015; Angulo et al. 2015; Bertolini et al. 2016), but this has not yet been extended to galaxy surveys, hampered by the necessity to include a full set of quartic (quintic) galaxy bias parameters for the 3PCF (4PCF) model.

Finally, even when the statistics are measured and modeled, extracting information from them is non-trivial due to their high dimension. Accurate analyses are likely to require data compression techniques (e.g., Scoccimarro 2000; Heavens et al. 2000; Munshi et al. 2011; Burkhart & Lazarian 2016; Joachimi 2017; Alsing & Wandelt 2018; Galdi et al. 2018; Slepian et al. 2018; Galdi et al. 2019; Philcox et al. 2021b) as well as analytic covariance matrices (e.g., Slepian & Eisenstein 2015, 2018; Philcox & Eisenstein 2019; Sugiyama et al. 2020; Umeh 2021 for the 3PCF). To date, the isotropic 3PCF or bispectrum has been included in a number of analyses (Jing & Börner 1998; Scoccimarro et al. 2001; Jing & Börner 2004; Kayo et al. 2004; Nichol et al. 2006; Marín 2011; Gil-Marín et al. 2015, 2017; Slepian et al. 2017b, 2018; Pearson & Samushia 2018), but few use any higher-order statistics (though see Sabiu et al. 2019, Peebles 1978 and the review of Peebles 2001 for some notable exceptions, as well as Münchmeyer & Smith (2019) for the CMB).

The present work tackles the following problem: how can we efficiently measure the isotropic NPCFs for $N > 3$? To do this, our starting point is the isotropic 3PCF algorithm of Slepian & Eisenstein (2015). By expanding the correlator’s angular dependence into a Legendre polynomial basis and further factorizing these into spherical harmonics, the former work was able to estimate the 3PCF of N_g particles with $\mathcal{O}(N_g^2)$ complexity; much faster than a naïve triple count, which has a complexity of $\mathcal{O}(N_g^3)$. This has facilitated a number of analyses, including detection of the 3PCF baryon acoustic oscillation (BAO) feature and a constraint on the baryon-dark matter relative velocity (Slepian et al. 2017a,b, 2018). Here, we consider whether such an approach is possible for higher-point functions. A crucial component is the choice of basis; for this we will use the separable isotropic $(N - 1)$ -point basis functions introduced in Cahn & Slepian (2020) (see also Mitsou et al. 2020), which are a generalization of the polypolar spherical harmonics of zero total angular momentum (Varshalovich et al. 1988). Indeed, the principal conclusion of this work is that an $\mathcal{O}(N_g^2)$ NPCF estimator is possible for *all* N ; furthermore, its practical implementation scales linearly with N_g if N is large. In fact, such an estimator can be constructed in arbitrary dimensions, as long as the underlying space is homogeneous and isotropic; generalization to such cases is discussed in our companion work Philcox & Slepian (2021). Below, we will present a detection of the galaxy 4PCF and 5PCF by applying our algorithm, implemented in the ENCORE code, to mock catalogs. It is important to note that, for $N > 3$, only part of the information in the NPCF is independent from that in lower-point functions; there is also a ‘disconnected’ term that does not add further information.

A key assumption underlying the above discussion is that of isotropy. Arising from the projection of galaxies’ peculiar motions onto our line of sight, redshift-space distortions (RSD; Kaiser 1987; Hamilton 1998; Zaroubi & Hoffman 1996) break this ideal, introducing additional dependence on the line-of-sight angle. However, this effect also adds cosmological information in the form of the growth-rate $f(z)$ and allows us to break degeneracies, for example those between linear galaxy bias and the primordial amplitude A_S in the 2PCF (e.g., Weinberg et al. 2013). Inclusion of such effects has become commonplace in any 2PCF analysis, but has been limited for higher-point functions, though Galdi & Verde (2020); Sugiyama et al. (2021) show anisotropy to contain useful information. In general, whilst there are a handful of estimators dedicated to the measurement of the anisotropic 3PCF or bispectrum (e.g., Gardner et al. 2007; Scoccimarro 2015; Slepian & Eisenstein 2018; Sugiyama et al. 2019; Garcia & Slepian 2020), they have been scarcely used. Whilst we consider only the isotropic NPCFs in this work, it is expected that the methodology can be extended further, utilizing basis functions with a non-zero total angular momentum; this is discussed in Philcox & Slepian (2021).

The remainder of this paper is structured as follows. We begin in §2 by introducing the basis functions underlying our algorithm (following Cahn & Slepian 2020), before presenting the NPCF estimators in §3, including discussion of the split into ‘connected’ and ‘disconnected’ components (§3.1), and the necessary corrections for non-uniform survey geometry (§3.3). §4 contains a detailed discussion of the ENCORE code, including its CPU and GPU implementations, before we present results in §5, both concerning the algorithm scalings and application to mock data. We conclude with a summary in §6, before presenting several useful mathematical results in Appendices A & B.

2 BASIS FUNCTIONS

To begin, we present the isotropic basis functions underlying our NPCF estimators and outline some of their key properties. Full details of these functions (which are similar to those of Mitsou et al. 2020) are presented in Cahn & Slepian (2020). Generalization anisotropic functions, as well as homogeneous and isotropic spaces of arbitrary dimension is presented in Philcox & Slepian (2021).

2.1 Definition

Consider a function, f , depending on $(N - 1)$ spatial coordinates in 3D Euclidean space, $\mathbf{r}_1, \dots, \mathbf{r}_{N-1}$.¹ If the function is invariant under simultaneous rotations of all $(N - 1)$ coordinates, *i.e.* $f(R\mathbf{r}_1, \dots, R\mathbf{r}_{N-1}) = f(\mathbf{r}_1, \dots, \mathbf{r}_{N-1})$ for any rotation matrix R , f can be expanded in a basis of isotropic functions as

$$f(\mathbf{r}_1, \dots, \mathbf{r}_{N-1}) = \sum_{\Lambda} f_{\Lambda}(r_1, \dots, r_{N-1}) \mathcal{P}_{\Lambda}(\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{N-1}), \quad (1)$$

where the set Λ indexes the basis function in question; the precise form of this will be given below. To fulfill the requirements of isotropy, the functions must satisfy

$$\mathcal{P}_{\Lambda}(R\hat{\mathbf{R}}) = \mathcal{P}_{\Lambda}(\hat{\mathbf{R}}) \quad (2)$$

for all rotations R , introducing the nomenclature $\mathbf{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_{N-1}\}$, $\hat{\mathbf{R}} = \{\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{N-1}\}$.

We now consider the functional form of such a basis, following [Cahn & Slepian \(2020\)](#) and [Philcox & Slepian \(2021\)](#). Firstly, we note that any scalar function of $(N - 1)$ variables in 3D Euclidean space can be expanded in terms of $(N - 1)$ spherical harmonics, regardless of isotropy:

$$g(\mathbf{r}_1, \dots, \mathbf{r}_{N-1}) = \sum_{\ell_1 \dots \ell_{N-1}} \sum_{m_1 \dots m_{N-1}} g^{m_1 \dots m_{N-1}}(r_1, \dots, r_{N-1}) Y_{\ell_1 m_1}(\hat{\mathbf{r}}_1) \dots Y_{\ell_{N-1} m_{N-1}}(\hat{\mathbf{r}}_{N-1}), \quad (3)$$

where $Y_{\ell m}(\hat{\mathbf{r}})$ is a spherical harmonic (in the Condon-Shortley phase convention), ℓ and m are the orbital and magnetic angular momentum quantum numbers, and we sum over all $\ell_i \geq 0$ and $|m_i| \leq \ell_i$. If we further assume g to be isotropic, relation (2) places constraints on the allowed values of m_i ; in particular, this enforces $m_1 + \dots + m_{N-1} = 0$. To obtain an isotropic basis from (3), we must consider the possible ways in which to combine products of $(N - 1)$ spherical harmonics to obtain a rotationally invariant state. As shown in [Cahn & Slepian \(2020\)](#), such an approach leads to the set of basis functions

$$\mathcal{P}_{\Lambda}(\hat{\mathbf{R}}) = \sum_{m_1 \dots m_{N-1}} C_M^{\Lambda} Y_{\ell_1 m_1}(\hat{\mathbf{r}}_1) \dots Y_{\ell_{N-1} m_{N-1}}(\hat{\mathbf{r}}_{N-1}), \quad (4)$$

where C_M^{Λ} is a set of coupling coefficients (which enforce $m_1 + \dots + m_{N-1} = 0$) and $M \equiv \{m_1 \dots m_{N-1}\}$. The coupling coefficients can be written either in terms of Clebsch-Gordan coefficients or Wigner 3- j symbols (e.g., [NIST DLMF](#), §34.2); in the latter case, we obtain

$$C_M^{\Lambda} \equiv C_{m_1 m_2 \dots m_{N-1}}^{\ell_1 \ell_2 (\ell_{12}) \dots \ell_{N-1}} = \mathcal{E}(\Lambda) \sqrt{2\ell_{12} + 1} \times \dots \times \sqrt{2\ell_{12 \dots N-3} + 1} \quad (5)$$

$$\times \sum_{m_{12} \dots} (-1)^{\kappa} \begin{pmatrix} \ell_1 & \ell_2 & \ell_{12} \\ m_1 & m_2 & -m_{12} \end{pmatrix} \begin{pmatrix} \ell_{12} & \ell_3 & \ell_{123} \\ m_{12} & m_3 & -m_{123} \end{pmatrix} \dots \begin{pmatrix} \ell_{12 \dots N-3} & \ell_{N-2} & \ell_{N-1} \\ m_{12 \dots N-3} & m_{N-2} & m_{N-1} \end{pmatrix},$$

with $\mathcal{E}(\Lambda) = (-1)^{\sum_i \ell_i}$ and $\kappa = \ell_{12} - m_{12} + \ell_{123} - m_{123} + \dots + \ell_{12 \dots N-3} - m_{12 \dots N-3}$. The basis (4) can be derived by considering the eigenstates of the angular momentum operators; in particular the basis is an eigenstate of \mathbf{L}_1 with eigenvalue $\ell_1(\ell_1 + 1)$, an eigenstate of $(\mathbf{L}_1 + \mathbf{L}_2)^2$ with eigenvalue $\ell_{12}(\ell_{12} + 1)$ *et cetera*, as well as a zero-eigenvalue eigenstate of the combined angular momentum operator $(\mathbf{L}_1 + \dots + \mathbf{L}_{N-1})^2$. We note that this basis is simply related to the polypolar spherical harmonics of [Varshalovich et al. \(1988\)](#), except that we restrict to zero total angular momentum. Furthermore, the $N = 4$ case is a rescaled version of a subset of the tripolar spherical harmonics, which have been used a number of times in cosmological contexts (e.g., [Shiraishi et al. 2017](#)).

Comparison to (1) shows that the basis functions are parametrized by sets of integers $\Lambda = \{\ell_1, \ell_2, (\ell_{12}), \ell_3, (\ell_{123}), \dots, \ell_{N-1}\}$, hereafter denoted *multiplets*. These contain two types of contributions: (a) principal angular momenta ℓ_i (which index the spherical harmonics), and (b) intermediate angular momenta ℓ_{12}, ℓ_{123} *et cetera* (appearing only in the coupling C_M^{Λ}).² For $N > 4$, intermediates are required since there is more than one way to combine the $(N - 1)$ spherical harmonics to obtain a state of zero combined angular momentum; put another way, due to the vectorial-nature of angular momentum addition, the principal angular momenta $\{\ell_1, \ell_2, \dots, \ell_{N-1}\}$ are insufficient to fully define the state. From (5), we note several constraints on the values of ℓ and m . In particular, the 3- j symbols enforce that $m_{12} = m_1 + m_2$, $m_{123} = m_{12} + m_3$ *et cetera* and thus that $\sum_i m_i = 0$, giving invariance under z -axis rotations. This ensures that the basis function (4) is a state of zero combined orbital angular momentum. Furthermore, each 3- j symbol implies a triangle inequality of the form $|\ell_1 - \ell_2| \leq \ell_{12} \leq \ell_1 + \ell_2$. This restricts the number of physical Λ states.

2.2 Basic Properties

The coupling coefficients, C_M^{Λ} , appearing in (4) satisfy:

$$C_{-M}^{\Lambda} = \mathcal{E}(\Lambda) C_M^{\Lambda}, \quad (6)$$

¹ Foreshadowing the application to N -point functions, we consider only basis functions involving $(N - 1)$ coordinates in this section. Statistical homogeneity dictates that the NPCFs are independent of the N -th spatial position.

² Throughout this work, we indicate intermediate angular momentum components by parentheses.

where $-M = \{-m_1, -m_2, \dots, -m_{N-1}\}$, and $\mathcal{E}(\Lambda) = +1$ if the sum of all principal angular momentum ($\sum_i \ell_i$) is even, and -1 else. Consequently the isotropic basis functions obey the conjugate condition

$$\mathcal{P}_\Lambda^*(\hat{\mathbf{R}}) = \mathcal{E}(\Lambda)\mathcal{P}_\Lambda(\hat{\mathbf{R}}) \quad (7)$$

(Cahn & Slepian 2020). Under parity transformations, $\mathbb{P}[Y_{\ell m}(\hat{\mathbf{r}})] \equiv Y_{\ell m}(-\hat{\mathbf{r}}) = (-1)^\ell Y_{\ell m}(\hat{\mathbf{r}})$, with parity operator \mathbb{P} . Thus

$$\mathbb{P}[\mathcal{P}_\Lambda(\hat{\mathbf{R}})] = \mathcal{E}(\Lambda)\mathcal{P}_\Lambda(\hat{\mathbf{R}}). \quad (8)$$

Due to this property, we will refer to the cases $\mathcal{E}(\Lambda) = +1$ and -1 as *parity-even* and *parity-odd* respectively. We further note that (7) implies that $\mathcal{P}_\Lambda(\hat{\mathbf{R}})$ is purely real for even-parity Λ and purely imaginary for odd-parity Λ .

Secondly, the basis functions are orthonormal when integrated over $\hat{\mathbf{R}}$:

$$\int d\hat{\mathbf{R}} \mathcal{P}_\Lambda(\hat{\mathbf{R}})\mathcal{P}_{\Lambda'}^*(\hat{\mathbf{R}}) = \delta_{\Lambda\Lambda'}^{\mathbf{K}}, \quad (9)$$

where $d\hat{\mathbf{R}} = d\hat{\mathbf{r}}_1 \dots d\hat{\mathbf{r}}_{N-1}$ and the Kronecker delta is unity if all components of Λ and Λ' match and zero else. This follows straightforwardly from (5) and the orthonormality of the spherical harmonics. Given the expansion (1), this allows the basis coefficients f_Λ to be written as

$$f_\Lambda(R) = \int d\hat{\mathbf{R}} f(\mathbf{R})\mathcal{P}_\Lambda^*(\hat{\mathbf{R}}), \quad (10)$$

adopting the notation $R = \{r_1, \dots, r_{N-1}\}$. For $\Lambda = \mathbf{0}$, the basis functions take the simple form

$$\mathcal{P}_0(\hat{\mathbf{R}}) = (4\pi)^{-(N-1)/2}, \quad (11)$$

using $Y_{00}(\hat{\mathbf{r}}) = (4\pi)^{-1/2}$ and $C_0^0 = 1$.

An additional property of interest is the contraction of two basis functions of the same arguments. This is the analog of the well-known spherical harmonic product-to-sum relation;

$$Y_{\ell_1 m_1}(\hat{\mathbf{r}})Y_{\ell_2 m_2}(\hat{\mathbf{r}}) = \sum_{\ell_3 m_3} \mathcal{G}_{m_1 m_2 (-m_3)}^{\ell_1 \ell_2 \ell_3} (-1)^{m_3} Y_{\ell_3 m_3}(\hat{\mathbf{r}}) \quad (12)$$

(NIST DLMF, Eq. 34.3.22), where $\mathcal{G}_{m_1 m_2 (-m_3)}^{\ell_1 \ell_2 \ell_3}$ is the Gaunt integral, defined by

$$\mathcal{G}_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \equiv \left[\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi} \right]^{1/2} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}. \quad (13)$$

In our context, we find

$$\mathcal{P}_\Lambda(\hat{\mathbf{R}})\mathcal{P}_{\Lambda'}(\hat{\mathbf{R}}) = \sum_{\Lambda''} \mathcal{E}(\Lambda'') \mathcal{G}^{\Lambda\Lambda'\Lambda''} \mathcal{P}_{\Lambda''}(\hat{\mathbf{R}}) \quad (14)$$

(Cahn & Slepian 2020, §6), where $\mathcal{G}^{\Lambda\Lambda'\Lambda''}$ is the generalized Gaunt integral, defined as

$$\mathcal{G}^{\Lambda\Lambda'\Lambda''} = \int d\hat{\mathbf{R}} \mathcal{P}_\Lambda(\hat{\mathbf{R}})\mathcal{P}_{\Lambda'}(\hat{\mathbf{R}})\mathcal{P}_{\Lambda''}(\hat{\mathbf{R}}). \quad (15)$$

In Appendix A, we provide expressions for $\mathcal{G}^{\Lambda\Lambda'\Lambda''}$ with $N = 3, 4$, and 5 , involving Wigner $9-j$ symbols.

2.3 Special Cases

For later use, we give the explicit forms of the basis functions for $N = 3, 4, 5$. For $N = 3$, the basis is specified by two components $\Lambda = \{\ell_1, \ell_2\}$, with

$$C_M^\Lambda = C_{m_1 m_2}^{\ell_1 \ell_2} = \frac{(-1)^{\ell_1 - m_1}}{\sqrt{2\ell_1 + 1}} \delta_{\ell_1 \ell_2}^{\mathbf{K}} \delta_{m_1 - m_2}^{\mathbf{K}}, \quad (16)$$

which enforces $\ell_1 = \ell_2$, $m_1 + m_2 = 0$.³ From (4), the full basis function can thus be written

$$\mathcal{P}_\ell(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2) = \sum_{m=-\ell}^{\ell} \frac{(-1)^{\ell-m}}{\sqrt{2\ell+1}} Y_{\ell m}(\hat{\mathbf{r}}_1) Y_{\ell-m}(\hat{\mathbf{r}}_2) = \frac{\sqrt{2\ell+1}}{4\pi} (-1)^\ell L_\ell(\hat{\mathbf{r}}_1 \cdot \hat{\mathbf{r}}_2), \quad (17)$$

where $L_\ell(\mu)$ is the Legendre polynomial of order ℓ . In this case, our basis is simply a rescaled version of a Legendre expansion into multipoles. Note that $\ell_1 + \ell_2 \equiv 2\ell$ is always even, thus all $N = 3$ basis functions are parity-even. This occurs since the $N = 3$ parity transformation is equivalent to a rotation, and the basis functions are invariant under rotations.

For $N = 4$, simplifying the coefficients of (5) gives

$$C_M^\Lambda = C_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} = (-1)^{\ell_1 + \ell_2 + \ell_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix}, \quad (18)$$

³ Note that these are equivalent to the $N = 2$ basis functions of Cahn & Slepian (2020).

depending on three principal angular momenta, $\Lambda = \{\ell_1, \ell_2, \ell_3\}$. The full basis function becomes

$$\mathcal{P}_{\ell_1 \ell_2 \ell_3}(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \hat{\mathbf{r}}_3) = (-1)^{\ell_1 + \ell_2 + \ell_3} \sum_{m_1 m_2 m_3} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix} Y_{\ell_1 m_1}(\hat{\mathbf{r}}_1) Y_{\ell_2 m_2}(\hat{\mathbf{r}}_2) Y_{\ell_3 m_3}(\hat{\mathbf{r}}_3). \quad (19)$$

As before this is just a simple combination of spherical harmonics with the coupling chosen to ensure rotational invariance. Here, the sum of the principal angular momenta, $\ell_1 + \ell_2 + \ell_3$, can be either even or odd, thus both even- and odd-parity functions are permitted.

Finally, we remark on the form for $N = 5$. Here, we require *five* sets of angular momentum coefficients to fully specify each basis function, $\Lambda = \{\ell_1, \ell_2, (\ell_{12}), \ell_3, \ell_4\}$. An intermediate angular momentum ℓ_{12} appears as there are multiple ways to combine four spherical harmonics with given principal angular momenta ℓ_i such that the total angular momentum is zero. In this case, the coupling coefficients and associated basis functions are

$$\begin{aligned} C_M^\Lambda &= C_{m_1 m_2 m_3 m_4}^{\ell_1 \ell_2 (\ell_{12}) \ell_3 \ell_4} = (-1)^{\ell_1 + \ell_2 + \ell_3 + \ell_4} \sum_{m_{12}} (-1)^{\ell_{12} - m_{12}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_{12} \\ m_1 & m_2 & -m_{12} \end{pmatrix} \begin{pmatrix} \ell_{12} & \ell_3 & \ell_4 \\ m_{12} & m_3 & m_4 \end{pmatrix}, \\ \mathcal{P}_{\ell_1 \ell_2 (\ell_{12}) \ell_3 \ell_4}(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \hat{\mathbf{r}}_3, \hat{\mathbf{r}}_4) &= (-1)^{\ell_1 + \ell_2 + \ell_3 + \ell_4} \sum_{m_{12}} (-1)^{\ell_{12} - m_{12}} \sum_{m_1 m_2 m_3 m_4} \begin{pmatrix} \ell_1 & \ell_2 & \ell_{12} \\ m_1 & m_2 & -m_{12} \end{pmatrix} \begin{pmatrix} \ell_{12} & \ell_3 & \ell_4 \\ m_{12} & m_3 & m_4 \end{pmatrix} \\ &\quad \times Y_{\ell_1 m_1}(\hat{\mathbf{r}}_1) Y_{\ell_2 m_2}(\hat{\mathbf{r}}_2) Y_{\ell_3 m_3}(\hat{\mathbf{r}}_3) Y_{\ell_4 m_4}(\hat{\mathbf{r}}_4). \end{aligned} \quad (20)$$

Higher orders proceed similarly, with the basis function of $(N - 1)$ coordinates involving $(N - 4)$ intermediate angular momenta.

3 THE NPCF ESTIMATOR

3.1 Definition

The 3D N -point correlation function is defined as a statistical average over the product of N overdensity fields, δ in some redshift bin:

$$\begin{aligned} \zeta(\mathbf{r}_1, \dots, \mathbf{r}_{N-1}) &= \langle \delta(\mathbf{s}) \delta(\mathbf{s} + \mathbf{r}_1) \dots \delta(\mathbf{s} + \mathbf{r}_{N-1}) \rangle \\ &= \int \frac{d\mathbf{s}}{V} \delta(\mathbf{s}) \delta(\mathbf{s} + \mathbf{r}_1) \dots \delta(\mathbf{s} + \mathbf{r}_{N-1}), \end{aligned} \quad (21)$$

using the ergodic theorem to convert the statistical average into a spatial integral in the second line, assuming the volume V to be large. By statistical homogeneity, the NPCF cannot depend on the absolute position \mathbf{s} , thus it is a function of only $(N - 1)$ sets of coordinates. Further, if one assumes δ to be statistically isotropic, the function must be invariant under any global rotation, *i.e.* one in which all separation vectors $\hat{\mathbf{r}}$ are simultaneously rotated about some common origin. This is exactly the condition for the function to be expressed in the isotropic basis of §2; thus we may write

$$\boxed{\zeta(\mathbf{r}_1, \dots, \mathbf{r}_{N-1}) = \sum_{\Lambda} \zeta_{\Lambda}(r_1, \dots, r_{N-1}) \mathcal{P}_{\Lambda}(\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{N-1})} \quad (22)$$

without loss of generality, where \mathcal{P}_{Λ} are the isotropic basis functions of $(N - 1)$ coordinates defined in §4. Λ specifies the basis function, and is composed of angular momentum coefficients $\ell_1, \ell_2, \ell_{12}, \dots$, which take integer values. The coefficients indicate the scales of the *internal* angles of the N -hedron of particles forming the NPCF (with ℓ_i corresponding to the side with length r_i). For $N = 3$, inserting the explicit basis function of (17) gives

$$\zeta(\mathbf{r}_1, \mathbf{r}_2) = \sum_{\ell} \frac{\sqrt{2\ell + 1}}{4\pi} (-1)^{\ell} \zeta_{\ell}(r_1, r_2) L_{\ell}(\hat{\mathbf{r}}_1 \cdot \hat{\mathbf{r}}_2), \quad (23)$$

thus the 3PCF basis coefficients ζ_{ℓ} are simply rescaled Legendre multipoles (matching Slepian & Eisenstein 2015).

The above equations involve a sum over an infinite number of Λ coefficients. However, provided that we analyze the basis coefficients directly (rather than using them to reconstruct the full ζ), it is fully valid to restrict to a subset of Λ values with $\ell_i \leq \ell_{\max}$, given that the basis functions are orthogonal (9). Here, ℓ_{\max} sets the resolution of the internal N -hedron angles; for example, restricting to $\ell_{\max} = 10$ allows us to resolve opening angles of $\approx 180^\circ/10 = 18^\circ$. Since all N points in the polyhedron are located at similar distances from the observer, there is no need for ℓ_{\max} to be particularly large; this differs from analysis of the CMB and photometric galaxy samples, where one point of the N -hedron is located at the observer. In the latter case, $(N - 1)$ side lengths are large (equal to the distance to the sample), thus the N -hedra are squeezed and a high ℓ_{\max} is typically required.

As noted in §2.2, the basis functions come in two forms; parity-odd (with even $\ell_1 + \dots + \ell_{N-1}$) and parity-even (with odd $\ell_1 + \dots + \ell_{N-1}$). To understand our notation, first consider the NPCF under parity transforms:

$$\begin{aligned} \mathbb{P}[\zeta(\mathbf{r}_1, \dots, \mathbf{r}_{N-1})] &= \zeta(-\mathbf{r}_1, \dots, -\mathbf{r}_{N-1}) \\ &= \sum_{\Lambda} \zeta_{\Lambda}(r_1, \dots, r_{N-1}) \mathcal{P}_{\Lambda}(-\hat{\mathbf{r}}_1, \dots, -\hat{\mathbf{r}}_{N-1}) \equiv \sum_{\Lambda} \mathcal{E}(\Lambda) \zeta_{\Lambda}(r_1, \dots, r_{N-1}) \mathcal{P}_{\Lambda}(\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{N-1}), \end{aligned} \quad (24)$$

where \mathbb{P} is the parity operator and we have used the basis decomposition (22) and the parity inversion formula (8) in the second line. Invoking orthogonality (9), we find that multiplets ζ_{Λ} with odd $\ell_1 + \dots + \ell_{N-1}$ are odd under parity-inversion, whilst those with even $\ell_1 + \dots + \ell_{N-1}$

are even under parity-inversion. Furthermore, given that the odd-parity basis functions are imaginary (§2.2), this implies that all ζ_Λ with odd-parity Λ are also purely imaginary (as found in Appendix B). Assuming the NPCF to be parity-symmetric, we expect ζ_Λ to vanish for odd $\ell_1 + \dots + \ell_{N-1}$. We further note that odd-parity terms appear only for $N > 3$, as noted in §2.2. This occurs since parity inversion is equivalent to a 3D rotation for $N \leq 3$, and the isotropic NPCF is rotationally invariant.

Finally, we note that the general NPCF is a sum of two contributions: ‘connected’ and ‘disconnected’. The first involves the joint probability of N points in space, whilst the second factorizes into multiple independent components. For example, the cosmological 4PCF contains contributions both from the intrinsic non-Gaussian 4PCF (the connected piece), and the product of two 2PCFs (the disconnected piece). In full:

$$\begin{aligned} \zeta(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) &= \langle \delta(\mathbf{s})\delta(\mathbf{s} + \mathbf{r}_1)\delta(\mathbf{s} + \mathbf{r}_2)\delta(\mathbf{s} + \mathbf{r}_3) \rangle_c + [\langle \delta(\mathbf{s})\delta(\mathbf{s} + \mathbf{r}_1) \rangle \langle \delta(\mathbf{s} + \mathbf{r}_2)\delta(\mathbf{s} + \mathbf{r}_3) \rangle + 2 \text{ perms.}] \\ &= \zeta_c(\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) + [\xi(\mathbf{r}_1)\xi(\mathbf{r}_2 - \mathbf{r}_3) + 2 \text{ perms.}], \end{aligned} \quad (25)$$

where the subscript ‘c’ indicates a connected contribution and $\xi(\mathbf{r})$ is the 2PCF. Similarly, the 6PCF contains a connected non-Gaussian term and disconnected terms involving (a) two 3PCFs, (b) a 4PCF and a 2PCF, and (c) three 2PCFs. A detection of the NPCF signal is thus not necessarily a detection of non-Gaussianity; most of the signal-to-noise is expected to be in the disconnected contribution, which is degenerate with the 2PCF for even N . To measure the *connected* part we must subtract off the disconnected contribution either at the estimator-level or using some theory model; this will be discussed in depth in Philcox et al. (2021a) for the 4PCF.

3.2 General Estimator

To extract the NPCF multiplets $\zeta_\Lambda(r_1, \dots, r_{N-1})$, we may use the orthogonality relation (10):

$$\begin{aligned} \zeta_\Lambda(r_1, \dots, r_{N-1}) &= \int d\hat{\mathbf{r}}_1 \dots d\hat{\mathbf{r}}_{N-1} \zeta(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{N-1}) \mathcal{P}_\Lambda^*(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \dots, \hat{\mathbf{r}}_{N-1}) \\ &= \mathcal{E}(\Lambda) \int d\hat{\mathbf{r}}_1 \dots d\hat{\mathbf{r}}_{N-1} \zeta(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{N-1}) \mathcal{P}_\Lambda(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \dots, \hat{\mathbf{r}}_{N-1}), \end{aligned} \quad (26)$$

using the conjugation rule (7) to obtain the second line. Inserting the NPCF definition (21) and the explicit forms of the basis functions (4) gives an estimator for ζ_Λ :

$$\hat{\zeta}_\Lambda(r_1, \dots, r_{N-1}) = \frac{1}{V} \int d\mathbf{s} \sum_{m_1 \dots m_{N-1}} \mathcal{E}(\Lambda) C_M^\Lambda \int d\hat{\mathbf{r}}_1 \dots d\hat{\mathbf{r}}_{N-1} \delta(\mathbf{s})\delta(\mathbf{s} + \mathbf{r}_1) \dots \delta(\mathbf{s} + \mathbf{r}_{N-1}) Y_{\ell_1 m_1}(\hat{\mathbf{r}}_1) \dots Y_{\ell_{N-1} m_{N-1}}(\hat{\mathbf{r}}_{N-1}), \quad (27)$$

where we have integrated the absolute position \mathbf{s} over a volume V , as in (21). Note that we have not yet introduced radial binning. The appeal of this decomposition is clear; the estimator is *exactly separable in $\hat{\mathbf{r}}_i$* . Explicitly, we may define

$$a_{\ell m}(\mathbf{s}; r) = \int d\hat{\mathbf{r}} \delta(\mathbf{s} + \mathbf{r}) Y_{\ell m}(\hat{\mathbf{r}}) \quad (28)$$

(which is simply a convolution of δ and $Y_{\ell m}$), and write (27) as

$$\hat{\zeta}_\Lambda(r_1, \dots, r_{N-1}) = \frac{1}{V} \int d\mathbf{s} \delta(\mathbf{s}) \sum_{m_1 \dots m_{N-1}} \mathcal{E}(\Lambda) C_M^\Lambda a_{\ell_1 m_1}(\mathbf{s}; r_1) \dots a_{\ell_{N-1} m_{N-1}}(\mathbf{s}; r_{N-1}). \quad (29)$$

Given the harmonic coefficients $a_{\ell m}(\mathbf{s}; r)$, the NPCF is formed by summing over the m indices, weighted by the coupling coefficients $\mathcal{E}(\Lambda) C_M^\Lambda$, and integrating over position \mathbf{s} . Restricting to $\ell_i \leq \ell_{\max}$ (and thus bounding the resolution of the internal N -hedron angles), an upper bound to the number of summation terms is given by $(1 + 2\ell_{\max})^{N-2}$. Whilst (29) is unbiased (as shown by inserting 21 and 22, then using orthonormality), it is not strictly optimal, *i.e.* it is not minimum variance. Optimal estimation of the higher-point functions is possible (and discussed for example in Smith & Zaldarriaga (2011); Smith et al. (2015); Fergusson & Shellard (2011); Fergusson et al. (2012); Philcox (2021a) for Fourier-space statistics); this generally involves extra additive terms. We do not explore such complications in this work.

In practice, one usually bins the NPCFs in radius. To do so, we introduce a binning function $\Theta^b(r)$ which is unity if r is in radial bin b and zero else. In this case, the radially-averaged NPCF becomes

$$\begin{aligned} \zeta_\Lambda^B &= \frac{1}{v_B} \int \left[\prod_{i=1}^{N-1} r_i^2 dr_i \Theta^{b_i}(r_i) \right] \zeta_\Lambda(r_1, \dots, r_{N-1}), \\ v_B &= \int \left[\prod_{i=1}^{N-1} r_i^2 dr_i \Theta^{b_i}(r_i) \right], \end{aligned} \quad (30)$$

where the second line gives the bin volume. We denote the set of all radial bin indices by $B = \{b_1, \dots, b_{N-1}\}$, with b_i referring to the N -hedron side with radius r_i and thus angular momentum coefficient ℓ_i . Since the binning is separable in r_i , the corresponding NPCF estimator may be written

$$\hat{\zeta}_\Lambda^B = \frac{1}{V} \int d\mathbf{s} \delta(\mathbf{s}) \sum_{m_1 \dots m_{N-1}} \mathcal{E}(\Lambda) C_M^\Lambda a_{\ell_1 m_1}^{b_1}(\mathbf{s}) \dots a_{\ell_{N-1} m_{N-1}}^{b_{N-1}}(\mathbf{s}), \quad (31)$$

with the new harmonic coefficients

$$a_{\ell m}^b(\mathbf{s}) = \frac{1}{v_b} \int d\mathbf{r} \delta(\mathbf{s} + \mathbf{r}) Y_{\ell m}(\hat{\mathbf{r}}) \Theta^b(r), \quad (32)$$

where $v_b \equiv \int d\mathbf{r} \Theta^b(r)$ is the bin volume. This is again a simple convolution of δ and $Y_{\ell m} \Theta^b$. In this approach, binning has an important function; it allows us to estimate only a finite number of $a_{\ell m}^b(\mathbf{s})$ harmonic coefficients, before performing the summation of (31).

3.2.1 Discrete Data

In most contexts, we are interested in the estimation of NPCFs from discrete data. Given a set of N_g particle positions at $\{\mathbf{r}_i\}$, the discrete density field may be written as a sum of Dirac delta functions and weights w_i :

$$\delta(\mathbf{r}) = \sum_{i=1}^{N_g} w_i \delta_{\text{D}}(\mathbf{r} - \mathbf{r}_i). \quad (33)$$

Using this, the \mathbf{s} integral in the NPCF estimator becomes a sum over particle positions as

$$\hat{\xi}_{\Lambda}^B = \frac{1}{V} \sum_i w_i \sum_{m_1 \dots m_{N-1}} \mathcal{E}(\Lambda) C_M^{\Lambda} a_{\ell_1 m_1}^{b_1}(\mathbf{s}_i) \dots a_{\ell_{N-1} m_{N-1}}^{b_{N-1}}(\mathbf{s}_i), \quad (34)$$

and the harmonic coefficients $a_{\ell m}^b$ can be written

$$a_{\ell m}^b(\mathbf{s}_i) = \frac{1}{v_b} \sum_j w_j Y_{\ell m}(\widehat{\mathbf{r}_j - \mathbf{s}_i}) \Theta^b(|\mathbf{r}_j - \mathbf{s}_i|). \quad (35)$$

Note that the $a_{\ell m}^b$ harmonic coefficients need to be evaluated only at the locations of the particles, $\{\mathbf{r}_i\}$. The interpretation of (35) is straightforward; for a given primary position \mathbf{s}_i , we sum over all secondary positions \mathbf{r}_j which lie within bin b of \mathbf{s}_i , weighted by the spherical harmonics evaluated at the separation vector $\mathbf{r}_j - \mathbf{s}_i$. Practically this is evaluated as a weighted sum of pairs of particles, *i.e.* a *pair-count*.

The full procedure for estimating the binned NPCF from a discrete set of particles is thus:

- (i) For a given primary particle i , compute $a_{\ell m}^b(\mathbf{s}_i)$ as a weighted sum of spherical harmonics.
- (ii) For each Λ multiplet and bin combination B , sum the product of $(N-1)$ harmonic coefficients over all m_i , weighting by the coupling coefficient $\mathcal{E}(\Lambda) C_M^{\Lambda}$.
- (iii) Repeat for each primary particle and sum.

Since it is a sum over secondary positions \mathbf{r}_j , computation of (35) scales as N_g for total particle number N_g . Given that each $a_{\ell m}^b$ function must be evaluated at the location of each particle, the algorithm has complexity $O(N_g^2)$. This is the same as for the 3PCF (Slepian & Eisenstein 2015), and *much* better than the naive $O(N_g^N)$ complexity expected from counting every possible N -tuple of particles and assigning them to bins. Additional details of the computational scalings can be found in §5.1.

For discrete data, an important subtlety should be noted. If the bins $B = \{b_1, \dots, b_{N-1}\}$ are not disjoint (*i.e.* we include scenarios in which multiple side-lengths of the N -hedron fall into the same radial bin), the same secondary particle can contribute to multiple $a_{\ell m}^b$ harmonic coefficients. Physically this corresponds to a scenario in which we have an N -point function with multiple coincident points. As discussed in Philcox (2021b), these ‘self-count’ contributions can be removed by subtracting a hierarchy of lower N -point functions.⁴ However, given that these NPCF regimes are difficult to model and interpret, a simpler approach is to enforce that all b_i radial bins are distinct, effectively enforcing some minimum separation on the side lengths of the N -hedron. In this work, we assume $b_1 < b_2 < \dots < b_{N-1}$ for this reason.

3.2.2 Gridded Data

Estimator (31) may be similarly applied to data discretized onto a grid, for example a matter density field in an N -body simulation, or some fluid in a hydrodynamic simulation. If the density field δ is defined on some grid, we may compute the harmonic coefficients using a Fast Fourier Transform (FFT):

$$a_{\ell m}^b(\mathbf{s}) = \mathcal{F}^{-1} \left[\mathcal{F}[\delta](\mathbf{k}) \mathcal{F} \left[Y_{\ell m}(\hat{\mathbf{r}}) \Theta^b(r) \right](\mathbf{k}) \right](\mathbf{s}) \left/ \left[\int d\mathbf{r} \Theta^b(r) \right] \right., \quad (36)$$

where \mathcal{F} (\mathcal{F}^{-1}) is the forward (reverse) Fourier operator. Assuming N_{FFT} grid-points, this has complexity $O(N_{\text{FFT}} \log N_{\text{FFT}})$ and allows estimation of $a_{\ell m}^b$ in all grid-cells simultaneously.⁵ Practically, we may use (31) to compute the NPCF of a gridded field δ in the following fashion:

⁴ Specifically, we must subtract off basis functions with two coincident points, e.g., $\mathcal{P}_{\Lambda}(\mathbf{r}_1, \mathbf{r}_1, \dots, \mathbf{r}_{N-2})$, those with three coincident points, e.g., $\mathcal{P}_{\Lambda}(\mathbf{r}_1, \mathbf{r}_1, \mathbf{r}_1, \dots, \mathbf{r}_{N-3})$, *et cetera*. Using the contraction formulae of Cahn & Slepian (2020, §5), these can be written as lower-order basis functions.

⁵ We caution that this approach requires many $a_{\ell m}^b(\mathbf{s})$ fields must be stored in order to compute (31); this results in significant memory usage.

- (i) For each ℓ, m up to some ℓ_{\max} and radial bin b , estimate $a_{\ell m}^b(\mathbf{s})$ using FFTs.
- (ii) Take the product of $N - 1$ $a_{\ell m}^b$ fields and integrate over space, weighted by $\delta(\mathbf{s})$.
- (iii) Sum over $M = \{m_1, \dots, m_{N-1}\}$, weighted by the coupling coefficients $\mathcal{E}(\Lambda) C_M^\Lambda$.

This generalizes the approach of [Portillo et al. \(2018\)](#), which computed the 3PCF of gridded data via Fourier methods. Whilst this approach may prove useful for the analysis of simulations and applications beyond cosmology, we focus on discrete data for the remainder of this work, though much of the discussion applies to both cases.

3.3 Edge Correction

In the analysis of real galaxy surveys, one does not have direct access to the galaxy overdensity field δ . Instead, our observables are the set of N_D galaxy positions $\{\mathbf{r}_i\}$ (hereafter ‘data’), as well as a set of N_R unclustered particles (hereafter ‘randoms’) $\{\mathbf{r}_j\}$, with the randoms distributed according to the survey geometry and selection function. Mathematically, the associated random fields are given by

$$D(\mathbf{r}) = \sum_{i=1}^{N_D} w_i^D \delta_D(\mathbf{r} - \mathbf{r}_i), \quad R(\mathbf{r}) = \sum_{j=1}^{N_R} w_j^R \delta_D(\mathbf{r} - \mathbf{r}_j), \quad (37)$$

where w_i^D and w_j^R are data and random weights, with the latter normalized such that $\sum_i w_i^D = \sum_j w_j^R$. Averaging over the (Poissonian) sampling distribution, the data and random fields have expectations $\langle D(\mathbf{r}) \rangle = n(\mathbf{r})[1 + \delta(\mathbf{r})]$, $\langle R(\mathbf{r}) \rangle = n(\mathbf{r})$ where $n(\mathbf{r})$ is some background number density. This motivates the following estimator for the galaxy NPCF, based on [Landy & Szalay \(1993\)](#) and [Szapudi & Szalay \(1998\)](#):

$$\hat{\zeta}(\mathbf{r}_1, \dots, \mathbf{r}_{N-1}) = \frac{\int d\mathbf{s} [D - R](\mathbf{s}) [D - R](\mathbf{s} + \mathbf{r}_1) \dots [D - R](\mathbf{s} + \mathbf{r}_{N-1})}{\int d\mathbf{s} R(\mathbf{s}) R(\mathbf{s} + \mathbf{r}_1) \dots R(\mathbf{s} + \mathbf{r}_{N-1})} \equiv \frac{\mathcal{N}(\mathbf{r}_1, \dots, \mathbf{r}_{N-1})}{\mathcal{R}(\mathbf{r}_1, \dots, \mathbf{r}_{N-1})}, \quad (38)$$

defining the functions \mathcal{N} and \mathcal{R} . In the shot-noise limit, this is simply an inverse-variance weighting (cf. [Slepian & Eisenstein 2015](#)). For a simple periodic-box geometry in the limit of infinite randoms, $R(\mathbf{r})$ takes the constant value \bar{n} everywhere, thus the denominator is trivial. Note that (38) applies also to gridded density fields with non-uniform geometry; in this case \mathcal{N}_Λ and \mathcal{R}_Λ are the multiplets of the *gridded* ‘data-minus-random’ and random density fields.

In general, we can expand each function appearing in (38) in the isotropic basis of §2, giving

$$\sum_{\Lambda} \hat{\zeta}_{\Lambda}(R) \mathcal{P}_{\Lambda}(\hat{\mathbf{R}}) = \frac{\sum_{\Lambda''} \mathcal{N}_{\Lambda''}(R) \mathcal{P}_{\Lambda''}(\hat{\mathbf{R}})}{\sum_{\Lambda'} \mathcal{R}_{\Lambda'}(R) \mathcal{P}_{\Lambda'}(\hat{\mathbf{R}})}, \quad (39)$$

where we have denoted $\hat{\mathbf{R}} \equiv \{\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{N-1}\}$ and $R \equiv \{r_1, \dots, r_{N-1}\}$. Defining $f_{\Lambda}(R) \equiv \mathcal{R}_{\Lambda}(R)/\mathcal{R}_0(R)$, we can write

$$\begin{aligned} \sum_{\Lambda''} \frac{\mathcal{N}_{\Lambda''}(R)}{\mathcal{R}_0(R)} \mathcal{P}_{\Lambda''}(\hat{\mathbf{R}}) &= \sum_{\Lambda\Lambda'} \hat{\zeta}_{\Lambda}(R) f_{\Lambda'}(R) \mathcal{P}_{\Lambda}(\hat{\mathbf{R}}) \mathcal{P}_{\Lambda'}(\hat{\mathbf{R}}) \\ &= \sum_{\Lambda\Lambda'} \hat{\zeta}_{\Lambda}(R) f_{\Lambda'}(R) \left[\sum_{\Lambda''} \mathcal{E}(\Lambda'') \mathcal{G}^{\Lambda\Lambda'\Lambda''} \mathcal{P}_{\Lambda''}(\hat{\mathbf{R}}) \right], \end{aligned} \quad (40)$$

where we have expanded the product of two isotropic basis functions using (14). This uses the generalized Gaunt integrals $\mathcal{G}^{\Lambda\Lambda'\Lambda''}$ defined in (15), with explicit low-order forms given in terms of Wigner 3- j and 9- j symbols in Appendix A. Invoking orthonormality (9), we can extract the $\mathcal{P}_{\Lambda''}(\hat{\mathbf{R}})$ basis coefficient as

$$\frac{\mathcal{N}_{\Lambda''}(R)}{\mathcal{R}_0(R)} = \sum_{\Lambda} \hat{\zeta}_{\Lambda} \left[\sum_{\Lambda'} f_{\Lambda'}(R) \mathcal{E}(\Lambda'') \mathcal{G}^{\Lambda\Lambda'\Lambda''} \right] \equiv \sum_{\Lambda} \hat{\zeta}_{\Lambda}(R) M_{\Lambda\Lambda''}(R). \quad (41)$$

In the second line we have noted that this is simply a linear equation with (radial-bin dependent) coupling matrix $M_{\Lambda\Lambda''}$; inverting yields the *edge-corrected* estimator for $\hat{\zeta}_{\Lambda}$:

$$\hat{\zeta}_{\Lambda}(r_1, \dots, r_{N-1}) = \sum_{\Lambda''} M_{\Lambda\Lambda''}^{-1}(r_1, \dots, r_{N-1}) \frac{\mathcal{N}_{\Lambda''}(r_1, \dots, r_{N-1})}{\mathcal{R}_0(r_1, \dots, r_{N-1})}, \quad (42)$$

writing out R in full. This applies also to the binned statistic with the substitutions $\mathcal{N}_{\Lambda}(R) \rightarrow \mathcal{N}_{\Lambda}^B$, $\mathcal{R}_{\Lambda}(R) \rightarrow \mathcal{R}_{\Lambda}^B$, $\zeta_{\Lambda}(R) \rightarrow \zeta_{\Lambda}^B$, as can be seen from first rewriting (38) in binned form.

The meaning of (42) can be understood as follows. For a uniform periodic geometry (e.g., an N -body simulation or some gridded periodic data), $\mathcal{R}(\mathbf{r}_1, \dots, \mathbf{r}_{N-1})$ is independent of positions \mathbf{r}_i , thus its only non-zero basis coefficient is

$$\mathcal{R}_0(r_1, \dots, r_{N-1}) = \frac{\bar{n}^N}{V} \int d\mathbf{s} \int d\hat{\mathbf{r}}_1 \dots d\hat{\mathbf{r}}_{N-1} \mathcal{P}_0(\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_{N-1}) = (4\pi)^{-(N-1)/2} \bar{n}^N, \quad (43)$$

using (11). In this instance, $f_{\Lambda'} \propto \delta_{\Lambda'0}^K$, and $M_{\Lambda\Lambda''} \propto \delta_{\Lambda\Lambda''}^K$, hence the matrix multiplication is trivial. In general, a non-uniform survey geometry

FOURPCF	Compute the 4PCF
FIVEPCF	Compute the 5PCF
SIXPCF	Compute the 6PCF
ALLPARITY	Include both odd- and even-multiplets, rather than just even
DISCONNECTED	Compute the disconnected (Gaussian) 4PCF
PERIODIC	Implement periodic boundary conditions (applicable for N -body simulations)
AVX	Accelerate the $a_{\ell m}^b$ computation (32) using Advanced Vector Extensions (AVX) instruction sets
OPENMP	Parallelize the main computation (steps 3-12 in Algorithm 1) using OPENMP on as many threads as are available
GPU	Perform the $a_{\ell m}^b$ summation (Algorithm 3) on a Graphics Processing Unit (GPU) rather than a CPU

Table 1. The various computation modes available in the *ENCORE* NPCF package. These allow control over which statistics are computed, as well as architecture choices. Any combination of these can be used in tandem, except for OPENMP and GPU. In the ALLPARITY mode, we include contributions from odd-parity spectra, *i.e.* those with odd $\sum_i \ell_i$. Assuming the Universe to be parity symmetric, these are unphysical. In DISCONNECTED mode, we estimate the Gaussian contribution to the 4PCF; this allows direct estimation of the purely non-Gaussian connected 4PCF, and will be discussed in Philcox et al. (2021a).

induces a coupling between multipliers, implying that the naïve estimator $\mathcal{N}_\Lambda/\mathcal{R}_0$ is insufficient.⁶ In this case, the matrix $M_{\Lambda\Lambda'}$ corrects for the effect of the survey window function, ensuring that $\hat{\zeta}_\Lambda$ is unbiased.

The edge-correction equation (42) is straightforward to implement; one first measures the multipliers \mathcal{N}_Λ and \mathcal{R}_Λ using the estimators of (29) or (34), with δ replaced by $[D - R](\mathbf{r})$ or $R(\mathbf{r})$. Given the set of all measured multipliers, the matrix inversion is a straightforward operation, and may be done for each radial bin combination separately. In practice, $M_{\Lambda\Lambda'}(R)$ is close to diagonal (cf. Fig. 4), thus truncation of the infinite matrix $M_{\Lambda\Lambda'}(R)$ at finite ℓ_{\max} is expected to yield an accurate estimate of $\zeta_\Lambda(R)$ for all multipliers containing ℓ near ℓ_{\max} . In practice, this is ensured by measuring \mathcal{N}_Λ and \mathcal{R}_Λ using an ℓ_{\max} that is one larger than that used for the output NPCF.⁷ We additionally note that $\mathcal{G}^{\Lambda\Lambda'}$ is only non-zero if $\sum_i \ell_i + \sum_j \ell'_j + \sum_k \ell''_k$ is even. This implies that a parity-odd NPCF can be sourced by a parity-even \mathcal{R} and a parity-odd \mathcal{N} or *vice versa*, whilst a parity-even NPCF requires both the data and random counts to have the same parity (cf. §3.1).

4 IMPLEMENTATION

Below, we discuss the implementation of the $O(N_g^2)$ discrete NPCF estimator given in (34). Discussion of the gridded NPCF estimator (which is of use for hydrodynamic simulations and gridded cosmological simulations) will follow in future work, building upon the 3PCF code of Portillo et al. (2018). The discrete estimator is implemented within the code package, *ENCORE*, publicly available at github.com/oliverphilcox/encore, and is loosely based on the isotropic 3PCF code of Slepian & Eisenstein (2015) and the HIPSTER power spectrum estimator package of Philcox & Eisenstein (2020); Philcox (2021b). *ENCORE* is written in C++ and CUDA and includes routines for computing the NPCF counts using both parallelized CPU resources and GPUs, as well as computing the edge-corrections as in §3.3. Currently, the code supports computation of the isotropic 2-, 3-, 4-, 5- and 6-point functions for both odd- and even-parity multipliers. Higher N and ℓ_{\max} can be straightforwardly included, though we note that the C_M^Λ matrix becomes sizeable for large ℓ_{\max} (scaling as $(1 + \ell_{\max})^{N-1} (1 + 2\ell_{\max})^{2N-6}$, cf. §5.1). The code can be run in various modes, as shown in Table 1 and discussed in greater detail below.

4.1 Main Algorithm

4.1.1 Overview

The main process involved in computing NPCFs from a galaxy survey dataset is the estimation of multipoles, \mathcal{N}_Λ^B and \mathcal{R}_Λ^B , from a discrete set of particle positions. As in (34), this may be framed as a weighted pair-count, and an overview of the procedure is given in Algorithm 1. Essentially there are three processes of interest: (1) computing the $a_{\ell m}^b(\mathbf{s}_i)$ harmonic coefficients at the location of each primary particle i , (2) accumulating the contributions to ζ_Λ^B , and (3) input/output operations, including loading the particles and coupling coefficients and saving the NPCF estimates. The latter is straightforward; the only complexity is in computation of the C_M^Λ coupling coefficients (defined in 5, with explicit forms given in 16, 18 & 20). Since the determination of C_M^Λ requires Wigner 3- j manipulations, they are precomputed in PYTHON using the SYMPY package⁸ and loaded at runtime.

Computing the $a_{\ell m}^b(\mathbf{s}_i)$ harmonic coefficients at each primary particle location \mathbf{s}_i is less trivial. From (35), each can be evaluated as a spherical-harmonic-weighted sum of all secondary particles, j , whose separation from the primary lies within bin b . To implement this, we use the approach of Slepian & Eisenstein (2015), whereupon spherical harmonics are obtained from their Cartesian representations. For each secondary particle j , we first define the unit separation vector $\widehat{\mathbf{r}}_{j - \mathbf{s}_i} \equiv \hat{\mathbf{A}} = (\Delta_x, \Delta_y, \Delta_z)$, and compute all (non-negative integer) powers

⁶ The naïve estimator is analogous to the window-convolved galaxy power spectrum estimator of Feldman et al. (1994), or the CMB *pseudo- C_ℓ* estimator.

⁷ An additional caveat concerns *anisotropy* of the window function. Strictly speaking, couplings between the anisotropic window function and anisotropic galaxy distribution can generate an isotropic NPCF, and ought to be removed via an anisotropic edge-correction equation. This is analogous to the mixing of power spectrum monopole and quadrupole in the presence of an anisotropic survey window. Here, we assume this effect to be small, such that the bias in our estimates of the isotropic functions can be ignored.

⁸ www.sympy.org

Algorithm 1 Overview of the main algorithm underlying computation of the N -point correlation functions. Given precomputed coupling coefficients C_M^Λ (5), this routine computes the correlation function estimates (with $N > 2$) from N_g particles with angular momentum coefficients up to $\ell = \ell_{\max}$ using N_r radial bins with minimum (maximum) separation r_{\min} (r_{\max}). This is the main procedure underlying the ENCORE code and implements the estimator of (34). For windowed data, this approach is applied to two datasets: a field of random particles (which define the window) and one of ‘data-minus-randoms’. These are combination via (42).

```

1: Load particle positions and weights from disk
2: Read-in precomputed coupling coefficients  $C_M^\Lambda$  defined in (5)
3: for primary particle  $i = 1, \dots, N_g$  do
4:   for secondary particle  $j = 1, \dots, N_g$  with  $j \neq i$  do
5:     if  $r_{\min} < |\mathbf{r}_i - \mathbf{r}_j| < r_{\max}$  then
6:       for  $\ell \leq \ell_{\max}$  do
7:         for  $m = 0, \dots, \ell$  do
8:           for  $b = 1, \dots, N_r$  do
9:             Accumulate the contribution to the harmonic coefficient  $a_{\ell m}^b(\mathbf{s}_i)$  using (32) (cf. Algorithm 2)
10:          end for
11:        end for
12:      end for
13:    end if
14:  end for
15:  Compute the NPCF contribution as the product of  $(N - 1)$  harmonic coefficients at for each multiplet  $\Lambda$  and set of bin indices  $B$ ,
  summed over  $\{m_i\}$  (cf. 34 & Algorithm 3).
16: end for
17: Save the NPCF counts to disk

```

$\Delta_x^p \Delta_y^q \Delta_z^r$ where $p + q + r \leq \ell_{\max}$. These can be combined to form the spherical harmonics. In practice, we compute histograms of $\Delta_x^p \Delta_y^q \Delta_z^r$ for all required j particles (skipping those with separations outside $[r_{\min}, r_{\max}]$), before combining these to form the harmonic coefficients. This requires $\mathcal{O}(N_g^2)$ operations and a total of $(1 + \ell_{\max})(2 + \ell_{\max})(3 + \ell_{\max})/6$ histograms per bin b .

As an example, consider $a_{2-1}^b(\mathbf{s}_i)$. This requires the spherical harmonic $Y_{2-1}(\hat{\mathbf{A}}) \propto (\Delta_x - i\Delta_y)\Delta_z$; to compute it, we first obtain the sum of $\Delta_x\Delta_z$ and $\Delta_y\Delta_z$ over all secondary galaxies within the bin b , then combine the sums and normalize to obtain $a_{2-1}^b(\mathbf{s}_i)$. *Pseudo-code* for this procedure is shown in Algorithm 2. Due to the spherical harmonic conjugation symmetry, $a_{\ell-m}^b = a_{\ell m}^{b*}$, allowing us to compute only harmonic coefficients with $m \geq 0$.

Algorithm 2 Procedure to calculate the $a_{\ell m}^b(\mathbf{s}_i)$ harmonic coefficients (32) for a given primary particle (either a galaxy or random) at position \mathbf{s}_i . The algorithm relies on two facts: (1) $a_{\ell m}^b(\mathbf{s}_i)$ is a sum over spherical harmonics $Y_{\ell m}(\widehat{\mathbf{r}_j - \mathbf{s}_i})$ for secondary \mathbf{r}_j within bin b , and (2) the spherical harmonics can be written as a sum of products of their Cartesian components. In the first part of the algorithm, we accumulate histograms of all relevant unit powers (here denoted f_{pqr}), before transforming these into spherical harmonics. In practice, the first loop is implemented using custom Advanced Vector Extension (AVX) instruction sets (Slepian & Eisenstein 2015), and operates on sets of eight secondary particles simultaneously.

```

1: for secondary particle  $j = 1, \dots, N_g$  do
2:   Compute the unit separation vector  $(\Delta_x, \Delta_y, \Delta_z) = (\mathbf{s}_i - \mathbf{r}_j)/|\mathbf{s}_i - \mathbf{r}_j|$ 
3:   Find the radial bin  $b$  corresponding to separation  $|\mathbf{s}_i - \mathbf{r}_j|$ 
4:   Accumulate all combinations  $f_{pqr}^b += w_j \Delta_x^p \Delta_y^q \Delta_z^r$  with  $p + q + r \leq \ell_{\max}$  for secondary weight  $w_j$ 
5: end for
6: for  $\ell = 0, \dots, \ell_{\max}$  do
7:   for  $m = 0, \dots, \ell$  do
8:     Construct  $a_{\ell m}^b(\mathbf{s}_i)$  as a weighted sum of  $f_{pqr}^b$  coefficients
9:   end for
10: end for

```

Finally, we must combine the $a_{\ell m}^b(\mathbf{s}_i)$ harmonic coefficients to form the NPCF contributions. This is done using (34), and must be performed N_g times. In essence, this is just an outer product of $(N - 1)$ harmonic coefficients, summed over the m_i indices and weighted by the coupling matrix $\mathcal{E}(\Lambda) C_M^\Lambda$. For large N , evaluation of this expression becomes expensive, since the dimensionality of Λ , M and B is large (cf. §3.2.1). For each of element in ζ_Λ^B (of which there are at most $\binom{N_r}{N-1}(1 + \ell_{\max})^{N-1}(1 + 2\ell_{\max})^{N-4}$ for $N > 3$), we must sum over a maximum of $(1 + \ell_{\max})^{N-2} M$ combinations, requiring a large number of operations to be performed for each primary particle.

In practice, this may be expedited. Firstly, we omit any Λ contributions violating the triangle conditions, (e.g., for $N > 3$, we can enforce $|\ell_1 - \ell_2| \leq \ell_{12} \leq \ell_1 + \ell_2$), and note that the C_M^Λ conditions require $m_{N-1} = -m_1 - \dots - m_{N-2}$, as well as $|m_{12}| \leq \ell_{12}$, $|m_{123}| \leq \ell_{123}$ *et cetera*.

Usually, we compute only correlators obeying $\zeta(\hat{\mathbf{R}}) = \zeta(-\hat{\mathbf{R}})$, *i.e.* those with $\mathcal{E}(\Lambda) = 1$. These correspond to even $\sum_i \ell_i$. Finally, we can use the conjugate symmetry of the harmonic coefficients to sum only terms with $m_{N-1} \geq 0$. This is detailed in Appendix B, and corresponds to rewriting the single-particle contribution to (34) as

$$w_i \sum_{m_1 \dots m_{N-1}} \mathcal{E}(\Lambda) C_M^\Lambda a_{\ell_1 m_1}^{b_1}(\mathbf{s}_1) \dots a_{\ell_{N-1} m_{N-1}}^{b_{N-1}}(\mathbf{s}_i) = 2w_i \sum_{m_1 + \dots + m_{N-2} \leq 0} S(m_{N-1}) \mathcal{E}(\Lambda) C_M^\Lambda \mathbb{Q}_\Lambda \left[a_{\ell_1 m_1}^{b_1}(\mathbf{s}_1) \dots a_{\ell_{N-1} m_{N-1}}^{b_{N-1}}(\mathbf{s}_i) \right] \quad (44)$$

where $S(m) = 1/2$ if $m = 0$ and unity else, and the operator \mathbb{Q}_Λ returns the real-part if Λ is parity-even and i multiplied by the imaginary-part else. This reduces the number of terms that must be summed approximately by a factor of two, and makes manifest the reality of ζ_Λ for even-parity multiplets. Algorithm 3 gives an example of the full summation strategy for the 4PCF.

Algorithm 3 Typical example of the summations required to compute the NPCF estimator, given the $a_{\ell m}^b$ harmonic coefficients for radial bins b and angular momentum indices ℓ, m . Here, we give the *psuedo*-code for the parity-even $N = 4$ summations; higher orders proceed similarly. The 4PCF is parametrized by three angular momenta, $\Lambda = \{\ell_1, \ell_2, \ell_3\}$ (obeying the conditions $|\ell_1 - \ell_2| \leq \ell_3 \leq \ell_1 + \ell_2$ and $\ell_i \leq \ell_{\max}$) and three radial bins $B = \{b_1, b_2, b_3\}$. Note that we store only bins with $b_1 < b_2 < b_3$; any other bins are either superfluous or contain zero-lag contributions. Here, we compute only even-parity multiplets, skipping any triplet with odd $\ell_1 + \ell_2 + \ell_3$, as indicated by the ‘continue’ statement. We further note that ‘load’ indicates that a quantity held in RAM is accessed and passed to the cache. $S(m)$ is defined below (44).

```

1: for  $\ell_1 = 0, \dots, \ell_{\max}$  do
2:   for  $\ell_2 = 0, \dots, \ell_{\max}$  do
3:     for  $\ell_3 = |\ell_1 - \ell_2|, \dots, \min(\ell_1 + \ell_2, \ell_{\max})$  do
4:       if  $\ell_1 + \ell_2 + \ell_3$  is odd continue
5:       for  $m_1 = -\ell_1, \dots, \ell_1$  do
6:         for  $m_2 = -\ell_2, \dots, \ell_2$  do
7:            $m_3 = -m_1 - m_2$ 
8:           if  $m_3 > \ell_3$  or  $m_3 < 0$  continue
9:           Load coupling weight  $w_i(-1)^{\ell_1 + \ell_2 + \ell_3} C_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3}$ 
10:          for  $b_1 = 1, \dots, N_r$  do
11:            Load  $a_{\ell_1 m_1}^{b_1}$  or, if  $m_1 < 0$ ,  $(-1)^{m_1} a_{\ell_1 - m_1}^{b_1, *}$ 
12:            for  $b_2 = b_1 + 1, \dots, N_r$  do
13:              Load  $a_{\ell_2 m_2}^{b_2}$  or, if  $m_2 < 0$ ,  $(-1)^{m_2} a_{\ell_2 - m_2}^{b_2, *}$ 
14:              for  $b_3 = b_2 + 1, \dots, N_r$  do
15:                Load  $a_{\ell_3 m_3}^{b_3}$ 
16:                Accumulate 4PCF;  $\zeta_{\ell_1 \ell_2 \ell_3}^{b_1 b_2 b_3} += 2S(m_3) w_i(-1)^{\ell_1 + \ell_2 + \ell_3} C_{m_1 m_2 m_3}^{\ell_1 \ell_2 \ell_3} \text{Re} \left[ a_{\ell_1 m_1}^{b_1} a_{\ell_2 m_2}^{b_2} a_{\ell_3 m_3}^{b_3} \right]$ .
17:              end for
18:            end for
19:          end for
20:        end for
21:      end for
22:    end for
23:  end for
24: end for

```

4.1.2 CPU Implementation

ENCORE provides a C++ implementation of the above routine. Whilst most of the code is straightforward, there are several points of note. Firstly, when computing the harmonic coefficients for each primary particle, we optionally employ custom Advanced Vector Extension (AVX) instruction sets to accumulate the necessary unit vector powers $\Delta_x^p \Delta_y^q \Delta_z^r$ (cf. §4.1.1). These were developed for the *ABACUS* simulation project (Metchnik 2009; Garrison et al. 2019, 2021) and used in the C++ implementation of the isotropic 3PCF algorithm of Slepian & Eisenstein (2015). Practically, this allows optimal vectorization, with eight secondary particles being analyzed simultaneously (assuming double precision). Whilst the $a_{\ell m}^b$ computation is not rate-limiting at large N , the addition of AVX instruction sets gives significant speed-boosts for the 3PCF (cf. §5.1).

Secondly, the code is (optionally) parallelized using *OPENMP*. For this, we distribute the iteration over N_g primary particles to as many threads as are available, parallelizing steps three to twelve of Algorithm 1. Each thread has its own value of the NPCF sum and weight functions, with the former then being combined at the conclusion of the algorithm. Only the list of particle positions is shared between threads. As shown in §5.1, the computation time scales as $1/N_{\text{threads}}$ for the higher-point functions, implying that the parallelization is almost perfect.

4.1.3 GPU Implementation

Given that the bulk of the NPCF algorithm requires simple additions and multiplications of harmonic coefficients, we expect that the performance can be significantly enhanced using Graphical Processing Units (GPUs). Modern GPUs contain thousands of individual processing cores, each of which is capable of handling ~ 1000 threads concurrently. Developers can take advantage of this by using platforms such as NVIDIA's Compute Unified Device Architecture (CUDA) to design massively parallel algorithms, resulting in significant speed gains for algorithms that are limited by computational resources rather than memory access. The `ENCORE` code includes GPU acceleration via CUDA, and it is found to be of significant use for the 5PCF and beyond (§5.3). Since the computation time of the algorithm is a strongly increasing function of N (due to the larger dimensionality of Λ and B), we focus on the step which is rate-limiting for the higher-point functions; combination of harmonic coefficients to form the NPCF summand. *i.e.* step 11 in Algorithm 1. Currently, we do not include a GPU implementation of the $a_{\ell m}^b$ estimation; this may be added in the future. We note also that the GPU code is not currently compatible with OPENMP, *i.e.* it should be run single-threaded.

To implement a GPU-based code, we must first consider what functions we wish each GPU thread to compute. Two options come to mind: (a) allow each thread to compute the ζ_{Λ}^B contribution from a single multiplet Λ , set of bin-indices B , summing over all choices of M , or (b) use a separate thread for each M element. Option (b) allows a greater number of threads to be utilized, maximizing the GPU capabilities; further, each thread must perform the same number of operations. However, since multiple M contribute to the same multiplet Λ , this requires careful consideration of memory access. CUDA requires thread blocks to be able to execute independently, in any order, and with many threads running simultaneously it is possible for a collision to occur during which two threads attempt to write to the same NPCF element at once, resulting in an undefined value. In order to avoid this error, the addition into the NPCF array ζ_{Λ}^B must be done atomically; *i.e.* only one thread can access a given NPCF element at once. This will lead to reduced performance. For option (a), each thread computes a different element of ζ_{Λ}^B thus there are no issues with atomicity, though the total number of threads utilized (equal to the dimension of the NPCF) is smaller. This approach does not achieve maximum parallelization if the dimension is small and the work per thread varies as a function of Λ .

In practice, we implement both approaches within `ENCORE` for the 4PCF and 5PCF, and compare their performance in §5.3. Schematically, the action of the GPU-parallelized code is the following:

- For each primary position i , compute $a_{\ell m}^b$ on the CPU.
- Copy the $a_{\ell m}^b$ array to the GPU.
- For each multiplet Λ and set of bin-indices B , compute the contribution to ζ on the GPU.
- After all primary particles have been processed, copy the result back to the CPU.

The GPU kernels are launched a total of N_g times. Due to the conditions on Λ , B and M , it is non-trivial to relate the GPU thread index to the individual angular momentum and bin components; for this reason we first generate a look-up table for Λ , B and (optionally) M , which is passed to the GPU on initialization. This additionally avoids any threads having to do only trivial work; for example we can ensure that any combinations that violate the triangle conditions are not included. Finally, we note that some commercially available GPUs are optimized for float-based calculations whilst others perform well using doubles. For this reason we consider three variants: (1) using double precision at all times (as for the CPU calculation), (2) using single precision only, and (3), using single precision for the $a_{\ell m}^b$ arrays, but adding to the NPCF as a double. The performance of each variant is compared in §5.3.

4.2 Reconstructing the Geometry-Corrected NPCF

§4.1 presented a detailed discussion of how to compute the NPCF multiplets for a given set of input particles. We now consider the practicalities of computing the *full* geometry-corrected NPCF, as outlined in §3.3.

For this purpose, we require two fields: the data (galaxy positions), $D(\mathbf{r})$, and randoms (unclustered particles), $R(\mathbf{r})$ (cf. 37). The latter are used to compute the normalization term \mathcal{R}_{Λ}^B appearing in (42); this requires a straightforward application of estimator (34). For the second term, \mathcal{N}_{Λ}^B , we first compute a joint ‘data-minus-randoms’ field by combining the data and random catalogs together (following Slepian & Eisenstein 2015). Randoms are assigned negative weights and normalized such that the summed weight is zero. The multiplets of this field give \mathcal{N}_{Λ}^B directly.⁹

In practice, the random catalog $R(\mathbf{r})$ contains many more particles than the data to minimize the inherent Poissonian errors. Slepian & Eisenstein (2015) and Keihänen et al. (2019) advocate splitting the random catalog into subsets, each of size $\sim 1.5 - 2\times$ that of the data, and co-adding the \mathcal{N}_{Λ}^B counts for each; if the computation time is quadratic in the number of particles, this minimizes the work required for a given

⁹ An additional method to compute \mathcal{N}_{Λ}^B is to expand the numerator of (38) into separate terms, each involving just data or randoms. Schematically, one writes

$$\mathcal{N}_{\Lambda}^B \equiv [(D - R)^N]_{\Lambda}^B = [D^N]_{\Lambda}^B - N [D^{N-1}R]_{\Lambda}^B + \dots - N [DR^{N-1}]_{\Lambda}^B + [R^N]_{\Lambda}^B, \quad (45)$$

where $[X_1 X_2 \dots X_N]_{\Lambda}^B$ are the multiplets of N input fields X_1, X_2, \dots, X_N . This is analogous to the standard procedure for the 2PCF computation (using DD , DR , and RR counts), but somewhat inefficient here, since it requires significant recomputation.

level of Poisson noise.¹⁰ Whilst appropriate for the 3PCF, the runtime of *ENCORE* when applied to higher-point statistics is usually found to scale linearly with the number of particles (cf. §5.1), thus there is no gain to such a split. It is retained in *ENCORE* for full generality however, leading to ~ 32 separate catalogs if one uses randoms with $50\times$ the galaxy density. Given that the $a_{\ell m}^b$ contributions from the data must be included in each split, the approach is somewhat inefficient; to avoid this, we include the functionality to save and reload the histogrammed Cartesian powers of positively-weighted (data) particles. Finally, we note that, since the \mathcal{R}_Λ^B counts use only positively weighted particles, their intrinsic noise is much reduced, thus it is sufficient to use only a single random partition for their computation.

Given the estimated \mathcal{R}_Λ^B and N_Λ^B quantities (averaged over the random subsets), we can compute the full edge-corrected NPCF via (42). To form the coupling matrix $M_{\Lambda\Lambda'}$ we require the generalized Gaunt integrals of (15); these involve Wigner $9-j$ symbols and can be straightforwardly precomputed for a given N and ℓ_{\max} using the `PYTHON SYMPY` package. For simplicity, the entire edge-correction procedure is performed in `PYTHON`, and involves a matrix inversion for each radial bin combination B . In the case of a periodic box geometry, the edge-correction is trivial, and we use the analytic \mathcal{R}_Λ^B functions of (43). To facilitate general use of the *ENCORE* code, we provide a `BASH` (or `SLURM`) script to automate the entire process of computing raw NPCF multiplets and applying edge-correction, allowing for the full NPCF estimates to be obtained given the data and random particle positions.

5 RESULTS

We now present results from applying the NPCF estimators of §3 to data, considering both scalings with computational resources and particle number density, as well as a demonstration of the edge-corrected 4PCF and 5PCF of realistic simulations.

5.1 Dimensionality and Scalings

Before presenting numerical results, we consider the dimension of the ζ_Λ^B NPCF statistic, which will facilitate later discussion of computational scalings. Assuming N_r radial bins per dimension, and asserting that radial bins do not overlap (cf. §3.2.1), B contains $\binom{N_r}{N-1}$ elements, which is asymptotically equal to N_r^{N-1} . The NPCF requires isotropic basis functions of $(N-1)$ coordinates; for $N > 3$ these contain $(N-1)$ principal angular momenta ℓ_i , each of which runs from 0 to ℓ_{\max} , and $(N-4)$ intermediate angular momenta $\ell_{12\dots}$, each of which runs from 0 to $2\ell_{\max}$ (cf. §2.3). This implies that the number of elements in Λ is bounded from above by $(1 + \ell_{\max})^{N-1} (1 + 2\ell_{\max})^{N-4}$; the true number is somewhat smaller due to the triangle conditions. Furthermore, we must sum over $(N-2)$ integers m_i , each of which can run from $-\ell_i$ to ℓ_i (with m_{N-1} set by $\sum_i m_i = 0$), giving a maximum of $(1 + 2\ell_{\max})^{N-2}$ summation terms in (31) for each choice of Λ and B . This further bounds the dimensionality of the coupling coefficient C_M^Λ to $(1 + \ell_{\max})^{N-1} (1 + 2\ell_{\max})^{2N-6}$. Asymptotically, the dimension of the NPCF statistic, ζ_Λ^B , scales as $N_r^{N-1} \ell_{\max}^{2N-5}$, which quickly becomes very large as N increases. Given that our eventual goal is to compute $\mathcal{O}(10)$ parameters from the statistic, it is likely that data compression methods (e.g., [Scoccimarro 2000](#); [Philcox et al. 2021b](#)) will prove useful in any higher-point NPCF analysis.

The algorithm described in §4 contains two computationally intensive steps. The runtimes of these have the following scalings:

- **Computation of $a_{\ell m}^b(\mathbf{s}_i)$:** This requires a count over all pairs of galaxies separated by distances within a spherical shell $[r_{\min}, r_{\max}]$; as such, the work scales as $N_g \times \bar{n} \left(r_{\max}^3 - r_{\min}^3 \right)$ for (galaxy or random) number density \bar{n} and total number N_g , *i.e.* it is quadratic in the total number of particles if \bar{n} is fixed. Computation of each count requires accumulation of $(1 + \ell_{\max})(2 + \ell_{\max})(3 + \ell_{\max})/6$ powers of $\Delta_x^p \Delta_y^q \Delta_z^r$ (cf. §4.1), thus runtime is cubic in ℓ_{\max} and linear in N_r . Note that this is independent of the NPCF order, N .
- **Accumulation of ζ_Λ^B :** For each Λ and B , we must sum over all M indices for $(N-1)$ harmonic coefficients. Since this must be performed for each primary position, computation time scales as $N_g \equiv \bar{n}V$. Restricting to bins with $b_1 < b_2 < \dots < b_{N-1}$, there are $\binom{N_r}{N-1}$ combinations contained in B , with $(N-2)$ independent M elements for each of $(2N-5)$ choices for Λ (as discussed above). This implies the asymptotic scaling is of the form $\ell_{\max}^{3N-7} N_r^{N-1}$, which is a strong function of ℓ_{\max} .

Given this, we expect the runtime of *ENCORE* to scale as N_g^2 for dense samples at low N (as found in [Friesen et al. 2017](#)), or N_g for higher-point functions with fewer particles, or those with a large number of bins.

To test the above, we compute the parity-even NPCF for a synthetic box of galaxies at different mean densities \bar{n} . We use a cubic box with $L_{\text{box}} = 1500h^{-1}\text{Mpc}$, populated with randomly placed galaxies of varying number density from one-tenth to ten times that of BOSS CMASS ([Alam et al. 2015](#), cf. §5.4). The pair counts are computed on a 16-core machine in `PERIODIC` mode (cf. Tab. 1), using $\ell_{\max} \in \{10, 10, 5, 3\}$ for $N \in \{3, 4, 5, 6\}$ respectively. The maximum multipole is varied between samples to keep the dimensionality of the statistics reasonable, but the number of radial bins is fixed to $N_r = 10$; here, the dimensionality of B and Λ is $\{46, 123, 215, 259\}$ and $\{6, 381, 1211, 1558\}$ respectively, incorporating all triangle and binning conditions. The total dimension is the product of these, *i.e.* approximately $\{280, 4.7 \times 10^4, 2.6 \times 10^5, 4.0 \times 10^5\}$ elements respectively. The run-times, T , of *ENCORE* for each setup are shown in Fig. 1.

Notably, the four and higher-point algorithms follow closely a $T \sim N_g$ scaling (noting that $N_g \propto \bar{n}$, since V is fixed). The 3PCF exhibits a

¹⁰ For periodic-box geometries, it is possible to partially remove the dependence on random catalogs, as discussed in [Pearson & Samushia \(2019\)](#), [Sosa Nuñez & Niz \(2020\)](#) and [Philcox \(2021b\)](#).




Figure 1. Scaling of the `ENCORE` code with the mean number density of the sample, when run in 3PCF, 4PCF, 5PCF and 6PCF mode with 10 radial bins. As described in §5.1, the galaxy sample is simply a set of uniformly distributed points in a periodic box. The scaling is close to quadratic (linear) with number density and hence particle number N_g for the 3PCF (4PCF and above), as expected. The dotted (dashed) lines show quadratic (linear) relationships, anchored through the highest number-density sample. Note that the maximum multipole, ℓ_{\max} is varied to keep the dimensionality of the output reasonable; its value is indicated by the caption in each case. All calculations are performed on sixteen 2.4 GHz Intel Skylake CPUs, parallelized using `OPENMP`. We show only results computable within 16 hours of wall-clock time.

quadratic scaling N_g^2 at high \bar{n} , though is closer to linear for low density samples. Given the above discussion on the limiting steps of each algorithm, this is as expected. Furthermore, the low- \bar{n} plateau of the 3PCF is due to other processes from Algorithm 1 that become important, such as the loading of galaxies and coupling weights into memory and the particle neighbor searches. At much larger \bar{n} , we may expect the 4PCF to scale as N_g^2 , as it becomes limited by the same processes as for the 3PCF. Due to the prohibitive $N_r^{N-1} \ell_{\max}^{3N-7}$ scaling of the ζ_{Λ}^B accumulation step, the quadratic scaling is unlikely to occur for higher-point functions at any reasonable \bar{n} .

It is also useful to quantify how well the `ENCORE` algorithm can be parallelized. To this end, we rerun the analysis used to create Fig. 1 at $\bar{n} \approx 1.5 \times 10^{-4} h^3 \text{Mpc}^{-3}$, but on a varying number of CPU threads. The results are shown in Fig. 2 and the conclusion is straightforward; the algorithm exhibits almost perfect scaling with the number of cores. This is as expected; as discussed in §4.1.2, the `OPENMP` implementation parallelizes the sum over all primary particles, each of which is independent (and holds its own copy of the NPCF sum), except for look-ups in the shared particle grid. We note that the code does not currently include Message Passing Interface (MPI) parallelization; whilst this could lead to further speed boosts and distribute the work across multiple nodes, it is expected that a more useful application of these resources would be to compute separate N_{Λ} counts (cf. 42), or the NPCF estimates of different simulations. Fig. 2 may also be used to examine the dependence of the runtime on ℓ_{\max} . The ratio of $\ell_{\max} = 10$ to $\ell_{\max} = 5$ takes the values 1.9 ± 0.1 , 11.5 ± 1.1 for the 3PCF and 4PCF; this is somewhat better than the scalings discussed above, indicating the asymptotic limit (*i.e.* the scalings for $\ell_{\max} \gg 1$) has not yet been reached.

5.2 Comparison to Brute-Force NPCF Estimators

Given the complexity of the `ENCORE` algorithm (§3 & §4), it is important to test whether the code returns accurate estimates of the N -point functions. For this purpose, we compared the NPCF multiplets output by `ENCORE` to those computed via a brute-force approach. To form such estimates, we compute ζ_{Λ}^B directly, as a sum over N_g particles weighted by the relevant basis function:

$$\zeta_{\Lambda}^B = \sum_{i_1 \dots i_N} \mathcal{P}_{\Lambda}^*(\widehat{\mathbf{r}}_{i_2} - \widehat{\mathbf{r}}_{i_1}, \dots, \widehat{\mathbf{r}}_{i_N} - \widehat{\mathbf{r}}_{i_1}) \Theta^{b_1}(|\mathbf{r}_{i_2} - \mathbf{r}_{i_1}|) \dots \Theta^{b_N}(|\mathbf{r}_{i_1} - \mathbf{r}_{i_N}|), \quad (46)$$

where \mathbf{r}_i is the position vector of particle i , and $\Theta^b(r)$ are binary binning functions, as in §3.2. This is simply the definition for the NPCF multiplets (26) for discrete data, with the addition of binning. Since it involves a sum over N particles, it has complexity $\mathcal{O}(N_g^N)$. To evaluate the angular basis functions, we use the Cartesian forms given in Cahn & Slepian (2020, Appendix A); for example, $\mathcal{P}_{110}(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \hat{\mathbf{r}}_3) = -\sqrt{3}(4\pi)^{-3/2}(\hat{\mathbf{r}}_1 \cdot \hat{\mathbf{r}}_2)$. This approach is prohibitively slow for large N_g , thus we use only $N_g = 20$ as a test case. In our case, the two approaches agreed to machine precision for all Λ considered and N in $\{3, 4, 5, 6\}$. Since this approach knows nothing about spherical harmonics or angular momenta, the precise agreement affords us confidence in the validity of `ENCORE`.

5.3 GPU Testing

We now consider the test `ENCORE`'s GPU implementation. As in §4.1.3, GPUs are currently used to expedite the accumulation of the ζ_{Λ}^B arrays for each primary particle, given $a_{\ell m}^b(s_i)$ coefficients computed on the CPU. Two approaches are possible: (a) each GPU thread can sum over




Figure 2. Strong scaling of the *ENCORE* code: dependence of runtime, T , on the number of CPU cores on a single node, for the test cases considered in Fig. 1, fixing $\bar{n} \approx 1.5 \times 10^{-4}$. Dashed lines indicate linear relationships, and are calibrated at the single-CPU time. Results are shown for $N = 3$ and $N = 4$ with two choices of ℓ_{\max} . The code exhibits close-to-ideal parallelization, as shown by the linear trends. This extends also to higher order. As expected, the runtime increases significantly when ℓ_{\max} , and thus the number of multiplets Λ , is increased. The 3PCF results deviate slightly from the linear relationship when using 32 CPUs; this occurs since we become limited by other processes such as data loading.

Mode	Wall-clock time (4PCF, $\ell_{\max} = 5$)	Wall-clock time (4PCF, $\ell_{\max} = 10$)	Wall-clock time (5PCF, $\ell_{\max} = 5$)
CPU (16-core)	4.8	58.8	741.8
GPU-a (double)	31.1	41.5	86.1
GPU-a (float)	21.9	38.4	74.3
GPU-a (mixed)	22.0	38.6	77.2
GPU-b (double)	30.1	43.3	177.7
GPU-b (float)	23.9	42.7	202.8
GPU-b (mixed)	30.6	41.2	159.6

Table 2. Runtimes of the *ENCORE* GPU code for the 4PCF and 5PCF using $N_r = 10$ radial bins and $\ell_{\max} = 5$ or 10. In each case, we use a synthetic dataset containing $\approx 5 \times 10^4$ uniformly distributed particles in a periodic box of side-length $L = 1500h^{-1}\text{Mpc}$ (corresponding to one-tenth of the BOSS CMASS density), and count pairs of particles up to $r_{\max} = 200h^{-1}\text{Mpc}$. All computations are performed on a single P100 Nvidia GPU coupled with an Intel Xeon Broadwell CPU. We give results for each of the GPU modes detailed in §4.1.3; these vary whether the M summation is distributed across threads and what level of precision is used for intermediate calculations. For comparison, we give also the timings from the CPU code using a 16-core Intel Skylake node, parallelized using *OPENMP*. All times are reported in seconds.

all M for a given Λ and B , or (b) each GPU thread can compute contributions from a single set of coefficients M (requiring atomic memory access). To test these, we use the same setup as §5.1, except with \bar{n} fixed to the lowest value of $1.5 \times 10^{-5} h^3 \text{Mpc}^{-3}$ for speed (recalling that the higher-point functions scale linearly with \bar{n}). The 4PCF and 5PCF are then computed both on the CPU and GPU, using an Nvidia P100 chip for the latter, and the runtimes are shown in Tab. 2.

For the 4PCF at $\ell_{\max} = 5$, computation is far faster for the 16-core CPU code than for the GPU-accelerated version. This occurs since the GPU code uses only a single CPU to compute the $a_{\ell m}^B$ harmonic coefficients, which become rate-limiting for low N and ℓ_{\max} . Testing the GPU implementation against such a parallelized CPU code is a reasonable test however, given that GPUs are currently a scarcer resource for most high performance computing systems, though this is subject to change. For $\ell_{\max} = 10$, we find a slight preference for the GPU code, which is much increased if one considers the 5PCF. Indeed, the GPU implementation is able to provide speed-ups of $\mathcal{O}(8\times)$ for the latter case.

On newer GPU chips such as the A100 series,¹¹ we expect the performance to increase significantly. This is possible since the primary particle loops entering Algorithm 3 are distributed across a large number of threads. For the P100 chip used herein, we have 3584 CUDA cores, each with up to 2048 threads, allowing a maximum of $\sim 7.3 \times 10^6$ threads to be run in parallel. The A100 chip has 6912 CUDA cores and moreover, as each individual core is faster, Nvidia claims an order of magnitude speed-up for HPC applications.¹²

Tab. 2 also allows us to compare the different GPU modes discussed in §4.1.3. For the 4PCF test-cases, we find comparable runtimes between cases (a) and (b). Case (a) uses a much smaller number of tasks of unequal length (8 280 or 45 720 for $\ell_{\max} = 5$ or 10), whilst case (b) launches a far greater number of tasks (135 120 or 2 499 480 respectively), but requires atomic memory access (to ensure that no two threads access the same memory location concurrently). Whilst the first case does not come close to fully utilizing the GPU resources (since the number of tasks

¹¹ www.nvidia.com/en-us/data-center/ampere-architecture/

¹² www.nvidia.com/en-us/data-center/a100/

launched is significantly below the number that can be simultaneously performed on the GPU), atomic addition does not have to be performed, compensating for the inefficiency. In fact, the degree of similarity suggests that memory access of the NPCF array is the main limiting factor in this case. For the 5PCF, we note a different story; case (a) is $\sim 2 - 3\times$ faster than case (b), even though the number of tasks is much reduced (295,260 instead of 29 606 430, cf. §5.1). This is unsurprising since the 5PCF array is of much higher dimension than the 4PCF, thus we utilize a good fraction of the available GPU threads in case (a), while in case (b) we vastly exceed maximum parallelization; *i.e.*, the number of tasks is larger than the maximum number of threads that can be processed in parallel. Above this threshold, we cease to realize additional speed gains so the difference in number of parallel threads is not large enough to overcome the speed penalty of atomicity.

We finally comment on the dependence of computation-time on precision (cf. §4.1.3). Most consumer-grade GPUs are optimized for float-based calculations, thus one might naïvely expect a significant performance increase when lower precision is used in the GPU kernels; indeed this was found in preliminary testing on commercially-available GeForce GTX 1660 architectures where ‘float’ mode was found to be $\sim 8\times$ faster than ‘double’ mode. However, many research-grade data center GPUs are optimized for double-precision calculations. We observe this in practice, as we do not find a particularly large improvement in efficiency on the Nvidia P100 GPUs; at most the runtime decreases by $\sim 30\%$, though this depends on the precise kernel. We caution that reducing the precision of the GPU calculation necessarily leads to errors in the final result; for the 4PCF in ‘float’ (‘mixed’) mode, the average fractional error is $\sim 3 \times 10^{-3}$ (1×10^{-5}), increasing to $\sim 2 \times 10^{-2}$ (7×10^{-5}) for the 5PCF.

5.4 Application to Mock Catalogs

Finally, we apply the NPCF estimators of this work to realistic simulations of galaxy surveys. Here, we use a set of 10 MULTIDARK-PATCHY mocks (Kitauro et al. 2016; Rodríguez-Torres et al. 2016) constructed for the Baryon Oscillation Spectroscopic Survey (BOSS) (Dawson et al. 2013), using the twelfth data release (DR12) of SDSS-III (Eisenstein et al. 2011; Alam et al. 2015). Presentation of the full BOSS 4PCFs will follow in Philcox et al. (2021a). The MULTIDARK-PATCHY mocks are constructed using approximate gravity solvers, calibrated to an N -body simulation (Klypin et al. 2015). Here, we use the northern Galactic cap (NGC) CMASS sample (hereafter CMASS-N), containing $\sim 6 \times 10^5$ galaxies with mean density $\bar{n} \sim 1.5 \times 10^{-4} h^3 \text{Mpc}^{-3}$ at an effective redshift of $z_{\text{eff}} = 0.57$. To use the NPCF estimator of (38), we require also the random field $R(\mathbf{r})$; for this we use a public random catalog, containing 50 times the galaxy density. As in §4.2, this is split into 32 disjoint subsets, enhancing the computational efficiency for small N .

To give insight into the performance of the algorithm on real data, we first consider the runtime of the various sections of the NPCF computation in Algorithm 1 for a single ‘data-minus-random’ chunk. This is performed for both the 3PCF and 4PCF in 10 linearly-spaced radial bins up to $r_{\text{max}} = 200h^{-1}\text{Mpc}$, using $\ell_{\text{max}} = 5$ or 10 and skipping any odd-parity modes. This gives a total of 270 (495) correlation function components for the $\ell_{\text{max}} = 5$ (10) 3PCF, and 13, 320 (80, 520) for the 4PCF, as in §5.1. The code is run on a 16-core node with a 2.4 GHz Intel Skylake processor (without GPU acceleration), and requires a wall-clock time of 36 s to compute both the 3PCF and 4PCF multipoles up to $\ell_{\text{max}} = 5$, or 291 s with $\ell_{\text{max}} = 10$.

Fig. 3 shows the timing breakdown for the above samples. We immediately note that the computation time is dominated by the two processes highlighted in §4.1.1: estimation of $a_{\ell m}^b(\mathbf{s}_i)$ at each primary particle position and their combination (summing over M) to form the \mathcal{N}_{Λ}^B basis coefficients. For the 3PCF, the former is rate-limiting (as in Slepian & Eisenstein 2015), whilst for the 4PCF, the bin-summations start to become important, particularly as ℓ_{max} (and thus the number of m indices) increases (cf. §5.1). Indeed, for moments beyond the 4PCF, virtually all computation time is spent in this procedure, thus the algorithm is expected to scale linearly with $N_{\mathbf{g}}$, as found in Fig. 1. In practice, much of the time is spent loading the various harmonic coefficients elements into memory, rather than performing the actual addition into the NPCF sum itself. For $\ell_{\text{max}} = 10$, and $N_r = 10$, we have 660 complex elements in $a_{\ell m}^b(\mathbf{s}_i)$, with a total size of 10.3 kB in double precision. This can comfortably fit within the 64 kB L1 cache of the Intel Skylake processors, which allows for the fastest memory access (though this is not true for the much larger NPCF array).

Next, we consider estimation of the NPCFs themselves, including edge-correction. Here, we will present only the isotropic 4PCF and 5PCF, given that the 3PCF has been presented elsewhere (e.g., Slepian et al. 2017b). As before, we use $N_r = 10$, but restrict to $\ell_{\text{max}} = 5$ ($\ell_{\text{max}} = 3$) for the 4PCF (5PCF) to keep computation times reasonable. This leads to 69 (204) unique multiplets Λ and 120 (210) bin combinations B for the 4PCF (5PCF), giving a total of 8, 280 (42, 840) elements. Whilst this is certainly large, it is still much smaller than the number of galaxies in the survey, which is a practical upper-limit to the number of useful bins, if all are independent. As in §4.2, we compute and average \mathcal{N}_{Λ}^B for all 32 data-minus-random splits as well as the normalization \mathcal{R}_{Λ}^B for a single chunk, before forming the edge-correction matrix $M_{\Lambda\Lambda'}$ of (41) and solving for ζ_{Λ}^B using (42). Using the CPU code on 16 Skylake cores, each split requires ~ 1 (4) CPU-hours for the 4PCF (5PCF) computation, thus the entire statistic can be estimated in ~ 30 (120) CPU-hours. The short computation time means that, in a real analysis, many mocks can be analyzed within the confines of standard computational allocations. Furthermore, estimating the 4PCF multiplets requires only twice as much computation time as estimating the 3PCF, for the same ℓ_{max} .

The edge-correction matrices, $M_{\Lambda\Lambda'}^B$ (41), corresponding to the above scenarios are shown in Fig. 4 for the 4PCF and 5PCF.¹³ The coupling matrix is a function of bin B ; as expected, the coupling strength increases with the radial bin value since the CMASS survey window function is smooth except on the largest scales. For this reason we plot only the final radial bin. In both the 4PCF and 5PCF, the off-diagonal couplings

¹³ For the 3PCF analogs of Figs. 4 & 5, see Slepian & Eisenstein (2015).



$\ell_{\max} = 5$

 $\ell_{\max} = 10$


Figure 3. Breakdown of computation time for NPCF estimation using a single MULTIDARK-PATCHY galaxy mock, for $N = 3$ (left) and $N = 4$ (right), with $\ell_{\max} = 5$ (top) or 10 (bottom). In both cases, we use 10 linearly spaced radial bins up to $r_{\max} = 200$, and run on a 16-core Intel Skylake processor. Runtime is broken down into four sections: (1) computation of the $a_{\ell m}^b$ harmonic coefficients around each primary particle, (2) combination of harmonic coefficients into NPCF multiplets, including summation over M indices, (3) file input/output, (4) other, including placing particles onto a Cartesian grid and all other pre- and post-processing steps. Note that the runtime is dominated by the first two operations (which scale as N_g^2 and N_g respectively, for particle number N_g), with the summation over harmonic coefficients being the limiting factor at large ℓ_{\max} and high order. This motivates using GPUs to accelerate step (2) in these regimes.

between different angular momenta are small, at most 7% (5%) for the 4PCF (5PCF), consistent with that found for the 3PCF in Slepian & Eisenstein (2015). These may also be understood intuitively by way of simple edge-correction models, as in the former work. Whilst Fig. 4 shows a matrix with non-trivial structure, almost all the largest couplings are found in closely separated multiplets Λ and Λ' , *i.e.* those with $|\ell_i - \ell'_i| \leq 1$. This motivates the earlier statement (§3.3) that, by measuring all multiplets up to some ℓ_{\max} , we obtain accurate NPCF measurements up to $(\ell_{\max} - 1)$.

Finally, we come to the NPCF measurements themselves. Given that the statistic dimensionality is large, we here plot only a selection of Λ multiplets, but include all radial bins, collapsed into one-dimension. Fig. 5 shows the measured 4PCF multiplets for the mean of 10 MULTIDARK-PATCHY mocks, and we immediately note that there is a strong signal in some components, particularly $\Lambda = \{0, 0, 0\}$, $\{0, 3, 3\}$ and $\{2, 0, 2\}$, but a much weaker one in others, *e.g.*, $\Lambda = \{3, 3, 2\}$. It is important to stress that the trends in Fig. 5 do *not* indicate a strong detection of non-Gaussianity; we must remember that the 4PCF contains both a disconnected contribution (from the product of two 2PCFs) and a connected component, as in (25). The former is sourced by Gaussian statistics, and would be present even in the absence of non-Gaussianity; before binning, this information is fully degenerate with the 2PCF. For an isotropic Gaussian universe, the disconnected term contributes only to multiplets of the form $\Lambda = \{0, \ell, \ell\}$ and permutations (Philcox et al. 2021a), thus the non-zero elements of $\Lambda = \{1, 1, 2\}$ observed at small

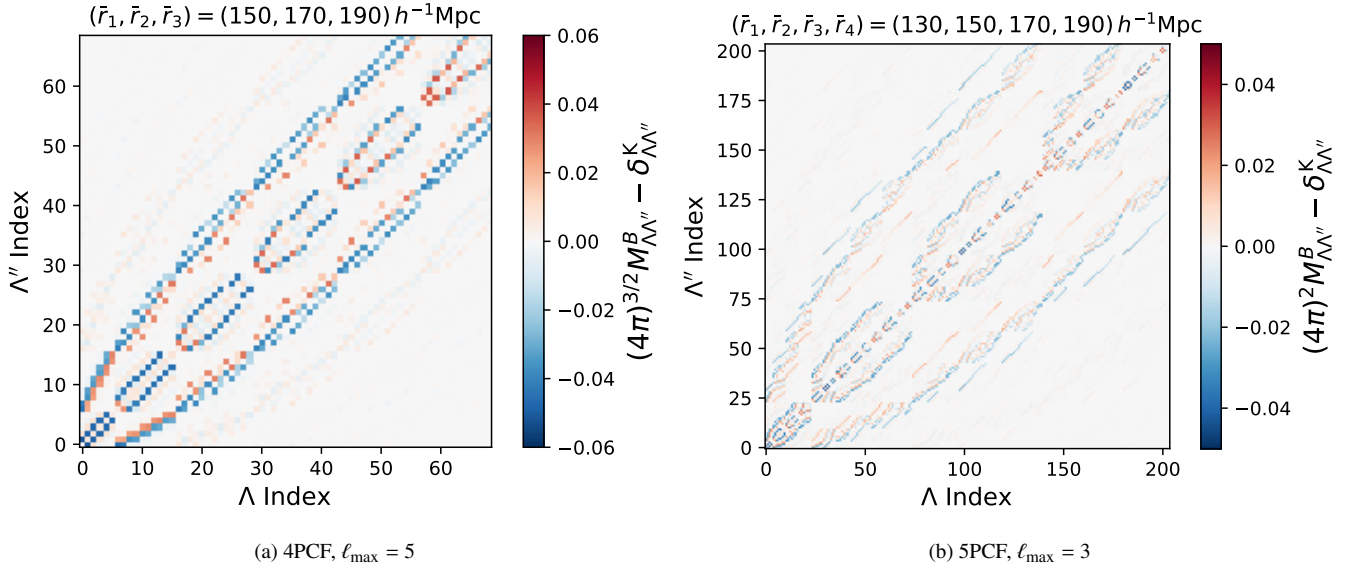


Figure 4. Edge-correction matrices $M_{\Lambda\Lambda'}^B$ for the 4PCF and 5PCF analyses of realistic MULTIDARK-PATCHY simulations. As described in §3.3, these matrices give the couplings between multipoles due to the non-uniform survey geometry and are independent of the data itself. The horizontal and vertical axes give the indices of the Λ and the Λ' multipoles respectively, each of which specifies the angular momentum coefficients $\ell_1, \ell_2, \dots, \ell_{N-1}$. These are arranged in ascending order with the last element of Λ varied first, and we omit any multipoles that do not satisfy the triangle conditions or are parity-odd, *i.e.* the first few terms are $\Lambda = \{0, 0, 0\}, \{0, 1, 1\}, \{0, 2, 2\}, \dots \{1, 0, 1\}, \{1, 1, 0\} \dots$ for the 4PCF. We show the matrices corresponding to the largest radial bin, since this contains the largest-amplitude correlations; the mean radius in each dimension is shown in the caption. The maximal off-diagonal coupling (in the $\Lambda - \Lambda'$ plane) ranges from 0.007 to 0.06 (0.009 to 0.05) for the 4PCF (5PCF). The edge-correction matrices are normalized by $(4\pi)^{(N-1)/2}$ with the unit matrix subtracted, such that all values would be zero for a periodic box geometry. Note that the couplings between multipoles are small in all cases, and further, that, for any given ℓ_{\max} , we do not find significant couplings between any multipoles with $|\ell_i - \ell'_i| > 1$.

radius give evidence for violations of isotropy or Gaussianity; here they are primarily an indication of redshift-space distortions. This will be discussed in depth in Philcox et al. (2021a), alongside a modified estimator that entirely remove the Gaussian contribution.

Fig. 6 shows the analogous results for a selection of 5PCF multipoles. Again we note that some multipoles show a strong signal whilst others are noise-dominated; in contrast to the 4PCF, the disconnected term now involves the product of a 2PCF and 3PCF, thus Fig. 6 *does* represent a robust detection of non-Gaussianity, albeit primarily in the 3PCF. A modified estimator may be derived to compute the connected-only 5PCF, as for the 4PCF. Analysis of the exact 5PCF shapes is beyond the scope of this work. However, we note a general trend of decreasing amplitude as the central bin values, \bar{r}_i increase. Since late-time non-Gaussianity is usually concentrated on small scales, this is as expected.

6 SUMMARY

Correlation functions form the cornerstone of modern cosmology, and their efficient computation is a task of utmost importance for the analysis of current and future galaxy surveys. In this work, we have presented new algorithms for the estimation of the isotropic N -point correlation functions (NPCFs) for arbitrary order N . Making use of recently developed isotropic basis functions (Cahn & Slepian 2020), the estimator can be written in factorizable form and has complexity $\mathcal{O}(N_g^2)$ for N_g particles. Furthermore, the output statistics are geometry-corrected and form a natural split into parity-even and parity-odd modes. The approach may be additionally generalized to homogeneous and isotropic spaces of arbitrary dimension, as discussed in our companion work (Philcox & Slepian 2021).

We provide a public code, ENCORE,¹⁴ with which the isotropic 3PCF, 4PCF, 5PCF and 6PCF can be efficiently computed. The algorithm performs two major steps; a weighted pair-count over galaxies to compute $a_{\ell m}^b$ harmonic coefficients in radial bins b for angular momentum coefficients ℓ, m , and a summation over the m indices of $(N-1)$ such harmonic coefficients to form the NPCF. Whilst the former scales as N_g^2 , the latter has $\mathcal{O}(N_g)$ complexity, and is often found to be rate-limiting for $N > 3$, particularly if the number of bins is large. In many cases, **our estimator scales linearly with the number of galaxies**. We have provided detailed discussion of the ENCORE code, including its implementation in C++ and CUDA, with the latter allowing for GPU acceleration. Furthermore, results for the 4PCF and 5PCF of realistic mock catalogs have been presented; these are amongst the first of their kind. The efficiency of our approach allows them to be quickly estimated; each required ~ 30 (150) CPU-hours in our test. Such methods will allow information to be extracted from the higher-point functions from upcoming surveys such as DESI and Euclid.

A number of avenues present themselves for future work. These include:

¹⁴ github.com/oliverphilcox/encore




Figure 5. Edge-corrected 4PCF multipliers ζ_{Λ}^B measured from 10 MULTIDARK-PATCHY simulations for five choices of the multiplet index $\Lambda = \{\ell_1, \ell_2, \ell_3\}$, setting the internal angles of the 4PCF tetrahedron. We show the mean value (solid lines), and 1σ errors (shaded regions) in all cases. For each Λ , there are 120 choices of radial bin indices $B = \{b_1, b_2, b_3\}$; the top panel shows the linearly-averaged value of the radial bin, \bar{r} , corresponding to each. Succeeding panels show the individual multiplet estimates (normalizing by $\bar{r}_1 \bar{r}_2 \bar{r}_3$), with the value of Λ shown in the caption. Using $\ell_{\max} = 5$, there are 69 such multipliers, of which only five are shown. The ‘sawtooth’ pattern of the 4PCF multipliers occurs due to the non-smooth variation in radial bins; a consequence of plotting a high-dimensional statistic in 1D. We note that the 4PCF contains both disconnected (Gaussian) and disconnected (non-Gaussian) contributions (cf. §3.1), thus much of this result is simply a detection of the 2PCF. All results are computed using 10 radial bins in the range $[0, 200]h^{-1}\text{Mpc}$ on a 16-core Intel Skylake machine. To compute and edge-correct the 4PCF in all multipliers Λ and bins B for ten simulations, a total of 284 CPU-hours was required.




Figure 6. As Fig. 5 but for the 5PCF. As before, the first panel shows the bin-centers for each of the 210 radial bin combinations $B = \{b_1, b_2, b_3, b_4\}$, whilst the remaining plots give the measured multipoles ζ_Λ^B with the internal angles specified by $\Lambda = \{\ell_1, \ell_2, (\ell_{12}), \ell_3, \ell_4\}$. Restricting to $\ell_{\max} = 3$, there are 204 such multipoles. As for the 4PCF, the signal is dominated by the disconnected contribution, here sourced by the product of a 2PCF and 3PCF. The total runtime for all 10 5PCF estimates was 1488 CPU-hours.

- **Theory Model:** To place constraints on cosmology from the NPCFs, we require a robust theory model. Whilst straightforward for the Gaussian contributions, the non-Gaussian pieces are more difficult to treat and must be modeled perturbatively or with simulations.
- **Covariance:** Due to the increase in complexity of the basis functions as N increases, the covariance matrices of the NPCF statistics are non-trivial, even in the isotropic and Gaussian limit. However, computing analytic covariances is of importance since the statistics are high-dimensional, prohibiting a full mock-based covariance. The Gaussian covariances will be presented in [Hou et al. \(2021\)](#).
- **Connected Estimators:** For $N > 3$, the NPCFs contain both *disconnected* and *connected* contributions (cf. §3.1). The former simply repeats information from the lower-point functions, and is not cosmologically relevant. Analyses of the higher-point functions must either model or subtract the disconnected pieces. A connected-only 4PCF estimator will be presented in [Philcox et al. \(2021a\)](#).
- **Application to Data:** The techniques developed in this work allow computation of the edge-corrected NPCFs of survey data. Coupled with the connected-only estimator and Gaussian covariance matrices, this will allow the first estimation of the non-Gaussian higher-point functions from data. Furthermore, consideration of the parity-odd multiplets provides a test for parity-violation in the late Universe. Application to the BOSS galaxy survey will be considered in [Philcox et al. \(2021a\)](#).
- **Gridded Estimator:** Whilst the efficient estimator developed in §3 applies to both gridded and discrete density fields, only the latter is included in ENCORE. Developing a gridded estimator (possible via FFTs) will allow fast analysis of high-density data and application to other contexts such as hydrodynamical simulations. This can be built upon the code of [Portillo et al. \(2018\)](#), which presented a gridded 3PCF estimator for analysis of the turbulent ISM.
- **Anisotropic NPCFs:** Whilst we have focused only on the isotropic case in this work, anisotropic NPCFs also contain information. We expect a similar approach to allow estimation of these statistics, using a generalized angular momentum basis. This is briefly discussed in [Philcox & Slepian \(2021\)](#).
- **CMB Estimators:** A similar prescription to the above may be used to compute N -point statistics on the 2-sphere. This will allow $O(N_{\text{pix}}^2)$ estimation of CMB N -point functions from N_{pix} pixels, and is detailed in [Philcox & Slepian \(2021\)](#).

To close, we briefly reflect on the justification for measuring higher-order NPCFs, and possible alternatives. As is well-known, the process of density field reconstruction (e.g., [Monaco & Efstathiou 1999](#); [Mohayaee et al. 2006](#); [Eisenstein et al. 2007](#); [Noh et al. 2009](#)) shifts information from the higher-point functions back to the 2PCF ([Schmittfull et al. 2015](#)) by reducing long-wavelength displacements. Given that it has become customary for survey analysis pipelines to implement BAO reconstruction (e.g., [Padmanabhan et al. 2012](#); [Beutler et al. 2017a](#); [Vargas-Magaña et al. 2018](#); [Philcox et al. 2020](#)), one must ask whether there remains any utility in the unreconstructed higher-point functions. As shown in [Samushia et al. \(2021\)](#), the information on parameters affecting the *evolution* of matter from its initial conditions, may be better measured by the combination of late-time NPCFs than the initial power spectrum (and hence reconstructed 2PCF). This occurs since these parameters appear in the forward model, rather than the initial conditions. In essence, our argument is that the combination of the reconstructed 2PCF with unreconstructed NPCFs of higher order provides more information, particularly on parameters such as the dark energy density, than either one alone.

Ours is not the only way to proceed. An alternative approach involves forward-modeling the entire density field from its initial conditions to today ([Jasche & Wandelt 2013](#); [Leclercq et al. 2019](#); [Seljak et al. 2017](#); [Feng et al. 2018](#); [Horowitz et al. 2019](#)), including the possibility of field-level inference ([Schmittfull et al. 2019, 2021](#); [Schmidt et al. 2020](#); [Cabass & Schmidt 2020](#); [Schmidt 2021](#)). Such approaches naturally capture information present both in the initial conditions and the forward model, and are a promising new way to perform inference. That said, analysis of this kind is computationally expensive due to the high dimensionality of the problem and necessity to run a large number of forward-model simulations to perform any parameter inference, though various tricks are used to expedite this. In contrast, the more classical approach of measuring the NPCFs allows fast inference once the summary statistics have been computed and stored. Indeed, any new cosmological model can be tested without remeasuring any statistics, unlike in the forward-modeling framework. With the methods discussed in this work, including higher-order statistics in future analyses, such as those of DESI, Roman, and Euclid, has become an attainable goal.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author. The ENCORE code is available at github.com/oliverphilcox/encore and the MULTIDARK-PATCHY mocks can be downloaded from data.sdss.org/sas/dr12/boos/lss/.

ACKNOWLEDGMENTS

We are indebted to many people for insightful discussions, including, but not limited to: Giovanni Cabass, Simone Ferraro, J. Richard Gott, III, Colin Hill, Chris Hirata, Benjamin Horowitz, Mikhail Ivanov, Alex Krolewski, Chung-Pei Ma, Marcel Schmittfull, Marko Simonović, David Spergel, James Sunseri, Martin White, Matias Zaldarriaga, and all members of the Slepian research group. We additionally thank the referee for useful feedback, which helped improve the clarity of the manuscript.

OP acknowledges funding from the WFIRST program through NNG26PJ30C and NNN12AA01C and thanks the Simons Foundation for additional support. ZS is grateful for hospitality at Laboratoire de Physique Nucléaire et des Hautes Energies in Paris for part of the period of this work, as well as at Lawrence Berkeley National Laboratory. Support for this work was provided by the National Aeronautics and Space Administration through Einstein Postdoctoral Fellowship Award Number PF7-180167 issued by the Chandra X-ray Observatory Center, which is operated by the Smithsonian Astrophysical Observatory for and on behalf of the National Aeronautics Space Administration under contract NAS8-03060. ZS also acknowledges support from a Chamberlain Fellowship at Lawrence Berkeley

National Laboratory (held prior to the Einstein) and from the Berkeley Center for Cosmological Physics. RNC also thanks Laboratoire de Physique Nucléaire et des Hautes Energies in Paris for their recurring hospitality. RNC's work is supported in part by the U.S. Department of Energy, Office of Science, Office of High Energy Physics under Contract No. DE-AC02-05CH11231. DJE was supported by DOE-SC0013718 and as a Simons Foundation Investigator.

The authors are pleased to acknowledge that the work reported on in this paper was substantially performed using the Princeton Research Computing resources at Princeton University which is consortium of groups led by the Princeton Institute for Computational Science and Engineering (PICSciE) and Office of Information Technology's Research Computing.

APPENDIX A: GENERALIZED GAUNT INTEGRALS

We give the explicit forms of the generalized Gaunt integrals (15) used to construct the edge-correction equations in §3.3. The forms at arbitrary order can be constructed following [Cahn & Slepian \(2020, §6\)](#).

For $N = 3$, $\Lambda = \{\ell, \ell\}$, giving

$$\mathcal{G}^{\Lambda\Lambda\Lambda} = \frac{1}{4\pi} \sqrt{(2\ell+1)(2\ell'+1)(2\ell''+1)} \begin{pmatrix} \ell & \ell' & \ell'' \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (\text{A1})$$

which is just a rescaled version of the integral of three Legendre polynomials, *viz.* (17). For $N = 4$, $\Lambda = \{\ell_1, \ell_2, \ell_3\}$, leading to

$$\mathcal{G}^{\Lambda\Lambda\Lambda} = \frac{1}{(4\pi)^{3/2}} \left[\prod_{i=1}^3 \sqrt{(2\ell_i+1)(2\ell'_i+1)(2\ell''_i+1)} \begin{pmatrix} \ell_i & \ell'_i & \ell''_i \\ 0 & 0 & 0 \end{pmatrix} \right] \times \begin{Bmatrix} \ell_1 & \ell'_1 & \ell''_1 \\ \ell_2 & \ell'_2 & \ell''_2 \\ \ell_3 & \ell'_3 & \ell''_3 \end{Bmatrix}, \quad (\text{A2})$$

where the 3×3 matrix is a Wigner 9- j symbol (e.g., [NIST DLMF, §34.6](#)). Finally for $N = 5$, $\Lambda = \{\ell_1, \ell_2, (\ell_{12}), \ell_3, \ell_4\}$, and we obtain

$$\mathcal{G}^{\Lambda\Lambda\Lambda} = \frac{1}{(4\pi)^2} \sqrt{(2\ell_{12}+1)(2\ell'_{12}+1)(2\ell''_{12}+1)} \left[\prod_{i=1}^4 \sqrt{(2\ell_i+1)(2\ell'_i+1)(2\ell''_i+1)} \begin{pmatrix} \ell_i & \ell'_i & \ell''_i \\ 0 & 0 & 0 \end{pmatrix} \right] \times \begin{Bmatrix} \ell_1 & \ell_2 & \ell_{12} \\ \ell'_1 & \ell'_2 & \ell'_{12} \\ \ell''_1 & \ell''_2 & \ell''_{12} \end{Bmatrix} \begin{Bmatrix} \ell_{12} & \ell_3 & \ell_4 \\ \ell'_{12} & \ell'_3 & \ell'_4 \\ \ell''_{12} & \ell''_3 & \ell''_4 \end{Bmatrix}, \quad (\text{A3})$$

now depending on two 9- j symbols.

APPENDIX B: SIMPLIFICATION OF THE M SUMMATIONS

Below, we outline a simplification of the M index summations necessary for the NPCF estimators. This reduces the computational requirements by a factor close to two. First, consider the summand of (34), which must be computed at the location of each primary particle:

$$\mathcal{H}_\Lambda^B(\mathbf{s}_i) \equiv \sum_{m_1 \dots m_{N-1}} \mathcal{E}(\Lambda) C_M^\Lambda \prod_{i=j}^{N-1} a_{\ell_j m_j}^{b_j}(\mathbf{s}_i). \quad (\text{B1})$$

As noted in §2.3, the coupling coefficients C_M^Λ impose $m_1 + \dots + m_{N-1} = 0$; this allows (B1) to be written as a sum of three pieces:

$$\mathcal{H}_\Lambda^B(\mathbf{s}_i) = \left(\sum_{m_1 + \dots + m_{n-2} < 0} + \sum_{m_1 + \dots + m_{n-2} = 0} + \sum_{m_1 + \dots + m_{n-2} > 0} \right) \mathcal{E}(\Lambda) C_M^\Lambda \prod_{j=1}^{N-1} a_{\ell_j m_j}^{b_j}(\mathbf{s}_i). \quad (\text{B2})$$

Taking the final term and redefining $m_i \rightarrow -m_i$, we find

$$\begin{aligned} \sum_{m_1 + \dots + m_{n-2} > 0} \mathcal{E}(\Lambda) C_M^\Lambda \prod_{j=1}^{N-1} a_{\ell_j m_j}^{b_j}(\mathbf{s}_i) &= \sum_{m_1 + \dots + m_{n-2} < 0} \mathcal{E}(\Lambda) C_{-M}^\Lambda \prod_{j=1}^{N-1} a_{\ell_j -m_j}^{b_j}(\mathbf{s}_i) \\ &= \sum_{m_1 + \dots + m_{n-2} < 0} C_M^\Lambda \prod_{j=1}^N (-1)^{m_j} a_{\ell_j m_j}^{b_j, *}, \end{aligned} \quad (\text{B3})$$

where we have used (6) and noted that $Y_{\ell-m}(\hat{\mathbf{r}}) = (-1)^m Y_{\ell m}^*(\hat{\mathbf{r}})$, thus $a_{\ell-m}^b(\mathbf{s}_i) = (-1)^m a_{\ell m}^{b*}(\mathbf{s}_i)$. The factor $\prod_{j=1}^{N-1} (-1)^{m_j} = 1$ by the summation rules. Combining (B3) with the first term in (B2) yields

$$\begin{aligned} \left(\sum_{m_1 + \dots + m_{n-2} < 0} + \sum_{m_1 + \dots + m_{n-2} > 0} \right) \mathcal{E}(\Lambda) C_M^\Lambda \prod_{j=1}^N a_{\ell_j m_j}^{b_j}(\mathbf{s}_i) &= \sum_{m_1 + \dots + m_{n-2} < 0} \mathcal{E}(\Lambda) C_M^\Lambda \left(\prod_{j=1}^N a_{\ell_j m_j}^{b_j}(\mathbf{s}_i) + \mathcal{E}(\Lambda) \prod_{j=1}^N a_{\ell_j m_j}^{b_j, *}(\mathbf{s}_i) \right) \\ &= 2 \sum_{m_1 + \dots + m_{n-2} < 0} \mathcal{E}(\Lambda) C_M^\Lambda \times \mathbb{Q}_\Lambda \left[\prod_{j=1}^N a_{\ell_j m_j}^{b_j}(r_j) \right], \end{aligned} \quad (\text{B4})$$

defining the operator

$$\mathbb{Q}_\Lambda[f] = \begin{cases} \text{Re}[f] & \text{if } \mathcal{E}(\Lambda) = +1 \\ i \text{Im}[f] & \text{if } \mathcal{E}(\Lambda) = -1. \end{cases} \quad (\text{B5})$$

Similarly, the second term in (B2) gives

$$\sum_{m_1+\dots+m_{n-2}=0} \mathcal{E}(\Lambda) C_M^\Lambda \prod_{j=1}^N a_{\ell_j m_j}^{b_j}(\mathbf{s}_i) = \sum_{m_1+\dots+m_{n-2}=0} C_M^\Lambda \prod_{j=1}^N a_{\ell_j m_j}^{b_{j,*}}(\mathbf{s}_i), \quad (\text{B6})$$

relabelling $m_i \rightarrow -m_i$ and simplifying. Averaging the two yields the equivalent of (B4) but without the factor of 2. In concert, we obtain

$$\mathcal{H}_\Lambda^B(\mathbf{s}_i) = 2 \sum_{m_1+\dots+m_{n-2} \leq 0} S(m_{N-1}) \mathcal{E}(\Lambda) C_M^\Lambda \mathbb{Q}_\Lambda \left[\prod_{j=1}^N a_{\ell_j m_j}(\mathbf{s}_i) \right], \quad (\text{B7})$$

where $S(m) = 1/2$ if $m = 0$ and unity else. This leads to a significant reduction in computation time and additionally makes manifest the notion that even-parity ($\mathcal{E}(\Lambda) = +1$) multiplets are real and odd-parity ($\mathcal{E}(\Lambda) = -1$) multiplets are imaginary.

REFERENCES

- Agarwal N., Desjacques V., Jeong D., Schmidt F., 2021, *J. Cosmology Astropart. Phys.*, 2021, 021
- Alam S., et al., 2015, *ApJS*, 219, 12
- Alam S., et al., 2017, *MNRAS*, 470, 2617
- Alam S., et al., 2020, arXiv e-prints, p. arXiv:2011.05771
- Alam S., et al., 2021, *Phys. Rev. D*, 103, 083533
- Alsing J., Wandelt B., 2018, *MNRAS*, 476, L60
- Angulo R. E., Foreman S., Schmittfull M., Senatore L., 2015, *J. Cosmology Astropart. Phys.*, 2015, 039
- Baldauf T., Mercolli L., Mirbabayi M., Pajer E., 2015, *J. Cosmology Astropart. Phys.*, 2015, 007
- Bernardeau F., Colombi S., Gaztañaga E., Scoccimarro R., 2002, *Phys. Rep.*, 367, 1
- Bertolini D., Schutz K., Solon M. P., Zurek K. M., 2016, *J. Cosmology Astropart. Phys.*, 2016, 052
- Beutler F., et al., 2017a, *MNRAS*, 464, 3409
- Beutler F., et al., 2017b, *MNRAS*, 466, 2242
- Burkhart B., Lazarian A., 2016, *ApJ*, 827, 26
- Cabass G., Schmidt F., 2020, *J. Cosmology Astropart. Phys.*, 2020, 042
- Cahn R. N., Slepian Z., 2020, arXiv e-prints, p. arXiv:2010.14418
- Chudaykin A., Ivanov M. M., 2019, *J. Cosmology Astropart. Phys.*, 2019, 034
- Colas T., d'Amico G., Senatore L., Zhang P., Beutler F., 2020, *J. Cosmology Astropart. Phys.*, 2020, 001
- Coles P., 2005, arXiv e-prints, pp astro-ph/0502088
- Dawson K. S., et al., 2013, *AJ*, 145, 10
- Eisenstein D. J., Seo H.-J., Sirko E., Spergel D. N., 2007, *ApJ*, 664, 675
- Eisenstein D. J., et al., 2011, *AJ*, 142, 72
- Feldman H. A., Kaiser N., Peacock J. A., 1994, *ApJ*, 426, 23
- Feng Y., Seljak U., Zaldarriaga M., 2018, *J. Cosmology Astropart. Phys.*, 2018, 043
- Fergusson J. R., Shellard E. P. S., 2011, arXiv e-prints, p. arXiv:1105.2791
- Fergusson J. R., Regan D. M., Shellard E. P. S., 2012, *Phys. Rev. D*, 86, 063511
- Friesen B., et al., 2017, arXiv e-prints, p. arXiv:1709.00086
- Gagrani P., Samushia L., 2017, *MNRAS*, 467, 928
- García K., Slepian Z., 2020, arXiv e-prints, p. arXiv:2011.03503
- Gardner J. P., Connolly A., McBride C., 2007, arXiv e-prints, p. arXiv:0709.1967
- Garrison L. H., Eisenstein D. J., Pinto P. A., 2019, *MNRAS*, 485, 3370
- Garrison L., Eisenstein D., Ferrer D., Maksimova N., Pinto P., 2021, The ABACUS Simulation Code, submitted
- Gaztañaga E., Scoccimarro R., 2005, *MNRAS*, 361, 824
- Gil-Marín H., Noreña J., Verde L., Percival W. J., Wagner C., Manera M., Schneider D. P., 2015, *MNRAS*, 451, 539
- Gil-Marín H., Percival W. J., Verde L., Brownstein J. R., Chuang C.-H., Kitaura F.-S., Rodríguez-Torres S. A., Olmstead M. D., 2017, *MNRAS*, 465, 1757
- Gray A. G., Moore A. W., Nichol R. C., Connolly A. J., Genovese C., Wasserman L., 2004, in Oshenbein F., Allen M. G., Egret D., eds, *Astronomical Society of the Pacific Conference Series Vol. 314, Astronomical Data Analysis Software and Systems (ADASS) XIII*. p. 249 (arXiv:astro-ph/0401121)
- Gualdi D., Verde L., 2020, *J. Cosmology Astropart. Phys.*, 2020, 041
- Gualdi D., Manera M., Joachimi B., Lahav O., 2018, *MNRAS*, 476, 4045
- Gualdi D., Gil-Marín H., Manera M., Joachimi B., Lahav O., 2019, *MNRAS*, 484, L29
- Gualdi D., Gil-Marín H., Verde L., 2021, *J. Cosmology Astropart. Phys.*, 2021, 008
- Hahn C., Villaescusa-Navarro F., 2021, *J. Cosmology Astropart. Phys.*, 2021, 029
- Hamilton A. J. S., 1998, *Linear Redshift Distortions: a Review*. p. 185, doi:10.1007/978-94-011-4960-0_17
- Heavens A. F., Jimenez R., Lahav O., 2000, *MNRAS*, 317, 965
- Horowitz B., Seljak U., Aslanyan G., 2019, *J. Cosmology Astropart. Phys.*, 2019, 035
- Hou J., Cahn R. N., Philcox O. H. E., Slepian Z., 2021, arXiv e-prints, p. arXiv:2108.01714
- Ivanov M. M., Simonović M., Zaldarriaga M., 2020, *J. Cosmology Astropart. Phys.*, 2020, 042
- Jasche J., Wandelt B. D., 2013, *MNRAS*, 432, 894

- Jing Y. P., Börner G., 1998, *ApJ*, 503, 37
- Jing Y. P., Börner G., 2004, *ApJ*, 607, 140
- Joachim B., 2017, *MNRAS*, 466, L83
- Kaiser N., 1987, *MNRAS*, 227, 1
- Kamalinejad F., Slepian Z., 2020, arXiv e-prints, p. arXiv:2011.00899
- Karagiannis D., Lazanu A., Liguori M., Raccanelli A., Bartolo N., Verde L., 2018, *MNRAS*, 478, 1341
- Kayo I., et al., 2004, *PASJ*, 56, 415
- Keihänen E., et al., 2019, *A&A*, 631, A73
- Kitaura F.-S., et al., 2016, *MNRAS*, 456, 4156
- Klypin A., Prada F., Yepes G., Heß S., Gottlöber S., 2015, *MNRAS*, 447, 3693
- Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
- Leclercq F., Enzi W., Jasche J., Heavens A., 2019, *MNRAS*, 490, 4237
- March W. B., et al., 2012, in SC'12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis. pp 1–12, doi:10.1109/SC.2012.89
- Marín F., 2011, *ApJ*, 737, 97
- Marín F. A., Wechsler R. H., Frieman J. A., Nichol R. C., 2008, *ApJ*, 672, 849
- Metchnik M. V. L., 2009, PhD Thesis. The University of Arizona
- Mitsou E., Yoo J., Durrer R., Scaccabarozzi F., Tansella V., 2020, *Physical Review Research*, 2, 033004
- Mohayaee R., Mathis H., Colombi S., Silk J., 2006, *MNRAS*, 365, 939
- Monaco P., Efstathiou G., 1999, *MNRAS*, 308, 763
- Moore A. W., et al., 2001, in Banday A. J., Zarubi S., Bartelmann M., eds, Mining the Sky. p. 71 (arXiv:astro-ph/0012333), doi:10.1007/10849171_5
- Moradinezhad Dizgah A., Biagetti M., Sefusatti E., Desjacques V., Noreña J., 2021, *J. Cosmology Astropart. Phys.*, 2021, 015
- Münchmeyer M., Smith K. M., 2019, *Phys. Rev. D*, 100, 123511
- Munshi D., Heavens A., Cooray A., Smidt J., Coles P., Serra P., 2011, *MNRAS*, 412, 1993
- NIST DLMF, NIST Digital Library of Mathematical Functions. <http://dlmf.nist.gov/>
- Nichol R. C., et al., 2006, *MNRAS*, 368, 1507
- Noh Y., White M., Padmanabhan N., 2009, *Phys. Rev. D*, 80, 123501
- Padmanabhan N., Xu X., Eisenstein D. J., Scalzo R., Cuesta A. J., Mehta K. T., Kazin E., 2012, *MNRAS*, 427, 2132
- Pearson D. W., Samushia L., 2018, *MNRAS*, 478, 4500
- Pearson D. W., Samushia L., 2019, *MNRAS*, 486, L105
- Peebles P. J. E., 1978, in Large Scale Structures in the Universe. p. 217
- Peebles P. J. E., 2001, in Martínez V. J., Trimble V., Pons-Bordería M. J., eds, Astronomical Society of the Pacific Conference Series Vol. 252, Historical Development of Modern Cosmology. p. 201 (arXiv:astro-ph/0103040)
- Peebles P. J. E., Groth E. J., 1975, *ApJ*, 196, 1
- Philcox O. H. E., 2021a, arXiv e-prints, p. arXiv:2107.06287
- Philcox O. H. E., 2021b, *MNRAS*, 501, 4004
- Philcox O. H. E., Eisenstein D. J., 2019, *MNRAS*, 490, 5931
- Philcox O. H. E., Eisenstein D. J., 2020, *MNRAS*, 492, 1214
- Philcox O. H. E., Slepian Z., 2021, arXiv e-prints, p. arXiv:2106.10278
- Philcox O. H. E., Ivanov M. M., Simonović M., Zaldarriaga M., 2020, *J. Cosmology Astropart. Phys.*, 2020, 032
- Philcox O. H. E., Hou J., Slepian Z., 2021a, arXiv e-prints, p. arXiv:2108.01670
- Philcox O. H. E., Ivanov M. M., Zaldarriaga M., Simonović M., Schmittfull M., 2021b, *Phys. Rev. D*, 103, 043508
- Portillo S. K. N., Slepian Z., Burkhart B., Kahraman S., Finkbeiner D. P., 2018, *ApJ*, 862, 119
- Rodríguez-Torres S. A., et al., 2016, *MNRAS*, 460, 1173
- Sabiu C. G., Hoyle B., Kim J., Li X.-D., 2019, *ApJS*, 242, 29
- Samushia L., Slepian Z., Villaescusa-Navarro F., 2021, *MNRAS*, 505, 628
- Schmidt F., 2021, *J. Cosmology Astropart. Phys.*, 2021, 032
- Schmidt F., Cabass G., Jasche J., Lavaux G., 2020, *J. Cosmology Astropart. Phys.*, 2020, 008
- Schmittfull M., Feng Y., Beutler F., Sherwin B., Chu M. Y., 2015, *Phys. Rev. D*, 92, 123522
- Schmittfull M., Simonović M., Assassi V., Zaldarriaga M., 2019, *Phys. Rev. D*, 100, 043514
- Schmittfull M., Simonović M., Ivanov M. M., Philcox O. H. E., Zaldarriaga M., 2021, *J. Cosmology Astropart. Phys.*, 2021, 059
- Scoccimarro R., 2000, *ApJ*, 544, 597
- Scoccimarro R., 2015, *Phys. Rev. D*, 92, 083532
- Scoccimarro R., Feldman H. A., Fry J. N., Frieman J. A., 2001, *ApJ*, 546, 652
- Seljak U., Aslanyan G., Feng Y., Modi C., 2017, *J. Cosmology Astropart. Phys.*, 2017, 009
- Shiraishi M., Sugiyama N. S., Okumura T., 2017, *Phys. Rev. D*, 95, 063508
- Simonović M., Baldauf T., Zaldarriaga M., Carrasco J. J., Kollmeier J. A., 2018, *J. Cosmology Astropart. Phys.*, 2018, 030
- Slepian Z., Eisenstein D. J., 2015, *MNRAS*, 454, 4142
- Slepian Z., Eisenstein D. J., 2016, *MNRAS*, 455, L31
- Slepian Z., Eisenstein D. J., 2017, *MNRAS*, 469, 2059
- Slepian Z., Eisenstein D. J., 2018, *MNRAS*, 478, 1468
- Slepian Z., et al., 2017a, *MNRAS*, 468, 1070
- Slepian Z., et al., 2017b, *MNRAS*, 469, 1738
- Slepian Z., et al., 2018, *MNRAS*, 474, 2109
- Smith K. M., Zaldarriaga M., 2011, *MNRAS*, 417, 2
- Smith K. M., Senatore L., Zaldarriaga M., 2015, arXiv e-prints, p. arXiv:1502.00635
- Sosa Nuñez F., Niz G., 2020, *J. Cosmology Astropart. Phys.*, 2020, 021
- Sugiyama N. S., Saito S., Beutler F., Seo H.-J., 2019, *MNRAS*, 484, 364

- Sugiyama N. S., Saito S., Beutler F., Seo H.-J., 2020, *MNRAS*, **497**, 1684
Sugiyama N. S., Saito S., Beutler F., Seo H.-J., 2021, *MNRAS*, **501**, 2862
Szapudi I., 2004, *ApJ*, **605**, L89
Szapudi I., Szalay A. S., 1998, *ApJ*, **494**, L41
Takada M., Jain B., 2003, *MNRAS*, **340**, 580
Tomlinson J., Jeong D., Kim J., 2019, *AJ*, **158**, 116
Umeh O., 2021, *J. Cosmology Astropart. Phys.*, **2021**, 035
Vargas-Magaña M., et al., 2018, *MNRAS*, **477**, 1153
Varshalovich D. A., Moskalev A. N., Khersonskii V. K., 1988, *Quantum Theory of Angular Momentum*, doi:10.1142/0270.
Wang Y., Yang X., Mo H. J., van den Bosch F. C., Chu Y., 2004, *MNRAS*, **353**, 287
Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2013, *Phys. Rep.*, **530**, 87
Zaroubi S., Hoffman Y., 1996, *ApJ*, **462**, 25
Zhang L. L., Pen U.-L., 2005, *New Astron.*, **10**, 569
d'Amico G., Gleyzes J., Kokron N., Markovic K., Senatore L., Zhang P., Beutler F., Gil-Marín H., 2020, *J. Cosmology Astropart. Phys.*, **2020**, 005

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