



The XFaster Power Spectrum and Likelihood Estimator for the Analysis of Cosmic Microwave Background Maps

A. E. Gambrel¹, A. S. Rahlin^{1,2} , X. Song³, C. R. Contaldi⁴ , P. A. R. Ade⁵, M. Amiri⁶, S. J. Benton³ , A. S. Bergman³, R. Bihary⁷, J. J. Bock^{8,9}, J. R. Bond¹⁰ , J. A. Bonetti⁹, S. A. Bryan¹¹, H. C. Chiang^{12,13} , A. J. Duivenvoorden^{3,14} , H. K. Eriksen¹⁵, M. Farhang^{10,16,17}, J. P. Filippini^{18,19} , A. A. Fraisse³, K. Freese^{14,20}, M. Galloway¹⁵, N. N. Gandilo²¹, R. Gualtieri²² , J. E. Gudmundsson¹⁴ , M. Halpern⁶ , J. Hartley²³, M. Hasselfield²⁴, G. Hilton²⁵, W. Holmes⁹, V. V. Hristov⁸, Z. Huang¹⁰, K. D. Irwin^{26,27}, W. C. Jones³ , A. Karacki¹⁵, C. L. Kuo²⁶, Z. D. Kermish³, J. S.-Y. Leung^{17,28} , S. Li^{3,29} , D. S. Y. Mak⁴, P. V. Mason⁸, K. Megerian⁹, L. Moncelsi⁸, T. A. Morford⁸, J. M. Nagy^{30,31} , C. B. Netterfield^{17,23}, M. Nolta¹⁰, R. O'Brient⁹, B. Osherson¹⁸, I. L. Padilla^{17,32} , B. Racine¹⁵, C. Reintsema²⁵, J. E. Ruhl⁷, T. M. Ruud¹⁵ , J. A. Shariff¹⁰, E. C. Shaw¹⁸ , C. Shiu³, J. D. Soler³³ , A. Trangsud^{8,9}, C. Tucker⁵, R. S. Tucker⁸, A. D. Turner⁹, J. F. van der List³, A. C. Weber⁹, I. K. Wehus¹⁵ , S. Wen⁷, D. V. Wiebe⁶, and E. Y. Young^{26,27}

¹ Kavli Institute for Cosmological Physics, University of Chicago, 5640 South Ellis Avenue, Chicago, IL 60637 USA; agambrel@kicp.uchicago.edu

² Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510-5011, USA

³ Department of Physics, Princeton University, Jadwin Hall, Princeton, NJ 08544, USA

⁴ Blackett Laboratory, Imperial College London, SW7 2AZ, London, UK

⁵ School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, CF24 3AA, UK

⁶ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1, Canada

⁷ Physics Department, Case Western Reserve University, 10900 Euclid Avenue, Rockefeller Building, Cleveland, OH 44106, USA

⁸ Division of Physics, Mathematics and Astronomy, California Institute of Technology, MS 367-17, 1200 East California Boulevard, Pasadena, CA 91125, USA

⁹ Jet Propulsion Laboratory, Pasadena, CA 91109, USA

¹⁰ Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, ON M5S 3H8, Canada

¹¹ School of Electrical, Computer, and Energy Engineering, Arizona State University, 650 East Tyler Mall, Tempe, AZ 85281, USA

¹² Department of Physics, McGill University, 3600 Rue University, Montreal, QC, H3A 2T8, Canada

¹³ School of Mathematics, Statistics and Computer Science, University of KwaZulu-Natal, Durban, South Africa

¹⁴ The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden

¹⁵ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, NO-0315 Oslo, Norway

¹⁶ Department of Physics, Shahid Beheshti University, 1983969411, Tehran, Iran

¹⁷ Department of Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4 Canada

¹⁸ Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, IL 61801, USA

¹⁹ Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 West Green Street, Urbana, IL 61801, USA

²⁰ Department of Physics, University of Texas, 2515 Speedway, C1600, Austin, TX 78712, USA

²¹ Steward Observatory, 933 North Cherry Avenue, Tucson, AZ, 85721, USA

²² High Energy Physics Division, Argonne National Laboratory, Argonne, IL 60439, USA

²³ Department of Physics, University of Toronto, 60 St. George Street, Toronto, ON M5S 3H4 Canada

²⁴ Department of Astronomy and Astrophysics, Pennsylvania State University, 520 Davey Lab, University Park, PA 16802, USA

²⁵ National Institute of Standards and Technology, 325 Broadway, Mailcode 817.03, Boulder, CO 80305, USA

²⁶ Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA

²⁷ SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA

²⁸ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4 Canada

²⁹ Department of Mechanical and Aerospace Engineering, Princeton University, Engineering Quadrangle, Princeton, NJ 08544, USA

³⁰ Department of Physics, Washington University in St. Louis, 1 Brookings Drive, St. Louis, MO 63130, USA

³¹ McDonnell Center for the Space Sciences, Washington University in St. Louis, 1 Brookings Drive, St. Louis, MO 63130, USA

³² Department of Physics and Astronomy, Johns Hopkins University, 3701 San Martin Drive, Baltimore, MD 21218 USA

³³ Max-Planck-Institute for Astronomy, Königstuhl 17, D-69117, Heidelberg, Germany

Received 2021 May 21; revised 2021 August 24; accepted 2021 August 24; published 2021 November 25

Abstract

We present the XFaster analysis package, a fast, iterative angular power spectrum estimator based on a diagonal approximation to the quadratic Fisher matrix estimator. It uses Monte Carlo simulations to compute noise biases and filter transfer functions and is thus a hybrid of both Monte Carlo and quadratic estimator methods. In contrast to conventional pseudo- C_ℓ -based methods, the algorithm described here requires a minimal number of simulations and does not require them to be precisely representative of the data to estimate accurate covariance matrices for the bandpowers. The formalism works with polarization-sensitive observations and also data sets with identical, partially overlapping, or independent survey regions. The method was first implemented for the analysis of BOOMERanG data and also used as part of the *Planck* analysis. Here we describe the full, publicly available analysis package, written in Python, as developed for the analysis of data from the 2015 flight of the SPIDER instrument. The package includes extensions for self-consistently estimating null spectra and estimating fits for Galactic foreground contributions. We show results from the extensive validation of XFaster using simulations and its application to the SPIDER data set.

Unified Astronomy Thesaurus concepts: Cosmic microwave background radiation (322); Maximum likelihood estimation (1901); Observational cosmology (1146)

1. Introduction

Analysis of cosmic microwave background (CMB) observations generally requires the compression of large time-domain data sets into the much smaller space of angular power spectra, C_ℓ . This compression is typically achieved in multiple stages. The first stage involves a compression of the time-domain data into a higher signal-to-noise ratio map of the sky (Borrill 1999). The second stage involves a compression from the map domain to the angular power spectrum (Bond et al. 1998). In order to achieve an unbiased, optimal, and lossless compression, both of these stages require a number of assumptions to hold. The first compression relies on the assumption that the residual between the time-domain data and a signal model is distributed as a Gaussian variate. In practice, the noise is also assumed to be stationary over sufficiently long timescales such that it can be modeled efficiently in the Fourier domain over a useful range of frequencies. The second stage assumes that both the signal and noise components of the map are Gaussian-distributed random variates with known pixel-to-pixel covariances. In principle, if these assumptions hold, optimal lossless maps of the sky can be obtained using a closed-form solution of the χ^2 minimization problem, and the maps can be compressed to a final set of C_ℓ values through a numerical maximization of a likelihood of the map.

In practice, a number of complications limit the validity of these assumptions. This is particularly the case for ground-based or suborbital observations, where scan strategies limit long-term stability compared to space-based observations and introduce a number of systematics. These include beam asymmetries, atmospheric effects, instrumental noise, and thermal instabilities. These systematics are difficult, or sometimes impossible, to account for using covariances describing stationary, statistically isotropic random variates. Instead, many observations are modeled using end-to-end simulations where nonidealities can be included more easily. The simulations are used to calibrate templates of systematic effects to be subtracted from the data, calculate noise biases, and determine the distribution of estimated quantities. The presence of sky cuts, foregrounds, and inhomogeneous pixel coverage can also be modeled easily in end-to-end simulations. Simulation-based methods have been employed successfully in the analysis of ground-based, suborbital, and space-based observations (e.g., Netterfield et al. 2002; Chiang et al. 2010; Planck Collaboration et al. 2020a).

One of the first methods to use simulations to estimate power spectra is the MASTER formalism (Hivon et al. 2002). In this approach, the full-sky C_ℓ values are estimated from the cut-sky pseudo- C_ℓ , \tilde{C}_ℓ , by subtracting a noise bias and dividing by a filter function, both of which are calibrated using noise- and signal-only ensembles of sky maps obtained by end-to-end simulations of the time-domain data. The \tilde{C}_ℓ values are related to the full-sky C_ℓ using coupling kernels that can be calculated from the weighted cuts imposed on the sky maps.

As long as the noise and signal simulations are representative of the data, this method gives an unbiased estimate of the bandpowers, and their covariance is determined from the simulations. However, it is nontrivial to produce accurate signal and noise simulations, and in general, the MASTER method requires iteration of the simulated map ensembles to produce an accurate covariance. Since map generation using end-to-end simulations of the time-domain data tends to be the most computationally expensive step in most CMB analysis

pipelines, this can be inefficient or intractable for modern data sets.

This paper details the XFaster method and demonstrates its application on the SPIDER 2015 flight data. The XFaster method blends the maximum-likelihood approach of Bond et al. (1998) with the MASTER approach by introducing an approximate likelihood for the data that is calibrated using simulation ensembles. There are a number of advantages to this approach. The definition of an approximate likelihood allows the use of a quadratic estimator to obtain the C_ℓ with a simultaneous estimate of a Fisher matrix. The likelihood method can be extended to include marginalization over additional signals, systematics, or noise biases. Prior constraints are also easily included when a likelihood is defined. Ancillary quantities such as bandpower window functions can be calculated from the maximum-likelihood estimator. It also reduces the number of simulations required to calibrate the noise and filter biases to a minimum set of noise- and signal-only simulations that do not necessarily need to be fiducial matches to the data. The introduction of an approximate likelihood can also be used to define a higher-level likelihood for the model parameters given the data. This allows XFaster to be used in likelihood samplers to obtain estimates of model parameters directly from the map-level data.

The XFaster method was originally conceived for the analysis of data from the BOOMERanG flights (Netterfield et al. 2002; Jones et al. 2006). The application to BOOMERanG data motivated the main development of the XFaster pipeline, originally written as a Fortran90 package. It was also applied to *Planck* data (Rocha et al. 2011). Here we comprehensively review a revamped and extended version, rewritten as a versatile Python package, for the analysis of SPIDER data. SPIDER is a balloon-borne polarimeter that was launched on 2015 January 1 from the NASA/NSF Long-Duration Balloon Facility near McMurdo Station, Antarctica. It mapped 2480 deg² at 95 and 150 GHz during its 16.5 day flight. Details relevant to the application of XFaster on this data set, as presented in Spider Collaboration (2021), are given throughout. However, while this paper uses SPIDER simulations and data to demonstrate the functionality of the pipeline, the methods and publicly available code base are intended for general use for any CMB observatory.

The paper is organized as follows. In Section 2, we introduce the XFaster algorithm for parameter likelihoods and bandpowers. In Section 3, we discuss extensions to the baseline algorithm: null tests and foregrounds. In Section 4, we show the results of running the pipeline on simulations. In Section 5, we show the results of running the pipeline on data from SPIDER's 2015 flight. We discuss details of the public code base and its computational requirements in Section 6, and we conclude in Section 7.

2. The XFaster Algorithm

2.1. Likelihood Approximation

When experiments observe only a fraction of the sky, or when a portion of the sky is excluded to avoid foreground biases, an expansion over full-sky spherical harmonic basis functions will no longer yield orthonormal modes. The spherical harmonic coefficients, or pseudo- $\tilde{a}_{\ell m}$, obtained in this way will be statistically correlated between modes m and ℓ in the sense that $\langle \tilde{a}_{\ell m} \tilde{a}_{\ell' m'}^* \rangle \neq \delta_{\ell\ell'} \delta_{mm'} \tilde{C}_\ell$, where \tilde{C}_ℓ is the cut-

sky, pseudoangular power spectrum, defined as

$$\tilde{C}_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |\tilde{a}_{\ell m}|^2. \quad (1)$$

The tilde indicates that the quantity is computed in the partial-sky, filtered, and beam-smoothed reference space of the data from the instrument.

Hivon et al. (2002) showed how the geometry of the mask applied to the data can be used to calculate the coupling between $\tilde{a}_{\ell m}$ coefficients. In turn, under an assumption of isotropy, this can be used to define a linear relationship between the ensemble average of cut- and full-sky angular power spectra,

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} K_{\ell\ell'} C_{\ell'}, \quad (2)$$

where $\langle a_{\ell m} a_{\ell' m'}^* \rangle = \delta_{\ell\ell'} \delta_{mm'} C_\ell$ holds on the full sky, and $K_{\ell\ell'}$ is a coupling kernel that can be computed from the sky mask. This expression can be generalized to include polarization (Challinor & Chon 2005; Rocha et al. 2011).

Assuming the full-sky $a_{\ell m}$ coefficients are Gaussian-distributed, the cut-sky $\tilde{a}_{\ell m}$ s must also be Gaussian, since they are related by a linear transformation. We can then write the likelihood, L , as

$$L(\tilde{\mathbf{d}} | \boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi} |\tilde{\mathbf{C}}|} \exp\left(-\frac{1}{2} \tilde{\mathbf{d}}^\dagger \cdot \tilde{\mathbf{C}}^{-1} \cdot \tilde{\mathbf{d}}\right), \quad (3)$$

where $\tilde{\mathbf{d}}$ is a generalized data vector containing the observed $\tilde{a}_{\ell m}$, $\boldsymbol{\theta}$ is a vector of model parameters, and the generalized covariance matrix $\tilde{\mathbf{C}}$ is the sum of the signal and noise components of the model,

$$\tilde{\mathbf{C}}(\boldsymbol{\theta}) = \tilde{\mathbf{S}}(\boldsymbol{\theta}) + \tilde{\mathbf{N}}, \quad (4)$$

where $\tilde{\mathbf{S}}$ is the signal, which depends on the model parameters, and $\tilde{\mathbf{N}}$ is the noise. The likelihood for the data given the parameters $\boldsymbol{\theta}$ can be interpreted as the likelihood for the parameters given the data $L(\boldsymbol{\theta} | \tilde{\mathbf{d}})$ assuming uniform Bayesian priors in $\boldsymbol{\theta}$.

In principle, the exact likelihood of Equation (3) can be used to estimate model parameters by defining the full, nondiagonal ℓ, m structure of the covariance $\tilde{\mathbf{C}}$. In practice, this is not feasible because of the size of the covariance matrix and the difficulty in defining the full anisotropic structure of the noise term.

The X Faster algorithm approximates the likelihood in Equation (3) using two simplifications. The first is an assumption of isotropy, using Equation (1) to assign equal variance to m modes for each multipole ℓ . The second is to construct the model covariance using power averaged over bins in multipoles, or “bandpowers.” Averaging the power in this way reduces the effect of correlations between multipoles induced by the partial coverage. Here X Faster uses bandpower parameters that retain the full ℓ, ℓ' coupling but approximates the likelihood as diagonal in ℓ . This results in an unbiased estimate of the model parameters, but the effective degrees of freedom in the X Faster likelihood must be calibrated using simulations in order to obtain a robust estimate of the likelihood curvature. This calibration, which is most important at the lowest multipoles, where the effect of mode coupling is strongest, is discussed further in Section 2.3.2.

2.2. Bandpower Model

We parameterize the signal portion of the model (Equation (4)) by introducing bandpower deviations, q_b^{XY} , where b is a generalized index indicating the multipole range and XY indicates the cross-spectrum polarization combination, i.e., TT, TE, EE, BB, EB , and TB . We then construct the signal model bandpowers that the q_b factors modify as described in Sections 2.2.1–2.2.5 below. The parameterization of the noise portion of the model is described in Section 2.2.6.

2.2.1. Shape Spectrum

The bandpower deviations q_b are defined with respect to a template, full-sky angular power spectral shape $C_\ell^{(S)}$ such that any sufficiently smooth model power spectrum can be expressed as

$$C_\ell^{XY} = \sum_b \chi_{b\ell}^{XY} q_b^{XY} C_\ell^{XY(S)}, \quad (5)$$

where $\chi_{b\ell}^{XY}$ is a binning operator that is nonzero only for the spectrum component XY and is usually assumed to be piecewise linear by multipole range but could be chosen to be a set of tapered, overlapping kernels. The shape spectrum $C_\ell^{(S)}$ for the CMB is computed using the CAMB package (Lewis et al. 2000). The shape spectrum could also include other sky components in addition to the CMB, such as foregrounds. Alternatively, the template shape can be flat, in which case the q_b parameters are interpreted as the traditional bandpowers, C_b^{XY} . The choice of a flat shape spectrum is appropriate for spectra that vary little within a bandpower and is nonetheless unbiased in the mean for any spectrum shape. However, this is a suboptimal weighting of the power if the signal varies substantially within each bandpower. The calculated coupling between multipoles due to the mask is also more accurate if a known template shape can be used.

2.2.2. Signal Bandpower Kernels

The model signal component $\tilde{\mathbf{S}}$ for the X Faster likelihood covariance $\tilde{\mathbf{C}}$ can be defined using so-called bandpower kernels $\tilde{C}_{b\ell}^{XY,ij}$:

$$\tilde{S}_\ell^{XY,ij} = \sum_b q_b^{XY} \tilde{C}_{b\ell}^{XY,ij}. \quad (6)$$

We have also introduced the index combination ij to indicate the cross-correlation of modes from separate maps i and j . This allows maps from different observations to be combined into a single estimate of a unified power spectrum. The maps can be of different sizes, have different geometries and/or weightings, be overlapping or nonoverlapping, and have different beam smoothing and transfer functions. The coupling between modes is propagated through the estimation by the mode-coupling kernels $K_{\ell\ell'}$ introduced in Equation (2). This structure can also allow for maps of observations at different frequencies when fitting for galactic foregrounds, as described in Section 3.3.

Following the MASTER formalism, the bandpower kernels are written as

$$\tilde{C}_{b\ell}^{XY,ij} = \sum_{\ell'} \chi_{b\ell'}^{XY} K_{\ell\ell'}^{ij} F_{\ell'}^{XY,ij} (B_{\ell'}^{XY,ij})^2 C_{\ell'}^{XY(S)}, \quad (7)$$

with a time-domain filter transfer function $F_\ell^{XY,ij}$, beam window function $B_\ell^{XY,ij}$, mode-coupling kernels $K_{\ell\ell'}^{ij}$, and shape spectrum $C_\ell^{XY(S)}$. When combined with Equation (6), it becomes clear that the signal model is simply Equation (5) reconstructed on the cut sky.

Equation (6) is valid for $XY \in \{TT, TE, TB, EB\}$. However, the cut-sky mask results in the mixing of E - and B -modes, which must be accounted for in their spectral models. The remaining spectral components are

$$\tilde{S}_\ell^{XX,ij} = \sum_b q_b^{XX} + \tilde{C}_{b\ell}^{XX,ij} + \sum_b q_b^{YY} - \tilde{C}_{b\ell}^{YY,ij}, \quad (8)$$

where XX and YY are now the combinations EE and BB , and the $-\tilde{C}_{b\ell}$ terms mix BB power into the EE signal model, and vice versa.

Equation (7) is valid as written for $XY = TT$ but must be modified to construct the remaining bandpower kernels. The EE kernel terms are

$$\begin{aligned} +\tilde{C}_{b\ell}^{EE,ij} &= \sum_{\ell'} \chi_{b\ell'}^{EE} + K_{\ell\ell'}^{ij} F_{\ell'}^{EE,ij} (B_{\ell'}^{EE,ij})^2 C_{\ell'}^{EE(S)}, \\ -\tilde{C}_{b\ell}^{EE,ij} &= \sum_{\ell'} \chi_{b\ell'}^{BB} - K_{\ell\ell'}^{ij} F_{\ell'}^{EE,ij} (B_{\ell'}^{EE,ij})^2 C_{\ell'}^{EE(S)}, \end{aligned} \quad (9)$$

and similarly for BB , where we now introduce the polarization mode-coupling kernels $\pm K_{\ell\ell'}^{ij}$. In particular, $-K_{\ell\ell'}^{ij}$ accounts for the $E-B$ mixing terms. The EB cross-spectrum bandpower kernels are

$$\tilde{C}_{b\ell}^{EB,ij} = \sum_{\ell'} \chi_{b\ell'}^{EB} (+K_{\ell\ell'}^{ij} - K_{\ell\ell'}^{ij}) F_{\ell'}^{EB,ij} (B_{\ell'}^{EB,ij})^2 C_{\ell'}^{EB(S)}. \quad (10)$$

Finally, for $XY \in \{TE, TB\}$, we have

$$\tilde{C}_{b\ell}^{XY,ij} = \sum_{\ell'} \chi_{b\ell'}^{XY} \times K_{\ell\ell'}^{ij} F_{\ell'}^{XY,ij} (B_{\ell'}^{XY,ij})^2 C_{\ell'}^{XY(S)}, \quad (11)$$

where the mode-coupling kernel $\times K_{\ell\ell'}^{ij}$ describes the coupling between temperature and polarization.

The $\tilde{C}_{b\ell}$ for TT , EE , and BB CMB shape spectra for the SPIDER 150 GHz cross-spectrum are shown in Figure 1. The computation of the components of the bandpower kernels and their values for SPIDER are given in Sections 2.2.3–2.2.5.

2.2.3. Beam Window Functions

The beam window functions B_ℓ are an input to the XFaster algorithm. One window function is required per map, and the estimated error on the beam can also be input to the pipeline. The error may be marginalized over in computing the cosmological parameter likelihoods to account for these uncertainties.

The beam terms in the bandpower kernels are constructed as the product of the individual beam windows for each of the two maps indexed by i and j ,

$$(B_\ell^{XY,ij})^2 = B_\ell^{XY,i} \cdot B_\ell^{XY,j}, \quad (12)$$

and the beam error terms are included by adding derivatives of the model with respect to each beam window to the signal covariance.

For SPIDER, the beam window functions are computed by cross-correlating SPIDER data maps at 95 and 150 GHz with

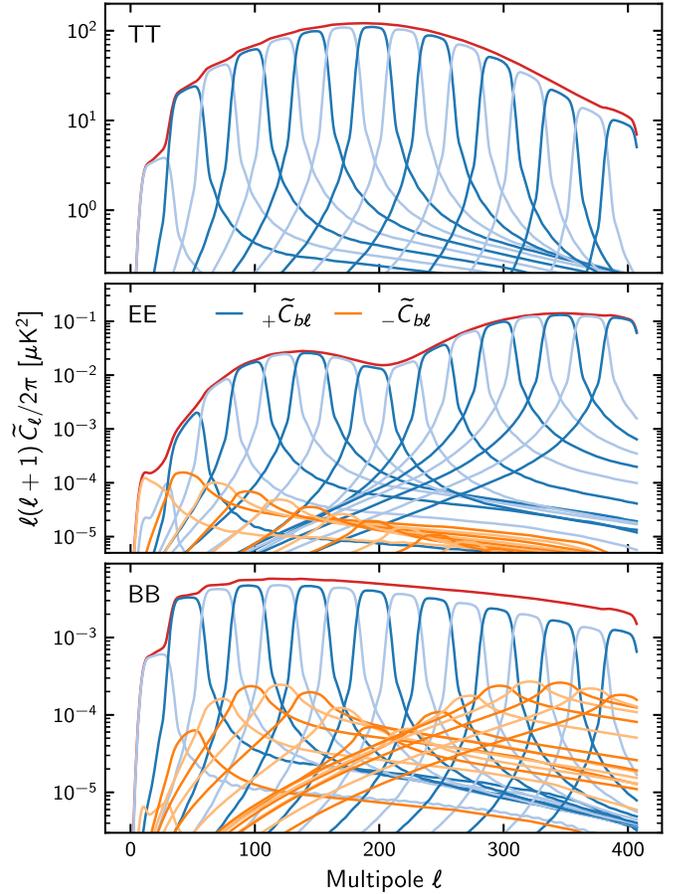


Figure 1. The CMB bandpower kernels ($\tilde{C}_{b\ell}$) for TT , EE , and BB cross-spectra, including SPIDER 150 GHz masking, filtering, and beam smoothing. The binning operator $\chi_{b\ell}$ is piecewise linear with equal-sized bins of width $\Delta\ell = 25$. The colors alternate by bin, with the sum of the contributions from each bin (i.e., the signal model \tilde{S}_ℓ with $q_b = 1$) given by the red line. The underlying BB shape spectrum is constant in $\ell(\ell + 1)C_\ell$ to have appreciable input power for determining the filter transfer function. The estimated data spectra have been found to be insensitive to the choice of shape spectra. The mode-coupling matrix terms that mix E and B polarizations ($-\tilde{C}_{b\ell}$) are also shown (orange); these mixing terms contribute additional power in the tails for each bin, most visible in the BB model.

Planck maps at 100 and 143 GHz, respectively.³⁴ The beam is modeled as a Gaussian function, with an approximate FWHM of $41'$ at 95 GHz and $29'$ at 150 GHz. The errors on the average beams are determined from the distribution of estimated detector beams at that frequency. This produces a 1σ Gaussian error per ℓ ($\sigma \approx 0.2\%$ at $\ell = 100$). To marginalize over this error, we introduce a single beam parameter per frequency that scales this error envelope (i.e., all multipoles are scaled by the same value). This beam parameter is assigned a Gaussian prior ($\mu = 0$, $\sigma = 1$), and the scaled error envelope is added to the beam window function in constructing the model within the likelihood. The `HEALPIX` $N_{\text{side}} = 512$ pixel window function is multiplied by the instrument beam window function to account for smoothing from pixelization in the process of producing pixelized maps from discretely sampled time-ordered data. The beam window functions and 1σ statistical errors for SPIDER are shown in Figure 2.

³⁴ Throughout this paper, we use release 3.01 of the *Planck* HFI maps (Planck Collaboration et al. 2020b).

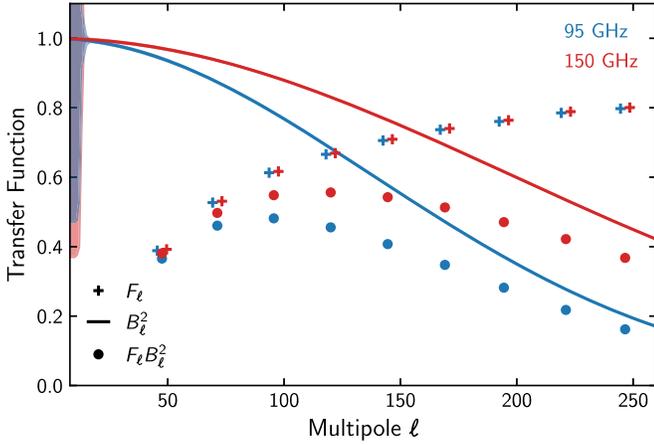


Figure 2. SPIDER’s filter transfer function (F_ℓ), beam window function (B_ℓ^2), and total transfer function for 95 and 150 GHz. Beam error envelopes at each frequency are also shown, most evident at low multipoles. The envelopes represent the 1σ statistical uncertainty in the beam transfer function as determined by the scatter in per-detector beam estimates. Quantities shown are the average of the EE and BB transfer functions, which are similar but not assumed to be identical.

2.2.4. Filter Transfer Functions

In practice, observations of the sky are binned into maps from time-ordered data, which must be filtered to remove systematics like scan-synchronous noise or noisy frequencies. This filtering suppresses signal modes at certain angular scales, and the resulting bias must be computed empirically by comparing an input model spectrum to the spectra of an ensemble of simulations that have been filtered identically to the on-sky data. As in the MASTER formalism, we approximate the filter transfer function F_ℓ as a spherically symmetric function of only ℓ -modes. We also assume that the transfer function is independent of the input signal spectrum used to compute it. We have verified that this is a good assumption for CMB spectra, and in particular that the BB CMB spectra are insensitive to reasonable changes to input spectra. However, using CMB-like input spectra to compute transfer functions results in a bias at low multipoles for maps with a significant dustlike foreground component. Further efforts to account for this bias in the estimation of dust bandpowers will be presented in future work.

The filter transfer functions F_ℓ are computed in the same way as bandpowers, described in Section 2.3.1, but substituting the average of signal-only simulations for the observed data and setting the F_ℓ term in the signal model to 1. The remaining nonunitarity of the q_b values is the binned transfer function. Transfer functions are computed for TT , EE , and BB spectra, independently for each map, since the filtering may differ significantly between maps. The TE , EB , and TB transfer functions are approximated as the geometric mean of their component transfer functions, e.g., $F_\ell^{TE} = \sqrt{F_\ell^{TT} F_\ell^{EE}}$.

When constructing the signal model using the binned transfer function, we expand F_b to the full ℓ range using a constant value in each bin. The transfer function term for each cross-spectrum in the signal model is then the geometric mean of the transfer functions for each of the two maps indexed by i and j :

$$F_\ell^{XY,ij} = \sqrt{F_\ell^{XY,i} \cdot F_\ell^{XY,j}}. \quad (13)$$

The SPIDER EE and BB transfer functions are shown in Figure 2.

2.2.5. Mode-coupling Kernels

Multiple mode-coupling kernels are required in the case of polarization-sensitive observations or maps with different masks. The weighting applied to masked maps can also be different for each Stokes parameter I , Q , and U . The kernels are computed from the cross-correlation power spectrum of the masks on the full sky, \mathcal{W}_L , and we compute separate kernels for each of the polarization combinations and each unique pair of maps i and j . Following Challinor & Chon (2005), the extended set of polarization-sensitive kernels is

$$\begin{aligned} K_{\ell\ell'}^{ij} &= \frac{2\ell' + 1}{4\pi} \sum_L (2L + 1) \mathcal{W}_L^{TT,ij} \begin{pmatrix} \ell & \ell' & L \\ 0 & 0 & 0 \end{pmatrix}^2, \\ \pm K_{\ell\ell'}^{ij} &= \frac{2\ell' + 1}{8\pi} \sum_L (2L + 1) \mathcal{W}_L^{PP,ij} \begin{pmatrix} \ell & \ell' & L \\ 2 & -2 & 0 \end{pmatrix}^2 \\ &\quad \times [1 \pm (-1)^{\ell+\ell'+L}], \\ \times K_{\ell\ell'}^{ij} &= \frac{2\ell' + 1}{4\pi} \sum_L (2L + 1) \mathcal{W}_L^{TP,ij} \begin{pmatrix} \ell & \ell' & L \\ 2 & -2 & 0 \end{pmatrix} \\ &\quad \times \begin{pmatrix} \ell & \ell' & L \\ 0 & 0 & 0 \end{pmatrix}, \end{aligned} \quad (14)$$

where the terms in parentheses are the Wigner 3- j symbols, T is the Stokes I mask, and P is the Stokes Q/U mask, which need not be identical. The \pm kernels are used to compute EE , BB , and EB power spectrum terms, with the $-$ term in particular used to account for mixing between E and B due to the mask. The \times kernel is used to compute the TE and TB spectra. The kernels for the SPIDER mask are shown in Figure 3.

2.2.6. Residual Noise Calibration

The XFAster likelihood approximation also enables an estimation of noise calibration parameters. Uncorrelated noise enters the covariance matrix as a diagonal term. To account for inaccuracies in our noise simulations, we fit a scalar parameter n_b^i per bin as a correction to the noise model,

$$\widetilde{N}_\ell^{XY,ij} = \delta^{ij} \sum_b \chi_{b\ell}^{XY} (1 + n_b^i) \langle \widetilde{N}_\ell^i \rangle, \quad (15)$$

where $\langle \widetilde{N}_\ell^i \rangle$ is the mean spectrum of an ensemble of noise simulations. Noise is treated differently from signal, in that we model it directly in the cut-sky power spectrum with no polarization coupling. In principle, noise is coupled across polarizations, but this is difficult to account for analytically in the cut-sky spectra.

Clearly, if the binning structures of q_b and n_b are too similar, and/or if the template biases $\langle \widetilde{N}_\ell^i \rangle$ have a similar ℓ dependence to the shape templates C_ℓ^{XY} , this extension will introduce significant degeneracies in an autospectrum analysis. The degeneracies are broken by including multiple cross-spectra, where noise biases do not contribute. This has been found for SPIDER to sufficiently decouple the signal and noise parameters, though increasing the bin width for the n_b parameters in comparison to that for q_b would further address this potential issue. The piecewise linear model for the noise calibration means the n_b “noise residual” parameters can be estimated jointly with the signal q_b and marginalized over using the full

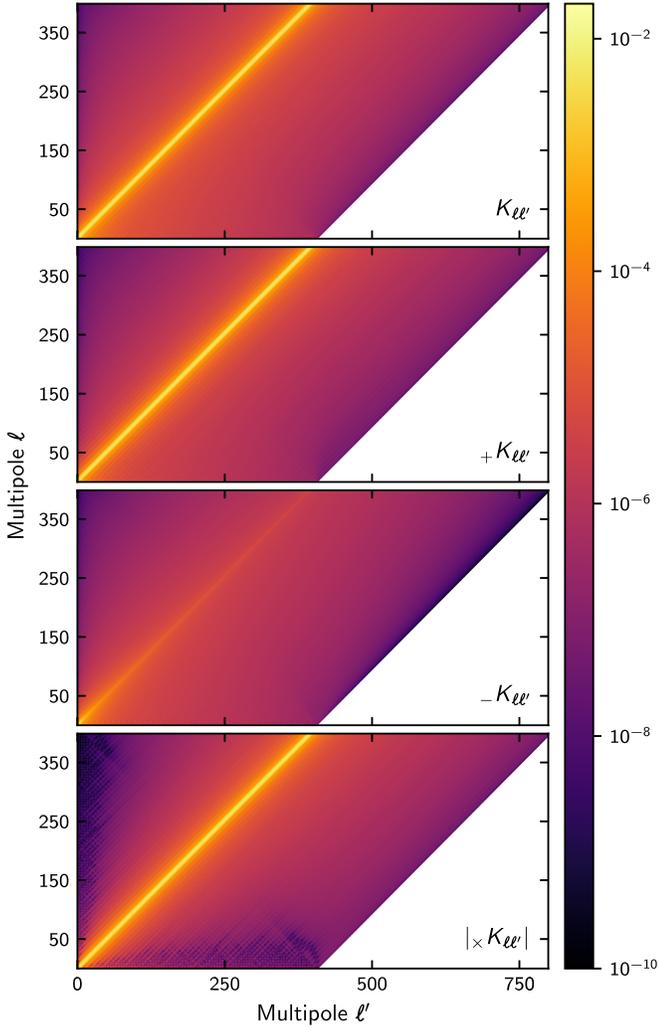


Figure 3. Mode-coupling kernels, $K_{\ell\ell'}$, for the SPIDER mask. Each subplot corresponds to a component in Equation (14). For plotting purposes, negative elements are set to their absolute value for $\times K_{\ell\ell'}$, resulting in the low- ℓ structure seen.

Fisher matrix once the estimator has converged to the maximum-likelihood solution.

2.3. Likelihood Computation

The components introduced above are used to construct the signal and noise covariances in Equation (3). In the XFaster approximation, the covariances are block-diagonal by ℓ . The subblocks are built from the cross-spectra of N maps and polarizations with each subblock as follows:

$$\tilde{\mathcal{C}}_{\ell}^{ij} = \begin{bmatrix} \tilde{\mathcal{C}}_{\ell}^{TT} & \tilde{\mathcal{C}}_{\ell}^{TE} & \tilde{\mathcal{C}}_{\ell}^{TB} \\ - & \tilde{\mathcal{C}}_{\ell}^{EE} & \tilde{\mathcal{C}}_{\ell}^{EB} \\ - & - & \tilde{\mathcal{C}}_{\ell}^{BB} \end{bmatrix}_{ij}, \quad (16)$$

where i and j index over the N independent maps. Formulated in this way, the matrix $\tilde{\mathcal{C}}_{\ell}$ is the covariance of the $\tilde{a}_{\ell m}^{X,i}$ that make up the generalized, observed data vector $\tilde{\mathbf{d}}$ in Equation (3). The block-diagonal form of the covariance means the data vector

can be precompressed into spectra:

$$\hat{\mathcal{C}}_{\ell}^{XY,ij} = \frac{1}{2\ell + 1} \sum_m \hat{a}_{\ell m}^{X,i} \hat{a}_{\ell m}^{*Y,j}. \quad (17)$$

Here and elsewhere, we use the hat symbol ($\hat{\cdot}$) to distinguish matrices of data pseudospectra from general matrices of pseudospectra. Then the log-likelihood, up to an overall constant, can be written as

$$\mathcal{L} \equiv \ln L = -\frac{1}{2} \sum_{tk} (2\ell + 1) g_{\ell}^k [\tilde{\mathcal{C}}_{\ell}^{-1} \cdot \hat{\mathcal{C}}_{\ell} + \ln \tilde{\mathcal{C}}_{\ell}]_{kk}, \quad (18)$$

where k indexes over polarization and independent maps in the subblocks (see Equation (2) in Bond et al. 2000). The factor $(2\ell + 1)$ appears as a degree-of-freedom count due to the block-diagonal form. The coefficient g_{ℓ}^k accounts for the effective number of modes from each map that contribute to the final trace for each multipole.

2.3.1. Likelihood Maximization

The XFaster likelihood (Equation (18)) can be maximized using an iterative quadratic estimator (Bond et al. 1998) to find the maximum-likelihood estimates for parameters θ that include all q_b and n_b that are allowed to vary freely. The method uses the Fisher matrix to approximate the curvature of the likelihood at each iteration. The Fisher matrix is given by

$$\mathcal{F}_{bb'} = \frac{1}{2} \sum_{tk} (2\ell + 1) g_{\ell}^k \left[\frac{\partial \tilde{\mathcal{C}}_{\ell}}{\partial \theta_b} \cdot \tilde{\mathcal{C}}_{\ell}^{-1} \cdot \frac{\partial \tilde{\mathcal{C}}_{\ell}}{\partial \theta_{b'}} \cdot \tilde{\mathcal{C}}_{\ell}^{-1} \right]_{kk}, \quad (19)$$

where we have used the same notation convention as in Equation (18), and the index b now runs across all parameters in the set of q_b and n_b . In practice, starting from an initial guess with $\theta_b = 1$, we calculate an updated solution at each iteration using

$$\theta_b = \frac{1}{2} \sum_{b'tk} \mathcal{F}_{bb'}^{-1} (2\ell + 1) g_{\ell}^k \left[\tilde{\mathcal{C}}_{\ell}^{-1} \cdot \frac{\partial \tilde{\mathcal{C}}_{\ell}}{\partial \theta_{b'}} \cdot \tilde{\mathcal{C}}_{\ell}^{-1} \cdot (\hat{\mathcal{C}}_{\ell} - \tilde{\mathcal{N}}_{\ell}) \right]_{kk} \quad (20)$$

This is equivalent to a Newton–Raphson minimization method. Iterations can be terminated when a convergence criterion is satisfied. We terminate when the maximum of the absolute fractional change in θ_b is below some threshold, typically 10^{-3} . Note that the mode-counting factor g_{ℓ} enters into the elements of the Fisher matrix (and therefore, the resulting uncertainties on the parameters) but is effectively divided out in Equation (20), so miscalibration may result in biased uncertainties but not biased parameter estimates.

In practice, the model matrix is not positive definite because the bandpower deviations are allowed to be negative. This is not strictly a problem because the likelihood is not evaluated during the optimization, but the problem becomes ill conditioned if the steps approach the threshold where the covariance becomes singular. This leads to spurious values for the gradient contribution, driven by numerical errors, and can slow down or prevent convergence. The problem can be solved by ensuring the matrices involved are better conditioned. This is achieved by adding a conditioner to the diagonals of the covariances and is equivalent to adding a numerical floor to the eigenvalues of

the matrices. In our implementation, the conditioning level is introduced if the bandpower deviations are failing to converge and then adjusted automatically to the minimum level required for convergence. The conditioning is then dropped when estimating the final Fisher matrix at the likelihood maximum, since the matrices are always invertible if the likelihood point is well defined.³⁵

The conditioning step is also a useful indicator of systematics. If the presence of the conditioner leads to significant fractional changes to the value of the likelihood at the maximum, this can indicate that inconsistencies in the observed data are driving the spurious contributions to the gradient term in Equation (20). This is a particularly useful check when analyzing multiple maps.

2.3.2. Mode-counting Factor, g_ℓ

In Equation (18), a trace over ℓ, m, k has been reduced to $\sum_{\ell k} (2\ell + 1) g_\ell^k$ by assuming isotropy. The factor of $(2\ell + 1)$ counts the number of m modes per ℓ , and the factor g_ℓ^k accounts for the change to the effective degrees of freedom in the likelihood induced by both map weighting and the “missing” contributions to the covariance in the block-diagonal likelihood approximation. It is important to accurately calibrate this factor when using the XFaster method in order to produce unbiased uncertainties on the fitted parameters in the signal model.

The mode loss is most pronounced at low multipoles, approaching the overall scales of each map, where ℓ -to- ℓ' couplings are significant. In dealing with this bias, one option is to limit the range in ℓ covered by the q_b parameters, but we have found that the addition of a g_ℓ^k factor significantly extends the useful range of multipoles where bandpowers can be estimated accurately.

To calculate g_ℓ , we first estimate an overall ℓ -independent starting amplitude based on the ratio of mask moments introduced in Hivon et al. (2002),

$$g = f_{\text{sky}} \frac{w_2^2}{w_4}, \quad (21)$$

where w_n is the n th moment of mask w_p with

$$w_n = \frac{1}{f_{\text{sky}}} \sum_p w_p^n, \quad (22)$$

where p is the pixel index for the mask and f_{sky} is the fraction of the sky covered by the mask. For past applications, such as Netterfield et al. (2002) and Montroy et al. (2006), this was found to be a sufficient approximation to accurately model the variance of the final bandpowers. We have found that the application to SPIDER requires a more accurate counting of modes, likely due to the presence of high signal-to-noise polarization modes at large scales for which the coupling structure is most complicated. The mode loss induced by the coupling can, in principle, be computed analytically to higher order in map moments (see Challinor & Chon 2005), but the calculation is not straightforward and relies on a number of simplifying assumptions. In practice, we find that a second-order correction to the overall amplitude g by a factor of $(1 + 4f_{\text{sky}})$ is an improvement but insufficient to bring the

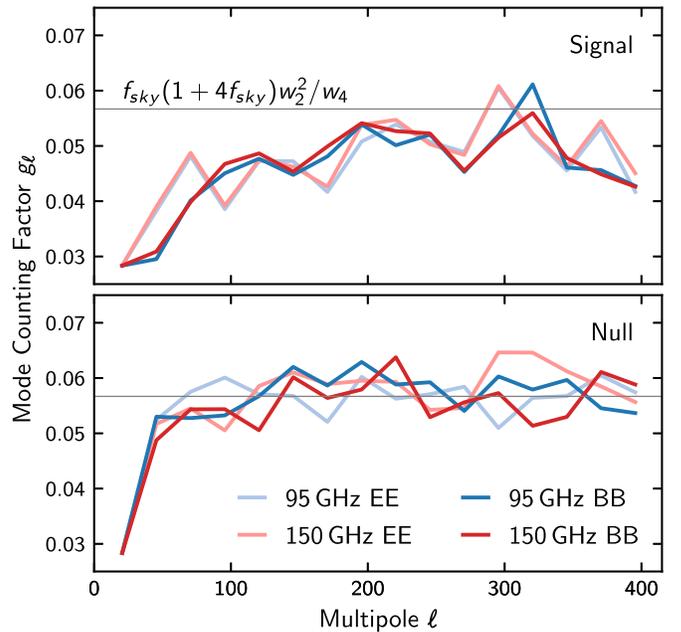


Figure 4. The g_ℓ effective mode-counting factor as estimated for SPIDER. The predominant effect is a constant approximate 6% retention of the full-sky power after masking (Equation (21)). An additional ℓ -dependent reduction results from “missing” contributions to the covariance due to approximations in the construction of the likelihood, as described in Section 2.3.2. This component is empirically calibrated with simulations.

Fisher matrix into agreement with the scatter of simulation bandpowers. Thus, we also compute an empirical Monte Carlo estimate of the ℓ -dependence, which results in accurate error bar estimation throughout the SPIDER ℓ range.

Since the mode-coupling factor is partially degenerate with the filter transfer function F_ℓ , we estimate the final g -correction iteratively by computing bandpowers for 1000 signal-only simulations that have been filtered in the same way as the data and compare the scatter of the ensemble of bandpowers to the diagonal of the Fisher matrix (i.e., XFaster’s estimate of the error bar). We use the ratio as an estimate of the g -correction, feed it back into the estimate of the ensemble of bandpowers, and repeat the process until the correction converges.

Figure 4 shows the amplitude of the total g_ℓ factor per bin using both signal and null (noise-dominated) simulations. The latter are discussed in more detail in Section 3.1. We note here that in the noise-dominated regime, the analytical equation for g_ℓ with the constant second-order correction appears to be a good approximation, on average, at all but the largest angular scales, while the signal-dominated regime shows more structure as a function of ℓ .

2.4. Prior Constraints

A likelihood-based estimator, such as the quadratic estimator of Equation (20), is easily modified to include prior constraints on any of the parameters θ_b . Priors can help improve convergence if poorly constrained parameter directions are included. They are also very useful to include self-consistent marginalization over prior constraints on parameters to obtain a final Fisher matrix that contains a full propagation and accumulation of errors. This is important for nuisance

³⁵ See Gjerløw et al. (2015) for a similar approach using a conditioning prior and an additional step regularizer.

parameters such as noise residual calibrations or foreground parameters (see Section 3).

Including Gaussian priors, which suffices for most practical applications, is particularly simple. At each iteration of the quadratic estimator, the Fisher matrix and estimate values can be modified as

$$\mathcal{F}_{bb'} \rightarrow \mathcal{F}_{bb'} + \delta_{bb'} \frac{1}{\sigma_b^2}, \quad (23)$$

$$\theta_b \rightarrow \theta_b + \frac{\mu_b}{\sigma_b^2}, \quad (24)$$

where μ_b and σ_b^2 are the Gaussian means and variances for each of the priors. In the limit of a tight prior $\sigma_b^2 \rightarrow 0$, this method thus recovers $\theta_b \rightarrow \mu_b$ upon convergence. Final marginalization over any subset of parameters can be achieved by excluding the rows and columns for those parameters in the final Fisher matrix before inversion to obtain the estimated covariances.

2.5. Bandpower Window Functions

The XFaster method is best suited for surveys where reduced sky coverage requires the use of a compression to bandpowers in order to reduce the effect of mode correlations. The final estimated quantities in this case are a set of bandpowers C_b or bandpower deviations q_b and their associated Fisher matrices. To compare these to any proposed model for the full-sky angular power spectrum, C_ℓ , one needs to calculate model C_b or q_b values. Bandpower window functions are needed for this step (Bond 1993; Knox 1999).

Bandpower window functions are linear operators that transform the full-sky spectrum into the estimated quantity. The window functions depend on the effective filtering induced by the observation strategy, the correlations induced by any sky cut, and the definition of the estimator. Different unbiased estimators acting on the same set of observations, for example, will produce estimates that are, in general, different and will only agree in the ensemble limit.

Given a model spectrum C_ℓ , the window functions are the weighting operators appearing in the logarithmic averaging of the power into generalized bandpowers. For bandpower parameters θ_b that correspond to an underlying spectrum C_ℓ in the ensemble limit, this generalizes to

$$\langle \theta_b^{XY} \rangle = \sum_\ell \mathcal{N}_\ell W_{b\ell}^{XY(\theta)} C_\ell^{XY}, \quad (25)$$

where the weighting function $\mathcal{N}_\ell = (2\ell + 1)/4\pi$, and the window functions are normalized as $\sum_\ell \mathcal{N}_\ell W_{b\ell}^{XY(\theta)} C_\ell^{XY(M)} = 1$, where $C_\ell^{XY(M)}$ is a set of model spectra. When the model spectrum is used to construct the signal covariance (i.e., $C_\ell^{(M)} = \sum_b \chi_{b\ell} C_\ell^{(S)}$), Equation (25) results in bandpower deviations q_b ; when the model spectrum is flat (i.e., $C_\ell^{(M)} = 1$), the result is a set of C_b bandpowers; and when the model spectrum is scale-invariant (i.e., $\ell(\ell + 1)C_\ell^{(M)}/2\pi = 1$), the result is a set of bandpowers with a scale-invariant weighting over ℓ .

Comparing Equation (25) to the XFaster estimator of Equation (20) in the ensemble limit (where $\langle \tilde{C}_\ell - \tilde{N}_\ell \rangle \rightarrow \tilde{S}_\ell$ and $\langle q_b \rangle \rightarrow 1$), we find that the window functions for the bandpower

deviations q_b are given by

$$W_{b\ell}^{XY(q)} = \frac{2\pi}{2\ell + 1} \sum_{b'\ell'k} \mathcal{F}_{bb'}^{-1}(2\ell' + 1) g_{\ell'}^k \times \left[\tilde{C}_{\ell'}^{-1} \cdot \frac{\partial \tilde{C}_{\ell'}}{\partial \theta_{b'}} \cdot \tilde{C}_{\ell'}^{-1} \cdot \tilde{M}_{\ell'\ell}^{XY} \right]_{kk}, \quad (26)$$

where all quantities are evaluated at their maximum-likelihood values. Here $\tilde{M}_{\ell'\ell}$ is a matrix of derivatives of the signal model \tilde{S}_ℓ (Equation (6)) with respect to the model spectrum $C_\ell^{(M)}$. Explicitly, each of the spectrum component blocks are

$$\begin{aligned} \tilde{M}_{\ell'\ell}^{TT,ij} &= \delta^{TT} K_{\ell'\ell}^{ij} F_\ell^{TT,ij} (B_\ell^{TT,ij})^2, \\ \tilde{M}_{\ell'\ell}^{TE,ij} &= \delta^{TE} \times K_{\ell'\ell}^{ij} F_\ell^{TE,ij} (B_\ell^{TE,ij})^2, \\ \tilde{M}_{\ell'\ell}^{TB,ij} &= \delta^{TB} \times K_{\ell'\ell}^{ij} F_\ell^{TB,ij} (B_\ell^{TB,ij})^2, \\ \tilde{M}_{\ell'\ell}^{EE,ij} &= (\delta^{EE} + K_{\ell'\ell}^{ij} + \delta^{BB} - K_{\ell'\ell}^{ij}) F_\ell^{EE,ij} (B_\ell^{EE,ij})^2, \\ \tilde{M}_{\ell'\ell}^{BB,ij} &= (\delta^{BB} + K_{\ell'\ell}^{ij} + \delta^{EE} - K_{\ell'\ell}^{ij}) F_\ell^{BB,ij} (B_\ell^{BB,ij})^2, \\ \tilde{M}_{\ell'\ell}^{EB,ij} &= \delta^{EB} (+K_{\ell'\ell}^{ij} - K_{\ell'\ell}^{ij}) F_\ell^{EB,ij} (B_\ell^{EB,ij})^2, \end{aligned} \quad (27)$$

where δ^{XY} is nonzero for only the XY spectrum components of the matrix.

To compute bandpowers C_b from bandpower deviations q_b and their associated window functions, we construct the underlying spectrum C_ℓ in terms of q_b using Equation (5) and normalize Equation (25) with respect to a flat model, so that

$$\begin{aligned} C_b^{XY} &= \sum_B q_B^{XY} \frac{\sum_\ell \mathcal{N}_\ell W_{b\ell}^{XY(q)} \chi_{B\ell}^{XY} C_\ell^{XY(S)}}{\sum_\ell \mathcal{N}_\ell W_{b\ell}^{XY(q)}} \\ &\equiv \sum_B q_B^{XY} \frac{\partial \langle C_b^{XY} \rangle}{\partial q_B^{XY}}, \end{aligned} \quad (28)$$

where the second equality defines the block-diagonal matrix for “rotating” deviations q_b into bandpowers C_b for each spectrum component. The bandpower covariance is then

$$C_{bb'} = \sum_{BB'} \mathcal{F}_{BB'}^{-1} \frac{\partial \langle C_b \rangle}{\partial q_B} \frac{\partial \langle C_{b'} \rangle}{\partial q_{B'}}, \quad (29)$$

where $\mathcal{F}_{BB'}$ is the Fisher information matrix of Equation (19). Finally, we obtain bandpower window functions as

$$W_{b\ell}^{XY(C)} = \sum_B W_{B\ell}^{XY(q)} \frac{\partial \langle C_b^{XY} \rangle}{\partial q_B^{XY}}, \quad (30)$$

with the normalization condition $\sum_\ell \mathcal{N}_\ell W_{b\ell}^{XY(C)} = 1$. Equations (25) and (30) can be used to evaluate a likelihood point for a spectrum C_ℓ from a set of cosmological parameters.

2.6. Parameter Likelihood

The definition of an approximation for the data likelihood (Equation (18)) introduces the possibility of circumventing the power spectrum and associated bandpowers altogether. Since the space of models can be scanned directly as a function of cosmological parameters \mathbf{a} , we can define a likelihood for \mathbf{a} , given the data $\tilde{\mathbf{d}}$, using the Bayesian chain formalism

$$L(\mathbf{a}|\tilde{\mathbf{d}}) \sim P(\mathbf{a})L(\tilde{\mathbf{d}}|\mathbf{a}), \quad (31)$$

where $P(\mathbf{a})$ is the prior in the cosmological parameters.

For the analysis of SPIDER data, the only cosmological parameter allowed to vary is the tensor-to-scalar ratio, r . All other parameters are fixed to the *Planck* best-fit values (Planck Collaboration et al. 2020a). Expanding this parameter space to allow variation in other important parameters, such as the scalar amplitude A_s , is left to future work. The *EE* and *BB* spectra are computed for a scalar-only case with lensing and a tensor-only case with $r = 1$ and tensor spectral index $n_t = 0$. The total CMB power is taken as the sum of the scalar modes with the tensor component scaled linearly with r . This treatment therefore does not assume slow-roll inflation.

The disadvantage of this approach to fitting parameters is that it is still limited by the approximations involved in the definition of the likelihood. However, an important advantage is that it avoids the requirement for defining a bandpower likelihood for model comparison and the approximations associated with that step. Instead, direct evaluation of Equation (31) can be used in numerical searches for maximum-likelihood parameter sets using Markov Chain Monte Carlo (MCMC) techniques. An additional advantage of this direct map-to-parameter likelihood evaluation is that it becomes straightforward to marginalize the final parameter estimates over nuisance parameters such as noise calibrations or even foreground parameters, as we discuss below.

3. Extensions

The XFaster likelihood-based approach can be extended in a number of ways in order to use the estimator for systematic checks, foreground reduction, and component separation. These steps are an important part of any CMB analysis, as they provide robustness and consistency checks, the ability to quantify systematic uncertainties, and the ability to determine the origin of any statistically significant signal. The advantage of likelihood-based estimates is that they always yield an estimate of the Fisher matrix for the parameters. This is useful for establishing the significance of any signal and correctly propagating all uncertainties into the final estimates.

In this section, we show how XFaster has been adapted as an estimator for null test validation and foreground minimization. In both cases, special consideration must be taken to properly treat sample variance, as the XFaster covariance construction assumes that all signal components are Gaussian random fields and thus susceptible to sample variance. This is not appropriate for null tests, where differencing two maps removes the signal component. Likewise, template-based foreground fitting relies on the assumption that the dust morphology is known, and thus sample variance is inappropriate to include. Modifications to the pipeline to address these considerations are detailed below.

3.1. Null Tests

Null tests are often used in CMB analyses to identify systematic noise biases and evaluate the general quality of the data. If undetected, any biases may be misinterpreted as signal. The technique involves taking differences between disjoint subsets of the data and comparing the residual spectra with those from simulations. In this case, we take the differences of maps produced using split halves of the time-stream data. There are a number of ways the splits can be defined to probe different potential systematic effects; the splits used for the SPIDER data set are described in Spider Collaboration (2021). It

is important to use a consistent estimation pipeline for each null split and for signal spectra.

With the XFaster estimator, null spectra can be evaluated using the same method one would use to calculate a signal power spectrum but using the difference of two sets of maps, with some of the components handled differently. The filter transfer function can only be estimated from ensembles where the signal is nonnegligible; it is therefore computed using the ensemble average of the two simulated signal-only half-maps. The model shape spectrum for null tests is flat, and, once the spectrum has converged, the final Fisher matrix for a null spectrum is calculated without the sample variance component. This is done by setting the final signal q_b parameters to a very small value, thus nulling out the signal covariance.

The noise component term for null tests (appearing in both the covariance and data debias terms in Equation (20)) includes both signal and noise residuals, along with their correlations, as these terms all contribute to the expected variance and biases in the data spectra. Unlike for total signal spectra, these terms are included for off-diagonal elements (cross-spectra of different half-maps) of the covariance as well; this is because all auto- and cross-spectra must account for expected residual signal and noise due to different filtering between the half-maps.

As an alternative to debiasing the data for expected signal and noise residuals in spectrum space, the pipeline also has the option to subtract residuals with known morphologies in map space. When such maps are used, the covariance matrix noise term does not include signal contributions, since there is no sample variance in the debias term. For example, to estimate signal residuals for the SPIDER null tests, we use *Planck* frequency maps processed and differenced in the same way as the data null maps instead of using CMB signal simulations. The frequency mismatch between the SPIDER and *Planck* bands could be accounted for but was found to negligibly affect the results. We find that this properly accounts for foreground residuals due to slight differences in the time-domain filtering between the two halves, which dominate our null signal residuals at large scales for some data splits. This method also allows us to accurately model the morphology of the residuals and eliminates the need to account for sample variance from the subtraction in the covariance matrix. The *Planck*-subtraction method has been tested with half-missions and half-rings, with negligible differences between the two. This confirms that the residuals subtracted in this manner are dominated by signal rather than noise in the *Planck* data.

The mode-counting factor g_ℓ is expected to be different for nulls in comparison to total signal spectra due to the different relative contribution of sample variance to the error (Figure 4). Sample variance only affects the null test bandpower errors through its contribution to the uncertainty of the expected signal residual that debiases the data spectrum. The mode-counting factor cannot be empirically calibrated for null tests using signal-only simulations because the remaining signal after debiasing with the expected signal spectrum for a signal-only simulation is nearly zero, which makes the covariance matrix singular and thus noninvertible. Instead, we add noise to the simulations and calibrate g_ℓ iteratively in the same fashion as for total signal spectra.

We have found that the resulting g_ℓ is somewhat sensitive to the noise level used in its calibration. Thus, we use the noise residual terms n_b^i (Equation (15)) calculated for the data to rescale the $a_{\ell m}$ coefficients of the simulated noise maps as

$\sqrt{(1 + n_b^i)} a_{lm}^i$. This modification affects both the $S \times N$ and $N \times N$ terms of the covariance matrices.

3.2. Polarized Foreground Template Fitting

On large angular scales, the polarized CMB signal is obscured by Galactic foregrounds. The contribution from foregrounds biases any estimate of power on the sky with respect to the underlying cosmological signal. This bias must be estimated and removed in order to recover the cosmological signal in a way that minimizes the impact on the final variance. In SPIDER’s observing region in both its 95 and 150 GHz bands, the dominant foreground component is polarized dust (Spider Collaboration 2021). We therefore focus on dust in this section, though the method could also be adapted for other foreground components, such as polarized synchrotron emission.

There are a number of approaches that can be used to remove the foreground bias and include the effect of the removal in the error propagation. If independent observations of the foreground signal exist, a template-subtraction method can be used. In the absence of accurate templates, both the estimate of foreground power and its subtraction from the data must be carried out internally. Here we describe a template-based method implemented as part of the XFaster SPIDER analysis.

We model the dust with a map-space template to include the dust signal in our estimate of CMB bandpowers and parameter likelihoods. We take advantage of the high signal-to-noise ratio measurements of dust by the *Planck* instrument at 217 and 353 GHz, where the dust is much brighter than the CMB (Planck Collaboration et al. 2020c). Because these maps also contain CMB, we subtract *Planck*’s 100 GHz measurement of the sky. Since the *Planck* maps are calibrated against each other using the common CMB component at these frequencies, the residual in the difference map contains only foregrounds and noise. We make the assumption that the morphology of the dust foreground does not depend on frequency, which is consistent with arcminute-resolution observations of polarized dust emission toward diffuse regions of the Milky Way between 353 GHz and 1.2 THz (Ashton et al. 2018). Under this assumption, the dust in the template map can be linearly scaled to match the dust in the SPIDER maps. We label these linear coefficients for the two frequencies as α_{95} and α_{150} .

Once the *Planck* templates are created, they are “reobserved” through the SPIDER pipeline such that they are subjected to the same filtering, beam, and cut-sky effects as the actual observations. We then subtract α -scaled templates from the data spectra using *Planck* half-missions in each half of every cross-correlation so that no *Planck* autocorrelation noise enters the spectra. The cleaned data spectra, $\hat{\Delta}$, are given by

$$\begin{aligned} \hat{\Delta}^{ij} &= \langle (\mathbf{m}_i - \alpha_i \mathbf{t}_i) \times (\mathbf{m}_j - \alpha_j \mathbf{t}_j) \rangle - \alpha_i \alpha_j \langle \mathbf{n}_i^t \times \mathbf{n}_j^t \rangle, \\ &= \hat{\mathbf{C}}^{ij} - \alpha_i \langle \mathbf{t}_i \times \mathbf{m}_j \rangle - \alpha_j \langle \mathbf{m}_i \times \mathbf{t}_j \rangle + \alpha_i \alpha_j \langle \mathbf{t}_i \times \mathbf{t}_j \rangle \\ &\quad - \alpha_i \alpha_j \langle \mathbf{n}_i^t \times \mathbf{n}_j^t \rangle, \end{aligned} \quad (32)$$

where $\langle \dots \times \dots \rangle$ indicates cross pseudospectra of two maps, i and j are map indices, \mathbf{m} represents a SPIDER map, \mathbf{t}_i and \mathbf{t}_j are different half-mission templates reobserved to match the SPIDER map, \mathbf{n}_i^t and \mathbf{n}_j^t are different reobserved half-mission template noise simulations, and α values are the linear scaling factors. Each of the terms in Equation (32) is computed once

and subsequently scaled with α values that are varied with each iteration in the likelihood.

We have chosen to subtract the scaled template from the data rather than to add it to the model covariance, since the latter method would add unnecessary sample variance to the covariance. To account for the error introduced from the template subtraction, such as from *Planck* noise or chance correlations between the CMB and the template, we run the algorithm on an ensemble of simulations.

The simulations include CMB signal realizations, SPIDER noise, a foreground template (either directly using a *Planck* template or with a scaled Gaussian realization of a power law), and *Planck* noise maps from the FFP10 simulation ensemble.³⁶ The distribution of parameters determined for this ensemble is taken to be the true covariance. The average covariance of the XFaster algorithm estimates is the total covariance without error from template fitting. The difference of the two is taken to be the additional contribution to the covariance due to the foreground subtraction and is added to each Monte Carlo sample in the data parameter likelihoods.

The amplitude of this additional covariance is found to be independent of simulated r and foreground morphology. However, it scales with α , as expected, since α scales the *Planck* noise contribution. Because of this, we compute the term with the precise data-preferred α values applied to the simulations.

The *Planck* FFP10 noise simulations are correlated with each other by a nonnegligible amount due to the method by which they were produced. To account for these correlations, we include the scaled average of *Planck* half-mission 1 \times half-mission 2 noise in the terms subtracted from the data.

3.3. Harmonic-domain Foreground Fitting

An alternative method of foreground fitting involves estimating the contribution of frequency-dependent foregrounds without requiring a map-space template. This method necessarily relies on an assumed model for the frequency dependence of the contributing foreground. The model can be included in either the map estimation step as a contribution to the model for the observed data or in the power spectrum estimation step as a contribution to the sample variance in the likelihood. Here we describe an extension to XFaster using the latter.

We model the spectral energy density (SED) of the dust component as

$$S_d(\nu) = A_d \frac{B_\nu(T_d)}{B_{\nu_0}(T_d)} \left(\frac{\nu}{\nu_0} \right)^{\beta_d - 2}, \quad (33)$$

where T_d is the blackbody temperature of the dust, A_d is the model amplitude at reference frequency ν_0 , $B_\nu(T)$ is the blackbody spectrum at temperature T , and β_d is a spectral index. The SED describes the brightness temperature of the dust. To relate this to the representation of CMB maps using a thermodynamic temperature Θ , in which blackbody sources are frequency-independent, we use the idealized conversion factor

$$\Theta_d(\nu) = \frac{(e^x - 1)^2}{x^2 e^x} S_d(\nu) \equiv g(\nu) S_d(\nu), \quad (34)$$

³⁶ *Planck* end-to-end “full focal plane” simulations (Planck Collaboration et al. 2020b).

where $x = h\nu/kT_{\text{CMB}}$. In practice, the $g(\nu)$ factors must be corrected for the specific frequency dependence of experimental window functions. These factors are color-corrected for the SPIDER data by integrating over the 95 and 150 GHz window functions (Shaw et al. 2020).

The dust contribution to the model cut-sky covariance (Equation (5)) is given by

$$\widetilde{S}_\ell^{d,ij} = \frac{\Theta_d(\nu_i)\Theta_d(\nu_j)}{\Theta_d^2(\nu_0)} \sum_{(\text{dust})b} q_b^d \widetilde{C}_{bl}^{d,ij}. \quad (35)$$

We have introduced a set of dust bandpower parameters q_b^d , with b a set of bands defined specifically for the dust component. We use bandpower kernels $\widetilde{C}_{bl}^{d,ij}$ —including all kernel, filter, and beam terms as in Equation (7)—in which we parameterize the full-sky dust shape spectrum as $\ell(\ell+1)C_\ell^{d(S)}/2\pi = A(\ell/80)^{\alpha+2}$. The amplitude A and angular spectral index α are set to the best-fit values reported in Planck Collaboration et al. (2020c). The bandpower parameters q_b^d are varied simultaneously with all other bandpower parameters to find a global maximum-likelihood fit for the combination of, e.g., CMB, dust, and noise residual bandpowers for each polarization combination. Since the dust bandpowers appear at linear order in the covariance, the estimator in Equation (20) is unchanged as long as other parameters, such as T_d, β_d , are fixed. If the assumed values for the additional parameters are not correct, the effect will be incorporated into the bandpower parameters q_b^d , and these can be regarded as the “effective” rescalings of the model.³⁷ We explore the use of harmonic space foreground fitting in future work.

4. Pipeline Validation

To validate the pipeline, we conduct a series of tests with simulated inputs that each require two criteria to be met.

1. The ensemble-averaged parameter estimates produced by XFaster must match the input values used to generate the simulations.
2. The error in the estimate must match the scatter of individual XFaster estimates over the ensemble.

These conditions, achieved to within set tolerance levels, ensure that the XFaster estimates are unbiased in both mean and variance. In practice, any recalibration of the Fisher matrix required to satisfy the second condition minimizes the effect of the approximation used to define the XFaster likelihood. This is achieved at the cost of some additional simulations used for the effective calibration of the likelihood for specific data sets. The pipeline must also behave well for reasonable changes to the simulation inputs, such as different cosmological parameters, foreground morphologies, or noise amplitudes within expected ranges for the data.

In this section, we first describe the simulation inputs, followed by the results of validation tests for each of the pipeline outputs.

³⁷ An alternative is to include a nonlinear maximization of the likelihood over, e.g., β_d at each step of the quadratic iteration.

4.1. Simulated Maps

Simulation ensembles include a CMB signal component, an instrument noise component, and (optionally) a foreground signal component. Each of these are described in detail in the following sections.

4.1.1. CMB Maps

CMB maps are produced using the `synfast` HEALPix routine (Gorski et al. 2005), which generates Gaussian realizations of input angular power spectra. The power spectra are computed by the software package `CAMB` (Lewis et al. 2000) from a fiducial set of Λ CDM best-fit parameters to the *Planck* data (Planck Collaboration et al. 2016). To generate appreciable signal for computing the transfer function, we construct the simulation ensemble by replacing the fiducial Λ CDM input BB spectrum with a spectrum that is flat in $\ell(\ell+1)C_\ell$ with amplitude $0.2 \mu\text{K}^2$.

Each signal realization is generated at HEALPix resolution $N_{\text{side}} = 2048$, then smoothed with a SPIDER beam per focal plane. Each realization of the Stokes T, Q , and U maps is then “reobserved” through the SPIDER mapmaking pipeline, applying all flagging and time-domain filtering identically to the measured data. The T component of each input realization is then passed through the filtering and mapmaking pipeline a second time. The resulting Q and U maps then contain only the temperature-to-polarization leakage induced by the pipeline and are subtracted from the first set of reobserved Q and U maps to remove any bias due to this leakage.

This leakage-subtraction procedure was deemed prohibitively costly for null tests, since signal simulations must be made for each of 10 null splits. Instead, the null signal simulations use CMB maps with the T map replaced with a *Planck* map at the frequency closest to its SPIDER analog and the E -mode power constrained by the TE cross-correlation spectrum. This means the T -to- P leakage map is the same for every seed in the signal ensemble. Because null tests use the difference of maps with approximately the same leakage correction, this further approximation to the leakage correction negligibly affects the results.

We use 1000 unconstrained CMB realizations for computing bandpowers and likelihoods and 500 constrained CMB realizations for computing bandpowers for each of 10 null splits. To preserve any data correlations between null tests that are present in the measured data, we use the same set of random number generator seeds for the signal components going into each null split.

4.1.2. Noise

Noise simulations are generated in the time domain by sampling from a single power spectral density (PSD) per channel. The input PSD is constructed by averaging over the entire flight in 10 minute chunks; in practice, this overestimates the SPIDER noise due to the asymmetric impact of high outliers in the distribution. As with the signal simulations, the noise realization for each channel uses the same set of random number generator seeds for each null split in order to preserve noise correlations that are similarly present in the data.

4.1.3. Foregrounds

We require simulated foreground templates to test the template-cleaning method of foreground separation. Each simulated template is constructed from a fiducial foreground template with added instrument noise from the *Planck* FFP10 simulation ensemble, where both the templates and noise realizations are reobserved with the SPIDER mapmaking pipeline. The foreground template ensemble is limited by the size of the FFP10 data set to 300 such realizations. The fiducial foreground template is constructed by differencing two *Planck* frequency maps as usual, e.g., *Planck* 353 minus 100 GHz. The fiducial template also includes *Planck* instrument noise by construction, but it is our best estimate of the morphology of the true foregrounds. Alternatively, we can construct the fiducial template using a Gaussian realization of a power law matching the *Planck* best-fit dust power-law model. We find that these two foreground simulation methods behave equally well in the validation results.

4.2. Bandpower Validation

To validate the XFaster bandpower estimation pipeline, we compute bandpowers for an ensemble of 500 CMB+noise simulations for both signal and null spectra. We verify that the average of the computed signal bandpowers matches the spectrum input to generate the CMB maps and is therefore unbiased, and that the average of the null bandpowers is consistent with zero. These bandpower distributions are shown in Figure 5.

To verify that the covariance computed by XFaster is accurate, we check that the scatter of the ensemble matches the covariance computed by XFaster for both signal and null simulations. We construct two distributions of χ^2 values:

1. the χ^2 of the spectrum for each of 500 realizations relative to the corresponding input model (fiducial Λ CDM for the signal ensemble or a null model for the null ensemble) and
2. the χ^2 of each of 100,000 spectrum realizations, sampled from the average of the bandpower covariance matrices computed for each of the realizations in the signal and null map ensemble.

The second distribution forms the expectation for the first distribution; as shown in Figure 6, the two distributions are in good agreement. The covariance is slightly overestimated in *EE* and *BB* for signal spectra, indicating that the error bars for these data spectra will be somewhat larger than is optimal. These results are insensitive to input signal shape and noise amplitude.

4.3. Likelihood Validation

The likelihood pipeline was tested on ensembles of CMB+noise+dust “fake data” maps and dust+*Planck* noise “fake template” maps. The parameters r , α_{95} , and α_{150} were varied using an MCMC sampler, and all other parameters were fixed to the *Planck* best-fit values from Table 1 in Planck Collaboration et al. (2020a). The pipeline was deemed to be validated if the output parameter likelihoods were unbiased with respect to the input parameters and if the likelihood widths matched the scatter of the best-fit parameters from the ensemble. The pipeline was tested using various values of r ranging from zero to 0.7, various noise amplitudes, and various

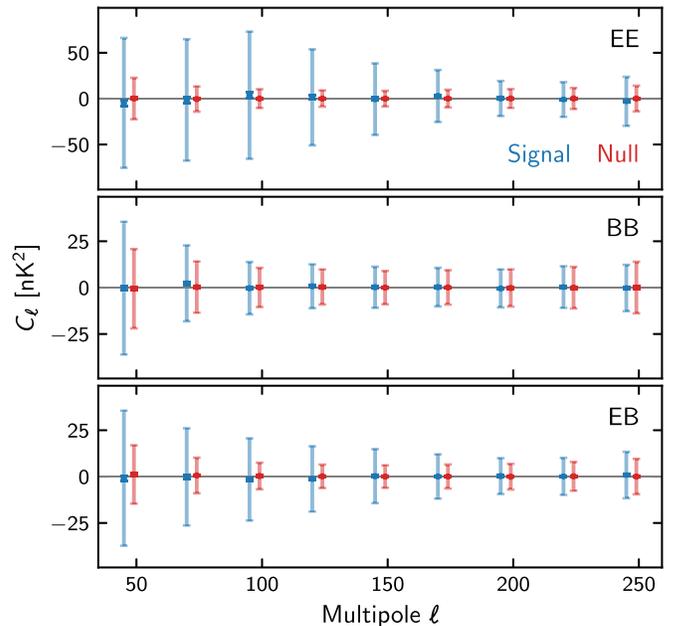


Figure 5. Residual spectra averaged over an ensemble of 500 signal (blue) and null test (red) simulations differenced with the corresponding input model spectrum. Dark error bars show the error on the mean. Light error bars indicate the total spectrum error and excluded for null spectra. The null test shown is the checkerboard split (see Spider Collaboration 2021 for null split definitions), though all other splits show similar results.

template morphologies. In addition, we fit for and marginalized over beam uncertainty and noise residual amplitudes; these nuisance parameters were found to have a negligible effect on the data posteriors for a SPIDER-like data set. The results for the nominal simulation data set are shown in Figure 7. The output parameter likelihoods are not biased with respect to the input parameters, and the likelihood widths match the scatter of the best-fit parameters from the ensemble, indicating that the XFaster likelihood is accurate and unbiased.

The recovered α_{95} and α_{150} parameters indicate a slight bias in the estimator, as evidenced by comparing the 1σ error on the distribution mean to the simulation input values. However, these biases are small with respect to the error on the parameters and evidently do not impact the cosmological result, as the recovered maximum-likelihood r value is within 1σ error on the mean (0.011) of the input r to the simulation.

5. Application to SPIDER Data

After extensive validations of the XFaster pipeline using simulations, the analysis was applied to the data from the 2015 flight of the SPIDER instrument. The comprehensive results of this analysis, using XFaster and additional pipelines, are presented in Spider Collaboration (2021). A subset of the XFaster results are reproduced here as a demonstration of the pipeline functionality.

5.1. Null Tests

Ten different null tests were carried out on the SPIDER data set, including seven detector-based splits and three time-based splits. Nine of the splits were designed to probe potential systematic errors in the data; one, the checkerboard pattern detector split, was chosen to be largely insensitive to systematics, instead probing the pipeline’s handling of the

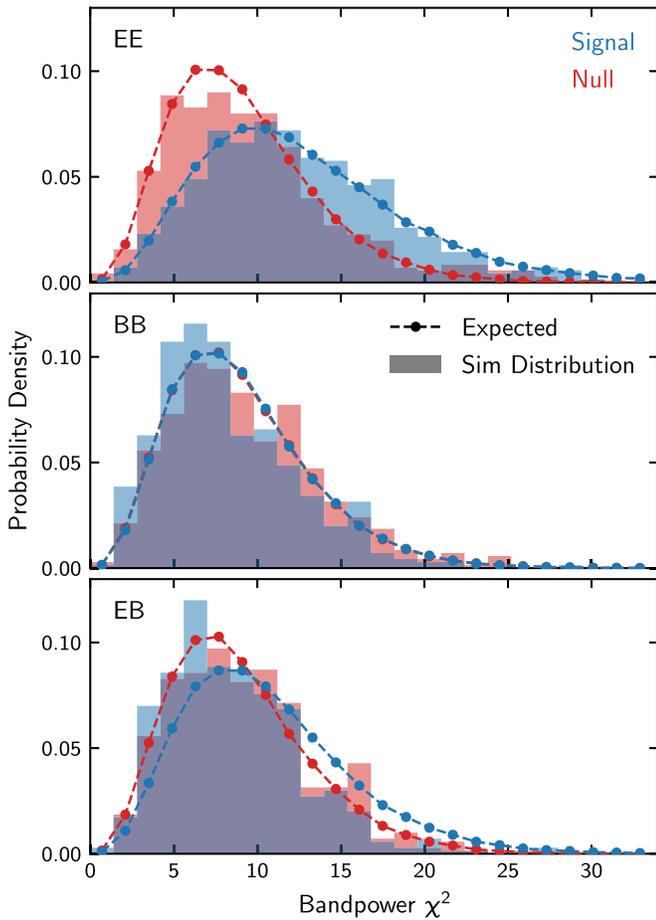


Figure 6. Distributions of bandpower χ^2 estimates for an ensemble of 500 simulations (histograms) compared to expectations from the corresponding covariance (dotted lines). Results are shown separately for CMB signal simulations (blue) and the checkerboard null split (red), which is representative of the other nine null splits. The χ^2 values are calculated for nine multipole bins in the range $33 < \ell < 257$ using a fiducial Λ CDM model for the signal ensemble and a null model for the null split ensemble. The expectation distribution is determined by generating 100,000 realizations of bandpowers from the Fisher covariance matrix and histogramming the resulting χ^2 values. Agreement between the histogram and expectation lines indicates that the covariance matrix is accurate. The covariance is slightly overestimated in *BB* and *EB* signal, producing lower χ^2 values than expected from random realizations of the covariance matrix.

noise model. The full list of splits is detailed in Spider Collaboration (2021). The estimated null spectra for some representative subsets of data splits are shown in Figure 8.

By construction, different null splits contain overlapping detector samples and are therefore correlated at some level. For example, the SPIDER “inner row” null half overlaps by 75% with the “inner radius” null half. This creates a challenge in assessing the statistics of the full null ensemble. These correlations are preserved among simulated maps generated using the same random number generator seed, so we can use simulations to inform our expectations for distribution shape statistics.

When such correlations are negligible, the expected distribution of the null χ^2 values can be determined by drawing simulated null bandpowers from the covariance matrix. We can then perform a Kolmogorov–Smirnov (KS) test that computes the likelihood that the data are drawn from the cumulative distribution extrapolated from the simulated χ^2 values. The distributions used in this test for the SPIDER

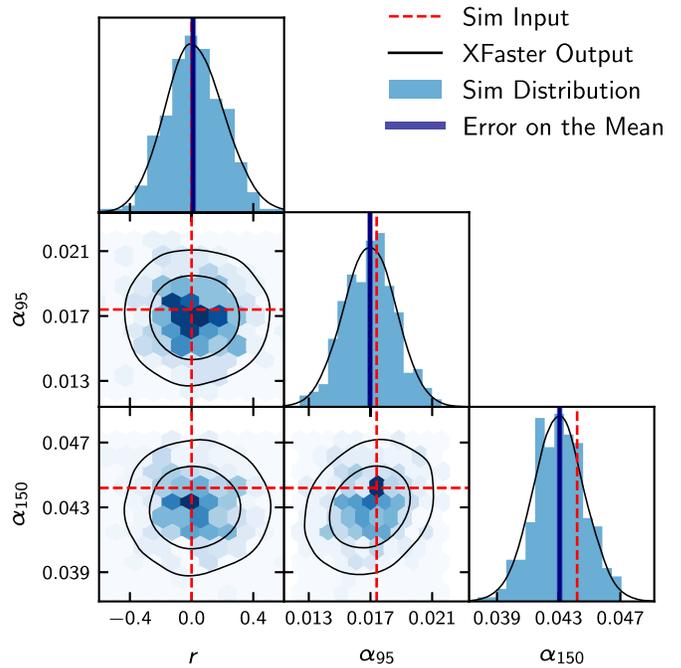


Figure 7. The XFastest likelihood results for an ensemble of 300 CMB+noise+foreground simulations. Red dotted lines indicate the input r and foreground template scalings, α_{95} and α_{150} , used to make simulated maps. The XFastest likelihood contours computed for the ensemble mean of the simulations are shown in black, and the maximum-likelihood parameter values computed for each realization in the ensemble are histogrammed in blue, with dark blue shading to indicate the mean and error on the mean of the distribution. The agreement between the histograms and the black contours shows that the XFastest pipeline is unbiased in the mean and its estimate of error for r . The biases in α , most evident at 150 GHz, are small compared to the error and do not affect the cosmological result.

combined data set are shown in the top panel of Figure 9; the resulting KS test p -value is 0.54, indicating good agreement.

To account for the presence of correlations, we then perform the same test on each of 500 end-to-end simulations. We compare the KS test p -value from data to the distribution of the simulation KS test p -values to test how unlikely our data are relative to simulations. These results are shown in the bottom panel of Figure 9, indicating that the SPIDER data exceeded the criterion of being at least 1% likely given the simulation ensemble. We further test the outliers of the data distribution by comparing the maximum data χ^2 to the maximum χ^2 in each end-to-end simulation. We again find that the outlier data χ^2 values are more than 1% likely, as shown in the middle panel of Figure 9.

The flexibility and reduced computational overhead introduced by the XFastest approximate-likelihood method is a key advantage of the pipeline. It allows a full exploration of null testing splits, internal calibration of systematics, and determination of goodness-of-fit statistics for any estimated quantity.

5.2. Bandpower Estimates

The XFastest procedure applied to the SPIDER data set requires a set of ~ 100 parameters, including both signal and noise components, and as many as ~ 300 parameters when the harmonic-domain foreground model is also included. We construct the XFastest estimator to compute 16 bandpowers, equally spaced between $\ell = 8$ and 407, for each of six

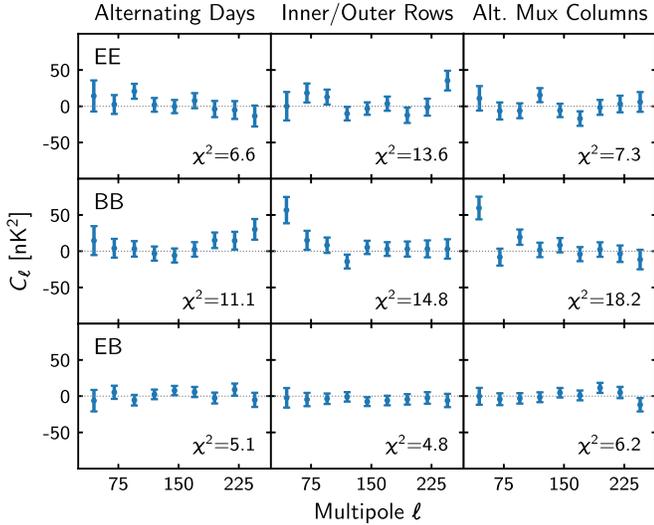


Figure 8. Three SPIDER null tests, showing *EE*, *BB*, and *EB* spectra for the combined 95 and 150 GHz data, with χ^2 values computed for the bins shown.

polarization spectra, with additional noise calibration residual bandpowers defined over the same bins. We have found that the choice of noise bin width does not significantly affect the result, as the SPIDER noise model is not very degenerate with the expected signal. The nine bins between $\ell = 33$ and 257 are used for subsequent analyses; the other seven bins are included in the estimator to accurately account for their contribution to the so-called science bins through bin-to-bin leakage. All spectra contain information from modes at $\ell \lesssim 8$ that are so heavily correlated due to the reduced sky coverage that the XFaster approximation breaks down, even after calibration of the effective mode count. To reduce the sensitivity to these correlations, we exclude all modes at $\ell < 8$ in all spectral decompositions. The second bin ($8 \leq \ell < 33$) is included in the iterative bandpower estimation procedure but discarded in later analysis steps due to the level of correlation with the unconstrained lowest multipoles.

Convergence of the XFaster estimator is reached at iteration i when

$$\max_b \left| \frac{\theta_b^{(i)} - \theta_b^{(i-1)}}{\theta_b^{(i-1)}} \right| < 0.005, \quad (36)$$

i.e., when all model parameters change by less than 0.5% relative to their values at the previous iteration. This requires $\mathcal{O}(10)$ iterations of the quadratic estimator in Equation (20).

Figure 10 shows the estimated bandpowers computed from the raw SPIDER data set (i.e., making no attempt to remove any foreground contamination), as well as the bandpowers obtained using the template-based dust-subtraction method of Section 3.2. The bandpower error bars are obtained from the diagonal of the bandpower covariance (Equation (29)), computed from the inverse of the Fisher matrix after discarding all nuisance parameters (noise residuals, in this case). Spectra are computed using a data set containing both 95 and 150 GHz observations with a common sky mask. The results show significant dust power in the raw spectra at large angular scales. Dust-cleaned spectra are in good agreement with expectations from Λ CDM, and the contribution to the angular power

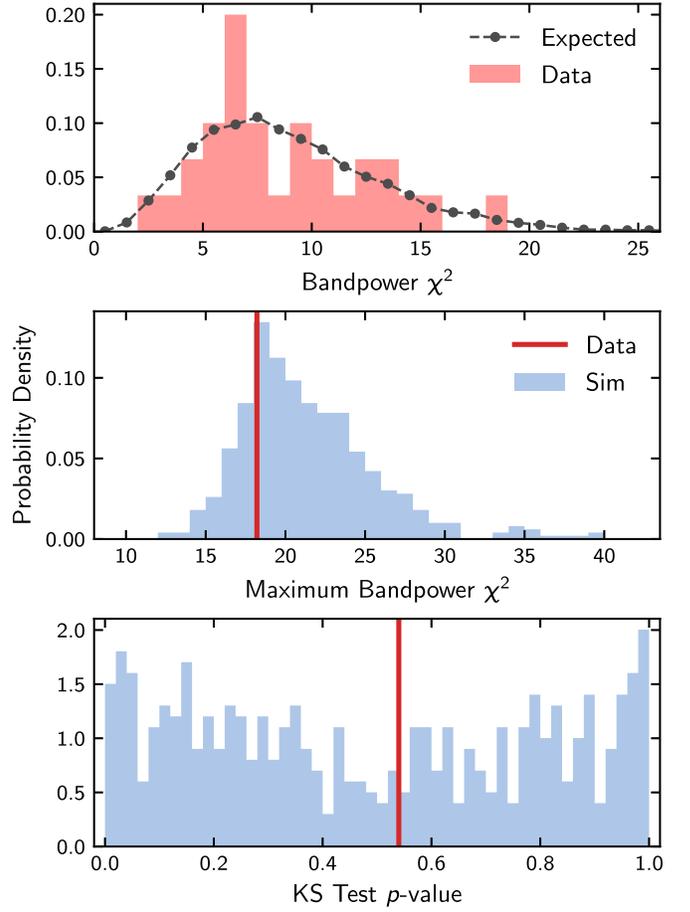


Figure 9. Illustration of null statistics computed from the combined SPIDER data set. (Top) Histogram of all null test bandpower χ^2 values and expected distribution of drawn χ^2 from the covariance matrix. The red histogram consists of 30 χ^2 values (three polarization spectra for 10 tests). The expectation histogram (dashed line) consists of 5000 random bandpower draws from each covariance matrix for each test, used to compute a total of 150,000 χ^2 values. A KS test is performed between the data and expectation distributions, resulting in a p -value of 0.54. (Middle) Histogram of the maximum bandpower χ^2 value from 500 end-to-end simulations. The maximum χ^2 of the data is 18.2, corresponding to a probability to exceed (PTE) of 0.78 given the ensemble of max- χ^2 values. (Bottom) Histogram of the KS test p -values from the same set of 500 end-to-end simulations. The p -value of the data (red line) corresponds to a PTE of 0.55 given the ensemble of KS tests. Together, these tests indicate that the combined SPIDER data set passes its suite of null tests.

spectrum from inflationary gravitational waves, parameterized by r , is not detected.

5.3. Likelihood Estimates

Finally, we demonstrate the application of XFaster as a direct map-to-parameter likelihood estimator for the SPIDER data set. The likelihood model includes foreground templates constructed from the *Planck* 353–100 GHz difference map parameterized with 95 and 150 GHz scalings, α_{95} and α_{150} ; a CMB signal component parameterized by the tensor-to-scalar ratio r and all other Λ CDM parameters held fixed; and rescaled noise components as determined by the iterative fitting procedure. The noise model parameters are marginalized over for the final result, along with uncertainties in the beam model; these nuisance parameters do not contribute significantly to the uncertainties on the cosmological and foreground parameters.

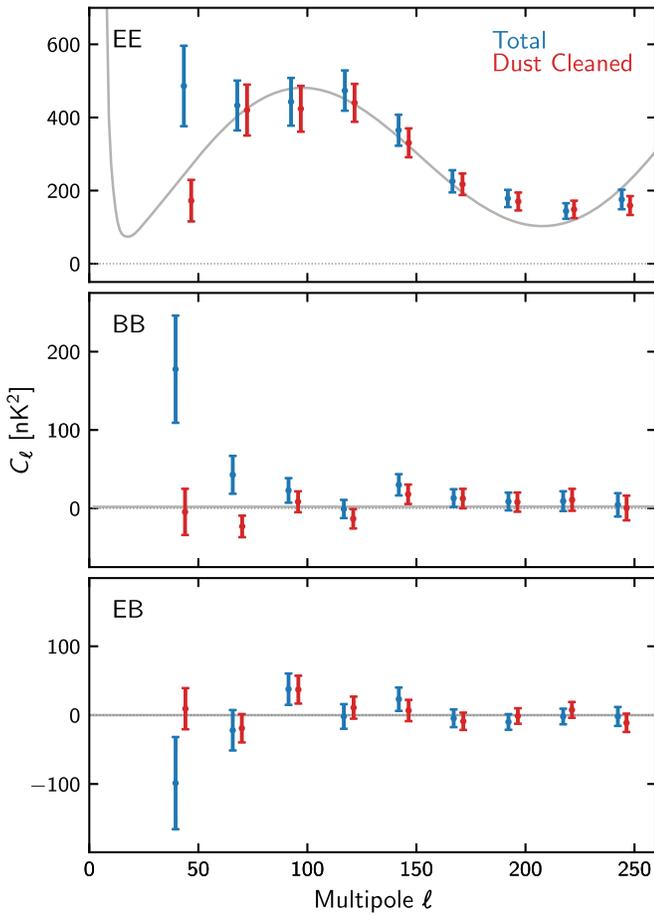


Figure 10. SPIDER bandpower estimates for total power (blue) and CMB only (red), computed from a combination of 95 and 150 GHz data. All error bars include contributions from the sample variance of the total estimated power along with instrument noise, and foreground-cleaned error bars additionally include errors from template fitting. A fiducial Λ CDM CMB spectrum with $r = 0$ is shown in gray.

Profile likelihoods in the directions along r , α_{95} , and α_{150} are shown in Figure 11.

Further discussion of this result, including construction of an appropriate upper limit on the tensor-to-scalar ratio r , can be found in Spider Collaboration (2021).

6. Using the Code

The pipeline is written entirely in Python and available on GitHub³⁸ (Gambrel et al. 2021). While it is possible to run the code on a single processor, it is greatly sped up by parallelizing matrix operations with OpenMP. Further speed gains could be achieved by using MPI or high-throughput computing to distribute pseudospectrum computation across multiple processors, but this functionality is not currently implemented. The code base also includes tools for submitting individual jobs on computing clusters using the Slurm workload manager.³⁹

The most time-consuming step in the analysis process is computing the pseudospectra of each of the simulated map crosses using the *anafast* method of the *healpy* package. However, this step only needs to be performed once per mask choice, as its results are stored to disk and read in for

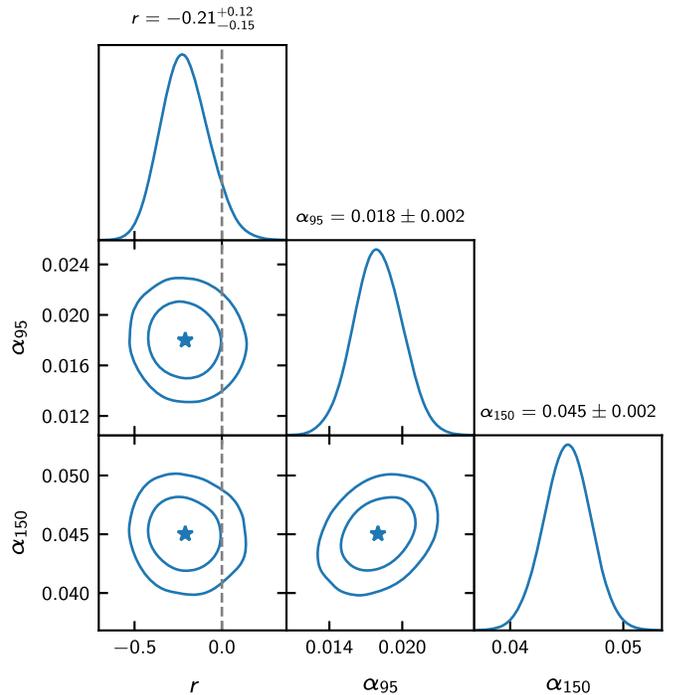


Figure 11. Combined XFaster likelihood for r and α , imposing no priors on these parameters and using a *Planck* 353–100 GHz template to estimate the foreground morphology. The 1σ constraints are shown in the panel titles.

bandpower and likelihood computations. The XFaster method applied to the SPIDER data set included 1000 simulations, $\ell_{\max} = 407$, $N_{\text{side}} = 512$, and eight independent maps (four at 95 and four at 150 GHz), which required approximately 5 hr on 20 cores and 2 GB of memory to complete the bandpower computation from the individual maps on disk. By contrast, the longest step of the full SPIDER analysis pipeline—generating the ensemble of 1000 signal and noise simulation maps—requires approximately 45 core yr.

The sequential steps of the XFaster algorithm are listed in Table 1, along with the time and memory required for each step. Intermediate results are written to disk after each step is completed, so that subsequent steps may be run starting from that checkpoint. For example, it is typical to construct the mask and simulation ensemble spectra once, then adjust the bins and nuisance parameters for the particular application. Due to the modular structure of the code, the total runtime required is the sum of the rows (modified by the number of CPUs provided for each step), and the memory required is simply the maximum among the rows, or 2.2 GB for the SPIDER fiducial case.

7. Conclusion

We have presented the XFaster power spectrum and parameter-likelihood estimation package and demonstrated its validation with simulations and application to the SPIDER 2015 data set. XFaster builds upon the MASTER formalism for estimating full-sky CMB power spectra and covariances, accounting for filtering and noise biases using an ensemble of simulations. Unlike the MASTER method, it estimates the covariance of bandpowers using an iterative calibration of the Fisher matrix and therefore only requires one set of signal and noise simulations that do not need to be precise representations of the data. The result is a pipeline that can produce fast, accurate power spectra and likelihoods for cosmological

³⁸ <https://github.com/SPIDER-CMB/xfaster>

³⁹ <https://slurm.schedmd.com/>

Table 1
Pipeline Steps and Computing Requirements along with Their Order-of-magnitude Scaling for Relevant Pipeline Parameters

Pipeline Step	CPU Time	Time Scaling	Memory (GB)	Memory Scaling	OMP Speed Up
Mask cross-spectra	1 minute	$N_{x\text{-spec}}, N_{\text{pix}}, \ell_{\text{max}}^2$	1.3	$N_{\text{map}}, N_{\text{pix}}, \ell_{\text{max}}$	Minimal
Mode-coupling kernel	6 minutes	$N_{x\text{-spec}}, \ell_{\text{max}}^2$	0.75	$N_{x\text{-spec}}, \ell_{\text{max}}^2$	None
Pseudospectra of simulated maps	16 hr	$N_{\text{sim}}, N_{x\text{-spec}}, N_{\text{pix}}, \ell_{\text{max}}^2$	2.2	$N_{\text{map}}, N_{\text{pix}}, \ell_{\text{max}}$	$\sqrt{\text{CPU}}$
Filter transfer function	10 s	$N_{x\text{-spec}}, \ell_{\text{max}}^2$	0.8	$N_{x\text{-spec}}, \ell_{\text{max}}^2$	None
Pseudospectra of data maps	100 s	$N_{x\text{-spec}}, N_{\text{pix}}, \ell_{\text{max}}^2$	1.3	$N_{\text{map}}, N_{\text{pix}}, \ell_{\text{max}}$	$\sqrt{\text{CPU}}$
Bandpowers	3 minutes	$N_{x\text{-spec}}, \ell_{\text{max}}^2$	1.9	$N_{x\text{-spec}}, \ell_{\text{max}}^2$	Minimal
Likelihoods	5 hr	$N_{\text{param}}, N_{x\text{-spec}}, \ell_{\text{max}}$	1.9	$N_{\text{param}}, N_{x\text{-spec}}, \ell_{\text{max}}$	None

Note. Fiducial numbers use eight $N_{\text{side}} = 512$ maps, $\ell_{\text{max}} = 407$, and a single core. Where changing these quantities affects the time or memory requirements, they are listed with their approximate scaling. Here N_{map} is the number of data maps, $N_{x\text{-spec}} = N_{\text{map}}(N_{\text{map}} + 1)/2$ is the number of cross-spectra, N_{pix} is the number of pixels in each map, N_{sim} is the number of simulations (assumed to be the same for signal and noise), ℓ_{max} is the maximum multipole used, and N_{param} is the number of parameters solved for in the likelihood.

parameters. It is additionally capable of computing null spectra and fitting for Galactic foregrounds, all within a self-consistent, self-contained framework. This pipeline has been thoroughly validated with simulations of the SPIDER data set and is publicly available for use on other CMB data sets.

We acknowledge the contribution to the development of the XFaster pipeline by all members of the BOOMERanG, *Planck*, and SPIDER collaborations. SPIDER is supported in the U.S. by the National Aeronautics and Space Administration under grants NNX07AL64G, NNX12AE95G, and NNX17AC55G issued through the Science Mission Directorate and by the National Science Foundation through PLR-1043515. Logistical support for the Antarctic deployment and operations is provided by the NSF through the U.S. Antarctic Program. Support in Canada is provided by the Natural Sciences and Engineering Research Council and the Canadian Space Agency. Support in Norway is provided by the Research Council of Norway. Support in Sweden is provided by the Swedish Research Council through the Oskar Klein Centre (contract No. 638-2013-8993), as well as a grant from the Swedish Research Council (dnr. 2019-93959) and a grant from the Swedish Space Agency (dnr. 139/17). The Dunlap Institute is funded through an endowment established by the David Dunlap family and the University of Toronto. The multiplexing readout electronics were developed with support from the Canada Foundation for Innovation and the British Columbia Knowledge Development Fund. A.E.G. is supported by the Kavli Institute for Cosmological Physics at the University of Chicago through an endowment from the Kavli Foundation and its founder, Fred Kavli. C.R.C. was supported by UKRI consolidated grants ST/P000762/1, ST/N000838/1, and ST/T000791/1. K.F. holds the Jeff and Gail Kodosky Endowed Chair at UT Austin and is grateful for that support. W.C.J. acknowledges the generous support of the David and Lucile Packard Foundation, which has been crucial to the success of the project.

Some of the results in this paper have been derived using the HEALPix package (Gorski et al. 2005). The computations described in this paper were performed on four computing clusters: Hippo at the University of KwaZulu-Natal, Feynman at Princeton University, and the GPC and Niagara supercomputers at the SciNet HPC Consortium (Loken et al. 2010; Ponce et al. 2019). SciNet is funded by the Canada Foundation for Innovation under the auspices of Compute Canada, the

Government of Ontario, Ontario Research Fund–Research Excellence, and the University of Toronto.

The collaboration is grateful to the British Antarctic Survey, particularly Sam Burrell, and the Alfred Wegener Institute and the crew of R.V. *Polarstern* for invaluable assistance with the recovery of the data and payload after the 2015 flight. Brendan Crill and Tom Montroy made significant contributions to SPIDER’s development. This project, like so many others that he founded and supported, owes much to the vision and leadership of the late Professor Andrew E. Lange.

ORCID iDs

- A. S. Rahlin  <https://orcid.org/0000-0003-3953-1776>
C. R. Contaldi  <https://orcid.org/0000-0001-7285-0707>
S. J. Benton  <https://orcid.org/0000-0002-4214-9298>
J. R. Bond  <https://orcid.org/0000-0003-2358-9949>
H. C. Chiang  <https://orcid.org/0000-0002-4098-9533>
A. J. Duivenvoorden  <https://orcid.org/0000-0003-2856-2382>
J. P. Filippini  <https://orcid.org/0000-0001-8217-6832>
R. Gualtieri  <https://orcid.org/0000-0003-4245-3215>
J. E. Gudmundsson  <https://orcid.org/0000-0003-1760-0355>
M. Halpern  <https://orcid.org/0000-0002-1760-0868>
W. C. Jones  <https://orcid.org/0000-0002-3636-1241>
J. S.-Y. Leung  <https://orcid.org/0000-0001-7116-3710>
S. Li  <https://orcid.org/0000-0002-8896-911X>
J. M. Nagy  <https://orcid.org/0000-0002-2036-7008>
I. L. Padilla  <https://orcid.org/0000-0002-0024-2662>
T. M. Ruud  <https://orcid.org/0000-0002-7615-1900>
E. C. Shaw  <https://orcid.org/0000-0001-5644-8750>
J. D. Soler  <https://orcid.org/0000-0002-0294-4465>
I. K. Wehus  <https://orcid.org/0000-0003-3821-7275>

References

- Ashton, P. C., Ade, P. A. R., Angilè, F. E., et al. 2018, *ApJ*, 857, 10
Bond, J. 1993, Proc. Les Houches Summer School on Cosmology and Large Scale Structure (Session 60) (Les Houches, France), 469
Bond, J. R., Jaffe, A. H., & Knox, L. 1998, *PhRvD*, 57, 2117
Bond, J. R., Jaffe, A. H., & Knox, L. 2000, *ApJ*, 533, 19
Borrill, J. 1999, arXiv:astro-ph/9911389
Challinor, A., & Chon, G. 2005, *MNRAS*, 360, 509
Chiang, H. C., Ade, P. A. R., Barkats, D., et al. 2010, *ApJ*, 711, 1123
Gambrel, A., Rahlin, A., & Contaldi, C. 2021, XFaster, doi:10.5281/zenodo.5248317
Gjerløw, E., Colombo, L., Eriksen, H., et al. 2015, *ApJS*, 221, 5
Gorski, K. M., Hivon, E., Banday, A., et al. 2005, *ApJ*, 622, 759

- Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, [ApJ](#), **567**, 2
- Jones, W. C., Ade, P. A. R., Bock, J. J., et al. 2006, [ApJ](#), **647**, 823
- Knox, L. 1999, [PhRvD](#), **60**, 103516
- Lewis, A., Challinor, A., & Lasenby, A. 2000, [ApJ](#), **538**, 473
- Loken, C., Gruner, D., Groer, L., et al. 2010, [JPhCS](#), **256**, 012026
- Montroy, T. E., Ade, P. A. R., Bock, J. J., et al. 2006, [ApJ](#), **647**, 813
- Netterfield, C. B., Ade, P. A. R., Bock, J. J., et al. 2002, [ApJ](#), **571**, 604
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, [A&A](#), **594**, A13
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, [A&A](#), **641**, A6
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020b, [A&A](#), **641**, A3
- Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020c, [A&A](#), **641**, A4
- Ponce, M., van Zon, R., Northrup, S., et al. 2019, in Proc. of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning), PEARC '19 (New York: Association for Computing Machinery)
- Rocha, G., Contaldi, C. R., Bond, J. R., & Górski, K. M. 2011, [MNRAS](#), **414**, 823
- Shaw, E. C., Ade, P. A. R., Akers, S., et al. 2020, [Proc. SPIE](#), **11453**, 114532F
- Spider Collaboration 2021, arXiv:2103.13334