



A Simulation-based Method for Correcting Mode Coupling in CMB Angular Power Spectra

J. S.-Y. Leung^{1,2} , J. Hartley³, J. M. Nagy^{4,5} , C. B. Netterfield^{1,3}, J. A. Shariff⁶, P. A. R. Ade⁷ , M. Amiri⁸, S. J. Benton⁹ , A. S. Bergman⁹, R. Bihary¹⁰, J. J. Bock^{11,12}, J. R. Bond⁶ , J. A. Bonetti¹², S. A. Bryan¹³, H. C. Chiang^{14,15} , C. R. Contaldi¹⁶ , O. Doré^{11,12} , A. J. Duivenvoorden⁹ , H. K. Eriksen¹⁷, M. Farhang^{1,6,18}, J. P. Filippini^{19,20} , A. A. Fraisse⁹, K. Freese^{21,22}, M. Galloway¹⁷, A. E. Gambrel²³, N. N. Gandilo²⁴, K. Ganga²⁵, R. Gualtieri²⁶ , J. E. Gudmundsson²² , M. Halpern⁸ , M. Hasselfield²⁷, G. Hilton²⁸, W. Holmes¹², V. V. Hristov¹¹, Z. Huang⁶ , K. D. Irwin^{29,30}, W. C. Jones⁹ , A. Karakci¹⁷, C. L. Kuo²⁹, Z. D. Kermish⁹, S. Li^{9,31} , D. S. Y. Mak¹⁶, P. V. Mason¹¹, K. Megerian¹², L. Moncelsi¹¹ , T. A. Morford¹¹, M. Nolta⁶, R. O'Brient¹², B. Osherson¹⁹, I. L. Padilla^{1,32} , B. Racine¹⁷, A. S. Rahlin³⁴ , C. Reintsema²⁸, J. E. Ruhl¹⁰, M. C. Runyan¹¹, T. M. Ruud¹⁷ , E. C. Shaw¹⁹ , C. Shiu⁹, J. D. Soler³⁴ , X. Song⁹, A. Trangsrud^{11,12}, C. Tucker⁷, R. S. Tucker¹¹, A. D. Turner¹², J. F. van der List⁹, A. C. Weber¹², I. K. Wehus¹⁷ , S. Wen¹⁰, D. V. Wiebe⁸, and E. Y. Young^{29,30}

¹Department of Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON, M5S 3H4, Canada; leung@astro.utoronto.ca

²Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON, M5S 3H4, Canada

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

⁴Department of Physics, Washington University in St. Louis, 1 Brookings Drive, St. Louis, MO 63130, USA

⁵McDonnell Center for the Space Sciences, Washington University in St. Louis, 1 Brookings Drive, St. Louis, MO 63130, USA

⁶Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, ON, M5S 3H8, Canada

⁷School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, CF24 3AA, UK

⁸Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC, V6T 1Z1, Canada

⁹Department of Physics, Princeton University, Jadwin Hall, Princeton, NJ 08544, USA

¹⁰Physics Department, Case Western Reserve University, 10900 Euclid Avenue, Rockefeller Building, Cleveland, OH 44106, USA

¹¹Division of Physics, Mathematics and Astronomy, California Institute of Technology, MS 367-17, 1200 E. California Boulevard, Pasadena, CA 91125, USA

¹²Jet Propulsion Laboratory, Pasadena, CA 91109, USA

¹³School of Electrical, Computer, and Energy Engineering, Arizona State University, 650 E Tyler Mall, Tempe, AZ 85281, USA

¹⁴Department of Physics, McGill University, 3600 Rue University, Montreal, QC, H3A 2T8, Canada

¹⁵School of Mathematics, Statistics and Computer Science, University of KwaZulu-Natal, Durban, South Africa

¹⁶Blackett Laboratory, Imperial College London, SW7 2AZ, London, UK

¹⁷Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, NO-0315 Oslo, Norway

¹⁸Department of Physics, Shahid Beheshti University, 1983969411, Tehran, Iran

¹⁹Department of Physics, University of Illinois at Urbana-Champaign, 1110 W. Green Street, Urbana, IL 61801, USA

²⁰Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 W. Green Street, Urbana, IL 61801, USA

²¹Department of Physics, University of Texas, 2515 Speedway, C1600, Austin, TX 78712, USA

²²The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden

²³Kavli Institute for Cosmological Physics, University of Chicago, 5640 S Ellis Avenue, Chicago, IL 60637, USA

²⁴Steward Observatory, 933 North Cherry Avenue, Tucson, AZ 85721, USA

²⁵APC, Univ. Paris Diderot, CNRS/IN2P3, CEA/Irfu, Obs de Paris, Sorbonne Paris Cité, France

²⁶High Energy Physics Division, Argonne National Laboratory, Argonne, IL 60439, USA

²⁷Department of Astronomy and Astrophysics, Pennsylvania State University, 520 Davey Lab, University Park, PA 16802, USA

²⁸National Institute of Standards and Technology, 325 Broadway Mailcode 817.03, Boulder, CO 80305, USA

²⁹Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA

³⁰SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA

³¹Department of Mechanical and Aerospace Engineering, Princeton University, Engineering Quadrangle, Princeton, NJ 08544, USA

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

³³Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510-5011, USA

³⁴Max-Planck-Institute for Astronomy, Königstuhl 17, D-69117, Heidelberg, Germany

Received 2021 November 2; revised 2022 February 1; accepted 2022 February 11; published 2022 March 30

Abstract

Modern cosmic microwave background (CMB) analysis pipelines regularly employ complex time-domain filters, beam models, masking, and other techniques during the production of sky maps and their corresponding angular power spectra. However, these processes can generate couplings between multipoles from the same spectrum and from different spectra, in addition to the typical power attenuation. Within the context of pseudo- C_ℓ based, MASTER-style analyses, the net effect of the time-domain filtering is commonly approximated by a multiplicative transfer function, F_ℓ , that can fail to capture mode mixing and is dependent on the spectrum of the signal. To address these shortcomings, we have developed a simulation-based spectral correction approach that constructs a two-dimensional transfer matrix, $J_{\ell\ell'}$, which contains information about mode mixing in addition to mode attenuation. We demonstrate the application of this approach on data from the first flight of the SPIDER balloon-borne CMB experiment.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Unified Astronomy Thesaurus concepts: Cosmic microwave background radiation (322); Observational cosmology (1146); Astronomy data analysis (1858)

1. Introduction

Producing well-characterized maps of the cosmic microwave background (CMB) from time-ordered data requires accurately accounting for the impact of instrumental effects and any signal processing on the underlying astrophysical signal (e.g., Jarosik et al. 2007; Planck Collaboration et al. 2016). These processing techniques typically also have a nontrivial impact on the fidelity of the cosmological signal in ways that are spatially anisotropic and inhomogeneous. These effects must be precisely and accurately characterized in order to avoid biasing the estimation of the CMB angular power spectra, and therefore the inference of cosmological parameters.

The impact of timestream processing and instrument response is commonly approximated in the harmonic domain with a filter window function unique to the instrument, derived from the analysis of large ensembles of signal simulations. In its simplest formulation, the processing pipeline is modeled as a power attenuation mechanism in each multipole, i.e., the filter window function is a transfer function—a simple one-dimensional vector of ratios of output to input power.

This paper addresses the shortcomings of the one-dimensional model and proposes an alternative approach through the construction of a two-dimensional transfer matrix. This simulation-based spectral correction approach takes into account both the mode mixing and attenuation from instrumental and data processing effects including beams, filtering, and masking. Section 2 of this paper describes the theoretical motivation and compares the transfer matrix to the one-dimensional transfer function approach. A concrete example is presented in Section 3 using data from the first flight of SPIDER, a balloon-borne telescope designed to measure CMB polarization on roughly degree angular scales (SPIDER Collaboration 2021). This section explores different techniques for constructing the transfer matrix to reduce the computational demand including using binned power spectra and performing Fourier-space interpolation. Section 4 presents several comparisons of these techniques, including tests of signal recovery on spectra with different shapes. While all tested approaches were found to accurately recover a target spectrum identical to that used for the transfer matrix construction, the performance varied when applied to a different target spectrum. As discussed in Section 5, this has implications for increasingly sensitive CMB polarization measurements where the cosmological signals are heavily obscured by Galactic foregrounds. As the foreground power spectra are less well constrained and vary substantially between different sky regions, understanding the signal dependence of potential analysis techniques becomes even more important.

2. Theoretical Description

Maps of the CMB temperature (T) and linear polarization (Q , U) anisotropies over the partial sky can be decomposed into linear combinations of spherical harmonics:

$$T(\mathbf{r})W(\mathbf{r}) = \sum_{\ell,m} \tilde{a}_{\ell m}^T Y_{\ell m}(\mathbf{r}), \quad (1)$$

$$[Q(\mathbf{r}) \pm iU(\mathbf{r})]W(\mathbf{r}) = \sum_{\ell,m} \pm 2\tilde{a}_{\ell m} \pm 2Y_{\ell m}(\mathbf{r}), \quad (2)$$

where $W(\mathbf{r})$ is the window representing the relative weights of the partial-sky mask. To avoid using spin-weighted spherical harmonic components, the $\pm 2\tilde{a}_{\ell m}$ coefficients are frequently expressed as a combination of scalar and fixed parity E and B components:

$$\pm 2\tilde{a}_{\ell m} = -(\tilde{a}_{\ell m}^E \pm i\tilde{a}_{\ell m}^B), \quad (3)$$

where the sign convention follows Zaldarriaga & Seljak (1997). For each $\tilde{a}_{\ell m}$, the pseudo-power spectrum \tilde{C}_ℓ is defined as

$$\tilde{C}_\ell = \frac{1}{2\ell + 1} \sum_m |\tilde{a}_{\ell m}|^2. \quad (4)$$

Also known as a pseudo- C_ℓ (PCL), it is related to the angular power spectrum specified by the theory of primordial perturbations, C_ℓ , via

$$\langle \tilde{C}_\ell^X \rangle = \sum_{\ell':X'} K_{\ell\ell'}^{XX'} C_{\ell'}^{X'}, \quad (5)$$

where $K_{\ell\ell'}^{XX'}$ is a mode coupling kernel (or “mixing matrix”) that accounts for the mixing within and between the observables $X \in \{TT, EE, BB, TE, EB, TB\}$ due to the partial-sky mask, and the brackets $\langle \cdot \rangle$ denote an ensemble average over infinite spectrum realizations. Because $K_{\ell\ell'}^{XX'}$ is entirely determined by the chosen pixel weighting and the geometry of the cut sky (Hivon et al. 2002), in the absence of instrumental effects and noise, PCL estimators can use the (finite) set of measured \tilde{C}_ℓ to solve for the underlying power spectrum C_ℓ . Examples of PCL estimators include MASTER (Hivon et al. 2002), NaMaster (Alonso Monge et al. 2019), and PolSpice (Chon et al. 2004).

A major challenge in interpreting CMB data is that the instruments cannot directly probe the true sky signal; the incoming signals are inevitably altered by instrumental systematics and noise. Therefore, as described in Section 1, an experiment’s raw datastreams must be processed to remove many different types of spurious signals. The act of observing and time-domain filtering both distort the signal estimate and present additional sources of mode coupling. Assuming that the coupling is homogeneous, isotropic, and linear, the impact of the experiment is captured by introducing another coupling matrix, $F_{\ell\ell'}^{XX'}$, such that

$$\langle \tilde{C}_\ell^X \rangle = \sum_{\ell'':X''} \sum_{\ell':X'} K_{\ell\ell'}^{XX'} F_{\ell'\ell''}^{X'X''} C_{\ell''}^{X''}. \quad (6)$$

In general, $F_{\ell\ell'}^{XX'}$ is unique to the experiment and cannot be determined analytically.

2.1. The Filter Transfer Function

To reduce complexity, $F_{\ell\ell'}^{XX'}$ is often approximated as a diagonal matrix whose entries represent the one-dimensional

transfer function $F_\ell^{XX'}$, such that

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} K_{\ell\ell'} F_{\ell'} C_{\ell'} \quad (7)$$

(here and hereafter, the superscripts X are suppressed for brevity). Using a set of pseudo-power spectra \tilde{C}_ℓ derived from an ensemble of simulations of a known power spectrum C_ℓ , the transfer function F_ℓ can be estimated through an iterative process to avoid inverting $K_{\ell\ell'}$ (Hivon et al. 2002; Dutcher et al. 2021). Note that F_ℓ is frequently decomposed further into a filter component f_ℓ and a beam component b_ℓ^2 , such that $F_\ell = f_\ell b_\ell^2$; additional instrumental effects can be inserted in a similar fashion.

This one-dimensional approximation implicitly assumes that each ℓ -mode remains independent throughout the entire filtering process. As long as this assumption holds, the mapping from input to output is one-to-one: F_ℓ is simply the ratio of output to input power for each ℓ . However, in practice, we expect modes to become coupled with one another, where the mapping becomes many-to-one and the contributions from the coupled input modes become impossible to disentangle using F_ℓ alone. The consequence is that changing the input power spectrum C_ℓ also changes the output \tilde{C}_ℓ in some nontrivial way due to this many-to-one mapping. In other words, F_ℓ is inextricably tied to the input used to compute it.

More concretely, the one-dimensional transfer function formulation conflates mode mixing with the in-mode filter gain. This introduces a sensitivity to the power spectrum of the simulated sky used to calibrate F_ℓ ; the final spectrum is correct only if the simulated sky is statistically similar to the true sky (i.e., a Gaussian sky realization with the same power spectra).

2.2. The Multipole–Multipole Transfer Matrix

To address this shortcoming of the F_ℓ formulation, we introduce a two-dimensional linear operator that encodes the (asymmetric) coupling between each $\ell\ell'$ pair:

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} J_{\ell\ell'} C_{\ell'}. \quad (8)$$

We refer to this coupling operator $J_{\ell\ell'}$ as the “transfer matrix” because it directly relates the input of the true power spectrum to the output of the spectrum estimator. This relation holds as long as the filtering process is approximately linear. Treating the entire pipeline as a single operator avoids the diagonal approximation and ensures that all multipole–multipole couplings induced by time-domain and map-domain filtering are properly taken into account, i.e., we are not locked into a specific input spectrum.

Note that these couplings extend to those between the six power spectra; both temperature-to-polarization (T -to- P) and E -to- B leakage are automatically included. Because standard PCL spectrum estimators provide the TT , EE , BB , TE , EB , and TB power spectra, any input mode can be readily related to an output mode from any of these six power spectra, resulting in 6^2 coupling matrices. We find it convenient to compile these individual matrices into a single 6×6 block matrix encapsulating every $\ell\ell'$ coupling between the six power spectra. Following the rules of matrix multiplication (Equation (8)), we arrange these blocks horizontally according to input spectra and vertically according to output spectra (see Figure 1). The ordering of the six spectra does not matter as long as it is

consistent between the two axes; likewise, the six C_ℓ s must be concatenated in the same order. We choose the above ordering based on convenience.

While the approach presented here has similarities to that used by the BICEP/Keck Collaboration for their CMB polarization analyses, the implementation varies due to key differences in the observing strategies and analysis pipelines. As described in Ade et al. (2016), the simplicity of the BICEP/Keck observing strategy allows for the construction of an observing matrix that renders large simulation ensembles more computationally tractable. Having determined their filtering operations to be linear, the observing matrix is used to construct the transfer matrix $J_{\ell\ell'}$, which is then used similarly to reconstruct and interpret the on-sky power spectra. Because the observing matrix formulation is less generally applicable to experiments at other sites with different observing strategies, the remainder of this paper is dedicated to estimating $J_{\ell\ell'}$ through other means.

3. Application to SPIDER Data

To illustrate the application of the transfer matrix, we present results from the SPIDER balloon-borne telescope. During its first Antarctic long-duration balloon (LDB) flight in January 2015, SPIDER mapped 4.8% of the sky with polarimeters operating at 95 and 150 GHz to constrain the B -mode power spectrum from primordial gravitational waves (SPIDER Collaboration 2021). An upcoming flight with additional 280 GHz receivers will provide improved characterization of the polarized Galactic dust foregrounds (Shaw et al. 2020).

SPIDER’s processing pipeline is described more fully in SPIDER Collaboration (2021) and SPIDER Collaboration (2022, in preparation); here we briefly highlight the most relevant steps and note that they are sufficiently linear to allow usage of Equation (8). Once features such as cosmic-ray hits, payload transmitter signals, and thermal transients have been removed from the raw detector timestreams, they are filtered to reduce quasi-stationary noise. Null test performance was used to identify the weakest filter that sufficiently removed quasi-stationary noise, resulting in a fifth-order polynomial fit to each detector’s data as a function of azimuth angle between scan turnarounds. The impact of scanning, filtering, and flagging are determined by applying the entire processing pipeline to an ensemble of time-domain signal simulations in a procedure known as reobservation. The reobserved timestreams are produced at the full data sample rate without any down-sampling or binning. Unlike the BICEP/Keck experiment, SPIDER’s observing strategy does not allow for the creation of an observing matrix, so obtaining the transfer matrix $J_{\ell\ell'}$ requires producing an appropriate set of reobserved CMB maps.

SPIDER’s measured and simulated timestreams are converted into two-dimensional maps of the sky with a binned mapmaker. This approach assembles the detector data into spatial pixels based on the telescope pointing and polarization sensitivity as described further in SPIDER Collaboration (2021). The computational simplicity of this method is crucial for enabling the generation of large simulation ensembles, including those used in this work.

The main SPIDER cosmological results presented in SPIDER Collaboration (2021) use two complementary pipelines for power spectrum estimation. Here we describe results only from the Noise Simulation Independent (NSI) pipeline, while the

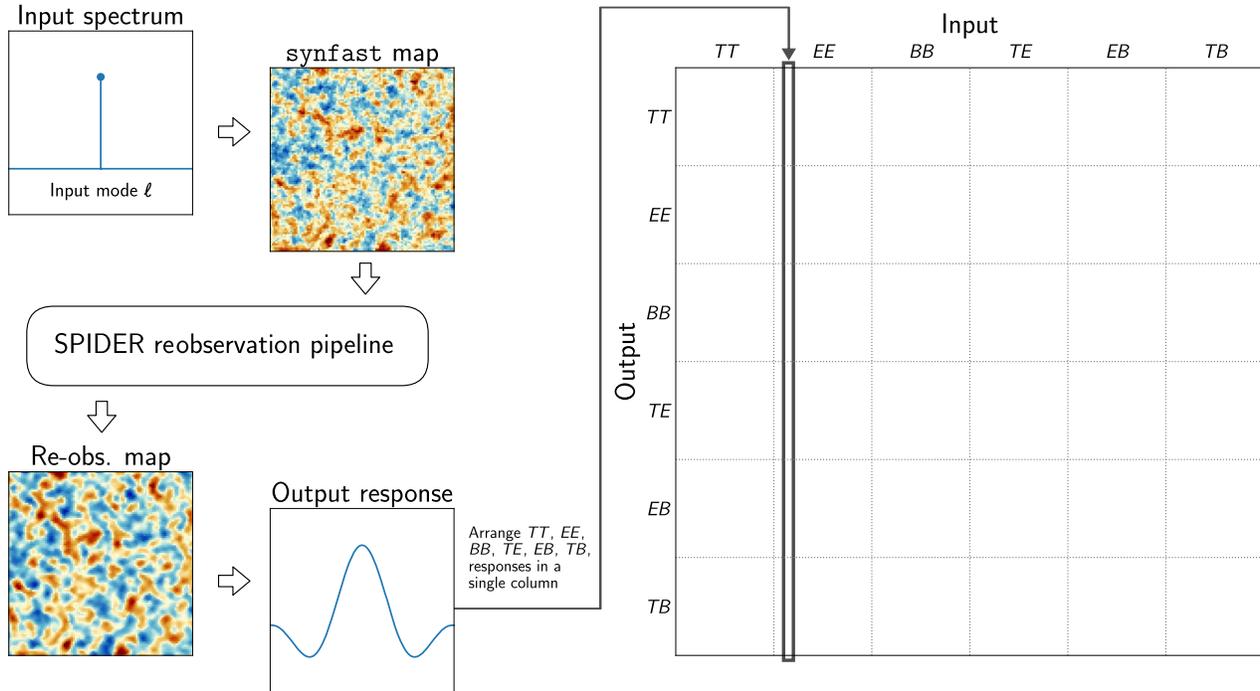


Figure 1. Procedure to create the transfer matrix $J_{\ell\ell'}$. For each input mode, the six output response spectra (gains) are arranged vertically to form the columns of the matrix. Each of the 36 blocks within the matrix spans a range of multipoles with $\ell_{\min} \leq \ell \leq \ell_{\max}$.

other pipeline is presented in Gambrel et al. (2021). The NSI pipeline is similar to Xspect (Tristram et al. 2005) and Xpol (Tristram 2006), as well as to those used by the South Pole Telescope (Lueker et al. 2010), CLASS (Padilla et al. 2020), and SPIDER’s circular polarization analysis (Nagy et al. 2017). It decomposes the 2015 flight data into 14 maps by splitting the timestreams into interleaved 3 minute segments for each of the six receivers. This timescale was chosen to maximize the number of complete and well-conditioned maps of the sky region, taking into account SPIDER’s scan rate and observing strategy. These maps are then co-added by frequency to produce 14 maps in each of the two SPIDER bands, 95 and 150 GHz. The cross-power spectra for all pairs of maps (neglecting the auto-spectra) are produced with PolSpice (Chon et al. 2004), and the distribution of the cross-spectra allows for an empirical measurement of the uncertainty without reliance on an instrumental noise model. In SPIDER’s NSI pipeline, these cross-spectra are used to estimate cosmological parameters by comparing them to theoretical models that depend on the parameters of interest. It is therefore SPIDER’s power spectra, rather than the maps, that we correct for the beam and filtering effects as described in the rest of this work. Since the PolSpice estimator is linear, the assumption of linearity in the definition of the transfer matrix (Equation (8)) is satisfied. Although the transfer matrix can be used to correct the power spectra for temperature-to-polarization leakage, we perform this correction on the maps based on simulations of Planck temperature-only maps (Planck Collaboration et al. 2020). This avoids the sample variance from the simulation ensemble, which is relatively large due to the amplitude of the temperature signals, by using existing measurements of the temperature anisotropies in SPIDER’s observing region.

As $J_{\ell\ell'}$ is merely the linear operator linking the input and output power spectra (Equation (8)), its computation is conceptually straightforward: pass a (simulated) CMB map with a known spectrum through the reobservation pipeline, then compute the power spectra of the output map. But to capture the asymmetric mode–mode coupling between each $\ell\ell'$ pair, this process must be performed on each mode individually. Thus we begin with a set of unit δ -functions—one at each ℓ of interest—and simulate a set of CMB maps using the synfast HEALPix utility³⁵ (Gorski et al. 2005). Upon obtaining the reobserved maps, we organize their power spectra into columns, arranged by ℓ , to form $J_{\ell\ell'}$. The whole process is illustrated schematically in Figure 1.

To reduce the effects of sample variance from the CMB map realization and partial-sky spectrum estimation procedures, we repeat the above process 100 times and average the 100 resulting matrices to obtain the final $J_{\ell\ell'}$. However, the computation cost is intensive: with 100 map realizations of a δ -function in each of the three CMB modes (TT , EE , BB) that must be reobserved by each of SPIDER’s six receivers, we require 1800 simulated maps per $\ell \in [2, 250 + \Delta\ell_b]$ (a $\Delta\ell_b$ “buffer” is necessary to account for couplings of higher multipoles with lower multipoles). To minimize the computational burden, we use full-mission SPIDER maps instead of the 14 NSI pipeline submaps because this was demonstrated to have a negligible effect on SPIDER’s signal recovery in simulations. Nevertheless, reobserving the entire set of required maps would take about 1000 core-years on the Niagara supercomputing cluster (Loken et al. 2010; Ponce et al. 2019), which is computationally infeasible given SPIDER Collaboration resources. Since the fidelity of the computation needs to be

³⁵ <https://healpix.sourceforge.io>

balanced against the time required, two practical options are investigated in the following sections: (1) use input spectra with multipole bins rather than evaluating each multipole individually, or (2) compute the transfer matrix for only a sample of input ℓ s and interpolate between them.

3.1. Bin–Bin Transfer Matrix

CMB angular power spectra are typically binned into multipole ranges of width $\Delta\ell$ in order to increase the signal-to-noise ratio and mitigate the correlative effects of adjacent multipole leakage. As a result, we consider the case where the input power spectra are also binned into the same multipole ranges, thereby reducing the computational demands of the calculation. By analogy with Equation (8), we define a bin–bin transfer matrix $J_{bb'}$ that describes the average power in output bin b given the average power in each of the input bins b' :

$$\langle \tilde{C}_b \rangle = \sum_{b'} J_{bb'} \sum_{\ell} P_{b'\ell} C_{\ell} = \sum_{b'} J_{bb'} C_{b'}, \quad (9)$$

where $P_{b\ell}$ is a binning operator.

At first glance, the δ -function inputs to $J_{\ell\ell'}$ might be replaced by unit-boxcars for $J_{bb'}$, but, as it turns out, that may not be the best choice. Unlike the case for the full $J_{\ell\ell'}$, the binning introduces a dependence on the shape of the input spectrum used to estimate the transfer matrix (see the Appendix). Consequently, the choice of input spectrum should be as close as possible to the anticipated output spectrum, as is further discussed in Section 4. Because our application is toward degree-scale CMB B -modes, we use a Λ cold dark matter (CDM) + ($r = 0.03$) spectrum, provided by CAMB (Lewis et al. 2000), as our base model. Windowing this model spectrum at each bin of interest—i.e., multiplying it by unit-boxcars centered on each bin—provides the set of input spectra needed to construct the bin–bin transfer matrix $J_{bb'}$.

Because the analysis of SPIDER data uses only 12 bins in the range $2 \leq \ell \leq 300$, the computational demand is decreased by a factor of ~ 30 ; months of cluster time is reduced to days. The price is a dependence on an input model, but unlike F_{ℓ} , the ability of $J_{bb'}$ to capture couplings between multipole bins provides a more general framework.

3.2. Fourier-space Interpolation

An alternative method for reducing computation time is to compute the transfer matrix $J_{\ell\ell'}$ only at select multipoles (i.e., columns) and interpolate in between. This requires the matrix to be smooth in the ℓ -range of interest, but, as we shall see, this assumption is often well justified in practice.

Because $J_{\ell\ell'}$ is a linear operator (Equation (8)), the output from a δ -function input can be interpreted as the impulse response of a linear filter. As shown in the left panel of Figure 2, SPIDER’s transfer components ($TT \rightarrow TT$, $EE \rightarrow EE$, etc.) have impulse responses resembling displaced sinc functions. This allows us to recast the problem from the ℓ domain to its dual frequency-like domain (hereafter simply the “frequency” domain), where the interpolation turns out to be considerably easier. For our application, we use the sign and normalization conventions established by NumPy’s `fft` module;³⁶ for a given input ℓ' , the discrete Fourier transform of its impulse response x_{ℓ} —the frequency response—is defined

by

$$X_f = \sum_{\ell=0}^{N-1} x_{\ell} \exp\left(-i\frac{2\pi f}{N}\ell\right). \quad (10)$$

We use $N = 801$ points ($0 \leq \ell \leq 800$) in our computations. The magnitudes $|X_f|$ and phases $\phi = \angle X_f$ of the resulting transforms are shown in the right panel of Figure 2.

In general, the properties of the responses are unique to each experiment and therefore often determined empirically. SPIDER’s response corresponds to a linear-phase low-pass filter: the magnitude rolls off until it hits the noise floor of the filter at $f = f_c$, while the phase ϕ is nearly linear up to that same frequency. The slope of the linear-phase component is proportional to the shift (in ℓ) of the impulse, namely, $m = \phi/f = -2\pi\ell$ when f is expressed in normalized frequency units (cycles per multipole). The magnitude of the noise floor encodes the level of background noise at the input multipole ℓ' .

This simple structure in the frequency domain permits a straightforward interpolation. The magnitudes, in particular, can be interpolated directly at all f , as can the unwrapped phases at $f < f_c$. For simplicity, we interpolate piecewiselinearly. In the noise region $f > f_c$, the *wrapped* phases exhibit a curious behavior: they settle onto one of two possible values separated by π . Thus, rather than interpolating, we implement a nearest-neighbor scheme for the phases at $f > f_c$.

One might be tempted to filter out the noise at $f > f_c$ completely. This requires care: any operation in the frequency domain carries consequences into the ℓ domain. For example, a simple low-pass filter with a rigid cutoff is equivalent to convolution with a sinc function in ℓ -space; this operation will pollute the impulse response with heavy ringing, and should be avoided. Because good filter design merits a study of its own, we do not carry out any filtering of this sort.

To complete the process, the interpolated values are inverse Fourier transformed back into ℓ -space to form our interpolated $J_{\ell\ell'}^{\text{interp}}$ (Figure 3). The choice of nodes for the interpolation, i.e., which columns of $J_{\ell\ell'}$ to compute rather than interpolate, reduces to a balance between fidelity and computation time. Naturally we will want to concentrate our efforts to regions where $J_{\ell\ell'}$ changes rapidly in order to obtain a good interpolation. For us, that is the low- ℓ regime. We thus compute the columns at $\ell = \{5, 10, 15\}$ in addition to SPIDER’s bandcenters.

We note that this particular choice of interpolation method is motivated by SPIDER’s sinc-function-like impulse responses. However, any sufficiently smooth $J_{\ell\ell'}$ can be interpolated with similar techniques.

3.3. Choice of Spectrum Estimator

The matrix-making process encodes the spectrum estimator into the matrix—that is, $J_{\ell\ell'}$ depends on the choice of estimator by design (Equation (8)). Nonetheless, this choice can be trivially changed because the spectrum estimation and reobservation steps are completely independent of each other.

The `anafast` utility (provided in the `healpy` package) is a simple spectrum estimator that directly computes the power spectra from a set of input maps using Equation (4). In general, the interpolated matrix constructed using `anafast` provides very good agreement with the expected results. However, we find that the “buffer” needs to be quite large ($\Delta\ell_b > 120$) in

³⁶ <https://numpy.org/doc/stable/reference/routines.fft.html>

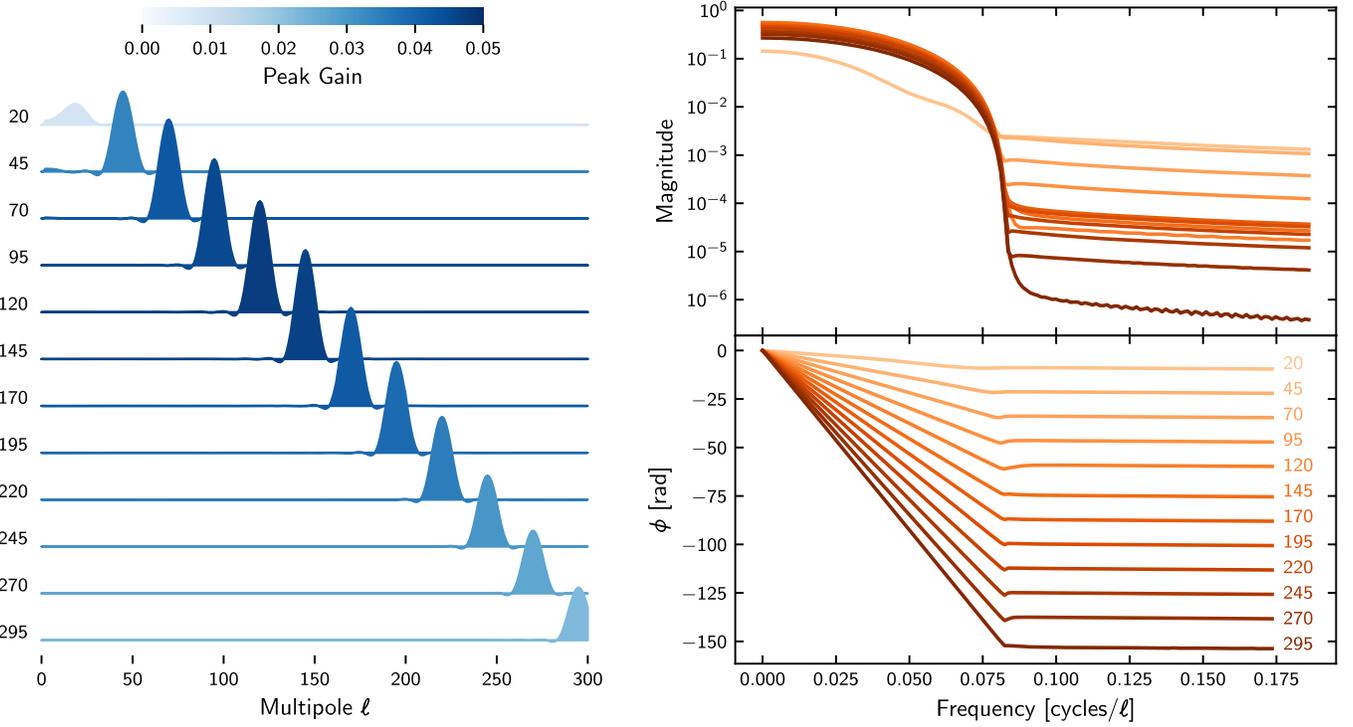


Figure 2. Left: impulse responses of unit δ -function inputs at $\ell = 20, 45, \dots, 295$ for the 150 GHz $EE \rightarrow EE$ coupling; the TT and BB counterparts are similar. Right: the corresponding magnitude and (unwrapped) phase responses. The “frequency” axis has units cycles/ ℓ and does not have a simple physical interpretation. These responses are interpolated in the Fourier domain to approximate the transfer matrix (Figure 3).

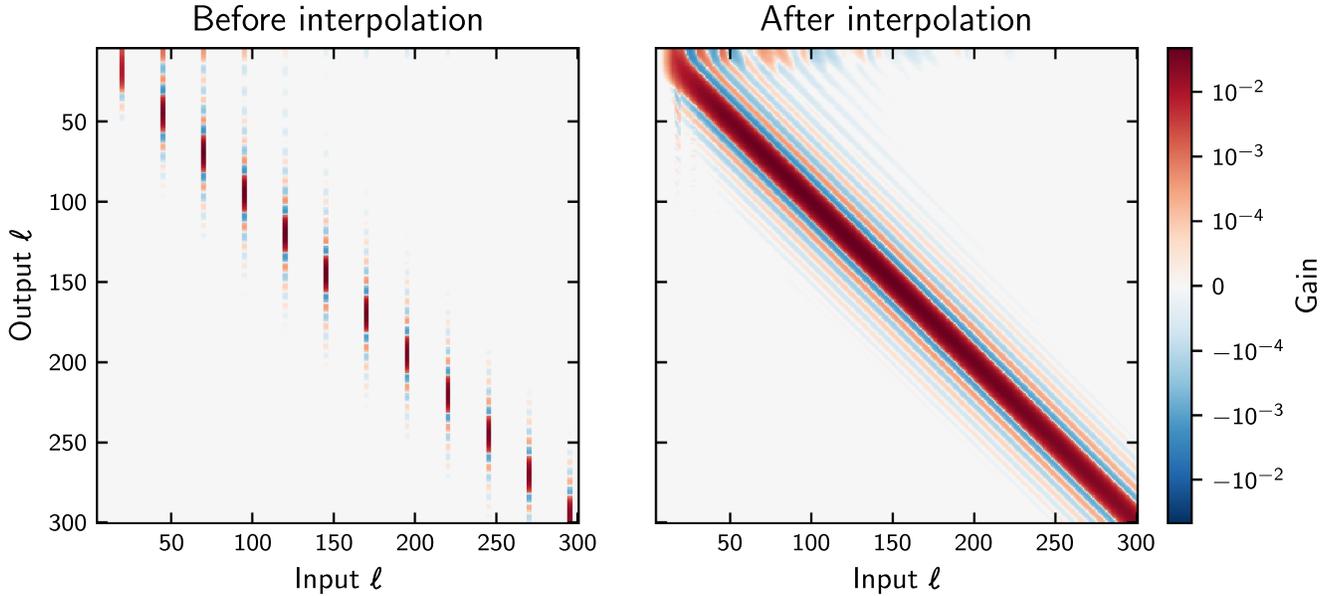


Figure 3. The impulse responses from the left panel of Figure 2, arranged into columns (left), are interpolated in Fourier space to form an approximated $J_{\ell\ell}^{\text{interp}}$ (right). Shown is the 150 GHz $EE \rightarrow EE$ coupling; the TT and BB counterparts are similar.

order to properly account for E -to- B leakage (Figure 4). To avoid the additional computational demand, in the following work, we proceed using the `PolSpice` estimator, which includes as part of its algorithm a correction for E -to- B leakage built-in. This reduces the “buffer” to a much more manageable $\Delta\ell_b = 50$.

We use the same `PolSpice` parameters as the SPIDERB-mode analysis for this work, including a cosine apodization

with $\sigma = 10^\circ$, $\theta_{\text{max}} = 50^\circ$, and `symmetric_cl` (SPIDER Collaboration 2021).

3.4. Cross-spectra: TE , EB , and TB

An outstanding problem in our discussion so far is the treatment of the cross-spectra TE , EB , and TB . While the interspectral couplings from auto- to cross-spectra ($TT \rightarrow TE$, etc.)

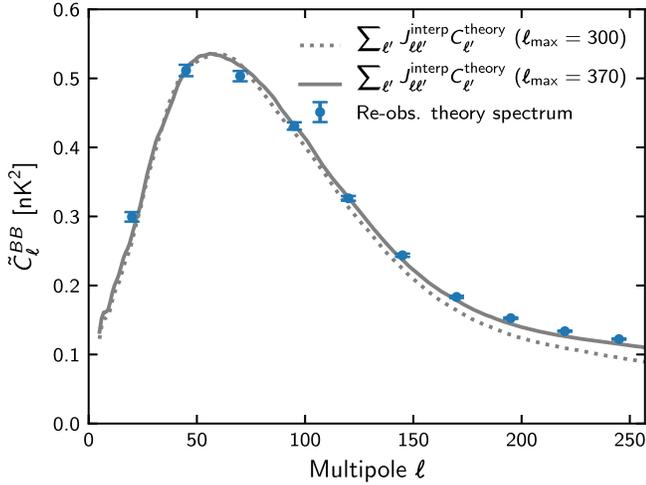


Figure 4. Although *anafast* does not automatically correct for *E*-to-*B* leakage, a transfer matrix built using this utility can still capture such mode mixing induced by SPIDER’s filter and mask. When the matrix is applied to a Λ CDM + ($r = 0.03$) theory spectrum, the result can predict the output from reobserving a map with the same theory spectrum (blue points; error bars indicate the error on the mean of the 500 simulations). However, replicating the full effects of the reobservation procedure in the higher bins requires computing a larger matrix than anticipated due to the wide tails of the $EE \rightarrow BB$ response. Namely, a matrix with $\ell_{\max} = 300$ (dotted line) underestimates the power in the higher bins; this is alleviated somewhat by increasing to $\ell_{\max} = 370$ (solid line), but that is still not quite enough. As a result, we proceed with a matrix constructed using the *PolSpice* estimator, as its built-in *E*-to-*B* leakage correction algorithm avoids the need to generate larger matrices, thereby reducing the computational demand.

are implicitly accounted for by the spectrum estimator (which returns all six power spectra), the method described here cannot be used to compute the matrix components with a cross-spectrum as input. To wit, it is impossible to generate maps of Gaussian-distributed $a_{\ell m}$ s with a nonzero correlation (the input mode) in *TE*, *EB*, or *TB* while the *TT*, *EE*, and *BB* spectra are imposed to be zero. Consequently, we approximate the diagonal cross-spectra transfer components (*TE* \rightarrow *TE*, etc.) as the geometric mean of the related auto-spectra components, $J_{\ell\ell'}^{TE,TE} = (|J_{\ell\ell'}^{TT,TT}| |J_{\ell\ell'}^{EE,EE}|)^{1/2}$, etc., and treat the off-diagonal components (*TE* \rightarrow *EB*, etc.) as negligible. Although this is not ideal, it was found not to bias simulation results for our application. An alternative method could be to construct the transfer matrix using the harmonic modes, rather than the power spectra, of the maps. As an example, see Choi et al. (2020), who postulate a transfer matrix built from (two-dimensional) Fourier transforms of *T*, *E*, and *B* maps.

4. Comparison of Methods

We quantify the effectiveness of the two methods described above (binning and interpolating) by two metrics: their ability to replicate the end-to-end SPIDER pipeline, and to recover the original signal. In the context of Equation (8), this corresponds to producing the correct \tilde{C}_ℓ when given C_ℓ , and vice versa.

4.1. Case 1: Solve for \tilde{C}_ℓ , Given C_ℓ

We demonstrate the first of these by applying our computed transfer function/matrices directly to a theory spectrum. This result is then compared to the output spectrum generated by passing a set of map realizations of the same input theory spectrum through the SPIDER reobservation procedure. To

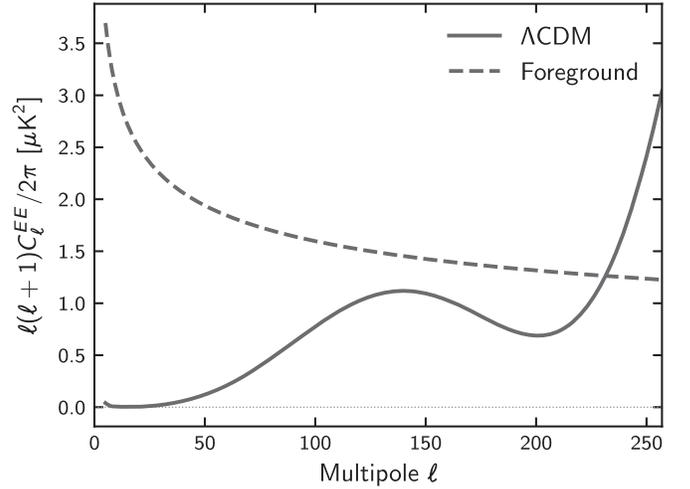


Figure 5. Two fiducial input theory spectra with very different shapes are used to evaluate the performance of different methods of computing the transfer function: a theoretical Λ CDM CMB spectrum and a power-law Galactic foreground model. Here only the CMB *EE* spectrum is shown, while the same foreground model is used for temperature and polarization.

ensure broad applicability, we use two qualitatively different theory spectra as input: a Λ CDM + ($r = 0.03$) spectrum (from CAMB; 500 realizations) and a power-law dust foreground spectrum (300 realizations). These two spectra are chosen for their contrasting behavior (Figure 5).

As discussed in Section 2.1, the input-dependence of F_ℓ makes it suitable only when the target spectrum does not deviate far from the input spectrum that was used to construct F_ℓ . This is demonstrated in Figure 6. Having derived our F_ℓ using a fiducial Λ CDM + ($r = 0.03$) spectrum as input, applying it to the same spectrum again outputs a trivial result (left panels). But when the target is the foreground spectrum (right panels), large deviations occur wherever mode–mode coupling dominates over in-mode gain—i.e., wherever the transfer matrix is least diagonal-like. In our case, this occurs at low ℓ s (Figure 3). The simple power attenuation model of F_ℓ is insufficient in capturing SPIDER’s filtering scheme in this regime.

The same kind of input-dependence is seen in the binned transfer matrix $J_{bb'}$ (see the Appendix). Like F_ℓ , the $J_{bb'}$ constructed from a fiducial Λ CDM + ($r = 0.03$) spectrum is highly effective when the target spectrum is also a Λ CDM + ($r = 0.03$) spectrum, and diverges when the target is the foreground dust spectrum (Figure 6). But this time, all foreground bins up to $\ell = 100$ are noticeably afflicted. This is due to the loss of granularity from binning the spectrum and the matrix, as we had used a bin width of $\Delta\ell = 25$ to match SPIDER’s *B*-mode analysis. In the limit where the bin width is 1—a single multipole—we recover the full, theoretically ideal transfer matrix $J_{\ell\ell'}$.

The interpolated transfer matrix $J_{\ell\ell'}^{\text{interp}}$ thus functions as a compromise allowing us to have a bin width of 1 without being computationally intractable. Despite the simple piecewise-linear interpolation approximation used, it reliably recovers both target spectra.

4.2. Case 2: Solve for C_ℓ , Given \tilde{C}_ℓ

To recover C_ℓ from \tilde{C}_ℓ , the linear system defined by Equation (8) must be solved numerically; directly inverting

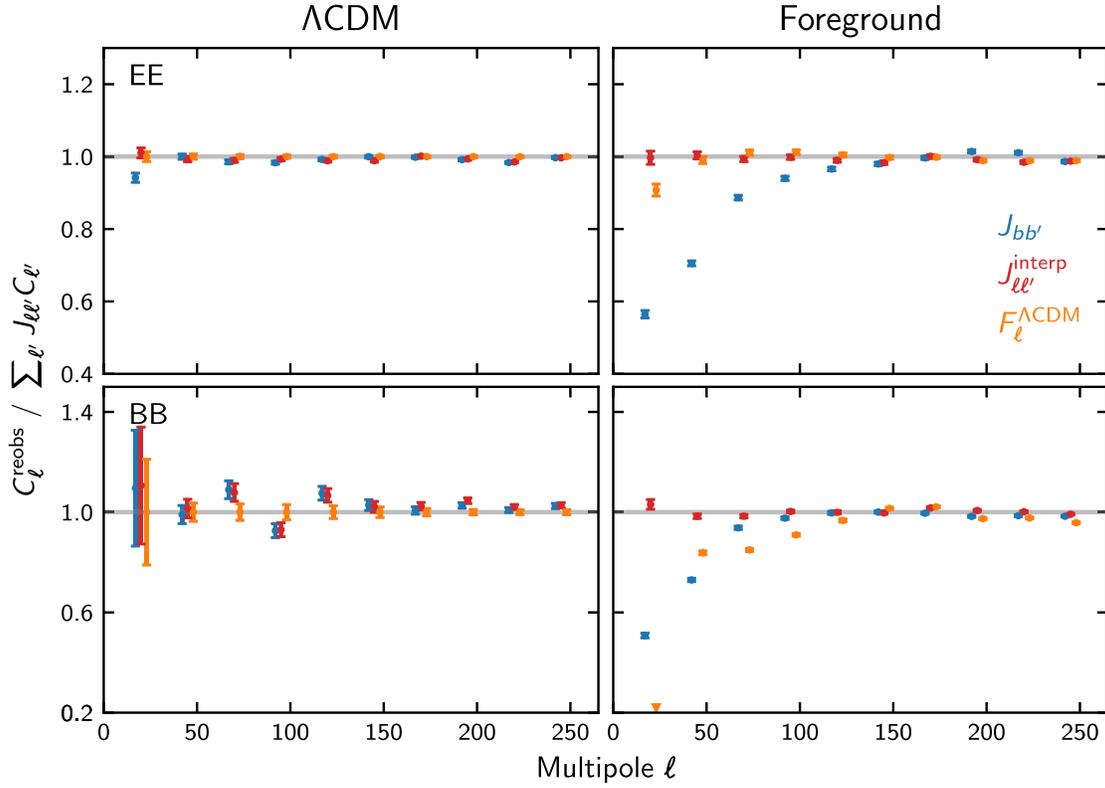


Figure 6. Average ratios of reobserved 150 GHz spectra to those obtained from applying the transfer function and transfer matrices to fiducial Λ CDM + ($r = 0.03$) and foreground input spectra (500 and 300 realizations, respectively). A ratio of 1 represents the ideal performance, signifying a perfect replication of SPIDER’s filtering process. As the transfer function F_ℓ was derived using the reobserved Λ CDM result, applying it to that same spectrum again returns 1 by definition (left panels). Error bars indicate the error on the mean of the simulation ensembles.

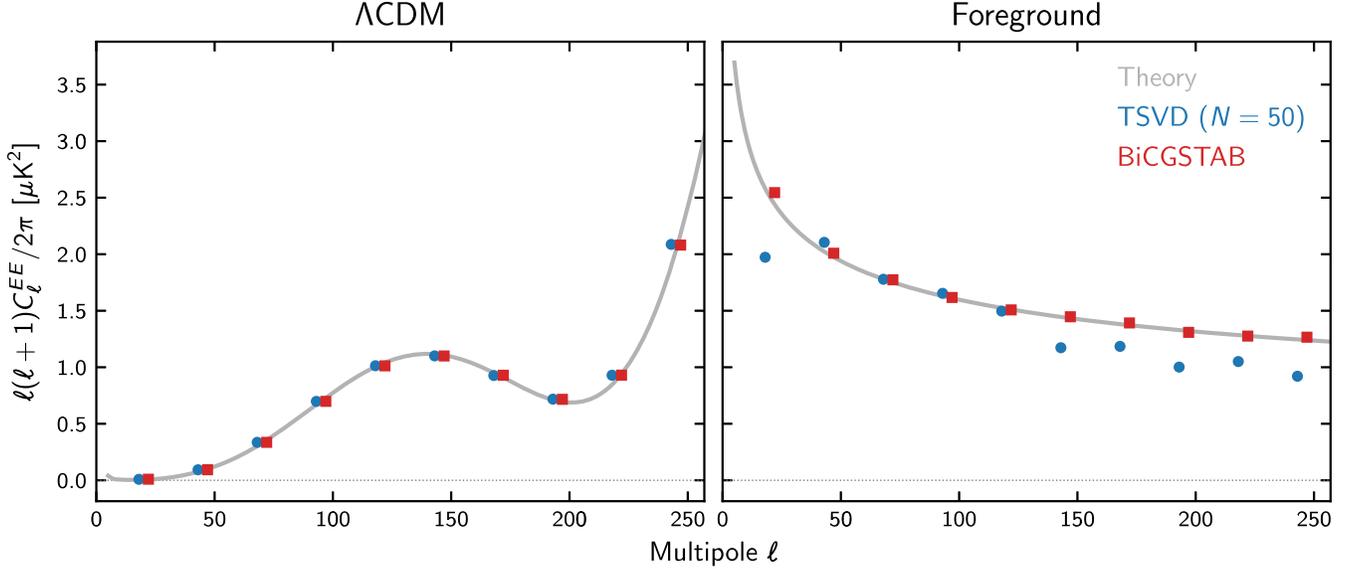


Figure 7. Solutions to the linear system $\sum_{l'} J_{ll'}^{\text{interp}} x_{l'} = \sum_{l'} J_{ll'}^{\text{interp}} C_l^{\text{theory}}$ for a Λ CDM EE spectrum (left) and a dust foreground spectrum (right) computed by the truncated singular value decomposition (TSVD) method with 50 singular values and the biconjugate gradient stabilized (BiCGSTAB) method. Near-optimal results can be obtained when using the TSVD matrix as a preconditioner and providing an initial guess close to C_l^{theory} (details provided in the text). Though not shown here, TT and BB spectra return similar results.

a linear system is not recommended except in special cases. One such case is the transfer function F_ℓ —computing its inverse is trivial since it is a diagonal matrix (a one-dimensional function). Likewise, SPIDER’s binned transfer matrix $J_{bb'}$ is also easily invertible because it is diagonally dominant and measures only 72×72 (12 bins in each of the

six correlation spectra); its condition number is very close to 1. The main challenge, therefore, is solving the system for the interpolated transfer matrix $J_{ll'}^{\text{interp}}$, whose condition number is $\mathcal{O}(10^{17})$.

The large condition number of $J_{ll'}^{\text{interp}}$ —and also of $J_{ll'}$ in general—means that it is highly susceptible to numerical noise

and cannot be solved using elementary linear algebraic techniques. Numerous algorithms of varying complexities exist to solve such ill-conditioned linear systems; the study of these techniques is beyond the scope of this paper. As a demonstration, below we focus on the truncated singular value decomposition (TSVD) method and the biconjugate gradient stabilized (BiCGSTAB) method (van der Vorst 1992). To test the efficacy of these two techniques, we compute $\tilde{C}_\ell = \sum_{\ell'} J_{\ell\ell'}^{\text{interp}} C_\ell^{\text{theory}}$, solve the system $\sum_{\ell'} J_{\ell\ell'}^{\text{interp}} x_{\ell'} = \tilde{C}_\ell$, and compare the solution x_ℓ with the original C_ℓ^{theory} . The results are shown in Figure 7. We use the same test spectra as the previous section (Figure 5) as C_ℓ^{theory} .

TSVD approximates the inverse matrix by rank-reducing the factorized matrices such that only the N largest singular values remain. Because singular values represent magnitudes along orthogonal axes, this truncation removes the low-amplitude components most susceptible to numerical instabilities. For our $J_{\ell\ell'}^{\text{interp}}$, we truncate the singular values at the quasi-arbitrary cutoff of 10^{-3} , which leaves the top 50 singular values; below this cutoff, the singular vectors are inundated with numerical noise. As shown in Figure 7, TSVD with $N = 50$ does well for a Λ CDM spectrum but remains noisy for a dust power law.

On the other hand, BiCGSTAB is an iterative algorithm. While we find that its standard configuration can recover the C_ℓ^{theory} to an acceptable tolerance, its convergence can be aided by a suitable preconditioner matrix and initial guess. For convenience, we use the truncated matrix from our TSVD tests above as the preconditioner, and, to replicate real conditions, we use a noisy version of C_ℓ^{theory} ($\pm 5\%$ noise) as the initial guess. This produces near-optimal results (Figure 7) in fewer than 10 iterations.

5. Conclusion and Discussion

The transfer function F_ℓ is an ansatz introduced by Hivon et al. (2002) to represent the effects of datastream filtering in the process of CMB mapmaking. It was noted at the time that F_ℓ cannot fully capture such filtering, on account of inconsistencies with the statistics governing the power spectrum C_ℓ (in particular, time-domain filtering violates the assumption of isotropy), but the authors showed that such inconsistencies had a negligible impact on their results. With filtering and analysis techniques having grown far more complex as CMB experiments have become more sensitive in the two decades since, and considering the crucial role F_ℓ fills in recovering cosmological signals, this ansatz merits a more thorough reexamination. This is particularly important for measurements of CMB polarization in the presence of large foreground components, where there is significant uncertainty in the shape of the power spectrum for a given sky region, and where foreground residuals may be a concern in component-separated maps.

In this study, we present a different approach, the transfer matrix $J_{\ell\ell'}$, that accounts for a key element missing in F_ℓ : the multipole–multipole coupling induced by the instrument response and timestream filtering. These couplings are not restricted to those within the same spectrum; also included are inter-spectral coupling components such as T -to- P and E -to- B leakage. Because the full matrix is computationally infeasible for most modern experiments, we additionally explore two approximations: a binned transfer matrix $J_{bb'}$ and an interpolated transfer matrix $J_{\ell\ell'}^{\text{interp}}$. Along with F_ℓ , we compare their

performance by applying them to the sky-to-spectrum pipeline from SPIDER’s first flight.

When the target is a Λ CDM-like spectrum, we find that all three approaches are effective at replicating the sky-to-spectrum pipeline. However, when the target is switched to a power-law dust foreground model, F_ℓ and $J_{bb'}$ —both being dependent on the fiducial model used to construct them—fail at large scales, and only $J_{\ell\ell'}^{\text{interp}}$ is able to reliably recover the underlying signal. However, like the full transfer matrix $J_{\ell\ell'}$ itself, $J_{\ell\ell'}^{\text{interp}}$ is ill-conditioned and care needs to be taken in solving the linear system. Λ CDM-like TT , EE , and BB spectra are easily recovered using a TSVD approach, but a dust foreground-like spectrum requires more sophisticated algorithms like the BiCGSTAB method to remain numerically stable.

Nonetheless, because SPIDER employs a map-based method of foreground removal, we expect the remaining signal to deviate, at most, only slightly from Λ CDM. As long as this is the case, the error from binning remains small (see the Appendix). We therefore chose to use the binned transfer matrix $J_{bb'}$ in the NSI pipeline for SPIDER’s main result (SPIDER Collaboration 2021), as it reliably recovers a Λ CDM-like spectrum (Figure 6) and is easily invertible. However, we caution that $J_{bb'}$ should only be pursued when the binning scheme used in analysis has been established, as its input bins must match the output bins; any changes to the binning scheme must be accompanied by a recomputation of $J_{bb'}$. On the other hand, $J_{\ell\ell'}^{\text{interp}}$ is only as accurate as the interpolation allows, but is attractive for the capacity to increase its resolution as desired. Additionally, $J_{\ell\ell'}^{\text{interp}}$ may be most advantageous for large angular scale foreground analyses, where it performs significantly better than $J_{bb'}$.

Of course, there is no need to adhere strictly to one technique. If computational resources are limited, one may employ a hybrid approach and compute the transfer matrix only in regimes dominated by mode–mode coupling, i.e., where the transfer matrix has significant off-diagonal components. For example, one may choose to commit to fully computing every column of the lowest bin, where the discrepancy between the input and output spectra is greatest (see Figure 6), and use a less expensive strategy elsewhere. We also leave open the exploration of different interpolation schemes, as the implementation described in this paper has not been found to significantly affect SPIDER’s B -mode results. Alternatively, the simple structure of the transfer matrix in Fourier space also suggests, in place of interpolation, a model with a polynomial fit in magnitude and a linear fit in phase. We note that the biggest bottleneck is often computational time for the production of reobserved maps; as long as this portion of the pipeline is fixed, downstream tasks such as masking and choice of spectrum estimator are trivial to change.

As CMB experiments become capable of more precise measurements, improving the accuracy of the signal recovery process becomes increasingly important. The transfer matrix decouples the estimation of power spectra—and, by extension, the cosmological parameters—from assumptions about the shape of the underlying sky signal. This is particularly useful when the cosmological signals are heavily obscured by foreground Galactic dust, and is thus relevant for upcoming CMB instruments searching for primordial B -modes (e.g., Abazajian et al. 2019; Ade et al. 2019; Hazumi et al. 2020; Shaw et al. 2020) or measuring the E -modes on the largest

angular scales (Allison et al. 2015; Watts et al. 2018; Hazumi et al. 2020). As foreground component separation becomes an increasingly important part of CMB analyses, future experiments can benefit from the development of new signal-independent analysis techniques to improve the power spectrum reconstruction.

SPIDER is supported in the U.S. by the National Aeronautics and Space Administration under grant Nos. NNX07AL64G, NNX12AE95G, NNX17AC55G, and 80NSSC21K1986 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. K.F. holds the Jeff & 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. C.R.C. was supported by UKRI Consolidated grant Nos. ST/P000762/1, ST/N000838/1, and ST/T000791/1.

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 to 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. Paul Steinhardt provided very helpful comments regarding the status of early universe models. 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.

Appendix

Relation between Bin–Bin and Multipole–Multipole Transfer Matrices

The power spectrum binning procedure can be described by an operator $P_{b\ell}$ that determines the relative weighting of each multipole ℓ within a bin b :

$$C_b = \sum_{\ell} P_{b\ell} C_{\ell}. \quad (\text{A1})$$

Likewise, we can bin the pseudo-power spectrum and write

$$\tilde{C}_b = \sum_{\ell} P_{b\ell} \tilde{C}_{\ell} = \sum_{\ell} P_{b\ell} \sum_{\ell'} J_{\ell\ell'} C_{\ell'} \quad (\text{A2})$$

as a consequence of Equation (8). Combining Equations (9) and (A2),

$$\sum_{b'} J_{bb'} \sum_{\ell'} P_{b'\ell'} C_{\ell'} = \sum_{\ell} P_{b\ell} \sum_{\ell'} J_{\ell\ell'} C_{\ell'}. \quad (\text{A3})$$

From this equation, we can identify the binned transfer matrix $J_{bb'}$ as taking the form

$$J_{bb'} = \frac{\sum_{\ell} P_{b\ell} \sum_{\ell' \in b'} J_{\ell\ell'} C_{\ell'}}{\sum_{\ell'} P_{b'\ell'} C_{\ell'}}. \quad (\text{A4})$$

Note that the inner sum in the numerator is confined to the input bin b' —recall each column of $J_{bb'}$ is the output of passing a single input bin through the reobservation pipeline (in analogy with the unit δ -functions used to construct $J_{\ell\ell'}$). Taking the sum over all input bins b' , as on the left-hand side of Equation (A3), recovers the sum over all input multipoles ℓ' , as on the right-hand side. For SPIDER, we employ a uniform weighting,

$$P_{b\ell} = \begin{cases} \frac{1}{\Delta\ell} & \text{if } \ell \in b, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A5})$$

which reduces Equation (A4) to

$$J_{bb'} = \sum_{\ell \in b} \left(\frac{\sum_{\ell' \in b'} J_{\ell\ell'} C_{\ell'}}{\sum_{\ell' \in b'} C_{\ell'}} \right). \quad (\text{A6})$$

The quantity in parentheses in Equation (A6) can be taken to be $J_{\ell b'}$, the proportion of the power in input bin b' that contributes to the power in output multipole ℓ . It is equal to the weighted average of all of the $J_{\ell\ell'}$ elements that lie within input bin b' , with the weights given by some input model spectrum C_{ℓ} . This highlights an important point: while $J_{\ell\ell'}$ is a function of the data analysis pipeline only, $J_{bb'}$ is additionally dependent on the input spectrum used to create it.

We can understand this point by noting how $J_{\ell\ell'}$ and $J_{bb'}$ are each applied to a target power spectrum c_{ℓ} :

1. First apply $J_{\ell\ell'}$, then bin the result:

$$\tilde{c}_b = \frac{1}{\Delta\ell} \sum_{\ell \in b} \sum_{\ell'} J_{\ell\ell'} c_{\ell'} \quad (\text{A7a})$$

2. First bin c_{ℓ} , then apply $J_{bb'}$:

$$\hat{c}_b = \sum_{b'} \left(\frac{\sum_{\ell \in b} \sum_{\ell' \in b'} J_{\ell\ell'} C_{\ell'}}{\sum_{\ell' \in b'} C_{\ell'}} \right) \left(\frac{1}{\Delta\ell} \sum_{\ell' \in b'} c_{\ell'} \right). \quad (\text{A7b})$$

If we let δ_{ℓ} represent the deviation of c_{ℓ} from C_{ℓ} (i.e., $c_{\ell} = C_{\ell} + \delta_{\ell}$; not to be confused with a δ -function centered at ℓ), then Equations A7(a) and A7(b) can be understood intuitively: whereas \tilde{c}_b represents the true amount of filtered power in bin b , \hat{c}_b is an estimate that differs from the true value by

$$\delta_b = \tilde{c}_b - \hat{c}_b = \frac{1}{\Delta\ell} \left(\sum_{\ell \in b} \sum_{\ell'} J_{\ell\ell'} \delta_{\ell'} - \sum_{b'} J_{bb'} \sum_{\ell' \in b'} \delta_{\ell'} \right). \quad (\text{A8})$$

In other words, the error from using $J_{bb'}$ is the difference between the average of transformed excess power in each

multipole ℓ and transform of average excess power in each bin b . Note that as $\delta_\ell \rightarrow 0$, the deviation $\delta_b \rightarrow 0$ as expected, i.e., $J_{bb'}$ returns the same result as $J_{\ell\ell'}$ when applied to the input spectrum C_ℓ used to create it. Consequently, the choice of input spectrum should be as close as possible to the target spectrum c_ℓ to be applied, such that δ_b is minimized.

ORCID iDs

J. S.-Y. Leung  <https://orcid.org/0000-0001-7116-3710>
 J. M. Nagy  <https://orcid.org/0000-0002-2036-7008>
 P. A. R. Ade  <https://orcid.org/0000-0002-5127-0401>
 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>
 C. R. Contaldi  <https://orcid.org/0000-0001-7285-0707>
 O. Doré  <https://orcid.org/0000-0001-7432-2932>
 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/0000-0003-4245-2315>
 J. E. Gudmundsson  <https://orcid.org/0000-0003-1760-0355>
 M. Halpern  <https://orcid.org/0000-0002-1760-0868>
 Z. Huang  <https://orcid.org/0000-0002-1506-1063>
 W. C. Jones  <https://orcid.org/0000-0002-3636-1241>
 S. Li  <https://orcid.org/0000-0002-8896-911X>
 L. Moncelsi  <https://orcid.org/0000-0002-4242-3015>
 I. L. Padilla  <https://orcid.org/0000-0002-0024-2662>
 A. S. Rahlin  <https://orcid.org/0000-0003-3953-1776>
 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

- Abazajian, K., Addison, G., Adshead, P., et al. 2019, arXiv:1907.04473
 Ade, P., Aguirre, J., Ahmed, Z., et al. 2019, *JCAP*, 2019, 056
 Ade, P., Ahmed, Z., Aikin, R., et al. 2016, *ApJ*, 825, 66
 Allison, R., Caucal, P., Calabrese, E., Dunkley, J., & Louis, T. 2015, *PhRvD*, 92, 123535
 Alonso Monge, D., Sanchez, J., & Slosar, A. 2019, *MNRAS*, 484, 4127
 Choi, S. K., Hasselfield, M., Ho, S.-P. P., et al. 2020, *JCAP*, 2020, 045
 Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, *MNRAS*, 350, 914
 Dutcher, D., Balkenhol, L., Ade, P., et al. 2021, *PhRvD*, 104, 022003
 Gambrel, A. E., Rahlin, A. S., Song, X., et al. 2021, *ApJ*, 922, 132
 Gorski, K. M., Hivon, E., Banday, A., et al. 2005, *ApJ*, 622, 759
 Hazumi, M., Ade, P. A., Adler, A., et al. 2020, *Proc. SPIE*, 11443, 114432F
 Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, 567, 2
 Jarosik, N., Barnes, C., Greason, M., et al. 2007, *ApJS*, 170, 263
 Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, 538, 473
 Loken, C., Gruner, D., Groer, L., et al. 2010, *JPhCS*, 256, 012026
 Lueker, M., Reichardt, C., Schaffer, K., et al. 2010, *ApJ*, 719, 1045
 Nagy, J., Ade, P., Amiri, M., et al. 2017, *ApJ*, 844, 151
 Padilla, I. L., Eimer, J. R., Li, Y., et al. 2020, *ApJ*, 889, 105
 Planck Collaboration, Adam, R., Ade, P., et al. 2016, *A&A*, 594, A7
 Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *A&A*, 641, A3
 Ponce, M., van Zon, R., Northrup, S., et al. 2019, in Proc. Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning), PEARC '19, ed. T. Furlani (New York: Association for Computing Machinery), 34
 Shaw, E. C., Ade, P. A. R., Akers, S., et al. 2020, *Proc. SPIE*, 11453, 114532F
 SPIDER Collaboration 2021, arXiv:2103.13334
 Tristram, M. 2006, in Proc. of the CMB and Physics of the Early Universe (CMB2006), ed. G. De Zotti et al. (Trieste: PoS), 63
 Tristram, M., Macías-Pérez, J. F., Renault, C., & Santos, D. 2005, *MNRAS*, 358, 833
 van der Vorst, H. A. 1992, *SIAM Journal on Scientific and Statistical Computing*, 13, 631
 Watts, D. J., Wang, B., Ali, A., et al. 2018, *ApJ*, 863, 121
 Zaldarriaga, M., & Seljak, U. 1997, *PhRvD*, 55, 1830